


53-1002269-02
26 April 2011



PowerConnect B-Series

TI24X

Configuration Guide

Information in this document is subject to change without notice.

© 2011 Dell Inc. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *Dell OpenManage* and *PowerConnect* are trademarks of Dell Inc.; *Microsoft*, *Windows* and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Regulatory Model Code: Turbolron 24X

Document History

Title	Publication number	Summary of changes	Date
<i>PowerConnect B-Series T124X Configuration Guide</i>	53-1002269-01	New document	March 2011
<i>PowerConnect B-Series T124X Configuration Guide</i>	53-1002269-02	Removed deprecated commands, web management and DVMRP references.	April 2011

Contents

About This Document

Introduction	xxxi
Audience	xxxi
Document conventions	xxxi
Text formatting	xxxi
Command syntax conventions	xxxii
Notes, cautions, and danger notices	xxxii
Notice to the reader	xxxii
Related publications	xxxii
Getting technical help or reporting errors	xxxiii
Contacting Dell	xxxiii

Chapter 1

Getting Familiar with Management Applications

Using the management port	1
How the management port works	1
CLI Commands for use with the management port	1
Logging on through the CLI	3
On-line help	3
Command completion	4
Scroll control	4
Line editing commands	4
Using and port number with CLI commands	5
CLI nomenclature on PowerConnect devices	5
Searching and filtering output from CLI commands	5
Using special characters in regular expressions	7
Creating an alias for a CLI command	9
Logging on through Brocade Network Advisor	10

Chapter 2

Configuring Basic Software Features

Configuring basic system parameters	11
Entering system administration information	12
Configuring Simple Network Management Protocol (SNMP) parameters	12
Disabling Syslog messages and traps for CLI access	16
Configuring an interface as the source for all Telnet packets	17
Cancelling an outbound Telnet session	18
Specifying a Simple Network Time Protocol (SNTP) server	18
Setting the system clock	19
Limiting broadcast, multicast, and unknown unicast traffic	21

Configuring basic port parameters	24
Assigning a port name	25
Modifying port speed and duplex mode	25
Auto speed detect	26
Modifying port duplex mode	26
Disabling or re-enabling a port	27
Disabling or re-enabling flow control	27
Auto-negotiation and advertisement of flow control	28
Configuring the Interpacket Gap (IPG)	29
Changing the Gbps fiber negotiation mode	30
Modifying port priority (QoS)	30
Configuring port flap dampening	30
Port loop detection	33

Chapter 3

Operations, Administration, and Maintenance

Overview	39
Determining the software versions installed and running on a device	39
Determining the flash image version running on the device ..	39
Determining the image versions installed in flash memory ..	40
Flash image verification	40
Image file types	42
Upgrading software	42
Upgrading the boot code	42
Upgrading the flash code	42
Boot code synchronization feature	43
Using SNMP to upgrade software	43
Changing the block size for TFTP file transfers	44
Rebooting	44
Displaying the boot preference	45
Loading and saving configuration files	45
Replacing the startup configuration with the running configuration	46
Replacing the running configuration with the startup configuration	46
Logging changes to the startup-config file	46
Copying a configuration file to or from a TFTP server	47
Dynamic configuration loading	47
Maximum file sizes for startup-config file and running-config ..	50
Scheduling a system reload	50
Reloading at a specific time	50
Reloading after a specific amount of time	51
Displaying the amount of time remaining before a scheduled reload	51
Canceling a scheduled reload	51
Diagnostic error codes and remedies for TFTP transfers	51

Chapter 4	Monitoring Hardware Components	
	Hardware support	53
	Digital optical monitoring	53
	Supported media	53
	Media not supported	53
	Supported media	54
	Media not supported	54
	Configuration limitations	54
	Enabling digital optical monitoring	54
	Setting the alarm interval	55
	Displaying information about installed media	55
	Viewing optical monitoring information	56
	Syslog messages	58
Chapter 5	Configuring IPv6 Connectivity	
	IPv6 addressing overview	59
	IPv6 address types	60
	IPv6 stateless autoconfiguration	62
	IPv6 CLI command support	62
	Configuring an IPv6 host address on a Layer 2 switch	63
	Configuring a global or site-local IPv6 address with a manually configured interface ID	64
	Configuring a link-local IPv6 address as a system-wide address for a switch	64
	Configuring the management port for an IPv6 automatic address configuration	65
	Configuring basic IPv6 connectivity on a Layer 3 switch	65
	Configuring IPv6 on each router interface	65
	IPv6 management (IPv6 host support)	68
	Restricting SNMP access to an IPv6 node	68
	Specifying an IPv6 SNMP trap receiver	68
	SNMP V3 over IPv6	69
	SNTP over IPv6	69
	Secure Shell, SCP, and IPv6	69
	IPv6 Telnet	69
	Configuring name-to-IPv6 address resolution using IPv6 DNS resolver	70
	Defining an IPv6 DNS entry	70
	Using the IPv6 copy command	71
	Using the IPv6 ncopy command	73
	IPv6 ping	74
	Configuring an IPv6 Syslog server	76
	Viewing IPv6 SNMP server addresses	76
	Disabling router advertisement and solicitation messages	77
	IPv6 debug	77
	Disabling IPv6 on a Layer 2 switch	77

Configuring IPv6 neighbor discovery	77
Configuration notes	78
Neighbor solicitation and advertisement messages	78
Configuring static neighbor entries	79
Clearing global IPv6 information	79
Clearing the IPv6 cache	80
Clearing IPv6 neighbor information	80
Clearing IPv6 traffic statistics	81
Displaying global IPv6 information	81
Displaying IPv6 cache information	81
Displaying IPv6 interface information	82
Displaying IPv6 neighbor information	84
Displaying IPv6 TCP information	85
Displaying IPv6 traffic statistics	88
.....	89

Chapter 6

Configuring Spanning Tree Protocol (STP) Related Features

STP overview	93
Configuring standard STP parameters	93
STP parameters and defaults	93
Enabling or disabling the Spanning Tree Protocol (STP)	95
Changing STP bridge and port parameters	96
STP protection enhancement	98
Displaying STP information	99
Configuring STP related features	106
802.1W Rapid Spanning Tree (RSTP)	107
802.1W Draft 3	144
Single Spanning Tree (SSTP)	148
PVST/PVST+ compatibility	150
Overview of PVST and PVST+	151
VLAN tags and dual mode	152
Configuring PVST+ support	153
Displaying PVST+ support information	153
Configuration examples	154
PVRST compatibility	157
BPDU guard	157
Enabling BPDU protection by port	157
Re-enabling ports disabled by BPDU guard	158
Displaying the BPDU guard status	158
Example console messages	159
Root guard	159
Enabling STP root guard	160
Displaying the STP root guard	160
Displaying the root guard by VLAN	160

802.1s Multiple Spanning Tree Protocol	161
Multiple spanning-tree regions	161
Configuration notes	163
Configuring MSTP mode and scope	163
Configuring additional MSTP parameters	164

Chapter 7

Configuring Basic Layer 2 Features

Enabling or disabling the Spanning Tree Protocol (STP)	175
Modifying STP bridge and port parameters	175
Changing the MAC age time and disabling MAC address learning	176
Disabling the automatic learning of MAC addresses	176
Displaying the MAC address table	177
Configuring static MAC entries	177
Multi-port static MAC address	178
Configuring VLAN-based static MAC entries	179
Clearing MAC address entries	179
Enabling port-based VLANs	180
Assigning IEEE 802.1Q tagging to a port	180
Defining MAC address filters	181
Configuration notes and limitations	181
Command syntax	181
Enabling logging of management traffic permitted by MAC filters	183
Displaying and modifying system parameter default settings	184
Configuration considerations	184
Displaying system parameter default values	184
Modifying system parameter default values	186
Egress buffer thresholds for QoS priorities	187
Cut-Through Switching Support on PowerConnect B-Series TI24X Switches	188
Default settings for egress buffer thresholds	188
Disabling and re-enabling the default settings for egress buffer thresholds	189
Setting the egress buffer threshold for all QoS priorities on a port or group of ports	189
Setting the egress buffer threshold for a specific QoS priority on a port or group of ports	190
Link Fault Signaling (LFS) for 10G	190
Jumbo frame support	191

Chapter 8

Configuring Metro Features

Topology groups	193
Master VLAN and member VLANs	193
Control ports and free ports	194
Configuration considerations	194
Configuring a topology group	194
Displaying topology group information	195
Metro Ring Protocol (MRP)	197
Configuration notes	199
MRP rings without shared interfaces (MRP Phase 1)	199
MRP rings with shared interfaces (MRP Phase 2)	200
Ring initialization	202
How ring breaks are detected and healed	205
Alarm RHP	208
Master VLANs and customer VLANs	209
Configuring MRP	211
Using MRP diagnostics	213
Displaying MRP information	214
MRP CLI example	216
Virtual Switch Redundancy Protocol (VSRP)	218
Configuration notes	220
Layer 2 and Layer 3 redundancy	220
Master election and failover	220
VSRP-Aware security features	225
VSRP parameters	225
Configuring basic VSRP parameters	228
Configuring optional VSRP parameters	229
Displaying VSRP information	238
VSRP fast start	241
VSRP and MRP signaling	242

Chapter 9

Configuring Uni-Directional Link Detection (UDLD) and Protected Link Groups

UDLD overview	245
Configuration considerations	246
Enabling UDLD	246
Changing the Keepalive interval	246
Changing the Keepalive retries	247
UDLD for tagged ports	247
Displaying UDLD information	247
Clearing UDLD statistics	249

Chapter 10

Configuring Virtual LANs (VLANs)

VLAN overview	251
Types of VLANs	251
Default VLAN	255
802.1Q tagging	256
Spanning Tree Protocol (STP)	258
Virtual routing interfaces	259
VLAN and virtual routing interface groups	260
Dynamic, static, and excluded port membership	261
Super aggregated VLANs	263
Trunk group ports and VLAN membership	263
Routing between VLANs	263
Virtual routing interfaces (Layer 3 Switches only)	263
Routing between VLANs using virtual routing interfaces (Layer 3 Switches only)	264
Dynamic port assignment (Layer 2 Switches and Layer 3 Switches)	265
Assigning a different VLAN ID to the default VLAN	265
Assigning different VLAN IDs to reserved VLANs 4091 and 4092	265
Assigning trunk group ports	266
Configuring port-based VLANs	267
Modifying a port-based VLAN	270
Enable spanning tree on a VLAN	271
Configuring IP subnet, IPX network and protocol-based VLANs	272
Configuration example	272
Configuring an IPv6 protocol VLAN	274
Routing between VLANs using virtual routing interfaces (Layer 3 Switches only)	275
Configuring uplink ports within a port-based VLAN	281
Configuration considerations	281
Configuration syntax	281
Configuring the same IP subnet address on multiple port-based VLANs	282
Configuring VLAN groups and virtual routing interface groups	285
Configuring a VLAN group	285
Configuring a virtual routing interface group	287
Displaying the VLAN group and virtual routing interface group information	288
Allocating memory for more VLANs or virtual routing interfaces	288
Configuring super aggregated VLANs	289
Configuration note	292
Configuring aggregated VLANs	292
Verifying the configuration	293
Complete CLI examples	293
Configuring 802.1Q-in-Q tagging	296
Configuration rules	297
Enabling 802.1Q-in-Q tagging	297
Example configuration	298

Configuring private VLANs	300
Configuration notes	301
Configuration notes and limitations for PowerConnect devices	302
Command syntax	302
CLI example for Figure 71	304
Enabling broadcast, unregistered multicast or unknown unicast traffic to the private VLAN on PowerConnect device	304
Dual-mode VLAN ports	305
Displaying VLAN information	307
Displaying VLANs in alphanumeric order	307
Displaying system-wide VLAN information	308
Displaying VLAN information for specific ports	309

Chapter 11

Configuring Trunk Groups and Dynamic Link Aggregation

Trunk group overview	311
Trunk group connectivity to a server	312
Trunk group rules	312
Trunk group configuration examples	313
Flexible trunk group membership	314
Trunk group load sharing	314
Configuring a trunk group	316
CLI syntax	316
Example 1: Configuring the trunk groups shown in Figure 75	317
Example 2: Configuring a trunk group that spans two Gbps Ethernet modules in a chassis device	317
Example 3: Configuring a multi-slot trunk group with one port per module	318
Example 4: Configuring a trunk group of 10 Gbps Ethernet ports	318
Additional trunking options	318
Displaying trunk group configuration information	323
Dynamic link aggregation	324
Examples of valid LACP trunk groups	325
Configuration notes and limitations	325
Adaptation to trunk disappearance	327
Flexible trunk eligibility	327
Enabling dynamic link aggregation	328
How changing the VLAN membership of a port affects trunk groups and dynamic keys	330
Link aggregation parameters	330
Displaying and determining the status of aggregate links	335
Events that affect the status of ports in an aggregate link	335
Displaying link aggregation and port status information	336
Displaying LACP status information	338
Clearing the negotiated aggregate links table	338

Configuring single link LACP	338
Configuration notes	339
CLI syntax	339

Chapter 12

Configuring GARP VLAN Registration Protocol

GVRP overview	341
Application examples	341
Dynamic core and fixed edge	342
Dynamic core and dynamic edge	343
Fixed core and dynamic edge	343
Fixed core and fixed edge	343
VLAN names	344
Configuration notes	344
Configuring GVRP	345
Changing the GVRP base VLAN ID	345
Increasing the maximum configurable value of the Leaveall timer	346
Enabling GVRP	346
Disabling VLAN advertising	347
Disabling VLAN learning	347
Changing the GVRP timers	347
Converting a VLAN created by GVRP into a statically-configured VLAN	349
Displaying GVRP information	349
Displaying GVRP configuration information	350
Displaying GVRP VLAN information	352
Displaying GVRP statistics	354
Displaying CPU utilization statistics	355
Displaying GVRP diagnostic information	356
Clearing GVRP statistics	357
CLI examples	357
Dynamic core and fixed edge	357
Dynamic core and dynamic edge	359
Fixed core and dynamic edge	359
Fixed core and fixed edge	359

Chapter 13

Configuring Rule-Based IP Access Control Lists

ACL overview	361
Types of IP ACLs	361
ACL IDs and entries	361
Numbered and named ACLs	362
Default ACL action	362
How hardware-based ACLs work	363
How fragmented packets are processed	363
Hardware aging of Layer 4 CAM entries	363
Configuration considerations	363

Configuring standard numbered ACLs	364
Standard numbered ACL syntax	364
Configuration example for standard numbered ACLs	366
Configuring standard named ACLs	366
Standard named ACL syntax	366
Configuration example for standard named ACLs	368
Configuring extended numbered ACLs	368
Extended numbered ACL syntax	369
Configuration examples for extended numbered ACLs	373
Configuring extended named ACLs	374
Extended named ACL syntax	375
Configuration example for extended named ACLs	379
Preserving user input for ACL TCP/UDP port numbers	379
Managing ACL comment text	379
Adding a comment to an entry in a numbered ACL	380
Applying an ACL to a virtual interface in a protocol- or subnet-based VLAN	380
Enabling ACL logging	381
Enabling strict control of ACL filtering of fragmented packets	383
Enabling ACL support for switched traffic in the router image	384
Enabling ACL filtering based on VLAN membership or VE port membership	384
Applying an IPv4 ACL to specific VLAN members on a port (Layer 2 devices only)	385
Applying an IPv4 ACL to a subset of ports on a virtual interface (Layer 3 devices only)	385
Filtering on IP precedence and ToS values	386
QoS options for IP ACLs	387
Using an IP ACL to mark DSCP values (DSCP marking)	387
DSCP matching	389
ACL-based rate limiting	389
Using ACLs to control multicast features	390
Enabling and viewing hardware usage statistics for an ACL	391
Displaying ACL information	391
.	392
Enabling and viewing hardware usage statistics for an ACL	392
Displaying ACL information	392
Troubleshooting ACLs	392

Chapter 14

Configuring Port Mirroring and Monitoring

Mirroring support by platform	395
---	-----

Configuring port mirroring and monitoring	395
Configuration notes	395
Monitoring a port	397
Monitoring an individual trunk port	397
ACL-based inbound mirroring	398
Creating an ACL-based inbound mirror clause for PowerConnect B-Series TI24X devices	398
MAC filter-based mirroring	402
Configuring MAC filter-based mirroring on PowerConnect B-Series TI24X devices	402

Chapter 15

Configuring Quality of Service

Classification	405
Processing of classified traffic	405
QoS queues	408
Assigning QoS priorities to traffic	408
Buffer allocation/threshold for QoS queues	410
Marking	410
Configuring DSCP-based QoS	410
Application notes	411
Using ACLs to honor DSCP-based QoS	411
Configuring the QoS mappings	411
Default DSCP -> Internal forwarding priority mappings	411
Changing the DSCP -> internal forwarding priority mappings	412
Changing the internal forwarding priority -> hardware forwarding queue mappings	413
Scheduling	414
QoS Queuing methods	414
Selecting the QoS queuing method	415
Configuring the QoS queues	415
Viewing QoS settings	418
Viewing DSCP-based QoS settings	418

Chapter 16

Configuring Rate Limiting and Rate Shaping on the PowerConnect B-Series TI24X

Rate limiting overview	421
Rate limiting in hardware	421
How Fixed Rate Limiting works	421
Configuration notes	422
Configuring a port-based rate limiting policy	422
Configuring an ACL-based rate limiting policy	423
Displaying the fixed rate limiting configuration	423

	Rate shaping overview	424
	Configuration notes	424
	Configuring outbound rate shaping for a port	424
	Configuring outbound rate shaping for a specific priority.	425
	Configuring outbound rate shaping for a trunk port	425
	Displaying rate shaping configurations	425
Chapter 17	Configuring Traffic Policies	
	About traffic policies	427
	Configuration notes and feature limitations	427
	Maximum number of traffic policies supported on a device	428
	Setting the maximum number of traffic policies supported on a Layer 3 device	429
	ACL-based rate limiting using traffic policies.	429
	Support for fixed rate limiting and adaptive rate limiting	430
	Configuring ACL-based fixed rate limiting.	430
	Configuring ACL-based adaptive rate limiting	431
	Specifying the action to be taken for packets that are over the limit.	433
	ACL and rate limit counting	434
	Enabling ACL statistics	434
	Enabling ACL statistics with rate limiting traffic policies.	435
	Viewing ACL and rate limit counters	436
	Clearing ACL and rate limit counters	437
	Viewing traffic policies	437
Chapter 18	Configuring IP Multicast Traffic Reduction for PowerConnect B- Series TI24X Switches	
	IGMP snooping overview.	439
	IGMP V1, V2, and V3 snooping support	440
	Queriers and non-queriers	440
	IGMP snooping enhancements.	441
	Configuration notes and feature limitations for PowerConnect B-Series TI24X devices	441
	PIM SM traffic snooping overview	442
	PIM SM snooping support.	443
	Application examples.	443
	Configuration notes and limitations	444

Configuring IGMP snooping	445
Enabling IGMP snooping globally on the device	447
Configuring the IGMP mode	447
Configuring the IGMP version	448
Disabling IGMP snooping on a VLAN	448
Disabling transmission and receipt of IGMP packets on a port	449
Modifying the age interval for group membership entries ..	449
Modifying the query interval (active IGMP snooping mode only)	449
Modifying the maximum response time	450
Configuring report control	450
Modifying the wait time before stopping traffic when receiving a leave message	450
Modifying the multicast cache age time	451
Enabling or disabling error and warning messages	451
Configuring static router ports	451
Turning off static group proxy	451
IGMP V3 membership tracking and fast leave	452
Fast leave for IGMP V2	452
Fast convergence	453
Configuring PIM SM snooping	453
Enabling or disabling PIM SM snooping	453
Enabling PIM SM snooping on a VLAN	454
Disabling PIM SM snooping on a VLAN	454
IGMP snooping show commands	454
Displaying the IGMP snooping configuration	454
Displaying IGMP snooping errors	455
Displaying IGMP group information	456
Displaying IGMP snooping mcache information	457
Displaying software resource usage for VLANs	458
Displaying the status of IGMP snooping traffic	459
PIM SM snooping show commands	460
Displaying PIM SM snooping information	460
Displaying PIM SM snooping information on a Layer 2 switch	460
Displaying PIM SM snooping information for a specific group or source group pair	461
Clear commands for IGMP snooping	462
Clearing the IGMP mcache	462
Clearing the mcache on a specific VLAN	462
Clearing traffic on a specific VLAN	463
Clearing IGMP counters on VLANs	463

Chapter 19

Configuring IP Multicast Protocols

Overview of IP multicasting	465
IPv4 multicast group addresses	465
Mapping of IPv4 Multicast group addresses to Ethernet MAC addresses	466
Supported Layer 3 multicast routing protocols	466
Multicast terms	466

Changing global IP multicast parameters	467
Changing dynamic memory allocation for IP multicast groups	467
Changing IGMP V1 and V2 parameters	468
PIM Dense	470
Initiating PIM multicasts on a network	470
Pruning a multicast tree	470
Grafts to a multicast Tree	472
PIM DM versions	472
Configuring PIM DM	473
Failover time in a multi-path topology	477
Modifying the TTL	477
PIM Sparse	478
PIM Sparse switch types	478
RP paths and SPT paths	479
Configuring PIM Sparse	479
Displaying PIM Sparse configuration information and statistics	489
Passive multicast route insertion	501
Multicast Source Discovery Protocol (MSDP)	501
Peer Reverse Path Forwarding (RPF) flooding	503
Source active caching	503
Configuring MSDP	504
Designating an interface IP address as the RP IP address	505
Filtering MSDP source-group pairs	506
MSDP mesh groups	509
Displaying MSDP information	515
Clearing MSDP information	519
Using ACLs to control multicast features	520
Using ACLs to limit static RP groups	520
Using ACLs to limit PIM RP candidate advertisement	522
Configuring a static multicast route	523
Tracing a multicast route	525
Displaying the multicast configuration for another multicast router	526
IGMP V3	527
Default IGMP version	528
Compatibility with IGMP V1 and V2	528
Globally enabling the IGMP version	528
Enabling the IGMP version per interface setting	528
Enabling the IGMP version on a physical port within a virtual routing interface	529
Enabling membership tracking and fast leave	529
Setting the query interval	530
Setting the group membership time	530
Setting the maximum response time	530
IGMP V3 and source specific multicast protocols	531
Displaying IGMP V3 information on Layer 3 Switches	531
Clearing IGMP statistics	535

IGMP Proxy	535
Configuration notes	535
Configuring IGMP Proxy	536
Displaying IGMP Proxy traffic	536

Chapter 20

Configuring LLDP

Terms used in this chapter	537
LLDP overview	538
Benefits of LLDP	538
General operating principles	539
Operating modes	539
LLDP packets	540
TLV support	540
MIB support	543
Syslog messages	543
Configuring LLDP	544
Configuration notes and considerations	544
Enabling and disabling LLDP	545
Changing a port LLDP operating mode	545
Specifying the maximum number of LLDP neighbors	546
Enabling LLDP SNMP notifications and syslog messages	547
Changing the minimum time between LLDP transmissions	548
Changing the interval between regular LLDP transmissions	548
Changing the holdtime multiplier for transmit TTL	549
Changing the minimum time between port reinitializations	549
LLDP TLVs advertised by the device	549
Displaying LLDP statistics and configuration settings	555
LLDP configuration summary	555
LLDP statistics	556
LLDP neighbors	557
LLDP neighbors detail	558
LLDP configuration details	560
Resetting LLDP statistics	561
Clearing cached LLDP neighbor information	561

Chapter 21

Configuring IP

Basic configuration	563
Overview	563
IP interfaces	564
IP packet flow through a Layer 3 Switch	564
IP route exchange protocols	569
IP multicast protocols	569
IP interface redundancy protocols	570
Access Control Lists and IP access policies	570

Basic IP parameters and defaults – Layer 3 Switches	570
When parameter changes take effect	571
IP global parameters – Layer 3 Switches	571
IP interface parameters – Layer 3 Switches	575
Basic IP parameters and defaults – Layer 2 Switches	576
IP global parameters – Layer 2 Switches	576
Interface IP parameters – Layer 2 Switches	578
Configuring IP parameters – Layer 3 Switches	578
Configuring IP addresses	579
Configuring packet parameters	581
Changing the router ID	584
Specifying a single source interface for Telnet, TACACS/TACACS+, or RADIUS Packets	585
Configuring ARP parameters	587
Configuring forwarding parameters	592
Disabling ICMP messages	594
Configuring static routes	596
Configuring a default network route	604
Configuring IP load sharing	605
Configuring IRDP	608
Configuring RARP	610
Configuring UDP broadcast and IP helper parameters	612
Configuring BootP/DHCP relay parameters	615
Configuring IP parameters – Layer 2 Switches	616
Configuring the management IP address and specifying the default gateway	616
Configuring Domain Name Server (DNS) resolver	617
Changing the TTL threshold	619
Configuring DHCP Assist	619
Displaying IP configuration information and statistics	623
Changing the network mask display to prefix format	623
Displaying IP information – Layer 3 Switches	623
Displaying IP information – Layer 2 Switches	637

Chapter 22

Configuring RIP

RIP overview	643
ICMP host unreachable message for undeliverable ARPs	643
RIP parameters and defaults	643
RIP global parameters	644
RIP interface parameters	644

Configuring RIP parameters	645
Enabling RIP	645
Configuring metric parameters	646
Changing the administrative distance	647
Configuring redistribution	647
Configuring route learning and advertising parameters	650
Changing the route loop prevention method	651
Suppressing RIP route advertisement on a VRRP or VRRPE backup interface	652
Configuring RIP route filters	652
Displaying RIP filters	653
Displaying CPU utilization statistics	654

Chapter 23

Configuring OSPF Version 2 (IPv4)

Overview of OSPF	657
OSPF point-to-point Links	658
Designated routers in multi-access networks	659
Designated router election in multi-access networks	659
OSPF RFC 1583 and 2178 compliance	660
Reduction of equivalent AS External LSAs	661
Support for OSPF RFC 2328 Appendix E	663
Dynamic OSPF activation and configuration	664

Configuring OSPF	665
Configuration rules	665
OSPF parameters	665
Enable OSPF on the router	666
Assign OSPF areas	667
Assigning an area range (optional)	671
Assigning interfaces to an area	671
Modify interface defaults	671
Change the timer for OSPF authentication changes	674
Block flooding of outbound LSAs on specific OSPF interfaces	675
Assign virtual links	675
Modify virtual link parameters	677
Changing the reference bandwidth for the cost on	
OSPF interfaces	678
Define redistribution filters	680
Prevent specific OSPF routes from being installed in the	
IP route table	682
Modify default metric for redistribution	685
Enable route redistribution	686
Disable or re-enable load sharing	687
Configure external route summarization	688
Configure default route origination	690
Modify SPF timers	691
Modify redistribution metric type	691
Modify administrative distance	692
Configure OSPF group Link State Advertisement	
(LSA) pacing	693
Modify OSPF traps generated	693
Modify OSPF standard compliance setting	694
Modify exit overflow interval	694
Specifying the types of OSPF Syslog messages to log	695
Clearing OSPF information	695
Clearing OSPF neighbor information	695
Clearing OSPF topology information	696
Clearing redistributed routes from the OSPF routing table ..	696
Clearing information for OSPF areas	696
Displaying OSPF information	697
Displaying general OSPF configuration information	697
Displaying CPU utilization statistics	698
Displaying OSPF area information	700
Displaying OSPF neighbor information	700
Displaying OSPF interface information	702
Displaying OSPF route information	704
Displaying OSPF external link state information	706
Displaying OSPF link state information	707
Displaying the data in an LSA	707
Displaying OSPF virtual neighbor information	708
Displaying OSPF virtual link information	708
Displaying OSPF ABR and ASBR information	708
Displaying OSPF trap status	709

Chapter 24

Configuring VRRP and VRRPE

Overview	711
Overview of VRRP	711
Overview of VRRPE	716
Configuration note	719
Comparison of VRRP and VRRPE	719
VRRP	719
VRRPE	719
Architectural differences	719
VRRP and VRRPE parameters	720
Configuring basic VRRP parameters	722
Configuring the Owner	723
Configuring a Backup	723
Configuration rules for VRRP	723
Configuring basic VRRPE parameters	723
Configuration rules for VRRPE	724
Note regarding disabling VRRP or VRRPE	724
Configuring additional VRRP and VRRPE parameters	724
Forcing a Master router to abdicate to a standby router	731
Displaying VRRP and VRRPE information	732
Displaying summary information	732
Displaying detailed information	734
Displaying statistics	739
Clearing VRRP or VRRPE statistics	740
Displaying CPU utilization statistics	740
Configuration examples	742
VRRP example	742
VRRPE example	743

Chapter 25

Configuring BGP4

Overview of BGP4	745
Relationship between the BGP4 route table and the IP route table	746
How BGP4 selects a path for a route	747
BGP4 message types	748
Basic configuration and activation for BGP4	750
Note regarding disabling BGP4	751
BGP4 parameters	751
When parameter changes take effect	752
Memory considerations	754
Memory configuration options obsoleted by dynamic memory	754

Basic configuration tasks	755
Enabling BGP4 on the router	755
Changing the router ID.	755
Setting the local AS number	756
Adding a loopback interface	756
Adding BGP4 neighbors.	756
Adding a BGP4 peer group	763
Optional configuration tasks	767
Changing the Keep Alive Time and Hold Time.	767
Changing the BGP4 next-hop update timer	768
Enabling fast external fallover.	768
Changing the maximum number of paths for BGP4 load sharing.	769
Customizing BGP4 load sharing	770
Specifying a list of networks to advertise.	771
Changing the default local preference.	772
Using the IP default route as a valid next hop for a BGP4 route	773
Advertising the default route.	773
Changing the default MED (Metric) used for route redistribution	773
Enabling next-hop recursion	774
Changing administrative distances	777
Requiring the first AS to be the neighbor AS	778
Disabling or re-enabling comparison of the AS-Path length	778
Enabling or disabling comparison of the router IDs	779
Configuring the Layer 3 Switch to always compare Multi-Exit Discriminators (MEDs)	779
Treating missing MEDs as the worst MEDs	780
Configuring route reflection parameters	780
Aggregating routes advertised to BGP4 neighbors	784
Modifying redistribution parameters	785
Redistributing connected routes.	785
Redistributing RIP routes.	786
Redistributing OSPF external routes.	786
Redistributing static routes.	787
Disabling or re-enabling re-advertisement of all learned BGP4 routes to all BGP4 neighbors	787
Redistributing IBGP routes into RIP and OSPF.	788
Filtering	788
Filtering specific IP addresses	788
Filtering AS-paths.	790
Filtering communities	793
Defining IP prefix lists	795
Defining neighbor distribute lists	796
Defining route maps	797
Using a table map to set the rag value.	805
Configuring cooperative BGP4 route filtering.	806

Configuring route flap dampening	809
Globally configuring route flap dampening	810
Using a route map to configure route flap dampening for specific routes	810
Using a route map to configure route flap dampening for a specific neighbor.	811
Removing route dampening from a route.	812
Removing route dampening from a neighbor routes suppressed due to aggregation	812
Displaying and clearing route flap dampening statistics	814
Generating traps for BGP	815
Displaying BGP4 information	816
Displaying summary BGP4 information	816
Displaying the active BGP4 configuration	818
Displaying CPU utilization statistics	819
Displaying summary neighbor information	820
Displaying BGP4 neighbor information.	822
Displaying peer group information	833
Displaying summary route information	834
Displaying the BGP4 route table.	835
Displaying BGP4 route-attribute entries.	841
Displaying the routes BGP4 has placed in the IP route table	842
Displaying route flap dampening statistics	843
Displaying the active route map configuration	844
Updating route information and resetting a neighbor session . . .	845
Using soft reconfiguration.	845
Dynamically requesting a route refresh from a BGP4 neighbor	848
Closing or resetting a neighbor session.	851
Clearing and resetting BGP4 routes in the IP route table. . .	851
Clearing traffic counters.	852
Clearing route flap dampening statistics.	852
Removing route flap dampening	852
Clearing diagnostic buffers.	853

Chapter 26

Securing Access to Management Functions

Securing access methods	855
-----------------------------------	-----

Restricting remote access to management functions	857
Using ACLs to restrict remote access	857
Defining the console idle time	859
Restricting remote access to the device to specific IP addresses	860
Restricting access to the device based on IP or MAC address	861
Specifying the maximum number of login attempts for Telnet access	861
Restricting remote access to the device to specific VLAN IDs	862
Designated VLAN for Telnet management sessions to a Layer 2 Switch	863
Device management security	863
Disabling specific access methods.	864
Setting passwords.	865
Setting a Telnet password	866
Setting passwords for management privilege levels	866
Recovering from a lost password	868
Displaying the SNMP community string	869
Disabling password encryption	869
Specifying a minimum password length.	869
Setting up local user accounts.	870
Enhancements to username and password	870
Configuring a local user account	874
Create password option.	876
Changing a local user password	876
Configuring TACACS/TACACS+ security	877
How TACACS+ differs from TACACS.	877
TACACS/TACACS+ authentication, authorization, and accounting	877
TACACS authentication	878
TACACS/TACACS+ configuration considerations	881
Enabling TACACS	881
Identifying the TACACS/TACACS+ servers.	882
Specifying different servers for individual AAA functions	883
Setting optional TACACS/TACACS+ parameters.	883
Configuring authentication-method lists for TACACS/TACACS+	884
Configuring TACACS+ authorization	886
Configuring TACACS+ accounting	889
Configuring an interface as the source for all TACACS/TACACS+ packets.	891
Displaying TACACS/TACACS+ statistics and configuration information	891

Configuring RADIUS security	892
RADIUS authentication, authorization, and accounting	893
RADIUS configuration considerations	896
RADIUS configuration procedure	896
Configuring Dell-specific attributes on the RADIUS server ...	896
Enabling SNMP to configure RADIUS	897
Identifying the RADIUS server to the device	898
Specifying different servers for individual AAA functions ...	898
Configuring a RADIUS server per port	898
Mapping a RADIUS server to individual ports	899
Setting RADIUS parameters	900
Configuring authentication-method lists for RADIUS	901
Configuring RADIUS authorization	903
Configuring RADIUS accounting	905
Configuring an interface as the source for all RADIUS packets	906
Displaying RADIUS configuration information	906
Configuring authentication-method lists	907
Configuration considerations for authentication- method lists	908
Examples of authentication-method lists	909

Chapter 27

Configuring SSH2 and SCP

SSH version 2 support	911
Tested SSH2 clients	911
Supported features	912
Unsupported features	912
AES encryption for SSH2	912
Configuring SSH2	913
Recreating SSH keys	914
Generating a host key pair	914
Configuring DSA challenge-response authentication	915
Setting optional parameters	917
Setting the number of SSH authentication retries	918
Deactivating user authentication	918
Enabling empty password logins	918
Setting the SSH port number	919
Setting the SSH login timeout value	919
Designating an interface as the source for all SSH packets (Layer 3 code only)	919
Configuring the maximum idle time for SSH sessions	920
Filtering SSH access using ACLs	920
Terminating an active SSH connection	920
Displaying SSH connection information	920
Using Secure copy with SSH2	922
Enabling and disabling SCP	922
Example file transfers using SCP	922

IETF RFC support.	925
How 802.1X port security works	925
Device roles in an 802.1X configuration	925
Communication between the devices	926
Controlled and uncontrolled ports	928
Message exchange during authentication.	929
Authenticating multiple hosts connected to the same port	931
802.1X port security and sFlow	933
Configuring 802.1X port security.	933
Configuring an authentication method list for 802.1X	934
Setting RADIUS parameters	934
Configuring dynamic VLAN assignment for 802.1X ports.	938
Dynamically applying IP ACLs and MAC filters to 802.1X ports	941
Enabling 802.1X port security.	945
Setting the port control	945
Configuring periodic re-authentication.	946
Re-authenticating a port manually.	947
Setting the quiet period.	947
Specifying the wait interval and number of EAP-request/ identity frame retransmissions from the PowerConnect device	947
Specifying the wait interval and number of EAP-request/ identity frame retransmissions from the RADIUS server	948
Specifying a timeout for retransmission of messages to the authentication server	949
Initializing 802.1X on a port	949
Allowing access to multiple hosts.	949
Configuring VLAN access for non-EAP-capable clients	952
Displaying 802.1X information.	953
Displaying 802.1X configuration information	953
Displaying 802.1X statistics	955
Clearing 802.1X statistics.	956
Displaying dynamically assigned VLAN information	957
Displaying information about dynamically applied MAC filters and IP ACLs	958
Displaying 802.1X multiple-host authentication information.	959
Sample 802.1X configurations.	963
Point-to-point configuration.	963
Hub configuration	964
802.1X Authentication with dynamic VLAN assignment.	965
Using multi-device port authentication and 802.1X security on the same port.	966
Configuring Dell-specific attributes on the RADIUS server	967
Example configurations.	968

Chapter 29	Using the MAC Port Security Feature	
	Overview	973
	Local and global resources	973
	Configuration notes and feature limitations	974
	Configuring the MAC port security feature	974
	Enabling the MAC port security feature	974
	Setting the maximum number of secure MAC addresses for an interface	975
	Setting the port security age timer	975
	Specifying secure MAC addresses	975
	Autosaving secure MAC addresses to the startup-config file	976
	Specifying the action taken when a security violation occurs	976
	Clearing port security statistics	978
	Clearing restricted MAC addresses	978
	Clearing violation statistics	978
	Displaying port security information	978
	Displaying port security settings	978
	Displaying the secure MAC addresses	979
	Displaying port security statistics	979
	Displaying restricted MAC addresses on a port	980
Chapter 30	Configuring Multi-Device Port Authentication	
	How multi-device port authentication works	981
	RADIUS authentication	981
	Authentication-failure actions	982
	Supported RADIUS attributes	982
	Support for dynamic VLAN assignment	983
	Support for dynamic ACLs	983
	Support for authenticating multiple MAC addresses on an interface	983
	Using multi-device port authentication and 802.1X security on the same port	983
	Configuring Dell-specific attributes on the RADIUS server	984

Configuring multi-device port authentication	985
Enabling multi-device port authentication	985
Specifying the format of the MAC addresses sent to the RADIUS server	986
Specifying the authentication-failure action	986
Generating traps for multi-device port authentication	987
Defining MAC address filters.	987
Configuring dynamic VLAN assignment	988
Dynamically applying IP ACLs to authenticated MAC addresses	990
Enabling denial of service attack protection	992
Clearing authenticated MAC addresses	993
Disabling aging for authenticated MAC addresses	993
Changing the hardware aging period for blocked MAC addresses	994
Specifying the aging time for blocked MAC addresses	995
Specifying the RADIUS timeout action	995
Multi-device port authentication password override	996
Limiting the number of authenticated MAC addresses.	997
Displaying multi-device port authentication information	997
Displaying authenticated MAC address information	997
Displaying multi-device port authentication configuration information	998
Displaying multi-device port authentication information for a specific MAC address or port	998
Displaying the authenticated MAC addresses	999
Displaying the non-authenticated MAC addresses	999
Displaying multi-device port authentication information for a port.	1000
Displaying multi-device port authentication settings and authenticated MAC addresses	1001

Chapter 31

Protecting Against Denial of Service Attacks

Protecting against Smurf attacks.	1005
Avoiding being an intermediary in a Smurf attack.	1005
Avoiding being a victim in a Smurf attack	1006
Protection against ICMP attacks in PowerConnect devices	1006
Protecting against TCP SYN attacks.	1007
Protection against TCP-SYN attacks in PowerConnect devices	1007
TCP security enhancement	1008
Displaying statistics about packets dropped because of DoS attacks	1009
Displaying statistics about packets dropped because of DoS attacks in PowerConnect devices	1010

Chapter 32

Securing SNMP Access

SNMP overview	1011
-------------------------	------

Establishing SNMP community strings	1011
Encryption of SNMP community strings	1012
Adding an SNMP community string	1012
Displaying the SNMP community strings	1013
Configuring your NMS	1014
Configuring SNMP version 3	1015
Defining the engine id	1015
Defining an SNMP group	1016
Defining an SNMP user account	1017
Defining SNMP views	1018
SNMP version 3 traps	1019
Defining an SNMP group and specifying which view is notified of traps	1019
Trap MIB changes	1021
Specifying an IPv6 host as an SNMP trap receiver	1021
Displaying SNMP Information	1022
Displaying the Engine ID	1022
Displaying SNMP groups	1022
Displaying user information	1022
Interpreting varbinds in report packets	1023
SNMP v3 Configuration examples	1023
Simple SNMP v3 configuration	1023
More detailed SNMP v3 configuration	1024

Chapter 33

Enabling the Foundry Discovery Protocol and Reading Cisco Discovery Protocol Packets

Using FDP	1025
Configuring FDP	1025
Displaying FDP information	1026
Clearing FDP and CDP information	1029
Reading CDP packets	1030
Enabling interception of CDP packets globally	1030
Enabling interception of CDP packets on an interface	1030
Displaying CDP information	1030
Clearing CDP information	1032

Chapter 34

Using Syslog

Overview	1035
Displaying Syslog messages	1036
Enabling real-time display of Syslog messages	1036
Enabling real-time display for a Telnet or SSH session	1036
Show log on all terminals	1037

Configuring the Syslog service	1037
Displaying the Syslog configuration	1037
Disabling or re-enabling Syslog	1041
Specifying a Syslog server	1041
Specifying an additional Syslog server	1041
Disabling logging of a message level	1042
Changing the number of entries the local buffer can hold .	1042
Changing the log facility	1042
Displaying Interface names in Syslog messages	1043
Displaying TCP or UDP port numbers in Syslog messages .	1044
Clearing the Syslog messages from the local buffer	1044
Syslog messages	1044

Appendix A

Network Monitoring

Basic management	1069
Viewing system information	1069
Viewing configuration information	1069
Viewing port statistics	1070
Viewing STP statistics	1072
Clearing statistics	1072
Traffic counters for outbound traffic	1073
RMON support	1076
Maximum number of entries allowed in the RMON control table	1076
Statistics (RMON group 1)	1076
History (RMON group 2)	1078
Alarm (RMON group 3)	1079
Event (RMON group 9)	1079
sFlow	1079
sFlow support for IPv6 packets	1080
Configuration considerations	1081
Configuring and enabling sFlow	1082
Displaying sFlow information	1087
Configuring a utilization list for an uplink port	1090
Command syntax	1090
Displaying utilization percentages for an uplink	1091

Appendix B

Software Specifications

IEEE compliance	1093
RFC support	1093
Internet drafts	1098

About This Document

Introduction

This guide describes the following product families from Dell:

- PowerConnect B-Series T124X Layer 2 switch

This guide includes procedures for configuring the software. The software procedures show how to perform tasks using the CLI. This guide also describes how to monitor Dell products using statistics and summary screens.

This guide applies to the PowerConnect B-Series T124X models.

Audience

This document is designed for system administrators with a working knowledge of Layer 2 and Layer 3 switching and routing.

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords
	Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis
	Identifies variables
	Identifies document titles
<code>code text</code>	Identifies CLI output

For readability, command names in the narrative portions of this guide are presented in bold: for example, **show version**.

Command syntax conventions

Command syntax in this manual follows these conventions:

command and parameters

Commands and parameters are printed in bold.

[]

Optional parameter.

variable

Variables are printed in italics enclosed in angled brackets < >.

...

Repeat the previous element, for example “member[:member...]”

|

Choose from one of the parameters.

Notes, cautions, and danger notices

The following notices and statements are used in this manual. They are listed below in order of increasing severity of potential hazards.

NOTE

A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.



CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Notice to the reader

This document may contain references to the trademarks of the following corporations. These trademarks are the properties of their respective companies and corporations.

Related publications

- PowerConnect B- Series TI24X Hardware Installation Guide

- PowerConnect B-MLXe MIB Reference

NOTE

For the latest edition of this document, which contains the most up-to-date information, refer to support.dell.com.

Getting technical help or reporting errors

Dell is committed to ensuring that your investment in our products remains cost-effective. If you need assistance or find errors in the manuals, contact Dell Technical Support. When contacting Dell Technical Support have the device configuration file and an output capture of show tech-support command available.

Contacting Dell

For customers in the United States, call 800-WWW.DELL (800.999.3355).

NOTE

If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues:

1. Visit <http://support.dell.com>.
2. Click your country or region at the bottom of the page. For a full listing of countries and regions, click **All**.
3. In the Support menu, click **All Support**.
4. Choose the method of contacting Dell that is convenient for you.

Getting Familiar with Management Applications

Using the management port

The management port is an out-of-band port that customers can use to manage their devices without interfering with the in-band ports. The management port is widely used to download images and configurations and for Telnet sessions.

The MAC address for the management port is derived from the base MAC address of the unit, plus the number of ports in the base module.

How the management port works

The following rules apply to management ports:

- Any packets that are specifically addressed to the management port MAC address or the broadcast MAC address are forwarded accordingly. All other packets are filtered out.
- No packet received on a management port is sent to any in-band ports, and no packets received on in-band ports are sent to a management port.
- A management port is not part of any VLAN
- Protocols are not supported on the management port.
- Creating a management VLAN disables the management port on the device.
- All features that can be configured from the global configuration mode can also be configured from the interface level of the management port. Features that are configured through the management port take effect globally, not on the management port itself (on switches only).

For switches, any in-band port may be used for management purposes. A Router sends Layer 3 packets using the MAC address of the port as the source MAC address.

CLI Commands for use with the management port

The following CLI commands can be used with a management port.

To display the current configuration, use the **show running-config interface management** command.

Syntax: **show running-config interface management** *num*

```
PowerConnect(config-if-mgmt)#ip addr 10.44.9.64/24
PowerConnect(config)#show running-config interface management 1
interface management 1
ip address 10.44.9/64 255.255.255.0
```

To display the current configuration, use the **show interfaces management** command.

Syntax: **show interfaces management** *num*

1 Using the management port

```
PowerConnect(config)#show interfaces management 1
GigEthernetmgmt1 is up, line protocol is up
Hardware is GigEthernet, address is 0000.9876.544a (bia 0000.9876.544a)
Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual none
BPRU guard is disabled, ROOT protect is disabled
Link Error Dampening is Disabled
STP configured to OFF, priority is level0, mac-learning is enabled
Flow Control is config disabled, oper enabled
Mirror disabled, Monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 0 bits-time, IPG GMII 0 bits-time
IP MTU 1500 bytes
300 second input rate: 83728 bits/sec, 130 packets/sec, 0.01% utilization
300 second output rate: 24 bits/sec, 0 packets/sec, 0.00% utilization
39926 packets input, 3210077 bytes, 0 no buffer
Received 4353 broadcasts, 32503 multicasts, 370 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runs, 0 giants
22 packets output, 1540 bytes, 0 underruns
Transmitted 0 broadcasts, 6 multicasts, 16 unicasts
0 output errors, 0 collisions
```

To display the management interface information in brief form, enter the **show interfaces brief management** command.

Syntax: **show interfaces brief management** *num*

Port	Link State	Dupl	Speed	Trunk	Tag	Pvid	Pri	MAC Name
mgmt1	Up	None	Full 100M	None	No	1	0	000d.5200.0118

To display management port statistics, enter the **show statistics management** command.

Syntax: **show statistics management** *num*

```
PowerConnect(config)# show statistics management 1
Port Link State Dupl Speed Trunk Tag Pvid Pri MAC Name
mgmt1 Up None Full 100M None No 1 0 000d.5200.0118

Port mgmt1 Counters:
  InOctets3210941OutOctets1540
  InPkts39939OutPackets22
InBroadcastPkts4355OutbroadcastPkts0
InMultiastPkts35214OutMulticastPkts6
InUnicastPkts370OutUnicastPkts16
InBadPkts0
InFragments0
InDiscards0OutErrors0
CRC 0 Collisions0
InErrors0 LateCollisions0
InGiantPkts0
InShortPkts0
InJabber0
InFlowCtrlPkts0OutFlowCtrlPkts0
InBitsPerSec83728OutBitsPerSec24
InPktsPerSec130OutPktsPerSec0
InUtilization0.01%OutUtilization0.00%
```


To display the management interface statistics in brief form, enter the **show statistics brief management** command.

Syntax: `show statistics brief management num`

```
PowerConnect(config)#show statistics brief management 1
PortIn PacketsOut PacketsTrunkIn ErrorsOut Errors
mgmt139946220 0

Total139945220 0
```

Logging on through the CLI

Once an IP address is assigned to a Dell device running Layer 2 software or to an interface on a Dell device running Layer 3 software, you can access the CLI either through the direct serial connection to the device or through a local or remote Telnet session.

You can initiate a local Telnet or SNMP connection by attaching a cable to a port and specifying the assigned management station IP address.

The commands in the CLI are organized into the following levels:

- **User EXEC** – Lets you display information and perform basic tasks such as pings and traceroutes.
- **Privileged EXEC** – Lets you use the same commands as those at the User EXEC level plus configuration commands that do not require saving the changes to the system-config file.
- **CONFIG** – Lets you make configuration changes to the device. To save the changes across reboots, you need to save them to the system-config file. The CONFIG level contains sub-levels for individual ports, for VLANs, for routing protocols, and other configuration areas.

NOTE

By default, any user who can open a serial or Telnet connection to the Dell device can access all these CLI levels. To secure access, you can configure Enable passwords or local user accounts, or you can configure the device to use a RADIUS or TACACS/TACACS+ server for authentication. Refer to [Chapter 26, “Securing Access to Management Functions”](#).

On-line help

To display a list of available commands or command options, enter “?” or press Tab. If you have not entered part of a command at the command prompt, all the commands supported at the current CLI level are listed. If you enter part of a command, then enter “?” or press Tab, the CLI lists the options you can enter at this point in the command string.

If you enter an invalid command followed by ?, a message appears indicating the command was unrecognized. An example is given below.

```
PowerConnect(config)#rooter ip
Unrecognized command
```

Command completion

The CLI supports command completion, so you do not need to enter the entire name of a command or option. As long as you enter enough characters of the command or option name to avoid ambiguity with other commands or options, the CLI understands what you are typing.

Scroll control

By default, the CLI uses a page mode to paginate displays that are longer than the number of rows in your terminal emulation window. For example, if you display a list of all the commands at the global CONFIG level but your terminal emulation window does not have enough rows to display them all at once, the page mode stops the display and lists your choices for continuing the display. An example is given below.

```
aaa
all-client
appletalk
arp
boot
some lines omitted for brevity...

ipx
lock-address
logging
mac
--More--, next page: Space, next line:
Return key, quit: Control-c
```

The software provides the following scrolling options:

- Press the **Space bar** to display the next page (one screen at a time).
- Press the **Return** or **Enter** key to display the next line (one line at a time).
- Press **Ctrl+C** or **Ctrl+Q** to cancel the display.

Line editing commands

The CLI supports the following line editing commands. To enter a line-editing command, use the CTRL+key combination for the command by pressing and holding the CTRL key, then pressing the letter associated with the command.

TABLE 1 CLI line editing commands

Ctrl+Key combination	Description
Ctrl+A	Moves to the first character on the command line.
Ctrl+B	Moves the cursor back one character.
Ctrl+C	Escapes and terminates command prompts and ongoing tasks (such as lengthy displays), and displays a fresh command prompt.
Ctrl+D	Deletes the character at the cursor.
Ctrl+E	Moves to the end of the current command line.
Ctrl+F	Moves the cursor forward one character.
Ctrl+K	Deletes all characters from the cursor to the end of the command line.

TABLE 1 CLI line editing commands (Continued)

Ctrl+Key combination	Description
Ctrl+L; Ctrl+R	Repeats the current command line on a new line.
Ctrl+N	Enters the next command line in the history buffer.
Ctrl+P	Enters the previous command line in the history buffer.
Ctrl+U; Ctrl+X	Deletes all characters from the cursor to the beginning of the command line.
Ctrl+W	Deletes the last word you typed.
Ctrl+Z	Moves from any CONFIG level of the CLI to the Privileged EXEC level; at the Privileged EXEC level, moves to the User EXEC level.

Using and port number with CLI commands

Many CLI commands require users to enter port numbers as part of the command syntax, and many **show** command outputs display port numbers. The port numbers are entered and displayed in one of the following formats.

CLI nomenclature on PowerConnect devices

The PowerConnect devices use port numbers only. When you enter CLI commands that require port numbers as part of the syntax, just specify the port number.

Here are some examples. The following commands change the CLI from the global CONFIG level to the configuration level for the first port on the device:

- PowerConnect commands

```
PowerConnect(config)#interface e1
PowerConnect(config-if-e10000-1)#
```

Searching and filtering output from CLI commands

You can filter CLI output from **show** commands and at the **–More–** prompt. You can search for individual characters, strings, or construct complex regular expressions to filter the output.

Searching and filtering output from Show commands

You can filter output from **show** commands to display lines containing a specified string, lines that do not contain a specified string, or output starting with a line containing a specified string. The search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. Refer to [“Using special characters in regular expressions”](#) on page 7 for information on special characters used with regular expressions.

Using include to display lines containing a specified string

The **include** modifier filters the output of the **show interface** command for port 11 so it displays only lines containing the word “Internet”. This command can be used to display the IP address of the interface.

1 Using and port number with CLI commands

```
PowerConnect#show interface e 11 | include Internet
Internet address is 192.168.1.11/24, MTU 1518 bytes, encapsulation ethernet
```

Syntax: *show-command | include regular-expression*

NOTE

The vertical bar (|) is part of the command.

The regular expression specified as the search string is case sensitive. In the example above, a search string of “Internet” would match the line containing the IP address, but a search string of “internet” would not.

Using exclude to display lines that do not contain a specified string

The **exclude** modifier filters the output of the **show who** command so it displays only lines that do not contain the word “closed”. This command can be used to display open connections to the device

```
PowerConnect#show who | exclude closed
Console connections:
    established
    you are connecting to this session
    2 seconds in idle
Telnet connections (inbound):
  1    established, client ip address 192.168.9.37
      27 seconds in idle
Telnet connection (outbound):
SSH connections:
```

Syntax: *show-command | exclude regular-expression*

Using begin to display lines starting with a specified string

The **begin** modifier filters the output of the **show who** command so it displays output starting with the first line that contains the word “SSH”. This command can be used to display information about SSH connections to the device.

```
PowerConnect#show who | begin SSH
SSH connections:
  1    established, client ip address 192.168.9.210
      7 seconds in idle
  2    closed
  3    closed
  4    closed
  5    closed
```

Syntax: *show-command | begin regular-expression*

Searching and filtering output at the --More-- prompt

The **--More--** prompt displays when output extends beyond a single page. From this prompt, you can press the Space bar to display the next page, the Return or Enter key to display the next line, or Ctrl+C or Q to cancel the display. In addition, you can search and filter output from this prompt.

At the **--More--** prompt, you can press the forward slash key (/) and then enter a search string. The device displays output starting from the first line that contains the search string, similar to the **begin** modifier for **show** commands. An example is given below.

```
--More--, next page: Space, next line: Return key, quit: Control-c
/telnet
```

The results of the search are displayed.

```
searching...
telnet      Telnet by name or IP address
temperature temperature sensor commands
terminal    display syslog
traceroute  TraceRoute to IP node
undebg      Disable debugging functions (see also 'debug')
undetele    Undetele flash card files
whois       WHOIS lookup
write       Write running configuration to flash or terminal
```

To display lines containing only a specified search string (similar to the **include** modifier for **show** commands) press the plus sign key (+) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
+telnet
```

The filtered results are displayed.

```
filtering...
telnet      Telnet by name or IP address
```

To display lines that do not contain a specified search string (similar to the **exclude** modifier for **show** commands) press the minus sign key (-) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
-telnet
```

The filtered results are displayed.

```
filtering...
temperature      temperature sensor commands
terminal         display syslog
traceroute       TraceRoute to IP node
undebg           Disable debugging functions (see also 'debug')
undetele         Undetele flash card files
whois            WHOIS lookup
write            Write running configuration to flash or terminal
```

As with the modifiers for filtering output from **show** commands, the search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. See the next section for information on special characters used with regular expressions.

Using special characters in regular expressions

You use a regular expression to specify a single character or multiple characters as a search string. In addition, you can include special characters that influence the way the software matches the output against the search string. These special characters are listed in the following table.

1 Using and port number with CLI commands

TABLE 2 Special characters for regular expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches “aaz”, “abz”, “acz”, and so on, but not just “az”: a.z
*	The asterisk matches on zero or more sequential instances of a pattern. For example, the following regular expression matches output that contains the string “abc”, followed by zero or more Xs: abcX*
+	The plus sign matches on one or more sequential instances of a pattern. For example, the following regular expression matches output that contains “de”, followed by a sequence of “g”s, such as “deg”, “degg”, “deggg”, and so on: deg+
?	The question mark matches on zero occurrences or one occurrence of a pattern. For example, the following regular expression matches output that contains “dg” or “deg”: de?g NOTE: Normally when you type a question mark, the CLI lists the commands or options at that CLI level that begin with the character or string you entered. However, if you enter Ctrl+V and then type a question mark, the question mark is inserted into the command line, allowing you to use it as part of a regular expression.
^	A caret (when not used within brackets) matches on the beginning of an input string. For example, the following regular expression matches output that begins with “deg”: ^deg
\$	A dollar sign matches on the end of an input string. For example, the following regular expression matches output that ends with “deg”: deg\$
_	An underscore matches on one or more of the following: <ul style="list-style-type: none">• , (comma)• { (left curly brace)• } (right curly brace)• ((left parenthesis)•) (right parenthesis)• The beginning of the input string• The end of the input string• A blank space For example, the following regular expression matches on “100” but not on “1002”, “2100”, and so on. _100_
[]	Square brackets enclose a range of single-character patterns. For example, the following regular expression matches output that contains “1”, “2”, “3”, “4”, or “5”: [1-5] You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets. <ul style="list-style-type: none">• ^ - The caret matches on any characters except the ones in the brackets. For example, the following regular expression matches output that does not contain “1”, “2”, “3”, “4”, or “5”: [^1-5]• - The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. See the example above.

TABLE 2 Special characters for regular expressions (Continued)

Character	Operation
	A vertical bar separates two alternative values or sets of values. The output can match one or the other value. For example, the following regular expression matches output that contains either “abc” or “defg”: abc defg
()	Parentheses allow you to create complex expressions. For example, the following complex expression matches on “abc”, “abcabc”, or “defg”, but not on “abcdefgdefg”: ((abc)+) ((defg)?)

If you want to filter for a special character instead of using the special character as described in the table above, enter “\” (backslash) in front of the character. For example, to filter on output containing an asterisk, enter the asterisk portion of the regular expression as “*”.

```
PowerConnect#show ip route bgp | include \*
```

Creating an alias for a CLI command

You can create **aliases** for CLI commands. An alias serves as a shorthand version of a longer CLI command. For example, you can create an alias called **shoro** for the CLI command **show ip route**. Then when you enter **shoro** at the command prompt, the **show ip route** command is executed.

To create an alias called **shoro** for **show ip route**, enter the following command.

```
PowerConnect(config)#alias shoro = show ip route
```

Syntax: [no] **alias** *alias-name* = *cli-command*

The <*alias-name*> must be a single word, without spaces.

After the alias is configured, entering **shoro** at either the Privileged EXEC or CONFIG levels of the CLI executes the **show ip route** command.

To create an alias called **wrsbc** for **copy running-config tftp 10.10.10.10 test.cfg**, enter the following command.

```
PowerConnect(config)#alias wrsbc = copy running-config tftp 10.10.10.10 test.cfg
```

To remove the **wrsbc** alias from the configuration, enter one of the following commands.

```
PowerConnect(config)#no alias wrsbc
```

or

```
PowerConnect(config)#unalias wrsbc
```

Syntax: **unalias** *alias-name*

The specified *alias-name* must be the name of an alias already configured on the device.

To display the aliases currently configured on the device, enter the following command at either the Privileged EXEC or CONFIG levels of the CLI.

```
PowerConnect#alias
      wrsbc      copy running-config tftp 10.10.10.10 test.cfg
      shoro      show ip route
```

Syntax: **alias**

1 Logging on through Brocade Network Advisor

Configuration notes

The following configuration notes apply to this feature:

- You cannot include additional parameters with the alias at the command prompt. For example, after you create the **shoro** alias, **shoro bgp** would not be a valid command.
- If configured on the device, authentication, authorization, and accounting is performed on the actual command, not on the alias for the command.
- To save an alias definition to the startup-config file, use the **write memory** command.

Logging on through Brocade Network Advisor

Refer to the *Brocade Network Advisor* manuals for information about using Brocade Network Advisor.

Configuring Basic Software Features

Configuring basic system parameters

Dell devices are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the advanced features such as VLANs or routing protocols for the device must first be enabled at the system (global) level before they can be configured. If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.

NOTE

Before assigning or modifying any router parameters, you must assign the IP subnet (interface) addresses for each port.

NOTE

For information about configuring IP addresses, DNS resolver, DHCP assist, and other IP-related parameters, refer to [Chapter 21, “Configuring IP”](#).

For information about the Syslog buffer and messages, refer to [Chapter 34, “Using Syslog”](#).

The procedures in this section describe how to configure the basic system parameters listed in [Table 3](#).

TABLE 3 Basic system parameters

Basic system parameter	See page
System name, contact, and location	page 12
SNMP trap receiver, trap source address, and other parameters	page 12
Single source address for all Telnet packets	page 17
Single source address for all TFTP packets	page 18
Single source address for all Syslog packets	page 18
Single source address for all SNTP packets	page 18
System time using a Simple Network Time Protocol (SNTP) server or local system counter	page 18
System clock	page 19
Broadcast, multicast, or unknown-unicast limits, if required to support slower third-party devices	page 21

NOTE

For information about the Syslog buffer and messages, refer to [Chapter 34, “Using Syslog”](#).

Entering system administration information

You can configure a system name, contact, and location for a device and save the information locally in the configuration file for future reference. This information is not required for system operation but is suggested. When you configure a system name, the name replaces the default system name in the CLI command prompt.

The name, contact, and location each can be up to 32 alphanumeric characters.

Here is an example of how to configure a system name, system contact, and location.

```
PowerConnect(config)# hostname zappa
zappa(config)#snmp-server contact Support Services
zappa(config)#snmp-server location Centerville
zappa(config)#end
zappa# write memory
```

Syntax: `hostname` *string*

Syntax: `snmp-server contact` *string*

Syntax: `snmp-server location` *string*

The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

NOTE

The **chassis name** command does not change the CLI prompt. Instead, the command assigns an administrative ID to the device.

Configuring Simple Network Management Protocol (SNMP) parameters

Use the procedures in this section to perform the following configuration tasks:

- Specify an SNMP trap receiver.
- Specify a source address and community string for all traps sent by the device.
- Change the holddown time for SNMP traps
- Disable individual SNMP traps. (All traps are enabled by default.)
- Disable traps for CLI access that is authenticated by a local user account, a RADIUS server, or a TACACS/TACACS+ server.

NOTE

To add and modify “get” (read-only) and “set” (read-write) community strings, refer to [Chapter 26, “Securing Access to Management Functions”](#).

Specifying an SNMP trap receiver

You can specify a trap receiver to ensure that all SNMP traps sent by the device go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. When you specify the host, you also specify a community string. The device sends all the SNMP traps to the specified hosts and includes the specified community string. Administrators can therefore filter for traps from a device based on IP address or community string.

When you add a trap receiver, the software automatically encrypts the community string you associate with the receiver when the string is displayed by the CLI. If you want the software to show the community string in the clear, you must explicitly specify this when you add a trap receiver. In either case, the software does not encrypt the string in the SNMP traps sent to the receiver.

To specify the host to which the device sends all SNMP traps, use one of the following methods.

To add a trap receiver and encrypt the display of the community string, enter commands such as the following.

To specify an SNMP trap receiver and change the UDP port that will be used to receive traps, enter a command such as the following.

```
PowerConnect(config)# snmp-server host 2.2.2.2 0 mypublic port 200
PowerConnect(config)# write memory
```

Syntax: `snmp-server host ip-addr [0 | 1] string [port value]`

The *ip-addr* parameter specifies the IP address of the trap receiver.

The **0 | 1** parameter specifies whether you want the software to encrypt the string (**1**) or show the string in the clear (**0**). The default is **0**.

The *string* parameter specifies an SNMP community string configured on the device. The string can be a read-only string or a read-write string. The string is not used to authenticate access to the trap host but is instead a useful method for filtering traps on the host. For example, if you configure each of your devices that use the trap host to send a different community string, you can easily distinguish among the traps from different devices based on the community strings.

The command in the example above adds trap receiver 2.2.2.2 and configures the software to encrypt display of the community string. When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server host 2.2.2.2 1 <encrypted-string>
```

To add a trap receiver and configure the software to encrypt display of the community string in the CLI, enter commands such as the following.

```
device(config)# snmp-server host 2.2.2.2 0 device-12
device(config)# write memory
```

The **port value** parameter allows you to specify which UDP port will be used by the trap receiver. This parameter allows you to configure several trap receivers in a system. With this parameter, device and another network management application can coexist in the same system. devices can be configured to send copies of traps to more than one network management application.

Specifying a single trap source

You can specify a single trap source to ensure that all SNMP traps sent by the device use the same source IP address. When you configure the SNMP source address, you specify the Ethernet port, loopback interface, or virtual interface that is the source for the traps. The device then uses the lowest-numbered IP address configured on the port or interface as the source IP address in the SNMP traps sent by the device.

Identifying a single source IP address for SNMP traps provides the following benefits:

- If your trap receiver is configured to accept traps only from specific links or IP addresses, you can use this feature to simplify configuration of the trap receiver by configuring the device to always send the traps from the same link or source address.

2 Configuring basic system parameters

- If you specify a loopback interface as the single source for SNMP traps, SNMP trap receivers can receive traps regardless of the states of individual links. Thus, if a link to the trap receiver becomes unavailable but the receiver can be reached through another link, the receiver still receives the trap, and the trap still has the source IP address of the loopback interface.

To specify a port, loopback interface, or virtual interface whose lowest-numbered IP address the device must use as the source for all SNMP traps sent by the device, use the following CLI method.

To configure the device to send all SNMP traps from the first configured IP address on port 4, enter the following commands.

```
PowerConnect(config)# snmp trap-source ethernet 4
PowerConnect(config)# write memory
```

Syntax: `snmp-server trap-source loopback num | ethernet <portnum> | ve num`

The *num* parameter is a loopback interface or virtual interface number.

To specify a loopback interface as the SNMP trap source for the device, enter commands such as the following.

```
PowerConnect(config)# int loopback 1
PowerConnect(config-lbif-1)# ip address 10.0.0.1/24
PowerConnect(config-lbif-1)# exit
PowerConnect(config)# snmp-server trap-source loopback 1
```

The commands in this example configure loopback interface 1, assign IP address 10.00.1/24 to the loopback interface, then designate the interface as the SNMP trap source for this device.

Regardless of the port the device uses to send traps to the receiver, the traps always arrive from the same source IP address.

Setting the SNMP trap holddown time

When a device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the device might not be able to reach the servers, in which case the messages are lost.

By default, a device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the device sends the traps, including traps such as “cold start” or “warm start” that occur before the holddown time expires.

You can change the holddown time to a value from one second to ten minutes.

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# snmp-server enable traps holddown-time 30
```

The command in this example changes the holddown time for SNMP traps to 30 seconds. The device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: `[no] snmp-server enable traps holddown-time secs`

The *secs* parameter specifies the number of seconds and can be from 1 – 600 (ten minutes). The default is 60 seconds.

Disabling SNMP traps

PowerConnect devices come with SNMP trap generation enabled by default for all traps. You can selectively disable one or more of the following traps.

NOTE

By default, all SNMP traps are enabled at system startup.

Layer 2 traps

The following traps are generated on devices running Layer 2 software:

- SNMP authentication keys
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation

Layer 3 traps

The following traps are generated on devices running Layer 3 software:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- BGP4
- OSPF
- VRRP
- VRRPE

To stop link down occurrences from being reported, enter the following.

```
PowerConnect(config)# no snmp-server enable traps link-down
```

Syntax: [no] snmp-server enable traps *trap-type*

Disabling Syslog messages and traps for CLI access

PowerConnect devices send Syslog messages and SNMP traps when a user logs into or out of the User EXEC or Privileged EXEC level of the CLI. The feature applies to users whose access is authenticated by an authentication-method list based on a local user account, RADIUS server, or TACACS/TACACS+ server.

NOTE

The Privileged EXEC level is sometimes called the “Enable” level, because the command for accessing this level is **enable**.

The feature is enabled by default.

Examples of Syslog messages for CLI access

When a user whose access is authenticated by a local user account, a RADIUS server, or a TACACS/TACACS+ server logs into or out of the CLI User EXEC or Privileged EXEC mode, the software generates a Syslog message and trap containing the following information:

- The time stamp
- The user name
- Whether the user logged in or out
- The CLI level the user logged into or out of (User EXEC or Privileged EXEC level)

NOTE

Messages for accessing the User EXEC level apply only to access through Telnet. The device does not authenticate initial access through serial connections but does authenticate serial access to the Privileged EXEC level. Messages for accessing the Privileged EXEC level apply to access through the serial connection or Telnet.

The following examples show login and logout messages for the User EXEC and Privileged EXEC levels of the CLI.

```
PowerConnect# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 12 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Oct 15 18:01:11:info:dg logout from USER EXEC mode
Oct 15 17:59:22:info:dg logout from PRIVILEGE EXEC mode
Oct 15 17:38:07:info:dg login to PRIVILEGE EXEC mode
Oct 15 17:38:03:info:dg login to USER EXEC mode
```

Syntax: show logging

The first message (the one on the bottom) indicates that user “dg” logged in to the CLI User EXEC level on October 15 at 5:38 PM and 3 seconds (Oct 15 17:38:03). The same user logged into the Privileged EXEC level four seconds later.

The user remained in the Privileged EXEC mode until 5:59 PM and 22 seconds. (The user could have used the CONFIG modes as well. Once you access the Privileged EXEC level, no further authentication is required to access the CONFIG levels.) At 6:01 PM and 11 seconds, the user ended the CLI session.

Disabling the Syslog messages and traps

Logging of CLI access is enabled by default. If you want to disable the logging, enter the following commands.

```
PowerConnect(config)# no logging enable user-login
PowerConnect(config)# write memory
PowerConnect(config)# end
PowerConnect#reload
```

Syntax: [no] logging enable user-login

Configuring an interface as the source for all Telnet packets

You can designate the lowest-numbered IP address configured on an interface as the source IP address for all Telnet packets from the device. Identifying a single source IP address for Telnet packets provides the following benefits:

- If your Telnet server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the Telnet server by configuring the device to always send the Telnet packets from the same link or source address.
- If you specify a loopback interface as the single source for Telnet packets, Telnet servers can receive the packets regardless of the states of individual links. Thus, if a link to the Telnet server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS/TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an interface as the source for all Telnet packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the interface as the source IP address for Telnet packets originated by the device.

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all Telnet packets, enter commands such as the following.

```
PowerConnect(config)# int loopback 2
PowerConnect(config-lbif-2)# ip address 10.0.0.2/24
PowerConnect(config-lbif-2)# exit
PowerConnect(config)# ip telnet source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all Telnet packets from the device.

Syntax: ip telnet source-interface ethernetportnum | loopback num | ve num

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the device.

2 Configuring basic system parameters

```
PowerConnect(config)# interface ethernet 4
PowerConnect(config-if-e10000-4)# ip address 209.157.22.110/24
PowerConnect(config-if-e10000-4)# exit
PowerConnect(config)# ip telnet source-interface ethernet 4
```

Cancelling an outbound Telnet session

If you want to cancel a Telnet session from the console to a remote Telnet server (for example, if the connection is frozen), you can terminate the Telnet session by doing the following.

1. At the console, press **Ctrl+^** (Ctrl+Shift-6).
2. Press the **X** key to terminate the Telnet session.

Pressing **Ctrl+^** twice in a row causes a single **Ctrl+^** character to be sent to the Telnet server. After you press **Ctrl+^**, pressing any key other than **X** or **Ctrl+^** returns you to the Telnet session.

Specifying a Simple Network Time Protocol (SNTP) server

You can configure the device to consult SNTP servers for the current time and date.

NOTE

PowerConnect devices do not retain time and date information across power cycles. Unless you want to reconfigure the system time counter each time the system is reset, Dell recommends that you use the SNTP feature.

To identify an SNTP server with IP address 208.99.8.95 to act as the clock reference for a device, enter the following.

```
PowerConnect(config)# sntp server 208.99.8.95
```

Syntax: `sntp server ip-addr | hostname [version]`

The *version* parameter specifies the SNTP version the server is running and can be from 1 – 4. The default is 1. You can configure up to three SNTP servers by entering three separate **sntp server** commands.

By default, the device polls its SNTP server every 30 minutes (1800 seconds). To configure the device to poll for clock updates from a SNTP server every 15 minutes, enter the following.

```
PowerConnect(config)# sntp poll-interval 900
```

Syntax: `[no] sntp poll-interval 1-65535`

To display information about SNTP associations, enter the following command.

```
PowerConnect# show sntp associations
  address      ref clock      st  when  poll  delay  disp
~207.95.6.102  0.0.0.0        16  202   4    0.0    5.45
~207.95.6.101  0.0.0.0        16  202   0    0.0    0.0
* synced, ~ configured
```

Syntax: `show sntp associations`

The following table describes the information displayed by the **show sntp associations** command.

TABLE 4 Output from the show sntp associations command

This field...	Displays...
(leading character)	One or both of the following: * Synchronized to this peer ~ Peer is statically configured
address	IP address of the peer
ref clock	IP address of the peer reference clock
st	NTP stratum level of the peer
when	Amount of time since the last NTP packet was received from the peer
poll	Poll interval in seconds
delay	Round trip delay in milliseconds
disp	Dispersion in seconds

To display information about SNTP status, enter the following command.

```
PowerConnect# show sntp status
Clock is unsynchronized, stratum = 0, no reference clock
precision is 2**0
reference time is 0 .0
clock offset is 0.0 msec, root delay is 0.0 msec
root dispersion is 0.0 msec, peer dispersion is 0.0 msec
```

Syntax: show sntp status

The following table describes the information displayed by the **show sntp status** command.

TABLE 5 Output from the show sntp status command

This field...	Indicates...
unsynchronized	System is not synchronized to an NTP peer.
synchronized	System is synchronized to an NTP peer.
stratum	NTP stratum level of this system
reference clock	IP Address of the peer (if any) to which the unit is synchronized
precision	Precision of this system's clock (in Hz)
reference time	Reference time stamp
clock offset	Offset of clock to synchronized peer
root delay	Total delay along the path to the root clock
root dispersion	Dispersion of the root path
peer dispersion	Dispersion of the synchronized peer

Setting the system clock

In addition to SNTP support, switches and routers also allow you to set the system time counter. The time counter setting is not retained across power cycles and is not automatically synchronized with an SNTP server. The counter merely starts the system time and date clock with the time and date you specify.

NOTE

You can synchronize the time counter with your SNTP server time by entering the **sntp sync** command from the Privileged EXEC level of the CLI.

NOTE

Unless you identify an SNTP server for the system time and date, you will need to re-enter the time and date following each reboot.

For more details about SNTP, refer to [“Specifying a Simple Network Time Protocol \(SNTP\) server”](#) on page 18.

To set the system time and date to 10:15:05 on October 15, 2003, enter the following command.

```
PowerConnect# clock set 10:15:05 10-15-2003
```

Syntax: [no] **clock set** *hh:mm:ss* | *mm-dd-yy* | *mm-dd-yyyy*

By default, switches and routers do not change the system time for daylight saving time. To enable daylight saving time, enter the following command.

```
PowerConnect# clock summer-time
```

Syntax: **clock summer-time**

Although SNTP servers typically deliver the time and date in Greenwich Mean Time (GMT), you can configure the device to adjust the time for any one-hour offset from GMT or for one of the following U.S. time zones:

- US Pacific
- Alaska
- Aleutian
- Arizona
- Central
- East-Indiana
- Eastern
- Hawaii
- Michigan
- Mountain
- Pacific
- Samoa

To change the time zone to Australian East Coast time (which is normally 10 hours ahead of GMT), enter the following command.

```
PowerConnect(config)# clock timezone gmt+10
```

Syntax: **clock timezone** *gmt* | *us time-zone*

You can enter one of the following values for *time-zone*:

- US time zones (**us**): alaska, aleutian, arizona, central, east-indiana, eastern, hawaii, michigan, mountain, pacific, samoa.
- GMT time zones (**gmt**): gmt+0:00 to gmt+12:00 in increments of 1, and gmt-0:00 to gmt-12:00 in decrements of 1 are supported.

New start and end dates for US daylight saving time

NOTE

This feature applies to US time zones only.

Starting in 2007, the system will automatically change the system clock to Daylight Saving Time (DST), in compliance with the new federally mandated start of daylight saving time, which is extended one month beginning in 2007. The DST will start at 2:00am on the second Sunday in March and will end at 2:00am on the first Sunday in November.

The DST feature is automatic, but to trigger the device to the correct time, the device must be configured to the US time zone, not the GMT offset. To configure your device to use the US time zone, enter the following command.

```
PowerConnect(config)# clock timezone us pacific
```

Syntax: [no] clock timezone us *timezone-type*

Enter pacific, eastern, central, or mountain for *timezone-type*.

This command must be configured on every device that follows the US DST.

To verify the change, run a **show clock** command.

```
PowerConnect# show clock
```

Refer to October 19, 2006 - Daylight Saving Time 2007 Advisory, posted on service.Brocade.com for more information

Limiting broadcast, multicast, and unknown unicast traffic

PowerConnect devices can forward all flooded traffic at wire speed within a VLAN. However, some third-party networking devices cannot handle high rates of broadcast, multicast, or unknown-unicast traffic. If high rates of traffic are being received by the device on a given port of that VLAN, you can limit the number of broadcast, multicast, or unknown-unicast packets or bytes received each second on that port. This can help to control the number of such packets or bytes that are flooded on the VLAN to other devices.

Byte-based limiting for broadcast, multicast, and unknown unicast traffic provides the ability to rate limit traffic based on byte count instead of packet count. When the byte mode is enabled, packets will be received on a port as long as the number of bytes received per second is less than the corresponding limit. Once the limit is reached, further packets will be dropped.

PowerConnect devices do not support packet-based and byte-based limiting simultaneously on the same port. For example, if you configure packet-based limiting for broadcast traffic, you must also configure packet-based limiting for multicast and unknown unicast traffic. Likewise, if you configure byte-based limiting for broadcast traffic, you must also configure byte-based limiting for multicast and unknown unicast traffic.

Command syntax for packet-based limiting

To enable broadcast limiting on a group of ports by counting the number of packets received, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1 to 8
PowerConnect(config-mif-e10000-1-8)# broadcast limit 65536
```

2 Configuring basic system parameters

These commands configure packet-based broadcast limiting on ports 1 – 8. On each port, the maximum number of broadcast packets per second cannot exceed 65,536 packets per second.

On PowerConnect devices, multicast limiting is independent of broadcast limiting. To enable multicast limiting on devices, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1 to 8
PowerConnect(config-mif-e10000-1-8)# multicast limit 65536
```

To enable unknown unicast limiting by counting the number of packets received, enter commands such as the following.

```
PowerConnect(config)# interface eth 1
PowerConnect(config-if-e10000-1)# unknown-unicast limit 65536
The combined number of inbound Unknown Unicast packets permitted
    for ports 1 to 12 is now set to 65536
PowerConnect((config-if-e10000-1)#
```

Syntax: [no] broadcast limit <num>

Syntax: [no] unknown-unicast limit <num>

Syntax: [no]

or

Syntax: [no] multicast limit <num>

NOTE

The **multicast limit** <num> command applies to devices only.

The <num> variable specifies the maximum number of packets per second. Acceptable values differ depending on the device you are configuring:

- On PowerConnect devices, <num> can be any number between 1 and 8388607 (packets per second). The actual value will be determined by the system. Once you enter the value, the CLI will display a message indicating the actual value. The following shows an example configuration.

```
PowerConnect(config)# interface ethernet 9
PowerConnectconfig-mif-e10000-9)# multicast limit 50
Multicast limit in pkts/sec set to 31
```

If you specify 0, limiting is disabled. Limiting is disabled by default.

Command syntax for byte-based limiting

PowerConnect devices limit traffic based on kilobits per second (kbps). To enable limiting, refer to the appropriate section, below.

PowerConnect devices

To enable broadcast limiting on a group of ports by counting the number of kilobits received, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 9 to 10
PowerConnect(config-mif-e10000-9-10)# broadcast limit 131072 kbps
Broadcast limit in kbits/sec set to 130000
```

These commands configure broadcast limiting on ports 9 and 10. On each port, the total number of kilobits received from broadcast packets cannot exceed 130,000 per second.

To enable multicast limiting, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 8
PowerConnect(config-if-e10000-1-8)# multicast limit 9000 kbps
Multicast limit in kbits/sec set to 8064
```

To enable unknown unicast limiting, enter commands such as the following.

```
PowerConnect(config)# int e 13
PowerConnect(config-if-e10000-13)# unknown-unicast limit 65536 kbps
Unknown unicast limit in kbits/sec set to 64000
```

Syntax: [no] broadcast limit *num* kbps

Syntax: [no] multicast limit *num* kbps

Syntax: [no] unknown-unicast limit *num* kbps

The *num* variable can be any number between 1 and 10000000. The actual value will be determined by the system. Once you enter the value, the CLI will display a message indicating the actual value, as shown in the configuration examples above. If you specify 0, limiting is disabled. Limiting is disabled by default.

Viewing broadcast, multicast, and unknown unicast limits

You can use the **show run interface** command to display the broadcast, multicast, and unknown-unicast limits configured on the device.

In addition to the **show run interface** command, to display the broadcast, multicast, and unknown-unicast limits configured on the device:

- **show rate-limit unknown-unicast**
- **show rate-limit broadcast**

Use the **show run interface** command to view the broadcast, multicast, and unknown-unicast limit configured on each port.

```
PowerConnect# show run interface
interface ethernet 4
broadcast limit 1245184 bytes
multicast limit
!
interface ethernet 5
broadcast limit 1245184 bytes
multicast limit
!
interface ethernet 12
unknown-unicast limit 524288
!
interface ethernet 13
unknown-unicast limit 65536 bytes
!
interface ethernet 14
broadcast limit 65536
!
interface ethernet 23
broadcast limit 131072
multicast limit
!
```

Syntax: show run interface

2 Configuring basic port parameters

Use the **show rate-limit unknown-unicast** command to display the unknown unicast limit for each port region to which it applies.

```
PowerConnect# show rate-limit unknown-unicast
Unknown Unicast Limit Settings:
Port Region Combined Limit Packets/Bytes
  1 - 12          524288      Packets
 13 - 24          65536       Bytes
```

Syntax: show rate-limit unknown-unicast

Use the **show rate-limit broadcast** command to display the broadcast limit or broadcast and multicast limit for each port to which it applies.

```
PowerConnect# show rate-limit broadcast
Broadcast/Multicast Limit Settings:
Port   Limit   Packets/Bytes   Packet Type(s)
 4     1245184      Bytes          Broadcast + Multicast
 5     1245184      Bytes          Broadcast + Multicast
14      65536       Packets        Broadcast only
23     131072      Packets        Broadcast + Multicast
```

Syntax: show rate-limit broadcast

Configuring basic port parameters

The procedures in this section describe how to configure the port parameters shown in [Table 6](#).

TABLE 6 Basic port parameters

Port parameter	See page
Name	page 25
Speed	page 25
Duplex mode	page 26
Port status (enable or disable)	page 27
Flow control	page 27
Auto-negotiation and advertisement of flow control	page 28
Configuring PHY FIFO Rx and TX Depth	page 29
Interpacket Gap (IPG)	page 29
Gbps fiber negotiate mode	page 30
QoS priority	page 30
Port flap dampening	page 30

All ports are pre-configured with default values that allow the device to be fully operational at initial startup without any additional configuration. However, in some cases, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

Assigning a port name

A port name can be assigned to help identify interfaces on the network. You can assign a port name to physical ports, virtual interfaces, and loopback interfaces.

To assign a name to a port.

```
PowerConnect(config)# interface e 2
PowerConnect(config-if-e10000-2)# port-name Marsha
```

Syntax: `port-name text`

The `text` parameter is an alphanumeric string. The name can be up to 64 characters long. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks.

Modifying port speed and duplex mode

This section describes how to modify port speed and duplex mode on PowerConnect devices.

Copper ports

The Gigabit Ethernet copper ports are designed to auto-sense and auto-negotiate the speed and duplex mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10, 100, or 1000 Mbps. The default and recommended setting is 10/100/1000 auto-sense.

NOTE

On PowerConnect devices, you can modify the port speed of copper ports and the 24 fiber ports.

NOTE

For optimal link operation, copper ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), and Flow Control.

Fiber ports on the PowerConnect

The fiber ports on the PowerConnect devices support 1 GbE and 10 GbE connections, depending on the SFP optic installed in the port. *SFP+* optics are used for 10 GbE fiber connections, and *SFP* optics are used for 1 GbE fiber connections. The default setting is 10 GbE full-duplex mode with *SFP+* optics. To use 1 GbE in a 10 GbE port, insert an *SFP* optic and change the speed-duplex to 1 GbE (**speed-duplex 1000**).

Configuration syntax

The following commands change the port speed of fiber interface 8 on a PowerConnect device from the default of 10 Gbps to 1 Gbps.

```
PowerConnect(config)# interface e 8
PowerConnect(config-if-e10000-8)# speed-duplex 1000
```

Syntax: `speed-duplex value`

where `value` can be one of the following:

- 10-full – 10 Mbps, full duplex

2 Configuring basic port parameters

- 10-half – 10 Mbps, half duplex
- 100-full – 100 Mbps, full duplex
- 100-half – 100 Mbps, half duplex
- 1000 – 1 Gbps, full duplex (supported on PowerConnect B-Series TI24X 10-GbE ports only)
- 1000-full-master – 1 Gbps, full duplex master (not supported on the PowerConnect B-Series TI24X)
- 1000-full-slave – 1 Gbps, full duplex slave (not supported on the PowerConnect B-Series TI24X)
- 10000 – 10 Gbps, full duplex (supported on PowerConnect B-Series TI24X 10-GbE ports only)
- auto – auto-negotiation

The default for copper ports is **auto** (auto-negotiation).

The default for fiber ports on the PowerConnect B-Series TI24X is 10000 (10 Gbps, full duplex).

Use the **no** form of the command to restore the default.

NOTE

On PowerConnect B-Series TI24X devices, when 10/100/1000 copper ports (ports 25 – 28) auto-negotiate to either 1 Gbps or 100 Mbps, the green and amber LEDs will be lit solid (ON) when the link is up, and the amber LED will blink when traffic flows through the port. On PowerConnect B-Series TI24X devices, if the speed is set to Auto for a 1G port, the port auto-negotiates the flow control with the neighboring port.

Auto speed detect

On PowerConnect B-Series TI24X devices, if you insert a 1G SFP, the device detects the media change and automatically change the speed to support 1G for that port. This happens when the configured speed is 10G. The configured speed continues to be 10G, but the port comes up with operational speed of 1G. This removes the need for explicitly configuring speed-duplex 1000 for SFPs where the device is able to detect the media type.

NOTE

All the ports with 1G SFPs which need to form a trunk (static or dynamic), need to use either the Auto speed detect feature to come up in 1G mode or use the **speed-duplex 1000** command. Configuring **speed-duplex 1000** on only few of the ports to be part of the trunk will prevent trunk creation.

Modifying port duplex mode

You can manually configure a 10/100 Mbps port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic.

NOTE

You can modify the port duplex mode of copper ports only. This feature does not apply to fiber ports.

Port duplex mode and port speed are modified by the same command.

Configuration syntax

To change the port speed of interface 8 from the default of 10/100/1000 auto-sense to 10 Mbps operating at full-duplex, enter the following.


```
PowerConnect(config)# interface e 8
PowerConnect(config-if-e10000-8)# speed-duplex 10-full
```

Syntax: `speed-duplex value`

The *value* can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto (default)

Disabling or re-enabling a port

A port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is enabled.

To disable port 8 of a device, enter the following.

```
PowerConnect(config)# interface e 8
PowerConnect(config-if-e10000-8)# disable
```

Syntax: `disable`

You also can disable or re-enable a virtual interface. To do so, enter commands such as the following.

```
PowerConnect(config)# interface ve v1
PowerConnect(config-vif-1)# disable
```

Syntax: `disable`

To re-enable a virtual interface, enter the **enable** command at the Interface configuration level. For example, to re-enable virtual interface v1, enter the following command.

```
PowerConnect(config-vif-1)#enable
```

Syntax: `enable`

Disabling or re-enabling flow control

You can configure full-duplex ports on a system to operate with or without flow control (802.3x). Flow control is enabled by default.

To disable flow control on full-duplex ports on a system, enter the following.

```
PowerConnect(config)# no flow-control
```

To turn the feature back on.

```
PowerConnect(config)# flow-control
```

Syntax: `[no] flow-control`

NOTE

For optimal link operation, link ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), and Flow Control.

Auto-negotiation and advertisement of flow control

Auto-negotiation of flow control can be enabled and advertised for 10/100/1000M ports. To enable and advertise flow control capability, enter the following commands.

```
PowerConnect(config)# interface ethernet 21
PowerConnect(config-if-e10000-21)# flow-control
```

To also enable auto-negotiation of flow control, enter the following commands.

```
PowerConnect(config)# interface ethernet 21
PowerConnect(config-if-e10000-21)# flow-control neg-on
```

Syntax: `#[no] flow-control [neg-on]`

- **flow-control** [default] - Enable flow control, advertise flow control and disable negotiation of flow control
- **flow-control neg-on** - Advertise flow control and enable negotiation of flow control
- **no flow-control** - Disable flow control, disable advertising flow control and also disable negotiation of flow control

Commands may be entered in IF (single port) or MIF (multiple ports at once) mode.

```
PowerConnect(config)# interface ethernet 21
PowerConnect(config-if-e10000-21)# flow-control
```

This command enables flow-control on port 21.

```
PowerConnect(config)# interface e 11 to 15
PowerConnect(config-mif-11-15)# flow-control
```

This command enables flow-control on ports 11 to 15.

Asynchronous flow control

The PowerConnect B-Series TI24X devices supports asynchronous flow control. By default, PAUSE frames are honored when you stops or starts sending traffic to a switch when another switch sends a PAUSE frame.

Flow control supports honor-only mode and generate-only mode. Honor-only is the default mode. To change the default mode to generate-only mode use the **flow-control generate-only** command.

To set the mode to support both honor and generate modes use the **flow-control both** command. To set the mode to default, use the **flow-control honor-only** command.

Displaying flow-control status

The **show interface port** command displays configuration, operation, and negotiation status where applicable.

NOTE

For 10 Gbps ports, the display shows Flow Control is enabled, or Flow Control is disabled, depending on the configuration.

NOTE

Auto-negotiation of flow control is not supported on 10 Gbps ports and copper/fiber combination ports.

NOTE

When any of the commands are applied to a port that is up, the port will be disabled and re-enabled.

NOTE

When flow-control is enabled, the hardware can only advertise Pause. It does not advertise Asym.

Configuring the Interpacket Gap (IPG)

IPG is the time delay, in bit time, between frames transmitted by the device. You configure IPG at the interface level. The command you use depends on the interface type on which IPG is being configured.

The default interpacket gap is 96 bits-time, which is 9.6 microseconds for 10 Mbps Ethernet, 960 nanoseconds for 100 Mbps Ethernet, 96 nanoseconds for 1 Gbps Ethernet, and 9.6 nanoseconds for 10 Gbps Ethernet.

Configuration notes

- When you enter a value for IPG, the device applies the closest valid IPG value for the port mode to the interface. For example, if you specify 120 for a 1 Gbps Ethernet port in 1 Gbps mode, the device assigns 112 as the closest valid IPG value to program into hardware.

Configuring IPG on a Gbps Ethernet port

On a Gbps Ethernet port, you can configure IPG for 10/100 mode and for Gbps Ethernet mode.

10/100M mode

To configure IPG on a Gbps Ethernet port for 10/100M mode, enter the following command.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipg-mii 120
IPG 120(120) has been successfully configured for ports 1 to 12
```

Syntax: [no] ipg-mii <bit time>

Enter 12-124 for <bit time>. The default is 96 bit time.

1G mode

To configure IPG on a Gbps Ethernet port for 1-Gbps Ethernet mode, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipg-gmii 120
IPG 120(112) has been successfully configured for ports 1 to 12
```

Syntax: [no] ipg-gmii bit time

Enter 48 - 112 for bit time. The default is 96 bit time.

Configuring IPG on a 10 Gbps Ethernet interface

To configure IPG on a 10 Gbps Ethernet interface, enter commands such as the following.

2 Configuring basic port parameters

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipg-xgmii 120
IPG 120(128) has been successfully configured for port 1
```

Syntax: [no] ipg-xgmii *bit time*

Enter 96-192 for *bit time*. The default is 96 bit time.

Changing the Gbps fiber negotiation mode

The globally configured Gbps negotiation mode is the default mode for all Gbps fiber ports. You can override the globally configured default and set individual ports to the following:

- **Negotiate-full-auto** – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default.
- **Auto-Gbps** – The port tries to perform a handshake with the other port to exchange capability information.
- **Negotiation-off** – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

To change the mode for individual ports, enter commands such as the following.

```
PowerConnect(config)# int ethernet 1 to 4
PowerConnect(config-mif-1-4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gbps for ports 1 – 4.

Syntax: gig-default neg-full-auto | auto-gig | neg-off

NOTE

When Gbps negotiation mode is turned off (CLI command **gig-default neg-off**), the device may inadvertently take down both ends of a link. This is a hardware limitation for which there is currently no workaround.

Modifying port priority (QoS)

You can give preference to the inbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. For information and procedures, refer to [Chapter 15, “Configuring Quality of Service”](#).

Configuring port flap dampening

Port Flap Dampening increases the resilience and availability of the network by limiting the number of port state transitions on an interface.

If the port link state toggles from up to down for a specified number of times within a specified period, the interface is physically disabled for the specified wait period. Once the wait period expires, the port link state is re-enabled. However, if the wait period is set to zero (0) seconds, the port link state will remain disabled until it is manually re-enabled.

Configuration notes

- When a flap dampening port becomes a member of a trunk group, that port, as well as all other member ports of that trunk group, will inherit the primary port configuration. This means that the member ports will inherit the primary port flap dampening configuration, regardless of any previous configuration.
- The device counts the number of times a port link state toggles from "up to down", and not from "down to up".
- The sampling time or window (the time during which the specified toggle threshold can occur before the wait period is activated) is triggered when the first "up to down" transition occurs.
- "Up to down" transitions include UDLD-based toggles, as well as the physical link state.

Configuring port flap dampening on an interface

This feature is configured at the interface level.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# link-error-disable 10 3 10
```

Syntax: [no] **link-error-disable** *toggle-threshold sampling-time-in-sec wait-time-in-sec*

The *toggle-threshold* is the number of times a port link state goes from up to down and down to up before the wait period is activated. The default is 0. Enter a valid value range from 1-50.

The *sampling-time-in-sec* is the amount of time during which the specified toggle threshold can occur before the wait period is activated. The default is 0 seconds. Enter 0 – 65535 seconds.

The *wait-time-in-sec* is the amount of time the port remains disabled (down) before it becomes enabled. Entering 0 – 65535 seconds; 0 indicates that the port will stay down until an administrative override occurs.

Configuring port flap dampening on a trunk

You can configure the port flap dampening feature on the primary port of a trunk using the **link-error-disable** command. Once configured on the primary port, the feature is enabled on all ports that are members of the trunk. You cannot configure port flap dampening on port members of the trunk.

Enter commands such as the following on the primary port of a trunk.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# link-error-disable 10 3 10
```

Re-enabling a port disabled by port flap dampening

A port disabled by port flap dampening is automatically re-enabled once the wait period expires; however, if the wait period is set to zero (0) seconds, you must re-enable the port by entering the following command on the disabled port.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# no link-error-disable 10 3 10
```

Displaying ports configured with port flap dampening

Ports that have been disabled due to the port flap dampening feature are identified in the output of the **show link-error-disable** command. The following shows an example output.

```
PowerConnect# show link-error-disable
Port 1 is forced down by link-error-disable.
```

Use the **show link-error-disable all** command to display the ports with the port flap dampening feature enabled.

For PowerConnect B-Series TI24X devices, the output of the command shows the following.

```
PowerConnect# show link-error-disable all
Port      -----Config-----      -----Oper-----
#      Threshold  Sampling-Time  Shutoff-Time  State  Counter
-----
  11           3           120           600      Idle   N/A
  12           3           120           500      Down   424
```

[Table 7](#) defines the port flap dampening statistics displayed by the **show link-error-disable all** command.

TABLE 7 Output of show link-error-disable

This column...	Displays...
Port #	The port number.
Threshold	The number of times the port link state will go from up to down and down to up before the wait period is activated.
Sampling-Time	The number of seconds during which the specified toggle threshold can occur before the wait period is activated.
Shutoff-Time	The number of seconds the port will remain disabled (down) before it becomes enabled. A zero (0) indicates that the port will stay down until an administrative override occurs.
State	The port state can be one of the following: <ul style="list-style-type: none"> • Idle – The link is normal and no link state toggles have been detected or sampled. • Down – The port is disabled because the number of sampled errors exceeded the configured threshold. • Err – The port sampled one or more errors.
Counter	<ul style="list-style-type: none"> • If the port state is Idle, this field displays N/A. • If the port state is Down, this field shows the remaining value of the shutoff timer. • If the port state is Err, this field shows the number of errors sampled.

Syntax: **show link-error-disable [all]**

Also, in PowerConnect B-Series TI24X devices, the **show interface** command indicates if the port flap dampening feature is enabled on the port.

```
PowerConnect# show interface ethernet 15
GigabitEthernet15 is up, line protocol is up
  Link Error Dampening is Enabled
  Hardware is GigabitEthernet, address is 00e0.5200.010e (bia 00e0.5200.010e)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX

PowerConnect# show interface ethernet 17
GigabitEthernet17 is ERR-DISABLED, line protocol is down
  Link Error Dampening is Enabled
  Hardware is GigabitEthernet, address is 00e0.5200.010e (bia 00e0.5200.010e)
  Configured speed auto, actual unknown, configured duplex fdx, actual unknown
```

The line “Link Error Dampening” displays “Enabled” if port flap dampening is enabled on the port or “Disabled” if the feature is disabled on the port. The feature is enabled on the ports in the two examples above. Also, the characters “ERR-DISABLED” is displayed for the “GbpsEthernet” line if the port is disabled because of link errors.

Syntax: `show interface ethernet <port-number>`

In addition to the show commands above, the output of the `show interface brief` command for PowerConnect B-Series T124X devices, indicates if a port is down due to link errors.

```
PowerConnect# show interface brief e17

Port  Link      State      Dupl Speed Trunk Tag Priori MAC          Name
17    ERR-DIS None      None None  15    Yes level0 00e0.5200.010e
```

The ERR-DIS entry under the “Link” column indicates the port is down due to link errors.

Port loop detection

This feature allows the device to disable a port that is on the receiving end of a loop by sending test packets. You can configure the time period during which test packets are sent.

Strict mode and loose mode

There are two types of loop detection; Strict Mode and Loose Mode. In Strict Mode, a port is disabled only if a packet is looped back to that same port. Strict Mode overcomes specific hardware issues where packets are echoed back to the input port. In Strict Mode, loop detection must be configured on the physical port.

In Loose Mode, loop detection is configured on the VLAN of the receiving port. Loose Mode disables the receiving port if packets originate from any port or VLAN on the same device. The VLAN of the receiving port must be configured for loop detection in order to disable the port.

Recovering disabled ports

Once a loop is detected on a port, it is placed in Err-Disable state. The port will remain disabled until one of the following occurs:

- You manually disable and enable the port at the Interface Level of the CLI.
- You enter the command `clear loop-detection`. This command clears loop detection statistics and enables all Err-Disabled ports.

2 Configuring basic port parameters

- The device automatically re-enables the port. To set your device to automatically re-enable Err-Disabled ports, refer to [“Configuring the device to automatically re-enable ports”](#) on page 35.

Configuration notes

- Loopback detection packets are sent and received on both tagged and untagged ports. Therefore, this feature cannot be used to detect a loop across separate devices.
- On PowerConnect devices, the port loop detection feature works only on untagged ports.

The following information applies to Loose Mode loop detection:

- With Loose Mode, two ports of a loop are disabled.
- Different VLANs may disable different ports. A disabled port affects every VLAN using it.
- Loose Mode floods test packets to the entire VLAN. This can impact system performance if too many VLANs are configured for Loose Mode loop detection.

NOTE

Dell recommends that you limit the use of Loose Mode. If you have a large number of VLANs, configuring loop detection on all of them can significantly affect system performance because of the flooding of test packets to all configured VLANs. An alternative to configuring loop detection in a VLAN-group of many VLANs is to configure a separate VLAN with the same tagged port and configuration, and enable loop detection on this VLAN only.

NOTE

When loop detection is used with L2 loop prevention protocols, such as spanning tree (STP), the L2 protocol takes higher priority. Loop detection cannot send or receive probe packets if ports are blocked by L2 protocols, so it does not detect L2 loops when STP is running because loops within a VLAN have been prevented by STP. Loop detection running in Loose Mode can detect and break L3 loops because STP cannot prevent loops across different VLANs. In these instances, the ports are not blocked and loop detection is able to send out probe packets in one VLAN and receive packets in another VLAN. In this way, loop detection running in Loose Mode disables both ingress and egress ports.

Enabling loop detection

Use the **loop-detection** command to enable loop detection on a physical port (Strict Mode) or a VLAN (Loose Mode). Loop detection is disabled by default. The following example shows a Strict Mode configuration.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# loop-detection
```

The following example shows a Loose Mode configuration.

```
PowerConnect(config)# vlan20
PowerConnect(config-vlan-20)# loop-detection
```

By default, the port will send test packets every one second, or the number of seconds specified by the **loop-detection-interval** command. Refer to [“Configuring a global loop detection interval”](#) on page 35.

Syntax: [no] loop-detection

Use the [no] form of the command to disable loop detection.

Configuring a global loop detection interval

The loop detection interval specifies how often a test packet is sent on a port. When loop detection is enabled, the loop detection time unit is 0.1 second, with a default of 10 (one second). The range is from 1 (one tenth of a second) to 100 (10 seconds). You can use the **show loop-detection status** command to view the loop detection interval.

To configure the global loop detection interval, enter a command similar to the following.

```
PowerConnect(config)# loop-detection-interval 50
```

This command sets the loop-detection interval to 5 seconds (50 x 0.1).

To revert to the default global loop detection interval of 10, enter one of the following.

```
PowerConnect(config)# loop-detection-interval 10
```

OR

```
PowerConnect(config)# no loop-detection-interval 50
```

Syntax: [no] loop-detection-interval *number*

where *number* is a value from 1 to 100. The system multiplies your entry by 0.1 to calculate the interval at which test packets will be sent.

Configuring the device to automatically re-enable ports

To configure the device to automatically re-enable ports that were disabled because of a loop detection, enter the following command.

```
PowerConnect(config)# errdisable recovery cause loop-detection
```

The above command will cause the device to automatically re-enable ports that were disabled because of a loop detection. By default, the device will wait 300 seconds before re-enabling the ports. You can optionally change this interval to a value from 10 to 65535 seconds. Refer to [“Specifying the recovery time interval”](#) on page 35.

Syntax: [no] errdisable recovery cause loop-detection

Use the [no] form of the command to disable this feature.

Specifying the recovery time interval

The recovery time interval specifies the number of seconds the device will wait before automatically re-enabling ports that were disabled because of a loop detection. (Refer to [“Configuring the device to automatically re-enable ports”](#) on page 35.) By default, the device will wait 300 seconds. To change the recovery time interval, enter a command such as the following.

```
PowerConnect(config)# errdisable recovery interval 120
```

This command configures the device to wait 120 seconds (2 minutes) before re-enabling the ports.

To revert to the default recovery time interval of 300 seconds (5 minutes), enter one of the following commands.

```
PowerConnect(config)# errdisable recovery interval 300
```

OR

```
PowerConnect(config)# no errdisable recovery interval 120
```

2 Configuring basic port parameters

Syntax: [no] errdisable recovery interval seconds

where seconds is a number from 10 to 65535.

Clearing loop-detection

To clear loop detection statistics and re-enable all ports that are in Err-Disable state because of a loop detection, enter the following command.

```
PowerConnect# clear loop-detection
```

Displaying loop-detection information

Use the **show loop-detection status** command to display loop detection status, as shown.

```
PowerConnect# show loop-detection status
loop detection packets interval: 10 (unit 0.1 sec)
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index port/vlan  status                                     #errdis  sent-pkts  recv-pkts
1          13      untag, LEARNING                                         0          0          0
2          15      untag, BLOCKING                                         0          0          0
3          17      untag, DISABLED                                         0          0          0
4          18      ERR-DISABLE by itself                                   1          6          1
5          19      ERR-DISABLE by vlan 12                                 0          0          0
6      vlan12  2 ERR-DISABLE ports                                    2         24          2
```

If a port is errdisabled in Strict mode, it shows "ERR-DISABLE by itself". If it is errdisabled due to its associated vlan, it shows "ERR-DISABLE by vlan ?"

The following command displays the current disabled ports, including the cause and the time.

```
PowerConnect# show loop-detection disable
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index  port          caused-by      disabled-time
1      18             itself         00:13:30
2      19             vlan 12        00:13:30
3      20             vlan 12        00:13:30
```

This example shows the disabled ports, the cause, and the time the port was disabled. If loop-detection is configured on a physical port, the disable cause will show "itself". For VLANs configured for loop-detection, the cause will be a VLAN.

The following command shows the hardware and software resources being used by the loop-detection feature.

```
Vlans configured loop-detection use 1 HW MAC
Vlans not configured but use HW MAC: 1 10

          alloc in-use  avail get-fail    limit  get-mem  size init
configuration pool      16     6     10     0     3712     6    15    16
linklist pool           16    10     6     0     3712    10    16    16
```

Syslog message

The following message is logged when a port is disabled due to loop detection. This message also appears on the console.

```
loop-detect: port ?\?\? vlan ?, into errdisable state
```

The Errdisable function logs a message whenever it re-enables a port.

2 Configuring basic port parameters

Operations, Administration, and Maintenance

Overview

For easy software image management, all devices support the download and upload of software images between the flash modules on the devices and a Trivial File Transfer Protocol (TFTP) server on the network.

PowerConnect devices have two flash memory modules:

- **Primary flash** – The default local storage device for image files and configuration files.
- **Secondary flash** – A second flash storage device. You can use the secondary flash to store redundant images for additional reload reliability or to preserve one software image while testing another one.

Only one flash device is active at a time. By default, the primary image will become active upon reload.

You can update the software contained on a flash module using TFTP to copy the update image from a TFTP server onto the flash module. In addition, you can copy software images and configuration files from a flash module to a TFTP server.

NOTE

PowerConnect devices are TFTP clients but not TFTP servers. You must perform the TFTP transaction from the device. You cannot “put” a file onto the device using the interface of your TFTP server.

NOTE

If you are attempting to transfer a file using TFTP but have received an error message, refer to [“Diagnostic error codes and remedies for TFTP transfers”](#) on page 51.

Determining the software versions installed and running on a device

Use the following methods to display the software versions running on the device and the versions installed in flash memory.

Determining the flash image version running on the device

To determine the flash image version running on a device, enter the **show version** command at any level of the CLI. Some examples are shown below.

3 Determining the software versions installed and running on a device

Compact devices

To determine the flash image version running on a Compact device, enter the **show version** command at any level of the CLI. The following shows an example output.

```
PowerConnect#show version
  SW: Version 4.2.00b Copyright (c) 1996-2010 Brocade Communications Systems,
  Inc.
    Compiled on Dec 02 2010 at 08:07:06 labeled as TIR04200b
    (6092645 bytes) from Secondary TIR04200b
    Compressed Boot-Monitor Image size = 373767, Version:04.1.00T205 (grz04100)
  HW: Stackable TurboIron-X24
=====
  Serial #: BFF2342E00X
  P-ASIC 0: type B820, rev 01 subrev 00
=====
  833 MHz Power PC processor 8541 (version 32/0020) 66 MHz bus
  512 KB boot flash memory
  31744 KB code flash memory
  512 MB DRAM
The system uptime is 5 minutes 34 seconds
The system : started=warm start reloaded=by "reload"
```

The version information is shown in bold type in this example:

- “4.2.00b” indicates the flash code version number.
- “labeled as TIR04200b” indicates the flash code image label. The label indicates the imagetype and version and is especially useful if you change the image file name.
- “Secondary TIR04200b” indicates the flash code image file name that was loaded.

Determining the image versions installed in flash memory

Enter the **show flash** command to display the boot and flash images installed on the device.

- The “Compressed Pri Code size” line lists the flash code version installed in the primary flash area.
- The “Compressed Sec Code size” line lists the flash code version installed in the secondary flash area.
- The “Boot Monitor Image size” line lists the boot code version installed in flash memory. The device does not have separate primary and secondary flash areas for the boot image. The flash memory module contains only one boot image.

Flash image verification

The Flash Image Verification feature allows you to verify boot images based on hash codes, and to generate hash codes where needed. This feature lets you select from three data integrity verification algorithms:

- **MD5** - Message Digest algorithm (RFC 1321)
- **SHA1** - US Secure Hash Algorithm (RFC 3174)
- **CRC** - Cyclic Redundancy Checksum algorithm

CLI commands

Use the following command syntax to verify the flash image:

Syntax: `verify md5 | sha1 | crc32 ASCII string | primary | secondary [hash code]`

- **md5** – Generates a 16-byte hash code
- **sha1** – Generates a 20-byte hash code
- **crc32** – Generates a 4 byte checksum
- **ascii string** – A valid image filename
- **primary** – The primary boot image (primary.img)
- **secondary** – The secondary boot image (secondary.img)
- **hash code** – The hash code to verify

The following examples show how the **verify** command can be used in a variety of circumstances.

To generate an MD5 hash value for the secondary image, enter the following command.

```
PowerConnect# verify md5 secondary
PowerConnect#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
```

To generate a SHA-1 hash value for the secondary image, enter the following command.

```
PowerConnect# verify sha secondary
PowerConnect#.....Done
Size = 2044830, SHA1 49d12d26552072337f7f5fcaef4cf4b742a9f525
```

To generate a CRC32 hash value for the secondary image, enter the following command.

```
PowerConnect# verify crc32 secondary
PowerConnect#.....Done
Size = 2044830, CRC32 b31fcbc0
```

To verify the hash value of a secondary image with a known value, enter the following commands.

```
PowerConnect# verify md5 secondary 01c410d6d153189a4a5d36c955653861
PowerConnect#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
Verification FAILED.
```

In the previous example, the codes did not match, and verification failed. If verification succeeds, the output will look like this.

```
PowerConnect# verify md5 secondary 01c410d6d153189a4a5d36c955653861
PowerConnect#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653861
Verification SUCCEEDED.
```

The following examples show this process for SHA-1 and CRC32 algorithms.

```
PowerConnect# verify sha secondary 49d12d26552072337f7f5fcaef4cf4b742a9f525
PowerConnect#.....Done
Size = 2044830, sha 49d12d26552072337f7f5fcaef4cf4b742a9f525
Verification SUCCEEDED.
```

and

```
PowerConnect# verify crc32 secondary b31fcbc0
PowerConnect#.....Done
Size = 2044830, CRC32 b31fcbc0
Verification SUCCEEDED.
```

Image file types

This section lists the boot and flash image file types supported on the PowerConnect B-Series TI24X of switches and how to install them. For information about a specific version of code, refer to the release notes.

TABLE 8 Software image files

Product	Boot image ¹	Flash image
PowerConnect B-Series TI24X	GRZxxxxx.bin	TISxxxxx.bin (Layer 2)

Upgrading software

Use the following procedures to upgrade the software.

Upgrading the boot code

Follow the steps given below to upgrade the boot code.

1. Place the new boot code on a TFTP server to which the device has access.
2. Enter the following command at the Privileged EXEC level of the CLI to copy the boot code from the TFTP server into flash memory:
 - **copy tftp flash ip-addr image-file-name bootrom**

NOTE

Use the **copy tftp flash** command to copy the boot code to the device only during a maintenance window. Attempting to do so during normal networking operations can cause disruption to the network.

3. Verify that the code has been successfully copied by entering the following command at any level of the CLI:
 - **show flash**

The output will display the compressed boot ROM code size and the boot code version.
4. Upgrade the flash code as instructed in the following section.

Upgrading the flash code

Follow the steps given below to upgrade the flash code.

1. Place the new flash code on a TFTP server to which the device has access.
2. Enter the following command at the Privileged EXEC level of the CLI to copy the flash code from the TFTP server into the flash memory.

copy tftp flash ip-addr image-file-name primary | secondary

3. Verify that the flash code has been successfully copied by entering the following command at any level of the CLI.

show flash

4. If the flash code version is correct, go to [step 5](#). Otherwise, go to [step 1](#).
5. Reload the software by entering one of the following commands:
 - **reload** (this command boots from the default boot source, which is the primary flash area by default)
 - **hitless-reload primary | secondary**
 - **boot system flash primary | secondary**

The `boot system flash` process occurs after a `boot system flash primary/secondary` command is entered and gives an administrator the opportunity to make last minute changes or corrections before performing a reload. The example below shows the confirmation step.

```
PowerConnect# boot system flash primary
Are you sure? (enter 'Y' or 'N'): y
```

Boot code synchronization feature

When the new boot image is copied into the active module, it is automatically synchronized with the redundant management module.

NOTE

There is currently no option for manual synchronization of the boot image.

To activate the boot synchronization process, enter the following command.

```
PowerConnect# copy tftp flash 192.168.255.102 GRZ04100.bin bootrom
```

The system responds with the following message.

```
PowerConnect# Load to buffer (8192 bytes per dot)
.....Write to boot flash.....
TFTP to Flash Done.
PowerConnect# Synchronizing with standby module...
Boot image synchronization done.
```

Using SNMP to upgrade software

You can use a third-party SNMP management application to upgrade software on a device.

NOTE

Dell recommends that you make a backup copy of the startup-config file before you upgrade the software. If you need to run an older release, you will need to use the backup copy of the startup-config file.

1. Configure a read-write community string on the device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI.

```
snmp-server community string ro | rw
```

where *string* is the community string and can be up to 32 characters long.

2. On the device, enter the following command from the global CONFIG level of the CLI.

```
no snmp-server pw-check
```

3 Changing the block size for TFTP file transfers

This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a device, by default the device rejects the request.

Changing the block size for TFTP file transfers

When you use TFTP to copy a file to or from a device, the device transfers the data in blocks of 8192 bytes by default. You can change the block size to one of the following if needed:

- 4096
- 2048
- 1024
- 512
- 256
- 128
- 64
- 32
- 16

To change the block size for TFTP file transfers, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# flash 2047  
set flash copy block size to 2048
```

Syntax: [no] flash *num*

The software rounds up the *num* value you enter to the next valid power of two, and displays the resulting value. In this example, the software rounds the value up to 2048.

NOTE

If the value you enter is one of the valid powers of two for this parameter, the software still rounds the value up to the next valid power of two. Thus, if you enter 2048, the software rounds the value up to 4096.

Rebooting

You can use boot commands to immediately initiate software boots from a software image stored in primary or secondary flash on a device or from a BootP or TFTP server. You can test new versions of code on a device or choose the preferred boot source from the console boot prompt without requiring a system reset.

NOTE

It is very important that you verify a successful TFTP transfer of the boot code **before** you reset the system. If the boot code is not transferred successfully but you try to reset the system, the system will not have the boot code with which to successfully boot.

By default, the device first attempts to boot from the image stored in its primary flash, then its secondary flash, and then from a TFTP server. You can modify this booting sequence at the global CONFIG level of the CLI using the **boot system...** command.

To initiate an immediate boot from the CLI, enter one of the **boot system...** commands.

NOTE

If you are booting the device from a TFTP server through a fiber connection, use the following command: **boot system tftp ip-address filename fiber-port**.

Displaying the boot preference

Use the **show boot-preference** command to display the boot sequence in the startup config and running config files. The boot sequence displayed is also identified as either user-configured or the default.

The following example shows the default boot sequence preference.

```
PowerConnect# show boot-preference
Boot system preference (Configured):
  Use Default
Boot system preference(Default):
  Boot system flash primary
  Boot system flash secondary
```

The following example shows a user-configured boot sequence preference.

```
PowerConnect# show boot-preference
Boot system preference(Configured):
  Boot system flash secondary
  Boot system tftp 10.1.1.1 TIX04200b1.bin
  Boot system flash primary
Boot system preference (Default):
  Boot system flash primary
  Boot system flash secondary
```

Syntax: show boot-preference

The results of the **show run** command for the configured example above appear as follows.

```
PowerConnect# show run
Current Configuration:
!
boot sys fl sec
boot sys df 10.1.1.1 TIX04200b1.bin
boot sys fl pri
ip address 10.1.1.4 255.255.255.0
snmp-client 10.1.1.1
!
end
```

Loading and saving configuration files

For easy configuration management, all devices support both the download and upload of configuration files between the devices and a TFTP server on the network.

3 Loading and saving configuration files

You can upload either the startup configuration file or the running configuration file to the TFTP server for backup and use in booting the system:

- **Startup configuration file** – This file contains the configuration information that is currently saved in flash. To display this file, enter the **show configuration** command at any CLI prompt.
- **Running configuration file** – This file contains the configuration active in the system RAM but not yet saved to flash. These changes could represent a short-term requirement or general configuration change. To display this file, enter the **show running-config** or **write terminal** command at any CLI prompt.

Each device can have one startup configuration file and one running configuration file. The startup configuration file is shared by both flash modules. The running configuration file resides in DRAM.

When you load the startup-config file, the CLI parses the file three times.

1. During the first pass, the parser searches for **system-max** commands. A **system-max** command changes the size of statically configured memory.
2. During the second pass, the parser implements the **system-max** commands if present and also implements trunk configuration commands (**trunk** command) if present.
3. During the third pass, the parser implements the remaining commands.

Replacing the startup configuration with the running configuration

After you make configuration changes to the active system, you can save those changes by writing them to flash memory. When you write configuration changes to flash memory, you replace the startup configuration with the running configuration.

To replace the startup configuration with the running configuration, enter the following command at any Enable or CONFIG command prompt.

```
PowerConnect# write memory
```

Replacing the running configuration with the startup configuration

If you want to back out of the changes you have made to the running configuration and return to the startup configuration, enter the following command at the Privileged EXEC level of the CLI.

```
PowerConnect# reload
```

Logging changes to the startup-config file

You can configure a device to generate a Syslog message when the startup-config file is changed. The trap is enabled by default.

The following Syslog message is generated when the startup-config file is changed.

```
startup-config was changed
```

If the startup-config file was modified by a valid user, the following Syslog message is generated.

```
startup-config was changed by <username>
```

To disable or re-enable Syslog messages when the startup-config file is changed, use the following command.

Syntax: [no] logging enable config-changed

Copying a configuration file to or from a TFTP server

To copy the startup-config or running-config file to or from a TFTP server, use one of the following methods.

NOTE

You can name the configuration file when you copy it to a TFTP server. However, when you copy a configuration file from the server to a device, the file is always copied as “startup-config” or “running-config”, depending on which type of file you saved to the server.

To initiate transfers of configuration files to or from a TFTP server using the CLI, enter one of the following commands:

- **copy startup-config tftp tftp-ip-addr filename** – Use this command to upload a copy of the startup configuration file from the device to a TFTP server.
- **copy running-config tftp tftp-ip-addr filename** – Use this command to upload a copy of the running configuration file from the device to a TFTP server.
- **copy tftp startup-config tftp-ip-addr filename** – Use this command to download a copy of the startup configuration file from a TFTP server to a device.

Dynamic configuration loading

You can load dynamic configuration commands (commands that do not require a reload to take effect) from a file on a TFTP server into the running-config on the device. You can make configuration changes off-line, then load the changes directly into the device running-config, without reloading the software.

Usage considerations

- Use this feature only to load configuration information that does not require a software reload to take effect. For example, you cannot use this feature to change statically configured memory (**system-max** command) or to enter trunk group configuration information into the running-config.
- Do not use this feature if you have deleted a trunk group but have not yet placed the changes into effect by saving the configuration and then reloading. When you delete a trunk group, the command to configure the trunk group is removed from the device running-config, but the trunk group remains active. To finish deleting a trunk group, save the configuration (to the startup-config file), then reload the software. After you reload the software, then you can load the configuration from the file.
- Do not load port configuration information for secondary ports in a trunk group. Since all ports in a trunk group use the port configuration settings of the primary port in the group, the software cannot implement the changes to the secondary port.

Preparing the configuration file

A configuration file that you create must follow the same syntax rules as the startup-config file the device creates.

- The configuration file is a script containing CLI configuration commands. The CLI reacts to each command entered from the file in the same way the CLI reacts to the command if you enter it. For example, if the command results in an error message or a change to the CLI configuration level, the software responds by displaying the message or changing the CLI level.
- The software retains the running-config that is currently on the device, and changes the running-config only by adding new commands from the configuration file. If the running config already contains a command that is also in the configuration file you are loading, the CLI rejects the new command as a duplicate and displays an error message. For example, if the running-config already contains a command that configures ACL 1, the software rejects ACL 1 in the configuration file, and displays a message that ACL 1 is already configured.
- The file can contain global CONFIG commands or configuration commands for interfaces, routing protocols, and so on. You cannot enter User EXEC or Privileged EXEC commands.
- The default CLI configuration level in a configuration file is the global CONFIG level. Thus, the first command in the file must be a global CONFIG command or “!”. The ! (exclamation point) character means “return to the global CONFIG level”.

NOTE

You can enter text following “!” as a comment. However, the “!” is not a comment marker. It returns the CLI to the global configuration level.

NOTE

If you copy-and-paste a configuration into a management session, the CLI ignores the “!” instead of changing the CLI to the global CONFIG level. As a result, you might get different results if you copy-and-paste a configuration instead of loading the configuration using TFTP.

- Make sure you enter each command at the correct CLI level. Since some commands have identical forms at both the global CONFIG level and individual configuration levels, if the CLI response to the configuration file results in the CLI entering a configuration level you did not intend, then you can get unexpected results.

For example, if a trunk group is active on the device, and the configuration file contains a command to disable STP on one of the secondary ports in the trunk group, the CLI rejects the commands to enter the interface configuration level for the port and moves on to the next command in the file you are loading. If the next command is a spanning-tree command whose syntax is valid at the global CONFIG level as well as the interface configuration level, then the software applies the command globally. Here is an example.

The configuration file contains these commands.

```
interface ethernet 2
no spanning-tree
```

The CLI responds like this.

```
PowerConnect(config)# interface ethernet 2
Error - cannot configure secondary ports of a trunk
PowerConnect(config)# no spanning-tree
PowerConnect(config)#
```

- If the file contains commands that must be entered in a specific order, the commands must appear in the file in the required order. For example, if you want to use the file to replace an IP address on an interface, you must first remove the old address using “no” in front of the **ip address** command, then add the new address. Otherwise, the CLI displays an error message and does not implement the command. Here is an example.

The configuration file contains these commands.

```
interface ethernet 11
ip address 10.10.10.69/24
```

The running-config already has a command to add an address to port 11, so the CLI responds like this.

```
PowerConnect(config)# interface ethernet 11
PowerConnect(config-if-e10000-11)# ip add 10.10.10.69/24
Error: can only assign one primary ip address per subnet
PowerConnect(config-if-e10000-11)#
```

To successfully replace the address, enter commands into the file as follows.

```
interface ethernet 11
no ip address 20.20.20.69/24
ip address 10.10.10.69/24
```

This time, the CLI accepts the command, and no error message is displayed.

```
PowerConnect(config)# interface ethernet 11
PowerConnect(config-if-e10000-11)# no ip add 20.20.20.69/24
PowerConnect(config-if-e10000-11)# ip add 10.10.10.69/24
PowerConnect(config-if-e10000-11)
```

- Always use the **end** command at the end of the file. The **end** command must appear on the last line of the file, by itself.

Loading the configuration information into the running-config

To load the file from a TFTP server, use either of the following commands:

- **copy tftp running-config ip-addr filename**
- **ncopy tftp ip-addr filename running-config**

NOTE

If you are loading a configuration file that uses a truncated form of the CLI command **access-list**, the software will not go into batch mode.

For example, the following command line *will initiate* batch mode.

```
access-list 131 permit host pc1 host pc2
```

The following command line *will not* initiate batch mode.

```
acc 131 permit host pc1 host pc2
```

Maximum file sizes for startup-config file and running-config

Each device has a maximum allowable size for the running-config and the startup-config file. If you use TFTP to load additional information into a device running-config or startup-config file, it is possible to exceed the maximum allowable size. If this occurs, you will not be able to save the configuration changes.

The maximum size for the running-config and the startup-config file is 64K each.

To determine the size of a running-config or startup-config file, copy it to a TFTP server, then use the directory services on the server to list the size of the copied file. To copy the running-config or startup-config file to a TFTP server, use one of the following commands:

- Commands to copy the running-config to a TFTP server:
 - **copy running-config tftp** *ip-addr filename*
 - **ncopy running-config tftp** *ip-addr from-name*
- Commands to copy the startup-config file to a TFTP server:
 - **copy startup-config tftp** *ip-addr filename*
 - **ncopy startup-config tftp** *ip-addr from-name*

Scheduling a system reload

In addition to reloading the system manually, you can configure the device to reload itself at a specific time or after a specific amount of time has passed.

NOTE

The scheduled reload feature requires the system clock. You can use a Simple Network Time Protocol (SNTP) server to set the clock or you can set the device clock manually. Refer to [“Specifying a Simple Network Time Protocol \(SNTP\) server”](#) on page 18 or [“Setting the system clock”](#) on page 19.

Reloading at a specific time

To schedule a system reload for a specific time, use the **reload at** command. For example, to schedule a system reload from the primary flash module for 6:00:00 AM, April 1, 2003, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect# reload at 06:00:00 04-01-03
```

Syntax: **reload at** *hh:mm:ss mm-dd-yy* [**primary** | **secondary**]

- *hh:mm:ss* is hours, minutes, and seconds.
- *mm-dd-yy* is month, day, and year.
- **primary** | **secondary** specifies whether the reload is to occur from the primary code flash module or the secondary code flash module. The default is **primary**.

Reloading after a specific amount of time

To schedule a system reload to occur after a specific amount of time has passed on the system clock, use **reload after** command. For example, to schedule a system reload from the secondary flash one day and 12 hours later, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect# reload after 01:12:00 secondary
```

Syntax: **reload after** *dd:hh:mm* [**primary** | **secondary**]

- *dd:hh:mm* is the number of days, hours, and minutes.
- **primary** | **secondary** specifies whether the reload is to occur from the primary code flash module or the secondary code flash module.

Displaying the amount of time remaining before a scheduled reload

To display how much time is remaining before a scheduled system reload, enter the following command from any level of the CLI.

```
PowerConnect# show reload
```

Canceling a scheduled reload

To cancel a scheduled system reload using the CLI, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect# reload cancel
```

Diagnostic error codes and remedies for TFTP transfers

If an error occurs with a TFTP transfer to or from a device one of the following error codes is displayed on the console.

Table 0.1:

Error code	Message	Explanation and action
1	Flash read preparation failed.	A flash error occurred during the download.
2	Flash read failed.	Retry the download. If it fails again, contact customer support.
3	Flash write preparation failed.	
4	Flash write failed.	
5	TFTP session timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.

3 Diagnostic error codes and remedies for TFTP transfers

Table 0.1:

Error code	Message	Explanation and action
6	TFTP out of buffer space.	The file is larger than the amount of room on the device or TFTP server. If you are copying an image file to flash, first copy the other image to your TFTP server, then delete it from flash. (Use the erase flash... CLI command at the Privileged EXEC level to erase the image in the flash.) If you are copying a configuration file to flash, edit the file to remove unneeded information, then try again.
7	TFTP busy, only one TFTP session can be active.	Another TFTP transfer is active on another CLI session or Brocade Network Advisor session. Wait, then retry the transfer.
8	File type check failed.	You accidentally attempted to copy the incorrect image code into the system. Retry the transfer using the correct image.
16	TFTP remote - general error.	The TFTP configuration has an error. The specific error message describes the error.
17	TFTP remote - no such file.	Correct the error, then retry the transfer.
18	TFTP remote - access violation.	
19	TFTP remote - disk full.	
20	TFTP remote - illegal operation.	
21	TFTP remote - unknown transfer ID.	
22	TFTP remote - file already exists.	
23	TFTP remote - no such user.	

Monitoring Hardware Components

Hardware support

The procedures in this chapter describe how to configure the software to monitor hardware components. You can configure the software to monitor temperature and signal power levels for optical transceivers

[Table 9](#) lists which devices support the features discussed in this chapter.

TABLE 9 Hardware components monitoring support for devices

Feature	PowerConnect B-Series T124X
Digital optical monitoring	Yes

Digital optical monitoring

You can configure your device to monitor optical transceivers in the system, either globally or by specified ports. When this feature is enabled, the system will monitor the temperature and signal power levels for the optical transceivers in the specified ports. Console messages and syslog messages are sent when optical operating conditions fall below or rise above the XFP or SFP manufacturer recommended thresholds.

Supported media

Digital optical monitoring is supported with the following Dell-qualified media types:

- 1000Base-LHA
- 1000Base-LHB
- 1000Base-LX
- 1000Base-SX
- 10GBase-LR
- 10GBase-SR

Media not supported

Digital optical monitoring is not supported for the following optics:

- 1000Base-SX 2
- 1000Base-BX-D
- 1000Base-BX-U
- E1MG-100BXU

4 Digital optical monitoring

- E1MG-100BXD
- E1MG-BXU
- E1MG-BXD

Supported media

Digital optical monitoring is supported with the following Dell-qualified media types:

- 1000Base-BX-D
- 1000Base-BX-U
- 1000Base-LHA
- 1000Base-LHB
- 1000Base-LX
- 1000Base-SX
- 1000Base-SX 2

Media not supported

Digital optical monitoring is not supported for the following optics:

- E1MG-100BXU
- E1MG-100BXD
- E1MG-BXU
- E1MG-BXD

Configuration limitations

A device can monitor a maximum of 24 SFPs and 12 XFPs.

Enabling digital optical monitoring

To enable optical monitoring on all Dell-qualified optics installed in the device, use the following command.

```
PowerConnect(config)# optical-monitor
```

To enable optical monitoring on a specific port, use the following command.

```
PowerConnect(config)# interface ethernet 1  
PowerConnect(config-if-e10000-1)# optical-monitor
```

To enable optical monitoring on a range of ports, use the following command.

```
PowerConnect(config)# interface ethernet 1 to 12  
PowerConnect(config-mif-e10000-1-12)# optical-monitor
```

Syntax: [no] optical-monitor

Use the **no** form of the command to disable digital optical monitoring.

Setting the alarm interval

You can optionally change the interval between which alarms and warning messages are sent. The default interval is three minutes. To change the interval, use the following command.

```
PowerConnect(config)# interface ethernet 1 to 2
PowerConnect(config-mif-e10000-1-2)# optical-monitor 10
```

Syntax: [no] **optical-monitor** [*alarm-interval*]

For *alarm-interval*, enter a value between 1 and 65535. Enter 0 to disable alarms and warning messages.

NOTE

The commands **no optical-monitor** and **optical-monitor 0** perform the same function. That is, they both disable digital optical monitoring.

Displaying information about installed media

Use the **show media**, **show media slot**, and **show media ethernet** commands to obtain information about the media devices installed per device, per slot, and per port. The results displayed from these commands provide the Type, Vendor, Part number, Version and Serial number of the SFP or XFP optical device installed in the port. **1G M-C** indicates 1b Gbps copper media. If no SFP or XFP device is installed in a port, the "Type" field will display "EMPTY".

Use the **show media** command to obtain information about the media devices installed in a device.

```
PowerConnect# show media
Port 1: Type : 1G M-SX2(SFP)
        Vendor:   Brocade Communications, Inc. Version: 0000
        Part#:   TRPAG1XRPBSS-FY   Serial#: 0635000468
Port 2: Type : EMPTY
Port 3: Type : EMPTY
Port 4: Type : 100M M-FX-SR(SFP)
        Vendor:   Brocade Communications, Inc. Version: A
        Part#:   FTLF1217P2BTL-F1   Serial#: UCQ003A
Port 5: Type : 1G M-C
Port 6: Type : 1G M-C
Port 7: Type : 1G M-C
Port 8: Type : 1G M-C
Port 9: Type : 1G M-C
Port 10: Type : 1G M-C
Port 11: Type : 1G M-C
Port 12: Type : 1G M-C
Port 13: Type : 1G M-C
Port 14: Type : 1G M-C
Port 15: Type : 1G M-C
Port 16: Type : 1G M-C
Port 17: Type : 1G M-C
Port 18: Type : 1G M-C
Port 19: Type : 1G M-C
Port 20: Type : 1G M-C
Port 21: Type : 1G M-C
Port 22: Type : 1G M-C
Port 23: Type : 1G M-C
Port 24: Type : 1G M-C
```

4 Digital optical monitoring

```
Port 25: Type : 10G XG-SR(XFP)
        Vendor:   Brocade Communications Inc. Version: 02
        Part# :   JXPR01SW05306   Serial#: F617604000A3
Port 26: Type : EMPTY
```

Use the **show media slot** command to obtain information about the media device installed in a slot.

```
PowerConnect# show media slot 1
Port 1: Type : 1G M-SX(SFP)
        Vendor: Brocade Communications, Inc. Version:
        Part# : PL-XPL-VC-S13-19   Serial#: 425HC109
Port 2: Type : 1G M-SX(SFP)
        Vendor: Brocade Communications, Inc. Version:
        Part# : PL-XPL-VC-S13-19   Serial#: 411HC0AH
Port 3: Type : EMPTY
Port 4: Type : 1G M-SX(SFP)
        Vendor: FINISAR CORP.      Version: X1
        Part# : FTRJ-8519-3        Serial#: H11654K
Port 5: Type : EMPTY
Port 6: Type : EMPTY
Port 7: Type : 100M M-FX-IR(SFP)
        Vendor: Brocade Communications, Inc. Version: A
        Part# : FTLF1323P1BTR-FD   Serial#: UCT000T
Port 8: Type : EMPTY
Port 9: Type : 100M M-FX-LR(SFP)
        Vendor: Brocade Communications, Inc. Version: A
        Part# : FTLF1323P1BTL-FD   Serial#: UD3085J
Port 10: Type : EMPTY
Port 11: Type : 100M M-FX-SR(SFP)
        Vendor: Brocade Communications, Inc. Version: A
        Part# : FTLF1217P2BTL-F1   Serial#: UCQ003J
Port 12: Type : EMPTY
Port 13: Type : 100M M-FX-IR(SFP)
        Vendor: Brocade Communications, Inc. Version: A
        Part# : FTLF1323P1BTR-F1   Serial#: PCA2XC5
```

Use the **show media ethernet** command to obtain information about the media device installed in a port.

```
PowerConnect# show media e 17
Port 17: Type : 1G M-SX(SFP)
        Vendor: Brocade Communications, Inc. Version:
        Part# : PL-XPL-VC-S13-19   Serial#: 425HC109
```

Syntax: `show media ethernet<port-num>`

Viewing optical monitoring information

To view temperature and power information for all qualified XFPs and SFPs in a particular slot, use the **show optic** command. The following shows an example output.

```
PowerConnect> show optic 4
Port Temperature Tx Power Rx Power Tx Bias Current
+-----+-----+-----+-----+-----+
1 30.8242 C -001.8822 dBm -002.5908 dBm 41.790 mA
   Normal Normal Normal Normal
2 31.7070 C -001.4116 dBm -006.4092 dBm 41.976 mA
   Normal Normal Normal Normal
3 30.1835 C -000.5794 dBm 0.000 mA
```

	Normal	Low-Alarm	Normal	Low-Alarm
4	0.0000 C			0.000 mA
	Normal	Normal	Normal	Normal

Syntax: `show optic slot number`

NOTE

This function takes advantage of information stored and supplied by the manufacturer of the XFP or SFP transceiver. This information is an optional feature of the Multi-Source Agreement standard defining the optical interface. Not all component suppliers have implemented this feature set. In such cases where the XFP or SFP transceiver does not supply the information, a “Not Available” message will be displayed for the specific port on which the module is installed.

The following table describes the information displayed by the **show optic** command.

TABLE 10 Output from the **show optic** command

This field...	Displays...
Port	The Dell port number.
Temperature	<ul style="list-style-type: none"> The operating temperature, in degrees Celsius, of the optical transceiver. The alarm status, as described in Table 11.
Tx Power	<ul style="list-style-type: none"> The transmit power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). The alarm status, as described in Table 11.
Rx Power	<ul style="list-style-type: none"> The receive power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). The alarm status, as described in Table 11.
Tx Bias Current	<ul style="list-style-type: none"> The transmit bias power signal, in milliamperes (mA). The alarm status, as described in Table 11.

For Temperature, Tx Power, Rx Power, and Tx Bias Current in the **show optic** command output, values are displayed along with one of the following alarm status values: Low-Alarm, Low-Warn, Normal, High-Warn or High-Alarm. The thresholds that determine these status values are set by the manufacturer of the optical transceivers. [Table 11](#) describes each of these status values.

TABLE 11 Alarm status value description

Status value	Description
Low-Alarm	Monitored level has dropped below the "low-alarm" threshold set by the manufacturer of the optical transceiver.
Low-Warn	Monitored level has dropped below the "low-warn" threshold set by the manufacturer of the optical transceiver.
Normal	Monitored level is within the "normal" range set by the manufacturer of the optical transceiver.
High-Warn	Monitored level has climbed above the "high-warn" threshold set by the manufacturer of the optical transceiver.
High-Alarm	Monitored level has climbed above the "high-alarm" threshold set by the manufacturer of the optical transceiver.

Viewing optical transceiver thresholds

The thresholds that determine the alarm status values for an optical transceiver are set by the manufacturer of the XFP or SFP. To view the thresholds for a qualified optical transceiver in a particular port, use the **show optic threshold** command as shown below.

```
PowerConnect> show optic threshold 2
Port 2 sfp monitor thresholds:
Temperature High alarm          5a00          90.0000 C
Temperature Low alarm           d300          -45.0000 C
Temperature High warning        5500          85.0000 C
Temperature Low warning         d800          -40.0000 C
Supply Voltage High alarm       9088
Supply Voltage Low alarm        7148
Supply Voltage High warning     8ca0
Supply Voltage Low warning      7530
TX Bias High alarm              7530          60.000 mA
TX Bias Low alarm               01f4          1.000 mA
TX Bias High warning            61a8          50.000 mA
TX Bias Low warning             05dc          3.000 mA
TX Power High alarm             1f07          -001.0001 dBm
TX Power Low alarm              02c4          -011.4996 dBm
TX Power High warning           18a6          -001.9997 dBm
TX Power Low warning            037b          -010.5012 dBm
RX Power High alarm             2710          000.0000 dBm
RX Power Low alarm              0028          -023.9794 dBm
RX Power High warning           1f07          -001.0001 dBm
RX Power Low warning            0032          -023.0102 dBm
```

Syntax: `show optic threshold port-num`

For Temperature, Supply Voltage, TX Bias, TX Power, and RX Power, values are displayed for each of the following four alarm and warning settings: High alarm, Low alarm, High warning, and Low warning. The hexadecimal values are the manufacturer internal calibrations, as defined in the SFF-8472 standard. The other values indicate at what level (above the high setting or below the low setting) the system should send a warning message or an alarm. Note that these values are set by the manufacturer of the optical transceiver, and cannot be configured.

Syslog messages

The system generates Syslog messages for optical transceivers in the following circumstances:

- The temperature, supply voltage, TX Bias, TX power, or TX power value goes above or below the high or low warning or alarm threshold set by the manufacturer.
- The optical transceiver does not support digital optical monitoring.
- The optical transceiver is not qualified, and therefore not supported by Dell.

For details about the above Syslog messages, refer to [Chapter 34, “Using Syslog”](#).

Configuring IPv6 Connectivity

IPv6 addressing overview


NOTE

This chapter does not describe IPv6 routing protocols, which are covered in separate chapters throughout this guide.

IPv6 was designed to replace IPv4, the Internet protocol that is most commonly used currently throughout the world. IPv6 increases the number of network address bits from 32 (IPv4) to 128 bits, which provides more than enough unique IP addresses to support all of the network devices on the planet into the future. IPv6 is expected to quickly become the network standard.

An IPv6 address is composed of 8 fields of 16-bit hexadecimal values separated by colons (:). [Figure 1](#) shows the IPv6 address format.

FIGURE 1 IPv6 address format



As shown in [Figure 1](#), HHHH is a 16-bit hexadecimal value, while H is a 4-bit hexadecimal value. The following is an example of an IPv6 address.

```
2001:0000:0000:0200:002D:D0FF:FE48:4672
```

Note that this IPv6 address includes hexadecimal fields of zeros. To make the address less cumbersome, you can do the following:

- Omit the leading zeros; for example, 2001:0:0:200:2D:D0FF:FE48:4672.
- Compress the successive groups of zeros at the beginning, middle, or end of an IPv6 address to two colons (::) once per address; for example, 2001::200:2D:D0FF:FE48:4672.

When specifying an IPv6 address in a command syntax, keep the following in mind:

- You can use the two colons (::) only once in the address to represent the longest successive hexadecimal fields of zeros
- The hexadecimal letters in IPv6 addresses are not case-sensitive

As shown in [Figure 1](#), the IPv6 network prefix is composed of the left-most bits of the address. As with an IPv4 address, you can specify the IPv6 prefix using the *prefix / prefix-length* format, where the following applies.

The *prefix* parameter is specified as 16-bit hexadecimal values separated by a colon.

The *prefix-length* parameter is specified as a decimal value that indicates the left-most bits of the IPv6 address.

The following is an example of an IPv6 prefix.

```
2001:FF08:49EA:D088::/64
```

IPv6 address types

As with IPv4 addresses, you can assign multiple IPv6 addresses to a switch interface. [Table 12](#) presents the three major types of IPv6 addresses that you can assign to a switch interface.

A major difference between IPv4 and IPv6 addresses is that IPv6 addresses support **scope**, which describes the topology in which the address may be used as a unique identifier for an interface or set of interfaces.

Unicast and multicast addresses support scoping as follows:

- Unicast addresses support two types of scope: global scope and local scope. In turn, local scope supports site-local addresses and link-local addresses. [Table 12](#) describes global, site-local, and link-local addresses and the topologies in which they are used.
- Multicast addresses support a scope field, which [Table 12](#) describes.

TABLE 12 IPv6 address types

Address type	Description	Address structure
Unicast	An address for a single interface. A packet sent to a unicast address is delivered to the interface identified by the address.	<p>Depends on the type of the unicast address:</p> <ul style="list-style-type: none"> • Aggregatable global address—An address equivalent to a global or public IPv4 address. The address structure is as follows: a fixed prefix of 2000::/3 (001), a 45-bit global routing prefix, a 16-bit subnet ID, and a 64-bit interface ID. • Site-local address—An address used within a site or intranet. (This address is similar to a private IPv4 address.) A site consists of multiple network links. The address structure is as follows: a fixed prefix of FEC0::/10 (1111 1110 11), a 16-bit subnet ID, and a 64-bit interface ID. • Link-local address—An address used between directly connected nodes on a single network link. The address structure is as follows: a fixed prefix of FE80::/10 (1111 1110 10) and a 64-bit interface ID. • IPv4-compatible address—An address used in IPv6 transition mechanisms that tunnel IPv6 packets dynamically over IPv4 infrastructures. The address embeds an IPv4 address in the low-order 32 bits and the high-order 96 bits are zeros. The address structure is as follows: 0:0:0:0:0:A.B.C.D. • Loopback address—An address (0:0:0:0:0:0:1 or ::1) that a switch can use to send an IPv6 packet to itself. You cannot assign a loopback address to a physical interface. • Unspecified address—An address (0:0:0:0:0:0:0 or ::) that a node can use until you configure an IPv6 address for it.
Multicast	An address for a set of interfaces belonging to different nodes. Sending a packet to a multicast address results in the delivery of the packet to all interfaces in the set.	A multicast address has a fixed prefix of FF00::/8 (1111 1111). The next 4 bits define the address as a permanent or temporary address. The next 4 bits define the scope of the address (node, link, site, organization, global).
Anycast	An address for a set of interfaces belonging to different nodes. Sending a packet to an anycast address results in the delivery of the packet to the closest interface identified by the address.	<p>An anycast address looks similar to a unicast address, because it is allocated from the unicast address space. If you assign a unicast address to multiple interfaces, it is an anycast address. An interface assigned an anycast address must be configured to recognize the address as an anycast address.</p> <p>An anycast address can be assigned to a switch only.</p> <p>An anycast address must not be used as the source address of an IPv6 packet.</p>

A switch automatically configures a link-local unicast address for an interface by using the prefix of FE80::/10 (1111 1110 10) and a 64-bit interface ID. The 128-bit IPv6 address is then subjected to duplicate address detection to ensure that the address is unique on the link. If desired, you can override this automatically configured address by explicitly configuring an address.

IPv6 stateless autoconfiguration

PowerConnect devices use the IPv6 stateless autoconfiguration feature to enable a host on a local link to automatically configure its interfaces with new and globally unique IPv6 addresses associated with its location. The automatic configuration of a host interface is performed without the use of a server, such as a Dynamic Host Configuration Protocol (DHCP) server, or manual configuration.

The automatic configuration of a host interface works in the following way: a switch on a local link periodically sends switch advertisement messages containing network-type information, such as the 64-bit prefix of the local link and the default route, to all nodes on the link. When a host on the link receives the message, it takes the local link prefix from the message and appends a 64-bit interface ID, thereby automatically configuring its interface. (The 64-bit interface ID is derived from the MAC address of the host NIC.) The 128-bit IPv6 address is then subjected to duplicate address detection to ensure that the address is unique on the link.

The duplicate address detection feature verifies that a unicast IPv6 address is unique before it is assigned to a host interface by the stateless auto configuration feature. Duplicate address detection uses neighbor solicitation messages to verify that a unicast IPv6 address is unique.

NOTE

For the stateless auto configuration feature to work properly, the advertised prefix length in switch advertisement messages must always be 64 bits.

The IPv6 stateless autoconfiguration feature can also automatically reconfigure a host interfaces if you change the ISP for the host network. (The host interfaces must be renumbered with the IPv6 prefix of the new ISP.)

The renumbering occurs in the following way: a switch on a local link periodically sends advertisements updated with the prefix of the new ISP to all nodes on the link. (The advertisements still contain the prefix of the old ISP.) A host can use the addresses created from the new prefix and the existing addresses created from the old prefix on the link. During this transition, the old prefix is removed from the switch advertisements. At this point, only addresses that contain the new prefix are used on the link.

IPv6 CLI command support

Table 13 lists the IPv6 CLI commands supported.

TABLE 13 IPv6 CLI command support

IPv6 command	Description	Switch code	Router code
clear ipv6 cache	Deletes all entries in the dynamic host cache.		X
clear ipv6 neighbor	Deletes all dynamic entries in the IPv6 neighbor table.	X	X
clear ipv6 traffic	Resets all IPv6 packet counters.	X	X
copy tftp	Downloads a copy of a Dell software image from a TFTP server into the system flash using IPv6.	X	X
debug ipv6	Displays IPv6 debug information.	X	X
ipv6 address	Configures an IPv6 address on an interface (router) or globally (switch)	X	X

TABLE 13 IPv6 CLI command support (Continued)

IPv6 command	Description	Switch code	Router code
ipv6 debug	Enables IPv6 debugging.	X	X
ipv6 dns domain-name	Configures an IPv6 domain name.	X	X
ipv6 dns server-address	Configures an IPv6 DNS server address.	X	X
ipv6 enable	Enables IPv6 on an interface.	X	X
ipv6 neighbor	Maps a static IPv6 address to a MAC address in the IPv6 neighbor table.		X
log host ipv6	Configures the IPv6 Syslog server.	X	X
ping ipv6	Performs an ICMP for IPv6 echo test.	X	X
show ipv6	Displays some global IPv6 parameters, such as IPv6 DNS server address.	X	X
show ipv6 cache	Displays the IPv6 host cache.		X
show ipv6 interface	Displays IPv6 information for an interface.		X
show ipv6 neighbor	Displays the IPv6 neighbor table.	X	X
show ipv6 tcp	Displays information about IPv6 TCP sessions.	X	X
show ipv6 traffic	Displays IPv6 packet counters.	X	X
snmp-client ipv6	Restricts SNMP access to a certain IPv6 node.	X	X
snmp-server host ipv6	Specifies the recipient of SNMP notifications.	X	X
sntp server ipv6	Enables the PowerConnect device to send SNTP packets over IPv6.	X	X
telnet	Enables a Telnet connection from the PowerConnect device to a remote IPv6 host using the console.	X	X

Configuring an IPv6 host address on a Layer 2 switch

In a Layer 3 (router) configuration, each port can be configured separately with an IPv6 address. This is accomplished using the interface configuration process that is described in [“Configuring IPv6 on each router interface”](#) on page 65.

In a Layer 2 (switch) configuration, individual ports cannot be configured with an IP address (IPv4 or IPv6). In this situation, the switch has one IP address for the management port and one IP address for the system. This has previously been supported for IPv4 but not for IPv6.

There is support for configuring an IPv6 address on the management port as described in [“Configuring the management port for an IPv6 automatic address configuration”](#) on page 65, and for configuring a system-wide IPv6 address on a Layer 2 switch. Configuration of the system-wide IPv6 address is exactly like configuration of an IPv6 address in router mode, except that the IPv6 configuration is at the Global Config level instead of at the Interface Config level.

The process for defining the system-wide interface for IPv6 is described in the following sections:

- [“Configuring a global or site-local IPv6 address with a manually configured interface ID”](#) on page 64
- [“Configuring a link-local IPv6 address as a system-wide address for a switch”](#) on page 64

NOTE

When configuring an IPv6 host address on a Layer 2 switch that has multiple VLANs, make sure the configuration includes a designated management VLAN that identifies the VLAN to which the global IP address belongs. Refer to [“Designated VLAN for Telnet management sessions to a Layer 2 Switch”](#) on page 863.

Configuring a global or site-local IPv6 address with a manually configured interface ID

To configure a global or site-local IPv6 address with a manually-configured interface ID, such as a system-wide address for a switch, enter a command similar to the following at the Global Config level.

```
PowerConnect(config)# ipv6 address 2001:200:12D:1300:240:D0FF:FE48:4000:1/64
```

Syntax: `ipv6 address ipv6-prefix / prefix-length`

You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

You must specify the *prefix-length* parameter in decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.

Configuring a link-local IPv6 address as a system-wide address for a switch

To enable IPv6 and automatically configure a global interface enter commands such as the following.

```
PowerConnect(config)# ipv6 enable
```

This command enables IPv6 on the switch and specifies that the interface is assigned an automatically computed link-local address.

Syntax: `[no] ipv6 enable`

To override a link-local address that is automatically computed for the global interface with a manually configured address, enter a command such as the following.

```
PowerConnect(config)# ipv6 address FE80::240:D0FF:FE48:4672 link-local
```

This command explicitly configures the link-local address FE80::240:D0FF:FE48:4672 for the global interface.

Syntax: `ipv6 address ipv6-address link-local`

You must specify the *ipv6-address* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The **link-local** keyword indicates that the router interface should use the manually configured link-local address instead of the automatically computed link-local address.

Configuring the management port for an IPv6 automatic address configuration

You can have the management port configured to automatically obtain an IPv6 address. This process is the same for any other port and is described in detail in the [“Configuring a global IPv6 address with an automatically computed EUI-64 interface ID”](#) on page 66

Configuring basic IPv6 connectivity on a Layer 3 switch

To configure basic IPv6 connectivity on a Layer 3 Switch, you must do the following:

- Configure an IPv6 address or explicitly enable IPv6 on each router interface.

All other configuration tasks in this chapter are optional..

Configuring IPv6 on each router interface

To forward IPv6 traffic on a router interface, the interface must have an IPv6 address, or IPv6 must be explicitly enabled. By default, an IPv6 address is not configured on a router interface.

If you choose to configure a global or site-local IPv6 address for an interface, IPv6 is also enabled on the interface. Further, when you configure a global or site-local IPv6 address, you must decide on one of the following in the low-order 64 bits:

- A manually configured interface ID.
- An automatically computed EUI-64 interface ID.

If you prefer to assign a link-local IPv6 address to the interface, you must explicitly enable IPv6 on the interface, which causes a link-local address to be automatically computed for the interface. If preferred, you can override the automatically configured link-local address with an address that you manually configure.

This section provides the following information:

- Configuring a global or site-local address with a manually configured or automatically computed interface ID for an interface.
- Automatically or manually configuring a link-local address for an interface.
- Configuring IPv6 anycast addresses

Configuring a global or site-local IPv6 address on an interface

Configuring a global or site-local IPv6 address on an interface does the following:

- Automatically configures an interface ID (a link-local address), if specified.
- Enables IPv6 on that interface.

Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:0:1:FF00::/104 for each unicast address assigned to the interface.
- Solicited-node for subnet anycast address for each unicast assigned address

5 Configuring basic IPv6 connectivity on a Layer 3 switch

- Solicited-node for anycast address FF02:0:0:0:1:FF00::0000
- All-nodes link-local multicast group FF02::1
- All-routers link-local multicast group FF02::2

The neighbor discovery feature sends messages to these multicast groups. For more information, refer to “[Configuring IPv6 neighbor discovery](#)” on page 77.

Configuring a global or site-local IPv6 address with a manually configured interface ID

To configure a global or site-local IPv6 address, including a manually configured interface ID, for an interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipv6 address 2001:200:12D:1300:240:D0FF:
FE48:4672:/64
```

These commands configure the global prefix 2001:200:12d:1300::/64 and the interface ID ::240:D0FF:FE48:4672, and enable IPv6 on Ethernet interface 1.

Syntax: `ipv6 address ipv6-prefix / prefix-length`

You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.

To configure a /122 address on a VE enter commands similar to the following.

```
PowerConnect(config-vlan-11)# int ve11
PowerConnect(config-vif-11)# ipv6 add 2020::1/122
PowerConnect(config-vif-11)# sh ipv6 int
Routing Protocols : R - RIP  O - OSPF
Interface      Status      Routing  Global Unicast Address
VE 11          up/up          2020::1/122
PowerConnect(config-vif-11)# sh ipv6 route
IPv6 Routing Table - 1 entries:
Type Codes:  C - Connected, S - Static, R - RIP, O - OSPF, B - BGP
OSPF Sub Type Codes:  O - Intra, Oi - Inter, Ol - Type1 external, O2 - Type2
external
Type IPv6 Prefix                Next Hop Router          Interface  Dis/Metric
C  2020::/122                    ::                       ve 11     0/0
```

Configuring a global IPv6 address with an automatically computed EUI-64 interface ID

To configure a global IPv6 address with an automatically computed EUI-64 interface ID in the low-order 64-bits, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipv6 address 2001:200:12D:1300::/64 eui-64
```

These commands configure the global prefix 2001:200:12d:1300::/64 and an interface ID, and enable IPv6 on Ethernet interface 1.

Syntax: `ipv6 address ipv6-prefix / prefix-length eui-64`

You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.

The **eui-64** keyword configures the global address with an EUI-64 interface ID in the low-order 64 bits. The interface ID is automatically constructed in IEEE EUI-64 format using the interface MAC address.

Configuring a link-local IPv6 address on an interface

To explicitly enable IPv6 on a router interface without configuring a global or site-local address for the interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipv6 enable
```

These commands enable IPv6 on Ethernet interface 1 and specify that the interface is assigned an automatically computed link-local address.

Syntax: [no] ipv6 enable

NOTE

When configuring VLANs that share a common tagged interface with a physical or Virtual Ethernet (VE) interface, Dell recommends that you override the automatically computed link-local address with a manually configured unique address for the interface. If the interface uses the automatically computed address, which in the case of physical and VE interfaces is derived from a global MAC address, all physical and VE interfaces will have the same MAC address.

To override a link-local address that is automatically computed for an interface with a manually configured address, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e10000-1)# ipv6 address FE80::240:D0FF:FE48:4672
link-local
```

These commands explicitly configure the link-local address FE80::240:D0FF:FE48:4672 for Ethernet interface 1.

Syntax: ipv6 address ipv6-address link-local

You must specify the *ipv6-address* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The **link-local** keyword indicates that the router interface should use the manually configured link-local address instead of the automatically computed link-local address.

Configuring an IPv6 anycast address on an interface

In IPv6, an **anycast** address is an address for a set of interfaces belonging to different nodes. Sending a packet to an anycast address results in the delivery of the packet to the closest interface configured with the anycast address.

An anycast address looks similar to a unicast address, because it is allocated from the unicast address space. If you assign an IPv6 unicast address to multiple interfaces, it is an anycast address. On the PowerConnect device, you configure an interface assigned an anycast address to recognize the address as an anycast address.

For example, the following commands configure an anycast address on interface 1.

```
PowerConnect(config)# int e 1
PowerConnect(config-if-e10000-1)# ipv6 address 2002::/64 anycast
```

Syntax: ipv6 address ipv6-prefix / prefix-length [anycast]

IPv6 anycast addresses are described in detail in RFC 1884. See RFC 2461 for a description of how the IPv6 Neighbor Discovery mechanism handles anycast addresses.

IPv6 management (IPv6 host support)

An **IPv6 host** has IPv6 addresses on its interfaces, but does not have full IPv6 routing enabled on it.

This section describes the following IPv6 host features:

- [“Restricting SNMP access to an IPv6 node”](#)
- [“Specifying an IPv6 SNMP trap receiver”](#)
- [“SNMP V3 over IPv6”](#)
- [“SNTP over IPv6”](#)
- [“Secure Shell, SCP, and IPv6”](#)
- [“IPv6 Telnet”](#)
- [“Configuring name-to-IPv6 address resolution using IPv6 DNS resolver”](#)
- [“Defining an IPv6 DNS entry”](#)
- [“Using the IPv6 copy command”](#)
- [“Using the IPv6 ncopy command”](#)
- [“IPv6 ping”](#)
- [“Configuring an IPv6 Syslog server”](#)
- [“Viewing IPv6 SNMP server addresses”](#)
- [“IPv6 debug”](#)
- [“Disabling IPv6 on a Layer 2 switch”](#)

The following IPv6 host feature is also supported:

- [“Configuring a link-local IPv6 address as a system-wide address for a switch”](#)

Restricting SNMP access to an IPv6 node

You can restrict SNMP access (which includes IronView Network Manager) to the device to the IPv6 host whose IP address you specify. To do so, enter a command such as the following.

```
PowerConnect(config)# snmp-client ipv6 2001:efff:89::23
```

Syntax: `snmp-client ipv6 ipv6-address`

- The *ipv6-address* you specify must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specifying an IPv6 SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter a command such as the following.

```
PowerConnect(config)# snmp-server host ipv6 2001:efff:89::13
```

Syntax: `snmp-server host ipv6 ipv6-address`

- The *ipv6-address* you specify must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

SNMP V3 over IPv6

PowerConnect devices support IPv6 for SNMP version 3. For more information about how to configure SNMP, refer to [Chapter 32, “Securing SNMP Access”](#).

SNTP over IPv6

To enable the PowerConnect device to send SNTP packets over IPv6, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)# sntp server ipv6 3000::400
```

Syntax: `sntp server ipv6 ipv6-address`

- The *ipv6-address* is the IPv6 address of the SNTP server. When you enter the IPv6 address, you do not need to specify the prefix length. A prefix length of 128 is implied.

Secure Shell, SCP, and IPv6

Secure Shell (SSH) is a mechanism that allows secure remote access to management functions on the PowerConnect device. SSH provides a function similar to Telnet. You can log in to and configure the PowerConnect device using a publicly or commercially available SSH client program, just as you can with Telnet. However, unlike Telnet, which provides no security, SSH provides a secure, encrypted connection to the device.

To open an SSH session between an IPv6 host running an SSH client program and the PowerConnect device, open the SSH client program and specify the IPv6 address of the device. For more information about configuring SSH on the PowerConnect device, refer to [“SSH version 2 support”](#) on page 911.

IPv6 Telnet

Telnet sessions can be established between a PowerConnect device to a remote IPv6 host, and from a remote IPv6 host to the PowerConnect device using IPv6 addresses.

The **telnet** command establishes a Telnet connection from a PowerConnect device to a remote IPv6 host using the console. Up to five **read-access** Telnet sessions are supported on the router at one time. **Write-access** through Telnet is limited to one session, and only one outgoing Telnet session is supported on the router at one time. To see the number of open Telnet sessions at any time, enter the **show telnet** command.

Example

To establish a Telnet connection to a remote host with the IPv6 address of 3001:2837:3de2:c37::6, enter the following command.

```
PowerConnect# telnet 3001:2837:3de2:c37::6
```

Syntax: `telnet ipv6-address [port-number | outgoing-interface ethernet port | ve number]`

- The *ipv6-address* parameter specifies the address of a remote host. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

- The *port-number* parameter specifies the port number on which the PowerConnect device establishes the Telnet connection. You can specify a value between 1 - 65535. If you do not specify a port number, the PowerConnect device establishes the Telnet connection on port 23.
- If the IPv6 address you specify is a link-local address, you must specify the **outgoing-interface** Ethernet *port | ve number* parameter. This parameter identifies the interface that must be used to reach the remote host. If you specify an Ethernet interface, you must also specify the port number associated with the interface. If you specify a VE interface, also specify the VE number.

Establishing a Telnet session from an IPv6 host

To establish a Telnet session from an IPv6 host to the PowerConnect device, open your Telnet application and specify the IPv6 address of the device.

Configuring name-to-IPv6 address resolution using IPv6 DNS resolver

The Domain Name Server (DNS) resolver feature lets you use a host name to perform Telnet and ping commands. You can also define a DNS domain on a PowerConnect device and thereby recognize all hosts within that domain. After you define a domain name, the PowerConnect device automatically appends the appropriate domain to the host and forwards it to the domain name server.

For example, if the domain “newyork.com” is defined on a PowerConnect device, and you want to initiate a ping to host “NYC01” on that domain, you need to reference only the host name in the command instead of the host name and its domain name. For example, you could enter either of the following commands to initiate the ping.

```
PowerConnect# ping ipv6 nyc01
PowerConnect# ping ipv6 nyc01.newyork.com
```

Defining an IPv6 DNS entry

IPv6 defines new DNS record types to resolve queries for domain names to IPv6 addresses, as well as IPv6 addresses to domain names. PowerConnect devices running IPv6 software support AAAA DNS records, which are defined in RFC 1886.

AAAA DNS records are analogous to the A DNS records used with IPv4. They store a complete IPv6 address in each record. AAAA records have a type value of 28.

To establish an IPv6 DNS entry for the device, enter the following command.

```
PowerConnect(config)# ipv6 dns domain-name companynet.com
```

Syntax: [no] **ipv6 dns domain-name** *domain name*

To define an IPv6 DNS server address, enter the following command.

```
PowerConnect(config)# ipv6 dns server-address 200::1
```

Syntax: [no] **ipv6 dns server-address** *ipv6-addr* [*ipv6-addr*] [*ipv6-addr*] [*ipv6-addr*]

As an example, in a configuration where ftp6.companynet.com is a server with an IPv6 protocol stack, when a user pings ftp6.companynet.com, the PowerConnect device attempts to resolve the AAAA DNS record. In addition, if the DNS server does not have an IPv6 address, as long as it is able to resolve AAAA records, it can still respond to DNS queries.

Using the IPv6 copy command

The **copy** command for IPv6 allows you to do the following:

- Copy a file from a specified source to an IPv6 TFTP server.
- Copy a file from an IPv6 TFTP server to a specified destination.

Copying a file to an IPv6 TFTP server

You can copy a file from the following sources to an IPv6 TFTP server:

- Flash memory.
- Running configuration.
- Startup configuration.

Copying a file from flash memory

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
PowerConnect# copy flash tftp 2001:7382:e0ff:7837::3 test.img secondary
```

This command copies the secondary boot image named test.img from flash memory to a TFTP server with the IPv6 address of 2001:7382:e0ff:7837::3.

Syntax: **copy flash tftp** *ipv6-address source-file-name primary | secondary*

- The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name* parameter specifies the name of the file you want to copy to the IPv6 TFTP server.
- The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.

Copying a file from the running or startup configuration

For example, to copy the running configuration to an IPv6 TFTP server, enter a command such as the following.

```
PowerConnect# copy running-config tftp 2001:7382:e0ff:7837::3 newrun.cfg
```

This command copies the running configuration to a TFTP server with the IPv6 address of 2001:7382:e0ff:7837::3 and names the file on the TFTP server newrun.cfg.

Syntax: **copy running-config | startup-config tftp** *ipv6-address destination-file-name*

- Specify the **running-config** keyword to copy the running configuration file to the specified IPv6 TFTP server.
- Specify the **startup-config** keyword to copy the startup configuration file to the specified IPv6 TFTP server.
- The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *destination-file-name* parameter specifies the name of the file that is copied to the IPv6 TFTP server.

Copying a file from an IPv6 TFTP server

You can copy a file from an IPv6 TFTP server to the following destinations:

- Flash memory.
- Running configuration.
- Startup configuration.

Copying a file to flash memory

For example, to copy a boot image from an IPv6 TFTP server to the primary or secondary storage location in the device flash memory, enter a command such as the following.

```
PowerConnect# copy tftp flash 2001:7382:e0ff:7837::3 test.img secondary
```

This command copies a boot image named test.img from an IPv6 TFTP server with the IPv6 address of 2001:7382:e0ff:7837::3 to the secondary storage location in the device flash memory.

Syntax: `copy tftp flash ipv6-address source-file-name primary | secondary`

- The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name* parameter specifies the name of the file you want to copy from the IPv6 TFTP server.
- The **primary** keyword specifies the primary storage location in the device flash memory, while the **secondary** keyword specifies the secondary storage location in the device flash memory.

Copying a file to the running or startup configuration

For example, to copy a configuration file from an IPv6 TFTP server to the router running or startup configuration, enter a command such as the following.

```
PowerConnect# copy tftp running-config 2001:7382:e0ff:7837::3 newrun.cfg  
overwrite
```

This command copies the new running .oncfg file from the IPv6 TFTP server and overwrites the router running configuration file with the contents of newrun.cfg.

NOTE

To activate this configuration, you must reload (reset) the device.

Syntax: `copy tftp running-config | startup-config ipv6-address source-file-name [overwrite]`

- Specify the **running-config** keyword to copy the running configuration from the specified IPv6 TFTP server.
- Specify the **startup-config** keyword to copy the startup configuration from the specified IPv6 TFTP server.
- The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name* parameter specifies the name of the file that is copied from the IPv6 TFTP server.
- The **overwrite** keyword specifies that the device should overwrite the current configuration file with the copied file. If you do not specify this parameter, the device copies the file into the current running or startup configuration but does not overwrite the current configuration.

NOTE

You cannot use the overwrite option from non-console sessions, because it will disconnect the session.

Using the IPv6 ncopy command

The **ncopy** command for IPv6 allows you to do the following:

- Copy a primary or secondary boot image from flash memory to an IPv6 TFTP server.
- Copy the running configuration to an IPv6 TFTP server.
- Copy the startup configuration to an IPv6 TFTP server
- Upload various files from an IPv6 TFTP server.

Copying a primary or secondary boot image from flash memory to an IPv6 TFTP server

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
PowerConnect# ncopy flash primary tftp 2001:7382:e0ff:7837::3 primary.img
```

This command copies the primary boot image named primary.img from flash memory to a TFTP server with the IPv6 address of 2001:7382:e0ff:7837::3.

Syntax: **ncopy flash primary** | **secondary tftp** *ipv6-address source-file-name*

- The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.
- The **tftp ipv6-address** parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name* parameter specifies the name of the file you want to copy from flash memory.

Copying the running or startup configuration to an IPv6 TFTP server

For example, to copy a device running or startup configuration to an IPv6 TFTP server, enter a command such as the following.

```
PowerConnect# ncopy running-config tftp 2001:7382:e0ff:7837::3 bakrun.cfg
```

This command copies a device running configuration to a TFTP server with the IPv6 address of 2001:7382:e0ff:7837::3 and names the destination file bakrun.cfg.

Syntax: **ncopy running-config** | **startup-config tftp** *ipv6-address destination-file-name*

- Specify the **running-config** keyword to copy the device running configuration or the **startup-config** keyword to copy the device startup configuration.
- The **tftp ipv6-address** parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *destination-file-name* parameter specifies the name of the running configuration that is copied to the IPv6 TFTP server.

Uploading files from an IPv6 TFTP server

You can upload the following files from an IPv6 TFTP server:

- Primary boot image.
- Secondary boot image.
- Running configuration.
- Startup configuration.

Uploading a primary or secondary boot image from an IPv6 TFTP server

For example, to upload a primary or secondary boot image from an IPv6 TFTP server to a device flash memory, enter a command such as the following.

```
PowerConnect# ncopy tftp 2001:7382:e0ff:7837::3 primary.img flash primary
```

This command uploads the primary boot image named `primary.img` from a TFTP server with the IPv6 address of `2001:7382:e0ff:7837::3` to the device primary storage location in flash memory.

Syntax: `ncopy tftp ipv6-address source-file-name flash primary | secondary`

- The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name* parameter specifies the name of the file you want to copy from the TFTP server.
- The **primary** keyword specifies the primary location in flash memory, while the **secondary** keyword specifies the secondary location in flash memory.

Uploading a running or startup configuration from an IPv6 TFTP server

For example to upload a running or startup configuration from an IPv6 TFTP server to a device, enter a command such as the following.

```
PowerConnect# ncopy tftp 2001:7382:e0ff:7837::3 newrun.cfg running-config
```

This command uploads a file named `newrun.cfg` from a TFTP server with the IPv6 address of `2001:7382:e0ff:7837::3` to the device.

Syntax: `ncopy tftp ipv6-address source-file-name running-config | startup-config`

- The **tftp** *ipv6-address*> parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *source-file-name*> parameter specifies the name of the file you want to copy from the TFTP server.
- Specify the **running-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The device copies the specified file into the current running configuration but does not overwrite the current configuration.

Specify the **startup-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The the device copies the specified file into the current startup configuration but does not overwrite the current configuration.

IPv6 ping

The **ping** command allows you to verify the connectivity from a PowerConnect device to an IPv6 device by performing an ICMP for IPv6 echo test.

For example, to ping a device with the IPv6 address of 2001:3424:847f:a385:34dd::45 from the PowerConnect device, enter the following command.

```
PowerConnect# ping ipv6 2001:3424:847f:a385:34dd::45
```

Syntax: `ping ipv6 ipv6-address [outgoing-interface [port | ve number]] [source ipv6-address] [{ count number } [timeout milliseconds] [ttl number] [size bytes] [quiet] [numeric] [verify] [data 1-to-4 byte hex] [brief]`

- The `ipv6-address` parameter specifies the address of the router. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The `outgoing-interface` keyword specifies a physical interface over which you can verify connectivity. If you specify a physical interface, such as an Ethernet interface, you must also specify the port number of the interface. If you specify a virtual interface, such as a VE, you must specify the number associated with the VE.
- The `source ipv6-address` parameter specifies an IPv6 address to be used as the origin of the ping packets.

NOTE

The `outgoing-interface` and `source` options are available only on router code and not on switch code.

- The `count number` parameter specifies how many ping packets the router sends. You can specify from 1 - 4294967296. The default is 1.
- The `timeout milliseconds` parameter specifies how many milliseconds the router waits for a reply from the pinged device. You can specify a timeout from 1 - 4294967294 milliseconds. The default is 5000 (5 seconds).
- The `ttl number` parameter specifies the maximum number of hops. You can specify a TTL from 1 - 255. The default is 64.
- The `size bytes` parameter specifies the size of the ICMP data portion of the packet. This is the payload and does not include the header. You can specify from 0 - 10173. The default is 16.
- The `quiet` keyword hides informational messages such as a summary of the ping parameters sent to the device, and instead only displays messages indicating the success or failure of the ping. This option is disabled by default.
- The `verify` keyword verifies that the data in the echo packet (the reply packet) is the same as the data in the echo request (the ping). By default the device does not verify the data.
- The `data 1 - 4 byte hex` parameter lets you specify a specific data pattern for the payload instead of the default data pattern, "abcd", in the packet's data payload. The pattern repeats itself throughout the ICMP message (payload) portion of the packet.

NOTE

For parameters that require a numeric value, the CLI does not check that the value you enter is within the allowed range. Instead, if you do exceed the range for a numeric value, the software rounds the value to the nearest valid value.

- The `brief` keyword causes ping test characters to be displayed. The following ping test characters are supported.
 - ! Indicates that a reply was received.
 - . Indicates that the network server timed out while waiting for a reply.
 - U Indicates that a destination unreachable error PDU was received.

I Indicates that the user interrupted ping.

Configuring an IPv6 Syslog server

To enable IPv6 logging, specify an IPv6 Syslog server. Enter a command such as the following.

```
PowerConnect(config)# log host ipv6 2000:2383:e0bb::4/128
```

Syntax: `log host ipv6 ipv6-address [udp-port-num]`

- The *ipv6-address* must be in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *udp-port-num* optional parameter specifies the UDP application port used for the Syslog facility.

Viewing IPv6 SNMP server addresses

Some of the **show** commands display IPv6 addresses for IPv6 SNMP servers. The following shows an example output for the **show snmp server** command.

```
PowerConnect# show snmp server
```

```

    Contact:
    Location:
Community(ro): .....
```

Traps

```

    Warm/Cold start: Enable
    Link up: Enable
    Link down: Enable
    Authentication: Enable
    Locked address violation: Enable
    Power supply failure: Enable
    Fan failure: Enable
    Temperature warning: Enable
    STP new root: Enable
    STP topology change: Enable
    vsrp: Enable
```

```
Total Trap-Receiver Entries: 4
```

Trap-Receiver	IP-Address	Port-Number	Community
1	192.147.201.100	162
2	4000::200	162
3	192.147.202.100	162
4	3000::200	162

Disabling router advertisement and solicitation messages

Router advertisement and solicitation messages enable a node on a link to discover the routers on the same link. By default, router advertisement and solicitation messages are permitted on the device. To disable these messages, configure an IPv6 access control list that denies them. The following shows an example configuration.

Example

```
PowerConnect(config)# ipv6 access-list rtradvert
PowerConnect(config)# deny icmp any any router-advertisement
PowerConnect(config)# deny icmp any any router-solicitation
PowerConnect(config)# permit ipv6 any any
```

IPv6 debug

The **debug ipv6** commands enable the collection of information about IPv6 configurations for troubleshooting.

Syntax: `debug ipv6 address cache icmp mld nd packet ra`

- *address* - IPv6 address
- *cache* - IPv6 cache entry
- *icmp* - ICMPv6
- *nd* - neighbor discovery
- *packet* - IPv6 packet
- *ra* - router add

Disabling IPv6 on a Layer 2 switch

IPv6 is enabled by default in the Layer 2 switch code. If desired, you can disable IPv6 on a global basis on a device running the switch code. To do so, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)# no ipv6 enable
```

Syntax: `no ipv6 enable`

To re-enable IPv6 after it has been disabled, enter **ipv6 enable**.

NOTE

IPv6 is disabled by default in the router code and must be configured on each interface that will support IPv6.

Configuring IPv6 neighbor discovery

The neighbor discovery feature for IPv6 uses IPv6 ICMP messages to do the following:

- Determine the link-layer address of a neighbor on the same link.
- Verify that a neighbor is reachable.
- Track neighbor routers.

An IPv6 host is required to listen for and recognize the following addresses that identify itself:

- Link-local address.
- Assigned unicast address.
- Loopback address.
- All-nodes multicast address.
- Solicited-node multicast address.
- Multicast address to all other groups to which it belongs.

You can adjust the following IPv6 neighbor discovery features:

- Neighbor solicitation messages for duplicate address detection.
- Amount of time during which an IPv6 node considers a remote node reachable (for use by all nodes on a given link).

Configuration notes

NOTE

For all solicitation and advertisement messages, Dell uses seconds as the unit of measure instead of milliseconds.

- If you add a port to a port-based VLAN, and the port has IPv6 neighbor discovery configuration, the system will clean up the neighbor discovery configuration from the port and display the following message on the console.

```
ND6 port config on the new member ports removed
```
- Neighbor discovery is not supported on tunnel interfaces.

Neighbor solicitation and advertisement messages

Neighbor solicitation and advertisement messages enable a node to determine the link-layer address of another node (neighbor) on the same link. (This function is similar to the function provided by the Address Resolution Protocol [ARP] in IPv4.) For example, node 1 on a link wants to determine the link-layer address of node 2 on the same link. To do so, node 1, the source node, multicasts a neighbor solicitation message. The neighbor solicitation message, which has a value of 135 in the Type field of the ICMP packet header, contains the following information:

- **Source address:** IPv6 address of node 1 interface that sends the message.
- **Destination address:** solicited-node multicast address (FF02:0:0:0:0:1:FF00::/104) that corresponds the IPv6 address of node 2.
- Link-layer address of node 1.
- A query for the link-layer address of node 2.

After receiving the neighbor solicitation message from node 1, node 2 replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header. The neighbor solicitation message contains the following information:

- **Source address:** IPv6 address of the node 2 interface that sends the message.
- **Destination address:** IPv6 address of node 1.
- Link-layer address of node 2.

After node 1 receives the neighbor advertisement message from node 2, nodes 1 and 2 can now exchange packets on the link.

After the link-layer address of node 2 is determined, node 1 can send neighbor solicitation messages to node 2 to verify that it is reachable. Also, nodes 1, 2, or any other node on the same link can send a neighbor advertisement message to the all-nodes multicast address (FF02::1) if there is a change in their link-layer address.

Configuring static neighbor entries

In some special cases, a neighbor cannot be reached using the neighbor discovery feature. In this situation, you can add a static entry to the IPv6 neighbor discovery cache, which causes a neighbor to be reachable at all times without using neighbor discovery. (A static entry in the IPv6 neighbor discovery cache functions like a static ARP entry in IPv4.)

NOTE

A port that has a statically assigned IPv6 entry cannot be added to a VLAN.

NOTE

Static neighbor configurations will be cleared on secondary ports when a trunk is formed.

For example, to add a static entry for a neighbor with the IPv6 address 3001:ffe0:2678:47b and link-layer address 0004.6a2b.8641 that is reachable through Ethernet interface 1, enter the following command.

```
PowerConnect(config)# ipv6 neighbor 3001:ffe0:2678:47b ethernet 1 0004.6a2b.8641
```

Syntax: `[no] ipv6 neighbor ipv6-address ethernet port | ve ve-number [ethernet port] link-layer-address`

The *ipv6-address* parameter specifies the address of the neighbor.

The **ethernet | ve** parameters specify the interface through which to reach a neighbor. If you specify an Ethernet interface, specify the port number of the Ethernet interface. If you specify a VE, specify the VE number and then the Ethernet port numbers associated with the VE. The link-layer address is a 48-bit hardware address of the neighbor.

If you attempt to add an entry that already exists in the neighbor discovery cache, the software changes the already existing entry to a static entry.

To remove a static IPv6 entry from the IPv6 neighbor discovery cache, use the **no** form of this command.

Clearing global IPv6 information

You can clear the following global IPv6 information:

- Entries from the IPv6 cache.
- Entries from the IPv6 neighbor table.
- IPv6 routes from the IPv6 route table.
- IPv6 traffic statistics.

Clearing the IPv6 cache

You can remove all entries from the IPv6 cache or specify an entry based on the following:

- IPv6 prefix.
- IPv6 address.
- Interface type.

For example, to remove entries for IPv6 address 2000:e0ff::1, enter the following command at the Privileged EXEC level or any of the Config levels of the CLI.

```
PowerConnect# clear ipv6 cache 2000:e0ff::1
```

Syntax: `clear ipv6 cache [ipv6-prefix / prefix-length | ipv6-address | ethernet port | tunnel number | ve number]`

You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373. You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.

You must specify the *ipv6-address* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The **ethernet** | **tunnel** | **ve** parameters specify the interfaces for which you can remove cache entries. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE or tunnel interface, also specify the VE or tunnel number, respectively.

Clearing IPv6 neighbor information

You can remove all entries from the IPv6 neighbor table or specify an entry based on the following:

- IPv6 prefix
- IPv6 address
- Interface type

For example, to remove entries for Ethernet interface 1, enter the following command at the Privileged EXEC level or any of the CONFIG levels of the CLI.

```
PowerConnect# clear ipv6 neighbor ethernet 1
```

Syntax: `clear ipv6 neighbor [ipv6-prefix / prefix-length | ipv6-address | ethernet port | ve number]`

- You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373. You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.
- You must specify the *ipv6-address* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The **ethernet** | **ve** parameters specify the interfaces for which you can remove cache entries. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE, also specify the VE number.

Clearing IPv6 traffic statistics

To clear all IPv6 traffic statistics (reset all fields to zero), enter the following command at the Privileged EXEC level or any of the Config levels of the CLI.

```
PowerConnect(config)# clear ipv6 traffic
```

Syntax: clear ipv6 traffic

Displaying global IPv6 information

You can display output for the following global IPv6 parameters:

- IPv6 cache
- IPv6 interfaces
- IPv6 neighbors
- IPv6 route table
- Local IPv6 routers
- IPv6 TCP connections and the status of individual connections
- IPv6 traffic statistics

Displaying IPv6 cache information

The IPv6 cache contains an IPv6 host table that has indices to the next hop gateway and the router interface on which the route was learned.

To display IPv6 cache information, enter the following command at any CLI level.

```
PowerConnect# show ipv6 cache
```

```
Total number of cache entries: 10
```

	IPv6 Address	Next Hop	Port
1	5000:2::2	LOCAL	tunnel 2
2	2000:4::106	LOCAL	ethe 2
3	2000:4::110	DIRECT	ethe 2
4	2002:c0a8:46a::1	LOCAL	ethe 2
5	5005::2e0:52ff:fe99:9737	LOCAL	ethe 2
6	5005::ffff:ffff:feff:ffff	LOCAL	loopback 2
7	5005::c0a8:46a	LOCAL	tunnel 2
8	5005::c0a8:46a	LOCAL	tunnel 6
9	2999::1	LOCAL	loopback 2
10	5005::2e0:52ff:fe99:9700	LOCAL	ethe 1

Syntax: show ipv6 cache [*index-number* | *ipv6-prefix / prefix-length* | *ipv6-address* | **ethernet** *port* | **ve** *number* | **tunnel** *number*]

- The *index-number* parameter restricts the display to the entry for the specified index number and subsequent entries.

5 Displaying global IPv6 information

- The `ipv6-prefix>/ prefix-length` parameter restricts the display to the entries for the specified IPv6 prefix. You must specify the `ipv6-prefix>` parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373. You must specify the `prefix-length` parameter as a decimal value. A slash mark (/) must follow the `ipv6-prefix` parameter and precede the `prefix-length` parameter.
- The `ethernet | ve | tunnel` parameters restrict the display to the entries for the specified interface. The `ipv6-address` parameter restricts the display to the entries for the specified IPv6 address. You must specify this parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE interface, also specify the VE number. If you specify a tunnel interface, also specify the tunnel number.

This display shows the following information.

TABLE 14 IPv6 cache information fields

This field...	Displays...
Total number of cache entries	The number of entries in the cache table.
IPv6 Address	The host IPv6 address.
Next Hop	The next hop, which can be one of the following: <ul style="list-style-type: none">• Direct – The next hop is directly connected to the router.• Local – The next hop is originated on this router.• <code>ipv6 address></code> – The IPv6 address of the next hop.
Port	The port on which the entry was learned.

Displaying IPv6 interface information

To display IPv6 interface information, enter the following command at any CLI level.

```
PowerConnect# show ipv6 interface
Interface      Status      Routing  Global Unicast Address
Ethernet 3     down/down  R
Ethernet 5     down/down
Ethernet 17    up/up      2017::c017:101/64
Ethernet 19    up/up      2019::c019:101/64
VE 4           down/down
VE 14         up/up      2024::c060:101/64
Loopback 1    up/up      ::1/128
Loopback 2    up/up      2005::303:303/128
Loopback 3    up/up
```

Syntax: `show ipv6 interface [interface [port-number | number]]`

- The `interface` parameter displays detailed information for a specified interface. For the interface, you can specify the **Ethernet**, **loopback**, **tunnel**, or **VE** keywords. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a loopback, tunnel, or VE interface, also specify the number associated with the interface.

This display shows the following information.

TABLE 15 General IPv6 interface information fields

This field...	Displays...
Interface	The interface type, and the port number or number of the interface.
Status	The status of the interface. The entry in the Status field will be either “up/up” or “down/down”.
Routing	The routing protocols enabled on the interface.
Global Unicast Address	The global unicast address of the interface.

To display detailed information for a specific interface, enter a command such as the following at any CLI level.

```
PowerConnect# show ipv6 interface ethernet 1
Interface Ethernet 1 is up, line protocol is up
  IPv6 is enabled, link-local address is fe80::2e0:52ff:fe99:97
  Global unicast address(es):
  Joined group address(es):
    ff02::9
    ff02::1:ff99:9700
    ff02::2
    ff02::1
  MTU is 1500 bytes
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 3
  ND reachable time is 30 seconds
  ND advertised reachable time is 0 seconds
  ND retransmit interval is 1 seconds
  ND advertised retransmit interval is 0 seconds
```

This display shows the following information.

TABLE 16 Detailed IPv6 interface information fields

This field...	Displays...
Interface/line protocol status	The status of interface and line protocol. If you have disabled the interface with the disable command, the status will be “administratively down”. Otherwise, the status is either “up” or “down”.
IPv6 status/link-local address	The status of IPv6. The status is either “enabled” or “disabled”. Displays the link-local address, if one is configured for the interface.
Global unicast address(es)	Displays the global unicast address(es), if one or more are configured for the interface.
Joined group address(es)	The multicast address(es) that a router interface listens for and recognizes.
MTU	The setting of the maximum transmission unit (MTU) configured for the IPv6 interface. The MTU is the maximum length an IPv6 packet can have to be transmitted on the interface. If an IPv6 packet is longer than an MTU, the host that originated the packet fragments the packet and transmits its contents in multiple packets that are shorter than the configured MTU.

TABLE 16 Detailed IPv6 interface information fields (Continued)

This field...	Displays...
ICMP	The setting of the ICMP redirect parameter for the interface.
ND	The setting of the various neighbor discovery parameters for the interface.

Displaying IPv6 neighbor information

You can display the IPv6 neighbor table, which contains an entry for each IPv6 neighbor with which the router exchanges IPv6 packets.

To display the IPv6 neighbor table, enter the following command at any CLI level.

```
PowerConnect(config)# show ipv6 neighbor
Total number of Neighbor entries: 3
 IPv6 Address                               LinkLayer-Addr State Age Port  vlan
IsR 5555::55                               0002.0002.0002 *REACH0 e 11 - 0
2000:4::110                                00e0.5291.bb37 REACH 20 e 1 5 1
fe80::2e0:52ff:fe91:bb37                   00e0.5291.bb37 DELAY 1 e2 4 1
fe80::2e0:52ff:fe91:bb40                   00e0.5291.bb40 STALE 5930e 3 5 1
```

Syntax: `show ipv6 neighbor [ipv6-prefix / prefix-length | ipv6-address | interface [port | number]]`

- The *ipv6-prefix* / *prefix-length* parameters restrict the display to the entries for the specified IPv6 prefix. You must specify the *ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373. You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.
- The *ipv6-address* parameter restricts the display to the entries for the specified IPv6 address. You must specify this parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373.
- The *interface* parameter restricts the display to the entries for the specified router interface. For this parameter, you can specify the **Ethernet** or **VE** keywords. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE interface, also specify the VE number.

This display shows the following information.

TABLE 17 IPv6 neighbor information fields

This field...	Displays...
Total number of neighbor entries	The total number of entries in the IPv6 neighbor table.
IPv6 Address	The 128-bit IPv6 address of the neighbor.
Link-Layer Address	The 48-bit interface ID of the neighbor.

TABLE 17 IPv6 neighbor information fields (Continued)

This field...	Displays...
State	The current state of the neighbor. Possible states are as follows: <ul style="list-style-type: none"> • INCOMPLETE – Address resolution of the entry is being performed. • *REACH – The static forward path to the neighbor is functioning properly. • REACH – The forward path to the neighbor is functioning properly. • STALE – This entry has remained unused for the maximum interval. While stale, no action takes place until a packet is sent. • DELAY – This entry has remained unused for the maximum interval, and a packet was sent before another interval elapsed. • PROBE – Neighbor solicitation are transmitted until a reachability confirmation is received.
Age	The number of seconds the entry has remained unused (the default is 30 seconds), the entry is removed from the table.
Port	The physical port on which the entry was learned.
vlan	The VLAN on which the entry was learned.
IsR	Determines if the neighbor is a router or host: 0 – Indicates that the neighbor is a host. 1 – Indicates that the neighbor is a router.

Displaying IPv6 TCP information

You can display the following IPv6 TCP information:

- General information about each TCP connection on the router, including the percentage of free memory for each of the internal TCP buffers.
- Detailed information about a specified TCP connection.

To display general information about each TCP connection on the router, enter the following command at any CLI level.

```
PowerConnect# show ipv6 tcp connections
Local IP address:port -> Remote IP address:port TCP state
192.168.182.110:23 -> 192.168.8.186:4933 ESTABLISHED
192.168.182.110:8218 -> 192.168.182.106:179 ESTABLISHED
192.168.182.110:8039 -> 192.168.2.119:179 SYN-SENT
192.168.182.110:8159 -> 192.168.2.102:179 SYN-SENT
2000:4::110:179 -> 2000:4::106:8222 ESTABLISHED (1440)
Total 5 TCP connections
```

```
TCP MEMORY USAGE PERCENTAGE
FREE TCP = 98 percent
FREE TCP QUEUE BUFFER = 99 percent
FREE TCP SEND BUFFER = 97 percent
FREE TCP RECEIVE BUFFER = 100 percent
FREE TCP OUT OF SEQUENCE BUFFER = 100 percent
```

Syntax: show ipv6 tcp connections

This display shows the following information.

TABLE 18 General IPv6 TCP connection fields

This field...	Displays...
Local IP address:port	The IPv4 or IPv6 address and port number of the local router interface over which the TCP connection occurs.
Remote IP address:port	The IPv4 or IPv6 address and port number of the remote router interface over which the TCP connection occurs.
TCP state	<p>The state of the TCP connection. Possible states include the following:</p> <ul style="list-style-type: none"> • LISTEN – Waiting for a connection request. • SYN-SENT – Waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED – Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED – Data can be sent and received over the connection. This is the normal operational state of the connection. • FIN-WAIT-1 – Waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 – Waiting for a connection termination request from the remote TCP. • CLOSE-WAIT – Waiting for a connection termination request from the local user. • CLOSING – Waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK – Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request). • TIME-WAIT – Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request. • CLOSED – There is no connection state.
FREE TCP = <i>percentage</i> >	The percentage of free TCP control block (TCP) space.
FREE TCP QUEUE BUFFER = <i>percentage</i> >	The percentage of free TCP queue buffer space.
FREE TCP SEND BUFFER = <i>percentage</i> >	The percentage of free TCP send buffer space.
FREE TCP RECEIVE BUFFER = <i>percentage</i> >	The percentage of free TCP receive buffer space.
FREE TCP OUT OF SEQUENCE BUFFER = <i>percentage</i> >	The percentage of free TCP out of sequence buffer space.

To display detailed information about a specified TCP connection, enter a command such as the following at any CLI level.

```
PowerConnect# show ipv6 tcp status 2000:4::110 179 2000:4::106 8222
TCP: TCP = 0x217fc300
TCP: 2000:4::110:179 -> 2000:4::106:8222: state: ESTABLISHED Port: 1
  Send: initial sequence number = 242365900
  Send: first unacknowledged sequence number = 242434080
  Send: current send pointer = 242434080
  Send: next sequence number to send = 242434080
  Send: remote received window = 16384
  Send: total unacknowledged sequence number = 0
  Send: total used buffers 0
  Receive: initial incoming sequence number = 740437769
  Receive: expected incoming sequence number = 740507227
  Receive: received window = 16384
  Receive: bytes in receive queue = 0
  Receive: congestion window = 1459
```

Syntax: `show ipv6 tcp status local-ip-address local-port-number remote-ip-address remote-port-number`

- The *local-ip-address* parameter can be the IPv4 or IPv6 address of the local interface over which the TCP connection is taking place.
- The *local-port-number* parameter is the local port number over which a TCP connection is taking place.
- The *remote-ip-address* parameter can be the IPv4 or IPv6 address of the remote interface over which the TCP connection is taking place.
- The *remote-port-number* parameter is the local port number over which a TCP connection is taking place.

This display shows the following information.

TABLE 19 Specific IPv6 TCP connection fields

This field...	Displays...
TCP = <i>location</i> >	The location of the TCP.
<i>local-ip-address</i> > <i>local-port-number</i> > <i>remote-ip-address</i> > <i>remote-port-number</i> > <i>state</i> > <i>port</i> >	This field provides a general summary of the following: <ul style="list-style-type: none"> • The local IPv4 or IPv6 address and port number. • The remote IPv4 or IPv6 address and port number. • The state of the TCP connection. For information on possible states, refer to Table 18 on page 86. • The port numbers of the local interface.
Send: initial sequence number = <i>number</i> >	The initial sequence number sent by the local router.
Send: first unacknowledged sequence number = <i>number</i> >	The first unacknowledged sequence number sent by the local router.
Send: current send pointer = <i>number</i> >	The current send pointer.
Send: next sequence number to send = <i>number</i> >	The next sequence number sent by the local router.
Send: remote received window = <i>number</i> >	The size of the remote received window.

5 Displaying global IPv6 information

TABLE 19 Specific IPv6 TCP connection fields (Continued)

This field...	Displays...
Send: total unacknowledged sequence number = <i>number</i> >	The total number of unacknowledged sequence numbers sent by the local router.
Send: total used buffers <i>number</i> >	The total number of buffers used by the local router in setting up the TCP connection.
Receive: initial incoming sequence number = <i>number</i> >	The initial incoming sequence number received by the local router.
Receive: expected incoming sequence number = <i>number</i> >	The incoming sequence number expected by the local router.
Receive: received window = <i>number</i> >	The size of the local router receive window.
Receive: bytes in receive queue = <i>number</i> >	The number of bytes in the local router receive queue.
Receive: congestion window = <i>number</i> >	The size of the local router receive congestion window.

Displaying IPv6 traffic statistics

To display IPv6 traffic statistics, enter the following command at any CLI level.

```
PowerConnect# show ipv6 traffic
IP6 Statistics
 36947 received, 66818 sent, 0 forwarded, 36867 delivered, 0 rawout
 0 bad vers, 23 bad scope, 0 bad options, 0 too many hdr
 0 no route, 0 can not forward, 0 redirect sent
 0 frag recv, 0 frag dropped, 0 frag timeout, 0 frag overflow
 0 reassembled, 0 fragmented, 0 ofragments, 0 can not frag
 0 too short, 0 too small, 11 not member
 0 no buffer, 66819 allocated, 21769 freed
 0 forward cache hit, 46 forward cache miss

ICMP6 Statistics
Received:
 0 dest unreachable, 0 pkt too big, 0 time exceeded, 0 param prob
 2 echo req, 1 echo reply, 0 mem query, 0 mem report, 0 mem red
 0 router soli, 2393 router adv, 106 nei soli, 3700 nei adv, 0 redirect
 0 bad code, 0 too short, 0 bad checksum, 0 bad len
 0 reflect, 0 nd toomany opt, 0 badhopcount
Sent:
 0 dest unreachable, 0 pkt too big, 0 time exceeded, 0 param prob
 1 echo req, 2 echo reply, 0 mem query, 0 mem report, 0 mem red
 0 router soli, 2423 router adv, 3754 nei soli, 102 nei adv, 0 redirect
 0 error, 0 can not send error, 0 too freq
Sent Errors:
 0 unreachable no route, 0 admin, 0 beyond scope, 0 address, 0 no port
 0 pkt too big, 0 time exceed transit, 0 time exceed reassembly
 0 param problem header, 0 nexthead, 0 option, 0 redirect, 0 unknown

UDP Statistics
 470 received, 7851 sent, 6 no port, 0 input errors

TCP Statistics
 57913 active opens, 0 passive opens, 57882 failed attempts
 159 active resets, 0 passive resets, 0 input errors
 565189 in segments, 618152 out segments, 171337 retransmission
```

Syntax: show ipv6 traffic

This display shows the following information.

TABLE 20 IPv6 traffic statistics fields

This field...	Displays...
IPv6 statistics	
received	The total number of IPv6 packets received by the router.
sent	The total number of IPv6 packets originated and sent by the router.
forwarded	The total number of IPv6 packets received by the router and forwarded to other routers.
delivered	The total number of IPv6 packets delivered to the upper layer protocol.
rawout	This information is used by Dell Technical Support.
bad vers	The number of IPv6 packets dropped by the router because the version number is not 6.

TABLE 20 IPv6 traffic statistics fields (Continued)

This field...	Displays...
bad scope	The number of IPv6 packets dropped by the router because of a bad address scope.
bad options	The number of IPv6 packets dropped by the router because of bad options.
too many hdr	The number of IPv6 packets dropped by the router because the packets had too many headers.
no route	The number of IPv6 packets dropped by the router because there was no route.
can not forward	The number of IPv6 packets the router could not forward to another router.
redirect sent	This information is used by Dell Technical Support.
frag rcv	The number of fragments received by the router.
frag dropped	The number of fragments dropped by the router.
frag timeout	The number of fragment timeouts that occurred.
frag overflow	The number of fragment overflows that occurred.
reassembled	The number of fragmented IPv6 packets that the router reassembled.
fragmented	The number of IPv6 packets fragmented by the router to accommodate the MTU of this router or of another device.
ofragments	The number of output fragments generated by the router.
can not frag	The number of IPv6 packets the router could not fragment.
too short	The number of IPv6 packets dropped because they are too short.
too small	The number of IPv6 packets dropped because they do not have enough data.
not member	The number of IPv6 packets dropped because the recipient is not a member of a multicast group.
no buffer	The number of IPv6 packets dropped because there is no buffer available.
forward cache miss	The number of IPv6 packets received for which there is no corresponding cache entry.
ICMP6 statistics	
Some ICMP statistics apply to both Received and Sent, some apply to Received only, some apply to Sent only, and some apply to Sent Errors only.	
Applies to received and sent	
dest unreachable	The number of Destination Unreachable messages sent or received by the router.
pkt too big	The number of Packet Too Big messages sent or received by the router.
time exceeded	The number of Time Exceeded messages sent or received by the router.
param prob	The number of Parameter Problem messages sent or received by the router.
echo req	The number of Echo Request messages sent or received by the router.
echo reply	The number of Echo Reply messages sent or received by the router.
mem query	The number of Group Membership Query messages sent or received by the router.
mem report	The number of Membership Report messages sent or received by the router.
mem red	The number of Membership Reduction messages sent or received by the router.

TABLE 20 IPv6 traffic statistics fields (Continued)

This field...	Displays...
router soli	The number of Router Solicitation messages sent or received by the router.
router adv	The number of Router Advertisement messages sent or received by the router.
nei soli	The number of Neighbor Solicitation messages sent or received by the router.
nei adv	The number of Router Advertisement messages sent or received by the router.
redirect	The number of redirect messages sent or received by the router.
Applies to received only	
bad code	The number of Bad Code messages received by the router.
too short	The number of Too Short messages received by the router.
bad checksum	The number of Bad Checksum messages received by the router.
bad len	The number of Bad Length messages received by the router.
nd toomany opt	The number of Neighbor Discovery Too Many Options messages received by the router.
badhopcount	The number of Bad Hop Count messages received by the router.
Applies to sent only	
error	The number of Error messages sent by the router.
can not send error	The number of times the node encountered errors in ICMP error messages.
too freq	The number of times the node has exceeded the frequency of sending error messages.
Applies to sent errors only	
unreach no route	The number of Unreachable No Route errors sent by the router.
admin	The number of Admin errors sent by the router.
beyond scope	The number of Beyond Scope errors sent by the router.
address	The number of Address errors sent by the router.
no port	The number of No Port errors sent by the router.
pkt too big	The number of Packet Too Big errors sent by the router.
time exceed transit	The number of Time Exceed Transit errors sent by the router.
time exceed reassembly	The number of Time Exceed Reassembly errors sent by the router.
param problem header	The number of Parameter Problem Header errors sent by the router.
nextheader	The number of Next Header errors sent by the router.
option	The number of Option errors sent by the router.
redirect	The number of Redirect errors sent by the router.
unknown	The number of Unknown errors sent by the router.
UDP statistics	
received	The number of UDP packets received by the router.
sent	The number of UDP packets sent by the router.
no port	The number of UDP packets dropped because the packet did not contain a valid UDP port number.

5 Displaying global IPv6 information

TABLE 20 IPv6 traffic statistics fields (Continued)

This field...	Displays...
input errors	This information is used by Dell Technical Support.
TCP statistics	
active opens	The number of TCP connections opened by the router by sending a TCP SYN to another device.
passive opens	The number of TCP connections opened by the router in response to connection requests (TCP SYNs) received from other devices.
failed attempts	This information is used by Dell Technical Support.
active resets	The number of TCP connections the router reset by sending a TCP RESET message to the device at the other end of the connection.
passive resets	The number of TCP connections the router reset because the device at the other end of the connection sent a TCP RESET message.
input errors	This information is used by Dell Technical Support.
in segments	The number of TCP segments received by the router.
out segments	The number of TCP segments sent by the router.
retransmission	The number of segments that the router retransmitted because the retransmission timer for the segment had expired before the device at the other end of the connection had acknowledged receipt of the segment.

Configuring Spanning Tree Protocol (STP) Related Features **6**

STP overview

The Spanning Tree Protocol (STP) eliminates Layer 2 loops in networks, by selectively blocking some ports and allowing other ports to forward traffic, based on global (bridge) and local (port) parameters you can configure.

STP related features, such as RSTP and PVST, extend the operation of standard STP, enabling you to fine-tune standard STP and avoid some of its limitations.

You can enable or disable STP on a global basis (for the entire device), a port-based VLAN basis (for the individual Layer 2 broadcast domain), or an individual port basis.

Configuration procedures are provided for the standard STP bridge and port parameters as well as features listed in [Table 26](#).

Configuring standard STP parameters

PowerConnect devices support standard STP as described in the IEEE 802.1D specification.

By default, each port-based VLAN on a PowerConnect device runs a separate spanning tree (a separate instance of STP). A PowerConnect device has one port-based VLAN (VLAN 1) by default that contains all the device ports. Thus, by default each PowerConnect device has one spanning tree. However, if you configure additional port-based VLANs on a PowerConnect device, then each of those VLANs on which STP is enabled and VLAN 1 all run separate spanning trees.

If you configure a port-based VLAN on the device, the VLAN has the same STP state as the default STP state on the device. Thus, on Layer 2 Switches, new VLANs have STP enabled by default. On Layer 3 Switches, new VLANs have STP disabled by default. You can enable or disable STP in each VLAN separately. In addition, you can enable or disable STP on individual ports.

STP parameters and defaults

[Table 21](#) lists the default STP states for PowerConnect devices.

TABLE 21 Default STP states

Device type	Default STP type	Default STP state	Default STP state of new VLANs ¹
Layer 2 Switch	MSTP ²	Enabled	Enabled
Layer 3 Switch	MSTP	Disabled	Disabled

1. When you create a port-based VLAN, the new VLAN STP state is the same as the default STP state on the device. The new VLAN does not inherit the STP state of the default VLAN.

6 Configuring standard STP parameters

2. MSTP stands for “Multiple Spanning Tree Protocol”. In this type of STP, each port-based VLAN, including the default VLAN, has its own spanning tree. References in this documentation to “STP” apply to MSTP. The Single Spanning Tree Protocol (SSTP) is another type of STP. SSTP includes all VLANs on which STP is enabled in a single spanning tree. Refer to [“Single Spanning Tree \(SSTP\)”](#) on page 148.

[Table 22](#) lists the default STP bridge parameters. The bridge parameters affect the entire spanning tree. If you are using MSTP, the parameters affect the VLAN. If you are using SSTP, the parameters affect all VLANs that are members of the single spanning tree.

TABLE 22 Default STP bridge parameters

Parameter	Description	Default and valid values
Forward Delay	The period of time spent by a port in the listening and learning state before moving on to the learning or forwarding state, respectively. The forward delay value is also used for the age time of dynamic entries in the filtering database, when a topology change occurs.	15 seconds Possible values: 4 – 30 seconds
Maximum Age	The interval a bridge will wait for a configuration BPDU from the root bridge before initiating a topology change.	20 seconds Possible values: 6 – 40 seconds
Hello Time	The interval of time between each configuration BPDU sent by the root bridge.	2 seconds Possible values: 1 – 10 seconds
Priority	A parameter used to identify the root bridge in a spanning tree (instance of STP). The bridge with the lowest value has the highest priority and is the root. A higher numerical value means a lower priority; thus, the highest priority is 0.	32768 Possible values: 0 – 65535

NOTE

If you plan to change STP bridge timers, Dell recommends that you stay within the following ranges, from section 8.10.2 of the IEEE STP specification.

$$2 * (\text{forward_delay} - 1) \geq \text{max_age}$$

$$\text{max_age} \geq 2 * (\text{hello_time} + 1)$$

[Table 23](#) lists the default STP port parameters. The port parameters affect individual ports and are separately configurable on each port.

TABLE 23 Default STP port parameters

Parameter	Description	Default and valid values
Priority	The preference that STP gives this port relative to other ports for forwarding traffic out of the spanning tree. A higher numerical value means a lower priority.	128 Possible values in PowerConnect B-Series T124X (configurable in increments of 16)
Path Cost	The cost of using the port to reach the root bridge. When selecting among multiple links to the root bridge, STP chooses the link with the lowest path cost and blocks the other paths. Each port type has its own default STP path cost.	10 Mbps – 100 100 Mbps – 19 Gbps – 4 10 Gbps – 2 Possible values are 0 – 65535

Enabling or disabling the Spanning Tree Protocol (STP)

STP is *enabled* by default on devices running Layer 2 code. STP is *disabled* by default on devices running Layer 3 code.

You can enable or disable STP on the following levels:

- **Globally** – Affects all ports and port-based VLANs on the device.
- **Port-based VLAN** – Affects all ports within the specified port-based VLAN. When you enable or disable STP within a port-based VLAN, the setting overrides the global setting. Thus, you can enable STP for the ports within a port-based VLAN even when STP is globally disabled, or disable the ports within a port-based VLAN when STP is globally enabled.
- **Individual port** – Affects only the individual port. However, if you change the STP state of the primary port in a trunk group, the change affects all ports in the trunk group.

NOTE

The CLI converts the STP groups into topology groups when you save the configuration. For backward compatibility, you can still use the STP group commands. However, the CLI converts the commands into the topology group syntax. Likewise, the **show stp-group** command displays STP topology groups.

Enabling or disabling STP globally

Use the following method to enable or disable STP on a device on which you have not configured port-based VLANs.

NOTE

When you configure a VLAN, the VLAN inherits the global STP settings. However, once you begin to define a VLAN, you can no longer configure standard STP parameters globally using the CLI. From that point on, you can configure STP only within individual VLANs.

To enable STP for all ports in all VLANs on a device, enter the following command.

```
PowerConnect(config)#spanning-tree
```

This command enables a separate spanning tree in each VLAN, including the default VLAN.

Syntax: [no] spanning-tree

Enabling or disabling STP in a port-based VLAN

Use the following procedure to disable or enable STP on a device on which you have configured a port-based VLAN. Changing the STP state in a VLAN affects only that VLAN.

To enable STP for all ports in a port-based VLAN, enter commands such as the following.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#spanning-tree
```

Syntax: [no] spanning-tree

Enabling or disabling STP on an individual port

Use the following procedure to disable or enable STP on an individual port.

NOTE

If you change the STP state of the primary port in a trunk group, it affects all ports in the trunk group.

To enable STP on an individual port, enter commands such as the following.

```
PowerConnect(config)#interface 1
PowerConnect(config-if-e10000-1)#spanning-tree
```

Syntax: [no] spanning-tree

Changing STP bridge and port parameters

[Table 22](#) on page 94 and [Table 23](#) on page 95 list the default STP parameters. If you need to change the default value for an STP parameter, use the following procedures.

Changing STP bridge parameters

NOTE

If you plan to change STP bridge timers, Dell recommends that you stay within the following ranges, from section 8.10.2 of the IEEE STP specification.

$$2 * (\text{forward_delay} - 1) \geq \text{max_age}$$
$$\text{max_age} \geq 2 * (\text{hello_time} + 1)$$

To change a STP bridge priority on a Dell device to the highest value to make the device the root bridge, enter the following command.

```
PowerConnect(config)#spanning-tree priority 0
```

The command in this example changes the priority on a device on which you have not configured port-based VLANs. The change applies to the default VLAN. If you have configured a port-based VLAN on the device, you can configure the parameters only at the configuration level for individual VLANs. Enter commands such as the following.

```
PowerConnect(config)#vlan 20
PowerConnect(config-vlan-20)#spanning-tree priority 0
```

To make this change in the default VLAN, enter the following commands.

```
PowerConnect(config)#vlan 1
PowerConnect(config-vlan-1)#spanning-tree priority 0
```

Syntax: [no] spanning-tree [forward-delay <value>] | [hello-time <value>] | [maximum-age <value>] | [priority <value>]

The **forward-delay** <value> parameter specifies the forward delay and can be a value from 4 – 30 seconds. The default is 15 seconds.

NOTE

You can configure a device for faster convergence (including a shorter forward delay) using Fast Span. Refer to “[Configuring STP related features](#)” on page 106.

The **hello-time** <value> parameter specifies the hello time and can be a value from 1 – 10 seconds. The default is 2 seconds.

NOTE

This parameter applies only when this device or VLAN is the root bridge for its spanning tree.

The **maximum-age** <value> parameter specifies the amount of time the device waits for receipt of a configuration BPDU from the root bridge before initiating a topology change. You can specify from 6 – 40 seconds. The default is 20 seconds.

The **priority** <value> parameter specifies the priority and can be a value from 0 – 65535. A higher numerical value means a lower priority. Thus, the highest priority is 0. The default is 32768.

You can specify some or all of these parameters on the same command line. If you specify more than one parameter, you must specify them in the order shown above, from left to right.

Changing STP port parameters

To change the path and priority costs for a port, enter commands such as the following.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#spanning-tree ethernet 5 path-cost 15 priority 64
```

Syntax: spanning-tree ethernet<portnum> path-cost <value> | priority <value> | disable | enable

The <portnum> parameter specifies the interface.

The **path-cost** <value> parameter specifies the port cost as a path to the spanning tree root bridge. STP prefers the path with the lowest cost. You can specify a value from 0 – 65535.

The default depends on the port type:

- 10 Mbps – 100
- 100 Mbps – 19
- Gbps – 4
- 10 Gbps – 2

The **priority** <value> parameter specifies the preference that STP gives this port relative to other ports for forwarding traffic out of the spanning tree. The value you can specify depends on the software version running on the device, as follows:

The **disable | enable** parameter disables or re-enables STP on the port. The STP state change affects only this VLAN. The port STP state in other VLANs is not changed.

STP protection enhancement

STP protection provides the ability to prohibit an end station from initiating or participating in an STP topology change.

The 802.1W Spanning Tree Protocol (STP) detects and eliminates logical loops in a redundant network by selectively blocking some data paths (ports) and allowing only the best data paths to forward traffic.

In an STP environment, switches, end stations, and other Layer 2 devices use Bridge Protocol Data Units (BPDUs) to exchange information that STP will use to determine the best path for data flow. When a Layer 2 device is powered ON and connected to the network, or when a Layer 2 device goes down, it sends out an STP BPDU, triggering an STP topology change.

In some instances, it is unnecessary for a connected device, such as an end station, to initiate or participate in an STP topology change. In this case, you can enable the STP Protection feature on the device port to which the end station is connected. STP Protection disables the connected device ability to initiate or participate in an STP topology change, by dropping all BPDUs received from the connected device.

Enabling STP protection

You can enable STP Protection on a per-port basis.

To prevent an end station from initiating or participating in STP topology changes, enter the following command at the Interface level of the CLI.

```
PowerConnect#(config) interface e 2  
PowerConnect#(config-if-e10000-2)#stp-protect
```

This command causes the port to drop STP BPDUs sent from the device on the other end of the link.

Syntax: [no] stp-protect

The *portnum* parameter is a valid port number.

Enter the **no** form of the command to disable STP protection on the port.

Clearing BPDU drop counters

For each port that has STP Protection enabled, the device counts and records the number of dropped BPDUs. You can use CLI commands to clear the BPDU drop counters for all ports on the device, or for a specific port on the device.

To clear the BPDU drop counters for all ports on the device that have STP Protection enabled, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)#clear stp-protect-statistics
```

To clear the BPDU drop counter for a specific port that has STP Protection enabled, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect#clear stp-protect-statistics e 2
```

Syntax: clear stp-protect-statistics [ethernet <port-num>]

The *portnum* parameter is a valid port number.

Viewing the STP Protection Configuration

You can view the STP Protection configuration for all ports on a device, or for a specific port only. The **show stp-protect** command output shows the port number on which STP Protection is enabled, and the number of BPDUs dropped by each port.

To view the STP Protection configuration for all ports on the device, enter the following command at any level of the CLI.

```
PowerConnect#show stp-protect
Port ID      BDU Drop Count
  3          478
  5          213
  6           0
 12          31
```

To view STP Protection configuration for a specific port, enter the following command at any level of the CLI.

```
PowerConnect#show stp-protect e 3
STP-protect is enabled on port 3.  BDU drop count is 478
```

If you enter the **show stp-protect** command for a port that does not have STP protection enabled, the following message displays on the console.

```
PowerConnect#show stp-protect e 4
STP-protect is not enabled on port 4.
```

Syntax: `show stp-protect [ethernet <portnum>]`

Displaying STP information

You can display the following STP information:

- All the global and interface STP settings
- Detailed STP information for each interface
- STP state information for a port-based VLAN
- STP state information for an individual interface

Displaying STP information for an entire device

To display STP information, enter the following command at any level of the CLI.

6 Configuring standard STP parameters

```
PowerConnect#show span
```

```
VLAN 1 BPDU cam_index is 3 and the Master DMA Are(HEX)
STP instance owned by VLAN 1
```

```
Global STP (IEEE 802.1D) Parameters:
```

VLAN ID	Root ID	Root Cost	Root Port	Prio rity	Max Age	He llo	Ho- ld	Fwd dly	Last Chang	Chg cnt	Bridge Address
				Hex	sec	sec	sec	sec	sec		
1	800000e0804d4a00	0	Root	8000	20	2	1	15	689	1	00e0804d4a00

```
Port STP Parameters:
```

Port Num	Prio rity	Path Cost	State	Fwd Trans	Design Cost	Designated Root	Designated Bridge
		Hex					
1	80	19	FORWARDING	1	0	800000e0804d4a00	800000e0804d4a00
2	80	0	DISABLED	0	0	0000000000000000	0000000000000000
3	80	0	DISABLED	0	0	0000000000000000	0000000000000000
4	80	0	DISABLED	0	0	0000000000000000	0000000000000000
5	80	19	FORWARDING	1	0	800000e0804d4a00	800000e0804d4a00
6	80	19	BLOCKING	0	0	800000e0804d4a00	800000e0804d4a00
7	80	0	DISABLED	0	0	0000000000000000	0000000000000000

<lines for remaining ports excluded for brevity>

Syntax: `show span [vlan <vlan-id>] | [pvst-mode] | [<num>] | [detail [vlan <vlan-id> [ethernet <portnum>] | <num>]]`

The **vlan <vlan-id>** parameter displays STP information for the specified port-based VLAN.

The **pvst-mode** parameter displays STP information for the device Per VLAN Spanning Tree (PVST+) compatibility configuration. Refer to “[PVST/PVST+ compatibility](#)” on page 150

The **portnum** parameter is a valid port number.

The **<num>** parameter displays only the entries after the number you specify. For example, on a device with three port-based VLANs, if you enter 1, then information for the second and third VLANs is displayed, but information for the first VLAN is not displayed. Information is displayed according to VLAN number, in ascending order. The entry number is not the same as the VLAN number. For example, if you have port-based VLANs 1, 10, and 2024, then the command output has three STP entries. To display information for VLANs 10 and 2024 only, enter **show span 1**.

The **detail** parameter and its additional optional parameters display detailed information for individual ports. Refer to “[Displaying detailed STP information for each interface](#)” on page 102.

The **show span** command shows the following information.

TABLE 24 CLI display of STP information

This field...	Displays...
Global STP parameters	
VLAN ID	The port-based VLAN that contains this spanning tree (instance of STP). VLAN 1 is the default VLAN. If you have not configured port-based VLANs on this device, all STP information is for VLAN 1.
Root ID	The ID assigned by STP to the root bridge for this spanning tree.

TABLE 24 CLI display of STP information (Continued)

This field...	Displays...
Root Cost	The cumulative cost from this bridge to the root bridge. If this device is the root bridge, then the root cost is 0.
Root Port	The port on this device that connects to the root bridge. If this device is the root bridge, then the value is "Root" instead of a port number.
Priority Hex	This device or VLAN STP priority. The value is shown in hexadecimal format. NOTE: If you configure this value, specify it in decimal format. Refer to "Changing STP bridge parameters" on page 96.
Max age sec	The number of seconds this device or VLAN waits for a configuration BPDU from the root bridge before deciding the root has become unavailable and performing a reconvergence.
Hello sec	The interval between each configuration BPDU sent by the root bridge.
Hold sec	The minimum number of seconds that must elapse between transmissions of consecutive Configuration BPDUs on a port.
Fwd dly sec	The number of seconds this device or VLAN waits following a topology change and consequent reconvergence.
Last Chang sec	The number of seconds since the last time a topology change occurred.
Chg cnt	The number of times the topology has changed since this device was reloaded.
Bridge Address	The STP address of this device or VLAN. NOTE: If this address is the same as the Root ID, then this device or VLAN is the root bridge for its spanning tree.
Port STP parameters	
Port Num	The port number.
Priority Hex	The port STP priority, in hexadecimal format. NOTE: If you configure this value, specify it in decimal format. Refer to "Changing STP port parameters" on page 97.
Path Cost	The port STP path cost.
State	The port STP state. The state can be one of the following: <ul style="list-style-type: none"> • BLOCKING – STP has blocked Layer 2 traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is FORWARDING. When a port is in this state, the port does not transmit or receive user frames, but the port does continue to receive STP BPDUs. • DISABLED – The port is not participating in STP. This can occur when the port is disconnected or STP is disabled on the port. • FORWARDING – STP is allowing the port to send and receive frames. • LISTENING – STP is responding to a topology change and this port is listening for a BPDU from neighboring bridges in order to determine the new topology. No user frames are transmitted or received during this state. • LEARNING – The port has passed through the LISTENING state and will change to the FORWARDING state, depending on the results of STP reconvergence. The port does not transmit or receive user frames during this state. However, the device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table.
Fwd Trans	The number of times STP has changed the state of this port between BLOCKING and FORWARDING.

TABLE 24 CLI display of STP information (Continued)

This field...	Displays...
Design Cost	The cost to the root bridge as advertised by the designated bridge that is connected to this port. If the designated bridge is the root bridge itself, then the cost is 0. The identity of the designated bridge is shown in the Design Bridge field.
Designated Root	The root bridge as recognized on this port. The value is the same as the root bridge ID listed in the Root ID field.
Designated Bridge	The designated bridge to which this port is connected. The designated bridge is the device that connects the network segment on the port to the root bridge.

Displaying the STP state of a port-based VLAN

When you display information for a port-based VLAN, that information includes the STP state of the VLAN.

To display information for a port-based VLAN, enter a command such as the following at any level of the CLI. The STP state is shown in bold type in this example.

```
PowerConnect#show vlans

Total PORT-VLAN entries: 2
Maximum PORT-VLAN entries: 16

legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree On
  Untagged Ports: (S3) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
  Untagged Ports: (S3) 17 18 19 20 21 22 23 24
  Untagged Ports: (S4) 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17
  Untagged Ports: (S4) 18 19 20 21 22 23 24
  Tagged Ports: None
  Uplink Ports: None

PORT-VLAN 2, Name greenwell, Priority level0, Spanning tree Off
  Untagged Ports: (S1) 1 2 3 4 5 6 7 8
  Untagged Ports: (S4) 1
  Tagged Ports: None
  Uplink Ports: None
```

Syntax: `show vlan [<vlan-id> | ethernet [<portnum>]`

The `<vlan-id>` parameter specifies a VLAN for which you want to display the configuration information.

The `<portnum>` parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port. .

Displaying detailed STP information for each interface

To display the detailed STP information, enter the following command at any level of the CLI.

```
PowerConnect#show span detail
=====
VLAN 1 - MULTIPLE SPANNING TREE (MSTP) ACTIVE
=====
Bridge identifier      - 0x800000e0804d4a00
Active global timers - Hello: 0

Port 1 is FORWARDING
  Port - Path cost: 19, Priority: 128, Root: 0x800000e052a9bb00
  Designated - Bridge: 0x800000e052a9bb00, Interface: 1, Path cost: 0
  Active Timers - None
  BPDUs - Sent: 11, Received: 0
Port 2 is DISABLED
Port 3 is DISABLED
Port 4 is DISABLED
<lines for remaining ports excluded for brevity>
```

If a port is disabled, the only information shown by this command is “DISABLED”. If a port is enabled, this display shows the following information.

Syntax: `show span detail [vlan <vlan-id> [ethernet []<portnum> | <num>]`

The `vlan <vlan-id>` parameter specifies a VLAN.

The `<portnum>` parameter specifies an individual port within the VLAN (if specified).

. The `<portnum>` parameter is a valid port number.

The `<num>` parameter specifies the number of VLANs you want the CLI to skip before displaying detailed STP information. For example, if the device has six VLANs configured (VLAN IDs 1, 2, 3, 99, 128, and 256) and you enter the command **show span detail 4**, detailed STP information is displayed for VLANs 128 and 256 only.

NOTE

If the configuration includes VLAN groups, the **show span detail** command displays the master VLANs of each group but not the member VLANs within the groups. However, the command does indicate that the VLAN is a master VLAN. The **show span detail vlan <vlan-id>** command displays the information for the VLAN even if it is a member VLAN. To list all the member VLANs within a VLAN group, enter the **show vlan-group [<group-id>]** command.

The **show span detail** command shows the following information.

TABLE 25 CLI display of detailed STP information for ports

This field...	Displays...
Active Spanning Tree protocol	The VLAN that contains the listed ports and the active Spanning Tree protocol. The STP type can be one of the following: <ul style="list-style-type: none"> • MULTIPLE SPANNNG TREE (MSTP) • GLOBAL SINGLE SPANNING TREE (SSTP) NOTE: If STP is disabled on a VLAN, the command displays the following message instead: “Spanning-tree of port-vlan <vlan-id> is disabled.”
Bridge identifier	The STP identity of this device.

6 Configuring standard STP parameters

TABLE 25 CLI display of detailed STP information for ports (Continued)

This field...	Displays...
Active global timers	<p>The global STP timers that are currently active, and their current values. The following timers can be listed:</p> <ul style="list-style-type: none"> • Hello – The interval between Hello packets. This timer applies only to the root bridge. • Topology Change (TC) – The amount of time during which the topology change flag in Hello packets will be marked, indicating a topology change. This timer applies only to the root bridge. • Topology Change Notification (TCN) – The interval between Topology Change Notification packets sent by a non-root bridge toward the root bridge. This timer applies only to non-root bridges.
Port number and STP state	<p>The internal port number and the port STP state.</p> <p>The internal port number is one of the following:</p> <ul style="list-style-type: none"> • The port interface number, if the port is the designated port for the LAN. • The interface number of the designated port from the received BPDU, if the interface is not the designated port for the LAN. <p>The state can be one of the following:</p> <ul style="list-style-type: none"> • BLOCKING – STP has blocked Layer 2 traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is FORWARDING. When a port is in this state, the port does not transmit or receive user frames, but the port does continue to receive STP BPDUs. • DISABLED – The port is not participating in STP. This can occur when the port is disconnected or STP is administratively disabled on the port. • FORWARDING – STP is allowing the port to send and receive frames. • LISTENING – STP is responding to a topology change and this port is listening for a BPDU from neighboring bridges in order to determine the new topology. No user frames are transmitted or received during this state. • LEARNING – The port has passed through the LISTENING state and will change to the BLOCKING or FORWARDING state, depending on the results of STP reconvergence. The port does not transmit or receive user frames during this state. However, the device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table. <p>NOTE: If the state is DISABLED, no further STP information is displayed for the port.</p>
Port Path cost	The STP path cost for the port.
Port Priority	This STP priority for the port. The value is shown as a hexadecimal number.
Root	The ID assigned by STP to the root bridge for this spanning tree.
Designated Bridge	The MAC address of the designated bridge to which this port is connected. The designated bridge is the device that connects the network segment on the port to the root bridge.
Designated Port	The port number sent from the designated bridge.
Designated Path Cost	The cost to the root bridge as advertised by the designated bridge that is connected to this port. If the bridge is the root bridge itself, then the cost is 0. The identity of the designated bridge is shown in the Designated Bridge field.

TABLE 25 CLI display of detailed STP information for ports (Continued)

This field...	Displays...
Active Timers	The current values for the following timers, if active: <ul style="list-style-type: none"> • Message age – The number of seconds this port has been waiting for a hello message from the root bridge. • Forward delay – The number of seconds that have passed since the last topology change and consequent reconvergence. • Hold time – The number of seconds that have elapsed since transmission of the last Configuration BPDU.
BPDUs Sent and Received	The number of BPDUs sent and received on this port since the software was reloaded.

Displaying detailed STP information for a single port in a specific VLAN

Enter a command such as the following to display STP information for an individual port in a specific VLAN.

```
PowerConnect#show span detail vlan 1 ethernet 1
Port 1 is FORWARDING
  Port - Path cost: 19, Priority: 128, Root: 0x800000e052a9bb00
  Designated - Bridge: 0x800000e052a9bb00, Interface: 7, Path cost: 0
  Active Timers - None
  BPDUs - Sent: 29, Received: 0
```

Syntax: `show span detail [vlan <vlan-id> [ethernet <portnum>] <num>]`

Displaying STP state information for an individual interface

To display STP state information for an individual port, you can use the methods in [“Displaying STP information for an entire device”](#) on page 99 or [“Displaying detailed STP information for each interface”](#) on page 102. You also can display STP state information for a specific port using the following method.

To display information for a specific port, enter a command such as the following at any level of the CLI.

6 Configuring STP related features

```
PowerConnect#show interface ethernet 11

FastEthernet11 is up, line protocol is up
  Hardware is FastEthernet, address is 00e0.52a9.bb49 (bia 00e0.52a9.bb49)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Member of L2 VLAN ID 1, port is untagged, port state is FORWARDING
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  MTU 1518 bytes, encapsulation ethernet
  5 minute input rate: 352 bits/sec, 0 packets/sec, 0.00% utilization
  5 minute output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  1238 packets input, 79232 bytes, 0 no buffer
  Received 686 broadcasts, 0 runts, 0 giants
  0 input errors, 0 CRC, 0 frame, 0 ignored
  529 multicast
  918 packets output, 63766 bytes, 0 underruns
  0 output errors, 0 collisions
```

The STP information is shown in bold type in this example.

Syntax: show interfaces [ethernet <portnum>] | [loopback <num>] | [ve <num>] | [brief]

You also can display the STP states of all ports by entering a command such as the following, which uses the **brief** parameter.

```
PowerConnect#show interface brief

Port  Link State      Dupl Speed Trunk Tag Priori MAC          Name
1     Down None          None None  None No  level0 00e0.52a9.bb00
2     Down None          None None  None No  level0 00e0.52a9.bb01
3     Down None          None None  None No  level0 00e0.52a9.bb02
4     Down None          None None  None No  level0 00e0.52a9.bb03
5     Down None          None None  None No  level0 00e0.52a9.bb04
6     Down None          None None  None No  level0 00e0.52a9.bb05
7     Down None          None None  None No  level0 00e0.52a9.bb06
8     Down None          None None  None No  level0 00e0.52a9.bb07

.
. some rows omitted for brevity
.
10    Down None          None None  None No  level0 00e0.52a9.bb4a
11    Up   Forward       Full 100M  None No  level0 00e0.52a9.bb49
```

In the example above, only one port, 11, is forwarding traffic toward the root bridge.

Configuring STP related features

STP features extend the operation of standard STP, enabling you to fine tune standard STP and avoid some of its limitations.

This section describes how to configure these parameters on Layer 3 Switches using the CLI.

802.1W Rapid Spanning Tree (RSTP)

Rapid Spanning Tree Protocol (RSTP), which was 802.1W Draft 3, provided only a subset of the IEEE 802.1W standard; whereas the 802.1W RSTP feature provides the full standard. The implementation of the 802.1W Draft 3 is referred to as RSTP Draft 3.

RSTP Draft3 will continue to be supported on devices for backward compatibility. However, customers who are currently using RSTP Draft 3 should migrate to 802.1W.

The 802.1W feature provides rapid traffic reconvergence for point-to-point links within a few milliseconds (0 – 500 milliseconds), following the failure of a bridge or bridge port. This reconvergence occurs more rapidly than the reconvergence provided by the 802.1D Spanning Tree Protocol (STP) or by RSTP Draft 3.

NOTE

This rapid convergence will not occur on ports connected to shared media devices, such as hubs. To take advantage of the rapid convergence provided by 802.1W, make sure to explicitly configure all point-to-point links in a topology.

The convergence provided by the standard 802.1W protocol occurs more rapidly than the convergence provided by previous spanning tree protocols because of the following:

- Classic or legacy 802.1D STP protocol requires a newly selected Root port to go through listening and learning stages before traffic convergence can be achieved. The 802.1D traffic convergence time is calculated using the following formula.
$$2 \times \text{FORWARD_DELAY} + \text{BRIDGE_MAX_AGE}.$$

If default values are used in the parameter configuration, convergence can take up to 50 seconds. (In this document STP will be referred to as 802.1D.)
- RSTP Draft 3 works only on bridges that have Alternate ports, which are the precalculated “next best root port”. (Alternate ports provide back up paths to the root bridge.) Although convergence occurs from 0 – 500 milliseconds in RSTP Draft 3, the spanning tree topology reverts to the 802.1D convergence if an Alternate port is not found.
- Convergence in 802.1w bridge is not based on any timer values. Rather, it is based on the explicit handshakes between Designated ports and their connected Root ports to achieve convergence in less than 500 milliseconds.

Bridges and bridge port roles

A bridge in an 802.1W rapid spanning tree topology is assigned as the root bridge if it has the highest priority (lowest bridge identifier) in the topology. Other bridges are referred to as non-root bridges.

Unique roles are assigned to ports on the root and non-root bridges. Role assignments are based on the following information contained in the Rapid Spanning Tree Bridge Packet Data Unit (RST BPDUs):

- Root bridge ID
- Path cost value
- Transmitting bridge ID
- Designated port ID

The 802.1W algorithm uses this information to determine if the RST BPDU received by a port is superior to the RST BPDU that the port transmits. The two values are compared in the order as given above, starting with the Root bridge ID. The RST BPDU with a lower value is considered superior. The superiority and inferiority of the RST BPDU is used to assign a role to a port.

If the value of the received RST BPDU is the same as that of the transmitted RST BPDU, then the port ID in the RST BPDUs are compared. The RST BPDU with the lower port ID is superior. Port roles are then calculated appropriately.

The port role is included in the BPDU that it transmits. The BPDU transmitted by an 802.1W port is referred to as an RST BPDU, while it is operating in 802.1W mode.

Ports can have one of the following roles:

- **Root** – Provides the lowest cost path to the root bridge from a specific bridge
- **Designated** – Provides the lowest cost path to the root bridge from a LAN to which it is connected
- **Alternate** – Provides an alternate path to the root bridge when the root port goes down
- **Backup** – Provides a backup to the LAN when the Designated port goes down
- **Disabled** – Has no role in the topology

Assignment of port roles

At system start-up, all 802.1W-enabled bridge ports assume a Designated role. Once start-up is complete, the 802.1W algorithm calculates the superiority or inferiority of the RST BPDU that is received and transmitted on a port.

On a root bridge, each port is assigned a **Designated port** role, except for ports on the same bridge that are physically connected together. In these type of ports, the port that receives the superior RST BPDU becomes the **Backup port**, while the other port becomes the **Designated port**.

On non-root bridges, ports are assigned as follows:

- The port that receives the RST BPDU with the lowest path cost from the root bridge becomes the **Root port**.
- If two ports on the same bridge are physically connected, the port that receives the superior RST BPDU becomes the **Backup port**, while the other port becomes the **Designated port**.
- If a non-root bridge already has a Root port, then the port that receives an RST BPDU that is superior to those it can transmit becomes the **Alternate port**.
- If the RST BPDU that a port receives is inferior to the RST BPDUs it transmits, then the port becomes a **Designated port**.
- If the port is down or if 802.1W is disabled on the port, that port is given the role of **Disabled port**. Disabled ports have no role in the topology. However, if 802.1W is enabled on a port with a link down and the link of that port comes up, then that port assumes one of the following port roles: Root, Designated, Alternate, or Backup.


The following example ([Figure 2](#)) explains role assignments in a simple RSTP topology.

NOTE

All examples in this document assume that all ports in the illustrated topologies are point-to-point links and are homogeneous (they have the same path cost value) unless otherwise specified.

The topology in Figure 2 contains four bridges. Switch 1 is the root bridge since it has the lowest bridge priority. Switch 2 through Switch 4 are non-root bridges.

FIGURE 2 Simple 802.1W topology



Ports on Switch 1

All ports on Switch 1, the root bridge, are assigned Designated port roles.

Ports on Switch 2

Port2 on Switch 2 directly connects to the root bridge; therefore, Port2 is the Root port.

The bridge priority value on Switch 2 is superior to that of Switch 3 and Switch 4; therefore, the ports on Switch 2 that connect to Switch 3 and Switch 4 are given the Designated port role.

Furthermore, Port7 and Port8 on Switch 2 are physically connected. The RST BPDUs transmitted by Port7 are superior to those Port8 transmits. Therefore, Port8 is the Backup port and Port7 is the Designated port.

Ports on Switch 3

Port2 on Switch 3 directly connects to the Designated port on the root bridge; therefore, it assumes the Root port role.

The root path cost of the RST BPDUs received on Port4/Switch 3 is inferior to the RST BPDUs transmitted by the port; therefore, Port4/Switch 3 becomes the Designated port.

Similarly Switch 3 has a bridge priority value inferior to Switch 2. Port3 on Switch 3 connects to Port 3 on Switch 2. This port will be given the Alternate port role, since a Root port is already established on this bridge.

Ports Switch 4

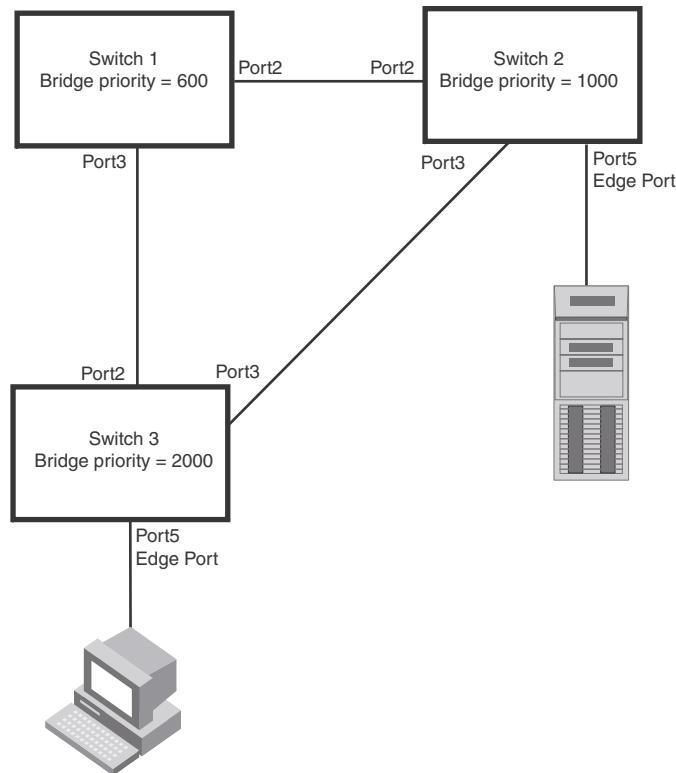
Switch 4 is not directly connected to the root bridge. It has two ports with superior incoming RST BPDUs from two separate LANs: Port3 and Port4. The RST BPDUs received on Port3 are superior to the RST BPDUs received on port 4; therefore, Port3 becomes the Root port and Port4 becomes the Alternate port.

Edge ports and edge port roles

The Dell implementation of 802.1W allows ports that are configured as Edge ports to be present in an 802.1W topology. (Figure 3). Edge ports are ports of a bridge that connect to workstations or computers. Edge ports do not register any incoming BPDU activities.

Edge ports assume Designated port roles. Port flapping does not cause any topology change events on Edge ports since 802.1W does not consider Edge ports in the spanning tree calculations.

FIGURE 3 Topology with edge ports



However, if any incoming RST BPDU is received from a previously configured Edge port, 802.1W automatically makes the port as a non-edge port. This is extremely important to ensure a loop free Layer 2 operation since a non-edge port is part of the active RSTP topology.

The 802.1W protocol can auto-detect an Edge port and a non-edge port. An administrator can also configure a port to be an Edge port using the CLI. It is recommended that Edge ports are configured explicitly to take advantage of the Edge port feature, instead of allowing the protocol to auto-detect them.

Point-to-point ports


To take advantage of the 802.1W features, ports on an 802.1W topology should be explicitly configured as point-to-point links using the CLI. Shared media should not be configured as point-to-point links.

NOTE

Configuring shared media or non-point-to-point links as point-to-point links could lead to Layer 2 loops.

The topology in Figure 4 is an example of shared media that should not be configured as point-to-point links. In Figure 4, a port on a bridge communicates or is connected to at least two ports.

FIGURE 4 Example of shared media



Bridge port states

Ports roles can have one of the following states:

- **Forwarding** – 802.1W is allowing the port to send and receive all packets.
- **Discarding** – 802.1W has blocked data traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is forwarding. When a port is in this state, the port does not transmit or receive data frames, but the port does continue to receive RST BPDUs. This state corresponds to the listening and blocking states of 802.1D.
- **Learning** – 802.1W is allowing MAC entries to be added to the filtering database but does not permit forwarding of data frames. The device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table.
- **Disabled** – The port is not participating in 802.1W. This can occur when the port is disconnected or 802.1W is administratively disabled on the port.

A port on a non-root bridge with the role of Root port is always in a forwarding state. If another port on that bridge assumes the Root port role, then the old Root port moves into a discarding state as it assumes another port role.

A port on a non-root bridge with a Designated role starts in the discarding state. When that port becomes elected to the Root port role, 802.1W quickly places it into a forwarding state. However, if the Designated port is an Edge port, then the port starts and stays in a forwarding state and it cannot be elected as a Root port.

A port with an Alternate or Backup role is always in a discarding state. If the port role changes to Designated, then the port changes into a forwarding state.

If a port on one bridge has a Designated role and that port is connected to a port on another bridge that has an Alternate or Backup role, the port with a Designated role cannot be given a Root port role until two instances of the forward delay timer expires on that port.

Edge port and non-edge port states

As soon as a port is configured as an Edge port using the CLI, it goes into a forwarding state instantly (in less than 100 msec).

When the link to a port comes up and 802.1W detects that the port is an Edge port, that port instantly goes into a forwarding state.

If 802.1W detects that port as a non-edge port, the port state is changed as determined by the result of processing the received RST BPDU. The port state change occurs within four seconds of link up or after two hello timer expires on the port.

Changes to port roles and states

To achieve convergence in a topology, a port role and state changes as it receives and transmits new RST BPDUs. Changes in a port role and state constitute a topology change. Besides the superiority and inferiority of the RST BPDU, bridge-wide and per-port state machines are used to determine a port role as well as a port state. Port state machines also determine when port role and state changes occur.

State machines

The bridge uses the Port Role Selection state machine to determine if port role changes are required on the bridge. This state machine performs a computation when one of the following events occur:

- New information is received on any port on the bridge
- The timer expires for the current information on a port on the bridge

Each port uses the following state machines:

- **Port Information** – This state machine keeps track of spanning-tree information currently used by the port. It records the origin of the information and ages out any information that was derived from an incoming BPDU.
- **Port Role Transition** – This state machine keeps track of the current port role and transitions the port to the appropriate role when required. It moves the Root port and the Designated port into forwarding states and moves the Alternate and Backup ports into discarding states.
- **Port Transmit** – This state machine is responsible for BPDU transmission. It checks to ensure only the maximum number of BPDUs per hello interval are sent every second. Based on what mode it is operating in, it sends out either legacy BPDUs or RST BPDUs. In this document legacy BPDUs are also referred to as STP BPDUs.

- **Port Protocol Migration** – This state machine deals with compatibility with 802.1D bridges. When a legacy BPDU is detected on a port, this state machine configures the port to transmit and receive legacy BPDUs and operate in the legacy mode.
- **Topology Change** – This state machine detects, generates, and propagates topology change notifications. It acknowledges Topology Change Notice (TCN) messages when operating in 802.1D mode. It also flushes the MAC table when a topology change event takes place.
- **Port State Transition** – This state machine transitions the port to a discarding, learning, or forwarding state and performs any necessary processing associated with the state changes.
- **Port Timers** – This state machine is responsible for triggering any of the state machines described above, based on expiration of specific port timers.

In contrast to the 802.1D standard, the 802.1W standard does not have any bridge specific timers. All timers in the CLI are applied on a per-port basis, even though they are configured under bridge parameters.

802.1W state machines attempt to quickly place the ports into either a forwarding or discarding state. Root ports are quickly placed in forwarding state when both of the following events occur:

- It is assigned to be the Root port.
- It receives an RST BPDU with a proposal flag from a Designated port. The proposal flag is sent by ports with a Designated role when they are ready to move into a forwarding state.

When a the role of Root port is given to another port, the old Root port is instructed to reroot. The old Root port goes into a discarding state and negotiates with its peer port for a new role and a new state. A peer port is the port on the other bridge to which the port is connected. For example, in [Figure 5](#), Port1 of Switch 200 is the peer port of Port2 of Switch 100.

A port with a Designated role is quickly placed into a forwarding state if one of the following occurs:

- The Designated port receives an RST BPDU that contains an agreement flag from a Root port
- The Designated port is an Edge port

However, a Designated port that is attached to an Alternate port or a Backup port must wait until the forward delay timer expires twice on that port while it is still in a Designated role, before it can proceed to the forwarding state.

Backup ports are quickly placed into discarding states.

Alternate ports are quickly placed into discarding states.

A port operating in 802.1W mode may enter a learning state to allow MAC entries to be added to the filtering database; however, this state is transient and lasts only a few milliseconds, if the port is operating in 802.1W mode and if the port meets the conditions for rapid transition.

Handshake mechanisms

To rapidly transition a Designated or Root port into a forwarding state, the Port Role Transition state machine uses handshake mechanisms to ensure loop free operations. It uses one type of handshake if no Root port has been assigned on a bridge, and another type if a Root port has already been assigned.

Handshake when no root port is elected

If a Root port has not been assigned on a bridge, 802.1W uses the *Proposing -> Proposed -> Sync -> Synced -> Agreed* handshake:

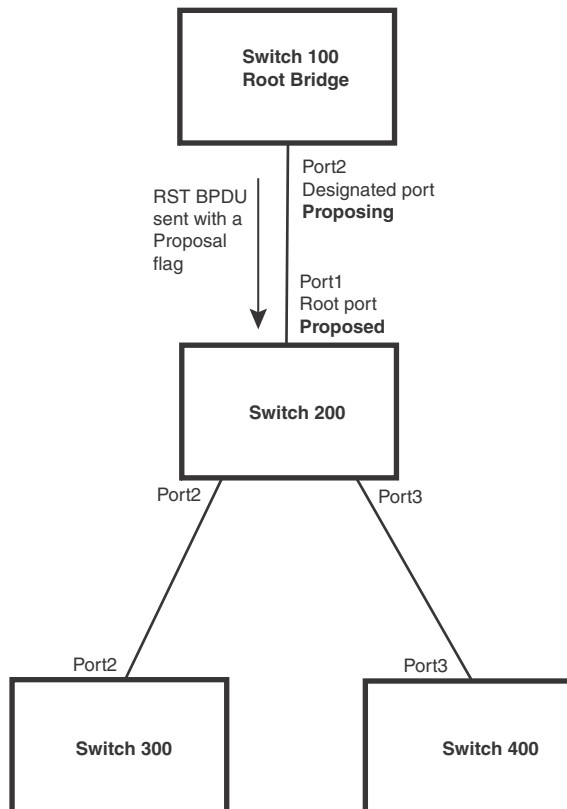
- **Proposing** – The Designated port on the root bridge sends an RST BPDU packet to its peer port that contains a proposal flag. The proposal flag is a signal that indicates that the Designated port is ready to put itself in a forwarding state (Figure 5). The Designated port continues to send this flag in its RST BPDU until it is placed in a forwarding state (Figure 8) or is forced to operate in 802.1D mode. (Refer to “Compatibility of 802.1W with 802.1D” on page 134).
- **Proposed** – When a port receives an RST BPDU with a proposal flag from the Designated port on its point-to-point link, it asserts the Proposed signal and one of the following occurs (Figure 5):
 - If the RST BPDU that the port receives is superior to what it can transmit, the port assumes the role of a Root port. (Refer to the section on “Bridges and bridge port roles” on page 107.)
 - If the RST BPDU that the port receives is inferior to what it can transmit, then the port is given the role of Designated port.

NOTE

Proposed will never be asserted if the port is connected on a shared media link.


In Figure 5, Port3/Switch 200 is elected as the Root port

FIGURE 5 Proposing and proposed stage



- **Sync** – Once the Root port is elected, it sets a sync signal on all the ports on the bridge. The signal tells the ports to synchronize their roles and states (Figure 6). Ports that are non-edge ports with a role of Designated port change into a discarding state. These ports have to negotiate with their peer ports to establish their new roles and states.

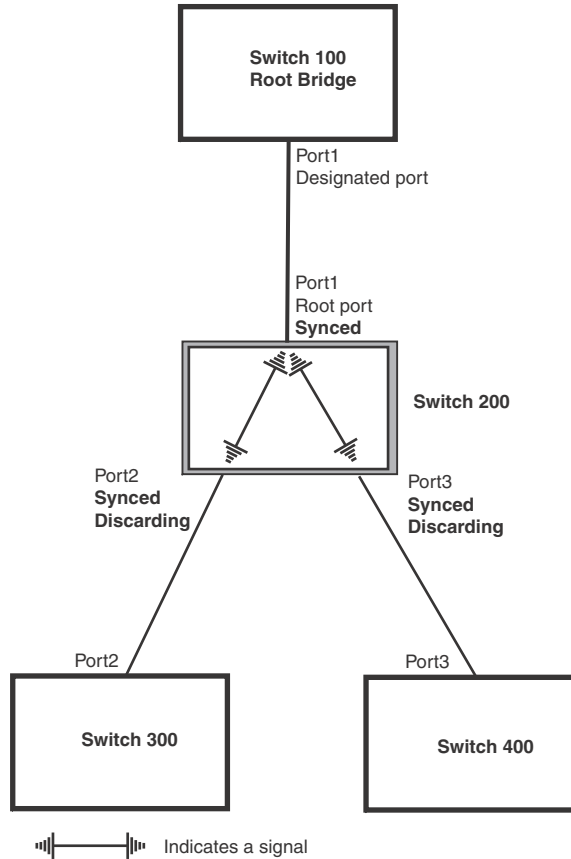
FIGURE 6 Sync stage



6 Configuring STP related features

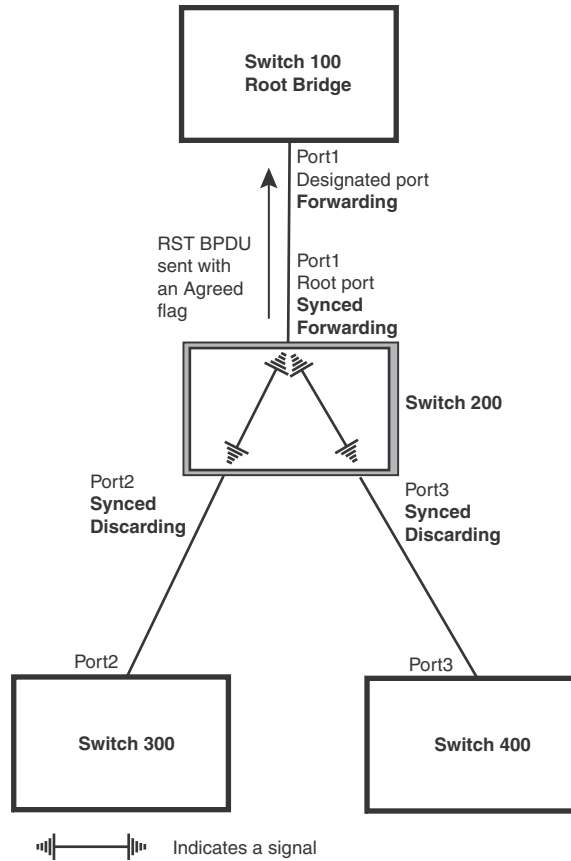
- **Synced** – Once the Designated port changes into a discarding state, it asserts a synced signal. Immediately, Alternate ports and Backup ports are synced. The Root port monitors the synced signals from all the bridge ports. Once all bridge ports asserts a synced signal, the Root port asserts its own synced signal (Figure 7).

FIGURE 7 Synced stage



- **Agreed** – The Root port sends back an RST BPDU containing an agreed flag to its peer Designated port and moves into the forwarding state. When the peer Designated port receives the RST BPDU, it rapidly transitions into a forwarding state.

FIGURE 8 Agree stage



At this point, the handshake mechanism is complete between Switch 100, the root bridge, and Switch 200.

Switch 200 updates the information on the Switch 200 Designated ports (Port2 and Port3) and identifies the new root bridge. The Designated ports send RST BPDUs, containing proposal flags, to their downstream bridges, without waiting for the hello timers to expire on them. This process starts the handshake with the downstream bridges.

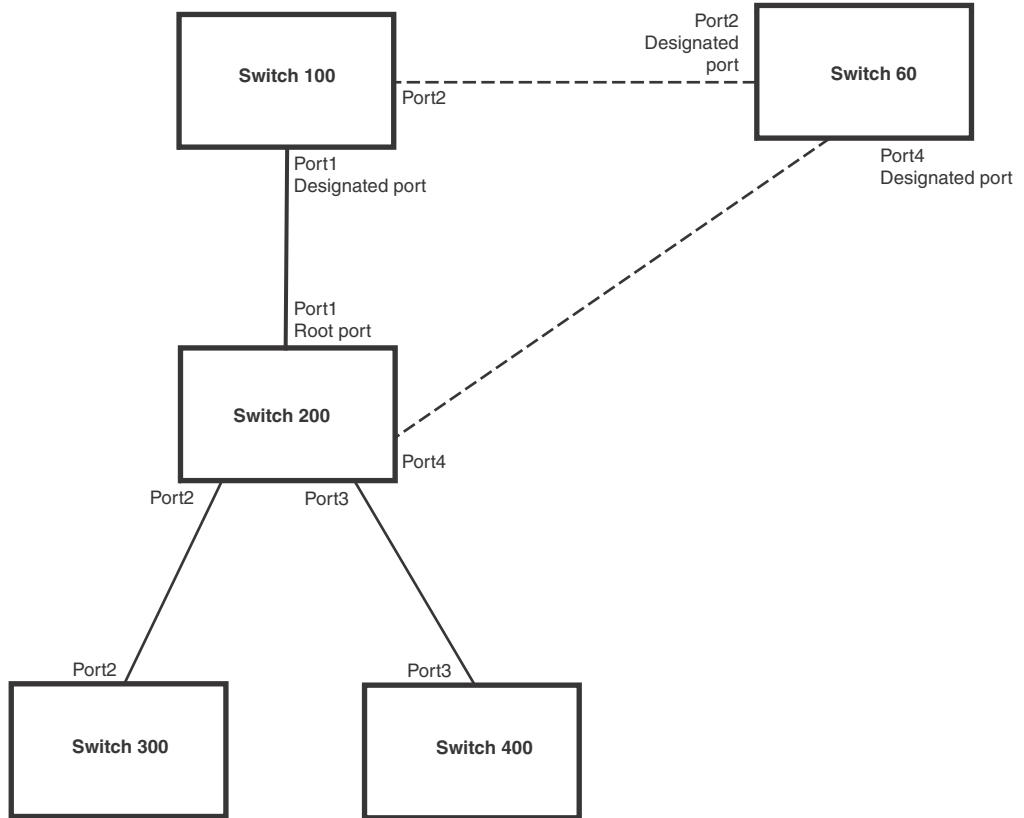
For example, Port2/Switch 200 sends an RST BPDU to Port2/Switch 300 that contains a proposal flag. Port2/Switch 300 asserts a proposed signal. Ports in Switch 300 then set sync signals on the ports to synchronize and negotiate their roles and states. Then the ports assert a synced signal and when the Root port in Switch 300 asserts its synced signal, it sends an RST BPDU to Switch 200 with an agreed flag.

This handshake is repeated between Switch 200 and Switch 400 until all Designated and Root ports are in forwarding states.

Handshake when a root port has been elected

If a non-root bridge already has a Root port, 802.1W uses a different type of handshake. For example, in [Figure 9](#), a new root bridge is added to the topology.

FIGURE 9 Addition of a new root bridge

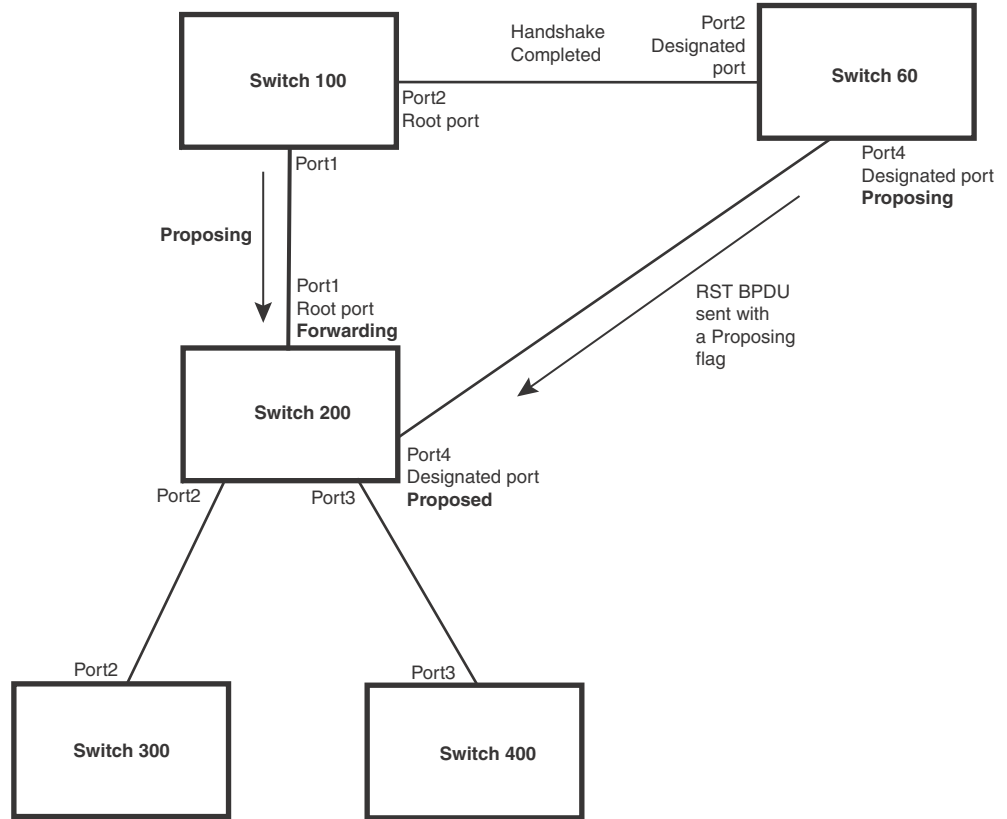


The handshake that occurs between Switch 60 and Switch 100 follows the one described in the previous section (“[Handshake when no root port is elected](#)” on page 113). The former root bridge becomes a non-root bridge and establishes a Root port ([Figure 10](#)).

However, since Switch 200 already had a Root port in a forwarding state, 802.1W uses the *Proposing -> Proposed -> Sync and Reroot -> Sync and Rerooted -> Rerooted and Synced -> Agreed handshake*:

- Proposing and Proposed** – The Designated port on the new root bridge (Port4/Switch 60) sends an RST BPDU that contains a proposing signal to Port4/Switch 200 to inform the port that it is ready to put itself in a forwarding state (Figure 10). 802.1W algorithm determines that the RST BPDU that Port4/Switch 200 received is superior to what it can generate, so Port4/Switch 200 assumes a Root port role.

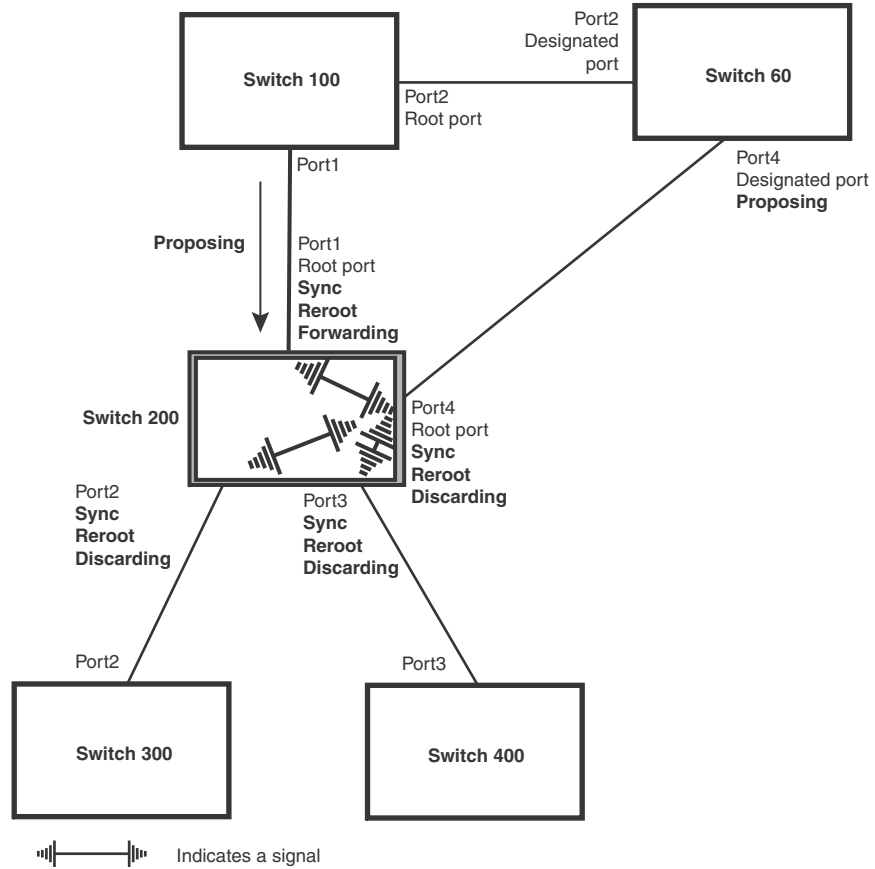
FIGURE 10 New root bridge sending a proposal flag



6 Configuring STP related features

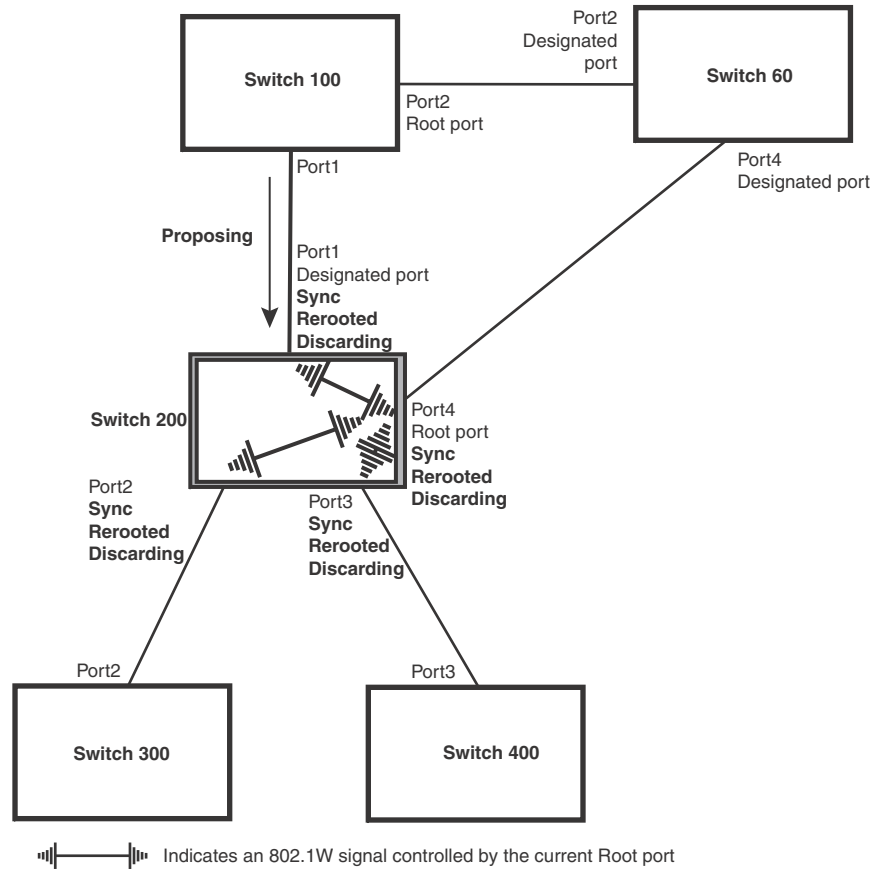
- **Sync and Reroot** – The Root port then asserts a sync and a reroot signal on all the ports on the bridge. The signal tells the ports that a new Root port has been assigned and they are to renegotiate their new roles and states. The other ports on the bridge assert their sync and reroot signals. Information about the old Root port is discarded from all ports. Designated ports change into discarding states (Figure 11).

FIGURE 11 Sync and reroot



- **Sync and Rerooted** – When the ports on Switch 200 have completed the reroot phase, they assert their rerooted signals and continue to assert their sync signals as they continue in their discarding states. They also continue to negotiate their roles and states with their peer ports (Figure 12).

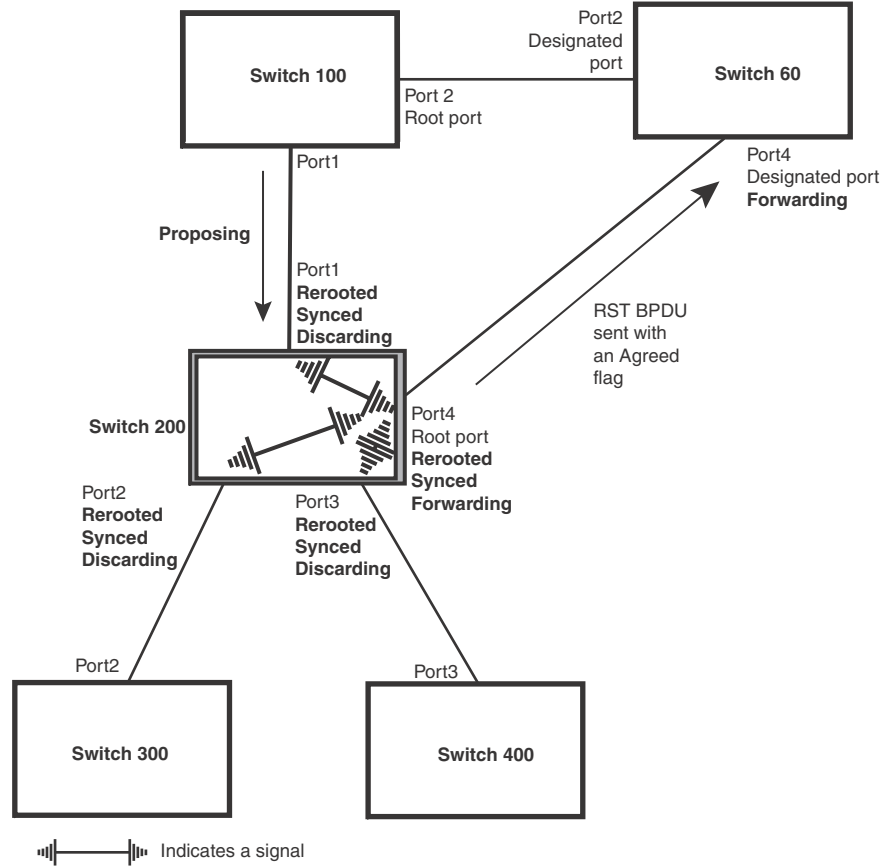
FIGURE 12 Sync and rerooted



6 Configuring STP related features

- **Synced and Agree** – When all the ports on the bridge assert their synced signals, the new Root port asserts its own synced signal and sends an RST BPDU to Port4/Switch 60 that contains an agreed flag (Figure 12). The Root port also moves into a forwarding state.

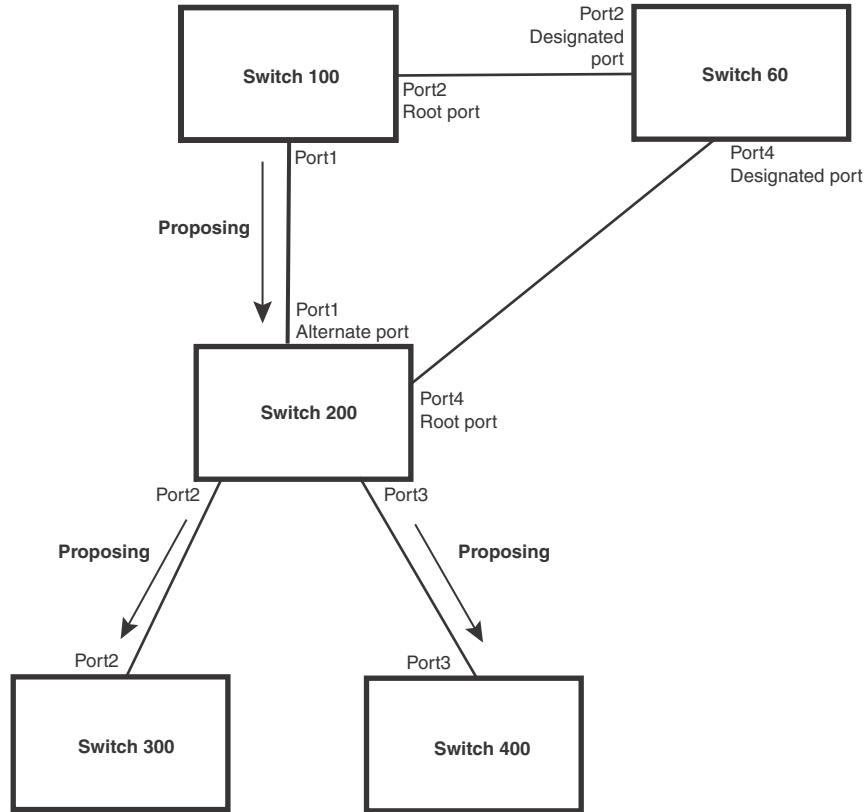
FIGURE 13 Rerouted, synced, and agreed



The old Root port on Switch 200 becomes an Alternate Port (Figure 14). Other ports on that bridge are elected to appropriate roles.

The Designated port on Switch 60 goes into a forwarding state once it receives the RST BPDU with the agreed flag.

FIGURE 14 Handshake completed after election of new root port



Recall that Switch 200 sent the agreed flag to Port4/Switch 60 and not to Port1/Switch 100 (the port that connects Switch 100 to Switch 200). Therefore, Port1/Switch 100 does not go into forwarding state instantly. It waits until two instances of the forward delay timer expires on the port before it goes into forwarding state.

At this point the handshake between the Switch 60 and Switch 200 is complete.

The remaining bridges (Switch 300 and Switch 400) may have to go through the reroot handshake if a new Root port needs to be assigned.

Convergence in a simple topology

The examples in this section illustrate how 802.1W convergence occurs in a simple Layer 2 topology at start-up.


NOTE

The remaining examples assume that the appropriate handshake mechanisms occur as port roles and states change.

Convergence at start up

In [Figure 15](#), two bridges Switch 2 and Switch 3 are powered up. There are point-to-point connections between Port3/Switch 2 and Port3/Switch 3.

FIGURE 15 Convergence between two bridges



At power up, all ports on Switch 2 and Switch 3 assume Designated port roles and are at discarding states before they receive any RST BPDU.

Port3/Switch 2, with a Designated role, transmits an RST BPDU with a proposal flag to Port3/Switch 3. A ports with a Designated role sends the proposal flag in its RST BPDU when they are ready to move to a forwarding state.

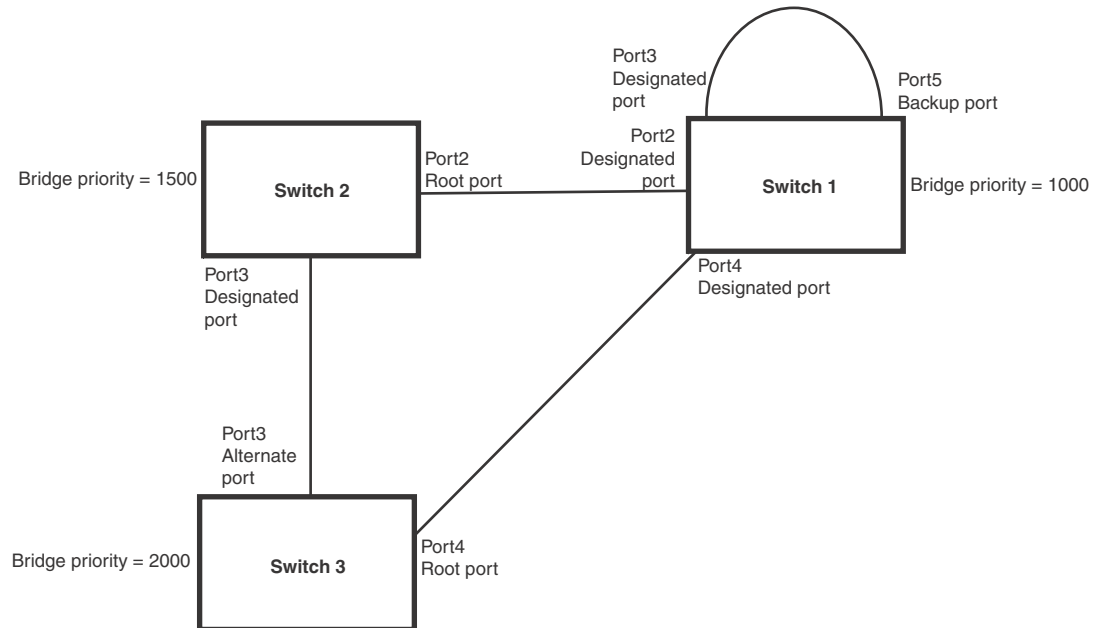
Port3/Switch 3, which starts with a role of Designated port, receives the RST BPDU and finds that it is superior to what it can transmit; therefore, Port3/Switch 3 assumes a new port role, that of a Root port. Port3/Switch 3 transmits an RST BPDU with an agreed flag back to Switch 2 and immediately goes into a forwarding state.

Port3/Switch 2 receives the RST BPDU from Port3/Switch 3 and immediately goes into a forwarding state.

Now 802.1W has fully converged between the two bridges, with Port3/Switch 3 as an operational root port in forwarding state and Port3/Switch 2 as an operational Designated port in forwarding state.

Next, Switch 1 is powered up (Figure 16).

FIGURE 16 Simple Layer 2 topology



The point-to-point connections between the three bridges are as follows:

- Port2/Switch 1 and Port2/Switch 2
- Port4/Switch 1 and Port4/Switch 3
- Port3/Switch 2 and Port3/Switch 3

Ports 3 and 5 on Switch 1 are physically connected together.

At start up, the ports on Switch 1 assume Designated port roles, which are in discarding state. They begin sending RST BPDUs with proposal flags to move into a forwarding state.

When Port4/Switch 3 receives these RST BPDUs 802.1W algorithm determines that they are better than the RST BPDUs that were previously received on Port3/Switch 3. Port4/Switch 3 is now selected as Root port. This new assignment signals Port3/Switch 3 to begin entering the discarding state and to assume an Alternate port role. As it goes through the transition, Port3/Switch 3 negotiates a new role and state with its peer port, Port3/Switch 2.

Port4/Switch 3 sends an RST BPDUs with an agreed flag to Port4/Switch 1. Both ports go into forwarding states.

Port2/Switch 2 receives an RST BPDUs. The 802.1W algorithm determines that these RST BPDUs that are superior to any that any port on Switch 2 can transmit; therefore, Port2/Switch 2 assumes the role of a Root port.

The new Root port then signals all ports on the bridge to start synchronization. Since none of the ports are Edge ports, they all enter the discarding state and assume the role of Designated ports. Port3/Switch 2, which previously had a Designated role with a forwarding state, starts the discarding state. They also negotiate port roles and states with their peer ports. Port3/Switch 2 also sends an RST BPDUs to Port3/Switch 3 with a proposal flag to request permission go into a forwarding state.

6 Configuring STP related features

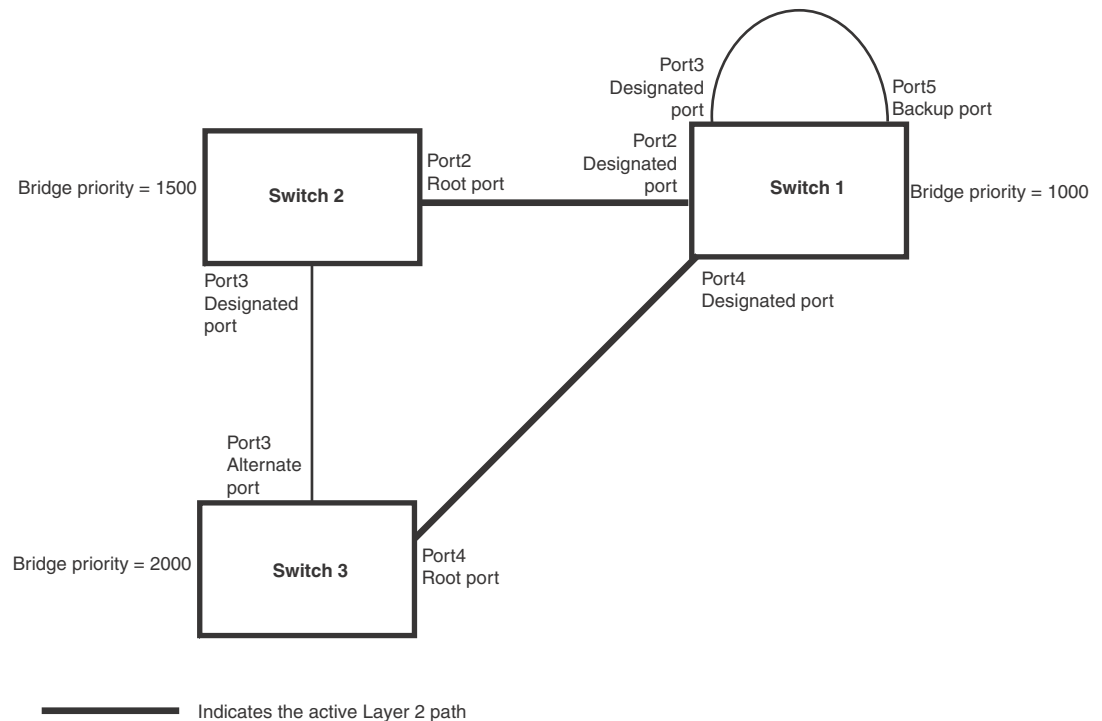
The Port2/Switch 2 bridge also sends an RST BPDUs with an agreed flag Port2/Switch 1 that Port2 is the new Root port. Both ports go into forwarding states.

Now, Port3/Switch 3 is currently in a discarding state and is negotiating a port role. It received RST BPDUs from Port3/Switch 2. The 802.1W algorithm determines that the RST BPDUs Port3/Switch 3 received are superior to those it can transmit; however, they are not superior to those that are currently being received by the current Root port (Port4). Therefore, Port3 retains the role of Alternate port.

Ports 3/Switch 1 and Port5/Switch 1 are physically connected. Port5/Switch 1 received RST BPDUs that are superior to those received on Port3/Switch 1; therefore, Port5/Switch 1 is given the Backup port role while Port3 is given the Designated port role. Port3/Switch 1, does not go directly into a forwarding state. It waits until the forward delay time expires twice on that port before it can proceed to the forwarding state.

Once convergence is achieved, the active Layer 2 forwarding path converges as shown in [Figure 17](#).

FIGURE 17 Active Layer 2 path




Convergence after a link failure

What happens if a link in the 802.1W topology fails?

For example, Port2/Switch, which is the port that connects Switch 2 to the root bridge (Switch 1), fails. Both Switch 2 and Switch 1 notice the topology change (Figure 18).

FIGURE 18 Link failure in the topology



Switch 1 sets its Port2 into a discarding state.

At the same time, Switch 2 assumes the role of a root bridge since its root port failed and it has no operational Alternate port. Port3/Switch 2, which currently has a Designated port role, sends an RST BPDU to Switch 3. The RST BPDU contains a proposal flag and a bridge ID of Switch 2 as its root bridge ID.

When Port3/Switch 3 receives the RST BPDUs, 802.1W algorithm determines that they are inferior to those that the port can transmit. Therefore, Port3/Switch 3 is given a new role, that of a Designated port. Port3/Switch 3 then sends an RST BPDU with a proposal flag to Switch 2, along with the new role information. However, the root bridge ID transmitted in the RST BPDU is still Switch 1.

When Port3/Switch 2 receives the RST BPDU, 802.1W algorithm determines that it is superior to the RST BPDU that it can transmit; therefore, Port3/Switch 2 receives a new role; that of a Root port. Port3/Switch 2 then sends an RST BPDU with an agreed flag to Port3/Switch 3. Port3/Switch 2 goes into a forwarding state.

When Port3/Switch 3 receives the RST BPDU that Port3/Switch 2 sent, Port3/Switch 3 changes into a forwarding state, which then completes the full convergence of the topology.

Convergence at link restoration

When Port2/Switch 2 is restored, both Switch 2 and Switch 1 recognize the change. Port2/Switch 1 starts assuming the role of a Designated port and sends an RST BPDU containing a proposal flag to Port2/Switch 2.

6 Configuring STP related features

When Port2/Switch 2 receives the RST BPDUs, 802.1W algorithm determines that the RST BPDUs the port received are better than those received on Port3/Switch 3; therefore, Port2/Switch 2 is given the role of a Root port. All the ports on Switch 2 are informed that a new Root port has been assigned which then signals all the ports to synchronize their roles and states. Port3/Switch 2, which was the previous Root port, enters a discarding state and negotiates with other ports on the bridge to establish its new role and state, until it finally assumes the role of a Designated port.

Next, the following happens:

- Port3/Switch 2, the Designated port, sends an RST BPDUs, with a proposal flag to Port3/Switch 3.
- Port2/Switch 2 also sends an RST BPDUs with an agreed flag to Port2/Switch 1 and then places itself into a forwarding state.

When Port2/Switch 1 receives the RST BPDUs with an agreed flag sent by Port2/Switch 2, it puts that port into a forwarding state. The topology is now fully converged.

When Port3/Switch 3 receives the RST BPDUs that Port3/Switch 2 sent, 802.1W algorithm determines that these RST BPDUs are superior to those that Port3/Switch 3 can transmit. Therefore, Port3/Switch 3 is given a new role, that of an Alternate port. Port3/Switch 3 immediately enters a discarding state.

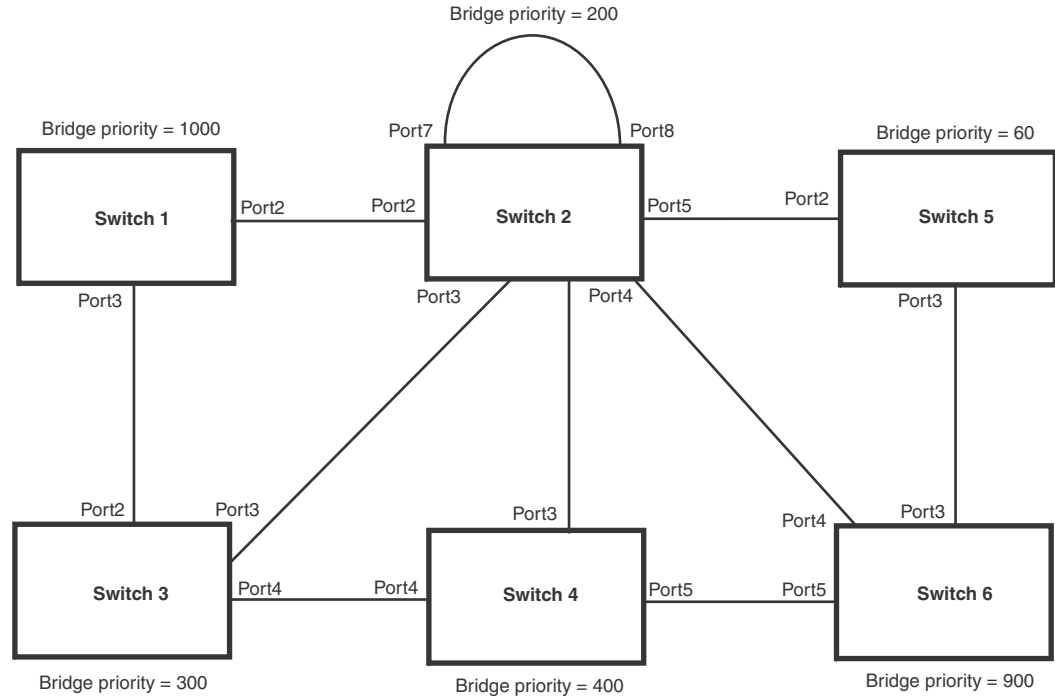
Now Port3/Switch 2 does not go into a forwarding state instantly like the Root port. It waits until the forward delay timer expires twice on that port while it is still in a Designated role, before it can proceed to the forwarding state. The wait, however, does not cause a denial of service, since the essential connectivity in the topology has already been established.

When fully restored, the topology is the same as that shown on [Figure 16](#).

Convergence in a complex 802.1W topology

The following is an example of a complex 802.1W topology.

FIGURE 19 Complex 802.1W topology



In [Figure 19](#), Switch 5 is selected as the root bridge since it is the bridge with the highest priority. Lines in the figure show the point-to-point connection to the bridges in the topology.

Switch 5 sends an RST BPDU that contains a proposal flag to Port5/Switch 2. When handshakes are completed in Switch 5, Port5/Switch 2 is selected as the Root port on Switch 2. All other ports on Switch 2 are given Designated port role with discarding states.

Port5/Switch 2 then sends an RST BPDU with an agreed flag to Switch 5 to confirm that it is the new Root port and the port enters a forwarding state. Port7 and Port8 are informed of the identity of the new Root port. 802.1W algorithm selects Port7 as the Designated port while Port8 becomes the Backup port.

Port3/Switch 5 sends an RST BPDU to Port3/Switch 6 with a proposal flag. When Port3/Switch 5 receives the RST BPDU, handshake mechanisms select Port3 as the Root port of Switch 6. All other ports are given a Designated port role with discarding states. Port3/Switch 6 then sends an RST BPDU with an agreed flag to Port3/Switch 5 to confirm that it is the Root port. The Root port then goes into a forwarding state.

Now, Port4/Switch 6 receives RST BPDUs that are superior to what it can transmit; therefore, it is given the Alternate port role. The port remains in discarding state.

Port5/Switch 6 receives RST BPDUs that are inferior to what it can transmit. The port is then given a Designated port role.

6 Configuring STP related features

Next Switch 2 sends RST BPDUs with a proposal flag to Port3/Switch 4. Port3 becomes the Root port for the bridge; all other ports are given a Designated port role with discarding states. Port3/Switch 4 sends an RST BPDU with an agreed flag to Switch 2 to confirm that it is the new Root port. The port then goes into a forwarding state.

Now Port4/Switch 4 receives an RST BPDU that is superior to what it can transmit. The port is then given an Alternate port role, and remains in discarding state.

Likewise, Port5/Switch 4 receives an RST BPDU that is superior to what it can transmit. The port is also given an Alternate port role, and remains in discarding state.

Port2/Switch 2 transmits an RST BPDU with a proposal flag to Port2/Switch 1. Port2/Switch 1 becomes the Root port. All other ports on Switch 1 are given Designated port roles with discarding states.

Port2/Switch 1 sends an RST BPDU with an agreed flag to Port2/Switch 2 and Port2/Switch 1 goes into a forwarding state.

Port3/Switch 1 receives an RST BPDUs that is inferior to what it can transmit; therefore, the port retains its Designated port role and goes into forwarding state only after the forward delay timer expires twice on that port while it is still in a Designated role.

Port3/Switch 2 sends an RST BPDU to Port3/Switch 3 that contains a proposal flag. Port3/Switch 3 becomes the Root port, while all other ports on Switch 3 are given Designated port roles and go into discarding states. Port3/Switch 3 sends an RST BPDU with an agreed flag to Port3/Switch 2 and Port3/Switch 3 goes into a forwarding state.

Now, Port2/Switch 3 receives an RST BPDUs that is superior to what it can transmit so that port is given an Alternate port state.

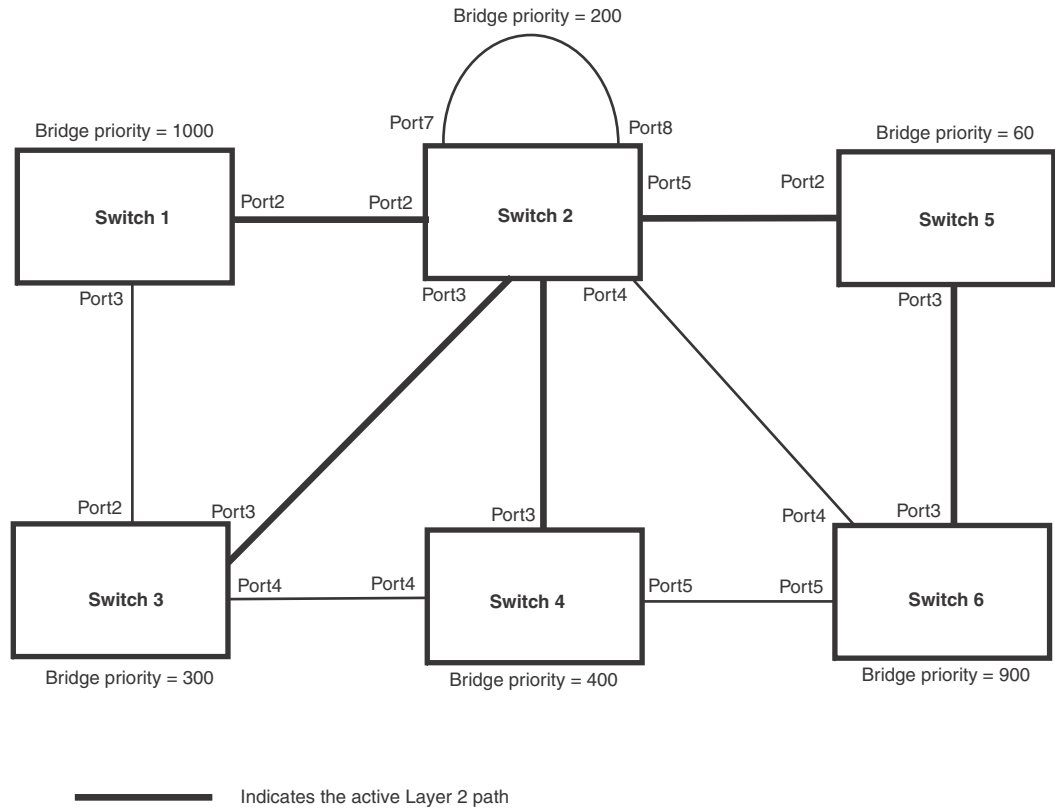
Port4/Switch 3 receives an RST BPDU that is inferior to what it can transmit; therefore, the port retains its Designated port role.

Ports on all the bridges in the topology with Designated port roles that received RST BPDUs with agreed flags go into forwarding states instantly. However, Designated ports that did not receive RST BPDUs with agreed flags must wait until the forward delay timer expires twice on those port. Only then will these port move into forwarding states.

The entire 802.1W topology converges in less than 300 msec and the essential connectivity is established between the designated ports and their connected root ports.

After convergence is complete, [Figure 20](#) shows the active Layer 2 path of the topology in [Figure 19](#).

FIGURE 20 Active Layer 2 path in complex topology



Propagation of topology change

The Topology Change state machine generates and propagates the topology change notification messages on each port. When a Root port or a Designated port goes into a forwarding state, the Topology Change state machine on those ports send a topology change notice (TCN) to all the bridges in the topology to propagate the topology change.

NOTE

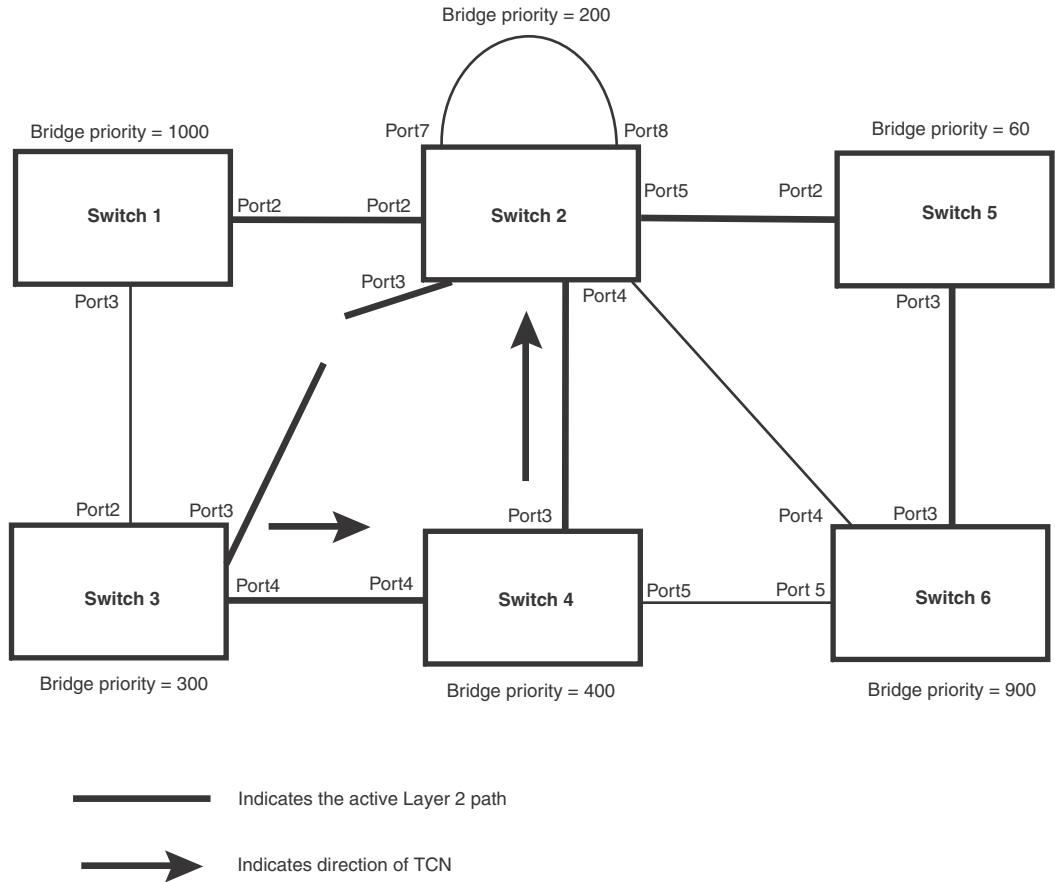
Edge ports, Alternate ports, or Backup ports do not need to propagate a topology change.

The TCN is sent in the RST BPDU that a port sends. Ports on other bridges in the topology then acknowledge the topology change once they receive the RST BPDU, and send the TCN to other bridges until all the bridges are informed of the topology change.

6 Configuring STP related features

For example, Port3/Switch 2 in [Figure 21](#), fails. Port4/Switch 3 becomes the new Root port. Port4/Switch 3 sends an RST BPDU with a TCN to Port4/Switch 4. To propagate the topology change, Port4/Switch 4 then starts a TCN timer on itself, on the bridge Root port, and on other ports on that bridge with a Designated role. Then Port3/Switch 4 sends RST BPDU with the TCN to Port4/Switch 2. (Note the new active Layer 2 path in [Figure 21](#).)

FIGURE 21 Beginning of topology change notice

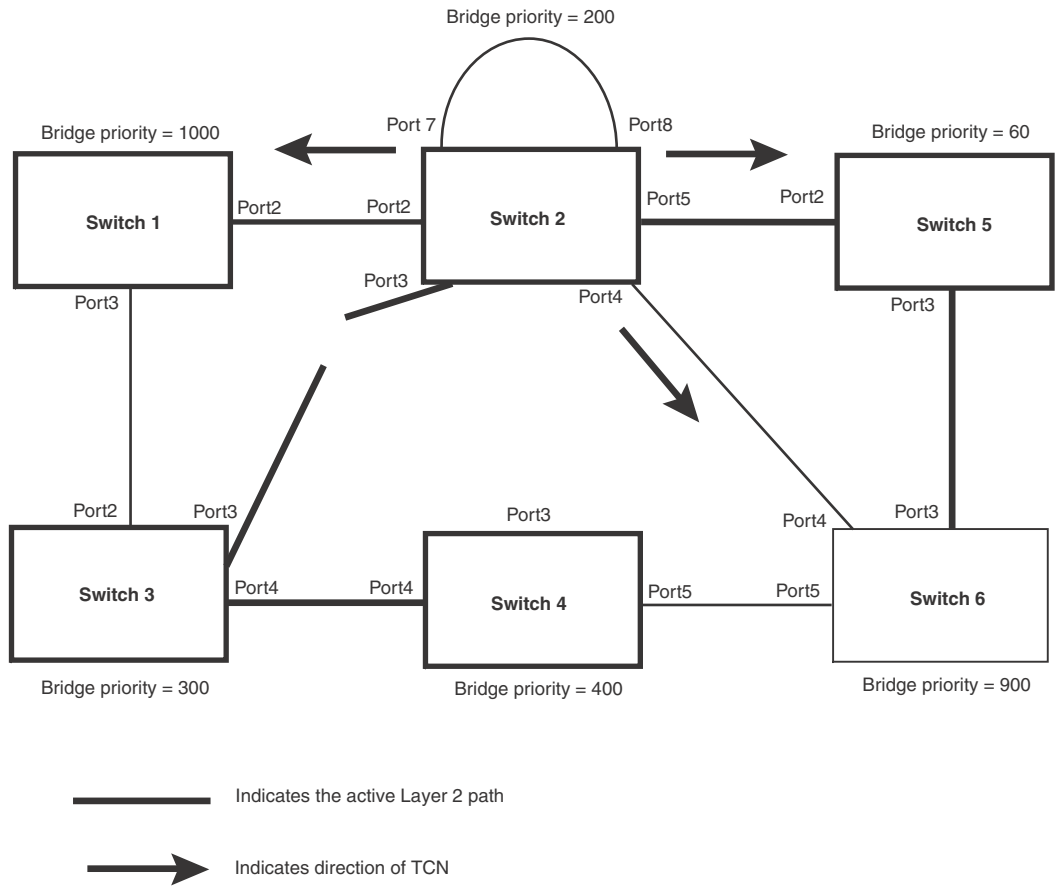


Switch 2 then starts the TCN timer on the Designated ports and sends RST BPDUs that contain the TCN as follows ([Figure 22](#)):

- Port5/Switch 2 sends the TCN to Port2/Switch 5
- Port4/Switch 2 sends the TCN to Port4/Switch 6

- Port2/Switch 2 sends the TCN to Port2/Switch 1

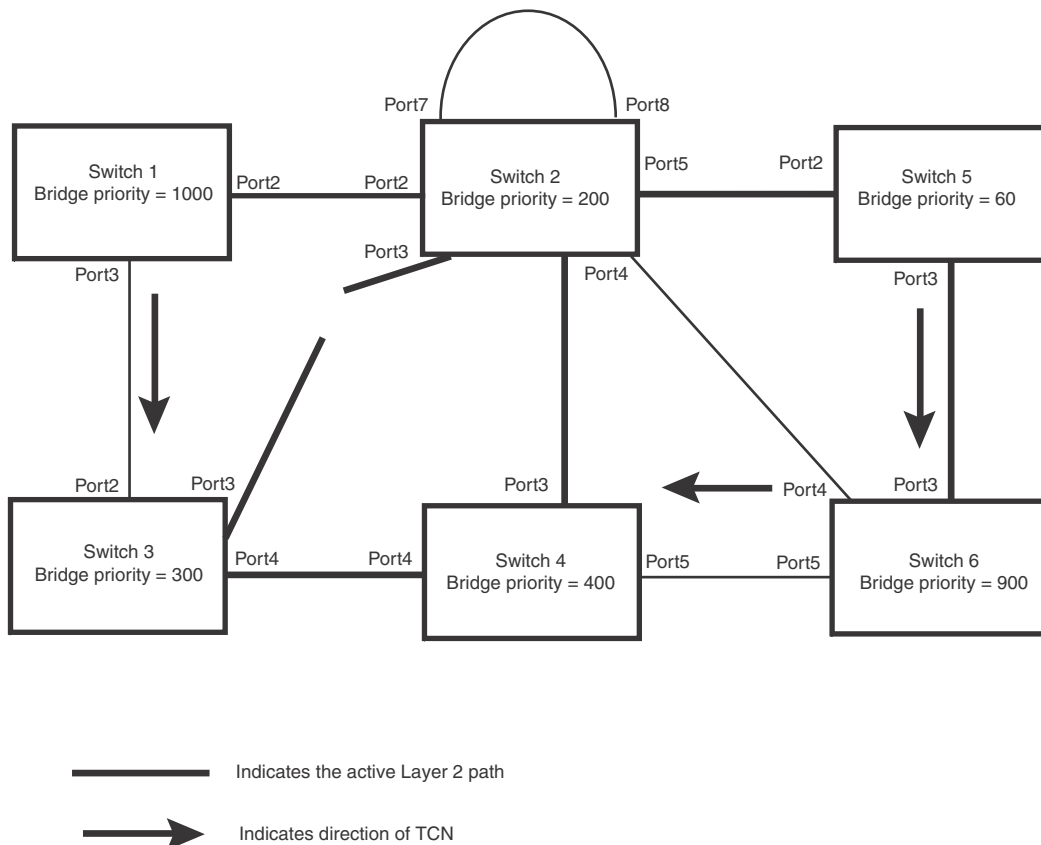
FIGURE 22 Sending TCN to bridges connected to Switch 2



6 Configuring STP related features

Then Switch 1, Switch 5, and Switch 6 send RST BPDUs that contain the TCN to Switch 3 and Switch 4 to complete the TCN propagation (Figure 23).

FIGURE 23 Completing the TCN propagation



Compatibility of 802.1W with 802.1D

802.1W-enabled bridges are backward compatible with IEEE 802.1D bridges. This compatibility is managed on a per-port basis by the Port Migration state machine. **However, intermixing the two types of bridges in the network topology is not advisable if you want to take advantage of the rapid convergence feature.**


Compatibility with 802.1D means that an 802.1W-enabled port can send BPDUs in the STP or 802.1D format when one of the following events occur:

- The port receives a legacy BPDU. A legacy BPDU is an STP BPDU or a BPDU in an 802.1D format. The port that receives the legacy BPDU automatically configures itself to behave like a legacy port. It sends and receives legacy BPDUs only.
- The entire bridge is configured to operate in an 802.1D mode when an administrator sets the bridge parameter to zero at the CLI, forcing all ports on the bridge to send legacy BPDUs only.

Once a port operates in the 802.1D mode, 802.1D convergence times are used and rapid convergence is not realized.

For example, in [Figure 24](#), Switch 10 and Switch 30 receive legacy BPDUs from Switch 20. Ports on Switch 10 and Switch 30 begin sending BPDUs in STP format to allow them to operate transparently with Switch 20.

FIGURE 24 802.1W bridges with an 802.1D bridge



Once Switch 20 is removed from the LAN, Switch 10 and Switch 30 receive and transmit BPDUs in the STP format to and from each other. This state will continue until the administrator enables the **force-migration-check** command to force the bridge to send RSTP BPDU during a migrate time period. If ports on the bridges continue to hear only STP BPDUs after this migrate time period, those ports will return to sending STP BPDUs. However, when the ports receive RST BPDUs during the migrate time period, the ports begin sending RST BPDUs. The migrate time period is non-configurable. It has a value of three seconds.

NOTE

The IEEE standards state that 802.1W bridges need to interoperate with 802.1D bridges. IEEE standards set the path cost of 802.1W bridges to be between 1 and 200,000,000; whereas path cost of 802.1D bridges are set between 1 and 65,535. In order for the two bridge types to be able to interoperate in the same topology, the administrator needs to configure the bridge path cost appropriately. Path costs for either 802.1W bridges or 802.1D bridges need to be changed; in most cases, path costs for 802.1W bridges need to be changed.

Configuring 802.1W parameters on a device

The remaining 802.1W sections explain how to configure the 802.1W protocol in a device.

NOTE

With RSTP running, enabling static trunk on ports that are members of VLAN 4000 will keep the system busy for 20 to 25 seconds.

PowerConnect devices are shipped from the factory with 802.1W disabled. Use the following methods to enable or disable 802.1W. You can enable or disable 802.1W at the following levels:

- **Port-based VLAN** – Affects all ports within the specified port-based VLAN. When you enable or disable 802.1W within a port-based VLAN, the setting overrides the global setting. Thus, you can enable 802.1W for the ports within a port-based VLAN even when 802.1W is globally disabled, or disable the ports within a port-based VLAN when 802.1W is globally enabled.
- **Individual port** – Affects only the individual port. However, if you change the 802.1W state of the primary port in a trunk group, the change affects all ports in the trunk group.

Enabling or disabling 802.1W in a port-based VLAN

Use the following procedure to disable or enable 802.1W on a device on which you have configured a port-based VLAN. Changing the 802.1W state in a VLAN affects only that VLAN.

6 Configuring STP related features

To enable 802.1W for all ports in a port-based VLAN, enter commands such as the following.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#spanning-tree 802-1w
```

Syntax: [no] spanning-tree 802-1w

Note regarding pasting 802.1W settings into the running configuration

If you paste 802.1W settings into the running configuration, and the pasted configuration includes ports that are already up, the ports will initially operate in STP legacy mode before operating in 802.1W RSTP mode. For example, the following pasted configuration will cause ports e 1 and e 2 to temporarily operate in STP legacy mode, because these ports are already up and running.

```
conf t
vlan 120
tag e 1 to e 2
spanning-tree 802-1w
spanning-tree 802-1w priority 1001
end
```

To avoid this issue, 802.1W commands/settings that are pasted into the configuration should be in the following order.

1. Ports that are not yet connected
2. 802.1W RSTP settings
3. Ports that are already up

Example

```
conf t
vlan 120
untag e 3
spanning-tree 802-1w
spanning-tree 802-1w priority 1001
tag e 1 to 2
end
```

In the above configuration, untagged port e3 is added to VLAN 120 *before* the 802.1W RSTP settings, and ports e1 and e2 are added *after* the 802.1W RSTP settings. When these commands are pasted into the running configuration, the ports will properly operate in 802.1W RSTP mode.

Enabling or disabling 802.1W on a single spanning tree

To enable 802.1W for all ports of a single spanning tree, enter a command such as the following.

```
PowerConnect(config-vlan-10)#spanning-tree single 802-1w
```

Syntax: [no] spanning-tree single 802-1w

Disabling or enabling 802.1W on an individual port

The **spanning-tree 802-1w** or **spanning-tree single 802-1w** command must be used to initially enable 802.1W on ports. Both commands enable 802.1W on all ports that belong to the VLAN or to the single spanning tree.

Once 802.1W is enabled on a port, it can be disabled on individual ports. 802.1W that have been disabled on individual ports can then be enabled as required.

NOTE

If you change the 802.1W state of the primary port in a trunk group, the change affects all ports in that trunk group.

To disable or enable 802.1W on an individual port, enter commands such as the following.

```
PowerConnect(config)#interface e 1
PowerConnect(config-if-e10000-1)#no spanning-tree
```

Syntax: [no] spanning-tree

Changing 802.1W bridge parameters

When you make changes to 802.1W bridge parameters, the changes are applied to individual ports on the bridge. To change 802.1W bridge parameters, use the following methods.

To designate a priority for a bridge, enter a command such as the following.

```
PowerConnect(config)#spanning-tree 802-1w priority 10
```

The command in this example changes the priority on a device on which you have not configured port-based VLANs. The change applies to the default VLAN. If you have configured a port-based VLAN on the device, you can configure the parameters only at the configuration level for individual VLANs. Enter commands such as the following.

```
PowerConnect(config)#vlan 20
PowerConnect(config-vlan-20)#spanning-tree 802-1w priority 0
```

To make this change in the default VLAN, enter the following commands.

```
PowerConnect(config)#vlan 1
PowerConnect(config-vlan-1)#spanning-tree 802-1w priority 0
```

Syntax: **spanning-tree 802-1w** [**forward-delay** <value>] | [**hello-time** <value>] | [**max-age** <time>] | [**force-version** <value>] | [**priority** <value>]

The **forward-delay** <value> parameter specifies how long a port waits before it forwards an RST BPDU after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. This parameter can have a value from 1 – 10 seconds. The default is 2 seconds.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value from 6 – 40 seconds. The default is 20 seconds.

The value of **max-age** must be greater than the value of **forward-delay** to ensure that the downstream bridges do not age out faster than the upstream bridges (those bridges that are closer to the root bridge).

The **force-version** <value> parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following values:

- **0** – The STP compatibility mode. Only STP (or legacy) BPDUs will be sent.
- **2** – The default. RST BPDUs will be sent unless a legacy bridge is detected. If a legacy bridge is detected, STP BPDUs will be sent instead.

The default is 2.

6 Configuring STP related features

The **priority** *<value>* parameter specifies the priority of the bridge. You can enter a value from 0 – 65535. A lower numerical value means the bridge has a higher priority. Thus, the highest priority is 0. The default is 32768.

You can specify some or all of these parameters on the same command line. If you specify more than one parameter, you must specify them in the order shown above, from left to right.

Changing port parameters

The 802.1W port commands can be enabled on individual ports or on multiple ports, such as all ports that belong to a VLAN.

The 802.1W port parameters are preconfigured with default values. If the default parameters meet your network requirements, no other action is required.

You can change the following 802.1W port parameters using the following method.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#spanning-tree 802-1w ethernet 5 path-cost 15
priority 64
```

Syntax: **spanning-tree 802-1w ethernet** *<portnum>* **path-cost** *<value>* | **priority** *<value>* | **[admin-edge-port]** | **[admin-pt2pt-mac]** | **[force-migration-check]**

The *<portnum>* parameter specifies the interface used.

The **path-cost** *<value>* parameter specifies the cost of the port path to the root bridge. 802.1W prefers the path with the lowest cost. You can specify a value from 1 – 20,000,000. [Table 26](#) shows the recommended path cost values from the IEEE standards.

TABLE 26 Recommended path cost values of 802.1W

Link speed	Recommended (Default) 802.1W path cost values	Recommended 802.1W patch cost range
Less than 100 kilobits per second	200,000,000	20,000,000 – 200,000,000
1 Megabit per second	20,000,000	2,000,000 – 200,000,000
10 Megabits per second	2,000,000	200,000 – 200,000,000
100 Megabits per second	200,000	20,000 – 200,000,000
1 Gbps per second	20,000	2,000 – 200,000,000
10 Gbps per second	2,000	200 – 20,000
100 Gbps per second	200	20 – 2,000
1 Terabits per second	20	2 – 200
10 Terabits per second	2	1 – 20

The **priority** *<value>* parameter specifies the preference that 802.1W gives to this port relative to other ports for forwarding traffic out of the topology. Y

- You can specify a value from 0 – 240, in increments of 16. If you enter a value that is not divisible by 16, the software returns an error message. The default value is 128. A higher numerical value means a lower priority; thus, the highest priority is 0.

Set the **admin-edge-port** to enabled or disabled. If set to enabled, then the port becomes an edge port in the domain.

Set the **admin-pt2pt-mac** to enabled or disabled. If set to enabled, then a port is connected to another port through a point-to-point link. The point-to-point link increases the speed of convergence. This parameter, however, does not auto-detect whether or not the link is a physical point-to-point link.

The **force-migration-check** parameter forces the specified port to sent one RST BPDU. If only STP BPDUs are received in response to the sent RST BPDU, then the port will go return to sending STP BPDUs.

Example

Suppose you want to enable 802.1W on a system with no active port-based VLANs and change the hello-time from the default value of 2 to 8 seconds. Additionally, suppose you want to change the path and priority costs for port 5 only. To do so, enter the following commands.

```
PowerConnect(config)#spanning-tree 802-1w hello-time 8
PowerConnect(config)#spanning-tree 802-1w ethernet 5 path-cost 15 priority 64
```

Displaying information about 802-1W

To display a summary of 802-1W, use the following command.

```
PowerConnect#show 802-1w
--- VLAN 1 [ STP Instance owned by VLAN 1 ] -----
VLAN 1 BPDU cam_index is 2 and the IGC and DMA master Are(HEX) 0 1 2 3
Bridge IEEE 802.1W Parameters:
Bridge          Bridge  Bridge Bridge Force   tx
Identifier      MaxAge Hello  FwdDly Version Hold
hex             sec    sec   sec      cnt
800000e080541700 20     2     15     Default 3

RootBridge      RootPath  DesignatedBri-   Root  Max  Fwd  Hel
Identifier      Cost      dge Identifier   Port  Age  Dly  lo
hex             hex      hex              sec  sec  sec
800000e0804c9c00 200000   800000e0804c9c00 1     20  15  2

Port IEEE 802.1W Parameters:
      <--- Config Params --->|<----- Current state ----->
Port  Pri PortPath P2P Edge Role      State      Designa- Designated
Num   Cost  Mac Port  Role      State      ted cost  bridge
1     128 200000  F  F  ROOT      FORWARDING 0      800000e0804c9c00
2     128 200000  F  F  DESIGNATED FORWARDING 200000 800000e080541700
3     128 200000  F  F  DESIGNATED FORWARDING 200000 800000e080541700
4     128 200000  F  F  BACKUP     DISCARDING 200000 800000e080541700
```

Syntax: `show 802-1w [vlan <vlan-id>]`

The **vlan <vlan-id>** parameter displays 802.1W information for the specified port-based VLAN.

The **show 802.1w** command shows the information listed in [Table 27](#).

TABLE 27 CLI display of 802.1W summary

This field...	Displays...
VLAN ID	The port-based VLAN that owns the STP instance. VLAN 1 is the default VLAN. If you have not configured port-based VLANs on this device, all 802.1W information is for VLAN 1.

TABLE 27 CLI display of 802.1W summary (Continued)

This field...	Displays...
Bridge IEEE 802.1W parameters	
Bridge Identifier	The ID of the bridge.
Bridge Max Age	The configured max age for this bridge. The default is 20.
Bridge Hello	The configured hello time for this bridge. The default is 2.
Bridge FwdDly	The configured forward delay time for this bridge. The default is 15.
Force-Version	The configured force version value. One of the following value is displayed: <ul style="list-style-type: none"> • 0 – The bridge has been forced to operate in an STP compatibility mode. • 2 – The bridge has been forced to operate in an 802.1W mode. (This is the default.)
txHoldCnt	The number of BPDUs that can be transmitted per Hello Interval. The default is 3.
Root Bridge Identifier	ID of the Root bridge that is associated with this bridge
Root Path Cost	The cost to reach the root bridge from this bridge. If the bridge is the root bridge, then this parameter shows a value of zero.
Designated Bridge Identifier	The bridge from where the root information was received. It can be from the root bridge itself, but it could also be from another bridge.
Root Port	The port on which the root information was received. This is the port that is connected to the Designated Bridge.
Max Age	<p>The max age is derived from the Root port. An 802.1W-enabled bridge uses this value, along with the hello and message age parameters to compute the effective age of an RST BPDU.</p> <p>The message age parameter is generated by the Designated port and transmitted in the RST BPDU. RST BPDUs transmitted by a Designated port of the root bridge contains a message value of zero.</p> <p>Effective age is the amount of time the Root port, Alternate port, or Backup port retains the information it received from its peer Designated port. Effective age is reset every time a port receives an RST BPDU from its Designated port. If a Root port does not receive an RST BPDU from its peer Designated port for a duration more than the effective age, the Root port ages out the existing information and recomputes the topology.</p> <p>If the port is operating in 802.1D compatible mode, then max age functionality is the same as in 802.1D (STP).</p>
Fwd Dly	<p>The number of seconds a non-edge Designated port waits until it can apply any of the following transitions, if the RST BPDU it receives does not have an agreed flag:</p> <ul style="list-style-type: none"> • Discarding state to learning state • Learning state to forwarding state <p>When a non-edge port receives the RST BPDU it goes into forwarding state within 4 seconds or after two hello timers expire on the port.</p> <p>Fwd Dly is also the number of seconds that a Root port waits for an RST BPDU with a proposal flag before it applies the state transitions listed above.</p> <p>If the port is operating in 802.1D compatible mode, then forward delay functionality is the same as in 802.1D (STP).</p>

TABLE 27 CLI display of 802.1W summary (Continued)

This field...	Displays...
Hello	The hello value derived from the Root port. It is the number of seconds between two Hello packets.
Port IEEE 802.1W parameters	
Port Num	The port number shown in a port# format.
Pri	The configured priority of the port. The default is 128 or 0x80.
Port Path Cost	The configured path cost on a link connected to this port.
P2P Mac	Indicates if the point-to-point-mac parameter is configured to be a point-to-point link: <ul style="list-style-type: none"> • T – The link is configured as a point-to-point link. • F – The link is not configured as a point-to-point link. This is the default.
Edge port	Indicates if the port is configured as an operational Edge port: <ul style="list-style-type: none"> • T – The port is configured as an Edge port. • F – The port is not configured as an Edge port. This is the default.
Role	The current role of the port: <ul style="list-style-type: none"> • Root • Designated • Alternate • Backup • Disabled Refer to “Bridges and bridge port roles” on page 107 for definitions of the roles.
State	The port current 802.1W state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled Refer to “Bridge port states” on page 111 and “Edge port and non-edge port states” on page 112.
Designated Cost	The best root path cost that this port received, including the best root path cost that it can transmit.
Designated Bridge	The ID of the bridge that sent the best RST BPDU that was received on this port.

To display detailed information about 802-1W, using the following command.

6 Configuring STP related features

```
PowerConnect#show 802-1w detail
=====
VLAN 1 - MULTIPLE SPANNING TREE (MSTP - IEEE 802.1W) ACTIVE
=====
BridgeId 800000e080541700, forceVersion 2, txHoldCount 3
Port 1 - Role: ROOT - State: FORWARDING
  PathCost 200000, Priority 128, AdminOperEdge F, AdminPt2PtMac F
  DesignatedPriority - Root: 0x800000e0804c9c00, Bridge: 0x800000e080541700
  ActiveTimers - rrWhile 4 rcvdInfoWhile 4
  MachineStates - PIM: CURRENT, PRT: ROOT_PORT, PST: FORWARDING
  TCM: ACTIVE, PPM: SENDING_STP, PTX: TRANSMIT_IDLE
  Received - RST BPDUs 0, Config BPDUs 1017, TCN BPDUs 0

Port 2 - Role: DESIGNATED - State: FORWARDING
  PathCost 200000, Priority 128, AdminOperEdge F, AdminPt2PtMac F
  DesignatedPriority - Root: 0x800000e0804c9c00, Bridge: 0x800000e080541700
  ActiveTimers - helloWhen 0
  MachineStates - PIM: CURRENT, PRT: DESIGNATED_PORT, PST: FORWARDING
  TCM: ACTIVE, PPM: SENDING_RSTP, PTX: TRANSMIT_IDLE
  Received - RST BPDUs 0, Config BPDUs 0, TCN BPDUs 0
```

Syntax: `show 802-1w detail [vlan <vlan-id>]`

The `vlan <vlan-id>` parameter displays 802.1W information for the specified port-based VLAN.

The `show spanning-tree 802.1W` command shows the following information.

TABLE 28 CLI display of show spanning-tree 802.1W

This field...	Displays...
VLAN ID	ID of the VLAN that owns the instance of 802.1W and whether or not it is active.
Bridge ID	ID of the bridge.
forceVersion	the configured version of the bridge: <ul style="list-style-type: none"> • 0 - The bridge has been forced to operate in an STP compatible mode. • 2 - The bridge has been forced to operate in an 802.1W mode.
txHoldCount	The number of BPDUs that can be transmitted per Hello Interval. The default is 3.
Port	ID of the port in sport#format.
Role	The current role of the port: <ul style="list-style-type: none"> • Root • Designated • Alternate • Backup • Disabled Refer to “Bridges and bridge port roles” on page 107 for definitions of the roles.

TABLE 28 CLI display of show spanning-tree 802.1W (Continued)

This field...	Displays...
State	<p>The port current 802.1W state. A port can have one of the following states:</p> <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled <p>Refer to “Bridge port states” on page 111 and “Edge port and non-edge port states” on page 112.</p>
Path Cost	The configured path cost on a link connected to this port.
Priority	The configured priority of the port. The default is 128 or 0x80.
AdminOperEdge	<p>Indicates if the port is an operational Edge port. Edge ports may either be auto-detected or configured (forced) to be Edge ports using the CLI:</p> <ul style="list-style-type: none"> • T – The port is and Edge port. • F – The port is not an Edge port. This is the default.
AdminP2PMac	<p>Indicates if the point-to-point-mac parameter is configured to be a point-to-point link:</p> <ul style="list-style-type: none"> • T – The link is a point-to-point link • F – The link is not a point-to-point link. This is the default.
DesignatedPriority	<p>Shows the following:</p> <ul style="list-style-type: none"> • Root – Shows the ID of the root bridge for this bridge. • Bridge – Shows the ID of the Designated bridge that is associated with this port.
ActiveTimers	<p>Shows what timers are currently active on this port and the number of seconds they have before they expire:</p> <ul style="list-style-type: none"> • rrWhile – Recent root timer. A non-zero value means that the port has recently been a Root port. • rcvdInfoWhile – Received information timer. Shows the time remaining before the information held by this port expires (ages out). This timer is initialized with the effective age parameter. (Refer to “Max Age” on page 140.) • rbWhile – Recent backup timer. A non-zero value means that the port has recently been a Backup port. • helloWhen – Hello period timer. The value shown is the amount of time between hello messages. • tcWhile – Topology change timer. The value shown is the interval when topology change notices can be propagated on this port. • fdWhile – Forward delay timer. • mdelayWhile – Migration delay timer. The amount of time that a bridge on the same LAN has to synchronize its migration state with this port before another BPDU type can cause this port to change the BPDU that it transmits.

TABLE 28 CLI display of show spanning-tree 802.1W (Continued)

This field...	Displays...
Machine States	<p>The current states of the various state machines on the port:</p> <ul style="list-style-type: none"> • PIM – State of the Port Information state machine. • PRT – State of the Port Role Transition state machine. • PST – State of the Port State Transition state machine. • TCM – State of the Topology Change state machine. • PPM – State of the Port Protocol Migration. • PTX – State of the Port Transmit state machine. <p>Refer to the section “State machines” on page 112 for details on state machines.</p>
Received	<p>Shows the number of BPDU types the port has received:</p> <ul style="list-style-type: none"> • RST BPDU – BPDU in 802.1W format. • Config BPDU – Legacy configuration BPDU (802.1D format). • TCN BPDU – Legacy topology change BPDU (802.1D format).

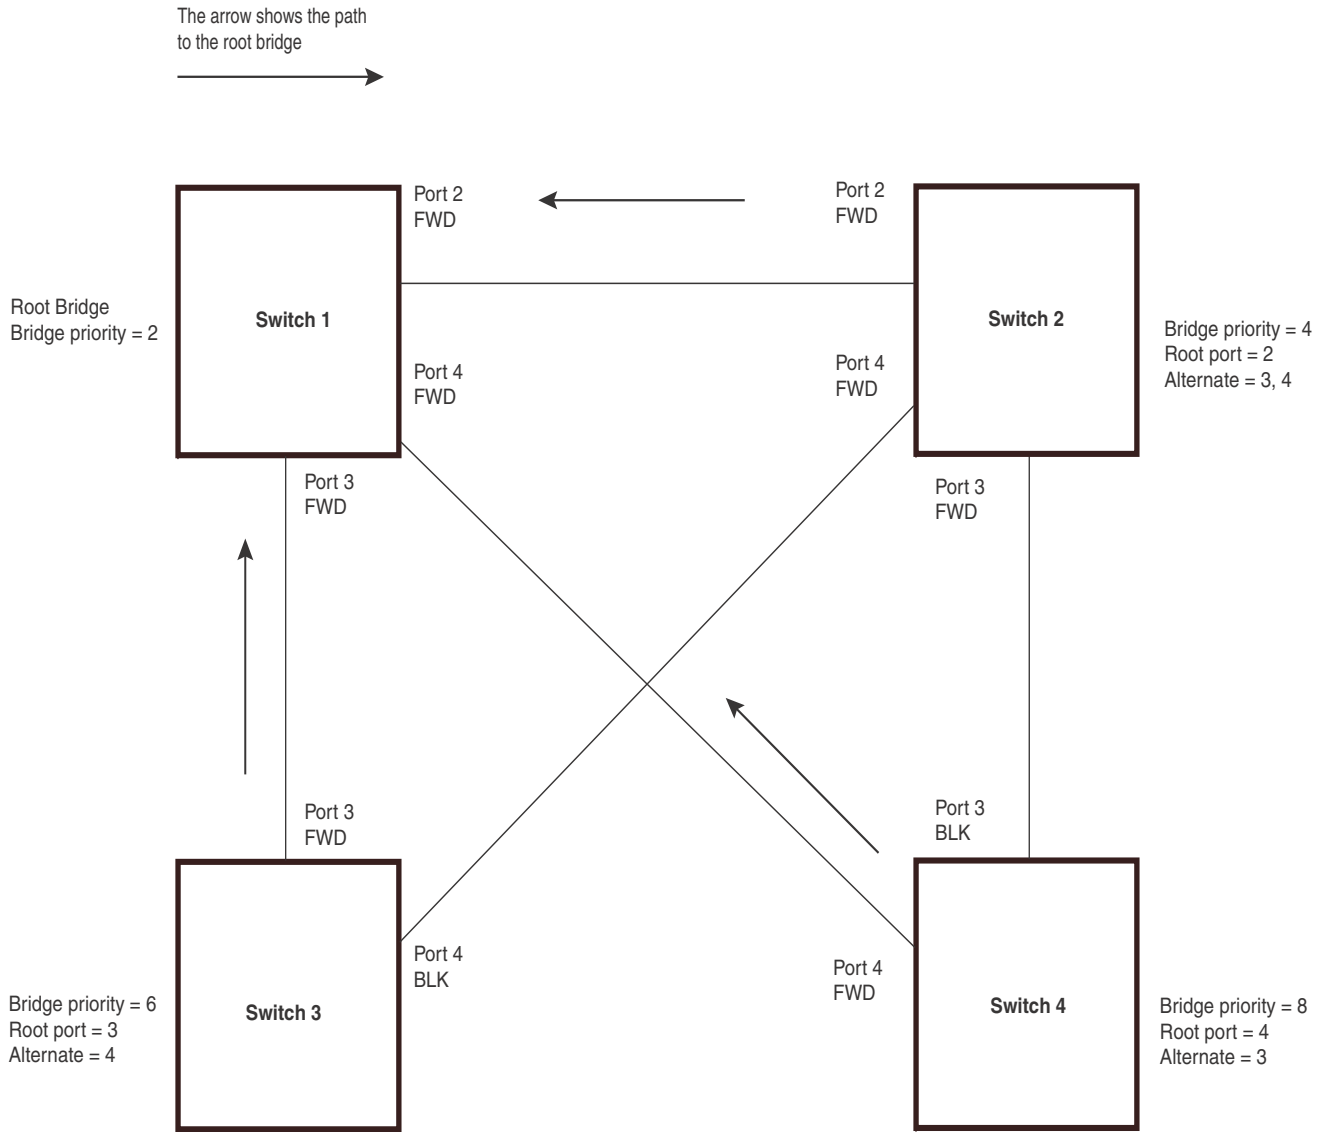
802.1W Draft 3

As an alternative to full 802.1W, you can configure 802.1W Draft 3. 802.1W Draft 3 provides a subset of the RSTP capabilities described in the 802.1W STP specification.

802.1W Draft 3 support is disabled by default. When the feature is enabled, if a root port on a device that is not the root bridge becomes unavailable, the device can automatically Switch over to an alternate root port, without reconvergence delays. 802.1W Draft 3 does not apply to the root bridge, since all the root bridge ports are always in the forwarding state.

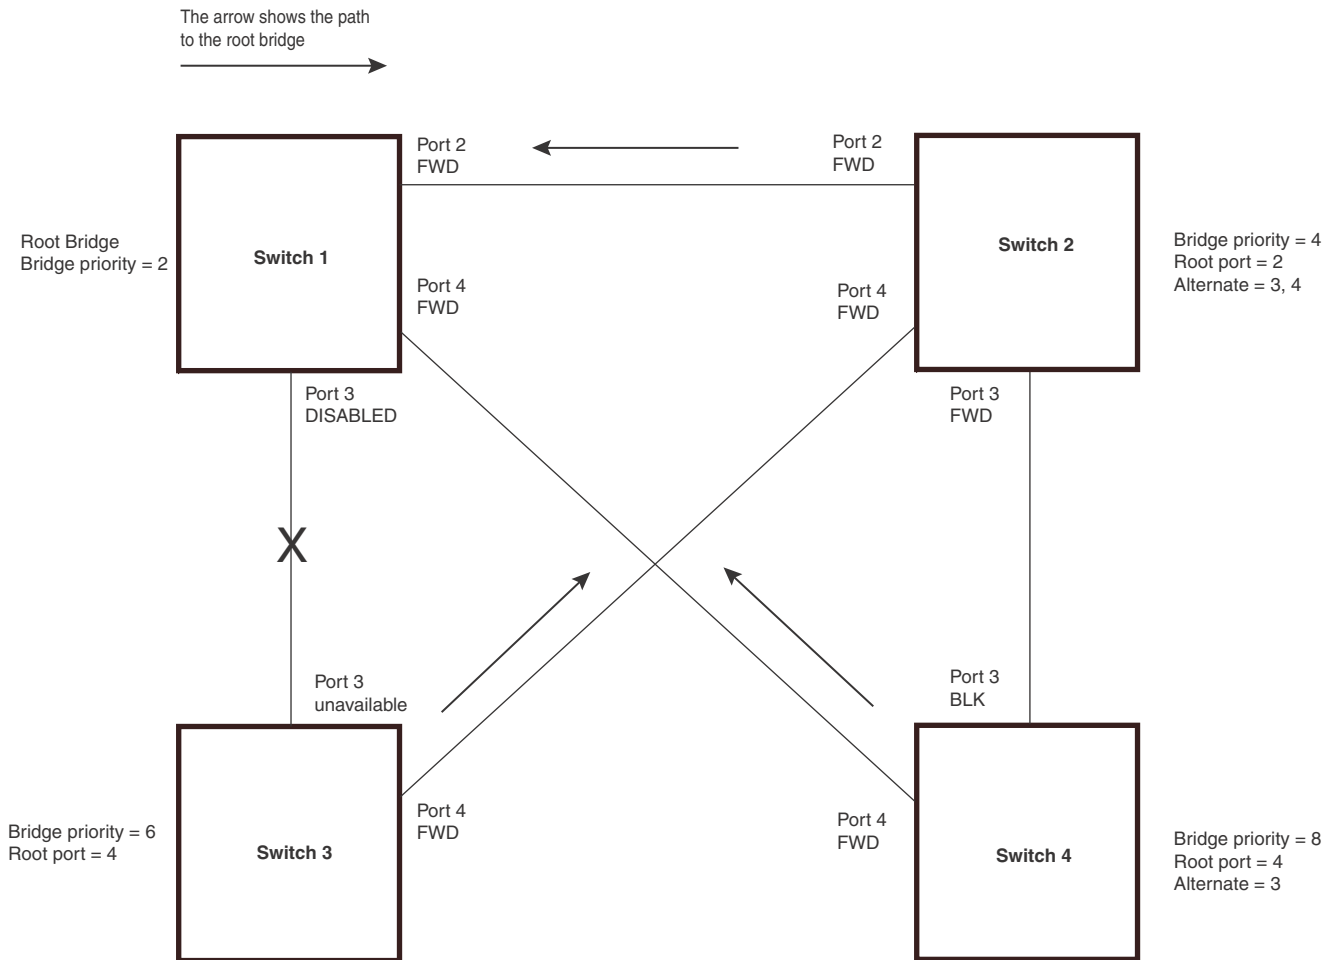
[Figure 25](#) shows an example of an optimal STP topology. In this topology, all the non-root bridges have at least two paths to the root bridge (Switch 1 in this example). One of the paths is through the root port. The other path is a backup and is through the alternate port. While the root port is in the forwarding state, the alternate port is in the blocking state.

FIGURE 25 802.1W Draft 3 RSTP ready for failover



If the root port on a Switch becomes unavailable, 802.1W Draft 3 immediately fails over to the alternate port, as shown in [Figure 26](#).

FIGURE 26 802.1W Draft 3 RSTP failover to alternate root port



In this example, port 3 on Switch 3 has become unavailable. In standard STP (802.1D), if the root port becomes unavailable, the Switch must go through the listening and learning stages on the alternate port to reconverge with the spanning tree. Thus, port 4 must go through the listening and learning states before entering the forwarding state and thus reconverging with the spanning tree.

802.1W Draft 3 avoids the reconvergence delay by calculating an alternate root port, and immediately failing over to the alternate port if the root port becomes unavailable. The alternate port is in the blocking state as long as the root port is in the forwarding state, but moves immediately to the active state if the root port becomes unavailable. Thus, using 802.1W Draft 3, Switch 3 immediately fails over to port 4, without the delays caused by the listening and learning states.

802.1W Draft 3 selects the port with the next-best cost to the root bridge. For example, on Switch 3, port 3 has the best cost to the root bridge and thus is selected by STP as the root port. Port 4 has the next-best cost to the root bridge, and thus is selected by 802.1W Draft 3 as the alternate path to the root bridge.

Once a failover occurs, the Switch no longer has an alternate root port. If the port that was an alternate port but became the root port fails, standard STP is used to reconverge with the network. You can minimize the reconvergence delay in this case by setting the forwarding delay on the root bridge to a lower value. For example, if the forwarding delay is set to 15 seconds (the default), change the forwarding delay to a value from 3 – 10 seconds.

During failover, 802.1W Draft 3 flushes the MAC addresses learned on the unavailable root port, selects the alternate port as the new root port, and places that port in the forwarding state. If traffic is flowing in both directions on the new root port, addresses are flushed (moved) in the rest of the spanning tree automatically.

Reconvergence time

Spanning tree reconvergence using 802.1W Draft 3 can occur within one second.

After the spanning tree reconverges following the topology change, traffic also must reconverge on all the bridges attached to the spanning tree. This is true regardless of whether 802.1W Draft 3 or standard STP is used to reconverge the spanning tree.

Traffic reconvergence happens after the spanning tree reconvergence, and is achieved by flushing the Layer 2 information on the bridges:

- Following 802.1W Draft 3 reconvergence of the spanning tree, traffic reconvergence occurs in the time it takes for the bridge to detect the link changes plus the STP maximum age set on the bridge.
- If standard STP reconvergence occurs instead, traffic reconvergence takes two times the forward delay plus the maximum age.

NOTE

802.1W Draft 3 does not apply when a failed root port comes back up. When this happens, standard STP is used.

Configuration considerations

802.1W Draft 3 is disabled by default. To ensure optimal performance of the feature before you enable it, do the following:

- Configure the bridge priorities so that the root bridge is one that supports 802.1W Draft 3. (Use a device or third-party device that supports 802.1W Draft 3.)
- Change the forwarding delay on the root bridge to a value lower than the default 15 seconds. Dell recommends a value from 3 – 10 seconds. The lower forwarding delay helps reduce reconvergence delays in cases where 802.1W Draft 3 is not applicable, such as when a failed root port comes back up.
- Configure the bridge priorities and root port costs so that each device has an active path to the root bridge if its root port becomes unavailable. For example, port 4 on Switch 3 is connected to port 4 on Switch 2, which has the second most favorable bridge priority in the spanning tree.

NOTE

If reconvergence involves changing the state of a root port on a bridge that supports 802.1D STP but not 802.1W Draft 3, then reconvergence still requires the amount of time it takes for the ports on the 802.1D bridge to change state to forwarding (as needed), and receive BPDUs from the root bridge for the new topology.

Enabling 802.1W Draft 3

802.1W Draft 3 is disabled by default. The procedure for enabling the feature differs depending on whether single STP is enabled on the device.

NOTE

STP must be enabled before you can enable 802.1W Draft 3.

Enabling 802.1W Draft 3 when single STP is not enabled

By default, each port-based VLAN on the device has its own spanning tree. To enable 802.1W Draft 3 in a port-based VLAN, enter commands such as the following.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#spanning-tree rstp
```

Syntax: [no] spanning-tree rstp

This command enables 802.1W Draft 3. You must enter the command separately in each port-based VLAN in which you want to run 802.1W Draft 3.

NOTE

This command does not also enable STP. To enable STP, first enter the **spanning-tree** command without the **rstp** parameter. After you enable STP, enter the **spanning-tree rstp** command to enable 802.1W Draft 3.

To disable 802.1W Draft 3, enter the following command.

```
PowerConnect(config-vlan-10)#no spanning-tree rstp
```

Enabling 802.1W Draft 3 when single STP is enabled

To enable 802.1W Draft 3 on a device that is running single STP, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)#spanning-tree single rstp
```

Syntax: [no] spanning-tree single rstp

This command enables 802.1W Draft 3 on the whole device.

NOTE

This command does not also enable single STP. To enable single STP, first enter the **spanning-tree single** command without the **rstp** parameter. After you enable single STP, enter the **spanning-tree single rstp** command to enable 802.1W Draft 3.

To disable 802.1W Draft 3 on a device that is running single STP, enter the following command.

```
PowerConnect(config)#no spanning-tree single rstp
```

Single Spanning Tree (SSTP)

By default, each port-based VLAN on a device runs a separate spanning tree, which you can enable or disable on an individual VLAN basis.

Alternatively, you can configure a device to run a single spanning tree across all ports and VLANs on the device. The Single STP feature (SSTP) is especially useful for connecting a device to third-party devices that run a single spanning tree in accordance with the 802.1Q specification.

SSTP uses the same parameters, with the same value ranges and defaults, as the default STP support on devices. Refer to “[STP parameters and defaults](#)” on page 93.

SSTP defaults

SSTP is disabled by default. When you enable the feature, all VLANs on which STP is enabled become members of a single spanning tree. All VLANs on which STP is disabled are excluded from the single spanning tree.

To add a VLAN to the single spanning tree, enable STP on that VLAN. To remove a VLAN from the single spanning tree, disable STP on that VLAN.

When you enable SSTP, all the ports that are in port-based VLANs with STP enabled become members of a single spanning tree domain. Thus, the ports share a single BPDU broadcast domain. The device places all the ports in a non-configurable VLAN, 4094, to implement the SSTP domain. However, this VLAN does not affect port membership in the port-based VLANs you have configured. Other broadcast traffic is still contained within the individual port-based VLANs. Therefore, you can use SSTP while still using your existing VLAN configurations without changing your network. In addition, SSTP does not affect 802.1Q tagging. Tagged and untagged ports alike can be members of the single spanning tree domain.

NOTE

When SSTP is enabled, the BPDUs on tagged ports go out untagged.

If you disable SSTP, all VLANs that were members of the single spanning tree run MSTP instead. In MSTP, each VLAN has its own spanning tree. VLANs that were not members of the single spanning tree were not enabled for STP. Therefore, STP remains disabled on those VLANs.

Enabling SSTP

To enable SSTP, use one of the following methods.

NOTE

If the device has only one port-based VLAN (the default VLAN), then the device is already running a single instance of STP. In this case, you do not need to enable SSTP. You need to enable SSTP only if the device contains more than one port-based VLAN and you want all the ports to be in the same STP broadcast domain.

To configure the device to run a single spanning tree, enter the following command at the global CONFIG level.

```
PowerConnect(config)#spanning-tree single
```

NOTE

If the device has only one port-based VLAN, the CLI command for enabling SSTP is not listed in the CLI. The command is listed only if you have configured a port-based VLAN.

To change a global STP parameter, enter a command such as the following at the global CONFIG level.

```
PowerConnect(config) spanning-tree single priority 2
```

This command changes the STP priority for all ports to 2.

To change an STP parameter for a specific port, enter commands such as the following.

6 PVST/PVST+ compatibility

```
PowerConnect(config) spanning-tree single ethernet 1 priority 10
```

The commands shown above override the global setting for STP priority and set the priority to 10 for port 1.

Here is the syntax for the global STP parameters.

Syntax: [no] spanning-tree single [forward-delay <value>] [hello-time <value>] | [maximum-age <time>] | [priority <value>]

Here is the syntax for the STP port parameters.

Syntax: [no] spanning-tree single [ethernet <portnum> path-cost <value> | priority <value>]

NOTE

Both commands listed above are entered at the global CONFIG level.

Displaying SSTP information

To verify that SSTP is in effect, enter the following commands at any level of the CLI.

```
PowerConnect#show span
```

Syntax: show span [vlan <vlan-id>] | [pvst-mode] | [<num>] | [detail [vlan <vlan-id> [ethernet <portnum>] | <num>]]

The **vlan <vlan-id>** parameter displays STP information for the specified port-based VLAN.

The **pvst-mode** parameter displays STP information for the device Per VLAN Spanning Tree (PVST+) compatibility configuration. Refer to [“PVST/PVST+ compatibility”](#) on page 150.

The **<num>** parameter displays only the entries after the number you specify. For example, on a device with three port-based VLANs, if you enter 1, then information for the second and third VLANs is displayed, but information for the first VLAN is not displayed. Information is displayed according to VLAN number, in ascending order. The entry number is not the same as the VLAN number. For example, if you have port-based VLANs 1, 10, and 2024, then the command output has three STP entries. To display information for VLANs 10 and 2024 only, enter **show span 1**.

The **detail** parameter and its additional optional parameters display detailed information for individual ports. Refer to [“Displaying detailed STP information for each interface”](#) on page 102.

PVST/PVST+ compatibility

The PowerConnect family of switches support Cisco's Per VLAN Spanning Tree plus (PVST+), by allowing the device to run multiple spanning trees (MSTP) while also interoperating with IEEE 802.1Q devices¹.

NOTE

Dell ports automatically detect PVST+ BPDUs and enable support for the BPDUs once detected. You do not need to perform any configuration steps to enable PVST+ support. However, to support the IEEE 802.1Q BPDUs, you might need to enable dual-mode support.

1. Cisco user documentation for PVST/PVST+ refers to the IEEE 802.1Q spanning tree as the **Common Spanning Tree (CST)**.

Support for Cisco's Per VLAN Spanning Tree plus (PVST+), allows a device to run multiple spanning trees (MSTP) while also interoperating with IEEE 802.1Q devices. Dell ports automatically detect PVST+ BPDUs and enable support for the BPDUs once detected. The enhancement allows a port that is in PVST+ compatibility mode due to auto-detection to revert to the default MSTP mode when one of the following events occurs:

- The link is disconnected or broken
- The link is administratively disabled
- The link is disabled by interaction with the link-keepalive protocol

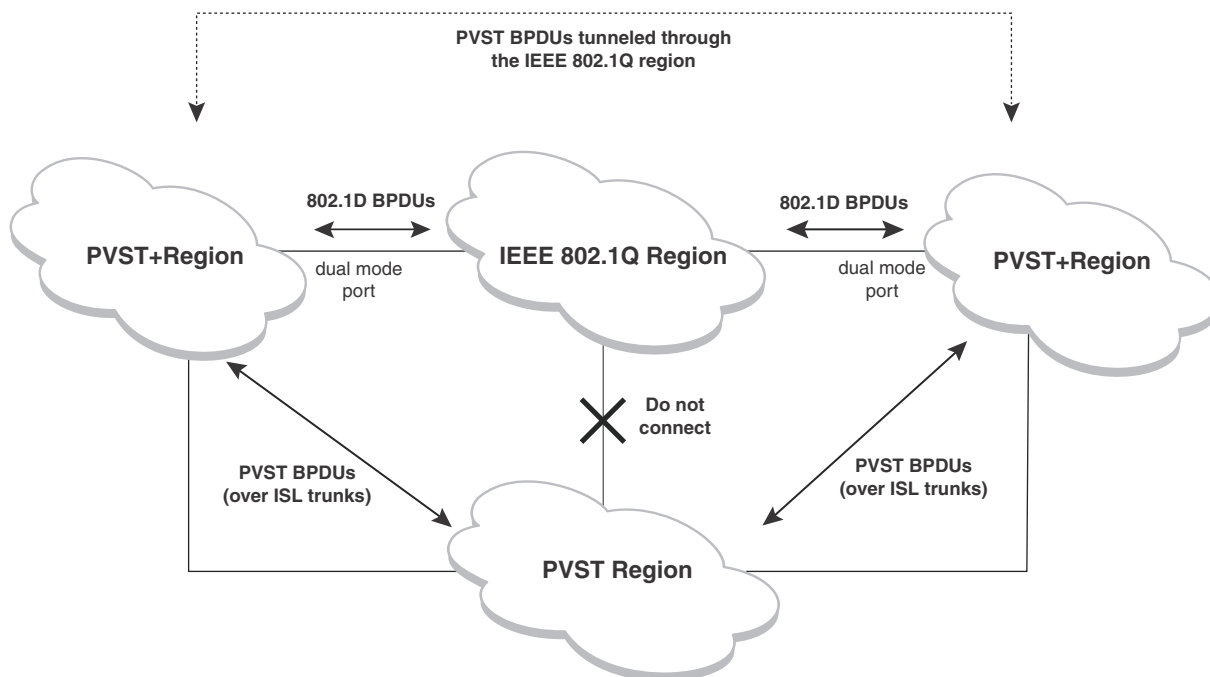
This enhancement allows a port that was originally interoperating with PVST+ to revert to MSTP when connected to a device.

Overview of PVST and PVST+

Per VLAN Spanning Tree (PVST) is a Cisco proprietary protocol that allows a Cisco device to have multiple spanning trees. The Cisco device can interoperate with spanning trees on other PVST devices but cannot interoperate with IEEE 802.1Q devices. An IEEE 802.1Q device has all its ports running a single spanning tree. **PVST+** is an extension of PVST that allows a Cisco device to also interoperate with devices that are running a single spanning tree (IEEE 802.1Q).

Enhanced PVST+ support allows a device to interoperate with PVST spanning trees and the IEEE 802.1Q spanning tree at the same time.

IEEE 802.1Q and PVST regions cannot interoperate directly but can interoperate indirectly through PVST+ regions. PVST BPDUs are tunneled through 802.1Q regions, while PVST BPDUs for VLAN 1 (the IEEE 802.1Q VLAN) are processed by PVST+ regions. [Figure 27](#) shows the interaction of IEEE 802.1Q, PVST, and PVST+ regions.

FIGURE 27 Interaction of IEEE 802.1Q, PVST, and PVST+ regions

VLAN tags and dual mode

The **dual-mode** feature enables a port to send and receive both tagged and untagged frames. When the dual-mode feature is enabled on a port, the port is an untagged member of one of its VLANs and is at the same time a tagged member of all its other VLANs. The untagged frames are supported on the port **Port Native VLAN**.

The dual-mode feature must be enabled on a Dell port in order to interoperate with another vendor device. Some vendors use VLAN 1 by default to support the IEEE 802.1Q-based standard spanning tree protocols, such as 802.1d and 802.1w for sending untagged frames on VLAN 1. On Dell switches, by default, the **Port Native VLAN** is the same as the **Default VLAN**, which is VLAN 1. Thus, to support IEEE 802.1Q in a typical configuration, a port must be able to send and receive untagged frames for VLAN 1 and tagged frames for the other VLANs, and interoperate with other vendor devices using VLAN 1.

If you want to use tagged frames on VLAN 1, you can change the default VLAN ID to an ID other than 1. You also can specify the VLAN on which you want the port to send and receive untagged frames (the Port Native VLAN). The Port Native VLAN ID does not need to be the same as the default VLAN. Make sure that the untagged (native) VLAN is also changed on the interoperating vendor side to match that on the Dell side.

To support the IEEE 802.1Q with non-standard proprietary protocols such as PVST and PVST+, a port must always send and receive untagged frames on VLAN 1 on both sides. In this case, enable the dual-mode 1 feature to allow untagged BPDUs on VLAN 1 and use Native VLAN 1 on the interoperating vendor side. You should not use VLAN 1 for tagged frames in this case.

Configuring PVST+ support

PVST+ support is automatically enabled when the port receives a PVST BPDU. You can manually enable the support at any time or disable the support if desired.

If you want a tagged port to also support IEEE 802.1Q BPDUs, you need to enable the dual-mode feature on the port. The dual-mode feature is disabled by default and must be enabled manually.

A port that is in PVST+ compatibility mode due to auto-detection reverts to the default MSTP mode when one of the following events occurs:

- The link is disconnected or broken
- The link is administratively disabled
- The link is disabled by interaction with the link-keepalive protocol

This allows a port that was originally interoperating with PVST+ to revert to MSTP when connected to a device.

Enabling PVST+ support manually

To immediately enable PVST+ support on a port, enter commands such as the following.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#pvst-mode
```

Syntax: [no] pvst-mode

NOTE

If you disable PVST+ support, the software still automatically enables PVST+ support if the port receives a BPDU with PVST+ format.

NOTE

If 802.1W and pvst-mode (either by auto-detection or by explicit configuration) are enabled on a tagged VLAN port, 802.1W will treat the PVST BPDUs as legacy 802.1D BPDUs.

Enabling dual-mode support

To enable the dual-mode feature on a port, enter the following command at the interface configuration level for the port.

```
PowerConnect(config-if-1)#dual-mode
```

Syntax: [no] dual-mode [<vlan-id>]

The <vlan-id> specifies the port Port Native VLAN. This is the VLAN on which the port will support untagged frames. By default, the Port Native VLAN is the same as the default VLAN (which is VLAN 1 by default).

For more information about the dual-mode feature, refer to [“Dual-mode VLAN ports”](#) on page 305.

Displaying PVST+ support information

To display PVST+ information for ports on a device, enter the following command at any level of the CLI.

6 PVST/PVST+ compatibility

```
PowerConnect#show span pvst-mode
PVST+ Enabled on:
Port      Method
1         Set by configuration
2         Set by configuration
10        Set by auto-detect
12        Set by configuration
24        Set by auto-detect
```

Syntax: show span pvst-mode

This command displays the following information.

TABLE 29 CLI display of PVST+ information

This field...	Displays...
Port	The Dell port number. NOTE: The command lists information only for the ports on which PVST+ support is enabled.
Method	The method by which PVST+ support was enabled on the port. The method can be one of the following: <ul style="list-style-type: none">• Set by configuration – You enabled the support.• Set by auto-detect – The support was enabled automatically when the port received a PVST+ BPDU.

Configuration examples


The following examples show configuration examples for two common configurations:

- Untagged IEEE 802.1Q BPDUs on VLAN 1 and tagged PVST+ BPDUs on other VLANs
- Tagged IEEE 802.1Q BPDUs on VLAN 1 and untagged BPDUs on another VLAN

Tagged port using default VLAN 1 as its port native VLAN

Figure 28 shows an example of a PVST+ configuration that uses VLAN 1 as the untagged default VLAN and VLANs 2, 3, and 4 as tagged VLANs.

FIGURE 28 Default VLAN 1 for untagged BPDU



To implement this configuration, enter the following commands.

Commands on the Dell Device

```
PowerConnect(config)#vlan-group 1 vlan 2 to 4
PowerConnect(config-vlan-group-1)#tagged ethernet 1
PowerConnect(config-vlan-group-1)#exit
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#dual-mode
PowerConnect(config-if-1)#pvst-mode
```

These commands configure a VLAN group containing VLANs 2, 3, and 4, add port 1 as a tagged port to the VLANs, and enable the dual-mode feature and PVST+ support on the port. The dual-mode feature allows the port to send and receive untagged frames for the default VLAN (VLAN 1 in this case) in addition to tagged frames for VLANs 2, 3, and 4. Enabling the PVST+ support ensures that the port is ready to send and receive PVST+ BPDUs. If you do not manually enable PVST+ support, the support is not enabled until the port receives a PVST+ BPDU.

The configuration leaves the default VLAN and the port Port Native VLAN unchanged. The default VLAN is 1 and the port Port Native VLAN also is 1. The dual-mode feature supports untagged frames on the default VLAN only. Thus, port 1 can send and receive untagged BPDUs for VLAN 1 and can send and receive tagged BPDUs for the other VLANs.


Port 1 will process BPDUs as follows:

- Process IEEE 802.1Q BPDUs for VLAN 1.
- Process tagged PVST BPDUs for VLANs 2, 3, and 4.
- Drop untagged PVST BPDUs for VLAN 1.

Untagged port using VLAN 2 as port native VLAN

Figure 29 shows an example in which a port Port Native VLAN is not VLAN 1. In this case, VLAN 1 uses tagged frames and VLAN 2 uses untagged frames.

FIGURE 29 Port Native VLAN 2 for Untagged BPDUs



To implement this configuration, enter the following commands.

Commands on the Dell Device

```
PowerConnect(config)#default-vlan-id 4000
PowerConnect(config)#vlan 1
PowerConnect(config-vlan-1)#tagged ethernet 1
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#vlan 2
PowerConnect(config-vlan-2)#tagged ethernet 1
PowerConnect(config-vlan-2)#exit
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#dual-mode 2
PowerConnect(config-if-1)#pvst-mode
PowerConnect(config-if-1)#exit
```

These commands change the default VLAN ID, configure port 1 as a tagged member of VLANs 1 and 2, and enable the dual-mode feature and PVST+ support on port 1. Since VLAN 1 is tagged in this configuration, the default VLAN ID must be changed from VLAN 1 to another VLAN ID.

Changing the default VLAN ID from 1 allows the port to process tagged frames for VLAN 1. VLAN 2 is specified with the **dual-mode** command, which makes VLAN 2 the port Port Native VLAN. As a result, the port processes untagged frames and untagged PVST BPDUs on VLAN 2.

NOTE

Although VLAN 2 becomes the port untagged VLAN, the CLI still requires that you add the port to the VLAN as a tagged port, since the port is a member of more than one VLAN.

Port 1 will process BPDUs as follows:

- Process IEEE 802.1Q BPDUs for VLAN 1.
- Process untagged PVST BPDUs for VLAN 2.
- Drop tagged PVST BPDUs for VLAN 1.

Note that when VLAN 1 is not the default VLAN, the ports must have the dual-mode feature enabled in order to process IEEE 802.1Q BPDUs.

For example, the following configuration is incorrect.

```
PowerConnect(config)#default-vlan-id 1000
PowerConnect(config)#vlan 1
PowerConnect(config-vlan-1)#tagged ethernet 1 to 2
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#pvst-mode
PowerConnect(config-if-1)#exit
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-2)#pvst-mode
PowerConnect(config-if-2)#exit
```

In the configuration above, all PVST BPDUs associated with VLAN 1 would be discarded. Since IEEE BPDUs associated with VLAN 1 are untagged, they are discarded because the ports in VLAN 1 are tagged. Effectively, the BPDUs are never processed by the Spanning Tree Protocol. STP assumes that there is no better bridge on the network and sets the ports to FORWARDING. This could cause a Layer 2 loop.

The following configuration is correct.

```
PowerConnect(config)#default-vlan-id 1000
PowerConnect(config)#vlan 1
PowerConnect(config-vlan-1)#tagged ethernet 1 to 2
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#pvst-mode
PowerConnect(config-if-1)#dual-mode
PowerConnect(config-if-1)#exit
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-2)#pvst-mode
PowerConnect(config-if-2)#dual-mode
PowerConnect(config-if-2)#exit
```

Setting the ports as dual-mode ensures that the untagged IEEE 802.1Q BPDUs reach the VLAN 1 instance.

PVRST compatibility

PVRST, the "rapid" version of per-VLAN spanning tree (PVST), is a Cisco proprietary protocol. PVRST corresponds to the Dell full implementation of IEEE 802.1w (RSTP). Likewise, PVST, also a Cisco proprietary protocol, corresponds to the Dell implementation of IEEE 802.1D (STP).

PowerConnect B-Series TI24X devices also support PVRST compatibility. When it receives PVRST BPDUs on a port configured to run 802.1w, it recognizes and processes these BPDUs and continues to operate in 802.1w mode.

PVRST compatibility is automatically enabled, when a port receives a PVRST BPDU.

BPDU guard

In an STP environment, switches, end stations, and other Layer 2 devices use Bridge Protocol Data Units (BPDUs) to exchange information that STP will use to determine the best path for data flow.

The BPDU guard, an enhancement to STP, removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP.

In some instances, it is unnecessary for a connected device, such as an end station, to initiate or participate in an STP topology change. In this case, you can enable the STP BPDU guard feature on the Dell port to which the end station is connected. STP BPDU guard shuts down the port and puts it into an errdisable state. This disables the connected device's ability to initiate or participate in an STP topology. A log message is then generated for a BPDU guard violation, and a CLI message is displayed to warn the network administrator of a severe invalid configuration. The BPDU guard feature provides a secure response to invalid configurations because the administrator must manually put the interface back in service if errdisable recovery is not enabled.

Enabling BPDU protection by port

You enable STP BPDU guard on individual interfaces. The feature is disabled by default.

To enable STP BPDU guard on a specific port, enter commands such as the following.

```
PowerConnect(config) interface ethe 1
PowerConnect(config-if-e10000-1)#stp-bpdu-guard
```

Syntax: [no] stp-bpdu-guard

The **no** parameter disables the BPDU guard on this interface.

You can also use the multiple interface command to enable this feature on multiple ports at once.

Example

```
PowerConnect(config)#interface ethernet 1 to 9
PowerConnect(config-mif-1-9)#stp-bpdu-guard
PowerConnect(config-mif-1-9)#
```

This will enable stp-bpdu-guard on ports 1 to 9

Re-enabling ports disabled by BPDU guard

When a BPSU Guard-enabled port is disabled by BPSU Guard, the device will place the port in **errdisable** state and display a message on the console indicating that the port is errdisabled (refer to “[Example console messages](#)” on page 159). In addition, the **show interface** command output will indicate that the port is errdisabled.

Example

```
PowerConnect#show int e 2
Gigabit Ethernet2 is ERR-DISABLED (bpduguard), line protocol is down
```

To re-enable a port that is in **errdisable** state, you must first disable the port then re-enable it. Enter commands such as the following.

```
PowerConnect(config)#int e 2
PowerConnect(config-if-e10000-2)#disable
PowerConnect(config-if-e10000-2)#enable
```

If you attempt to enable an errdisabled port without first disabling it, the following error message will appear on the console.

```
PowerConnect(config-if-e10000-2)#enable
Port 2 is errdisabled, do disable first and then enable to enable it
```

Displaying the BPDU guard status

To display the BPDU guard state, enter the **show running configuration** or the **show stp-bpdu-guard** command.

Example configurations

Example

The following example shows how to configure BPDU guard at the interface level and to verify the configuration by issuing the **show stp-bpdu-guard** and the **show interface** commands.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#stp-bpdu-guard
PowerConnect(config-if-e10000-1)#
PowerConnect(config-if-e10000-1)#show stp-bpdu-guard
BPDU Guard Enabled on:
Port
1
PowerConnect(config-if-e10000-1)#
PowerConnect(config-if-e10000-1)#show interfaces ethernet 1
GigabitEthernet1 is up, line protocol is up
Hardware is GigabitEthernet, address is 000c.dba0.7100 (bia 000c.dba0.7100)
Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual MDI
Member of L2 VLAN ID 2, port is untagged, port state is FORWARDING
BPDU guard is Enabled, ROOT protect is Disabled
STP configured to ON, priority is level0, flow control enabled
mirror disabled, monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 96 bits-time, IPG GMII 96 bits-time
IP MTU 1500 bytes
```

```
300 second input rate: 8 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 256 bits/sec, 0 packets/sec, 0.00% utilization
88 packets input, 15256 bytes, 0 no buffer
Received 75 broadcasts, 13 multicasts, 0 unicasts
1 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
4799 packets output, 313268 bytes, 0 underruns
Transmitted 90 broadcasts, 4709
```

Example console messages

A console message such as the following is generated after a BPDU guard violation occurs on a system that is running MSTP.

```
PowerConnect(config-if-e10000-23)#MSTP: Received BPDU on BPDU guard enabled Port
23,errdisable Port 23
```

A console message such as the following is generated after a BPDU guard violation occurs on a system that is running STP.

```
PowerConnect(config)#STP: Received BPDU on BPDU guard enabled Port 23 (vlan=1),
errdisable Port 23
```

A console message such as the following is generated after a BPDU guard violation occurs on a system that is running RSTP.

```
PowerConnect(config-vlan-1)#RSTP: Received BPDU on BPDU guard enabled Port 23
(vlan=1),errdisable Port 23
```

Root guard

The standard STP (802.1D), RSTP (802.1W) or 802.1S does not provide any way for a network administrator to securely enforce the topology of a switched layer 2 network. The forwarding topology of a switched network is calculated based on the root bridge position, along with other parameters. This means any switch can be the root bridge in a network as long as it has the lowest bridge ID. The administrator cannot enforce the position of the root bridge. A better forwarding topology comes with the requirement to place the root bridge at a specific predetermined location. Root Guard can be used to predetermine a root bridge location and prevent rogue or unwanted switches from becoming the root bridge.

When root guard is enabled on a port, it keeps the port in a designated role. If the port receives a superior STP Bridge Protocol Data Units (BPDU), it puts the port into a ROOT-INCONSISTANT state and triggers a log message and an SNMP trap. The ROOT-INCONSISTANT state is equivalent to the BLOCKING state in 802.1D and to the DISCARDING state in 802.1W. No further traffic is forwarded on this port. This allows the bridge to prevent traffic from being forwarded on ports connected to rogue or misconfigured STP bridges.

Once the port stops receiving superior BPDUs, root guard automatically sets the port back to learning, and eventually to a forwarding state through the spanning-tree algorithm.

Configure root guard on all ports where the root bridge should not appear. This establishes a protective network perimeter around the core bridged network, cutting it off from the user network.

NOTE

Root guard may prevent network connectivity if it is improperly configured. Root guard must be configured on the perimeter of the network rather than the core.

NOTE

Root guard is not supported when MSTP is enabled.

Enabling STP root guard

An STP root guard is configured on an interface by entering commands similar to the following.

```
PowerConnect(config)#interface ethernet 5
PowerConnect(config-if-e10000-5)spanning-tree root-protect
```

Syntax: [no] spanning-tree root-protect

Enter the **no** form of the command to disable STP root guard on the port.

Displaying the STP root guard

To display the STP root guard state, enter the **show running configuration** or the **show spanning-tree root-protect** command.

```
PowerConnect#show spanning-tree root-protect
Root Protection Enabled on:
Port 1
```

Syntax: show spanning-tree root-protect

Displaying the root guard by VLAN

You can display root guard information for all VLANs or for a specific VLAN. For example, to display root guard violation information for VLAN 7.

Syntax: show spanning-tree [*<vlan-id>*]

If you do not specify a *<vlan-id>*, information for all VLANs is displayed. For example, to display root guard violation information for VLAN 7.

```
PowerConnect#show spanning-tree vlan 7
STP instance owned by VLAN 7
Global STP (IEEE 802.1D) Parameters:
VLAN Root Root Root Prio Max He- Ho- Fwd Last Chg Bridge
ID ID Cost Port rity Age llo ld dly Chang cnt Address
Hex sec sec sec sec sec
7 a000000011112220 0 Root a000 20 2 1 15 4 4 000011112220
Port STP Parameters:
Port Prio Path State Fwd Design Designated Designated
Num rity Cost Trans Cost Root Bridge
Hex
1 80 19 ROOT-INCONS 2 0 a000000011112220 a000000011112220
```

802.1s Multiple Spanning Tree Protocol

Multiple Spanning Tree Protocol (MSTP), as defined in IEEE 802.1s, allows multiple VLANs to be managed by a single STP instance and supports per-VLAN STP. As a result, several VLANs can be mapped to a reduced number of spanning-tree instances. This ensures loop-free topology for one or more VLANs that have the similar layer-2 topology. The Dell implementation supports up to 16 spanning tree instances in an MSTP enabled bridge which means that it can support up to 16 different Layer 2 topologies. The spanning tree algorithm used by MSTP is RSTP which provides quick convergence.

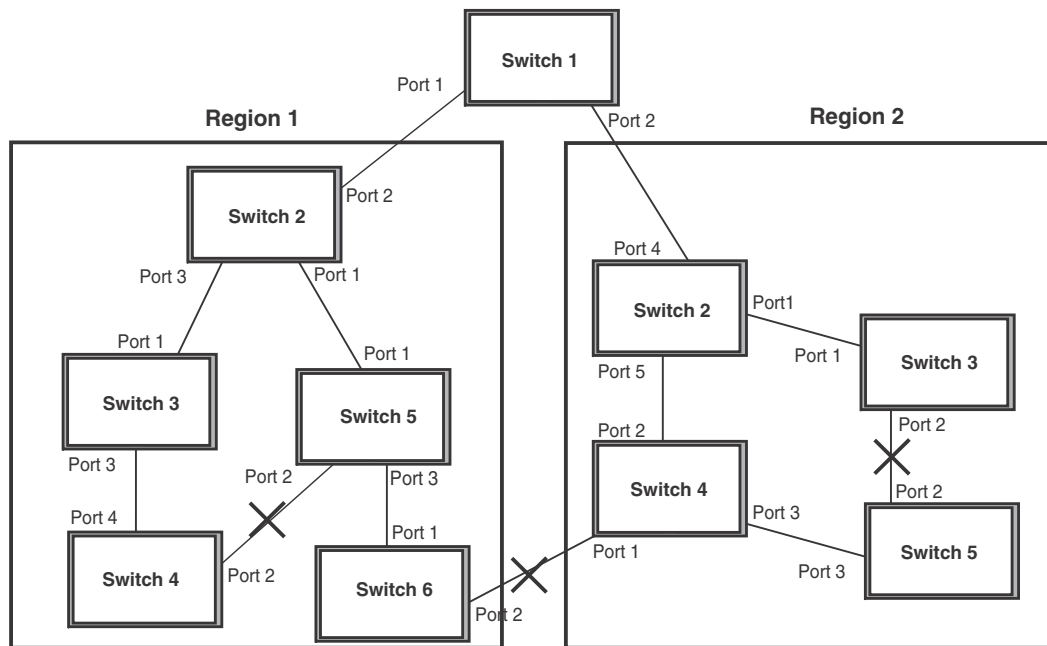
Multiple spanning-tree regions

Using MSTP, the entire network runs a common instance of RSTP. Within that common instance, one or more VLANs can be individually configured into distinct regions. The entire network runs the common spanning tree instance (CST) and the regions run a local instance. The local instance is known as Internal Spanning Tree (IST). The CST treats each instance of IST as a single bridge. Consequently, ports are blocked to prevent loops that might occur within an IST and also throughout the CST. With the exception of the provisions for multiple instances, MSTP operates exactly like RSTP.

For example, in [Figure 30](#) a network is configured with two regions: Region 1 and Region 2. The entire network is running an instance of CST. Each of the regions is running an instance of IST. In addition, this network contains Switch 1 running MSTP that is not configured in a region and consequently is running in the CIST instance. In this configuration, the regions are each regarded as a single bridge to the rest of the network, as is Switch 1. The CST prevents loops from occurring across the network. Consequently, a port is blocked at port 2 of switch 4.

Additionally, loops must be prevented in each of the IST instances. Within the IST Region 1, a port is blocked at port 2 of switch 4 to prevent a loop in that region. Within Region 2, a port is blocked at port 2 of switch 3 to prevent a loop in that region.

FIGURE 30 MSTP configured network



The following definitions describe the STP instances that define an MSTP configuration.

Common Spanning (CST) – CST is defined in 802.1q and assumes one spanning-tree instance for the entire bridged network regardless of the number of VLANs. In MSTP, an MSTP region appears as a virtual bridge that runs CST.

Internal Spanning Tree (IST) – IST is a new terminology introduced in 802.1s. An MSTP bridge must handle at least these two instances: one IST and one or more MSTIs (Multiple Spanning Tree Instances). Within each MST region, the MSTP maintains multiple spanning-tree instances. Instance 0 is a special instance known as IST, which extends CST inside the MST region. IST always exists if the switch runs MSTP. Besides IST, this implementation supports up to 15 MSTIs, numbered from 1 to 4094.

Common and Internal Spanning Trees (CIST) – CIST is a collection of the ISTs in each MST region and the CST that interconnects the MST regions and single spanning trees.

Multiple Spanning Tree Instance (MSTI) – The MSTI is identified by an MST identifier (MSTid) value between 1 and 4094.

MSTP Region – These are clusters of bridges that run multiple instances of the MSTP protocol. Multiple bridges detect that they are in the same region by exchanging their configuration (instance to VLAN mapping), name, and revision-level. Therefore, if you need to have two bridges in the same region, the two bridges must have identical configurations, names, and revision-levels. Also, one or more VLANs can be mapped to one MSTP instance (IST or MSTI) but a VLAN cannot be mapped to multiple MSTP instances.

NOTE

One or more VLANs can be mapped to one MSTP instance (IST or MSTI) but a VLAN cannot be mapped to multiple MSTP instances.

Configuration notes

When configuring MSTP, note the following:

- With MSTP running, enabling static trunk on ports that are members of many VLANs (4000 or more VLANs) will keep the system busy for 20 to 25 seconds.

Configuring MSTP mode and scope

With the introduction of MSTP, a system can be either under MSTP mode or not under MSTP mode. The default state is to **not** be under MSTP mode. MSTP configuration can only be performed in a system under MSTP mode.

With a system configured under MSTP mode, there is a concept called MSTP scope. MSTP scope defines the VLANs that are under direct MSTP control. You cannot run 802.1D or 802.1w on any VLAN (even outside of MSTP scope) and you cannot create topology groups when a system is under MSTP mode. While a VLAN group will still be supported when a system is under MSTP mode, the member VLAN should either be all in the MSTP scope or all out of the MSTP scope.

When a system is configured from non-MSTP mode to MSTP mode, the following changes are made to the system configuration:

- All 802.1D and 802.1w STP instances are deleted regardless of whether the VLAN is inside the MSTP scope or not
- All topology groups are deleted
- Any GVRP configuration is deleted
- Any VSRP configuration is deleted
- Single-span (if configured) is deleted
- MRP running on a VLAN inside MSTP scope is deleted
- The CIST is created and all VLANs inside the MSTP scope are attached with the CIST

Make sure that no physical layer-2 loops exist prior to switching from non-MSTP mode to MSTP mode. If, for example, you have an L2 loop topology configured as a redundancy mechanism before you perform the switch, a Layer 2 storm should be expected.

To configure a system into MSTP mode, use the following command at the Global Configuration level.

```
PowerConnect(config)#mstp scope all
```

Syntax: [no] mstp scope all

NOTE

MSTP is not operational however until the **mstp start** command is issued as described in [“Activating MSTP on a switch”](#) on page 167.

Once the system is configured into MSTP mode, CIST (sometimes referred to as “instance 0”) is created and all existing VLANs inside the MSTP scope are controlled by CIST. In addition, whenever you create a new VLAN inside MSTP scope, it is put under CIST control by default. In the Dell MSTP implementation however, a VLAN ID can be pre-mapped to another MSTI as described in [“Configuring an MSTP instance”](#) on page 165. A VLAN whose ID is pre-mapped, will attach to the specified MSTI instead of to the CIST when created.

NOTE

Once under MSTP mode, CIST always controls all ports in the system. If you do not want a port to run MSTP, configure the **no spanning-tree** command under the specified interface configuration.

Using the **[no]** option on a system that is configured for MSTP mode changes the system to non-MSTP mode. When this switch is made, all MSTP instances are deleted together with all MSTP configurations. ALL VLANs inside the original MSTP scope will not run any Layer-2 protocols after the switch.

Configuring additional MSTP parameters

To configure a switch for MSTP, you could configure the name and the revision on each switch that is being configured for MSTP. You must then create an MSTP Instance and assign an ID. VLANs are then assigned to MSTP instances. These instances must be configured on all switches that interoperate with the same VLAN assignments. Port cost, priority and global parameters can then be configured for individual ports and instances. In addition, operational edge ports and point-to-point links can be created and MSTP can be disabled on individual ports.

Each of the commands used to configure and operate MSTP are described in the following:

- “Setting the MSTP name”
- “Setting the MSTP revision number”
- “Configuring an MSTP instance”
- “Configuring bridge priority for an MSTP instance”
- “Setting the MSTP global parameters”
- “Setting ports to be operational edge ports”
- “Setting automatic operational edge ports”
- “Setting point-to-point link”
- “Disabling MSTP on a port”
- “Forcing ports to transmit an MSTP BPDU”
- “Activating MSTP on a switch”

Setting the MSTP name

Each switch that is running MSTP is configured with a name. It applies to the switch which can have many different VLANs that can belong to many different MSTP regions.

To configure an MSTP name, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp name Dell
```

Syntax: **[no] mstp name** <name>

The **name** parameter defines an ASCII name for the MSTP configuration. The default name is for the name variable to be blank.

Setting the MSTP revision number

Each switch that is running MSTP is configured with a revision number. It applies to the switch which can have many different VLANs that can belong to many different MSTP regions.

To configure an MSTP revision number, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp revision 4
```

Syntax: [no] mstp revision <revision-number>

The **revision** parameter specifies the revision level for MSTP that you are configuring on the switch. It can be a number from 0 and 65535. The default revision number is 0.

Configuring an MSTP instance

An MSTP instance is configured with an MSTP ID for each region. Each region can contain one or more VLANs. The Dell implementation of MSTP allows you to assign VLANs or ranges of VLANs to an MSTP instance before or after they have been defined. If pre-defined, a VLAN will be placed in the MSTI that it was assigned to immediately when the VLAN is created. Otherwise, the default operation is to condition of assign all new VLANs to the CIST. VLANs assigned to the CIST by default can be moved later to a specified MSTI.

To configure an MSTP instance and map one or more VLANs to that MSTI, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp instance 7 vlan 4 to 7
```

Syntax: [no] mstp instance <instance-number> [vlan <vlan-id> | vlan-group <group-id>]

The **instance** parameter defines the number for the instance of MSTP that you are configuring. The value 0 (which identifies the CIST) cannot be used. You can have up to 15 instances, number 1 – 4094.

The **vlan** parameter assigns one or more VLANs or a range of VLANs to the instance defined in this command.

The **vlan-group** parameter assigns one or more VLAN groups to the instance defined in this command.

The **no** option moves a VLAN or VLAN group from its assigned MSTI back into the CIST.

NOTE

The system does not allow an MSTI without any VLANs mapped to it. Consequently, removing all VLANs from an MSTI, deletes the MSTI from the system. The CIST by contrast will exist regardless of whether or not any VLANs are assigned to it or not. Consequently, if all VLANs are moved out of a CIST, the CIST will still exist and functional.

Configuring bridge priority for an MSTP instance

Priority can be configured for a specified instance. To configure priority for an MSTP instance, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp instance 1 priority 8192
```

Syntax: [no] mstp instance <instance-number> priority <priority-value>

The <instance-number> variable is the number for the instance of MSTP that you are configuring.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 61440 in increments of 4096. The default value is 32768.

Setting the MSTP global parameters

MSTP has many of the options available in RSTP as well as some unique options. To configure MSTP Global parameters for all instances on a switch.

```
PowerConnect(config)#mstp force-version 0 forward-delay 10 hello-time 4 max-age
12 max-hops 9
```

Syntax: [no] mstp force-version <mode-number> forward-delay <value> hello-time <value>
max-age <value> max-hops <value>

The **force-version** parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following <mode-number> values:

- 0 – The STP compatibility mode. Only STP BPDUs will be sent. This is equivalent to single STP.
- 2 – The RSTP compatibility mode. Only RSTP BPDUS will be sent. This is equivalent to single STP.
- 3 – MSTP mode. In this default mode, only MSTP BPDUS will be sent.

The **forward-delay** <value> specifies how long a port waits before it forwards an RST BPDU after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. The parameter can have a value from 1 – 10 seconds. The default is 2 seconds.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value from 6 – 40 seconds, where the value adheres to the following formula.

max age equal to or greater than 2 x (hello-time + 1) AND max age equal to or greater than 2 x (forward-delay – 1)

The default max-age is 20 seconds.

The **max-hops** <value> parameter specifies the maximum hop count. You can specify a value from 1 – 40 hops. The default value is 20 hops.

Setting ports to be operational edge ports

You can define specific ports as edge ports for the region in which they are configured to connect to devices (such as a host) that are not running STP, RSTP, or MSTP. If a port is connected to an end device such as a PC, the port can be configured as an edge port. To configure ports as operational edge ports enter a command such as the following.

```
PowerConnect(config)#mstp admin-edge-port ethernet 1
```

Syntax: [no] mstp admin-edge-port ethernet <portnum>

The <portnum> parameter specifies a port or range of ports as edge ports in the instance in which they are configured.

Setting automatic operational edge ports

You can configure a Layer 3 switch to automatically set a port as an operational edge port if the port does not receive any BPDUs since link-up. If the port receives a BPDU later, it is automatically reset to become an operational non-edge port. This feature is set globally to apply to all ports on a router where it is configured. This feature is configured as shown in the following.

```
PowerConnect(config)#mstp edge-port-auto-detect
```

Syntax: [no] mstp edge-port-auto-detect

NOTE

If this feature is enabled, it takes the port about 3 seconds longer to come to the enable state.

Setting point-to-point link

You can set a point-to-point link between ports to increase the speed of convergence. To create a point-to-point link between ports, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 5
```

Syntax: [no] mstp admin-pt2pt-mac ethernet <portnum>

The <portnum> parameter specifies a port or range of ports as edge ports in the instance in which they are configured.

Disabling MSTP on a port

To disable MSTP on a specific port, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp disable ethernet 1
```

Syntax: [no] mstp disable ethernet <portnum>

The <portnum> variable specifies the location of the port for which you want to disable MSTP.

NOTE

When a port is disabled for MSTP, it behaves as blocking for all the VLAN traffic that is controlled by MSTIs and the CIST.

Forcing ports to transmit an MSTP BPDU

To force a port to transmit an MSTP BPDU, use a command such as the following at the Global Configuration level.

```
PowerConnect(config)#mstp force-migration-check ethernet 1
```

Syntax: [no] mstp force-migration-check ethernet <portnum>

The <portnum> variable specifies the port or ports from which you want to transmit an MSTP BPDU.

Activating MSTP on a switch

MSTP scope must be enabled on the switch as described in [“Configuring MSTP mode and scope”](#) on page 163 before MSTP can be enabled.

To enable MSTP on your switch, use the following at the Global Configuration level.

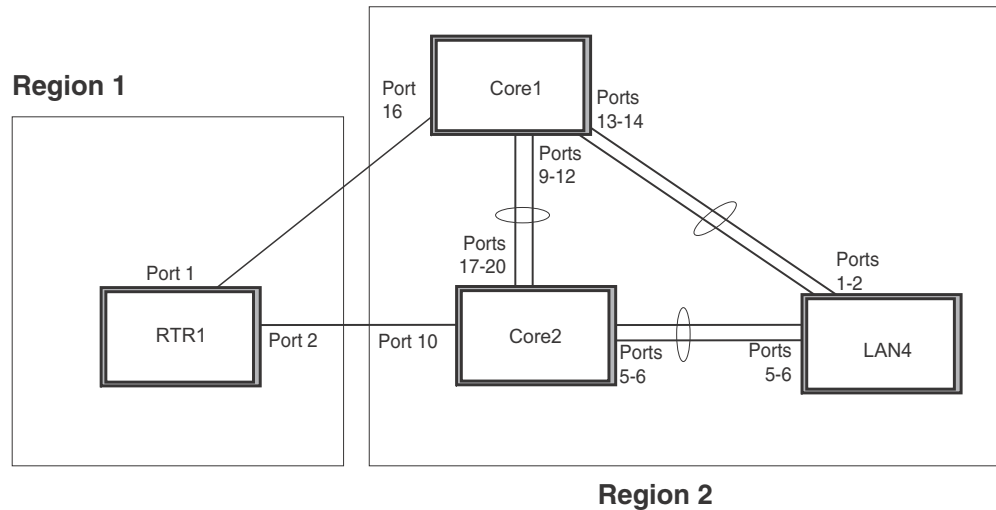
```
PowerConnect(config)#mstp start
```

Syntax: [no] mstp start

The [no] option disables MSTP from operating on a switch.

Example

In Figure 31 four device routers are configured in two regions. There are four VLANs in four instances in Region 2. Region 1 is in the CIST.

FIGURE 31 Sample MSTP configuration**RTR1 configuration**

```
PowerConnect(config-vlan-4093)#tagged ethernet 1 to 2
PowerConnect(config-vlan-4093)#exit
PowerConnect(config)#mstp scope all
PowerConnect(config)#mstp name Reg1
PowerConnect(config)#mstp revision 1
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 1 to 2
PowerConnect(config)#mstp start
PowerConnect(config)#hostname RTR1
```

Core 1 configuration

```
PowerConnect(config)#trunk ethernet 9 to 12 ethernet 13 to 14
PowerConnect(config-vlan-1)#name DEFAULT-VLAN by port
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#vlan 20 by port
PowerConnect(config-vlan-20)#tagged ethernet 9 to 14 ethernet 16
PowerConnect(config-vlan-20)#exit
PowerConnect(config)#vlan 21 by port
PowerConnect(config-vlan-21)#tagged ethernet 9 to 14 ethernet 16
PowerConnect(config-vlan-21)#exit
PowerConnect(config)#vlan 22 by port
PowerConnect(config-vlan-22)#tagged ethernet 9 to 14 ethernet 16
PowerConnect(config-vlan-22)#exit
PowerConnect(config)#vlan 23 by port
PowerConnect(config)#mstp scope all
PowerConnect(config)#mstp name HR
PowerConnect(config)#mstp revision 2
PowerConnect(config)#mstp instance 20 vlan 20
PowerConnect(config)#mstp instance 21 vlan 21
PowerConnect(config)#mstp instance 22 vlan 22
PowerConnect(config)#mstp instance 0 priority 8192
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 9 to 14
```

```
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 16
PowerConnect(config)#mstp disable ethernet 240
PowerConnect(config)#mstp start
PowerConnect(config)#hostname CORE1
```

Core2 configuration

```
PowerConnect(config)#trunk ethernet 5 to 6 ethernet 17 to 20
PowerConnect(config)#vlan 1 name DEFAULT-VLAN by port
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#vlan 20 by port
PowerConnect(config-vlan-20)#tagged ethernet 5 to 6 ethernet 17 to 20
PowerConnect(config-vlan-20)#exit
PowerConnect(config)#vlan 21 by port
PowerConnect(config-vlan-21)#tagged ethernet 5 to 6 ethernet 17 to 20
PowerConnect(config-vlan-21)#exit
PowerConnect(config)#vlan 22 by port
PowerConnect(config-vlan-22)#tagged ethernet 5 to 6 ethernet 17 to 20
PowerConnect(config-vlan-22)#exit
PowerConnect(config)#mstp scope all
PowerConnect(config)#mstp name HR
PowerConnect(config)#mstp revision 2
PowerConnect(config)#mstp instance 20 vlan 20
PowerConnect(config)#mstp instance 21 vlan 21
PowerConnect(config)#mstp instance 22 vlan 22
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 17 to 20 ethernet 5 to 6
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 10
PowerConnect(config)#mstp disable ethernet 7 ethernet 24
PowerConnect(config)#mstp start
PowerConnect(config)#hostname CORE2
```

LAN 4 configuration

```
PowerConnect(config)#trunk ethernet 5 to 6 ethernet 1 to 2
PowerConnect(config)#vlan 1 name DEFAULT-VLAN by port
PowerConnect(config-vlan-1)#exit
PowerConnect(config)#vlan 20 by port
PowerConnect(config-vlan-20)#tagged ethernet 1 to 2 ethernet 5 to 6
PowerConnect(config-vlan-20)#exit
PowerConnect(config)#vlan 21 by port
PowerConnect(config-vlan-21)#tagged ethernet 1 to 2 ethernet 5 to 6
PowerConnect(config-vlan-21)#exit
PowerConnect(config)#vlan 22 by port
PowerConnect(config-vlan-22)#tagged ethernet 1 to 2 ethernet 5 to 6
PowerConnect(config-vlan-22)#exit
PowerConnect(config)#mstp scope all
PowerConnect(config)#mstp config name HR
PowerConnect(config)#mstp revision 2
PowerConnect(config)#mstp instance 20 vlan 20
PowerConnect(config)#mstp instance 21 vlan 21
PowerConnect(config)#mstp instance 22 vlan 22
PowerConnect(config)#mstp admin-pt2pt-mac ethernet 5 to 6 ethernet 1 to 2
PowerConnect(config)#mstp start
PowerConnect(config)#hostname LAN4
```

Displaying MSTP statistics

MSTP statistics can be displayed using the commands shown below.

To display all general MSTP information, enter the following command.

6 802.1s Multiple Spanning Tree Protocol

```
PowerConnect#show mstp
MSTP Instance 0 (CIST) - VLANs: 1
-----
Bridge Identifier      Bridge MaxAge Hello FwdDly Hop Root MaxAge Hello FwdDly Hop
hex                   sec    sec    sec    cnt    sec    sec    sec    cnt
8000000cdb80af01 20     2     15    20    20    2     15    19

Root Bridge           ExtPath Cost RegionalRoot Bridge IntPath Cost Designated Bridge Root Port
hex                   hex                   hex                   hex
8000000480bb9876 2000   8000000cdb80af01 0 8000000480bb9876 1

Port Pri PortPath P2P Edge Role State Designa- Designated
Num   Cost   Mac Port      State   ted cost  bridge
1 128 2000 T F ROOT FORWARDING 0 8000000480bb9876

MSTP Instance 1 - VLANs: 2
-----
Bridge Identifier      Max RegionalRoot IntPath Designated Root Root
Hop Bridge           Cost Bridge         Bridge         Port Hop
cnt hex              hex                   hex                   hex cnt
8001000cdb80af01 20 8001000cdb80af01 0 8001000cdb80af01 Root 20

Port Pri PortPath Role State Designa- Designated
Num   Cost   Mac      Role      State   ted cost  bridge
1 128 2000 MASTER FORWARDING 0 8001000cdb80af01
```

Syntax: `show mstp <instance-number>`

The *<instance-number>* variable specifies the MSTP instance that you want to display information for.

TABLE 30 Output from Show MSTP

This field...	Displays...
MSTP Instance	The ID of the MSTP instance whose statistics are being displayed. For the CIST, this number is 0.
VLANs	The number of VLANs that are included in this instance of MSTP. For the CIST this number will always be 1.
Bridge Identifier	The MAC address of the bridge.
Bridge MaxAge sec	Displays configured Max Age.
Bridge Hello sec	Displays configured Hello variable.
Bridge FwdDly sec	Displays configured FwdDly variable.
Bridge Hop cnt	Displays configured Max Hop count variable.
Root MaxAge sec	Max Age configured on the root bridge.
Root Hello sec	Hello interval configured on the root bridge.
Root FwdDly sec	FwdDly interval configured on the root bridge.
Root Hop Cnt	Current hop count from the root bridge.
Root Bridge	Bridge identifier of the root bridge.

TABLE 30 Output from Show MSTP (Continued)

This field...	Displays...
ExtPath Cost	The configured path cost on a link connected to this port to an external MSTP region.
Regional Root Bridge	The Regional Root Bridge is the MAC address of the Root Bridge for the local region.
IntPath Cost	The configured path cost on a link connected to this port within the internal MSTP region.
Designated Bridge	The MAC address of the bridge that sent the best BPDU that was received on this port.
Root Port	Port indicating shortest path to root. Set to "Root" if this bridge is the root bridge.
Port Num	The port number of the interface.
Pri	The configured priority of the port. The default is 128.
PortPath Cost	Configured or auto detected path cost for port.
P2P Mac	Indicates if the port is configured with a point-to-point link: <ul style="list-style-type: none"> • T – The port is configured in a point-to-point link • F – The port is not configured in a point-to-point link
Edge	Indicates if the port is configured as an operational edge port: <ul style="list-style-type: none"> • T – indicates that the port is defined as an edge port. • F – indicates that the port is not defined as an edge port
Role	The current role of the port: <ul style="list-style-type: none"> • Master • Root • Designated • Alternate • Backup • Disabled
State	The port current spanning tree state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled
Designated Cost	Port path cost to the root bridge.
Max Hop cnt	The maximum hop count configured for this instance.
Root Hop cnt	Hop count from the root bridge.

Displaying MSTP information for a specified instance

The following example displays MSTP information specified for an MSTP instance.

6 802.1s Multiple Spanning Tree Protocol

```
PowerConnect#show mstp 1
MSTP Instance 1 - VLANs: 2
-----
Bridge          Max RegionalRoot  IntPath  Designated  Root  Root
Identifier      Hop Bridge         Cost      Bridge      Port  Hop
hex             cnt hex            hex         hex         cnt
8001000cdb80af01 20 8001000cdb80af01 0           8001000cdb80af01 Root 20

Port  Pri PortPath  Role      State      Designa-  Designated
Num   Cost                State      ted cost  bridge
1    128 2000      MASTER   FORWARDING 0      8001000cdb80af01
```

Refer to [Table 30](#) for details about the display parameters.

Displaying MSTP information for CIST instance 0

Instance 0 is the Common and Internal Spanning Tree Instance (CIST). When you display information for this instance there are some differences with displaying other instances. The following example displays MSTP information for CIST Instance 0.

```
PowerConnect#show mstp 0
MSTP Instance 0 (CIST) - VLANs: 1
-----
Bridge          Bridge Bridge Bridge Bridge Root  Root  Root  Root
Identifier      MaxAge Hello FwdDly Hop  MaxAge Hello FwdDly Hop
hex             sec   sec   sec   cnt  sec   sec   sec   cnt
8000000cdb80af01 20    2    15   20   20    2    15   19

Root           ExtPath  RegionalRoot  IntPath  Designated  Root
Bridge         Cost      Bridge         Cost      Bridge      Port
hex            hex
8000000480bb9876 2000    8000000cdb80af01 0      8000000480bb9876 1

Port  Pri PortPath  P2P Edge Role      State      Designa-  Designated
Num   Cost  Mac Port  Role      State      ted cost  bridge
1    128 2000    T  F  ROOT      FORWARDING 0      8000000480bb9876
```

To display details about the MSTP configuration, enter the following command.

```
PowerConnect#show mstp conf
MSTP CONFIGURATION
-----
Name       : Reg1
Revision  : 1
Version   : 3 (MSTP mode)
Status    : Started

Instance  VLANs
-----
0         4093
```

To display details about the MSTP that is configured on the device, enter the following command.

```
PowerConnect#show mstp detail
MSTP Instance 0 (CIST) - VLANs: 4093
-----
Bridge: 800000b000c00000 [Priority 32768, SysId 0, Mac 00b000c00000]
FwdDelay 15, HelloTime 2, MaxHops 20, TxHoldCount 6
Port 54 - Role: DESIGNATED - State: FORWARDING
PathCost 20000, Priority 128, OperEdge T, OperPt2PtMac F, Boundary T
Designated - Root 800000b000c00000, RegionalRoot 800000b000c00000,
Bridge 800000b000c00000, ExtCost 0, IntCost 0
ActiveTimers - helloWhen 1
MachineState - PRX-DISCARD, PTX-IDLE, PPM-SENDING_RSTP, PIM-CURRENT
PRT-ACTIVE_PORT, PST-FORWARDING, TCM-INACTIVE
BPDUs - Rcvd MST 0, RST 0, Config 0, TCN 0
Sent MST 6, RST 0, Config 0, TCN 0
```

Refer to [Table 30](#) for explanation about the parameters in the output.

Syntax: `show mstp [<mstp-id> | configuration | detail] [| begin <string> | exclude <string> | include <string>]`

Enter an MSTP ID for <mstp-id>.

6 802.1s Multiple Spanning Tree Protocol

Configuring Basic Layer 2 Features

The procedures in this chapter describe how to configure basic Layer 2 parameters.

PowerConnect devices are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the advanced features such as VLANs or routing protocols for the device must first be enabled at the system (global) level before they can be configured. If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.

NOTES:

- Before assigning or modifying any router parameters, you must assign the IP subnet (interface) addresses for each port.
- For information about configuring IP addresses, DNS resolver, DHCP assist, and other IP-related parameters, refer to [Chapter 21, “Configuring IP”](#).
- For information about the Syslog buffer and messages, refer to [Chapter 34, “Using Syslog”](#).

Enabling or disabling the Spanning Tree Protocol (STP)

STP (IEEE 802.1D bridge protocol) is supported on all devices. STP detects and eliminates logical loops in the network. STP also ensures that the least cost path is taken when multiple paths exist between ports or VLANs. If the selected path fails, STP searches for and then establishes an alternate path to prevent or limit retransmission of data.

NOTE

This section provides instructions for enabling and disabling STP. For configuration procedures and more information about STP, refer to [Chapter 6, “Configuring Spanning Tree Protocol \(STP\) Related Features”](#) in this guide.

STP must be enabled at the system level to allow assignment of this capability on the VLAN level. On devices running Layer 2 code, STP is enabled by default. On devices running Layer 3 code, STP is disabled by default.

To enable STP for all ports on a device, enter the following command.

```
PowerConnect(config)#spanning tree
```

Syntax: [no] spanning-tree

You can also enable and disable spanning tree on a port-based VLAN and on an individual port basis, and enable advanced STP features. Refer to [Chapter 6, “Configuring Spanning Tree Protocol \(STP\) Related Features”](#).

Modifying STP bridge and port parameters

You can modify the following STP Parameters:

- Bridge parameters – forward delay, maximum age, hello time, and priority
- Port parameters – priority and path cost

For configuration details, refer to “[Changing STP bridge and port parameters](#)” on page 96.

Changing the MAC age time and disabling MAC address learning

To change the MAC address age timer, enter a command such as the following.

```
PowerConnect(config)#mac-age-time 60
```

Syntax: [no] **mac-age-time** <secs>

<secs> specifies the number of seconds. Possible values differ depending on the version of software running on your device, as follows:

- On PowerConnect B-Series TI24X devices, you can configure 0 or a value from 10 – 86,400 (seconds), in 1-second intervals. If you set the MAC age time to 0, aging is disabled.

NOTES: Usually, the actual MAC age time is from one to two times the configured value. For example, if you set the MAC age timer to 60 seconds, learned MAC entries age out after remaining unused for between 60 – 120 seconds. However, if all of the following conditions are met, then the MAC entries age out after a longer than expected duration:

- The MAC age timer is greater than 630 seconds.
- The number of MAC entries is over 6000.
- All MAC entries are learned from the same packet processor.
- All MAC entries age out at the same time.

Disabling the automatic learning of MAC addresses

By default, when a packet with an unknown Source MAC address is received on a port, the device learns this MAC address on the port.

You can prevent a physical port from learning MAC addresses by entering the following command.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#mac-learn-disable
```

Syntax: [no] **mac-learn disable**

Use the no form of the command to allow a physical port to learn MAC addresses.

Configuration notes and feature limitations

- This command is not available on virtual routing interfaces. Also, if this command is configured on the primary port of a trunk, MAC address learning will be disabled on all the ports in the trunk.
- Entering the **mac-learn-disable** command on tagged ports disables MAC learning for that port in all VLANs to which that port is a member. For example, if tagged port 1 is a member of VLAN 10, 20, and 30 and you issue the **mac-learn-disable** command on port 1, port 1 will not learn MAC addresses, even if it is a member of VLAN 10, 20, and 30.

Displaying the MAC address table

To display the MAC table, enter the following command.

```
PowerConnect#show mac-address
Total active entries from all ports = 3
Total static entries from all ports = 1
  MAC-Address      Port      Type      VLAN
1234.1234.1234     15      Static      1
0004.8038.2f24     14      Dynamic     1
0004.8038.2f00     13      Dynamic     1
0010.5a86.b159     10      Dynamic     1
```

In the output of the **show mac-address** command, the *Type* column indicates whether the MAC entry is static or dynamic. A static entry is one you create using the **static-mac-address** command. A dynamic entry is one that is learned by the software from network traffic.

The output of the **show mac-address** command include an *Index* column which indicates the index where the entry exists in the hardware MAC table.

NOTE

The **show mac-address** command output does not include MAC addresses for management ports, since these ports do not support typical MAC learning and MAC-based forwarding.

Configuring static MAC entries

Static MAC addresses can be assigned to devices.

NOTE

PowerConnect devices running Layer 3 code also support the assignment of static IP Routes, static ARP, and static RARP entries. For details on configuring these types of static entries, refer to [“Configuring static routes”](#) on page 596 and [“Creating static ARP entries”](#) on page 590.

NOTE

PowerConnect B-Series TI24X devices support static MAC entries with unicast addresses only.

You can manually input the MAC address of a device to prevent it from being aged out of the system address table.

This option can be used to prevent traffic for a specific device, such as a server, from flooding the network with traffic when it is down. Additionally, the static MAC address entry is used to assign higher priorities to specific MAC addresses.

You can specify traffic priority (QoS) and VLAN membership (VLAN ID) for the MAC Address as well as specify the device type of either router or host.

The default and maximum configurable MAC table sizes can differ depending on the device. To determine the default and maximum MAC table sizes for your device, display the system parameter values. Refer to [“Displaying and modifying system parameter default settings”](#) on page 184.

Multi-port static MAC address

Many applications, such as Microsoft NLB, Juniper IPS, and Netscreen Firewall, use the same MAC address to announce load-balancing services. As a result, a switch must be able to learn the same MAC address on several ports. Multi-port static MAC allows you to statically configure a MAC address on multiple ports using a single command.

Configuration notes

- On PowerConnect B-Series T124X devices, this feature can be used to configure *unicast* MAC addresses on one or more ports; Multicast addresses are not supported.
- PowerConnect B-Series T124X devices support a maximum of 15 multi-port static MAC addresses.
- Hosts or physical interfaces normally join multicast groups dynamically, but you can also statically configure a host or an interface to join a multicast group.
- The following limitations apply to PowerConnect B-Series T124X devices have the **source_port_group_suppression_enable** flag enabled:
 - This feature is not supported on ports that are already deployed on a trunk. If attempted, the system will return a **port overlap** error.
 - External trunks cannot be created and deployed on ports on which static MAC addresses are configured.
 - This feature is not supported on a set of ports that overlap with a set of ports on which static MAC entries are configured, unless the port list in both sets is identical

NOTE

For more information about **source port group suppression**, refer to [“Configuring VLAN-based static MAC entries”](#) on page 179.

Configuring a multi-port static MAC address

For example, to add a static entry for a server with a MAC address of 0045.5563.67ff and a priority of 7, enter the following command.

```
PowerConnect(config)#static-mac-address 0045.5563.67ff ethernet 2 ethernet 3
ethernet 4 priority 7
```

To specify a range of ports, enter the following command.

```
PowerConnect(config)#static-mac-address 0045.5563.67ff ethernet 2 to 6 priority 7
```

Syntax: **[no] static-mac-address** *<mac-addr>* **ethernet** *<portnum>* **ethernet** *<portnum>* **ethernet** *<portnum>* **[priority** *<num>* **]**

or

Syntax: **[no] static-mac-address** *<mac-addr>* **ethernet***<portnum>* **to ethernet** *<portnum>* **[priority** *<num>* **]**

The *<portnum>* parameter is a valid port number.

The priority *<num>* is optional and can be a value from 0 – 7 (0 is lowest priority and 7 is highest priority). The default priority is 0.

NOTE

The location of the **static-mac-address** command in the CLI depends on whether you configure port-based VLANs on the device. If the device does not have more than one port-based VLAN (VLAN 1, which is the default VLAN that contains all the ports), the **static-mac-address** command is at the global CONFIG level of the CLI. If the device has more than one port-based VLAN, then the **static-mac-address** command is not available at the global CONFIG level. In this case, the command is available at the configuration level for each port-based VLAN.

Configuring VLAN-based static MAC entries

You can configure a VLAN to drop packets that have a particular source or destination MAC address.

You can configure a maximum of 2048 static MAC address drop entries on a device.

Use the CLI command **show running-config** to view the static MAC address drop entries currently configured on the device.

Command syntax

To configure a VLAN to drop packets with a source or destination MAC address of 1145.5563.67FF, enter the following commands.

```
PowerConnect(config)#vlan 2
PowerConnect(config-vlan-2)#static-mac-address 1145.5563.67FF drop
```

Syntax: [no] static-mac-address <mac-addr> drop

Use the **no** form of the command to remove the static MAC address drop configuration.

Clearing MAC address entries

You can remove learned MAC address entries from the MAC address table. The types of MAC address can be removed are as follows:

- All MAC address entries
- All MAC address entries for a specified Ethernet port
- All MAC address entries for a specified VLAN
- All specified MAC address entry in all VLANs

For example, to remove entries for the MAC address 000d.cd80.00d0 in all VLANs, enter the following command at the Privilege EXEC level of the CLI.

```
PowerConnect#clear mac-address 000d.cb80.00d0
```

Syntax: clear mac-address <mac-address> | ethernet <port-num> | vlan <vlan-num>

If you enter **clear mac-address** without any parameter, the software removes all MAC address entries.

Use the <mac-address> parameter to remove a specific MAC address from all VLANs. Specify the MAC address in the following format: HHHH.HHHH.HHHH.

Use the **ethernet** <port-num> parameter to remove all MAC addresses for a specific Ethernet port.

Use the **vlan <num>** parameter to remove all MAC addresses for a specific VLAN.

Enabling port-based VLANs

When using the CLI, port and protocol-based VLANs are created by entering one of the following commands at the global CONFIG level of the CLI.

To create a port-based VLAN, enter commands such as the following.

```
PowerConnect(config)#vlan 222 by port  
PowerConnect(config)#vlan 222 name Mktg
```

Syntax: **vlan <num> by port**

Syntax: **vlan <num> name <string>**

The **<num>** parameter specifies the VLAN ID. The valid range for VLAN IDs starts at 1 on all systems but the upper limit of the range differs depending on the device. In addition, you can change the upper limit on some devices using the **system max-vlans...** command.

The **<string>** parameter is the VLAN name and can be a string up to 32 characters. You can use blank spaces in the name if you enclose the name in double quotes (for example, "Product Marketing".)

You can configure up to 4063 port-based VLANs on a device running Layer 2 code or 4061 port-based VLANs on a device running Layer 3 code. Each port-based VLAN can contain either tagged or untagged ports. A port cannot be a member of more than one port-based VLAN unless the port is tagged. On both device types, valid VLAN IDs are 1 – 4095. You can configure up to the maximum number of VLANs within that ID range.

NOTE

If you want to use VLANs 4091 and 4092 as configurable VLANs, you can assign them to different VLAN IDs. For more information, refer to ["Assigning different VLAN IDs to reserved VLANs 4091 and 4092"](#) on page 265.

NOTE

The second command is optional and also creates the VLAN if the VLAN does not already exist. You can enter the first command after you enter the second command if you first exit to the global CONFIG level of the CLI.

Assigning IEEE 802.1Q tagging to a port

When a port is tagged, it allows communication among the different VLANs to which it is assigned. A common use for this might be to place an email server that multiple groups may need access to on a tagged port, which in turn, is resident in all VLANs that need access to the server.

NOTE

Tagging does not apply to the default VLAN.

When using the CLI, ports are defined as either tagged or untagged at the VLAN level.

Command syntax

Suppose you want to make port 5 a member of port-based VLAN 4, a tagged port. To do so, enter the following.

```
PowerConnect(config)#vlan 4
PowerConnect(config-vlan-4)#tagged e 5
```

Syntax: `tagged ethernet <portnum> [to <portnum> [ethernet <portnum>...]]`

Defining MAC address filters

MAC layer filtering enables you to build access lists based on MAC layer headers in the Ethernet/IEEE 802.3 frame. You can filter on the source and destination MAC addresses. The filters apply to incoming traffic only.

You configure MAC filters globally, then apply them to individual interfaces. To apply MAC filters to an interface, you add the filters to that interface MAC filter group.

The device takes the action associated with the first matching filter. If the packet does not match any of the filters in the access list, the default action is to drop the packet. If you want the system to permit traffic by default, you must specifically indicate this by making the last entry in the access list a permit filter. An example is given below.

Syntax: `mac filter <last-index-number> permit any any.`

For devices running Layer 3 code, the MAC filter is applied to all inbound Ethernet packets, including routed traffic. This includes those port associated with a virtual routing interface. However, the filter is not applied to the virtual routing interface. It is applied to the physical port.

When you create a MAC filter, it takes effect immediately. You do not need to reset the system. However, you do need to save the configuration to flash memory to retain the filters across system resets.

Configuration notes and limitations

- MAC filters that have a global **deny** statement can cause the device to block all BPDUs. In this case, include exception statements for control protocols in the MAC filter configuration. On PowerConnect B-Series TI24X devices, BPDUs are not blocked by MAC filters.

The following configuration notes apply to Layer 3 devices:

- MAC filters apply to both switched and routed traffic. If a routing protocol (for example, OSPF) is configured on an interface, the configuration must include a MAC filter rule that allows the routing protocol MAC and the neighbor system MAC address.
- You cannot use Layer 2 filters to filter Layer 4 information.

Command syntax

To configure and apply a MAC filter, enter commands such as the following.

```
PowerConnect(config)#mac filter 1 deny 3565.3475.3676 ffff.0000.0000
PowerConnect(config)#mac filter 1024 permit any any
PowerConnect(config)#int e 1
PowerConnect(config-if-e10000-1)#mac filter-group 1 1024
```

These commands configure a filter to deny traffic with a source MAC address that begins with “3565” to any destination. The second filter permits all traffic that is not denied by another filter.

NOTE

Once you apply a MAC filter to a port, the device drops all Ethernet traffic on the port that does not match a MAC permit filter on the port.

Syntax: [no] mac filter <filter-num> permit | deny <src-mac> <mask> | any <dest-mac> <mask> | any

The **permit | deny** argument determines the action the software takes when a match occurs.

The <src-mac> <mask> | **any** parameter specifies the source MAC address. You can enter a specific address value and a comparison mask or the keyword **any** to filter on all MAC addresses. Specify the mask using (ones) and zeros. For example, to match on the first two bytes of the address aabb.ccdd.eeff, use the mask ffff.0000.0000. In this case, the filter matches on all MAC addresses that contain "aabb" as the first two bytes. The filter accepts any value for the remaining bytes of the MAC address. If you specify **any**, do not specify a mask. In this case, the filter matches on all MAC addresses.

The <dest-mac> <mask> | **any** parameter specifies the destination MAC address. The syntax rules are the same as those for the <src-mac> <mask> | **any** parameter.

Syntax: [no] mac filter log-enable

Globally enables logging for filtered packets.

Syntax: [no] mac filter-group log-enable

Enables logging for filtered packets on a specific port.

Syntax: [no] mac filter-group <filter-list>

Applies MAC filters to a port.

On PowerConnect devices, you can filter packets based on the Ethernet type using the new **etype** optional keyword.

Syntax: [no] mac filter <filter-num> permit | deny <src-mac><mask> | any <dest-mac><mask> | any [etype <eq |gt |lt | neq> <type>]

eq - Matches packet with a given Ethernet type

gt - Matches packet with a greater Ethernet type

lt - Matches packet with a lower Ethernet type

neq - Matches packet not with a given Ethernet type

NOTE

The filters must be applied as a group. For example, if you want to apply four filters to an interface, they must all appear on the same command line.

NOTE

You cannot add or remove individual filters in the group. To add or remove a filter on an interface, apply the filter group again containing all the filters you want to apply to the port.

NOTE

If you apply a filter group to a port that already has a filter group applied, the older filter group is replaced by the new filter group.

When a MAC filter is applied to or removed from an interface, a Syslog message such as the following is generated.

```
SYSLOG: <14>Jan 1 00:00:00 10.44.9.11 MAC Filter applied to port 2 by tester from
telnet session (filter id=5 ).
SYSLOG: <14>Jan 1 00:00:00 10.44.9.11 MAC Filter removed from port 2 by tester
from telnet session (filter id=5 ).
```

The Syslog messages indicate that a MAC filter was applied to the specified port by the specified user during the specified session type. Session type can be Console, Telnet, SSH, SNMP, or others. The filter IDs that were added or removed are listed.

Enabling logging of management traffic permitted by MAC filters

You can configure the device to generate Syslog entries and SNMP traps for management traffic that is permitted by MAC filters. **Management traffic** applies to packets that are destined for the CPU, such as control packets. You can enable logging of permitted management traffic on a global basis or an individual port basis.

The first time an entry in a MAC filter permits a management packet and logging is enabled for that entry, the software generates a Syslog message and an SNMP trap. Messages for management packets permitted by MAC filters are at the warning level of the Syslog.

When the first Syslog entry for a management packet permitted by a MAC filter is generated, the software starts a five-minute timer. After this, the software sends Syslog messages every five minutes. The messages list the number of management packets permitted by each MAC filter during the previous five-minute interval. If a MAC filter does not permit any packets during the five-minute interval, the software does not generate a Syslog entry for that MAC filter.

NOTE

For a MAC filter to be eligible to generate a Syslog entry for permitted management packets, logging must be enabled for the filter. The Syslog contains entries only for the MAC filters that permit packets and have logging enabled.

When the software places the first entry in the log, the software also starts the five-minute timer for subsequent log entries. Thus, five minutes after the first log entry, the software generates another log entry and SNMP trap for permitted management packets.

Configuration notes

MAC filter logging is supported in the PowerConnect B-Series TI24X devices.

These devices support MAC filter logging of management traffic only.

Command syntax

To configure MAC filter logging globally, enter the following CLI commands at the global CONFIG level.

7 Displaying and modifying system parameter default settings

```
PowerConnect(config)#mac filter log-enable  
PowerConnect(config)#write memory
```

Syntax: [no] mac filter log-enable

To configure MAC filter logging for MAC filters applied to ports 1 and 3, enter the following CLI commands.

```
PowerConnect(config)#int ethernet 1  
PowerConnect(config-if-e10000-1)#mac filter-group log-enable  
PowerConnect(config-if-e10000-1)#int ethernet 3  
PowerConnect(config-if-e10000-3)#mac filter-group log-enable  
PowerConnect(config-if-e10000-3)#write memory
```

Syntax: [no] mac filter-group log-enable

Displaying and modifying system parameter default settings

PowerConnect devices have default table sizes for the system parameters shown in the following display outputs. The table sizes determine the maximum number of entries the tables can hold. You can adjust individual table sizes to accommodate your configuration needs.

The tables you can configure, as well as the default values and valid ranges for each table, differ depending on the device you are configuring. To display the adjustable tables on your device, use the `show default values` command. The following shows example outputs.

Configuration considerations

- Changing the table size for a parameter reconfigures the device memory. Whenever you reconfigure the memory on a device, you must save the change to the startup-config file, then reload the software to place the change into effect.
- Configurable tables and their defaults and maximum values differ on IPv4 devices versus IPv6-capable devices.

Displaying system parameter default values

To display the configurable tables and their defaults and maximum values, enter the `show default values` command at any level of the CLI.

The following shows an example output of the `show default values` command on a PowerConnect device.

```
PowerConnect# show default values  
  
          Minimum   Maximum   Current  
multicast-route      1024      2048      1024  
pim-mcache           1024      4096      4096  
igmp-max-group-addr 4096      8192      4096  
igmp-snoop-mcache   512       2048      512  
msdp-sa-cache        1024      4096      1024
```

[Table 31](#) defines the system parameters in the `show default values` command output.

TABLE 31 System parameters in show default values command

This system parameter...	Defines the maximum number of...
atalk-route	Appletalk routes
atalk-zone-port	Appletalk zones per port
atalk-zone-sys	Appletalk zones per system
hw-ip-mcast-ml	Multicast output interfaces (clients)
hw-ip-next-hop	IP next hops and routes, including unicast next hops and multicast route entries
hw-logical-interface	Hardware logical interface pairs (physical port and VLAN pairs)
hw-traffic-conditioner	Traffic policies
ip-arp	ARP entries
ip-cache	IP forwarding cache entries
ip-filter-port	IP ACL entries per port
ip-filter-sys	IP ACL entries per system
ip-qos-session	Layer 4 session table entries
ip-route	Learned IP routes
ip-static-arp	Static IP ARP entries
ip-static-route	Static IP routes
ip-subnet-port	IP subnets per port
ipx-forward-filter	IPX forward filter entries
ipx-rip-entry	IPX RIP entries
ipx-rip-filter	IPX RIP filter entries
ipx-sap-entry	IPX SAP entries
ipx-sap-filter	IPX SAP filter entries
l3-vlan	Layer 3 VLANs
mac	MAC entries
mac-filter-port	MAC filter entries per port
mac-filter-sys	MAC filter entries per system
multicast-route	Multicast routes
pim-mcache	PIM multicast cache entries
rmon-entries	RMON control table entries
session-limit	Session entries
spanning-tree	Spanning tree instances
view	SNMP views
virtual-interface	Virtual routing interfaces
vlan	VLANs
mld-max-group-addr	MLD group limit

TABLE 31 System parameters in show default values command (Continued)

This system parameter...	Defines the maximum number of...
igmp-snoop-mcache	IGMP snooping cache entries
mld-snoop-mcache	MLD snooping cache entries

Modifying system parameter default values

Information for the configurable tables appears under the columns that are shown in bold type in the above examples. To simplify configuration, the command parameter you enter to configure the table is used for the table name. For example, to increase the capacity of the IP route table, enter the following commands.

```
PowerConnect(config)#system-max ip-route 120000
PowerConnect(config)#write memory
PowerConnect(config)#exit
PowerConnect#reload
```

Syntax: **system-max ip-route** <num>

The <num> parameter specifies the maximum number of routes in the IP route table. The minimum value is 4096. The default is 80000 IP routes.

NOTE

If you accidentally enter a value that is not within the valid range of values, the CLI will display the valid range for you.

To increase the number of IP subnet interfaces you can configure on each port on a device running Layer 3 code from 24 to 64, then increase the total number of IP interfaces you can configure on the device from 256 to 512, enter the following commands.

```
PowerConnect(config)#system-max subnet-per-interface 64
PowerConnect(config)#write memory
PowerConnect(config)#exit
PowerConnect#reload
```

Syntax: **system-max subnet-per-interface** <num>

The <num> parameter specifies the maximum number of subnet addresses per port and can be from 1 – 64. The default is 24.

```
PowerConnect(config)#system-max subnet-per-system 512
PowerConnect(config)#write memory
PowerConnect(config)#exit
PowerConnect#reload
```

Syntax: **system-max subnet-per-system** <num>

The <num> parameter specifies the maximum number of subnet addresses for the entire device and can be from 1 – 512. The default is 256.

NOTE

If you increase the number of configurable subnet addresses on each port, you might also need to increase the total number of subnets that you can configure on the device.

Egress buffer thresholds for QoS priorities

NOTES:

- The terms *QoS priority* and *traffic class* are used interchangeably in this section and mean the same thing.
- Buffer threshold level-3 to maximum does not change the buffering behavior of the device.

The PowerConnect switch uses egress buffer threshold levels to dynamically adjust each port egress queue (outbound transmit queue) based on QoS priority (traffic class) and jumbo frame support. When the egress buffer queue is full, any new packets destined for the egress queue will be dropped.


The following egress buffer thresholds can be adjusted by an administrator:

- **Egress Buffer Threshold for all Traffic Classes for a port** – When this is configured, the specified egress buffer threshold will be applied to *all* traffic classes (0 – 7) on a port.
- **Egress Buffer Threshold for a given Traffic Class for a port** – When this is configured, the specified egress buffer threshold will be applied to a *specific* traffic class (0 – 7) on a port.

The egress buffer thresholds can be modified to various levels. The default settings are described in the section “[Default settings for egress buffer thresholds](#)” on page 188. Note that based on the egress buffer threshold assigned, the system will dynamically control the egress CoS queue size for a port based on the current buffer usage information from the system.

[Figure 32](#) shows the egress buffer thresholds.

FIGURE 32 Egress buffer thresholds on PowerConnect devices



Manually increasing buffer thresholds may be useful in situations where applications have intermittent bursts of oversubscription. For example, by increasing a port egress buffer threshold, the PowerConnect will be able to forward oversubscribed packets instead of dropping them.

Cut-Through Switching Support on PowerConnect B-Series TI24X Switches

The PowerConnect B-Series TI24X operates in cut-through switching mode, meaning it starts forwarding a frame even before the whole frame has been received. However, if there is any oversubscription on the egress port, either due to speed mismatch (10-Gbps to 1-Gbps) or network topology (traffic from two 10-Gig ports going out of one 10-Gig port), the switch will buffer the packets and the forwarding behavior will be similar to store-and-forward mode.

Cut-Through Switching Mode and CRC Error Packet Handling

In the cut-through switching mode, if a packet is detected with a CRC error, the whole packet will be dropped if the packet size is less than or equal to 384 bytes. If the packet size is larger than 384 bytes, the last few bytes will be trimmed off the frames.

Default settings for egress buffer thresholds

Figure 32 illustrates the egress buffer threshold levels. For both non-jumbo and jumbo frame forwarding, the default egress buffer threshold is level-4 for all QoS priorities.

The PowerConnect B-Series TI24X is automatically configured to use the default thresholds for egress buffers. If you want to change the default threshold values, you must first disable the default values. The next section shows how.

Disabling and re-enabling the default settings for egress buffer thresholds

To disable the device from using the default settings for egress buffer thresholds, enter the following command.

```
PowerConnect(config)#no enable egress-buffer-default
```

This command disables the default values for all traffic classes on all ports. Once disabled, you can configure new threshold values as instructed in [“Setting the egress buffer threshold for all QoS priorities on a port or group of ports”](#) on page 189 and [“Setting the egress buffer threshold for a specific QoS priority on a port or group of ports”](#) on page 190.

Syntax: no enable egress-buffer-default

To re-enable the default settings for egress buffer thresholds once they have been disabled, enter the following command.

```
PowerConnect(config)#enable egress-buffer-default
```

The PowerConnect B-Series T124X will revert to using the default values for all traffic classes on all ports, as described in [“Default settings for egress buffer thresholds”](#) on page 188.

Syntax: enable egress-buffer-default

Setting the egress buffer threshold for all QoS priorities on a port or group of ports

NOTE

Be sure to disable the default settings before performing the tasks in this section. Refer to [“Disabling and re-enabling the default settings for egress buffer thresholds”](#).

To set the egress buffer threshold for all QoS priorities (0 – 7) on a port, enter commands such as the following.

```
PowerConnect(config)#int e 3
PowerConnect(conf-if-e10000-3)#egress-buffer-threshold min
```

To set the egress buffer threshold for all QoS priorities on multiple ports, enter commands such as the following.

```
PowerConnect(config)#int e 4 to 5
PowerConnect(config-mif-e10000-4-5)#egress-buffer-threshold min
```

These commands set the egress buffer threshold for all QoS priorities on ports e 3, and ports e 4 and 5, to the minimum (**min**) threshold.

Syntax: [no] egress-buffer-threshold min|level-1|level-2|level-3|level-4|level-5|level-6|max

Use the **no** form of the command to revert back to the default values for all QoS priorities on the port. For the default values, refer to [“Default settings for egress buffer thresholds”](#) on page 188.

Setting the egress buffer threshold for a specific QoS priority on a port or group of ports

NOTE

Be sure to disable the default settings before performing the tasks in this section. Refer to “Disabling and re-enabling the default settings for egress buffer thresholds”.

To set the egress buffer threshold for a specific QoS priority on a port, enter commands such as the following.

```
PowerConnect(config)#int e 3
PowerConnect(config-if-e10000-3)#egress-buffer-threshold level-6 6
```

These commands set the egress buffer threshold for packets with QoS priority 6 to the **level-6** threshold.

To set the egress buffer threshold for a specific QoS priority on multiple ports, enter commands such as the following.

```
PowerConnect(config)#int e 4 to 5
PowerConnect(config-mif-4-5)#egress-buffer-threshold high 6
```

These commands set the egress buffer threshold for packets with QoS priority 6 to the **high** threshold.

Syntax: `[no] egress-buffer-threshold min|level-1|level-2|level-3|level-4|level-5|level-6|max <QoS-priority>`

For `<QoS-priority>`, enter a specific QoS priority (0 – 7).

Use the **no** form of the command to revert back to the default value for the QoS priority. For the default values, refer to “Default settings for egress buffer thresholds” on page 188.

Link Fault Signaling (LFS) for 10G

Link Fault Signaling (LFS) is a physical layer protocol that enables communication on a link between two 10 Gbps Ethernet devices. When configured on a 10 Gbps Ethernet port, the port can detect and report fault conditions on transmit and receive ports. Dell recommends enabling LFS on both ends of a link.

NOTE

LFS is always enabled on PowerConnect B-Series TI24X devices and cannot be disabled. However, as long as LFS is enabled on each end of the link regardless of the device type, a fault will be properly detected and the link not used.

When LFS is enabled on an interface, the following Syslog messages are generated when the link goes up or down, or when the TX or RX fiber is removed from one or both sides of the link that has LFS enabled.

```
Interface ethernet1, state down - link down
Interface ethernet1, state up
```

When a link fault occurs, the Link and Activity LEDs turn OFF.

The Link and Activity LEDs turn ON when there is traffic traversing the link after the fiber is installed.

Enabling LFS

To enable LFS between two 10 Gbps Ethernet devices, enter commands such as the following on both ends of the link.

```
PowerConnect(config)#interface e 1  
PowerConnect(config-if-e10000-1)#link-fault-signal
```

Syntax: link-fault-signal

LFS is OFF by default.

Jumbo frame support

Ethernet traffic moves in units called frames. The maximum size of frames is called the Maximum Transmission Unit (MTU). When a network device receives a frame larger than its MTU, the data is either fragmented or dropped. Historically, Ethernet has a maximum frame size of 1500 bytes, so most devices use 1500 as their default MTU.

Jumbo frames are Ethernet frames with more than 1,500 bytes MTU. Conventionally, jumbo frames can carry up to 9,000 bytes MTU. PowerConnect devices also support jumbo frames per VLAN.

7 Jumbo frame support

Configuring Metro Features

Topology groups

A topology group is a named set of VLANs that share a Layer 2 topology. Topology groups simplify configuration and enhance scalability of Layer 2 protocols by allowing you to run a single instance of a Layer 2 protocol on multiple VLANs.

You can use topology groups with the following Layer 2 protocols:

- STP
- MRP
- VSRP
- 802.1W

Topology groups simplify Layer 2 configuration and provide scalability by enabling you to use the same instance of a Layer 2 protocol for multiple VLANs. For example, if a device is deployed in a Metro network and provides forwarding for two MRP rings that each contain 128 VLANs, you can configure a topology group for each ring. If a link failure in a ring causes a topology change, the change is applied to all the VLANs in the ring topology group. Without topology groups, you would need to configure a separate ring for each VLAN.

NOTE

If you plan to use a configuration saved under an earlier software release and the configuration contains STP groups, the CLI converts the STP groups into topology groups when you save the configuration. For backward compatibility, you can still use the STP group commands. However, the CLI converts the commands into the topology group syntax. Likewise, the **show stp-group** command displays STP topology groups.

Master VLAN and member VLANs

Each topology group contains a master VLAN and can contain one or more member VLANs and VLAN groups:

- **Master VLAN** – The master VLAN contains the configuration information for the Layer 2 protocol. For example, if you plan to use the topology group for MRP, the topology group master VLAN contains the ring configuration information.
- **Member VLANs** – The member VLANs are additional VLANs that share ports with the master VLAN. The Layer 2 protocol settings for the ports in the master VLAN apply to the same ports in the member VLANs. A change to the master VLAN Layer 2 protocol configuration or Layer 2 topology affects all the member VLANs. Member VLANs do not independently run a Layer 2 protocol.
- **Member VLAN groups** – A VLAN group is a named set of VLANs. The VLANs within a VLAN group have the same ports and use the same values for other VLAN parameters.

When a Layer 2 topology change occurs on a port in the master VLAN, the same change is applied to that port in all the member VLANs that contain the port. For example, if you configure a topology group whose master VLAN contains ports 1 and 2, a Layer 2 state change on port 1 applies to port 1 in all the member VLANs that contain that port. However, the state change does not affect port 1 in VLANs that are not members of the topology group.

Control ports and free ports

A port that is in a topology group can be a control port or a free port:

- **Control port** – A control port is a port in the master VLAN, and is therefore controlled by the Layer 2 protocol configured in the master VLAN. The same port in all the member VLANs is controlled by the master VLAN Layer 2 protocol. Each member VLAN must contain all of the control ports and can contain additional ports.
- **Free port** – A free port is not controlled by the master VLAN Layer 2 protocol. The master VLAN can contain free ports. (In this case, the Layer 2 protocol is disabled on those ports.) In addition, any ports in the member VLANs that are not also in the master VLAN are free ports.

NOTE

Since free ports are not controlled by the master port Layer 2 protocol, they are assumed to always be in the Forwarding state.

Configuration considerations

- Topology groups are supported on PowerConnect.
- You must configure the master VLAN and member VLANs or member VLAN groups before you configure the topology group.
- You can configure up to 256 topology groups. Each group can control up to 4096 VLANs. A VLAN cannot be controlled by more than one topology group.
- The topology group must contain a master VLAN and can also contain individual member VLANs, VLAN groups, or a combination of individual member VLANs and VLAN groups.
- If you add a new master VLAN to a topology group that already has a master VLAN, the new master VLAN replaces the older master VLAN. All member VLANs and VLAN groups follow the Layer 2 protocol settings of the new master VLAN.
- If you remove the master VLAN (by entering `no master-vlan <vlan-id>`), the software selects the new master VLAN from member VLANs. A new candidate master-vlan will be in configured order to a member VLAN so that the first added member VLAN will be a new candidate master VLAN. Once you saved and reloaded, a member VLAN with the newest VLAN ID will be the new candidate master. The new master VLAN inherits the Layer 2 protocol settings of the older master VLAN.
- Once you add a VLAN as a member of a topology group, all the Layer 2 protocol information on the VLAN is deleted.

Configuring a topology group

To configure a topology group, enter commands such as the following.


```
PowerConnect(config)# topology-group 2
PowerConnect(config-topo-group-2)# master-vlan 2
PowerConnect(config-topo-group-2)# member-vlan 3
PowerConnect(config-topo-group-2)# member-vlan 4
PowerConnect(config-topo-group-2)# member-vlan 5
PowerConnect(config-topo-group-2)# member-group 2
```

These commands create topology group 2 and add the following:

- Master VLAN 2
- Member VLANs 2, 3, and 4
- Member VLAN group 2

Syntax: **[no] topology-group** <group-id>

The <group-id> parameter specifies the topology group ID and can be from 1 – 256.

Syntax: **[no] master-vlan** <vlan-id>

This command adds the master VLAN. The VLAN must already be configured. Make sure all the Layer 2 protocol settings in the VLAN are correct for your configuration before you add the VLAN to the topology group. A topology group can have only one master VLAN.

NOTE

If you remove the master VLAN (by entering **no master-vlan** <vlan-id>), the software selects the new master VLAN from member VLANs. For example, if you remove master VLAN 2 from the example above, the CLI converts member VLAN 3 into the new master VLAN. The new master VLAN inherits the Layer 2 protocol settings of the older master VLAN.

NOTE

If you add a new master VLAN to a topology group that already has a master VLAN, the new master VLAN replaces the older master VLAN. All member VLANs and VLAN groups follow the Layer 2 protocol settings of the new master VLAN.

Syntax: **[no] member-vlan** <vlan-id>

The <vlan-id> parameter specifies a VLAN ID. The VLAN must already be configured.

Syntax: **[no] member-group** <num>

The <num> specifies a VLAN group ID. The VLAN group must already be configured.

NOTE

Once you add a VLAN or VLAN group as a member of a topology group, all the Layer 2 protocol configuration information for the VLAN or group is deleted. For example, if STP is configured on a VLAN and you add the VLAN to a topology group, the STP configuration is removed from the VLAN. Once you add the VLAN to a topology group, the VLAN uses the Layer 2 protocol settings of the master VLAN.

If you remove a member VLAN or VLAN group from a topology group, you will need to reconfigure the Layer 2 protocol information in the VLAN or VLAN group.

Displaying topology group information

The following sections show how to display STP information and topology group information for VLANs.

Displaying STP information

To display STP information for a VLAN, enter a command such as the following.

```
PowerConnect# show span vlan 4
VLAN 4 BPDU cam_index is 14344 and the Master DMA Are(HEX) 18 1A
STP instance owned by VLAN 2
```

This example shows STP information for VLAN 4. The line shown in bold type indicates that the VLAN STP configuration is controlled by VLAN 2. This information indicates that VLAN 4 is a member of a topology group and VLAN 2 is the master VLAN in that topology group.

Displaying topology group information

To display topology group information, enter the following command.

```
PowerConnect# show topology-group

Topology Group 3
=====
  master-vlan 2
  member-vlan none

Common control ports          L2 protocol
ethernet 1                    MRP
ethernet 2                    MRP
ethernet 5                    VSRP
ethernet 22                   VSRP
Per vlan free ports
ethernet 3                    Vlan 2
ethernet 4                    Vlan 2
ethernet 11                   Vlan 2
ethernet 12                   Vlan 2
```

Syntax: `show topology-group [<group-id>]`

This display shows the following information.

TABLE 32 CLI display of topology group information

This field...	Displays...
master-vlan	The master VLAN for the topology group. The settings for STP, MRP, or VSRP on the control ports in the master VLAN apply to all control ports in the member VLANs within the topology group.
member-vlan	The member VLANs in the topology group.
Common control ports	The master VLAN ports that are configured with Layer 2 protocol information. The Layer 2 protocol configuration and state of these ports in the master VLAN applies to the same port numbers in all the member VLANs.
L2 protocol	The Layer 2 protocol configured on the control ports. The Layer 2 protocol can be one of the following: <ul style="list-style-type: none"> • MRP • STP • VSRP
Per vlan free ports	The ports that are not controlled by the Layer 2 protocol information in the master VLAN.

Metro Ring Protocol (MRP)

Metro Ring Protocol (MRP) was introduced in two phases:

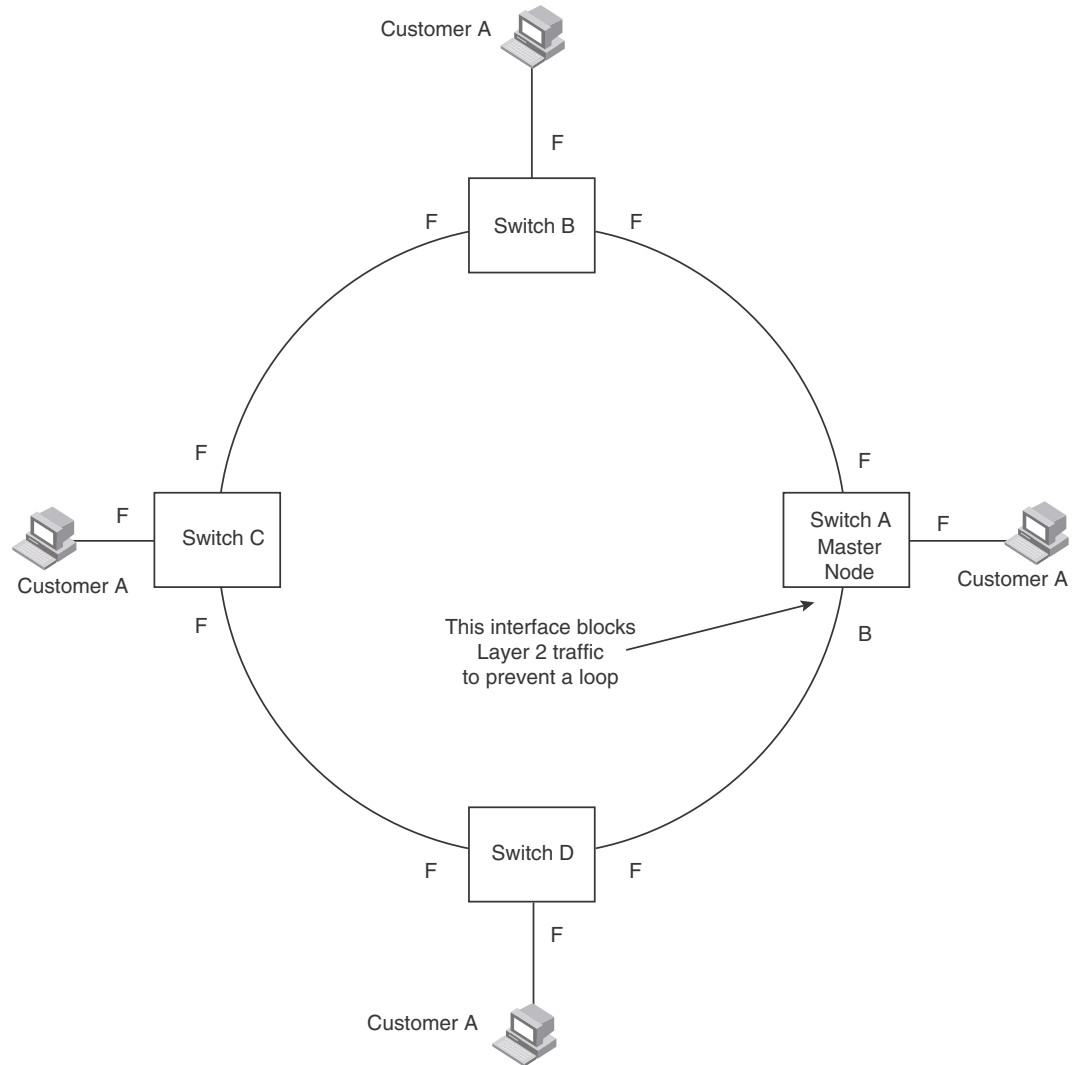
- **MRP Phase 1** is supported in all PowerConnect B-Series TI24X devices. Refer to [“MRP rings without shared interfaces \(MRP Phase 1\)”](#) on page 199.
- **MRP Phase 2** is supported in PowerConnect B-Series TI24X devices. Refer to [“MRP rings with shared interfaces \(MRP Phase 2\)”](#) on page 200.

MRP protocol prevents Layer 2 loops and provides fast reconvergence in Layer 2 ring topologies. It is an alternative to STP and is especially useful in Metropolitan Area Networks (MANs) where using STP has the following drawbacks:

- STP allows a maximum of seven nodes. Metro rings can easily contain more nodes than this.
- STP has a slow reconvergence time, taking many seconds or even minutes. MRP can detect and heal a break in the ring in sub-second time.

Figure 33 shows an example of an MRP metro ring.

FIGURE 33 Metro ring – normal state



The ring in this example consists of four MRP nodes. Each node has two interfaces with the ring. Each node also is connected to a separate customer network. The nodes forward Layer 2 traffic to and from the customer networks through the ring. The ring interfaces are all in one port-based VLAN. Each customer interface can be in the same VLAN as the ring or in a separate VLAN.

One node is configured as the master node of the MRP ring. One of the two interfaces on the master node is configured as the primary interface; the other is the secondary interface. The primary interface originates Ring Health Packets (RHPs), which are used to monitor the health of the ring. An RHP is forwarded on the ring to the next interface until it reaches the secondary interface of the master node. The secondary interface blocks the packet to prevent a Layer 2 loops. When the master node transmits or receives RHPs, the CPU goes high on the master node.

Configuration notes


- When you configure MRP, Dell recommends that you disable one of the ring interfaces before beginning the ring configuration. Disabling an interface prevents a Layer 2 loop from occurring while you are configuring MRP on the ring nodes. Once MRP is configured and enabled on all the nodes, you can re-enable the interface.
- MRP 1 and MRP 2 support are added for the PowerConnect B-Series TI24X devices.
- The above configurations can be configured as MRP masters or MRP members (for different rings).
- On PowerConnect platforms, you can configure maximum of 32 metro rings.

MRP rings without shared interfaces (MRP Phase 1)

MRP Phase 1 allows you to configure multiple MRP rings, as shown in [Figure 34](#), but the rings cannot share the same link. For example, you cannot configure ring 1 and ring 2 to each have interfaces 1 and 2.

Also, when you configure an MRP ring, any node on the ring can be designated as the master node for the ring. A master node can be the master node of more than one ring. (Refer to [Figure 34](#).) Each ring is an independent ring and RHP packets are processed within each ring.

FIGURE 34 Metro ring – multiple rings

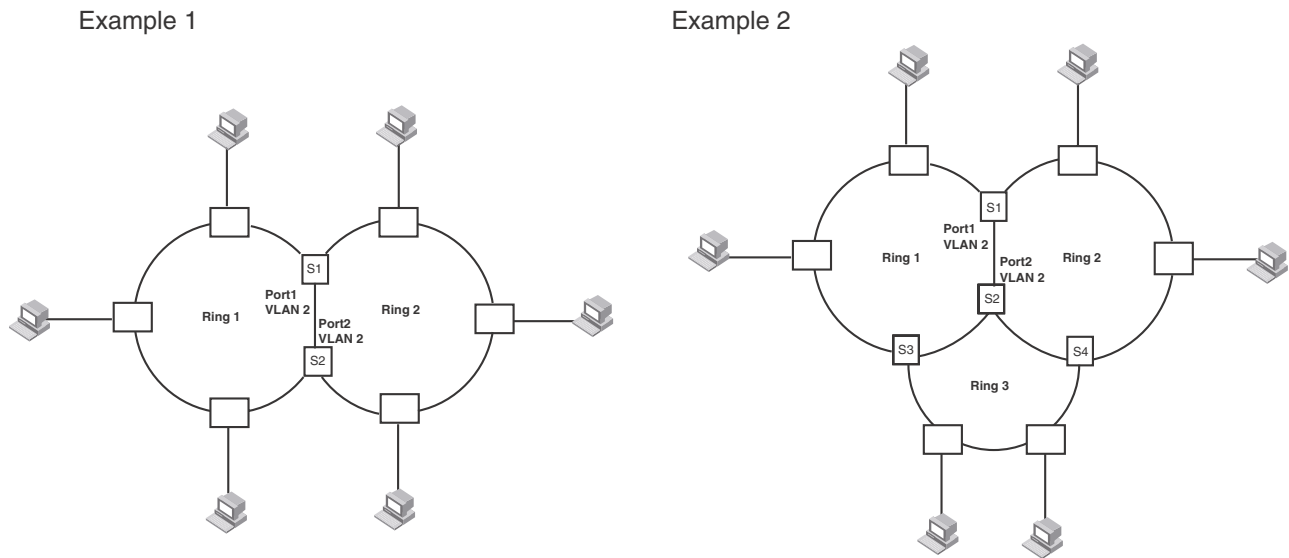


In this example, two nodes are each configured with two MRP rings. Any node in a ring can be the master for its ring. A node also can be the master for more than one ring.

MRP rings with shared interfaces (MRP Phase 2)


With MRP Phase 2, MRP rings can be configured to share the same interfaces as long as the interfaces belong to the same VLAN. [Figure 35](#) shows examples of multiple MRP rings that share the same interface.

FIGURE 35 Examples of multiple rings sharing the same interface - MRP Phase 2



On each node that will participate in the ring, you specify the ring ID and the interfaces that will be used for ring traffic. In a multiple ring configuration, a ring ID determines its priority. The lower the ring ID, the higher priority of a ring.

A ring ID is also used to identify the interfaces that belong to a ring.

FIGURE 36 Interface IDs and types

For example, in [Figure 36](#), the ID of all interfaces on all nodes on Ring 1 is 1 and all interfaces on all nodes on Ring 2 is 2. Port 1 on node S1 and Port 2 on S2 have the IDs of 1 and 2 since the interfaces are shared by Rings 1 and 2.

The ring ID is also used to determine an interface priority. Generally, a ring ID is also the ring priority and the priority of all interfaces on that ring. However, if the interface is shared by two or more rings, then the highest priority (lowest ID) becomes the priority of the interface. For example, in [Figure 36](#), all interfaces on Ring 1, except for Port 1 on node S1 and Port 2 on node S2 have a priority of 1. Likewise, all interfaces on Ring 2, except for Port 1 on node S1 and Port 2 on node S2 have a priority of 2. Port 1 on S1 and Port 2 on S2 have a priority of 1 since 1 is the highest priority (lowest ID) of the rings that share the interface.

If a node has interfaces that have different IDs, the interfaces that belong to the ring with the highest priority become regular ports. Those interfaces that do not belong to the ring with the highest priority become tunnel ports. In [Figure 36](#), nodes S1 and S2 have interfaces that belong to Rings 1 and 2. Those interfaces with a priority of 1 are regular ports. The interfaces with a priority of 2 are the tunnel ports since they belong to Ring 2, which has a lower priority than Ring 1.

Selection of master node


Allowing MRP rings to share interfaces limits the nodes that can be designated as the master node. Any node on an MRP ring that does not have a shared interface can be designated as the ring master node. However, if all nodes on the ring have shared interfaces, nodes that do not have tunnel ports can be designated as the master node of that ring. If none of the nodes meet these criteria, you must change the rings' priorities by reconfiguring the rings' ID.

In [Figure 36](#), any of the nodes on Ring 1, even S1 or S2, can be a master node since none of its interfaces are tunnel ports. However in Ring 2, neither S1 nor S2 can be a master node since these nodes contain tunnel ports.

Ring initialization

The ring shown in [Figure 33](#) shows the port states in a fully initialized ring without any broken links. [Figure 37](#) shows the initial state of the ring, when MRP is first enabled on the ring switches. All ring interfaces on the master node and member nodes begin in the Preforwarding state (PF).

FIGURE 37 Metro ring – initial state



MRP uses Ring Health Packets (RHPs) to monitor the health of the ring. An RHP is an MRP protocol packet. The source address is the MAC address of the master node and the destination MAC address is a protocol address for MRP. The master node generates RHPs and sends them on the ring. The state of a ring port depends on the RHPs.

RHP processing in MRP Phase 1

A ring interface can have one of the following MRP states:

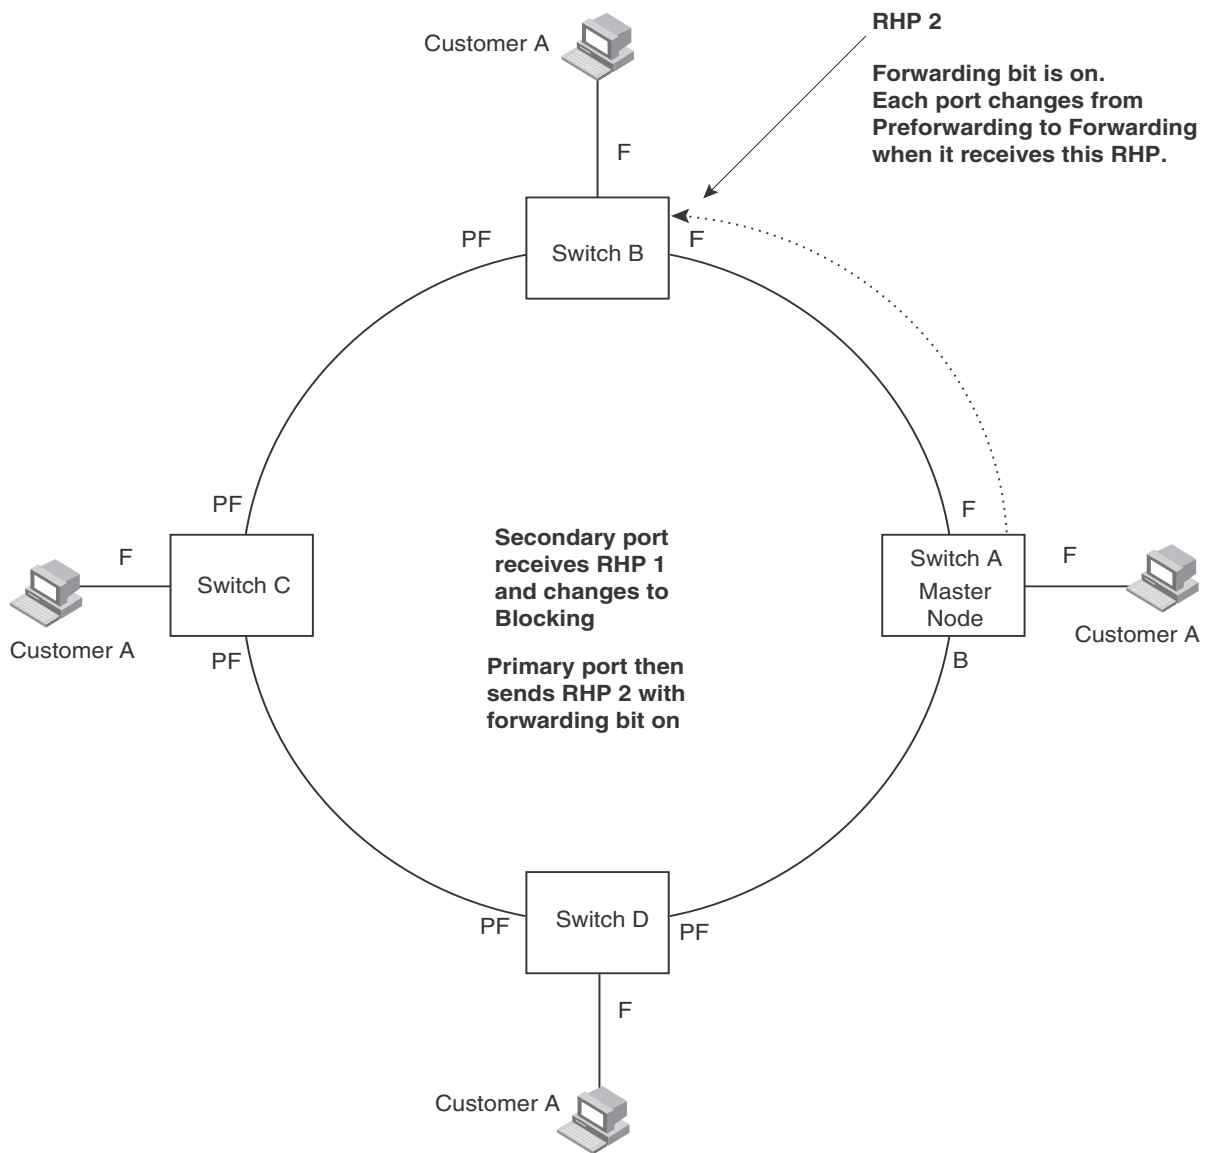
- **Preforwarding (PF)** – The interface can forward RHPs but cannot forward data. All ring ports begin in this state when you enable MRP.
- **Forwarding (F)** – The interface can forward data as well as RHPs. An interface changes from Preforwarding to Forwarding when the port preforwarding time expires. This occurs if the port does not receive an RHP from the master, or if the forwarding bit in the RHPs received by the port is off. This indicates a break in the ring. The port heals the ring by changing its state to Forwarding. The preforwarding time is the number of milliseconds the port will remain in the Preforwarding state before changing to the Forwarding state, even without receiving an RHP.
- **Blocking (B)** – The interface cannot forward data. Only the secondary interface on the master node can be Blocking.

When MRP is enabled, all ports begin in the Preforwarding state. The primary interface on the master node, although it is in the Preforwarding state like the other ports, immediately sends an RHP onto the ring. The secondary port on the master node listens for the RHP.

- If the secondary port receives the RHP, all links in the ring are up and the port changes its state to Blocking. The primary port then sends another MRP with its forwarding bit set on. As each of the member ports receives the RHP, the ports change their state to Forwarding. Typically, this occurs in sub-second time. The ring very quickly enters the fully initialized state.
- If the secondary port does not receive the RHP by the time the preforwarding time expires, a break has occurred in the ring. The port changes its state to Forwarding. The member ports also change their states from Preforwarding to Forwarding as their preforwarding timers expire. The ring is not intact, but data can still travel among the nodes using the links that are up.

Figure 38 shows an example.

FIGURE 38 Metro ring – from preforwarding to forwarding




Each RHP also has a sequence number. MRP can use the sequence number to determine the round-trip time for RHPs in the ring. Refer to [“Using MRP diagnostics”](#) on page 213.

RHP processing in MRP Phase 2

Figure 39 shows an example of how RHP packets are processed normally in MRP rings with shared interfaces.

FIGURE 39 Flow of RHP packets on MRP rings with shared interfaces



Port 1 on Ring 1 master node is the primary interface of the master node. The primary interface forwards an RHP packet on the ring. Since all the interfaces on Ring 1 are regular ports, the RHP packet is forwarded to all the interfaces until it reaches Port 2, the secondary interface of the master node. Port 2 then blocks the packet to complete the process.

On Ring 2, Port 3, is the primary interface of the master node. It sends an RHP packet on the ring. Since all ports on S4 are regular ports, the RHP packet is forwarded on those interfaces. When the packet reaches S2, the receiving interface is a tunnel port. The port compares the packet priority to its priority. Since the packet priority is the same as the tunnel port priority, the packet is forwarded up the link shared by Rings 1 and 2.

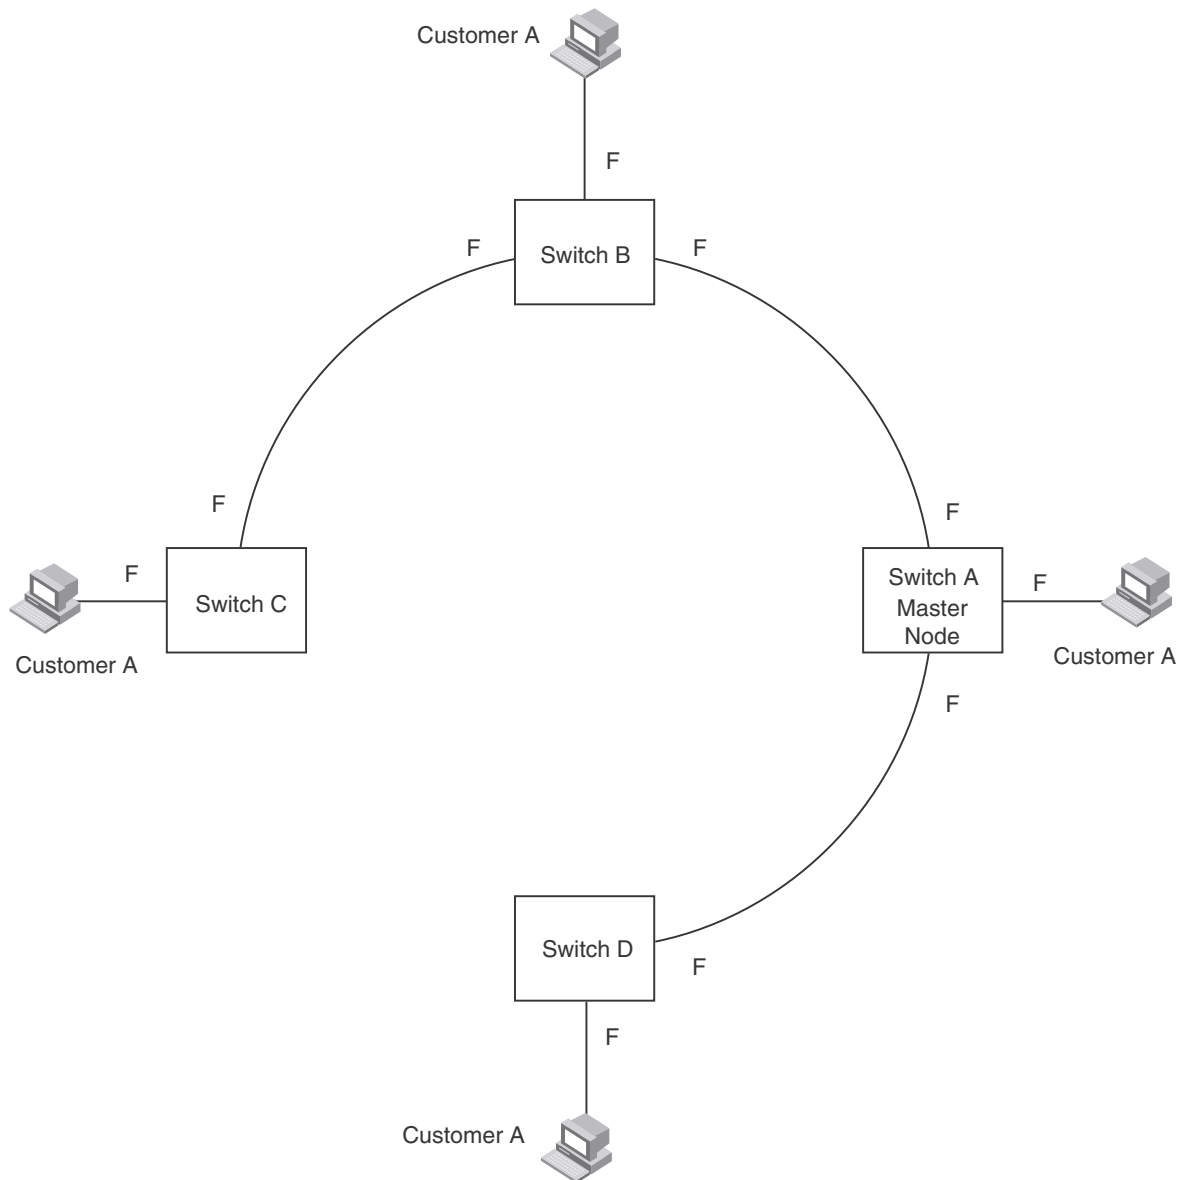
When the RHP packet reaches the interface on node S2 shared by Rings 1 and 2, the packet is forwarded since its priority is less than the interface priority. The packet continues to be forwarded to node S1 until it reaches the tunnel port on S1. That tunnel port determines that the RHP packet priority is equal to the port priority and forwards the packet. The RHP packet is forwarded to the remaining interfaces on Ring 2 until it reaches port 4, the secondary interface of the master node. Port 4 then blocks the packet to prevent a loop.

When the RHP packet from Ring 2 reached S2, it was also forwarded from S2 to S3 on Ring 1 since the port on S2 has a higher priority than the RHP packet. The packets is forwarded around Ring 1 until it reaches port 2, Ring 1 the secondary port. The RHP packet is then blocked by that port.

How ring breaks are detected and healed

Figure 40 shows ring interface states following a link break. MRP quickly heals the ring and preserves connectivity among the customer networks.

FIGURE 40 Metro ring – ring break



If a break in the ring occurs, MRP heals the ring by changing the states of some of the ring interfaces:

- **Blocking interface** – The Blocking interface on the master node has a dead timer. If the dead time expires before the interface receives one of its ring RHPs, the interface changes state to Preforwarding. Once the secondary interface changes state to Preforwarding:
 - If the interface receives an RHP, the interface changes back to the Blocking state and resets the dead timer.
 - If the interface does not receive an RHP for its ring before the Preforwarding time expires, the interface changes to the Forwarding state, as shown in [Figure 40](#).
- **Forwarding interfaces** – Each member interface remains in the Forwarding state.


When the broken link is repaired, the link interfaces come up in the Preforwarding state, which allows RHPs to travel through the restored interfaces and reach the secondary interface on the master node:

- If an RHP reaches the master node secondary interface, the ring is intact. The secondary interface changes to Blocking. The master node sets the forwarding bit on in the next RHP. When the restored interfaces receive this RHP, they immediately change state to Forwarding.
- If an RHP does not reach the master node secondary interface, the ring is still broken. The master node does not send an RHP with the forwarding bit on. In this case, the restored interfaces remain in the Preforwarding state until the preforwarding timer expires, then change to the Forwarding state.

If the link between **shared interfaces** breaks (Figure 41), the secondary interface on Ring 1 master node changes to a preforwarding state. The RHP packet sent by port 3 on Ring 2 is forwarded through the interfaces on S4, then to S2. The packet is then forwarded through S2 to S3, but not from S2 to S1 since the link between the two nodes is not available. When the packet reaches Ring 1 master node, the packet is forwarded through the secondary interface since it is currently in a preforwarding state. A secondary interface in preforwarding mode ignores any RHP packet that is not from its ring. The secondary interface changes to blocking mode only when the RHP packet forwarded by its primary interface is returned.

The packet then continues around Ring 1, through the interfaces on S1 to Ring 2 until it reaches Ring 2 master node. Port 4, the secondary interface on Ring 2 changes to blocking mode since it received its own packet, then blocks the packet to prevent a loop.

FIGURE 41 Flow of RHP packets when a link for shared interfaces breaks



RHP packets follow this flow until the link is restored; then the RHP packet returns to its normal flow as shown in Figure 39.

Alarm RHP

Previously, detection of MRP ring breaks was completely timer based. An absence of Ring Health Packets (RHP) for a period of 3 "hello times" indicated to the MRP master that the ring is broken. This initiated the transition to a topology change as described in the previous section. The convergence time associated with such an event could take several hundreds of milliseconds.

Now, each MRP node is made a more active participant in detecting link failures. When a link is detected to be down by its "downstream" neighbor, a special packet (called the Alarm RHP packet) is sent to the MRP master, indicating that the link is down.

This MRP packet is sent from the MRP member to the MRP master only when the secondary link goes down and it is sent on the primary link. The destination MAC address in the packet is the ring MAC address. This allows the packet to be hardware forwarded all the way to the MRP master. When the Master switch in the ring receives this packet, it is notified of a break in the ring. At that point, the secondary interface is immediately transitioned from "Blocked" to "Forwarding" .

NOTE

The Alarm RHP packet is only sent by the secondary link owner ring to prevent multiple MRP masters going into forward where shared rings are configured.

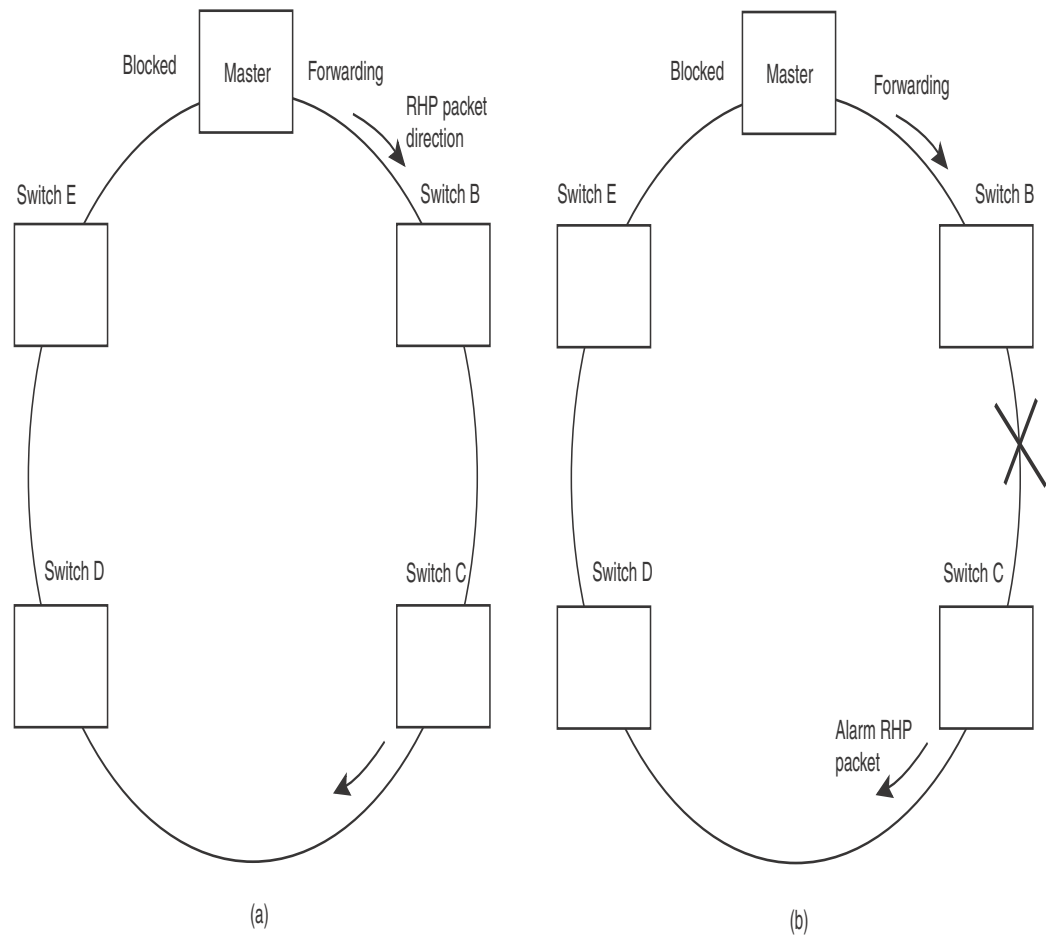
Operation of the MRP alarm RHP is as follows.

When the link between Switch B and Switch C fails, the "downstream" neighbor (Switch C) detects the failure of the link and triggers corrective action. The following is the complete sequence of events that occurs.

- The downstream neighbor (Switch C) detects a link down event of the link between Switch B and Switch C.
- Switch C sends a single RHP packet with a special Alarm bit set. The RHP packet is sent in the same direction of flow as that of the normal RHP packets (i.e. on the link to Switch D)
- Switch A receives the special RHP packet (on the secondary interface) that was sent by Switch C. It is now aware that the ring is broken even though the dead_interval may not have expired.
- Switch A immediately transitions its secondary interface (previously in Blocked state) to the Forwarding state.
- RHP packets continue to be sent on the primary interface by Switch A to detect if the ring has been healed.

From a user perspective, there is no difference in the behavior of the ring. The only noticeable difference is a rapid convergence in the event of ring failure. There is no CLI command required to enable this feature.

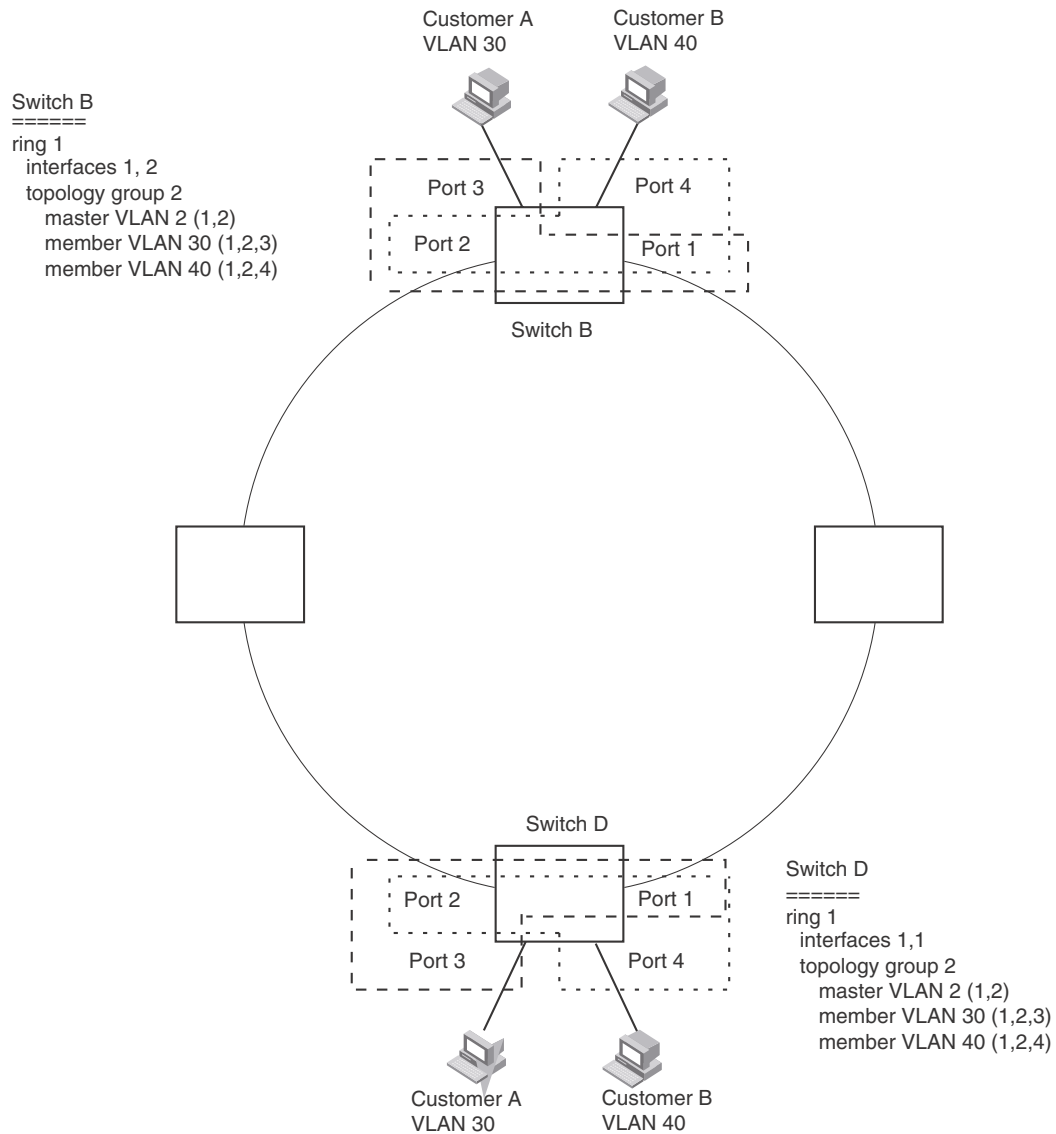
FIGURE 42 A MRP ring under normal operation (A) and after detection of a failure in the ring (B)



Master VLANs and customer VLANs

All the ring ports must be in the same VLAN. Placing the ring ports in the same VLAN provides Layer 2 connectivity for a given customer across the ring. [Figure 43](#) shows an example.

FIGURE 43 Metro ring – ring VLAN and customer VLANs



Notice that each customer has their own VLAN. Customer A has VLAN 30 and Customer B has VLAN 40. Customer A host attached to Switch D can reach the Customer A host attached to Switch B at Layer 2 through the ring. Since Customer A and Customer B are on different VLANs, they will not receive each other traffic.

You can configure MRP separately on each customer VLAN. However, this is impractical if you have many customers. To simplify configuration when you have a lot of customers (and therefore a lot of VLANs), you can use a topology group.

A topology group enables you to control forwarding in multiple VLANs using a single instance of a Layer 2 protocol such as MRP. A topology group contains a master VLAN and member VLANs. The master VLAN contains all the configuration parameters for the Layer 2 protocol (STP, MRP, or VSRP). The member VLANs use the Layer 2 configuration of the master VLAN.

In [Figure 43](#), VLAN 2 is the master VLAN and contains the MRP configuration parameters for ring 1. VLAN 30 and VLAN 40, the customer VLANs, are member VLANs in the topology group. Since a topology group is used, a single instance of MRP provides redundancy and loop prevention for both the customer VLANs.

If you use a topology group:

- The master VLAN must contain the ring interfaces. The ports must be tagged, since they will be shared by multiple VLANs.
- The member VLAN for a customer must contain the two ring interfaces and the interfaces for the customer. Since these interfaces are shared with the master VLAN, they must be tagged. Do not add another customer interfaces to the VLAN.

For more information about topology groups, refer to [“Topology groups”](#) on page 193.

Refer to [“MRP CLI example”](#) on page 216 for the configuration commands required to implement the MRP configuration shown in [Figure 43](#).

Configuring MRP

To configure MRP, perform the following tasks. You need to perform the first task on only one of the nodes. Perform the remaining tasks on all the nodes.

NOTE

There are no new commands or parameters to configure MRP with shared interfaces (MRP Phase 2).

- Disable one of the ring interfaces. This prevents a Layer 2 loop from occurring while you are configuring the devices for MRP.
- Add an MRP ring to a port-based VLAN. When you add a ring, the CLI changes to the configuration level for the ring, where you can perform the following tasks.
 - Optionally, specify a name for the ring.
 - On the master node only, enable the device to be the master for the ring. Each ring can have only one master node.
 - Specify the MRP interfaces. Each device has two interfaces to an MRP ring.
 - Optionally, change the hello time and the preforwarding time. These parameters control how quickly failover occurs following a change in the state of a link in the ring.
 - Enable the ring.
- Optionally, add the ring VLAN to a topology group to add more VLANs to the ring. If you use a topology group, make sure you configure MRP on the group master VLAN. Refer to [“Topology groups”](#) on page 193.
- Re-enable the interface you disabled to prevent a Layer 2 loop. Once MRP is enabled, MRP will prevent the Layer 2 loop.

Adding an MRP ring to a VLAN

To add an MRP ring to a VLAN, enter commands such as the following.

NOTE

If you plan to use a topology group to add VLANs to the ring, make sure you configure MRP on the topology group master VLAN.

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name CustomerA
PowerConnect(config-vlan-2-mrp-1)# master
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
```

These commands configure an MRP ring on VLAN 2. The ring ID is 1, the ring name is CustomerA, and this node (this PowerConnect device) is the master for the ring. The ring interfaces are 1 and 2. Interface 1 is the primary interface and 2 is the secondary interface. The primary interface will initiate RHPs by default. The ring takes effect in VLAN 2.

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name CustomerA
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
PowerConnect(config-vlan-2-mrp-1)# metro-ring 2
PowerConnect(config-vlan-2-mrp-2)# name CustomerB
PowerConnect(config-vlan-2-mrp-2)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-2)# enable
```

Syntax: [no] metro-ring <ring id>

The <ring-id> parameter specifies the ring ID. The <ring-id> can be from 1 - 1023; ID 256 is reserved for VSRP.

Syntax: [no] name <string>

The <string> parameter specifies a name for the ring. The name is optional, but it can be up to 20 characters long and can include blank spaces. If you use a name that has blank spaces, enclose the name in double quotation marks (for example: "Customer A").

Syntax: [no] master

Configures this node as the master node for the ring. Enter this command only on one node in the ring. The node is a member (non-master) node by default.

Syntax: [no] ring-interface ethernet <primary-if> ethernet <secondary-if>

The **ethernet** <primary-if> parameter specifies the primary interface. On the master node, the primary interface is the one that originates RHPs. Ring control traffic and Layer 2 data traffic will flow in the outward direction from this interface by default. On member nodes, the direction of traffic flow depends on the traffic direction selected by the master node. Therefore, on a member node, the order in which you enter the interfaces does not matter.

The **ethernet** <secondary-if> parameter specifies the secondary interface.

NOTE

To take advantage of every interface in a Metro network, you can configure another MRP ring and either configure a different master node for the ring or reverse the configuration of the primary and secondary interfaces on the master node. Configuring multiple rings enables you to use all the ports in the ring. The same port can forward traffic one ring while blocking traffic for another ring.

Syntax: [no] enable

The **enable** command enables the ring.

Changing the hello and preforwarding times

You also can change the RHP hello time and preforwarding time. To do so, enter commands such as the following.

```
PowerConnect(config-vlan-2-mrp-1)# hello-time 200
PowerConnect(config-vlan-2-mrp-1)# preforwarding-time 400
```

These commands change the hello time to 200 ms and change the preforwarding time to 400 ms.

Syntax: [no] hello-time <ms>

Syntax: [no] preforwarding-time <ms>

The <ms> specifies the number of milliseconds. For the hello time, you can specify from 100 – 1000 (one second). The default hello time is 100 ms. The preforwarding time can be from 200 – 5000 ms, but must be at least twice the value of the hello time and must be a multiple of the hello time. The default preforwarding time is 300 ms. A change to the hello time or preforwarding time takes effect as soon as you enter the command.

Configuration notes

- The preforwarding time must be at least twice the value of the hello time and must be a multiple of the hello time.
- If UDLD is also enabled on the device, Dell recommends that you set the MRP preforwarding time slightly higher than the default of 300 ms; for example, to 400 or 500 ms.
- You can use MRP ring diagnostics to determine whether you need to change the hello time and preforwarding time. Refer to [“Using MRP diagnostics”](#).

Using MRP diagnostics

The MRP diagnostics feature calculates how long it takes for RHP packets to travel through the ring. When you enable MRP diagnostics, the software tracks RHP packets according to their sequence numbers and calculates how long it takes an RHP packet to travel one time through the entire ring. When you display the diagnostics, the CLI shows the average round-trip time for the RHP packets sent since you enabled diagnostics. The calculated results have a granularity of 1 microsecond.

Enabling MRP diagnostics

To enable MRP diagnostics for a ring, enter the following command on the master node, at the configuration level for the ring.

```
PowerConnect(config-vlan-2-mrp-1)# diagnostics
```

Syntax: [no] diagnostics

NOTE

This command is valid only on the master node.

Displaying MRP diagnostics

To display MRP diagnostics results, enter the following command on the master node.

```
PowerConnect# show metro 1 diag

Metro Ring 1 - CustomerA
=====
diagnostics results

Ring      Diag      RHP average   Recommended   Recommended
id        state     time(microsec) hello time(ms) Prefwing time(ms)
2         enabled   125           100           300

Diag frame sent   Diag frame lost
1230              0
```

Syntax: `show metro <ring-id> diag`

This display shows the following information.

TABLE 33 CLI display of MRP ring diagnostic information

This field...	Displays...
Ring id	The ring ID.
Diag state	The state of ring diagnostics.
RHP average time	The average round-trip time for an RHP packet on the ring. The calculated time has a granularity of 1 microsecond.
Recommended hello time	The hello time recommended by the software based on the RHP average round-trip time.
Recommended Prefwing time	The preforwarding time recommended by the software based on the RHP average round-trip time.
Diag frame sent	The number of diagnostic RHPs sent for the test.
Diag frame lost	The number of diagnostic RHPs lost during the test.

If the recommended hello time and preforwarding time are different from the actual settings and you want to change them, refer to [“Configuring MRP”](#) on page 211.

Displaying MRP information

You can display the following MRP information:

- Topology group configuration information
- Ring configuration information and statistics

Displaying topology group information

To display topology group information, enter the following command.

Syntax: `show topology-group [<group-id>]`

Refer to [“Displaying topology group information”](#) on page 195 for more information.

Displaying ring information

To display ring information, enter the following command.

```
PowerConnect# show metro
```

```
Metro Ring 1
=====
Ring      State      Ring      Master      Topo      Hello      Prefwing
id        state      role      vlan        group     time(ms)   time(ms)
2         enabled   member    2           not conf  100        300

Ring interfaces      Interface role      Forwarding state      Active interface
Interface Type
ethernet 1          primary             disabled              none                  Regular
ethernet 2          secondary           forwarding            ethernet 2            Tunnel

RHPs sent          RHPs rcvd          TC RHPs rcvd          State changes
3                  0                  0                      4
```

Syntax: `show metro <ring-id>`

This display shows the following information.

TABLE 34 CLI display of MRP ring information

This field...	Displays...
Ring id	The ring ID
State	The state of MRP. The state can be one of the following: <ul style="list-style-type: none"> • enabled – MRP is enabled • disabled – MRP is disabled
Ring role	Whether this node is the master for the ring. The role can be one of the following: <ul style="list-style-type: none"> • master • member
Master vlan	The ID of the master VLAN in the topology group used by this ring. If a topology group is used by MRP, the master VLAN controls the MRP settings for all VLANs in the topology group. <p>NOTE: The topology group ID is 0 if the MRP VLAN is not the master VLAN in a topology group. Using a topology group for MRP configuration is optional.</p>
Topo group	The topology group ID.
Hello time	The interval, in milliseconds, at which the Forwarding port on the ring master node sends Ring hello Packets (RHPs).
Prefwing time	The number of milliseconds an MRP interface that has entered the Preforwarding state will wait before changing to the Forwarding state. If a member port in the Preforwarding state does not receive an RHP within the Preforwarding time (Prefwing time), the port assumes that a topology change has occurred and changes to the Forwarding state. The secondary port on the master node changes to Blocking if it receives an RHP, but changes to Forwarding if the port does not receive an RHP before the preforwarding time expires. <p>NOTE: A member node Preforwarding interface also changes from Preforwarding to Forwarding if it receives an RHP whose forwarding bit is on.</p>

TABLE 34 CLI display of MRP ring information (Continued)

This field...	Displays...
Ring interfaces	The device two interfaces with the ring. NOTE: If the interfaces are trunk groups, only the primary ports of the groups are listed.
Interface role	The interface role can be one of the following: <ul style="list-style-type: none"> • primary <ul style="list-style-type: none"> • Master node – The interface generates RHPs. • Member node – The interface forwards RHPs received on the other interface (the secondary interface). • secondary – The interface does not generate RHPs. <ul style="list-style-type: none"> • Master node – The interface listens for RHPs. • Member node – The interface receives RHPs.
Forwarding state	Whether MRP Forwarding is enabled on the interface. The forwarding state can be one of the following: <ul style="list-style-type: none"> • blocking – The interface is blocking Layer 2 data traffic and RHPs • disabled – The interface is down • forwarding – The interface is forwarding Layer 2 data traffic and RHPs • preforwarding – The interface is listening for RHPs but is blocking Layer 2 data traffic
Active interface	The physical interfaces that are sending and receiving RHPs. NOTE: If a port is disabled, its state is shown as “disabled”. NOTE: If an interface is a trunk group, only the primary port of the group is listed.
Interface Type	Shows if the interface is a regular port or a tunnel port.
RHPs sent	The number of RHPs sent on the interface. NOTE: This field applies only to the master node. On non-master nodes, this field contains 0. This is because the RHPs are forwarded in hardware on the non-master nodes.
RHPs rcvd	The number of RHPs received on the interface. NOTE: On most devices, this field applies only to the master node. On non-master nodes, this field contains 0. This is because the RHPs are forwarded in hardware on the non-master nodes.
TC RHPs rcvd	The number of Topology Change RHPs received on the interface. A Topology Change RHP indicates that the ring topology has changed.
State changes	The number of MRP interface state changes that have occurred. The state can be one of the states listed in the Forwarding state field.
Interface Type	Shows if the interface is a regular port or a tunnel port.

MRP CLI example

The following examples show the CLI commands required to implement the MRP configuration shown in [Figure 43](#) on page 210.

NOTE

For simplicity, the figure shows the VLANs on only two switches. The CLI examples implement the ring on all four switches.

Commands on Switch A (master node)

The following commands configure a VLAN for the ring. The ring VLAN must contain both of the node interfaces with the ring. Add these interfaces as tagged interfaces, since the interfaces also must be in each of the customer VLANs configured on the node.

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# tag ethernet 1 to 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name "Metro A"
PowerConnect(config-vlan-2-mrp-1)# master
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
PowerConnect(config-vlan-2-mrp-1)# exit
PowerConnect(config-vlan-2)# exit
```

The following commands configure the customer VLANs. The customer VLANs must contain both the ring interfaces as well as the customer interfaces.

```
PowerConnect(config)# vlan 30
PowerConnect(config-vlan-30)# tag ethernet 1 to 2
PowerConnect(config-vlan-30)# tag ethernet 3
PowerConnect(config-vlan-30)# exit
PowerConnect(config)#vlan 40
PowerConnect(config-vlan-40)# tag ethernet 1 to 2
PowerConnect(config-vlan-40)# tag ethernet 4
PowerConnect(config-vlan-40)# exit
```

The following commands configure topology group 1 on VLAN 2. The master VLAN is the one that contains the MRP configuration. The member VLANs use the MRP parameters of the master VLAN. The control interfaces (the ones shared by the master VLAN and member VLAN) also share MRP state.

```
PowerConnect(config)# topology-group 1
PowerConnect(config-topo-group-1)# master-vlan 2
PowerConnect(config-topo-group-1)# member-vlan 30
PowerConnect(config-topo-group-1)# member-vlan 40
```

Commands on Switch B

The commands for configuring Switches B, C, and D are similar to the commands for configuring Switch A, with two differences: the nodes are not configured to be the ring master. Omitting the **master** command is required for non-master nodes.

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# tag ethernet 1 to 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name "Metro A"
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
PowerConnect(config-vlan-2)# exit
PowerConnect(config)#vlan 30
PowerConnect(config-vlan-30)# tag ethernet 1 to 2
PowerConnect(config-vlan-30)# tag ethernet 3
PowerConnect(config-vlan-30)# exit
PowerConnect(config)#vlan 40
PowerConnect(config-vlan-40)# tag ethernet 1 to 2
PowerConnect(config-vlan-40)# tag ethernet 4
PowerConnect(config-vlan-40)# exit
```

8 Virtual Switch Redundancy Protocol (VSRP)

```
PowerConnect(config)# topology-group 1
PowerConnect(config-topo-group-1)# master-vlan 2
PowerConnect(config-topo-group-1)# member-vlan 30
PowerConnect(config-topo-group-1)# member-vlan 40
```

Commands on Switch C

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# tag ethernet 1 to 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name "Metro A"
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
PowerConnect(config-vlan-2)# exit
PowerConnect(config)# vlan 30
PowerConnect(config-vlan-30)# tag ethernet 1 to 2
PowerConnect(config-vlan-30)# tag ethernet 3
PowerConnect(config-vlan-30)# exit
PowerConnect(config)# vlan 40
PowerConnect(config-vlan-40)# tag ethernet 1 to 2
PowerConnect(config-vlan-40)# tag ethernet 4
PowerConnect(config-vlan-40)# exit
PowerConnect(config)# topology-group 1
PowerConnect(config-topo-group-1)# master-vlan 2
PowerConnect(config-topo-group-1)# member-vlan 30
PowerConnect(config-topo-group-1)# member-vlan 40
```

Commands on Switch D

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# tag ethernet 1 to 2
PowerConnect(config-vlan-2)# metro-ring 1
PowerConnect(config-vlan-2-mrp-1)# name "Metro A"
PowerConnect(config-vlan-2-mrp-1)# ring-interface ethernet 1 ethernet 2
PowerConnect(config-vlan-2-mrp-1)# enable
PowerConnect(config-vlan-2)# exit
PowerConnect(config)# vlan 30
PowerConnect(config-vlan-30)# tag ethernet 1 to 2
PowerConnect(config-vlan-30)# tag ethernet 3
PowerConnect(config-vlan-30)# exit
PowerConnect(config)# vlan 40
PowerConnect(config-vlan-40)# tag ethernet 1 to 2
PowerConnect(config-vlan-40)# tag ethernet 4
PowerConnect(config-vlan-40)# exit
PowerConnect(config)# topology-group 1
PowerConnect(config-topo-group-1)# master-vlan 2
PowerConnect(config-topo-group-1)# member-vlan 30
PowerConnect(config-topo-group-1)# member-vlan 40
```

Virtual Switch Redundancy Protocol (VSRP)

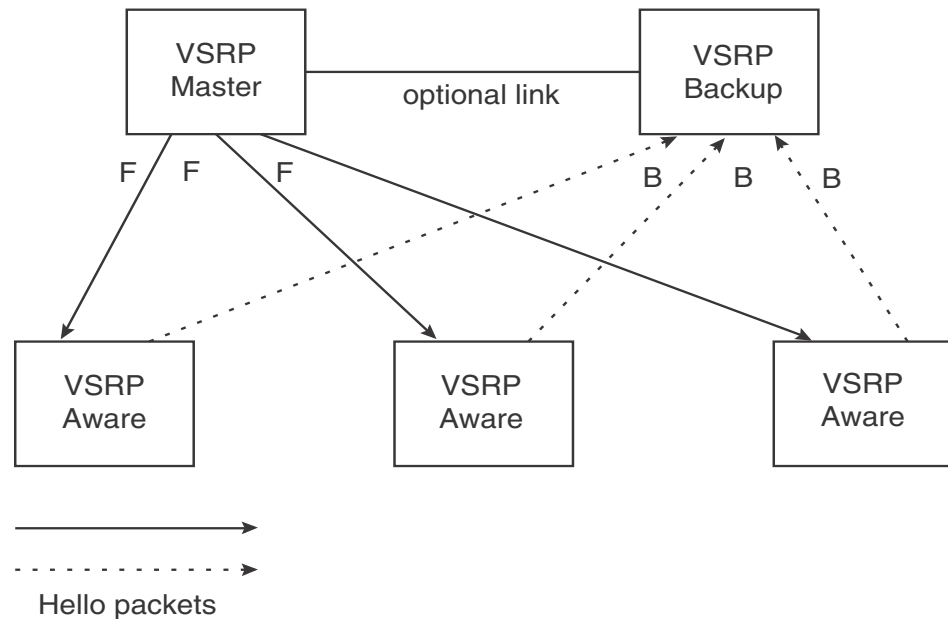
Virtual Switch Redundancy Protocol (VSRP) protocol provides redundancy and sub-second failover in Layer 2 and Layer 3 mesh topologies. Based on the Virtual Router Redundancy Protocol Extended (VRRPE), VSRP provides one or more backups for a Layer 2 Switch or Layer 3 Switch. If the active Layer 2 Switch or Layer 3 Switch becomes unavailable, one of the backups takes over as the active device and continues forwarding traffic for the network.

The PowerConnect support full VSRP and **VSRP-awareness**. A PowerConnect device that is not itself configured for VSRP, but is connected to a PowerConnect device that is configured for VSRP, is **VSRP aware**.

You can use VSRP for Layer 2, Layer 3, or for both layers. On Layer 3 Switches, Layer 2 and Layer 3 share the same VSRP configuration information. On Layer 2 Switches, VSRP applies only to Layer 2.

Figure 44 shows an example of a VSRP configuration.

FIGURE 44 VSRP mesh – redundant paths for Layer 2 and Layer 3 traffic



In this example, two devices are configured as redundant paths for VRID 1. On each of the devices, a Virtual Router ID (VRID) is configured on a port-based VLAN. Since VSRP is primarily a Layer 2 redundancy protocol, the VRID applies to the entire VLAN. However, you can selectively remove individual ports from the VRID if needed.

Following master election (described below), one of the devices becomes the master for the VRID and sets the state of all the VLAN ports to Forwarding. The other device is a backup and sets all the ports in its VRID VLAN to Blocking.

If a failover occurs, the backup becomes the new master and changes all its VRID ports to the Forwarding state.

Other devices can use the redundant paths provided by the VSRP devices. In this example, three devices use the redundant paths. A device that is not itself configured for VSRP but is connected to a device that is configured for VSRP, is **VSRP aware**. In this example, the three devices connected to the VSRP devices are VSRP aware. A device that is VSRP aware can failover its link to the new master in sub-second time, by changing the MAC address associated with the redundant path.

When you configure VSRP, make sure each of the non-VSRP devices connected to the VSRP devices has a separate link to each of the VSRP devices.

Configuration notes

- VSRP and 802.1Q-n-Q tagging are not supported together on the same device.
- VSRP and Super Aggregated VLANs are not supported together on the same device.
- VSRP does not work on a VLAN which has multicast enabled.
- PowerConnect support VSRP awareness, and VSRP-aware security features.

Layer 2 and Layer 3 redundancy

You can configure VSRP to provide redundancy for Layer 2 only or also for Layer 3:

- **Layer 2 only** – The Layer 2 links are backed up but specific IP addresses are not backed up.
- **Layer 2 and Layer 3** – The Layer 2 links are backed up and a specific IP address is also backed up. Layer 3 VSRP is the same as VRRPE. However, using VSRP provides redundancy at both layers at the same time.

Layer 2 Switches support Layer 2 VSRP only. Layer 3 Switches support Layer 2 and Layer 3 redundancy. You can configure a Layer 3 Switch for either Layer 2 only or Layer 2 and Layer 3. To configure for Layer 3, specify the IP address you are backing up.

NOTE

If you want to provide Layer 3 redundancy only, disable VSRP and use VRRPE.

Master election and failover

Each VSRP device advertises its VSRP priority in hello messages. During master election, the VSRP device with the highest priority for a given VRID becomes the master for that VRID. After master election, the master sends hello messages at regular intervals to inform the backups that the master is healthy.

If there is a tie for highest VSRP priority, the tie is resolved as follows:

- **Layer 2 Switches** – The Layer 2 Switch with the higher management IP address becomes the master.
 - Switches with management IP addresses are preferred over switches without management IP addresses.
 - If neither of the switches has a management IP address, then the switch with the higher MAC address becomes the master. (VSRP compares the MAC addresses of the ports configured for the VRID, not the base MAC addresses of the switches.)
- **Layer 3 Switches** – The Layer 3 Switch whose virtual routing interface has a higher IP address becomes the master.

VSRP failover

Each backup listens for hello messages from the master. The hello messages indicate that the master is still available. If the backups stop receiving hello messages from the master, the election process occurs again and the backup with the highest priority becomes the new master.

Each backup waits for a specific period of time, the dead Interval, to receive a new hello message from the master. If the backup does not receive a hello message from the master by the time the dead interval expires, the backup sends a hello message of its own, which includes the backup's VSRP priority, to advertise the backup's intent to become the master. If there are multiple backups for the VRID, each backup sends a hello message.

When a backup sends a hello message announcing its intent to become the master, the backup also starts a hold-down timer. During the hold-down time, the backup listens for a hello message with a higher priority than its own.

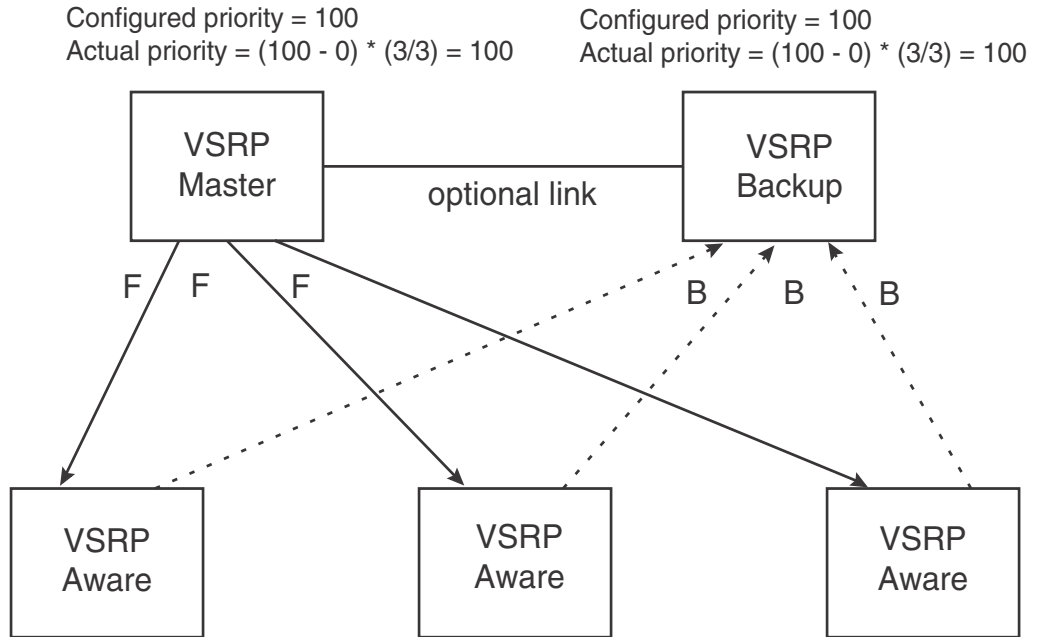
- If the backup receives a hello message with a higher priority than its own, the backup resets its dead interval and returns to normal backup status.
- If the backup does not receive a hello message with a higher priority than its own by the time the hold-down timer expires, the backup becomes the new master and starts forwarding Layer 2 traffic on all ports.

If you increase the timer scale value, each timer value is divided by the scale value. To achieve sub-second failover times, you can change the scale to a value up to 10. This shortens all the VSRP timers to 10 percent of their configured values.

VSRP priority calculation

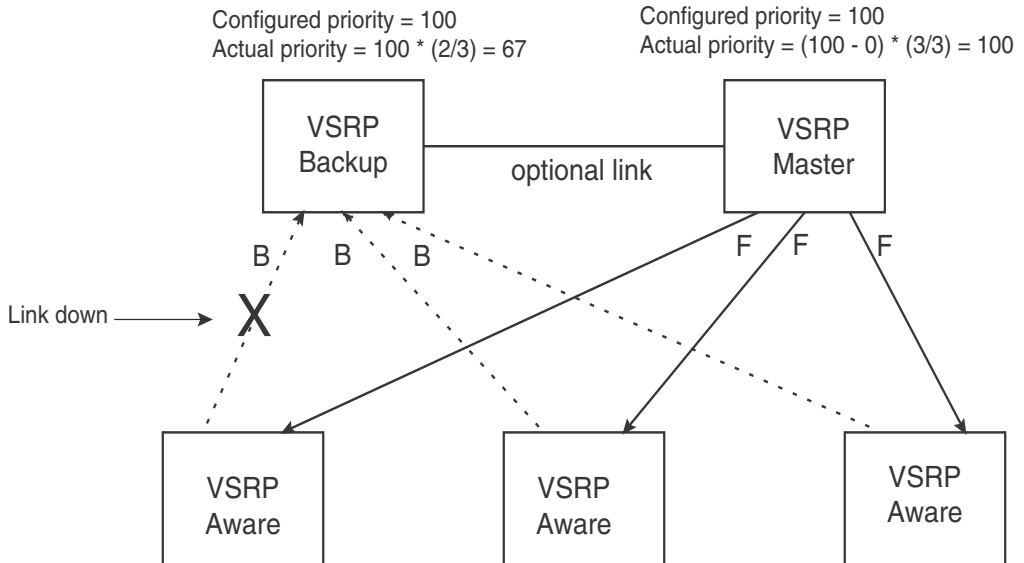
Each VSRP device has a VSRP priority for each VRID and its VLAN. The VRID is used during master election for the VRID. By default, a device VSRP priority is the value configured on the device (which is 100 by default). However, to ensure that a backup with a high number of up ports for a given VRID is elected, the device reduces the priority if a port in the VRID VLAN goes down. For example, if two backups each have a configured priority of 100, and have three ports in VRID 1 in VLAN 10, each backup begins with an equal priority, 100. This is shown in [Figure 45](#)

FIGURE 45 VSRP priority

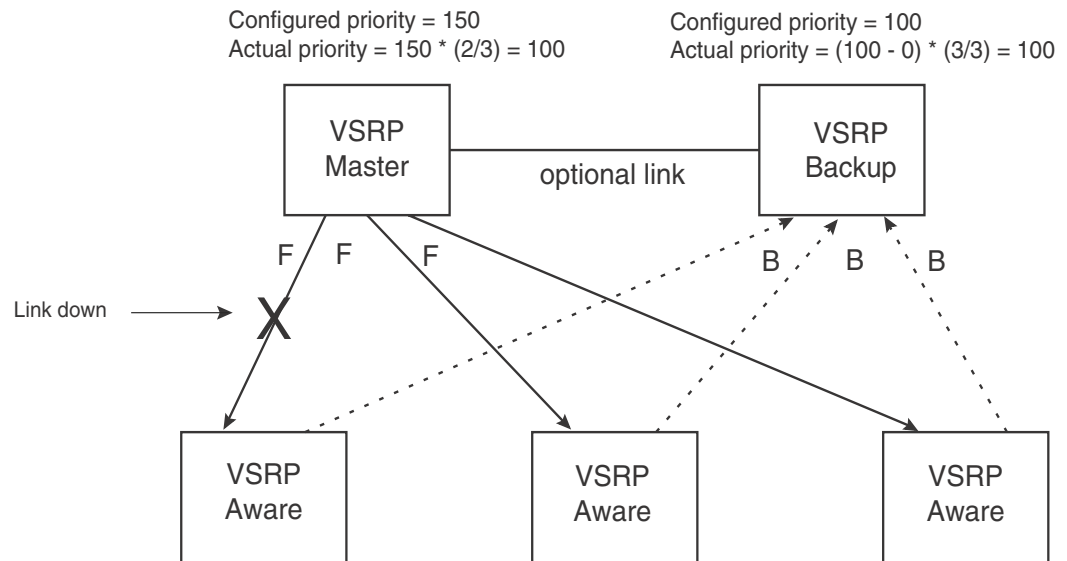


However, if one of the VRID ports goes down on one of the backups, that backup priority is reduced. If the master priority is reduced enough to make the priority lower than a backup priority, the VRID fails over to the backup. Figure 46 shows an example.

FIGURE 46 VSRP priority recalculation



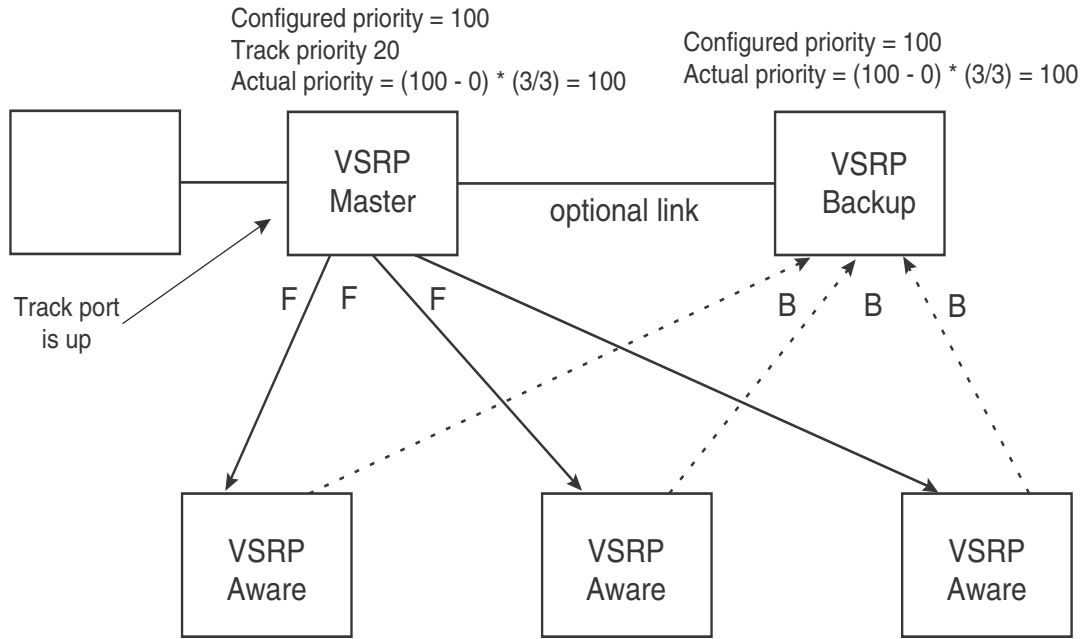
You can reduce the sensitivity of a VSRP device to failover by increasing its configured VSRP priority. For example, you can increase the configured priority of the VSRP device on the left in Figure 46 to 150. In this case, failure of a single link does not cause failover. The link failure caused the priority to be reduced to 100, which is still equal to the priority of the other device. This is shown in Figure 47.

FIGURE 47 VSRP priority bias**Track ports**

Optionally, you can configure track ports to be included during VSRP priority calculation. In VSRP, a **track port** is a port that is not a member of the VRID VLAN, but whose state is nonetheless considered when the priority is calculated. Typically, a track port represents the exit side of traffic received on the VRID ports. By default, no track ports are configured.

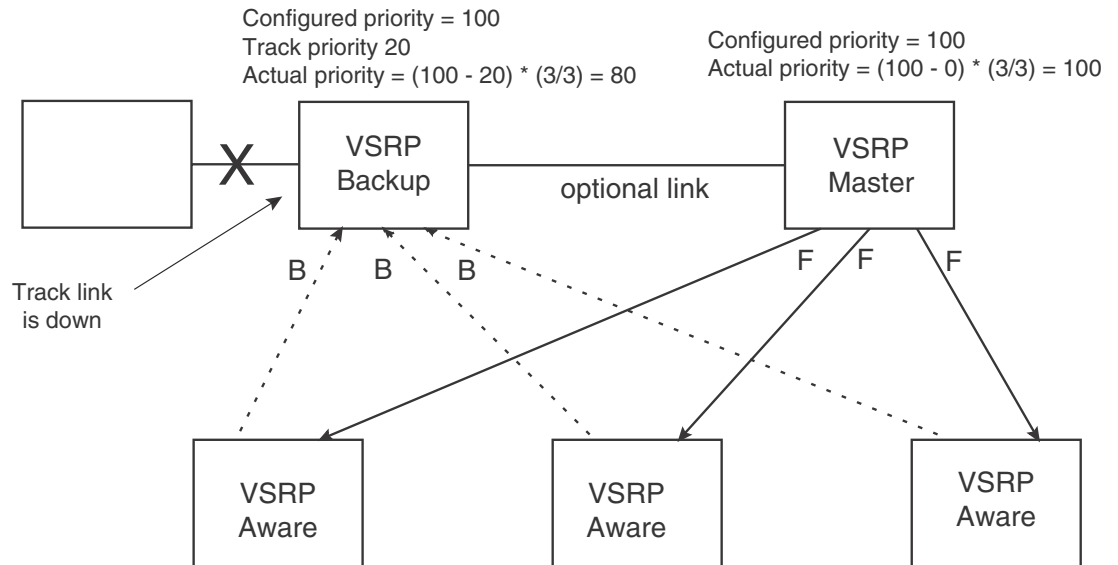
When you configure a track port, you assign a priority value to the port. If the port goes down, VSRP subtracts the track port priority value from the configured VSRP priority. For example, if you configure a track port with priority 20 and the configured VSRP priority is 100, the software subtracts 20 from 100 if the track port goes down, resulting in a VSRP priority of 80. The new priority value is used when calculating the VSRP priority. [Figure 48](#) shows an example.

FIGURE 48 Track port priority



In [Figure 48](#), the track port is up. Since the port is up, the track priority does not affect the VSRP priority calculation. If the track port goes down, the track priority does affect VSRP priority calculation, as shown in [Figure 49](#).

FIGURE 49 Track port priority subtracted during priority calculation



MAC address failover on VSRP-aware devices

VSRP-aware devices maintain a record of each VRID and its VLAN. When the device has received a hello message for a VRID in a given VLAN, the device creates a record for that VRID and VLAN and includes the port number in the record. Each subsequent time the device receives a hello message for the same VRID and VLAN, the device checks the port number:

- If the port number is the same as the port that previously received a hello message, the VSRP-aware device assumes that the message came from the same VSRP master that sent the previous message.
- If the port number does not match, the VSRP-aware device assumes that a VSRP failover has occurred to a new master, and moves the MAC addresses learned on the previous port to the new port.

The VRID records age out if unused. This can occur if the VSRP-aware device becomes disconnected from the master. The VSRP-aware device will wait for a hello message for the period of time equal to the following.

$$\text{VRID Age} = \text{Dead Interval} + \text{Hold-down Interval} + (3 \times \text{Hello Interval})$$

The values for these timers are determined by the VSRP device sending the hello messages. If the master uses the default timer values, the age time for VRID records on the VSRP-aware devices is as follows.

$$3 + 2 + (3 \times 1) = 8 \text{ seconds}$$

In this case, if the VSRP-aware device does not receive a new hello message for a VRID in a given VLAN, on any port, the device assumes the connection to the master is unavailable and removes the VRID record.

VSRP-Aware security features

This feature protects against unauthorized VSRP hello packets by enabling you to configure VSRP-aware security parameters. Without VSRP-aware security, a VSRP-aware device passively learns the authentication method conveyed by the received VSRP hello packet. The VSRP-aware device then stores the authentication method until it ages out with the aware entry.

The VSRP-aware security feature enables you to perform the following:

- Define the specific authentication parameters that a VSRP-aware device will use on a VSRP backup switch. The authentication parameters that you define will not age out.
- Define a list of ports that have authentic VSRP backup switch connections. For ports included in the list, the VSRP-aware switch will process VSRP hello packets using the VSRP-aware security configuration. Conversely, for ports not included in the list, the VSRP-aware switch will not use the VSRP-aware security configuration.

If VSRP hello packets do not meet the acceptance criteria, the VSRP-aware device forwards the packets normally, without any VSRP-aware security processing.

To configure VSRP-Aware Security features, refer to [“Configuring security features on a VSRP-aware device”](#) on page 231.

VSRP parameters

[Table 35](#) lists the VSRP parameters.

TABLE 35 VSRP parameters

Parameter	Description	Default	See page...
Protocol	VSRP state NOTE: On a Layer 3 Switch, you must disable VSRP to use VRRPE or VRRP.	Enabled	page 229
Virtual Router ID (VRID)	The ID of the virtual switch you are creating by configuring multiple devices as redundant links. You must configure the same VRID on each device that you want to use to back up the links.	None	page 228
Timer scale	The value used by the software to calculate all VSRP timers. Increasing the timer scale value decreases the length of all the VSRP timers equally, without changing the ratio of one timer to another.	1	page 230
Interface parameters			
Authentication type	The type of authentication the VSRP devices use to validate VSRP packets. On Layer 3 Switches, the authentication type must match the authentication type the VRID port uses with other routing protocols such as OSPF. <ul style="list-style-type: none"> • No authentication – The interfaces do not use authentication. This is the VRRP default. • Simple – The interface uses a simple text-string as a password in packets sent on the interface. If the interface uses simple password authentication, the VRID configured on the interface must use the same authentication type and the same password. NOTE: MD5 is not supported.	No authentication	page 231
VSRP-Aware Security Parameters			
VSRP-Aware Authentication type	The type of authentication the VSRP-aware devices will use on a VSRP backup switch: <ul style="list-style-type: none"> • No authentication – The device does not accept incoming packets that have authentication strings. • Simple – The device uses a simple text-string as the authentication string for accepting incoming packets. 	Not configured	page 231
VRID parameters			
VSRP device type	Whether the device is a VSRP backup for the VRID. All VSRP devices for a given VRID are backups.	Not configured	page 228
VSRP ports	The ports in the VRID VLAN that you want to use as VRID interfaces. You can selectively exclude individual ports from VSRP while allowing them to remain in the VLAN.	All ports in the VRID VLAN	page 232

TABLE 35 VSRP parameters (Continued)

Parameter	Description	Default	See page...
VRID IP address	<p>A gateway address you are backing up. Configuring an IP address provides VRRPE Layer 3 redundancy in addition to VSRP Layer 2 redundancy.</p> <p>The VRID IP address must be in the same subnet as a real IP address configured on the VSRP interface, but cannot be the same as a real IP address configured on the interface.</p> <p>NOTE: This parameter is valid only on Layer 3 Switches.</p>	None	page 232
Backup priority	<p>A numeric value that determines a backup preferability for becoming the master for the VRID. During negotiation, the device with the highest priority becomes the master.</p> <p>In VSRP, all devices are backups and have the same priority by default.</p> <p>If two or more backups are tied with the highest priority, the backup with the highest IP address becomes the master for the VRID.</p>	100 for all backups	page 233
Preference of timer source	<p>When you save a backup configuration, the software can save the configured VSRP timer values or the VSRP timer values received from the master.</p> <p>Saving the current timer values instead of the configured ones helps ensure consistent timer usage for all the VRID devices.</p> <p>NOTE: The backup always gets its timer scale value from the master.</p>	Configured timer values are saved	page 233
Time-to-Live (TTL)	The maximum number of hops a VSRP hello packet can traverse before being dropped. You can specify from 1 – 255.	2	page 234
Hello interval	The amount of time between hello messages from the master to the backups for a given VRID. The interval can be from 1 – 84 seconds.	One second	page 234
Dead interval	The amount of time a backup waits for a hello message from the master for the VRID before determining that the master is no longer active. If the master does not send a hello message before the dead interval expires, the backups negotiate (compare priorities) to select a new master for the VRID.	Three times the hello Interval	page 235
Backup hello state and interval	<p>The amount of time between hello messages from a backup to the master.</p> <p>The message interval can be from 60 – 3600 seconds.</p> <p>You must enable the backup to send the messages. The messages are disabled by default on backups. The current master sends hello messages by default.</p>	Disabled 60 seconds when enabled	page 235

TABLE 35 VSRP parameters (Continued)

Parameter	Description	Default	See page...
Hold-down interval	The amount of time a backup that has sent a hello packet announcing its intent to become master waits before beginning to forward traffic for the VRID. The hold-down interval prevents Layer 2 loops from occurring during VSRP rapid failover. The interval can from 1 – 84 seconds.	2 seconds	page 235
Track priority	A VSRP priority value assigned to the tracked ports. If a tracked port link goes down, the VRID port VSRP priority is reduced by the amount of the tracked port priority.	5	page 236
Track port	A track port is a port or virtual routing interface that is outside the VRID but whose link state is tracked by the VRID. Typically, the tracked interface represents the other side of VRID traffic flow through the device. If the link for a tracked interface goes down, the VSRP priority of the VRID interface is changed, causing the devices to renegotiate for master.	None	page 236
Backup preempt mode	Prevents a backup with a higher VSRP priority from taking control of the VRID from another backup that has a lower priority but has already assumed control of the VRID.	Enabled	page 237
VRID active state	The active state of the VSRP VRID.	Disabled	page 228
RIP parameters			
Suppression of RIP advertisements	A Layer 3 Switch that is running RIP normally advertises routes to a backed up VRID even when the Layer 3 Switch is not currently the active Layer 3 Switch for the VRID. Suppression of these advertisements helps ensure that other Layer 3 Switches do not receive invalid route paths for the VRID. This parameter is valid only on Layer 3 Switches.	Disabled (routes are advertised)	page 237

NOTE

To configure VSRP, ensure that spanning tree is disabled on the VLAN.

Configuring basic VSRP parameters

To configure VSRP, perform the following required tasks:

- Configure a port-based VLAN containing the ports for which you want to provide VSRP service.

NOTE

If you already have a port-based VLAN but only want to use VSRP on a sub-set of the VLANs ports, you can selectively remove ports from VSRP service in the VLAN. Refer to [“Removing a port from the VRID VLAN”](#) on page 232.

- Configure a VRID:

- Specify that the device is a backup. Since VSRP, like VRRPE, does not have an “owner”, all VSRP devices are backups. The active device for a VRID is elected based on the VRID priority, which is configurable.
- Activate the VRID.

The following example shows a simple VSRP configuration.

```
PowerConnect(config)# vlan 200
PowerConnect(config-vlan-200)# tag ethernet 1 to 8
PowerConnect(config-vlan-200)# vsrp vrid 1
PowerConnect(config-vlan-200-vrid-1)# backup
PowerConnect(config-vlan-200-vrid-1)# activate
```

Syntax: [no] vsrp vrid <num>

The <num> parameter specifies the VRID and can be from 1 – 255.

Syntax: [no] backup [priority <value>] [track-priority <value>]

This command is required. In VSRP, all devices on which a VRID are configured are backups. The master is then elected based on the VSRP priority of each device. There is no “owner” device as there is in VRRP.

For information about the command optional parameters, refer to the following:

- [“Changing the backup priority”](#) on page 233
- [“Changing the default track priority”](#) on page 236

Syntax: [no] activate

or

Syntax: enable | disable

Configuring optional VSRP parameters

The following sections describe how to configure optional VSRP parameters.

Disabling or re-enabling VSRP

VSRP is enabled by default on Layer 2 Switches and Layer 3 Switches. On a Layer 3 Switch, if you want to use VRRP or VRRPE for Layer 3 redundancy instead of VSRP, you need to disable VSRP first. To do so, enter the following command at the global CONFIG level.

```
PowerConnect(config)#no router vsrp
router vsrp is disabled. All vsrp config data will be lost when writing to flash
```

To re-enable the protocol, enter the following command.

```
PowerConnect(config)# router vsrp
```

Syntax: [no] router vsrp

Since VRRP and VRRPE do not apply to Layer 2 Switches, there is no need to disable VSRP and there is no command to do so. The protocol is always enabled.

Timer scale

The VSRP hello interval, dead interval, backup hello interval, and hold-down interval timers are individually configurable. You also can easily change all the timers at the same time while preserving the ratios among their values. To do so, change the timer scale. The timer scale is a value used by the software to calculate the timers. The software divides a timer value by the timer scale value. By default, the scale is 1. This means the VSRP timer values are the same as the values in the configuration. For more information on timer scale parameters and how it affects the timers, refer to Table 35.

Changing the timer scale

To achieve sub-second failover times, you can shorten the duration of all scale timers for VSRP, VRRP, and VRRP-E by adjusting the timer scale. The **timer scale** is a value used by the software to calculate the timers. By default, the scale value is 1. If you increase the timer scale, each timer value is divided by the scale value. Using the timer scale to adjust timer values enables you to easily change all the timers while preserving the ratios among their values. Here is an example.

Table 1:

Timer	Timer scale	Timer value
Hello interval	1	1 second
	2	0.5 seconds
Dead interval	1	3 seconds
	2	1.5 seconds
Backup Hello interval	1	60 seconds
	2	30 seconds
Hold-down interval	1	2 seconds
	2	1 second

If you configure the device to receive its timer values from the master, the backup also receives the timer scale value from the master.

NOTE

The backups always use the value of the timer scale received from the master, regardless of whether the timer values that are saved in the configuration are the values configured on the backup or the values received from the master.

To change the timer scale, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# scale-timer 2
```

This command changes the scale to 2. All VSRP, VRRP, and VRRP-E timer values will be divided by 2.

Syntax: [no] **scale-timer** <num>

The <num> parameter specifies the multiplier. You can specify a timer scale from 1 – 10.

Configuring authentication

If the interfaces on which you configure the VRID use authentication, the VSRP packets on those interfaces also must use the same authentication. VSRP supports the following authentication types:

- **No authentication** – The interfaces do not use authentication. This is the default.
- **Simple** – The interfaces use a simple text-string as a password in packets sent on the interface. If the interfaces use simple password authentication, the VRID configured on the interfaces must use the same authentication type and the same password.

To configure a simple password, enter a command such as the following at the interface configuration level.

```
PowerConnect(config-if-6)# ip vsrp auth-type simple-text-auth ourpword
```

This command configures the simple text password “ourpword”.

Syntax: [no] ip vsrp auth-type no-auth | simple-text-auth <auth-data>

The **auth-type no-auth** parameter indicates that the VRID and the interface it is configured on do not use authentication.

The **auth-type simple-text-auth <auth-data>** parameter indicates that the VRID and the interface it is configured on use a simple text password for authentication. The <auth-data> value is the password. If you use this parameter, make sure all interfaces on all the devices supporting this VRID are configured for simple password authentication and use the same password.

Configuring security features on a VSRP-aware device

This section shows how to configure security features on a VSRP-aware device. For an overview of this feature, refer to “[VSRP-Aware security features](#)” on page 225.

Specifying an authentication string for VSRP hello packets

The following configuration defines **pri-key** as the authentication string for accepting incoming VSRP hello packets. In this example, the VSRP-aware device will accept all incoming packets that have this authorization string.

```
PowerConnect(config)# vlan 10
PowerConnect(config-vlan-10)# vsrp-aware vrid 3 simple-text-auth pri-key
```

Syntax: vsrp-aware vrid <vrid number> simple text auth <string>

Specifying no authentication for VSRP hello packets

The following configuration specifies no authentication as the preferred VSRP-aware security method. In this case, the VSRP device will not accept incoming packets that have authentication strings.

```
PowerConnect(config)# vlan 10
PowerConnect(config-vlan-10)# vsrp-aware vrid 2 no-auth
```

Syntax: vsrp-aware vrid <vrid number> no-auth

The following configuration specifies no authentication for VSRP hello packets received on ports 1, 2, 3, and 4 in VRID 4. For these ports, the VSRP device will not accept incoming packets that have authentication strings.

```
PowerConnect(config)# vlan 10
PowerConnect(config-vlan-10)# vsrp-aware vrid 4 no-auth port-list ethe 1 to 4
```

Syntax: `vsrp-aware vrid <vrid number> no-auth port-list <port range>`

`<vrid number>` is a valid VRID (from 1 to 255).

no-auth specifies no authentication as the preferred VSRP-aware security method. The VSRP device will not accept incoming packets that have authentication strings.

simple-text-auth <string> specifies the authentication string for accepting VSRP hello packets, where `<string>` can be up to 8 characters.

port-list <port range> specifies the range of ports to include in the configuration.

Removing a port from the VRID VLAN

By default, all the ports on which you configure a VRID are interfaces for the VRID. You can remove a port from the VRID while allowing it to remain in the VLAN.

Removing a port is useful in the following cases:

- There is no risk of a loop occurring, such as when the port is attached directly to an end host.
- You plan to use a port in an MRP ring.

To remove a port from a VRID, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# no include-port ethernet 2
```

Syntax: `[no] include-port ethernet <portnum>`

The `<portnum>` parameter specifies the port you are removing from the VRID. The port remains in the VLAN but its forwarding state is not controlled by VSRP.

Configuring a VRID IP address

If you are configuring a Layer 3 Switch for VSRP, you can specify an IP address to back up. When you specify an IP address, VSRP provides redundancy for the address. This is useful if you want to back up the gateway address used by hosts attached to the VSRP backups.

VSRP does not require you to specify an IP address. If you do not specify an address, VSRP provides Layer 2 redundancy. If you do specify an address, VSRP provides Layer 2 and Layer 3 redundancy.

The Layer 3 redundancy support is the same as VRRPE support. For information, refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).

NOTE

The VRID IP address must be in the same subnet as a real IP address configured on the VSRP interface, but cannot be the same as a real IP address configured on the interface.

NOTE

Failover applies to both Layer 2 and Layer 3.

To specify an IP address to back up, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# ip-address 10.10.10.1
```

Syntax: `[no] ip-address <ip-addr>`

or

Syntax: [no] ip address <ip-addr>

Changing the backup priority

When you enter the backup command to configure the device as a VSRP backup for the VRID, you also can change the backup priority and the track priority:

- The backup priority is used for election of the master. The VSRP backup with the highest priority value for the VRID is elected as the master for that VRID. The default priority is 100. If two or more backups are tied with the highest priority, the backup with the highest IP address becomes the master for the VRID.
- The track priority is used with the track port feature. Refer to [“VSRP priority calculation”](#) on page 221 and [“Changing the default track priority”](#) on page 236.

To change the backup priority, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# backup priority 75
```

Syntax: [no] backup [priority <value>] [track-priority <value>]

The **priority <value>** parameter specifies the VRRP priority for this interface and VRID. You can specify a value from 3 – 254. The default is 100.

For a description of the **track-priority <value>** parameter, refer to [“Changing the default track priority”](#) on page 236.

Example

Following is an example to configure a simple VSRP interface.

Enter the following commands to add tagged ports to a VLAN.

```
PowerConnect(config)# vlan 30
PowerConnect(config-vlan-30)# tag ethernet 28
Enter the following command to added the IP address.
PowerConnect(config-vlan-30)# router-int ve 30
PowerConnect(config-vlan-30)# interface ve 30
PowerConnect(config-vif-30)# ip address 30.0.0.2/24
```

Enter the following commands to change the VSRP interface state and initiate VRID state.

```
PowerConnect(config-vif-30)# vlan 30
PowerConnect(config-vlan-30)# vsrp vri 30
PowerConnect(config-vlan-30-vrid-30)# backup
PowerConnect(config-vlan-30-vrid-30)# ip-address 30.0.0.1
PowerConnect(config-vlan-30-vrid-30)# enable
```

Saving the timer values received from the master

The hello messages sent by a VRID master contain the VRID values for the following VSRP timers:

- Hello interval
- Dead interval
- Backup hello interval
- Hold-down interval

By default, each backup saves the configured timer values to its startup-config file when you save the device configuration.

You can configure a backup to instead save the current timer values received from the master when you save the configuration. Saving the current timer values instead of the configured ones helps ensure consistent timer usage for all the VRID devices.

NOTE

The backups always use the value of the timer scale received from the master, regardless of whether the timer values that are saved in the configuration are the values configured on the backup or the values received from the master.

To configure a backup to save the VSRP timer values received from the master instead of the timer values configured on the backup, enter the following command.

```
PowerConnect(config-vlan-200-vrid-1)# save-current-values
```

Syntax: [no] save-current-values

Changing the Time-To-Live (TTL)

A VSRP hello packet TTL specifies how many hops the packet can traverse before being dropped. A hop can be a Layer 3 Switch or a Layer 2 Switch. You can specify from 1 – 255. The default TTL is 2. When a VSRP device (master or backup) sends a VSRP HELLO packet, the device subtracts one from the TTL. Thus, if the TTL is 2, the device that originates the hello packet sends it out with a TTL of 1. Each subsequent device that receives the packet also subtracts one from the packet TTL. When the packet has a TTL of 1, the receiving device subtracts 1 and then drops the packet because the TTL is zero.

NOTE

An MRP ring is considered to be a single hop, regardless of the number of nodes in the ring.

To change the TTL for a VRID, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# initial-ttl 5
```

Syntax: [no] initial-ttl <num>

The <num> parameter specifies the TTL and can be from 1 – 255. The default TTL is 2.

Changing the hello interval

The master periodically sends hello messages to the backups. To change the hello interval, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# hello-interval 10
```

Syntax: [no] hello-interval <num>

The <num> parameter specifies the interval and can be from 1 – 84 seconds. The default is 1 second.

NOTE

The default dead interval is three times the hello interval plus one-half second. Generally, if you change the hello interval, you also should change the dead interval on the backups.

NOTE

If you change the timer scale, the change affects the actual number of seconds.

Changing the dead interval

The dead interval is the number of seconds a backup waits for a hello message from the master before determining that the master is dead. The default is 3 seconds. This is three times the default hello interval.

To change the dead interval, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# dead-interval 30
```

Syntax: [no] dead-interval <num>

The <num> parameter specifies the interval and can be from 1 – 84 seconds. The default is 3 seconds.

NOTE

If you change the timer scale, the change affects the actual number of seconds.

Changing the backup hello state and interval

By default, backups do not send hello messages to advertise themselves to the master. You can enable these messages if desired and also change the message interval.

To enable a backup to send hello messages to the master, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# advertise backup
```

Syntax: [no] advertise backup

When a backup is enabled to send hello messages, the backup sends a hello message to the master every 60 seconds by default. You can change the interval to be up to 3600 seconds.

To change the backup hello interval, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# backup-hello-interval 180
```

Syntax: [no] backup-hello-interval <num>

The <num> parameter specifies the message interval and can be from 60 – 3600 seconds. The default is 60 seconds.

NOTE

If you change the timer scale, the change affects the actual number of seconds.

Changing the hold-down interval

The hold-down interval prevents Layer 2 loops from occurring during failover, by delaying the new master from forwarding traffic long enough to ensure that the failed master is really unavailable.

To change the Hold-down interval, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# hold-down-interval 4
```

Syntax: [no] hold-down-interval <num>

The <num> parameter specifies the hold-down interval and can be from 1 – 84 seconds. The default is 2 seconds.

NOTE

If you change the timer scale, the change affects the actual number of seconds.

Changing the default track priority

When you configure a VRID to track the link state of other interfaces, if one of the tracked interface goes down, the software changes the VSRP priority of the VRID interface.

The software reduces the VRID priority by the amount of the priority of the tracked interface that went down. For example, if the VSRP interface priority is 100 and a tracked interface with track priority 60 goes down, the software changes the VSRP interface priority to 40. If another tracked interface goes down, the software reduces the VRID priority again, by the amount of the tracked interface track priority.

The default track priority for all track ports is 1. You can change the default track priority or override the default for an individual track port.

- To change the default track priority, use the **backup track-priority** command, described below.
- To override the default track priority for a specific track port, use the **track-port** command. Refer to “[Specifying a track port](#)” on page 236.

To change the track priority, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# backup track-priority 2
```

Syntax: [no] backup [priority <value>] [track-priority <value>]

Specifying a track port

You can configure the VRID on one interface to track the link state of another interface on the device. This capability is useful for tracking the state of the exit interface for the path for which the VRID is providing redundancy. Refer to “[VSRP priority calculation](#)” on page 221.

To configure a VRID to track an interface, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# track-port e 4
```

Syntax: [no] track-port ethernet <portnum> | ve <num> [priority <num>]

The **priority <num>** parameter changes the VSRP priority of the interface. If this interface goes down, the VRID VSRP priority is reduced by the amount of the track port priority you specify here.

NOTE

The priority <num> option changes the priority of the specified interface, overriding the default track port priority. To change the default track port priority, use the **backup track-priority <num>** command.

Disabling or re-enabling backup pre-emption

By default, a backup that has a higher priority than another backup that has become the master can preempt the master, and take over the role of master. If you want to prevent this behavior, disable preemption.

Preemption applies only to backups and takes effect only when the master has failed and a backup has assumed ownership of the VRID. The feature prevents a backup with a higher priority from taking over as master from another backup that has a lower priority but has already become the master of the VRID.

Preemption is especially useful for preventing flapping in situations where there are multiple backups and a backup with a lower priority than another backup has assumed ownership, because the backup with the higher priority was unavailable when ownership changed.

If you enable the non-preempt mode (thus disabling the preemption feature) on all the backups, the backup that becomes the master following the disappearance of the master continues to be the master. The new master is not preempted.

To disable preemption on a backup, enter a command such as the following at the configuration level for the VRID.

```
PowerConnect(config-vlan-200-vrid-1)# non-preempt-mode
```

Syntax: [no] non-preempt-mode

Suppressing RIP advertisement from backups

Normally, for Layer 3 a VSRP backup includes route information for a backed up IP address in RIP advertisements. As a result, other Layer 3 Switches receive multiple paths for the backed up interface and might sometimes unsuccessfully use the path to the backup rather than the path to the master.

You can prevent the backups from advertising route information for the backed up interface by enabling suppression of the advertisements.

NOTE

This parameter applies only if you specified an IP address to back up and is valid only on Layer 3 Switches.

To suppress RIP advertisements, enter the following commands.

```
Router2(config)#router rip  
Router2(config-rip-router)# use-vrrp-path
```

Syntax: [no] use-vrrp-path

VSRP-aware interoperability

The **vsrp-aware tc-vlan-flush** command should be used in network configurations in which the switch operates as the VSRP-Aware device connecting to the other devices configured as a VSRP master.

The command is available at the VLAN level, and is issued per a specific VRID, as shown here for VRID 11.

```
PowerConnect(config-vlan-10)# vsrp-aware vrid 11 tc-vlan-flush
```

Syntax: `vsrp-aware vrid <num> tc-vlan-flush`

When this command is enabled, MAC addresses will be flushed at the VLAN level, instead of at the port level. MAC addresses will be flushed for every topology change (TC) received on the VSRP-aware ports.

When this command is enabled, the results of the `show vsrp-aware vlan` command resemble the following.

```
PowerConnect(config-vlan-10)#vsrp-aware vrid 11 tc-vlan-flush
PowerConnect(config-vlan-10)#show vsrp aware vlan 10
Aware Port Listing
  VLAN ID VRID Last Port Auth Type      Mac-Flush Age
    10      11 N/A   no-auth Configured Enabled  00:00:00.0
```

Displaying VSRP information

You can display the following VSRP information:

- Configuration information and current parameter values for a VRID or VLAN
- The interfaces on a VSRP-aware device that are active for the VRID

Displaying VRID information

To display VSRP information, enter the following command.

```
PowerConnect# show vsrp vrid 1
Total number of VSRP routers defined: 2
VLAN 200
auth-type no authentication
VRID 1
  State      Administrative-status Advertise-backup Preempt-mode save-current
  standby    enabled              disabled          true          false

  Parameter      Configured Current Unit
  priority        100      80   (100-0)*(4.0/5.0)
  hello-interval 1         1    sec/1
  dead-interval  3         3    sec/1
  hold-interval  3         3    sec/1
  initial-ttl    2         2    hops

  next hello sent in 00:00:00.8
  Member ports:     ethe 1 to 5
  Operational ports: ethe 1 to 4
  Forwarding ports: ethe 1 to 4
```

Syntax: `show vsrp [vrid <num> | vlan <vlan-id>]`

This display shows the following information when you use the `vrid <num>` or `vlan <vlan-id>` parameter. For information about the display when you use the `aware` parameter, refer to [“Displaying the active interfaces for a VRID”](#) on page 240.

TABLE 36 CLI display of VSRP VRID or VLAN information

This field...	Displays...
Total number of VSRP routers defined	The total number of VRIDs configured on this device.
VLAN	The VLAN on which VSRP is configured.
auth-type	The authentication type in effect on the ports in the VSRP VLAN.
VRID parameters	
VRID	The VRID for which the following information is displayed.
state	<p>This device VSRP state for the VRID. The state can be one of the following:</p> <ul style="list-style-type: none"> initialize – The VRID is not enabled (activated). If the state remains “initialize” after you activate the VRID, make sure that the VRID is also configured on the other routers and that the routers can communicate with each other. <p>NOTE: If the state is “initialize” and the mode is incomplete, make sure you have specified the IP address for the VRID.</p> <ul style="list-style-type: none"> standby – This device is a backup for the VRID. master – This device is the master for the VRID.
Administrative-status	<p>The administrative status of the VRID. The administrative status can be one of the following:</p> <ul style="list-style-type: none"> disabled – The VRID is configured on the interface but VSRP or VRRPE has not been activated on the interface. enabled – VSRP has been activated on the interface.
Advertise-backup	<p>Whether the device is enabled to send VSRP hello messages when it is a backup. This field can have one of the following values:</p> <ul style="list-style-type: none"> disabled – The device does not send hello messages when it is a backup. enabled – The device does send hello messages when it is a backup.
Preempt-mode	<p>Whether the device can be pre-empted by a device with a higher VSRP priority after this device becomes the master. This field can have one of the following values:</p> <ul style="list-style-type: none"> disabled – The device cannot be pre-empted. enabled – The device can be pre-empted.
save-current	<p>The source of VSRP timer values preferred when you save the configuration. This field can have one of the following values:</p> <ul style="list-style-type: none"> false – The timer values configured on this device are saved. true – The timer values most recently received from the master are saved instead of the locally configured values.
<p>NOTE: For the following fields:</p> <ul style="list-style-type: none"> Configured – indicates the parameter value configured on this device. Current – indicates the parameter value received from the master. Unit – indicates the formula used for calculating the VSRP priority and the timer scales in effect for the VSRP timers. A timer true value is the value listed in the Configured or Current field divided by the scale value. 	
priority	<p>The device preferability for becoming the master for the VRID. During negotiation, the backup with the highest priority becomes the master.</p> <p>If two or more backups are tied with the highest priority, the backup interface with the highest IP address becomes the master for the VRID.</p>
hello-interval	The number of seconds between hello messages from the master to the backups for a given VRID.

TABLE 36 CLI display of VSRP VRID or VLAN information (Continued)

This field...	Displays...
dead-interval	The configured value for the dead interval. The dead interval is the number of seconds a backup waits for a hello message from the master for the VRID before determining that the master is no longer active. If the master does not send a hello message before the dead interval expires, the backups negotiate (compare priorities) to select a new master for the VRID. NOTE: If the value is 0, then you have not configured this parameter.
hold-interval	The number of seconds a backup that intends to become the master will wait before actually beginning to forward Layer 2 traffic for the VRID. If the backup receives a hello message with a higher priority than its own before the hold-down interval expires, the backup remains in the backup state and does not become the new master.
initial-ttl	The number of hops a hello message can traverse after leaving the device before the hello message is dropped. NOTE: An MRP ring counts as one hop, regardless of the number of nodes in the ring.
next hello sent in	The amount of time until the master dead interval expires. If the backup does not receive a hello message from the master by the time the interval expires, either the IP address listed for the master will change to the IP address of the new master, or this Layer 3 Switch itself will become the master. NOTE: This field applies only when this device is a backup.
Member ports	The ports in the VRID.
Operational ports	The member ports that are currently up.
Forwarding ports	The member ports that are currently in the Forwarding state. Ports that are forwarding on the master are listed. Ports on the Standby, which are in the Blocking state, are not listed.

Displaying the active interfaces for a VRID

On a VSRP-aware device, you can display VLAN and port information for the connections to the VSRP devices (master and backups).

To display the active VRID interfaces, enter the following command on the VSRP-aware device.

```
PowerConnect# show vsrp aware
Aware port listing
VLAN ID  VRID  Last Port
100      1      2
200      2      1
```

Syntax: show vsrp aware

This display shows the following information when you use the **aware** parameter. For information about the display when you use the **vrid <num>** or **vlan <vlan-id>** parameter, refer to “[Displaying VRID information](#)” on page 238.

TABLE 37 CLI display of VSRP-aware information

This field...	Displays...
VLAN ID	The VLAN that contains the VSRP-aware device connection with the VSRP master and backups.

TABLE 37 CLI display of VSRP-aware information (Continued)

This field...	Displays...
VRID	The VRID.
Last Port	The most recent active port connection to the VRID. This is the port connected to the current master. If a failover occurs, the VSRP-aware device changes the port to the port connected to the new master. The VSRP-aware device uses this port to send and receive data through the backed up node.

VSRP fast start

VSRP fast start allows non-Dell or non-VSRP aware devices that are connected to a Dell device that is the VSRP master to quickly switchover to the new master when a VSRP failover occurs

This feature causes the port on a VSRP master to restart when a VSRP failover occurs. When the port shuts down at the start of the restart, ports on the non-VSRP aware devices that are connected to the VSRP master flush the MAC address they have learned for the VSRP master. After a specified time, the port on the previous VSRP master (which now becomes the backup) returns back online. Ports on the non-VSRP aware devices switch over to the new master and learn its MAC address.

Configuring VSRP fast start

The VSRP fast start feature can be enabled on a VSRP-configured Dell device, either on the VLAN to which the VRID of the VSRP-configured device belongs (globally) or on a port that belongs to the VRID.

To globally configure a VSRP-configured device to shut down its ports when a failover occurs, then restart after five seconds, enter the following command.

```
PowerConnect(configure)# vlan 100
PowerConnect(configure-vlan-100)# vsrp vrid 1
PowerConnect(configure-vlan-100-vrid-1)# restart-ports 5
```

Syntax: [no] restart-ports <seconds>

This command shuts down all the ports that belong to the VLAN when a failover occurs. All the ports will have the specified VRID.

To configure a single port on a VSRP-configured device to shut down when a failover occurs, then restart after a period of time, enter the following command.

```
PowerConnect(configure)# interface ethernet 1
PowerConnect(configure-if-1)# vsrp restart-port 5
```

Syntax: [no] vsrp restart-port <seconds>

In both commands, the <seconds> parameter instructs the VSRP master to shut down its port for the specified number of seconds before it starts back up. Enter a value between 1 – 120 seconds. The default is 1 second.

Displaying ports that Have the VSRP fast start feature enabled

The **show vsrp vrid** command shows the ports on which the VSRP fast start feature is enabled.

8 Virtual Switch Redundancy Protocol (VSRP)


```
PowerConnect# show vsrp vrid 100
VLAN 100
  auth-type no authentication
  VRID 100
  =====
  State      Administrative-status  Advertise-backup  Preempt-mode  save-current
  master     enabled                 disabled          true          false
  Parameter  Configured  Current  Unit/Formula
  priority   100         50      (100-0)*(2.0/4.0)
  hello-interval  1          1      sec/1
  dead-interval  3          3      sec/1
  hold-interval  3          3      sec/1
  initial-ttl   2          2      hops
  next hello sent in 00:00:00.3
  Member ports:  ethe 5 to 8
  Operational ports: ethe 5 ethe 8
  Forwarding ports: ethe 5 ethe 8
  Restart ports: 5(1) 6(1) 7(1) 8(1)
```

The "Restart ports:" line lists the ports that have the VSRP fast start enabled, and the downtime for each port. Refer to [Table 36](#) on page 239 to interpret the remaining information on the display.

VSRP and MRP signaling


A device may connect to an MRP ring through VSRP to provide a redundant path between the device and the MRP ring. VSRP and MRP signaling ensures rapid failover by flushing MAC addresses appropriately. The host on the MRP ring learns the MAC addresses of all devices on the MRP ring and VSRP link. From these MAC addresses, the host creates a MAC database (table), which is used to establish a data path from the host to a VSRP-linked device. [Figure 50](#) below shows two possible data paths from the host to Device 1.

FIGURE 50 Two data paths from host on an MRP ring to a VSRP-linked device



If a VSRP failover from master to backup occurs, VSRP needs to inform MRP of the topology change; otherwise, data from the host continues along the obsolete learned path and never reach the VSRP-linked device, as shown in [Figure 51](#).

FIGURE 51 VSRP on MRP rings that failed over

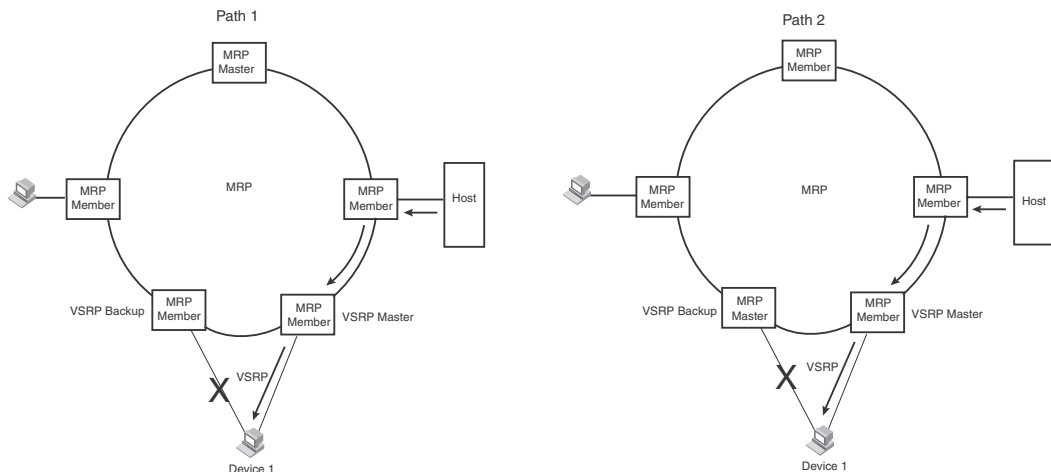


A signaling process for the interaction between VSRP and MRP ensures that MRP is informed of the topology change and achieves convergence rapidly. When a VSRP node fails, a new VSRP master is selected. The new VSRP master finds all MRP instances impacted by the failover. Then each MRP instance does the following:

- The MRP node sends out an MRP PDU with the mac-flush flag set three times on the MRP ring.
- The MRP node that receives this MRP PDU empties all the MAC entries from its interfaces that participate on the MRP ring.
- The MRP node then forwards the MRP PDU with the mac-flush flag set to the next MRP node that is in forwarding state.

The process continues until the master MRP node secondary (blocking) interface blocks the packet. Once the MAC address entries have been flushed, the MAC table can be rebuilt for the new path from the host to the VSRP-linked device ([Figure 52](#)).

FIGURE 52 New path established



8 Virtual Switch Redundancy Protocol (VSRP)


There are no CLI commands used to configure this process.

Configuring Uni-Directional Link Detection (UDLD) and Protected Link Groups

UDLD overview

Uni-Directional Link Detection (UDLD) monitors a link between two devices and brings the ports on both ends of the link down if the link goes down at any point between the two devices. This feature is useful for links that are individual ports and for trunk links. [Figure 53](#) shows an example.

FIGURE 53 UDLD example



Normally, a device load balances traffic across the ports in a trunk group. In this example, each device load balances traffic across two ports. Without the UDLD feature, a link failure on a link that is not directly attached to one of the devices is undetected by the devices. As a result, the devices continue to send traffic on the ports connected to the failed link.

When UDLD is enabled on the trunk ports on each device, the devices detect the failed link, disable the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Ports enabled for UDLD exchange proprietary health-check packets once every second (the keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for two more intervals. If the port still does not receive a health-check packet after waiting for three intervals, the port concludes that the link has failed and takes the port down.

Configuration considerations

- This feature is supported only on Ethernet ports.
- To configure UDLD on a trunk group, you must enable and configure the feature on each port of the group individually. Configuring UDLD on a trunk group primary port enables the feature on that port only.
- When UDLD is enabled on a trunk port, trunk threshold is not supported.
- Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.
- If MRP is also enabled on the device, Dell recommends that you set the MRP preforwarding time slightly higher than the default of 300 ms; for example, to 400 or 500 ms. Refer to [“Changing the hello and preforwarding times”](#) on page 213.

Enabling UDLD

To enable UDLD on a port, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#link-keepalive ethernet 1
```

To enable the feature on a trunk group, enter commands such as the following.

```
PowerConnect(config)#link-keepalive ethernet 1 ethernet 2
PowerConnect(config)#link-keepalive ethernet 3 ethernet 4
```

Syntax: [no] link-keepalive ethernet [portnum] ethernet [portnum]

These commands enable UDLD on ports 1 – 4.

```
PowerConnect(config)#link-keepalive ethernet 1 to 5 ethernet 7
```

Syntax: [no] link-keepalive ethernet [portlist]

Changing the Keepalive interval

By default, ports enabled for UDLD send a link health-check packet once every 500 ms. You can change the interval to a value from 1 – 60, where 1 is 100 ms, 2 is 200 ms, and so on. To change the interval, enter a command such as the following.

```
PowerConnect(config)#link-keepalive interval 3
```

Syntax: [no] link-keepalive interval <num>

The <num> parameter specifies how often the ports send a UDLD packet. You can specify from 1 – 60, in 100 ms increments. The default is 5 (500 ms).

Changing the Keepalive retries

By default, a port waits one second to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 – 64. To change the maximum number of attempts, enter a command such as the following.

```
PowerConnect(config)#link-keepalive retries 4
```

Syntax: [no] link-keepalive retries <num>

The <num> parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 – 64. The default is 7.

UDLD for tagged ports

The default implementation of UDLD sends the packets untagged, even across tagged ports. If the untagged UDLD packet is received by a third-party switch, that switch may reject the packet. As a result, UDLD may be limited only to devices, since UDLD may not function on third-party switches.

You can configure ports to send out UDLD control packets that are tagged with a specific VLAN ID as tagged UDLD control packets. This feature also enables third party switches to receive the control packets that are tagged with the specified VLAN.

To enable ports to receive and send UDLD control packets tagged with a specific VLAN ID, enter commands such as the following.

```
PowerConnect(config)#link-keepalive ethernet 18 vlan 22
```

This command enables UDLD on port 18 and allows UDLD control packet tagged with VLAN 22 to be received and sent on port 18.

Syntax: [no] link-keepalive ethernet <portnum> [vlan <vlan-ID>]

Enter the ID of the VLAN that the UDLD control packets can contain to be received and sent on the port. If a VLAN ID is not specified, then UDLD control packets are sent out of the port as untagged packets.

NOTE

You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.

Displaying UDLD information

This section describes the commands used to display information about a UDLD configuration.

Displaying information for all ports

To display UDLD information for all ports, enter the following command.

```
PowerConnect#show link-keepalive
Total link-keepalive enabled ports: 4
Keepalive Retries: 3    Keepalive Interval: 1 Sec.
```

Port	Physical Link	Logical Link	State	Link-vlan
1	up	up	FORWARDING	3
2	up	up	FORWARDING	
3	down	down	DISABLED	
4	up	down	DISABLED	

Syntax: show link-keepalive

TABLE 38 CLI display of UDLD information

This field...	Displays...
Total link-keepalive enabled ports	The total number of ports on which UDLD is enabled.
Keepalive Retries	The number of times a port will attempt the health check before concluding that the link is down.
Keepalive Interval	The number of seconds between health check packets.
Port	The port number.
Physical Link	The state of the physical link. This is the link between the port and the directly connected device.
Logical Link	The state of the logical link. This is the state of the link between this port and the port on the other end of the link.
State	The traffic state of the port.
Link-vlan	The ID of the tagged VLAN in the UDLD packet.

If a port is disabled by UDLD, the change also is indicated in the output of the **show interfaces brief** command. An example is given below.

```
PowerConnect#show interfaces brief

Port  Link State      Dupl Speed Trunk Tag Priori MAC          Name
1     Up   LK-DISABLE None None  None No  level0 00e0.52a9.bb00
2     Down None      None None  None No  level0 00e0.52a9.bb01
3     Down None      None None  None No  level0 00e0.52a9.bb02
4     Down None      None None  None No  level0 00e0.52a9.bb03
```

If the port was already down before you enabled UDLD for the port, the port state is listed as None.

Syntax: show interfaces brief

Displaying information for a single port

To display detailed UDLD information for a specific port, enter a command such as the following.

```
PowerConnect#show link-keepalive ethernet 1

Current State      : up          Remote MAC Addr   : 00e0.52d2.5100
Local Port         : 1            Remote Port       : 2
Local System ID    : e0927400  Remote System ID  : e0d25100
Packets sent       : 254        Packets received  : 255
Transitions        : 1          Link-vlan         : 100
Port blocking      : No         BM disabled       : No
```

Syntax: `show link-keepalive [ethernet<portnum>]`

TABLE 39 CLI display of detailed UDLD information

This field...	Displays...
Current State	The state of the logical link. This is the link between this port and the port on the other end of the link.
Remote MAC Addr	The MAC address of the port or device at the remote end of the logical link.
Local Port	The port number on this device.
Remote Port	The port number on the device at the remote end of the link.
Local System ID	A unique value that identifies this device. The ID can be used by Dell technical support for troubleshooting.
Remote System ID	A unique value that identifies the device at the remote end of the link.
Packets sent	The number of UDLD health-check packets sent on this port.
Packets received	The number of UDLD health-check packets received on this port.
Transitions	The number of times the logical link state has changed between up and down.
Port blocking	Information used by Dell technical support for troubleshooting.
Link-vlan	The ID of the tagged VLAN in the UDLD packet.
BM disabled	Information used by Dell technical support for troubleshooting.

The `show interface ethernet` command also displays the UDLD state for an individual port. In addition, the line protocol state listed in the first line will say “down” if UDLD has brought the port down. An example is given below.

```
PowerConnect#show interface ethernet 1
FastEthernet1 is down, line protocol is down, link keepalive is enabled
  Hardware is FastEthernet, address is 00e0.52a9.bbca (bia 00e0.52a9.bbca)
  Configured speed auto, actual unknown, configured duplex fdx, actual unknown
  Member of L2 VLAN ID 1, port is untagged, port state is DISABLED
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants, DMA received 0 packets
  19 packets output, 1216 bytes, 0 underruns
  Transmitted 0 broadcasts, 19 multicasts, 0 unicasts
  0 output errors, 0 collisions, DMA transmitted 19 packets
```

In this example, the port has been brought down by UDLD. Notice that in addition to the information in the first line, the port state on the fourth line of the display is listed as DISABLED.

Clearing UDLD statistics

To clear UDLD statistics, enter the following command.

```
PowerConnect#clear link-keepalive statistics
```

9 UDLD overview

Syntax: clear link-keepalive statistics

This command clears the Packets sent, Packets received, and Transitions counters in the **show link keepalive ethernet**<portnum> display.

Configuring Virtual LANs (VLANs)

VLAN overview

The following sections provide details about the VLAN types and features supported on the PowerConnect B-Series T124X family of switches.

Types of VLANs

This section describes the VLAN types supported on devices.

VLAN support on PowerConnect devices

The first software release for the PowerConnect B-Series T124X supports Layer 2 port-based VLANs only. A Layer 2 port-based VLAN is a set of physical ports that share a common, exclusive Layer 2 broadcast domain. The next section provides more details.

Layer 2 port-based VLANs

On all devices, you can configure port-based VLANs. A port-based VLAN is a subset of ports on a device that constitutes a Layer 2 broadcast domain.

By default, all the ports on a device are members of the default VLAN. Thus, all the ports on the device constitute a single Layer 2 broadcast domain. You can configure multiple port-based VLANs. When you configure a port-based VLAN, the device automatically removes the ports you add to the VLAN from the default VLAN.

You can configure up to 4094 port-based VLANs on a Layer 2 Switch or Layer 3 Switch. On both device types, valid VLAN IDs are 1 – 4095. You can configure up to the maximum number of VLANs within that ID range.

NOTE

If you want to use VLANs 4091 and 4092 as configurable VLANs, you can assign them to different VLAN IDs. For more information, refer to [“Assigning different VLAN IDs to reserved VLANs 4091 and 4092”](#) on page 265.

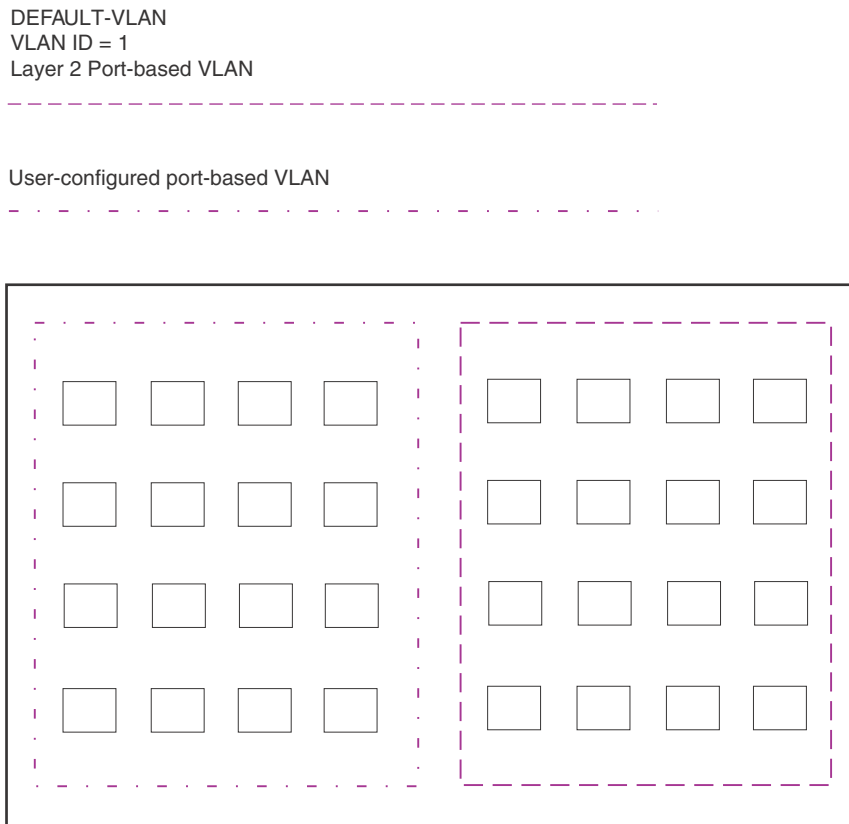
Each port-based VLAN can contain either tagged or untagged ports. A port cannot be a member of more than one port-based VLAN unless the port is tagged. **802.1Q tagging** allows the port to add a four-byte tag field, which contains the VLAN ID, to each packet sent on the port. You also can configure port-based VLANs that span multiple devices by tagging the ports within the VLAN. The tag enables each device that receives the packet to determine the VLAN the packet belongs to. 802.1Q tagging applies only to Layer 2 VLANs, not to Layer 3 VLANs.

Since each port-based VLAN is a separate Layer 2 broadcast domain, by default each VLAN runs a separate instance of the Spanning Tree Protocol (STP).

Layer 2 traffic is bridged within a port-based VLAN and Layer 2 broadcasts are sent to all the ports within the VLAN.

Figure 54 shows an example of a device on which a Layer 2 port-based VLAN has been configured.

FIGURE 54 Device containing user-defined Layer 2 port-based VLAN



When you add a port-based VLAN, the device removes all the ports in the new VLAN from DEFAULT-VLAN.

Layer 3 protocol-based VLANs

If you want some or all of the ports within a port-based VLAN to be organized according to Layer 3 protocol, you must configure a Layer 3 protocol-based VLAN within the port-based VLAN.

You can configure each of the following types of protocol-based VLAN within a port-based VLAN. All the ports in the Layer 3 VLAN must be in the same Layer 2 VLAN.

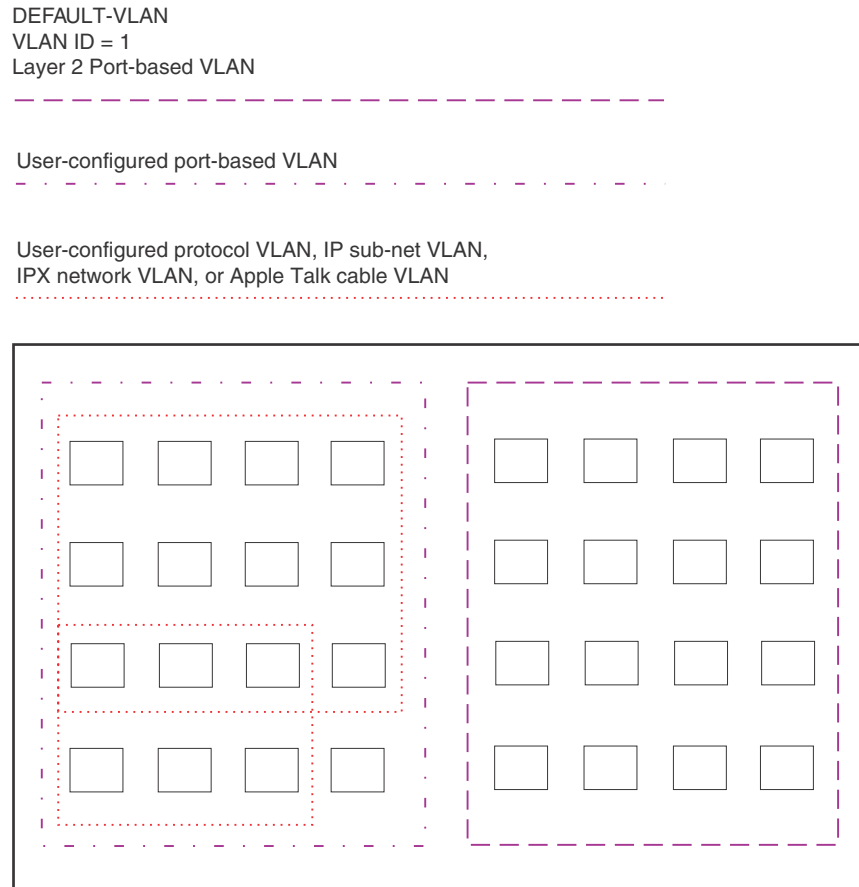
Layer 3 protocol-based VLANs are as follows:

- **AppleTalk** – The device sends AppleTalk broadcasts to all ports within the AppleTalk protocol VLAN.
- **IP** – The device sends IP broadcasts to all ports within the IP protocol VLAN.
- **IPv6** – The device sends IPv6 broadcasts to all ports within the IPv6 protocol VLAN.
- **IPX** – The device sends IPX broadcasts to all ports within the IPX protocol VLAN.
- **DECnet** – The device sends DECnet broadcasts to all ports within the DECnet protocol VLAN.

- **NetBIOS** – The device sends NetBIOS broadcasts to all ports within the NetBIOS protocol VLAN.
- **Other** – The device sends broadcasts for all protocol types other than those listed above to all ports within the VLAN.

Figure 55 shows an example of Layer 3 protocol VLANs configured within a Layer 2 port-based VLAN.

FIGURE 55 Layer 3 protocol VLANs within a Layer 2 port-based VLAN



You can add Layer 3 protocol VLANs or IP sub-net, IPX network, and AppleTalk cable VLANs to port-based VLANs.

Layer 3 VLANs cannot span Layer 2 port-based VLANs.

However, Layer 3 VLANs can overlap within a Layer 2 port-based VLAN.

Integrated Switch Routing (ISR)

The **Integrated Switch Routing (ISR)** feature enables VLANs configured on Layer 3 Switches to route Layer 3 traffic from one protocol VLAN or IP subnet, IPX network, or AppleTalk cable VLAN to another. Normally, to route traffic from one IP subnet, IPX network, or AppleTalk cable VLAN to another, you would need to forward the traffic to an external router. The VLANs provide Layer 3 broadcast domains for these protocols but do not in themselves provide routing services for these protocols. This is true even if the source and destination IP subnets, IPX networks, or AppleTalk cable ranges are on the same device.

ISR eliminates the need for an external router by allowing you to route between VLANs using virtual routing interfaces (ves). A **virtual routing interface** is a logical port on which you can configure Layer 3 routing parameters. You configure a separate virtual routing interface on each VLAN that you want to be able to route from or to. For example, if you configure two IP subnet VLANs on a Layer 3 Switch, you can configure a virtual routing interface on each VLAN, then configure IP routing parameters for the subnets. Thus, the Layer 3 Switch forwards IP subnet broadcasts within each VLAN at Layer 2 but routes Layer 3 traffic between the VLANs using the virtual routing interfaces.

NOTE

The Layer 3 Switch uses the lowest MAC address on the device (the MAC address of port 1) as the MAC address for all ports within all virtual routing interfaces you configure on the device.

The routing parameters and the syntax for configuring them are the same as when you configure a physical interface for routing. The logical interface allows the Layer 3 Switch to internally route traffic between the protocol-based VLANs without using physical interfaces.

All the ports within a protocol-based VLAN must be in the same port-based VLAN. The protocol-based VLAN cannot have ports in multiple port-based VLANs, unless the ports in the port-based VLAN to which you add the protocol-based VLAN are 802.1Q tagged.

You can configure multiple protocol-based VLANs within the same port-based VLAN. In addition, a port within a port-based VLAN can belong to multiple protocol-based VLANs of the same type or different types. For example, if you have a port-based VLAN that contains ports 1 – 10, you can configure port 5 as a member of an AppleTalk protocol VLAN, an IP protocol VLAN, and an IPX protocol VLAN, and so on.

IP subnet, IPX network, and AppleTalk cable VLANs

The protocol-based VLANs described in the previous section provide separate protocol broadcast domains for specific protocols. For IP, IPX, and AppleTalk, you can provide more granular broadcast control by instead creating the following types of VLAN:

- **IP subnet VLAN** – An IP subnet broadcast domain for a specific IP subnet.
- **IPX network VLAN** – An IPX network broadcast domain for a specific IPX network.
- **AppleTalk cable VLAN** – An AppleTalk broadcast domain for a specific cable range.

You can configure these types of VLANs on Layer 3 Switches only. The Layer 3 Switch sends broadcasts for the IP subnet, IPX network, or AppleTalk cable range to all ports within the IP subnet, IPX network, or AppleTalk cable VLAN at Layer 2.

The Layer 3 Switch routes packets between VLANs at Layer 3. To configure an IP subnet, IPX network, or AppleTalk cable VLAN to route, you must add a virtual routing interface to the VLAN, then configure the appropriate routing parameters on the virtual routing interface.

NOTE

The Layer 3 Switch routes packets between VLANs of the same protocol. The Layer 3 Switch cannot route from one protocol to another.

NOTE

IP subnet VLANs are not the same thing as IP protocol VLANs. An IP protocol VLAN sends all IP broadcasts on the ports within the IP protocol VLAN. An IP subnet VLAN sends only the IP subnet broadcasts for the subnet of the VLAN. You cannot configure an IP protocol VLAN and an IP subnet VLAN within the same port-based VLAN.

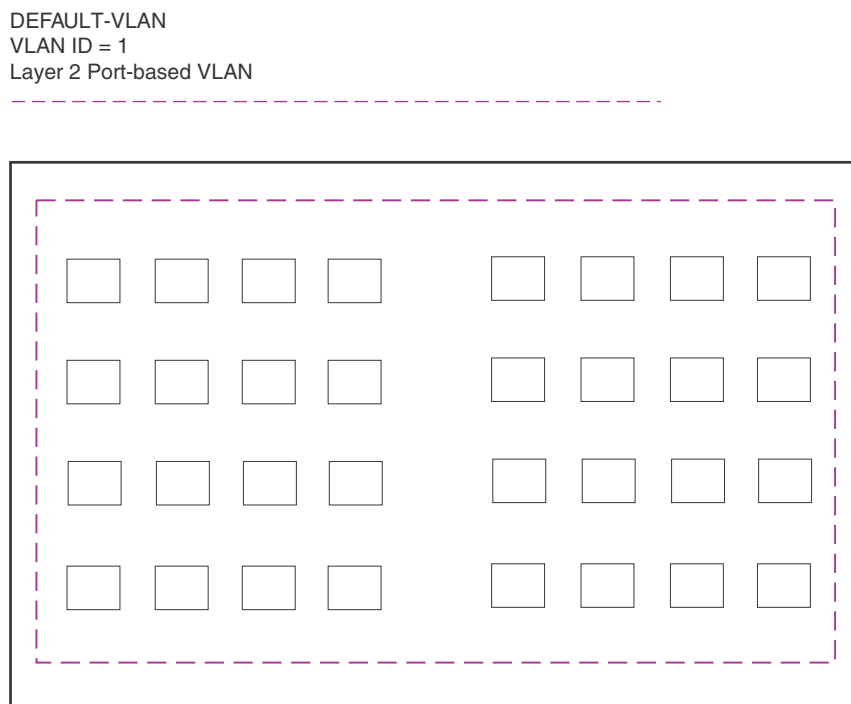
This note also applies to IPX protocol VLANs and IPX network VLANs, and to AppleTalk protocol VLANs and AppleTalk cable VLANs.

Default VLAN

By default, all the ports on a PowerConnect device are in a single port-based VLAN. This VLAN is called the DEFAULT-VLAN and is VLAN number 1.

Figure 56 shows an example of the default Layer 2 port-based VLAN.

FIGURE 56 Default Layer 2 port-based VLAN



By default, all ports belong to a single port-based VLAN, DEFAULT-VLAN. Thus, all ports belong to a single Layer 2 broadcast domain.

When you configure a port-based VLAN, one of the configuration items you provide is the ports that are in the VLAN. When you configure the VLAN, the device automatically removes the ports that you place in the VLAN from DEFAULT-VLAN. By removing the ports from the default VLAN, the device ensures that each port resides in only one Layer 2 broadcast domain.

NOTE

Information for the default VLAN is available only after you define another VLAN.

Some network configurations may require that a port be able to reside in two or more Layer 2 broadcast domains (port-based VLANs). In this case, you can enable a port to reside in multiple port-based VLANs by tagging the port. See the following section.

If your network requires that you use VLAN ID 1 for a user-configured VLAN, you can reassign the default VLAN to another valid VLAN ID. Refer to [“Assigning a different VLAN ID to the default VLAN”](#) on page 265.

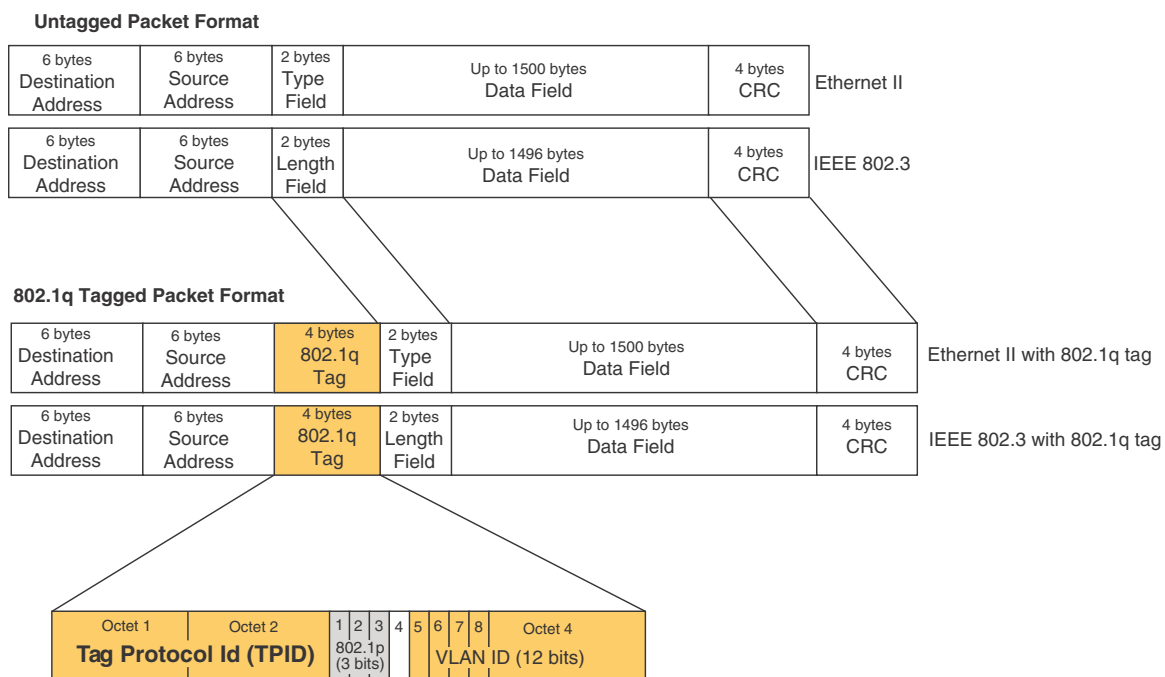
802.1Q tagging

802.1Q tagging is an IEEE standard that allows a networking device to add information to a Layer 2 packet in order to identify the VLAN membership of the packet. devices tag a packet by adding a four-byte tag to the packet. The tag contains the tag value, which identifies the data as a tag, and also contains the VLAN ID of the VLAN from which the packet is sent.

- The default tag value is 8100 (hexadecimal). This value comes from the 802.1Q specification. You can change this tag value on a global basis on devices if needed to be compatible with other vendors' equipment.
- The VLAN ID is determined by the VLAN on which the packet is being forwarded.

[Figure 57](#) shows the format of packets with and without the 802.1Q tag. The tag format is vendor-specific. To use the tag for VLANs configured across multiple devices, make sure all the devices support the same tag format.

FIGURE 57 Packet containing a 802.1Q VLAN tag

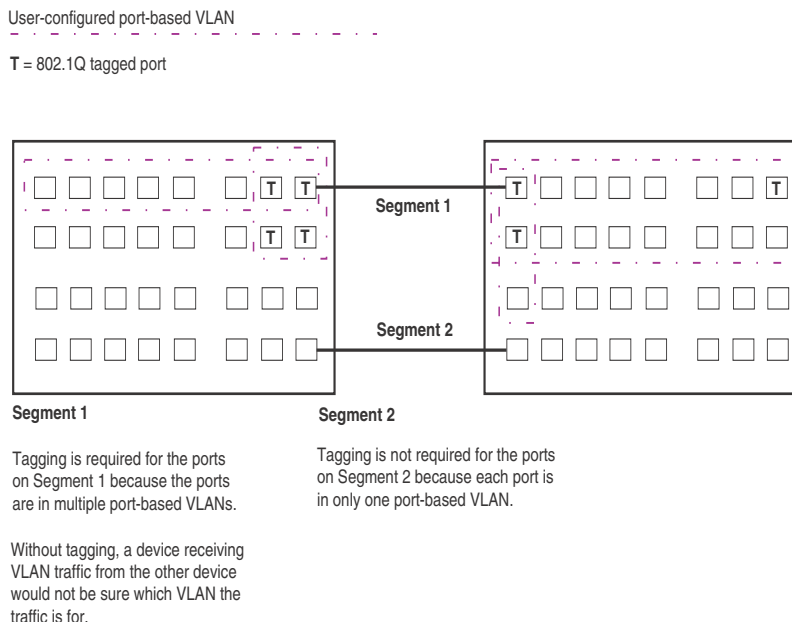


If you configure a VLAN that spans multiple devices, you need to use tagging only if a port connecting one of the devices to the other is a member of more than one port-based VLAN. If a port connecting one device to the other is a member of only a single port-based VLAN, tagging is not required.

If you use tagging on multiple devices, each device must be configured for tagging and must use the same tag value. In addition, the implementation of tagging must be compatible on the devices. The tagging on all devices is compatible with other devices.

Figure 58 shows an example of two devices that have the same Layer 2 port-based VLANs configured across them. Notice that only one of the VLANs requires tagging.

FIGURE 58 VLANs configured across multiple devices



Support for 802.1Q-in-Q tagging

Devices provide finer granularity for configuring 802.1Q tagging, enabling you to configure 802.1Q tag-types on a group of ports, thereby enabling the creation of two identical 802.1Q tags (802.1Q-in-Q tagging) on a single device. This enhancement improves SAV interoperability between devices and other vendors' devices that support the 802.1Q tag-types, but are not very flexible with the tag-types they accept.

For example applications and configuration details, refer to [“Configuring 802.1Q-in-Q tagging”](#) on page 296.

Spanning Tree Protocol (STP)

The default state of STP depends on the device type:

- STP is disabled by default on Layer 3 Switches.
- STP is enabled by default on Layer 2 Switches.

Also by default, each port-based VLAN has a separate instance of STP. Thus, when STP is globally enabled, each port-based VLAN on the device runs a separate spanning tree.

You can enable or disable STP on the following levels:

- **Globally** – Affects all ports on the device.

NOTE

If you configure a port-based VLAN on the device, the VLAN has the same STP state as the default STP state on the device. Thus, on Layer 2 Switches, new VLANs have STP enabled by default. On Layer 3 Switches, new VLANs have STP disabled by default. You can enable or disable STP in each VLAN separately. In addition, you can enable or disable STP on individual ports.

- **Port-based VLAN** – Affects all ports within the specified port-based VLAN.

STP is a Layer 2 protocol. Thus, you cannot enable or disable STP for individual protocol VLANs or for IP subnet, IPX network, or AppleTalk cable VLANs. The STP state of a port-based VLAN containing these other types of VLANs determines the STP state for all the Layer 2 broadcasts within the port-based VLAN. This is true even though Layer 3 protocol broadcasts are sent on Layer 2 within the VLAN.

It is possible that STP will block one or more ports in a protocol VLAN that uses a virtual routing interface to route to other VLANs. For IP protocol and IP subnet VLANs, even though some of the physical ports of the virtual routing interface are blocked, the virtual routing interface can still route so long as at least one port in the virtual routing interface protocol VLAN is not blocked by STP.

If you enable Single STP (SSTP) on the device, the ports in all VLANs on which STP is enabled become members of a single spanning tree. The ports in VLANs on which STP is disabled are excluded from the single spanning tree.

For more information, refer to [Chapter 6, “Configuring Spanning Tree Protocol \(STP\) Related Features”](#).

Virtual routing interfaces

A virtual routing interface is a logical routing interface that devices use to route Layer 3 protocol traffic between protocol VLANs.

Devices send Layer 3 traffic at Layer 2 within a protocol VLAN. However, Layer 3 traffic from one protocol VLAN to another must be routed.

If you want the device to be able to send Layer 3 traffic from one protocol VLAN to another, you must configure a virtual routing interface on each protocol VLAN, then configure routing parameters on the virtual routing interfaces. For example, to enable a Layer 3 Switch to route IP traffic from one IP subnet VLAN to another, you must configure a virtual routing interface on each IP subnet VLAN, then configure the appropriate IP routing parameters on each of the virtual routing interfaces.

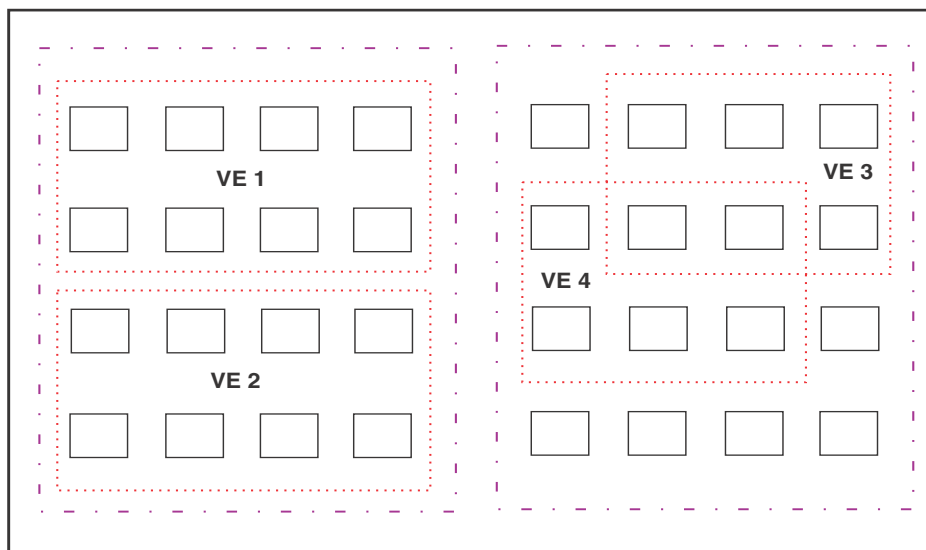
[Figure 59](#) shows an example of Layer 3 protocol VLANs that use virtual routing interfaces for routing.

FIGURE 59 Use virtual routing interfaces for routing between Layer 3 protocol VLANs

User-configured port-based VLAN

User-configured protocol VLAN, IP sub-net VLAN, IPX network VLAN, or AppleTalk cable VLAN

VE = virtual interface
 (“VE” stands for “Virtual Ethernet”)



Layer 2 and Layer 3 traffic within a VLAN is bridged at Layer 2.

Layer 3 traffic between protocol VLANs is routed using virtual interfaces (VE). To route to one another, each protocol VLAN must have a virtual interface.

VLAN and virtual routing interface groups

PowerConnect devices support the configuration of VLAN groups. To simplify configuration, you can configure VLAN groups and virtual routing interface groups. When you create a VLAN group, the VLAN parameters you configure for the group apply to all the VLANs within the group. Additionally, you can easily associate the same IP subnet interface with all the VLANs in a group by configuring a virtual routing interface group with the same ID as the VLAN group.

For configuration information, refer to [“Configuring VLAN groups and virtual routing interface groups”](#) on page 285.

Dynamic, static, and excluded port membership

When you add ports to a protocol VLAN, IP subnet VLAN, IPX network VLAN, or AppleTalk cable VLAN, you can add them dynamically or statically:

- Dynamic ports
- Static ports

You also can explicitly exclude ports.

Dynamic ports

Dynamic ports are added to a VLAN when you create the VLAN. However, if a dynamically added port does not receive any traffic for the VLAN protocol within ten minutes, the port is removed from the VLAN. However, the port remains a candidate for port membership. Thus, if the port receives traffic for the VLAN protocol, the device adds the port back to the VLAN.

After the port is added back to the VLAN, the port can remain an active member of the VLAN up to 20 minutes without receiving traffic for the VLAN protocol. If the port ages out, it remains a candidate for VLAN membership and is added back to the VLAN when the VLAN receives protocol traffic. At this point, the port can remain in the VLAN up to 20 minutes without receiving traffic for the VLAN protocol, and so on.

Unless you explicitly add a port statically or exclude a port, the port is a dynamic port and thus can be an active member of the VLAN, depending on the traffic it receives.

NOTE

You cannot configure dynamic ports in an AppleTalk cable VLAN. The ports in an AppleTalk cable VLAN must be static. However, ports in an AppleTalk protocol VLAN can be dynamic or static.


Figure 60 shows an example of a VLAN with dynamic ports. Dynamic ports not only join and leave the VLAN according to traffic, but also allow some broadcast packets of the specific protocol to “leak” through the VLAN. Refer to “Broadcast leaks” on page 263.

FIGURE 60 VLAN with dynamic ports—all ports are active when you create the VLAN

A = active port

C = candidate port

When you add ports dynamically, all the ports are added when you add the VLAN.




SUBNET Ports in a new protocol VLAN that do not receive traffic for the VLAN protocol age out after 10 minutes and become candidate ports. Figure 61 shows what happens if a candidate port receives traffic for the VLAN protocol.

FIGURE 61 VLAN with dynamic ports—candidate ports become active again if they receive protocol traffic

Ports that time out remain candidates for membership in the VLAN and become active again if they receive traffic for the VLAN's protocol, IP sub-net, IPX network, or AppleTalk cable range.

When a candidate port rejoins a VLAN, the timeout for that port becomes 20 minutes. Thus, the port remains an active member of the VLAN even if it does not receive traffic for 20 minutes. After that, the port becomes a candidate port again.



Static ports

Static ports are permanent members of the protocol VLAN. The ports remain active members of the VLAN regardless of whether the ports receive traffic for the VLAN protocol. You must explicitly identify the port as a static port when you add it to the VLAN. Otherwise, the port is dynamic and is subject to aging out.

Excluded ports

If you want to prevent a port in a port-based VLAN from ever becoming a member of a protocol, IP subnet, IPX network, or AppleTalk cable VLAN configured in the port-based VLAN, you can explicitly exclude the port. You exclude the port when you configure the protocol, IP subnet, IPX network, or AppleTalk cable VLAN.

Excluded ports do not leak broadcast packets. Refer to “Broadcast leaks” on page 263.

Broadcast leaks

A dynamic port becomes a member of a Layer 3 protocol VLAN when traffic from the VLAN's protocol is received on the port. After this point, the port remains an active member of the protocol VLAN, unless the port does not receive traffic from the VLAN's protocol for 20 minutes. If the port does not receive traffic for the VLAN's protocol for 20 minutes, the port ages out and is no longer an active member of the VLAN.

To enable a host that has been silent for awhile to send and receive packets, the dynamic ports that are currently members of the Layer 3 protocol VLAN "leak" Layer 3 broadcast packets to the ports that have aged out. When a host connected to one of the aged out ports responds to a leaked broadcast, the port is added to the protocol VLAN again.

To "leak" Layer 3 broadcast traffic, an active port sends 1/8th of the Layer 3 broadcast traffic to the inactive (aged out) ports.

Static ports do not age out and do not leak broadcast packets.

Super aggregated VLANs

You can aggregate multiple VLANs within another VLAN. This feature allows you to construct Layer 2 paths and channels. This feature is particularly useful for Virtual Private Network (VPN) applications in which you need to provide a private, dedicated Ethernet connection for an individual client to transparently reach its subnet across multiple networks.

For an application example and configuration information, refer to ["Configuring super aggregated VLANs"](#) on page 289.

Trunk group ports and VLAN membership

A trunk group is a set of physical ports that are configured to act as a single physical interface. Each trunk group port configuration is based on the configuration of the lead port, which is the lowest numbered port in the group.

If you add a trunk group lead port to a VLAN, all of the ports in the trunk group become members of that VLAN.

Routing between VLANs

Layer 3 Switches can locally route IP, IPX, and Appletalk between VLANs defined within a single router. All other routable protocols or protocol VLANs (for example, DecNet) must be routed by another external router capable of routing the protocol.

Virtual routing interfaces (Layer 3 Switches only)

You need to configure virtual routing interfaces if an IP, IPX, or Appletalk protocol VLAN, IP subnet VLAN, AppleTalk cable VLAN, or IPX network VLAN needs to route protocols to another port-based VLAN on the same router. A virtual routing interface can be associated with the ports in only a single port-based VLAN. Virtual router interfaces must be defined at the highest level of the VLAN hierarchy.

If you do not need to further partition the port-based VLAN by defining separate Layer 3 VLANs, you can define a single virtual routing interface at the port-based VLAN level and enable IP, IPX, and Appletalk routing on a single virtual routing interface.

Some configurations may require simultaneous switching and routing of the same single protocol across different sets of ports on the same router. When IP, IPX, or Appletalk routing is enabled on a Layer 3 Switch, you can route these protocols on specific interfaces while bridging them on other interfaces. In this scenario, you can create two separate backbones for the same protocol, one bridged and one routed.

To bridge IP, IPX, or Appletalk at the same time these protocols are being routed, you need to configure an IP protocol, IP subnet, IPX protocol, IPX network, or Appletalk protocol VLAN and not assign a virtual routing interface to the VLAN. Packets for these protocols are bridged or switched at Layer 2 across ports on the router that are included in the Layer 3 VLAN. If these VLANs are built within port-based VLANs, they can be tagged across a single set of backbone fibers to create separate Layer 2 switched and Layer 3 routed backbones for the same protocol on a single physical backbone.

Routing between VLANs using virtual routing interfaces (Layer 3 Switches only)

Dell calls the ability to route between VLANs with virtual routing interfaces **Integrated Switch Routing (ISR)**. There are some important concepts to understand before designing an ISR backbone.

Virtual router interfaces can be defined on port-based, IP protocol, IP subnet, IPX protocol, IPX network, AppleTalk protocol, and AppleTalk cable VLANs.

To create any type of VLAN on a Layer 3 Switch, Layer 2 forwarding must be enabled. When Layer 2 forwarding is enabled, the Layer 3 Switch becomes a Switch on all ports for all non-routable protocols.

If the router interfaces for IP, IPX, or AppleTalk are configured on physical ports, then routing occurs independent of the Spanning Tree Protocol (STP). However, if the router interfaces are defined for any type VLAN, they are virtual routing interfaces and are subject to the rules of STP.

If your backbone consists of virtual routing interfaces all within the same STP domain, it is a bridged backbone, not a routed one. This means that the set of backbone interfaces that are blocked by STP will be blocked for routed protocols as well. The routed protocols will be able to cross these paths only when the STP state of the link is FORWARDING. This problem is easily avoided by proper network design.

When designing an ISR network, pay attention to your use of virtual routing interfaces and the spanning-tree domain. If Layer 2 switching of your routed protocols (IP, IPX, AppleTalk) is not required across the backbone, then the use of virtual routing interfaces can be limited to edge switch ports within each router. Full backbone routing can be achieved by configuring routing on each physical interface that connects to the backbone. Routing is independent of STP when configured on a physical interface.

If your ISR design requires that you switch IP, IPX, or Appletalk at Layer 2 while simultaneously routing the same protocols over a single backbone, then create multiple port-based VLANs and use VLAN tagging on the backbone links to separate your Layer 2 switched and Layer 3 routed networks.

There is a separate STP domain for each port-based VLAN. Routing occurs independently across port-based VLANs or STP domains. You can define each end of each backbone link as a separate tagged port-based VLAN. Routing will occur independently across the port-based VLANs. Because each port-based VLAN STP domain is a single point-to-point backbone connection, you are guaranteed to never have an STP loop. STP will never block the virtual router interfaces within the tagged port-based VLAN, and you will have a fully routed backbone.

Dynamic port assignment (Layer 2 Switches and Layer 3 Switches)

All Switch ports are dynamically assigned to any Layer 3 VLAN on Layer 2 Switches and any non-routable VLAN on Layer 3 Switches. To maintain explicit control of the VLAN, you can explicitly exclude ports when configuring any Layer 3 VLAN on a Layer 2 Switch or any non-routable Layer 3 VLAN on a Layer 3 Switch.

If you do not want the ports to have dynamic membership, you can add them statically. This eliminates the need to explicitly exclude the ports that you do not want to participate in a particular Layer 3 VLAN.

Assigning a different VLAN ID to the default VLAN

When you enable port-based VLANs, all ports in the system are added to the default VLAN. By default, the default VLAN ID is “VLAN 1”. The default VLAN is not configurable. If you want to use the VLAN ID “VLAN 1” as a configurable VLAN, you can assign a different VLAN ID to the default VLAN.

To reassign the default VLAN to a different VLAN ID, enter the following command.

```
PowerConnect(config)# default-vlan-id 4095
```

Syntax: [no] default-vlan-d <vlan-id>

You must specify a valid VLAN ID that is not already in use. For example, if you have already defined VLAN 10, do not try to use “10” as the new VLAN ID for the default VLAN. Valid VLAN IDs are numbers from 1 – 4095.

NOTE

Changing the default VLAN name does not change the properties of the default VLAN. Changing the name allows you to use the VLAN ID “1” as a configurable VLAN.

Assigning different VLAN IDs to reserved VLANs 4091 and 4092

VLAN 4094 is reserved for use by Single STP. If you want to use VLANs 4091 and 4092 as configurable VLANs, you can assign them to different VLAN IDs.

For example, to reassign reserved VLAN 4091 to VLAN 10, enter the following commands.

```
PowerConnect(config)# reserved-vlan-map vlan 4091 new-vlan 10
Reload required. Please write memory and then reload or power cycle.
PowerConnect(config)# write mem
PowerConnect(config)# exit
PowerConnect# reload
```

NOTE

You must save the configuration (write mem) and reload the software to place the change into effect.

The above configuration changes the VLAN ID of 4091 to 10. After saving the configuration and reloading the software, you can configure VLAN 4091 as you would any other VLAN.

Syntax: [no] reserved-vlan-map vlan 4091 | 4092 new-vlan <vlan-id>

For <vlan-id>, enter a valid VLAN ID that is not already in use. For example, if you have already defined VLAN 20, do not try to use “20 as the new VLAN ID. Valid VLAN IDs are numbers from 1 – 4090, 4093, and 4095. VLAN ID 4094 is reserved for use by the Single Spanning Tree feature.

Viewing reassigned VLAN IDs for reserved VLANs 4091 and 4092

To view the assigned VLAN IDs for reserved VLANs 4091 and 4092, use the **show reserved-vlan-map** command. The reassigned VLAN IDs also display in the output of the **show running-config** and **show config** commands.

The following shows example output for the **show reserved-vlan-map** command.

```
PowerConnect # show reserved-vlan-map
Reserved Purpose      Default      Re-assign    Current
CPU VLAN              4091        10           10
All Ports VLAN       4092        33           33
```

Syntax: show reserved-vlan-map

The following table defines the fields in the output of the **show reserved-vlan-map** command.

TABLE 40 Output of the **show reserved-vlan-map** command

This field	Displays
Reserved Purpose	Describes for what the VLAN is reserved. Note that the description is for Dell internal VLAN management.
Default	The default VLAN ID of the reserved VLAN.
Re-assign	The VLAN ID to which the reserved VLAN was reassigned. ¹
Current	The current VLAN ID for the reserved VLAN. ¹

1. If you reassign a reserved VLAN without saving the configuration and reloading the software, the reassigned VLAN ID will display in the **Re-assign** column. However, the previously configured or default VLAN ID will display in the **Current** column until the configuration is saved and the device reloaded.

Assigning trunk group ports

When a “lead” trunk group port is assigned to a VLAN, all other members of the trunk group are automatically added to that VLAN. A lead port is the first port of a trunk group port range; for example, “1” in 1 – 4 or “5” in

5 – 8. Refer to “Trunk group rules” on page 312 for more information.

Configuring port-based VLANs

Port-based VLANs allow you to provide separate spanning tree protocol (STP) domains or broadcast domains on a port-by-port basis.

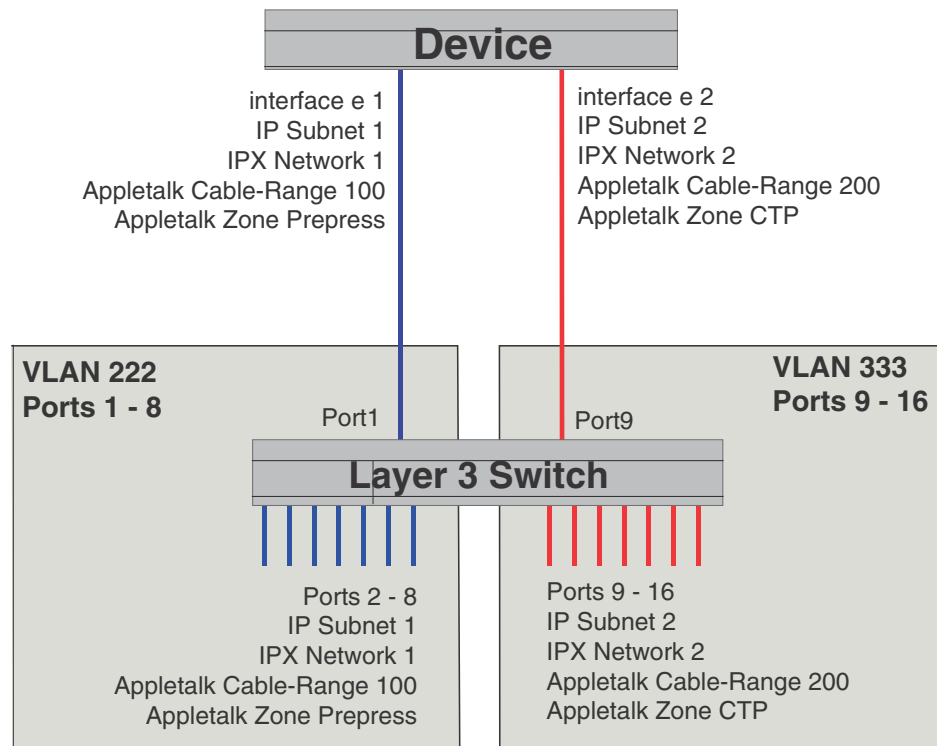
This section describes how to perform the following tasks for port-based VLANs using the CLI:

- Create a VLAN
- Delete a VLAN
- Modify a VLAN
- Change a VLAN priority
- Enable or disable STP on the VLAN

Example 1

Figure 62 shows a simple port-based VLAN configuration using a single Layer 2 Switch. All ports within each VLAN are untagged. One untagged port within each VLAN is used to connect the Layer 2 Switch to a Layer 3 Switch for Layer 3 connectivity between the two port-based VLANs.

FIGURE 62 Port-based VLANs 222 and 333




To create the two port-based VLANs shown in Figure 62, enter the following commands.

```
PowerConnect(config)# vlan 222 by port
PowerConnect(config-vlan-222)# untag e 1 to 8
PowerConnect(config-vlan-222)# vlan 333 by port
PowerConnect(config-vlan-333)# untag e 9 to 16
```

Syntax: `vlan <vlan-id> by port`

Syntax: `untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]`

Example 2



To configure the Port-based VLANs on the Layer 2 Switches in [Figure Syntax](#), use the following method.

Configuring device-A

Enter the following commands to configure device-A.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# hostname PowerConnect-A
PowerConnect-A(config)# vlan 2 name BROWN
PowerConnect-A(config-vlan-2)# untag ethernet 1 to 4 ethernet 17
PowerConnect-A(config-vlan-2)# tag ethernet 25 to 26
PowerConnect-A(config-vlan-2)# spanning-tree
PowerConnect-A(config-vlan-2)# vlan 3 name GREEN
PowerConnect-A(config-vlan-3)# untag ethernet 5 to 8 ethernet 18
PowerConnect-A(config-vlan-3)# tag ethernet 25 to 26
PowerConnect-A(config-vlan-3)# spanning-tree
PowerConnect-A(config-vlan-3)# vlan 4 name BLUE
PowerConnect-A(config-vlan-4)# untag ethernet 9 to 12 ethernet 19
PowerConnect-A(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-A(config-vlan-4)# spanning-tree
PowerConnect-A(config-vlan-4)# spanning-tree priority 500
PowerConnect-A(config-vlan-4)# vlan 5 name RED
PowerConnect-A(config-vlan-5)# untag ethernet 13 to 16 ethernet 20
PowerConnect-A(config-vlan-5)# tag ethernet 25 to 26
```

```
PowerConnect-A(config-vlan-5)# spanning-tree
PowerConnect-A(config-vlan-5)# spanning-tree priority 500
PowerConnect-A(config-vlan-5)# end
PowerConnect-A# write memory
```

Configuring device-B

Enter the following commands to configure device-B.

```
PowerConnect> en
PowerConnect# configure terminal
PowerConnect(config)# hostname PowerConnect-B
PowerConnect-B(config)# vlan 2 name BROWN
PowerConnect-B(config-vlan-2)# untag ethernet 1 to 4
PowerConnect-B(config-vlan-2)# tag ethernet 25 to 26
PowerConnect-B(config-vlan-2)# spanning-tree
PowerConnect-B(config-vlan-2)# spanning-tree priority 500
PowerConnect-B(config-vlan-2)# vlan 3 name GREEN
PowerConnect-B(config-vlan-3)# untag ethernet 5 to 8
PowerConnect-B(config-vlan-3)# tag ethernet 25 to 26
PowerConnect-B(config-vlan-3)# spanning-tree
PowerConnect-B(config-vlan-3)# spanning-tree priority 500
PowerConnect-B(config-vlan-3)# vlan 4 name BLUE
PowerConnect-B(config-vlan-4)# untag ethernet 9 to 12
PowerConnect-B(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-B(config-vlan-4)# vlan 5 name RED
PowerConnect-B(config-vlan-5)# untag ethernet 13 to 16
PowerConnect-B(config-vlan-5)# tag ethernet 25 to 26
PowerConnect-B(config-vlan-5)# end
PowerConnect-B# write memory
```

Configuring device-C

Enter the following commands to configure device-C.

```
PowerConnect> en
PowerConnect# configure terminal
PowerConnect(config)# hostname PowerConnect-C
PowerConnect-C(config)# vlan 2 name BROWN
PowerConnect-C(config-vlan-2)# untag ethernet 1 to 4
PowerConnect-C(config-vlan-2)# tag ethernet 25 to 26
PowerConnect-C(config-vlan-2)# vlan 3 name GREEN
PowerConnect-C(config-vlan-3)# untag ethernet 5 to 8
PowerConnect-C(config-vlan-3)# tag ethernet 25 to 26
PowerConnect-C(config-vlan-3)# vlan 4 name BLUE
PowerConnect-C(config-vlan-4)# untag ethernet 9 to 12
PowerConnect-C(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-C(config-vlan-4)# vlan 5 name RED
PowerConnect-C(config-vlan-5)# untag ethernet 13 to 16
PowerConnect-C(config-vlan-5)# tag ethernet 25 to 26
PowerConnect-C(config-vlan-5)# end
PowerConnect-C# write memory
```

Syntax: `vlan <vlan-id> by port`

Syntax: `untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]`

Syntax: `tagged ethernet<portnum> [to <portnum> | ethernet <portnum>]`

Syntax: `[no] spanning-tree`

Syntax: `spanning-tree [ethernet<portnum> path-cost <value> priority <value>] forward-delay <value> hello-time <value> maximum-age <time> priority <value>`

Modifying a port-based VLAN

You can make the following modifications to a port-based VLAN:

- Add or delete a VLAN port.
- Enable or disable STP.

Removing a port-based VLAN

Suppose you want to remove VLAN 5 from the example in [Figure Syntax](#). To do so, use the following procedure.

1. Access the global CONFIG level of the CLI on device-A by entering the following commands.

```
PowerConnect-A> enable
No password has been assigned yet...
PowerConnect-A# configure terminal
PowerConnect-A(config)#
```

2. Enter the following command.

```
PowerConnect-A(config)# no vlan 5
PowerConnect-A(config)#
```

3. Enter the following commands to exit the CONFIG level and save the configuration to the system-config file on flash memory.

```
PowerConnect-A(config)#
PowerConnect-A(config)# end
PowerConnect-A# write memory
PowerConnect-A#
```

4. Repeat steps 1 – 3 on device-B.

Syntax: `no vlan <vlan-id> by port`

Removing a port from a VLAN

Suppose you want to remove port 11 from VLAN 4 on device-A shown in [Figure Syntax](#). To do so, use the following procedure.

1. Access the global CONFIG level of the CLI on device-A by entering the following command.

```
PowerConnect-A> enable
No password has been assigned yet...
PowerConnect-A# configure terminal
PowerConnect-A(config)#
```

2. Access the level of the CLI for configuring port-based VLAN 4 by entering the following command.

```
PowerConnect-A(config)#
PowerConnect-A(config)# vlan 4
PowerConnect-A(config-vlan-4)#
```

3. Enter the following commands.

```
PowerConnect-A(config-vlan-4)#
PowerConnect-A(config-vlan-4)# no untag ethernet 11
deleted port ethe 11 from port-vlan 4.
PowerConnect-A(config-vlan-4)#
```

4. Enter the following commands to exit the VLAN CONFIG mode and save the configuration to the system-config file on flash memory.

```
PowerConnect-A(config-vlan-4)#
PowerConnect-A(config-vlan-4)# end
PowerConnect-A# write memory
```

You can remove all the ports from a port-based VLAN without losing the rest of the VLAN configuration. However, you cannot configure an IP address on a virtual routing interface unless the VLAN contains ports. If the VLAN has a virtual routing interface, the virtual routing interface IP address is deleted when the ports associated with the interface are deleted. The rest of the VLAN configuration is retained.

Enable spanning tree on a VLAN

The spanning tree bridge and port parameters are configurable using one CLI command set at the Global Configuration Level of each Port-based VLAN. Suppose you want to enable the IEEE 802.1D STP across VLAN 3. To do so, use the following method.

NOTE

When port-based VLANs are not operating on the system, STP is set on a system-wide level at the global CONFIG level of the CLI.

1. Access the global CONFIG level of the CLI on device-A by entering the following commands.

```
PowerConnect-A> enable
No password has been assigned yet...
PowerConnect-A# configure terminal
PowerConnect-A(config)#
```

2. Access the level of the CLI for configuring port-based VLAN 3 by entering the following command.

```
PowerConnect-A(config)#
PowerConnect-A(config)# vlan 3
PowerConnect-A(config-vlan-3)#
```

3. From VLAN 3 configuration level of the CLI, enter the following command to enable STP on all tagged and untagged ports associated with VLAN 3.

```
PowerConnect-B(config-vlan-3)#
PowerConnect-B(config-vlan-3)# spanning-tree
PowerConnect-B(config-vlan-3)#
```

4. Enter the following commands to exit the VLAN CONFIG mode and save the configuration to the system-config file on flash memory.

```
PowerConnect-B(config-vlan-3)#
PowerConnect-B(config-vlan-3)# end
PowerConnect-B# write memory
PowerConnect-B#
```

5. Repeat steps 1 – 4 on device-B.

NOTE

You do not need to configure values for the STP parameters. All parameters have default values as noted below. Additionally, all values will be globally applied to all ports on the system or on the port-based VLAN for which they are defined.

To configure a specific path-cost or priority value for a given port, enter those values using the key words in the brackets [] shown in the syntax summary below. If you do not want to specify values for any given port, this portion of the command is not required.

Syntax: `vlan <vlan-id> by port`

Syntax: `[no] spanning-tree`

Syntax: `spanning-tree [ethernet <portnum> path-cost <value> priority <value>] forward-delay <value> hello-time <value> maximum-age <time> priority <value>`

Bridge STP parameters (applied to all ports within a VLAN):

- Forward Delay – the period of time a bridge will wait (the listen and learn period) before forwarding data packets. Possible values: 4 – 30 seconds. Default is 15.
- Maximum Age – the interval a bridge will wait for receipt of a hello packet before initiating a topology change. Possible values: 6 – 40 seconds. Default is 20.
- Hello Time – the interval of time between each configuration BPDU sent by the root bridge. Possible values: 1 – 10 seconds. Default is 2.
- Priority – a parameter used to identify the root bridge in a network. The bridge with the lowest value has the highest priority and is the root. Possible values: 1 – 65,535. Default is 32,678.

Port parameters (applied to a specified port within a VLAN):

- Path Cost – a parameter used to assign a higher or lower path cost to a port. Possible values: 1 – 65535. Default is (1000/Port Speed) for Half-Duplex ports and is (1000/Port Speed)/2 for Full-Duplex ports.
- Priority – value determines when a port will be rerouted in relation to other ports. Possible values: 0 – 255. Default is 128.

Configuring IP subnet, IPX network and protocol-based VLANs

Protocol-based VLANs provide the ability to define separate broadcast domains for several unique Layer 3 protocols within a single Layer 2 broadcast domain. Some applications for this feature might include security between departments with unique protocol requirements. This feature enables you to limit the amount of broadcast traffic end-stations, servers, and routers need to accept.

Configuration example


Suppose you want to create five separate Layer 3 broadcast domains within a single Layer 2 STP broadcast domain:

- Three broadcast domains, one for each of three separate IP subnets
- One for IPX Network 1
- One for the Appletalk protocol

Also suppose you want a single router interface to be present within all of these separate broadcast domains, without using IEEE 802.1Q VLAN tagging or any proprietary form of VLAN tagging.

Figure 63 shows this configuration.

FIGURE 63 Protocol-based (Layer 3) VLANs



To configure the VLANs shown in Figure 63, use the following procedure.

1. To permanently assign ports 1 – 8 and port 25 to IP subnet VLAN 1.1.1.0, enter the following commands.

```
PowerConnect> en
No password has been assigned yet...
PowerConnect# config t
PowerConnect(config)#
PowerConnect(config)# ip-subnet 1.1.1.0/24 name Green
PowerConnect(config-ip-subnet)# no dynamic
PowerConnect(config-ip-subnet)# static ethernet 1 to 8 ethernet 25
```

2. To permanently assign ports 9 – 16 and port 25 to IP subnet VLAN 1.1.2.0, enter the following commands.

```
PowerConnect(config-ip-subnet)# ip-subnet 1.1.2.0/24 name Yellow
PowerConnect(config-ip-subnet)# no dynamic
PowerConnect(config-ip-subnet)# static ethernet 9 to 16 ethernet 25
```

3. To permanently assign ports 17 – 25 to IP subnet VLAN 1.1.3.0, enter the following commands.

```
PowerConnect(config-ip-subnet)# ip-subnet 1.1.3.0/24 name Brown
PowerConnect(config-ip-subnet)# no dynamic
PowerConnect(config-ip-subnet)# static ethernet 17 to 25
```

4. To permanently assign ports 1 – 12 and port 25 to IPX network 1 VLAN, enter the following commands.

10 Configuring an IPv6 protocol VLAN

```
PowerConnect(config-ip-subnet)# ipx-network 1 ethernet_802.3 name Blue
PowerConnect(config-ipx-network)# no dynamic
PowerConnect(config-ipx-network)# static ethernet 1 to 12 ethernet 25
PowerConnect(config-ipx-network)#
```

5. To permanently assign ports 12 – 25 to Appletalk VLAN, enter the following commands.

```
PowerConnect(config-ipx-proto)# atalk-proto name Red
PowerConnect(config-ataalk-proto)# no dynamic
PowerConnect(config-ataalk-proto)# static ethernet 13 to 25
PowerConnect(config-ataalk-proto)# end
PowerConnect# write memory
PowerConnect#
```

Syntax: `ip-subnet <ip-addr> <ip-mask> [name <string>]`

Syntax: `ipx-network <ipx-network-number> <frame-encapsulation-type> netbios-allow | netbios-disallow [name <string>]`

Syntax: `ip-proto | ipx-proto | atalk-proto | decnet-proto | netbios-proto | other-proto static | exclude | dynamic ethernet <portnum> [to <portnum>] [name <string>]`

Configuring an IPv6 protocol VLAN

You can configure a protocol-based VLAN as a broadcast domain for IPv6 traffic. When the Layer 3 Switch receives an IPv6 multicast packet (a packet with 06 in the version field and 0xFF as the beginning of the destination address), the Layer 3 Switch forwards the packet to all other ports.

NOTE

The Layer 3 Switch forwards all IPv6 multicast packets to all ports except the port that received the packet, and does not distinguish among subnet directed multicasts.

You can add the VLAN ports as static ports or dynamic ports. A static port is always an active member of the VLAN. Dynamic ports within any protocol VLAN age out after 10 minutes if no member protocol traffic is received on a port within the VLAN. The aged out port, however, remains as a candidate dynamic port for that VLAN. The port becomes active in the VLAN again if member protocol traffic is received on that port.

Once a port is re-activated, the aging out period for the port is reset to 20 minutes. Each time a member protocol packet is received by a candidate dynamic port (aged out port) the port becomes active again and the aging out period is reset for 20 minutes.

To configure an IPv6 VLAN, enter commands such as the following.

```
PowerConnect(config)# vlan 2
PowerConnect(config-vlan-2)# untag ethernet 1 to 8
PowerConnect(config-vlan-2)# ipv6-proto name V6
PowerConnect(config-ipv6-subnet)# static ethernet 1 to 6
PowerConnect(config-ipv6-subnet)# dynamic
```

The first two commands configure a port-based VLAN and add ports 1 – 8 to the VLAN. The remaining commands configure an IPv6 VLAN within the port-based VLAN. The **static** command adds ports 1 – 6 as static ports, which do not age out. The **dynamic** command adds the remaining ports, 7 – 8, as dynamic ports. These ports are subject to aging as described above.

Syntax: `[no] ipv6-proto [name <string>]`

Routing between VLANs using virtual routing interfaces (Layer 3 Switches only)

Layer 3 Switches offer the ability to create a virtual routing interface within a Layer 2 STP port-based VLAN or within each Layer 3 protocol, IP subnet, or IPX network VLAN. This combination of multiple Layer 2 or Layer 3 broadcast domains, or both, and virtual routing interfaces are the basis for Dell's very powerful Integrated Switch Routing (ISR) technology. ISR is very flexible and can solve many networking problems. The following example is meant to provide ideas by demonstrating some of the concepts of ISR.


Example

Suppose you want to move routing out to each of three buildings in a network. Remember that the only protocols present on VLAN 2 and VLAN 3 are IP and IPX. Therefore, you can eliminate tagged ports 25 and 26 from both VLAN 2 and VLAN 3 and create new tagged port-based VLANs to support separate IP subnets and IPX networks for each backbone link.

You also need to create unique IP subnets and IPX networks within VLAN 2 and VLAN 3 at each building. This will create a fully routed IP and IPX backbone for VLAN 2 and VLAN 3. However, VLAN 4 has no protocol restrictions across the backbone. In fact there are requirements for NetBIOS and DecNet to be bridged among the three building locations. The IP subnet and IPX network that exists within VLAN 4 must remain a flat Layer 2 switched STP domain. You enable routing for IP and IPX on a virtual routing interface only on device-A. This will provide the flat IP and IPX segment with connectivity to the rest of the network. Within VLAN 4 IP and IPX will follow the STP topology. All other IP subnets and IPX networks will be fully routed and have use of all paths at all times during normal operation.

[Figure 64](#) shows the configuration described above.

FIGURE 64 Routing between protocol-based VLANs



To configure the Layer 3 VLANs and virtual routing interfaces on the Layer 3 Switch in [Figure 64](#), use the following procedure.

Configuring device-A

Enter the following commands to configure device-A. The following commands enable OSPF or RIP routing.

```
PowerConnect>en
No password has been assigned yet...
PowerConnect# configure terminal
PowerConnect(config)# hostname PowerConnect-A
PowerConnect-A(config)# router ospf
PowerConnect-A(config-ospf-router)# area 0.0.0.0 normal
Please save configuration to flash and reboot.
PowerConnect-A(config-ospf-router)#
```

The following commands create the port-based VLAN 2. In the previous example, an external device defined the router interfaces for VLAN 2. With ISR, routing for VLAN 2 is done locally within each device. Therefore, there are two ways you can solve this problem. One way is to create a unique IP subnet and IPX network VLAN, each with its own virtual routing interface and unique IP or IPX address within VLAN 2 on each device. In this example, this is the configuration used for VLAN 3. The second way is to split VLAN 2 into two separate port-based VLANs and create a virtual router interface within each port-based VLAN. Later in this example, this second option is used to create a port-based VLAN 8 to show that there are multiple ways to accomplish the same task with ISR.

You also need to create the Other-Protocol VLAN within port-based VLAN 2 and 8 to prevent unwanted protocols from being Layer 2 switched within port-based VLAN 2 or 8. Note that the only port-based VLAN that requires STP in this example is VLAN 4. You will need to configure the rest of the network to prevent the need to run STP.

```
PowerConnect-A(config-ospf-router)# vlan 2 name IP-Subnet_1.1.2.0/24
PowerConnect-A(config-vlan-2)# untag e 1 to 4
PowerConnect-A(config-vlan-2)# no spanning-tree
PowerConnect-A(config-vlan-2)# router-interface ve1
PowerConnect-A(config-vlan-2)# other-proto name block_other_protocols
PowerConnect-A(config-vlan-other-proto)# no dynamic
PowerConnect-A(config-vlan-other-proto)# exclude e 1 to 4
```

Once you have defined the port-based VLAN and created the virtual routing interface, you need to configure the virtual routing interface just as you would configure a physical interface.

```
PowerConnect-A(config-vlan-other-proto)# interface ve1
PowerConnect-A(config-vif-1)# ip address 1.1.2.1/24
PowerConnect-A(config-vif-1)# ip ospf area 0.0.0.0
```

Do the same thing for VLAN 8.

```
PowerConnect-A(config-vif-1)# vlan 8 name IPX_Network2
PowerConnect-A(config-vlan-8)# untag ethernet 5 to 8
PowerConnect-A(config-vlan-8)# no spanning-tree
PowerConnect-A(config-vlan-8)# router-interface ve 2
PowerConnect-A(config-vlan-8)# other-proto name block-other-protocols
PowerConnect-A(config-vlan-other-proto)# no dynamic
PowerConnect-A(config-vlan-other-proto)# exclude ethernet 5 to 8
PowerConnect-A(config-vlan-other-proto)# int ve2
PowerConnect-A(config-vif-2)# ipx network 2 ethernet_802.3
PowerConnect-A(config-vif-2)#
```

The next thing you need to do is create VLAN 3. This is very similar to the previous example with the addition of virtual routing interfaces to the IP subnet and IPX network VLANs. Also there is no need to exclude ports from the IP subnet and IPX network VLANs on the router.

```
PowerConnect-A(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
PowerConnect-A(config-vlan-3)# untag e 9 to 16
PowerConnect-A(config-vlan-3)# no spanning-tree
PowerConnect-A(config-vlan-3)# ip-subnet 1.1.1.0/24
PowerConnect-A(config-vlan-ip-subnet)# static e 9 to 12
PowerConnect-A(config-vlan-ip-subnet)# router-interface ve3
PowerConnect-A(config-vlan-ip-subnet)# ipx-network 1 ethernet_802.3
PowerConnect-A(config-vlan-ipx-network)# static e 13 to 16
PowerConnect-A(config-vlan-ipx-network)# router-interface ve4
PowerConnect-A(config-vlan-ipx-network)# other-proto name block-other-protocols
PowerConnect-A(config-vlan-other-proto)# exclude e 9 to 16
PowerConnect-A(config-vlan-other-proto)# no dynamic
PowerConnect-A(config-vlan-other-proto)# interface ve 3
PowerConnect-A(config-vif-3)# ip addr 1.1.1.1/24
PowerConnect-A(config-vif-3)# ip ospf area 0.0.0.0
PowerConnect-A(config-vif-3)# int ve4
PowerConnect-A(config-vif-4)# ipx network 1 ethernet_802.3
PowerConnect-A(config-vif-4)#
```

Now configure VLAN 4. Remember this is a flat segment that, in the previous example, obtained its IP default gateway and IPX router services from an external device. In this example, device-A will provide the routing services for VLAN 4. You also want to configure the STP priority for VLAN 4 to make device-A the root bridge for this VLAN.

```
PowerConnect-A(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
PowerConnect-A(config-vlan-4)# untag ethernet 17 to 24
PowerConnect-A(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-A(config-vlan-4)# spanning-tree
PowerConnect-A(config-vlan-4)# spanning-tree priority 500
```

10 Routing between VLANs using virtual routing interfaces (Layer 3 Switches only)

```
PowerConnect-A(config-vlan-4)# router-interface ve5
PowerConnect-A(config-vlan-4)# int ve5
PowerConnect-A(config-vif-5)# ip address 1.1.3.1/24
PowerConnect-A(config-vif-5)# ip ospf area 0.0.0.0
PowerConnect-A(config-vif-5)# ipx network 3 ethernet_802.3
PowerConnect-A(config-vif-5)#
```

It is time to configure a separate port-based VLAN for each of the routed backbone ports (Ethernet 25 and 26).

If you do not create a separate tagged port-based VLAN for each point-to-point backbone link, you need to include tagged interfaces for Ethernet 25 and 26 within VLANs 2, 3, and 8. This type of configuration makes the entire backbone a single STP domain for each VLAN 2, 3, and 8. This is the configuration used in the example in [“Configuring IP subnet, IPX network and protocol-based VLANs”](#) on page 272. In this scenario, the virtual routing interfaces within port-based VLANs 2, 3, and 8 will be accessible using only one path through the network. The path that is blocked by STP is not available to the routing protocols until it is in the STP FORWARDING state.

```
PowerConnect-A(config-vif-5)# vlan 5 name Rtr_BB_to_Bldg.2
PowerConnect-A(config-vlan-5)# tag e 25
PowerConnect-A(config-vlan-5)# no spanning-tree
PowerConnect-A(config-vlan-5)# router-interface ve6
PowerConnect-A(config-vlan-5)# vlan 6 name Rtr_BB_to_Bldg.3
PowerConnect-A(config-vlan-6)# tag ethernet 26
PowerConnect-A(config-vlan-6)# no spanning-tree
PowerConnect-A(config-vlan-6)# router-interface ve7
PowerConnect-A(config-vlan-6)# int ve6
PowerConnect-A(config-vif-6)# ip addr 1.1.4.1/24
PowerConnect-A(config-vif-6)# ip ospf area 0.0.0.0
PowerConnect-A(config-vif-6)# ipx network 4 ethernet_802.3
PowerConnect-A(config-vif-6)# int ve7
PowerConnect-A(config-vif-7)# ip addr 1.1.5.1/24
PowerConnect-A(config-vif-7)# ip ospf area 0.0.0.0
PowerConnect-A(config-vif-7)# ipx network 5 ethernet_802.3
PowerConnect-A(config-vif-7)#
```

This completes the configuration for device-A. The configuration for device-B and C is very similar except for a few issues which are as follows:

- IP subnets and IPX networks configured on device-B and device-C must be unique across the entire network, except for the backbone port-based VLANs 5, 6, and 7 where the subnet is the same but the IP address must change.
- There is no need to change the default priority of STP within VLAN 4.
- There is no need to include a virtual router interface within VLAN 4.
- The backbone VLAN between device-B and device-C must be the same at both ends and requires a new VLAN ID. The VLAN ID for this port-based VLAN is VLAN 7.

Configuration for device-B

Enter the following commands to configure device-B.

```
PowerConnect> en
No password has been assigned yet...
PowerConnect# config t
PowerConnect(config)# hostname PowerConnect-B
PowerConnect-B(config)# router ospf
PowerConnect-B(config-ospf-router)# area 0.0.0.0 normal
PowerConnect-B(config-ospf-router)# router ipx
```

```
PowerConnect-B(config-ospf-router)# vlan 2 name IP-Subnet_1.1.6.0/24
PowerConnect-B(config-vlan-2)# untag e 1 to 4
PowerConnect-B(config-vlan-2)# no spanning-tree
PowerConnect-B(config-vlan-2)# router-interface ve1
PowerConnect-B(config-vlan-2)# other-proto name block-other-protocols
PowerConnect-B(config-vlan-other-proto)# no dynamic
PowerConnect-B(config-vlan-other-proto)# exclude e 1 to 4
PowerConnect-B(config-vlan-other-proto)# int ve1
PowerConnect-B(config-vif-1)# ip addr 1.1.6.1/24
PowerConnect-B(config-vif-1)# ip ospf area 0.0.0.0
PowerConnect-B(config-vif-1)# vlan 8 name IPX_Network6
PowerConnect-B(config-vlan-8)# untag e 5 to 8
PowerConnect-B(config-vlan-8)# no span
PowerConnect-B(config-vlan-8)# router-int ve2
PowerConnect-B(config-vlan-8)# other-proto name block-other-protocols
PowerConnect-B(config-vlan-other-proto)# no dynamic
PowerConnect-B(config-vlan-other-proto)# exclude e 5 to 8
PowerConnect-B(config-vlan-other-proto)# int ve2
PowerConnect-B(config-vif-2)# ipx net 6 ethernet_802.3
PowerConnect-B(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
PowerConnect-B(config-vlan-3)# untag e 9 to 16
PowerConnect-B(config-vlan-3)# no spanning-tree
PowerConnect-B(config-vlan-3)# ip-subnet 1.1.7.0/24
PowerConnect-B(config-vlan-ip-subnet)# static e 9 to 12
PowerConnect-B(config-vlan-ip-subnet)# router-interface ve3
PowerConnect-B(config-vlan-ip-subnet)# ipx-network 7 ethernet_802.3
PowerConnect-B(config-vlan-ipx-network)# static e 13 to 16
PowerConnect-B(config-vlan-ipx-network)# router-interface ve4
PowerConnect-B(config-vlan-ipx-network)# other-proto name block-other-protocols
PowerConnect-B(config-vlan-other-proto)# exclude e 9 to 16
PowerConnect-B(config-vlan-other-proto)# no dynamic
PowerConnect-B(config-vlan-other-proto)# interface ve 3
PowerConnect-B(config-vif-3)# ip addr 1.1.7.1/24
PowerConnect-B(config-vif-3)# ip ospf area 0.0.0.0
PowerConnect-B(config-vif-3)# int ve4
PowerConnect-B(config-vif-4)# ipx network 7 ethernet_802.3
PowerConnect-B(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
PowerConnect-B(config-vlan-4)# untag ethernet 17 to 24
PowerConnect-B(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-B(config-vlan-4)# spanning-tree
PowerConnect-B(config-vlan-4)# vlan 5 name Rtr_BB_to_Bldg.1
PowerConnect-B(config-vlan-5)# tag e 25
PowerConnect-B(config-vlan-5)# no spanning-tree
PowerConnect-B(config-vlan-5)# router-interface ve5
PowerConnect-B(config-vlan-5)# vlan 7 name Rtr_BB_to_Bldg.3
PowerConnect-B(config-vlan-7)# tag ethernet 26
PowerConnect-B(config-vlan-7)# no spanning-tree
PowerConnect-B(config-vlan-7)# router-interface ve6
PowerConnect-B(config-vlan-7)# int ve5
PowerConnect-B(config-vif-5)# ip addr 1.1.4.2/24
PowerConnect-B(config-vif-5)# ip ospf area 0.0.0.0
PowerConnect-B(config-vif-5)# ipx network 4 ethernet_802.3
PowerConnect-B(config-vif-5)# int ve6
PowerConnect-B(config-vif-6)# ip addr 1.1.8.1/24
PowerConnect-B(config-vif-6)# ip ospf area 0.0.0.0
PowerConnect-B(config-vif-6)# ipx network 8 ethernet_802.3
PowerConnect-B(config-vif-6)#
```

Configuration for device-C

Enter the following commands to configure device-C.

```
PowerConnect> en
No password has been assigned yet...
PowerConnect# config t
PowerConnect-C(config)# hostname PowerConnect-C
PowerConnect-C(config)# router ospf
PowerConnect-C(config-ospf-router)# area 0.0.0.0 normal
PowerConnect-C(config-ospf-router)# router ipx
PowerConnect-C(config-ospf-router)# vlan 2 name IP-Subnet_1.1.9.0/24
PowerConnect-C(config-vlan-2)# untag e 1 to 4
PowerConnect-C(config-vlan-2)# no spanning-tree
PowerConnect-C(config-vlan-2)# router-interface ve1
PowerConnect-C(config-vlan-2)# other-proto name block-other-protocols
PowerConnect-C(config-vlan-other-proto)# no dynamic
PowerConnect-C(config-vlan-other-proto)# exclude e 1 to 4
PowerConnect-C(config-vlan-other-proto)# int ve1
PowerConnect-C(config-vif-1)# ip addr 1.1.9.1/24
PowerConnect-C(config-vif-1)# ip ospf area 0.0.0.0
PowerConnect-C(config-vif-1)# vlan 8 name IPX_Network9
PowerConnect-C(config-vlan-8)# untag e 5 to 8
PowerConnect-C(config-vlan-8)# no span
PowerConnect-C(config-vlan-8)# router-int ve2
PowerConnect-C(config-vlan-8)# other-proto name block-other-protocols
PowerConnect-C(config-vlan-other-proto)# no dynamic
PowerConnect-C(config-vlan-other-proto)# exclude e 5 to 8
PowerConnect-C(config-vlan-other-proto)# int ve2
PowerConnect-C(config-vif-2)# ipx net 9 ethernet_802.3
PowerConnect-C(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
PowerConnect-C(config-vlan-3)# untag e 9 to 16
PowerConnect-C(config-vlan-3)# no spanning-tree
PowerConnect-C(config-vlan-3)# ip-subnet 1.1.10.0/24
PowerConnect-C(config-vlan-ip-subnet)# static e 9 to 12
PowerConnect-C(config-vlan-ip-subnet)# router-interface ve3
PowerConnect-C(config-vlan-ip-subnet)# ipx-network 10 ethernet_802.3
PowerConnect-C(config-vlan-ipx-network)# static e 13 to 16
PowerConnect-C(config-vlan-ipx-network)# router-interface ve4
PowerConnect-C(config-vlan-ipx-network)# other-proto name block-other-protocols
PowerConnect-C(config-vlan-other-proto)# exclude e 9 to 16
PowerConnect-C(config-vlan-other-proto)# no dynamic
PowerConnect-C(config-vlan-other-proto)# interface ve 3
PowerConnect-C(config-vif-3)# ip addr 1.1.10.1/24
PowerConnect-C(config-vif-3)# ip ospf area 0.0.0.0
PowerConnect-C(config-vif-3)# int ve4
PowerConnect-C(config-vif-4)# ipx network 10 ethernet_802.3
PowerConnect-C(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
PowerConnect-C(config-vlan-4)# untag ethernet 17 to 24
PowerConnect-C(config-vlan-4)# tag ethernet 25 to 26
PowerConnect-C(config-vlan-4)# spanning-tree
PowerConnect-C(config-vlan-4)# vlan 7 name Rtr_BB_to_Bldg.2
PowerConnect-C(config-vlan-7)# tag e 25
PowerConnect-C(config-vlan-7)# no spanning-tree
PowerConnect-C(config-vlan-7)# router-interface ve5
PowerConnect-C(config-vlan-7)# vlan 6 name Rtr_BB_to_Bldg.1
PowerConnect-C(config-vlan-6)# tag ethernet 26
PowerConnect-C(config-vlan-6)# no spanning-tree
PowerConnect-C(config-vlan-6)# router-interface ve6
PowerConnect-C(config-vlan-6)# int ve5
```

```

PowerConnect-C(config-vif-5)# ip addr 1.1.8.2/24
PowerConnect-C(config-vif-5)# ip ospf area 0.0.0.0
PowerConnect-C(config-vif-5)# ipx network 8 ethernet_802.3
PowerConnect-C(config-vif-5)# int ve6
PowerConnect-C(config-vif-6)# ip addr 1.1.5.2/24
PowerConnect-C(config-vif-6)# ip ospf area 0.0.0.0
PowerConnect-C(config-vif-6)# ipx network 5 ethernet_802.3
PowerConnect-C(config-vif-6)#

```

Configuring uplink ports within a port-based VLAN

You can configure a subset of the ports in a port-based VLAN as uplink ports. When you configure uplink ports in a port-based VLAN, the device sends all broadcast and unknown-unicast traffic from a port in the VLAN to the uplink ports, but not to other ports within the VLAN. Thus, the uplink ports provide tighter broadcast control within the VLAN.

This uplink port feature behaves the same as the private VLAN feature, but with the ability to support tagged ports. This feature also supports two private VLAN modes: the Primary ports (uplink ports) and Isolated ports (host ports).

For example, if two ports within a port-based VLAN are Gbps ports attached to the network and the other ports are 10/100 ports attached to clients, you can configure the two ports attached to the network as uplink ports. In this configuration, broadcast and unknown-unicast traffic in the VLAN does not go to all ports. The traffic goes only to the uplink ports. The clients on the network do not receive broadcast and unknown-unicast traffic from other ports, including other clients.

Configuration considerations

- On PowerConnect devices, flooded traffic (unknown unicast, unregistered multicast, and broadcast traffic) is hardware forwarded.
- This feature should not be enabled with protocol VLANs or private VLANs in the same VLAN.

Configuration syntax

To configure a port-based VLAN containing uplink ports, enter commands such as the following.

```

PowerConnect(config)# vlan 10 by port
PowerConnect(config-vlan-10)# untag ethernet 1 to 24
PowerConnect(config-vlan-10)# untag ethernet 1 to 2
PowerConnect(config-vlan-10)# uplink-switch ethernet 1 to 2

```

Syntax: [no] uplink-switch ethernet <portnum> [to <portnum> | ethernet <portnum>]

In this example, 24 ports on a 10/100 module and two Gbps ports on a Gbps module are added to port-based VLAN 10. The two Gbps ports are then configured as uplink ports.

To configure a port-based VLAN containing uplink ports on PowerConnect devices, enter the following commands.

```

PowerConnect(config)# vlan 500
PowerConnect(config-vlan-500)# tagged ethernet 11 to 21
PowerConnect(config-vlan-500)# uplink-switch ethernet 20 to 21

```

Syntax: [no] uplink-switch ethernet <portnum> to <portnum>

Configuring the same IP subnet address on multiple port-based VLANs

For a device to route between port-based VLANs, you must add a virtual routing interface to each VLAN. Generally, you also configure a unique IP subnet address on each virtual routing interface. For example, if you have three port-based VLANs, you add a virtual routing interface to each VLAN, then add a separate IP subnet address to each virtual routing interface. The IP address on each of the virtual routing interfaces must be in a separate subnet. The device routes Layer 3 traffic between the subnets using the subnet addresses.

NOTE


This feature applies only to Layer 3 Switches.

NOTE

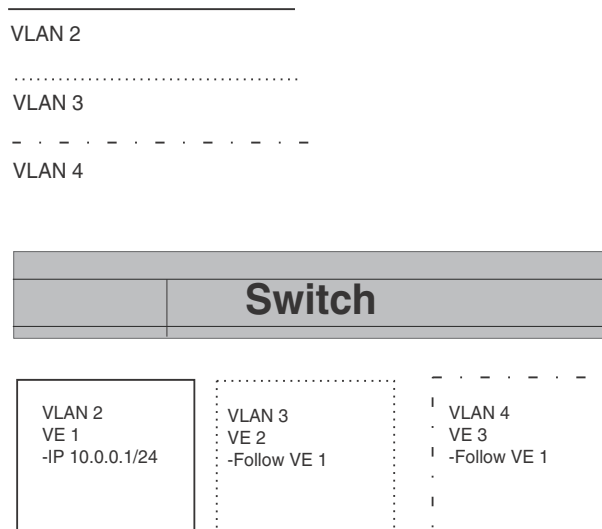
Before using the method described in this section, refer to “[Configuring VLAN groups and virtual routing interface groups](#)” on page 285. You might be able to achieve the results you want using the methods in that section instead.

[Figure 65](#) shows an example of this type of configuration.

FIGURE 65 Multiple port-based VLANs with separate protocol addresses



As shown in this example, each VLAN has a separate IP subnet address. If you need to conserve IP subnet addresses, you can configure multiple VLANs with the same IP subnet address, as shown in [Figure 66](#).

FIGURE 66 Multiple port-based VLANs with the same protocol address

Each VLAN still requires a separate virtual routing interface. However, all three VLANs now use the same IP subnet address.

In addition to conserving IP subnet addresses, this feature allows containment of Layer 2 broadcasts to segments within an IP subnet. For ISP environments where the same IP subnet is allocated to different customers, placing each customer in a separate VLAN allows all customers to share the IP subnet address, while at the same time isolating them from one another Layer 2 broadcasts.

NOTE

You can provide redundancy to an IP subnet address that contains multiple VLANs using a pair of Layer 3 Switches configured for VRRP (Virtual Router Redundancy Protocol).

The device performs proxy Address Resolution Protocol (ARP) for hosts that want to send IP traffic to hosts in other VLANs that are sharing the same IP subnet address. If the source and destination hosts are in the same VLAN, the device does not need to use ARP:

- If a host attached to one VLAN sends an ARP message for the MAC address of a host in one of the other VLANs using the same IP subnet address, the device performs a proxy ARP on behalf of the other host. The device then replies to the ARP by sending the virtual routing interface MAC address. The device uses the same MAC address for all virtual routing interfaces.

When the host that sent the ARP then sends a unicast packet addressed to the virtual routing interface MAC address, the device switches the packet on Layer 3 to the destination host on the VLAN.

NOTE

If the device ARP table does not contain the requested host, the device forwards the ARP request on Layer 2 to the same VLAN as the one that received the ARP request. Then the device sends an ARP for the destination to the other VLANs that are using the same IP subnet address.

10 Configuring the same IP subnet address on multiple port-based VLANs

- If the destination is in the same VLAN as the source, the device does not need to perform a proxy ARP.

To configure multiple VLANs to use the same IP subnet address:

- Configure each VLAN, including adding tagged or untagged ports.
- Configure a separate virtual routing interface for each VLAN, but do not add an IP subnet address to more than one of the virtual routing interfaces.
- Configure the virtual routing interfaces that do not have the IP subnet address to “follow” the virtual routing interface that does have the address.

To configure the VLANs shown in [Figure 66](#), you could enter the following commands.

```
PowerConnect(config)# vlan 1 by port
PowerConnect(config-vlan-1)# untag ethernet 1
PowerConnect(config-vlan-1)# tag ethernet 8
PowerConnect(config-vlan-1)# router-interface ve 1
```

Syntax: `router-interface ve <number>`

The commands above configure port-based VLAN 1. The VLAN has one untagged port (1) and a tagged port (8). In this example, all three VLANs contain port 8 so the port must be tagged to allow the port to be in multiple VLANs. You can configure VLANs to share a Layer 3 protocol interface regardless of tagging. A combination of tagged and untagged ports is shown in this example to demonstrate that sharing the interface does not change other VLAN features.

Notice that each VLAN still requires a unique virtual routing interface.

The following commands configure port-based VLANs 2 and 3.

```
PowerConnect(config-vlan-1)# vlan 2 by port
PowerConnect(config-vlan-2)# untag ethernet 2
PowerConnect(config-vlan-2)# tag ethernet 8
PowerConnect(config-vlan-2)# router-interface ve 2
PowerConnect(config-vlan-2)# vlan 3 by port
PowerConnect(config-vlan-3)# untag ethernet 5 to 6
PowerConnect(config-vlan-3)# tag ethernet 8
PowerConnect(config-vlan-3)# router-interface ve 3
```

The following commands configure an IP subnet address on virtual routing interface 1.

```
PowerConnect(config-vlan-3)# interface ve 1
PowerConnect(config-vif-1)# ip address 10.0.0.1/24
```

The following commands configure virtual routing interfaces 2 and 3 to “follow” the IP subnet address configured on virtual routing interface 1.

```
PowerConnect(config-vif-1)# interface ve 2
PowerConnect(config-vif-2)# ip follow ve 1
PowerConnect(config-vif-2)# interface ve 3
PowerConnect(config-vif-3)# ip follow ve 1
```

NOTE

Since virtual routing interfaces 2 and 3 do not have their own IP subnet addresses but instead are “following” virtual routing interface a IP address, you still can configure an IPX or AppleTalk interface on virtual routing interfaces 2 and 3.

Configuring VLAN groups and virtual routing interface groups

NOTE

On PowerConnect B-Series TI24X devices, VLAN groups are supported.

To simplify configuration when you have many VLANs with the same configuration, you can configure VLAN groups and virtual routing interface groups.

NOTE

VLAN groups are supported on Layer 3 Switches and Layer 2 Switches. Virtual routing interface groups are supported only on Layer 3 Switches.

When you create a VLAN group, the VLAN parameters you configure for the group apply to all the VLANs within the group. Additionally, you can easily associate the same IP subnet interface with all the VLANs in a group by configuring a virtual routing interface group with the same ID as the VLAN group.

- The VLAN group feature allows you to create multiple port-based VLANs with identical port members. Since the member ports are shared by all the VLANs within the group, you must add the ports as tagged ports. This feature not only simplifies VLAN configuration but also allows you to have a large number of identically configured VLANs in a startup-config file on the device flash memory module. Normally, a startup-config file with a large number of VLANs might not fit on the flash memory module. By grouping the identically configured VLANs, you can conserve space in the startup-config file so that it fits on the flash memory module.
- The virtual routing interface group feature is useful when you want to configure the same IP subnet address on all the port-based VLANs within a VLAN group. You can configure a virtual routing interface group only after you configure a VLAN group with the same ID. The virtual routing interface group automatically applies to the VLANs in the VLAN group that has the same ID and cannot be applied to other VLAN groups or to individual VLANs.

You can create up to 32 VLAN groups and 32 virtual routing interface groups. A virtual routing interface group always applies only to the VLANs in the VLAN group with the same ID.

NOTE

Depending on the size of the VLAN ID range you want to use for the VLAN group, you might need to allocate additional memory for VLANs. On Layer 3 Switches, if you allocate additional memory for VLANs, you also need to allocate the same amount of memory for virtual routing interfaces. This is true regardless of whether you use the virtual routing interface groups. To allocate additional memory, refer to [“Allocating memory for more VLANs or virtual routing interfaces”](#) on page 288.

Configuring a VLAN group

To configure a VLAN group, enter commands such as the following.

```
PowerConnect(config)# vlan-group 1 vlan 2 to 1000
PowerConnect(config-vlan-group-1)# tagged 1 to 2
```

The first command in this example begins configuration for VLAN group 1, and assigns VLANs 2 through 1000 to the group. The second command adds ports 1 and 2 as tagged ports. Since all the VLANs in the group share the ports, you must add the ports as tagged ports.

Syntax: `vlan-group <num> vlan <vlan-id> to <vlan-id>`

Syntax: `tagged ethernet [<portnum> [to<portnum> | ethernet <portnum>]`

The `<num>` parameter with the **vlan-group** command specifies the VLAN group ID and can be from 1 - 32. The **vlan <vlan-id> to <vlan-id>** parameters specify a contiguous range (a range with no gaps) of individual VLAN IDs. Specify the low VLAN ID first and the high VLAN ID second. The command adds all the specified VLANs to the VLAN group.

NOTE

The device memory must be configured to contain at least the number of VLANs you specify for the higher end of the range. For example, if you specify 2048 as the VLAN ID at the high end of the range, you first must increase the memory allocation for VLANs to 2048 or higher. Additionally, on Layer 3 Switches, if you allocate additional memory for VLANs, you also need to allocate the same amount of memory for virtual routing interfaces, before you configure the VLAN groups. This is true regardless of whether you use the virtual routing interface groups. The memory allocation is required because the VLAN groups and virtual routing interface groups have a one-to-one mapping. Refer to [“Allocating memory for more VLANs or virtual routing interfaces”](#) on page 288.

If a VLAN within the range you specify is already configured, the CLI does not add the group but instead displays an error message. In this case, create the group by specifying a valid contiguous range. Then add more VLANs to the group after the CLI changes to the configuration level for the group. See the following example.

You can add and remove individual VLANs or VLAN ranges from at the VLAN group configuration level. For example, if you want to add VLANs 1001 and 1002 to VLAN group 1 and remove VLANs 900 through 1000, enter the following commands.

```
PowerConnect(config-vlan-group-1)# add-vlan 1001 to 1002
PowerConnect(config-vlan-group-1)# remove-vlan 900 to 1000
```

Syntax: **add-vlan** `<vlan-id>` [**to** `<vlan-id>`]

Syntax: **remove-vlan** `<vlan-id>` [**to** `<vlan-id>`]

Displaying information about VLAN groups

To display VLAN group configuration information for devices, use the **show vlan-group** command.

```
PowerConnect# show vlan-group
vlan-group 1 vlan 2 to 20
  tagged ethe 1 to 2
!
vlan-group 2 vlan 21 to 40
  tagged ethe 1 to 2
!
```

Syntax: **show vlan-group** [`<group-id>`]

To display VLAN group configuration information for PowerConnect devices, enter the following command.

```
PowerConnect# show vlan-group
vlan-group 1 vlan 2 to 500
  tagged ethe 2 ethe 5
```

Syntax: **show vlan-group**

The `<group-id>` specifies a VLAN group. If you do not use this parameter, the configuration information for all the configured VLAN groups is displayed.

Configuring a virtual routing interface group

A virtual routing interface group allows you to associate the same IP subnet interface with multiple port-based VLANs. For example, if you associate a virtual routing interface group with a VLAN group, all the VLANs in the group have the IP interface of the virtual routing interface group.

Configuration notes and feature limitations

- When you configure a virtual routing interface group, all members of the group have the same IP subnet address. This feature is useful in collocation environments where the device has many IP addresses and you want to conserve the IP address space.
- The **group-router-interface** command creates router interfaces for each VLAN in the VLAN group by using the VLAN IDs of each of the VLANs as the corresponding virtual interface number. Therefore, if a VLAN group contains VLAN IDs greater than the maximum virtual interface number allowed, the **group-router-interface** command will be rejected.

CLI syntax

To configure a virtual routing interface group, enter commands such as the following.

```
PowerConnect(config)# vlan-group 1
PowerConnect(config-vlan-group-1)# group-router-interface
PowerConnect(config-vlan-group-1)# exit
PowerConnect(config)# interface group-ve 1
PowerConnect(config-vif-group-1)# ip address 10.10.10.1/24
```

These commands enable VLAN group 1 to have a group virtual routing interface, then configure virtual routing interface group 1. The software always associates a virtual routing interface group only with the VLAN group that has the same ID. In this example, the VLAN group ID is 1, so the corresponding virtual routing interface group also must have ID 1.

Syntax: group-router-interface

Syntax: interface group-ve <num>

Syntax: [no] ip address <ip-addr> <ip-mask> [secondary]

or

Syntax: [no] ip address <ip-addr>/<mask-bits> [secondary]

The **router-interface-group** command enables a VLAN group to use a virtual routing interface group. Enter this command at the configuration level for the VLAN group. This command configures the VLAN group to use the virtual routing interface group that has the same ID as the VLAN group. You can enter this command when you configure the VLAN group for the first time or later, after you have added tagged ports to the VLAN and so on.

The <num> parameter in the **interface group-ve <num>** command specifies the ID of the VLAN group with which you want to associate this virtual routing interface group. The VLAN group must already be configured and enabled to use a virtual routing interface group. The software automatically associates the virtual routing interface group with the VLAN group that has the same ID. You can associate a virtual routing interface group only with the VLAN group that has the same ID.

NOTE

IPv6 is not supported with **group-ve**.

The syntax and usage for the **ip address** command is the same as when you use the command at the interface level to add an IP interface.

Displaying the VLAN group and virtual routing interface group information

To verify configuration of VLAN groups and virtual routing interface groups, display the running-config file. If you have saved the configuration to the startup-config file, you also can verify the configuration by displaying the startup-config file. The following example shows the running-config information for the VLAN group and virtual routing interface group configured in the previous examples. The information appears in the same way in the startup-config file.

```
PowerConnect# show running-config

lines not related to the VLAN group omitted...

vlan-group 1 vlan 2 to 900
  add-vlan 1001 to 1002
  tagged ethe 1 to 2
  router-interface-group

lines not related to the virtual routing interface group omitted...

interface group-ve 1
  ip address 10.10.10.1 255.255.255.0
```

NOTE

If you have enabled display of subnet masks in CIDR notation, the IP address information is shown as follows: 10.10.10.1/24.

Allocating memory for more VLANs or virtual routing interfaces

Layer 2 and Layer 3 Switches support up to 4095 VLANs. In addition, Layer 3 switches support up to 512 virtual routing interfaces.

The number of VLANs and virtual routing interfaces supported on your product depends on the device and, for Chassis devices, the amount of DRAM on the management module. [Table 41](#) lists the default and configurable maximum numbers of VLANs and virtual routing interfaces for Layer 2 and Layer 3 Switches. Unless otherwise noted, the values apply to both types of switches.

TABLE 41 VLAN and virtual routing interface support

VLANs		Virtual routing interfaces	
Default maximum	Configurable maximum	Default maximum	Configurable maximum
64	4094	255	512

NOTE

If many of your VLANs will have an identical configuration, you might want to configure VLAN groups and virtual routing interface groups after you increase the system capacity for VLANs and virtual routing interfaces. Refer to [“Configuring VLAN groups and virtual routing interface groups”](#) on page 285.

Increasing the number of VLANs you can configure

NOTE

Although you can specify up to 4095 VLANs, you can configure only 4094 VLANs. VLAN ID 4094 is reserved for use by the Single Spanning Tree feature.

To increase the maximum number of VLANs you can configure, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# system-max vlan 2048
PowerConnect(config)# write memory
PowerConnect(config)# end
PowerConnect# reload
```

Syntax: `system-max vlan <num>`

The <num> parameter indicates the maximum number of VLANs. The range of valid values depends on the device you are configuring. Refer to [Table 41](#).

Increasing the number of virtual routing interfaces you can configure

To increase the maximum number of virtual routing interfaces you can configure, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# system-max virtual-interface 512
PowerConnect(config)# write memory
PowerConnect(config)# end
PowerConnect# reload
```

Syntax: `system-max virtual-interface <num>`

The <num> parameter indicates the maximum number of virtual routing interfaces. The range of valid values depends on the device you are configuring. Refer to [Table 41](#).

Configuring super aggregated VLANs

NOTE

On PowerConnect B-Series TI24X devices, this feature is supported .

You can aggregate multiple VLANs within another VLAN. This feature allows you to construct Layer 2 paths and channels. This feature is particularly useful for Virtual Private Network (VPN) applications in which you need to provide a private, dedicated Ethernet connection for an individual client to transparently reach its subnet across multiple networks.

Conceptually, the paths and channels are similar to Asynchronous Transfer Mode (ATM) paths and channels. A path contains multiple channels, each of which is a dedicated circuit between two end points. The two devices at the end points of the channel appear to each other to be directly attached. The network that connects them is transparent to the two devices.


You can aggregate up to 4094 VLANs within another VLAN. This provides a total VLAN capacity on one device of 16,760,836 channels (4094 * 4094).

The devices connected through the channel are not visible to devices in other channels. Therefore, each client has a private link to the other side of the channel.

The feature allows point-to-point and point-to-multipoint connections.

Figure 67 shows a conceptual picture of the service that aggregated VLANs provide. Aggregated VLANs provide a path for multiple client channels. The channels do not receive traffic from other channels. Thus, each channel is a private link.

FIGURE 67 Conceptual model of the super aggregated VLAN application

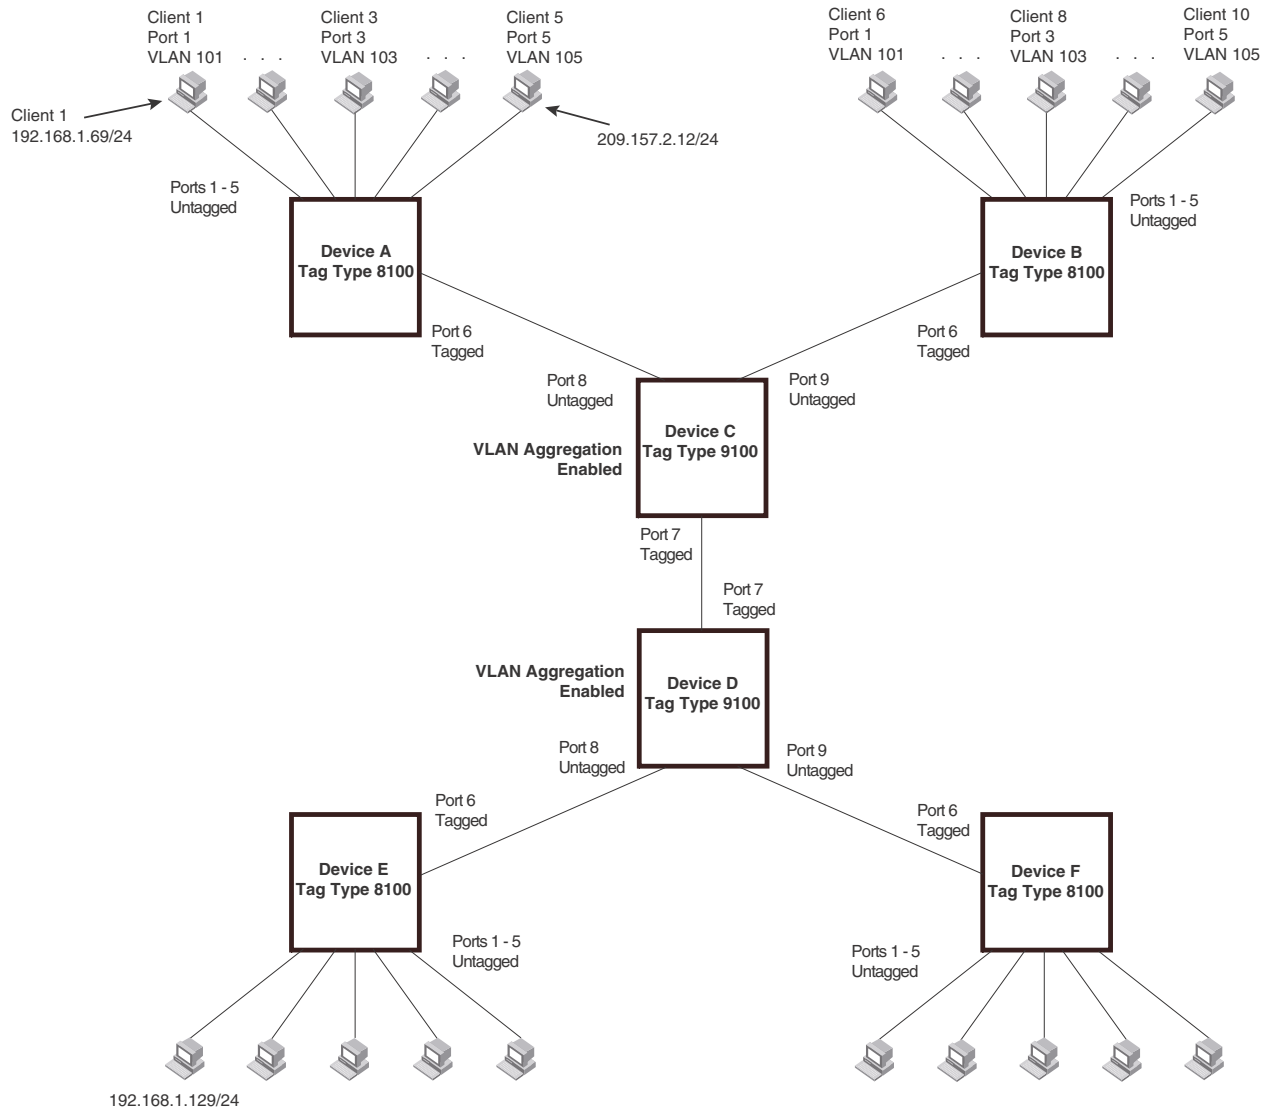


Each client connected to the edge device is in its own port-based VLAN, which is like an ATM channel. All the clients' VLANs are aggregated by the edge device into a single VLAN for connection to the core. The single VLAN that aggregates the clients' VLANs is like an ATM path.

The device that aggregates the VLANs forwards the aggregated VLAN traffic through the core. The core can consist of multiple devices that forward the aggregated VLAN traffic. The edge device at the other end of the core separates the aggregated VLANs into the individual client VLANs before forwarding the traffic. The edge devices forward the individual client traffic to the clients. For the clients' perspective, the channel is a direct point-to-point link.

Figure 68 shows an example application that uses aggregated VLANs. This configuration includes the client connections shown in Figure 67.

FIGURE 68 Example of a super aggregated VLAN application



In this example, a collocation service provides private channels for multiple clients. Although the same devices are used for all the clients, the VLANs ensure that each client receives its own Layer 2 broadcast domain, separate from the broadcast domains of other clients. For example, client 1 cannot ping client 5.

The clients at each end of a channel appear to each other to be directly connected and thus can be on the same subnet and use network services that require connection to the same subnet. In this example, client 1 is in subnet 192.168.1.0/24 and so is the device at the other end of client 1 channel.

Since each VLAN configured on the core devices is an aggregate of multiple client VLANs, the aggregated VLANs greatly increase the number of clients a core device can accommodate.

This example shows a single link between the core devices. However, you can use a trunk group to add link-level redundancy.

Configuration note

- Super Aggregated VLANs and VSRP are not supported together on the same device.

Configuring aggregated VLANs

To configure aggregated VLANs, perform the following tasks:

- On each edge device, configure a separate port-based VLAN for each client connected to the edge device. In each client VLAN:
 - Add the port connected to the client as an untagged port.
 - Add the port connected to the core device (the device that will aggregate the VLANs) as a tagged port. This port must be tagged because all the client VLANs share the port as an uplink to the core device.
- On each core device:
 - Enable VLAN aggregation. This support allows the core device to add an additional tag to each Ethernet frame that contains a VLAN packet from the edge device. The additional tag identifies the aggregate VLAN (the path). However, the additional tag can cause the frame to be longer than the maximum supported frame size. The larger frame support allows Ethernet frames up to 1530 bytes long.
 - Configure a VLAN tag type (tag ID) that is different than the tag type used on the edge devices. If you use the default tag type (8100) on the edge devices, set the tag type on the core devices to another value, such as 9100. The tag type must be the same on all the core devices. The edge devices also must have the same tag type but the type must be different from the tag type on the core devices.

NOTE

You can enable the Spanning Tree Protocol (STP) on the edge devices or the core devices, but not both. If you enable STP on the edge devices and the core devices, STP will prevent client traffic from travelling through the core to the other side.

Configuring aggregated VLANs on an edge device

To configure the aggregated VLANs on device A in [Figure 68](#) on page 291, enter the following commands.

```
PowerConnect(config)# vlan 101 by port
PowerConnect(config-vlan-101)# tagged ethernet 6
PowerConnect(config-vlan-101)# untagged ethernet 1
PowerConnect(config-vlan-101)# exit
PowerConnect(config)# vlan 102 by port
PowerConnect(config-vlan-102)# tagged ethernet 6
PowerConnect(config-vlan-102)# untagged ethernet 2
PowerConnect(config-vlan-102)# exit
PowerConnect(config)# vlan 103 by port
PowerConnect(config-vlan-103)# tagged ethernet 6
PowerConnect(config-vlan-103)# untagged ethernet 3
PowerConnect(config-vlan-103)# exit
PowerConnect(config)# vlan 104 by port
PowerConnect(config-vlan-104)# tagged ethernet 6
PowerConnect(config-vlan-104)# untagged ethernet 4
PowerConnect(config-vlan-104)# exit
PowerConnect(config)# vlan 105 by port
```

```
PowerConnect(config-vlan-105)# tagged ethernet 6
PowerConnect(config-vlan-105)# untagged ethernet 5
PowerConnect(config-vlan-105)# exit
PowerConnect(config)# write memory
```

Syntax: [no] vlan <vlan-id> [by port]

Syntax: [no] tagged ethernet <portnum> [to [<portnum> | ethernet <portnum>]]

Syntax: [no] untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]]

Use the **tagged** command to add the port that the device uses for the uplink to the core device. Use the **untagged** command to add the ports connected to the individual clients.

Configuring aggregated VLANs on a core device

To configure the aggregated VLANs on device C in [Figure 68](#) on page 291, enter the following commands.

```
PowerConnect(config)# tag-type 9100
PowerConnect(config)# aggregated-vlan
PowerConnect(config)# vlan 101 by port
PowerConnect(config-vlan-101)# tagged ethernet 7
PowerConnect(config-vlan-101)# untagged ethernet 8
PowerConnect(config-vlan-101)# exit
PowerConnect(config)# vlan 102 by port
PowerConnect(config-vlan-102)# tagged ethernet 7
PowerConnect(config-vlan-102)# untagged ethernet 9
PowerConnect(config-vlan-102)# exit
PowerConnect(config)# write memory
```

Syntax: [no] tag-type <num>

Syntax: [no] aggregated-vlan

The <num> parameter specifies the tag type can be a hexadecimal value from 0 – ffff. The default is 8100.

Verifying the configuration

You can verify the VLAN, VLAN aggregation option, and tag configuration by viewing the running-config. To display the running-config, enter the **show running-config** command from any CLI prompt. After you save the configuration changes to the startup-config, you also can display the settings in that file by entering the **show configuration** command from any CLI prompt.

Complete CLI examples

The following sections show all the Aggregated VLAN configuration commands on the devices in [Figure 68](#) on page 291.

NOTE

In these examples, the configurations of the edge devices (A, B, E, and F) are identical. The configurations of the core devices (C and D) also are identical. The aggregated VLAN configurations of the edge and core devices on one side must be symmetrical (in fact, a mirror image) to the configurations of the devices on the other side. For simplicity, the example in [Figure 68](#) on page 291 is symmetrical in terms of the port numbers. This allows the configurations for both sides of the link to be the same. If your configuration does not use symmetrically arranged port numbers, the configurations should not be identical but must use the correct port numbers.

Commands for device A

```
PowerConnectA(config)# vlan 101 by port
PowerConnectA(config-vlan-101)# tagged ethernet 6
PowerConnectA(config-vlan-101)# untagged ethernet 1
PowerConnectA(config-vlan-101)# exit
PowerConnectA(config)# vlan 102 by port
PowerConnectA(config-vlan-102)# tagged ethernet 6
PowerConnectA(config-vlan-102)# untagged ethernet 2
PowerConnectA(config-vlan-102)# exit
PowerConnectA(config)# vlan 103 by port
PowerConnectA(config-vlan-103)# tagged ethernet 6
PowerConnectA(config-vlan-103)# untagged ethernet 3
PowerConnectA(config-vlan-103)# exit
PowerConnectA(config)# vlan 104 by port
PowerConnectA(config-vlan-104)# tagged ethernet 6
PowerConnectA(config-vlan-104)# untagged ethernet 4
PowerConnectA(config-vlan-104)# exit
PowerConnectA(config)# vlan 105 by port
PowerConnectA(config-vlan-105)# tagged ethernet 6
PowerConnectA(config-vlan-105)# untagged ethernet 5
PowerConnectA(config-vlan-105)# exit
vA(config)# write memory
```

Commands for device B

The commands for configuring device B are identical to the commands for configuring device A. Notice that you can use the same channel VLAN numbers on each device. The devices that aggregate the VLANs into a path can distinguish between the identically named channel VLANs based on the ID of the path VLAN.

```
PowerConnectB(config)# vlan 101 by port
PowerConnectB(config-vlan-101)# tagged ethernet 6
PowerConnectB(config-vlan-101)# untagged ethernet 1
PowerConnectB(config-vlan-101)# exit
PowerConnectB(config)# vlan 102 by port
PowerConnectB(config-vlan-102)# tagged ethernet 6
PowerConnectB(config-vlan-102)# untagged ethernet 2
PowerConnectB(config-vlan-102)# exit
PowerConnectB(config)# vlan 103 by port
PowerConnectB(config-vlan-103)# tagged ethernet 6
PowerConnectB(config-vlan-103)# untagged ethernet 3
PowerConnectB(config-vlan-103)# exit
PowerConnectB(config)# vlan 104 by port
PowerConnectB(config-vlan-104)# tagged ethernet 6
PowerConnectB(config-vlan-104)# untagged ethernet 4
PowerConnectB(config-vlan-104)# exit
```

```
PowerConnectB(config)# vlan 105 by port
PowerConnectB(config-vlan-105)# tagged ethernet 6
PowerConnectB(config-vlan-105)# untagged ethernet 5
PowerConnectB(config-vlan-105)# exit
PowerConnectB(config)# write memory
```

Commands for device C

Since device C is aggregating channel VLANs from devices A and B into a single path, you need to change the tag type and enable VLAN aggregation.

```
PowerConnectC(config)# tag-type 9100
PowerConnectC(config)# aggregated-vlan
PowerConnectC(config)# vlan 101 by port
PowerConnectC(config-vlan-101)# tagged ethernet 7
PowerConnectC(config-vlan-101)# untagged ethernet 8
PowerConnectC(config-vlan-101)# exit
PowerConnectC(config)# vlan 102 by port
PowerConnectC(config-vlan-102)# tagged ethernet 7
PowerConnectC(config-vlan-102)# untagged ethernet 9
PowerConnectC(config-vlan-102)# exit
PowerConnectC(config)# write memory
```

Commands for device D

Device D is at the other end of path and separates the channels back into individual VLANs. The tag type must be the same as tag type configured on the other core device (Device C). In addition, VLAN aggregation also must be enabled.

```
PowerConnectD(config)# tag-type 9100
PowerConnectD(config)# aggregated-vlan
PowerConnectD(config)# vlan 101 by port
PowerConnectD(config-vlan-101)# tagged ethernet 7
PowerConnectD(config-vlan-101)# untagged ethernet 8
PowerConnectD(config-vlan-101)# exit
PowerConnectD(config)# vlan 102 by port
PowerConnectD(config-vlan-102)# tagged ethernet 7
PowerConnectD(config-vlan-102)# untagged ethernet 9
PowerConnectD(config-vlan-102)# exit
PowerConnectD(config)# write memory
```

Commands for device E

Since the configuration in [Figure 68](#) on page 291 is symmetrical, the commands for configuring device E are identical to the commands for configuring device A.

```
PowerConnectE(config)# vlan 101 by port
PowerConnectE(config-vlan-101)# tagged ethernet 6
PowerConnectE(config-vlan-101)# untagged ethernet 1
PowerConnectE(config-vlan-101)# exit
PowerConnectE(config)# vlan 102 by port
PowerConnectE(config-vlan-102)# tagged ethernet 6
PowerConnectE(config-vlan-102)# untagged ethernet 2
PowerConnectE(config-vlan-102)# exit
PowerConnectE(config)# vlan 103 by port
PowerConnectE(config-vlan-103)# tagged ethernet 6
PowerConnectE(config-vlan-103)# untagged ethernet 3
PowerConnectE(config-vlan-103)# exit
```

```
PowerConnectE(config)# vlan 104 by port
PowerConnectE(config-vlan-104)# tagged ethernet 6
PowerConnectE(config-vlan-104)# untagged ethernet4
PowerConnectE(config-vlan-104)# exit
PowerConnectE(config)# vlan 105 by port
PowerConnectE(config-vlan-105)# tagged ethernet 6
PowerConnectE(config-vlan-105)# untagged ethernet 5
PowerConnectE(config-vlan-105)# exit
PowerConnectE(config)# write memory
```

Commands for device F

The commands for configuring device F are identical to the commands for configuring device E. In this example, since the port numbers on each side of the configuration in [Figure 68](#) on page 291 are symmetrical, the configuration of device F is also identical to the configuration of device A and device B.

```
PowerConnectF(config)# vlan 101 by port
PowerConnectF(config-vlan-101)# tagged ethernet 6
PowerConnectF(config-vlan-101)# untagged ethernet 1
PowerConnectF(config-vlan-101)# exit
PowerConnectF(config)# vlan 102 by port
PowerConnectF(config-vlan-102)# tagged ethernet 6
PowerConnectF(config-vlan-102)# untagged ethernet 2
PowerConnectF(config-vlan-102)# exit
PowerConnectF(config)# vlan 103 by port
PowerConnectF(config-vlan-103)# tagged ethernet 6
PowerConnectF(config-vlan-103)# untagged ethernet 3
PowerConnectF(config-vlan-103)# exit
PowerConnectF(config)# vlan 104 by port
PowerConnectF(config-vlan-104)# tagged ethernet 6
PowerConnectF(config-vlan-104)# untagged ethernet 4
PowerConnectF(config-vlan-104)# exit
PowerConnectF(config)# vlan 105 by port
PowerConnectF(config-vlan-105)# tagged ethernet 6
PowerConnectF(config-vlan-105)# untagged ethernet 5
PowerConnectF(config-vlan-105)# exit
PowerConnectF(config)# write memory
```

Configuring 802.1Q-in-Q tagging

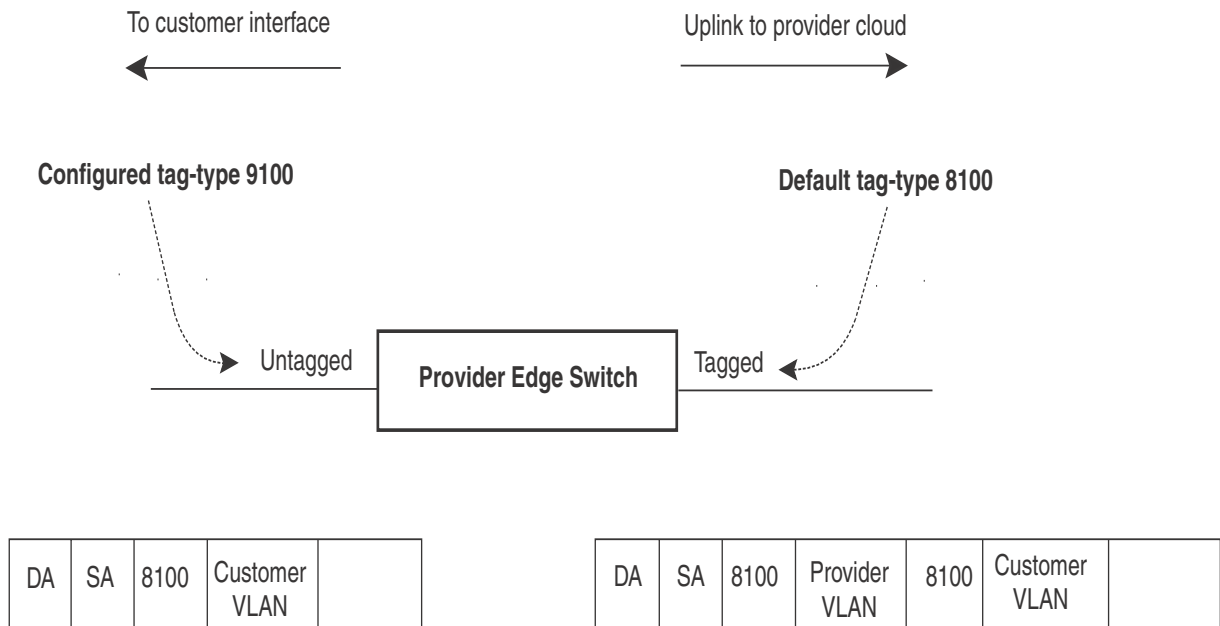
802.1Q-in-Q tagging provides finer granularity for configuring 802.1Q tagging, enabling you to configure 802.1Q tag-types on a group of ports. This feature allows you to create two identical 802.1Q tags (802.1Q-in-Q tagging) on a single device. This enhancement improves SAV interoperability between devices and other vendors' devices that support the 802.1Q tag-types, but are not very flexible with the tag-types they accept.

NOTE

Devices treat a double-tagged Ethernet frame as a Layer 2 only frame. The packets are not inspected for Layer 3 and Layer 4 information, and operations are not performed on the packet utilizing Layer 3 or Layer 4 information.

[Figure 69](#) shows an example application with 802.1Q-in-Q tagging.

FIGURE 69 802.1Q-in-Q configuration example



In [Figure 69](#), the untagged ports (to customer interfaces) accept frames that have any 802.1Q tag other than the configured tag-type 9100. These packets are considered untagged on this incoming port and are re-tagged when they are sent out of the uplink towards the provider. The 802.1Q tag-type on the uplink port is 8100, so the device will switch the frames to the uplink device with an additional 8100 tag, thereby supporting devices that only support this method of VLAN tagging.

Configuration rules

- Since the uplink (to the provider cloud) and the edge link (to the customer port) must have different 802.1Q tags, make sure the uplink and edge link are in different port regions. Refer to [“Enabling or disabling the Spanning Tree Protocol \(STP\)”](#) on page 175 for a list of valid port regions.
- If you configure a port with an 802.1Q tag-type, the device automatically applies the 802.1Q tag-type to all ports within the same port region. Likewise, if you remove the 802.1Q tag-type from a port, the device automatically removes the 802.1Q tag-type from all ports within the same port region.
- PowerConnect B-Series TI24X devices support one configured tag-type per port along with the default tag-type of 8100. PowerConnect devices do not have the port region concept and do not support tag-profile.
- 802.1Q-in-Q tagging and VSRP are not supported together on the same device.

Enabling 802.1Q-in-Q tagging

To enable 802.1Q-in-Q tagging, configure an 802.1Q tag on the untagged edge links (the customer ports) to any value other than the 802.1Q tag for incoming traffic. For example, in [Figure 70](#), the 802.1Q tag on the untagged edge links (ports 11 and 12) is 9100, whereas, the 802.1Q tag for incoming traffic is 8100.

10 Configuring 802.1Q-in-Q tagging

To configure 802.1 Q-in-Q tagging as shown in [Figure 70](#), enter commands such as the following on the untagged edge links of devices C and D.

```
PowerConnect(config)# tag-type 9100 e 11 to 12
```

Syntax: **[no] tag-type** <num> **[ethernet** <port number> **[to** <port number>]]

The <num> parameter specifies the tag-type number and can be a hexadecimal value from 0 - ffff. The default is 8100.

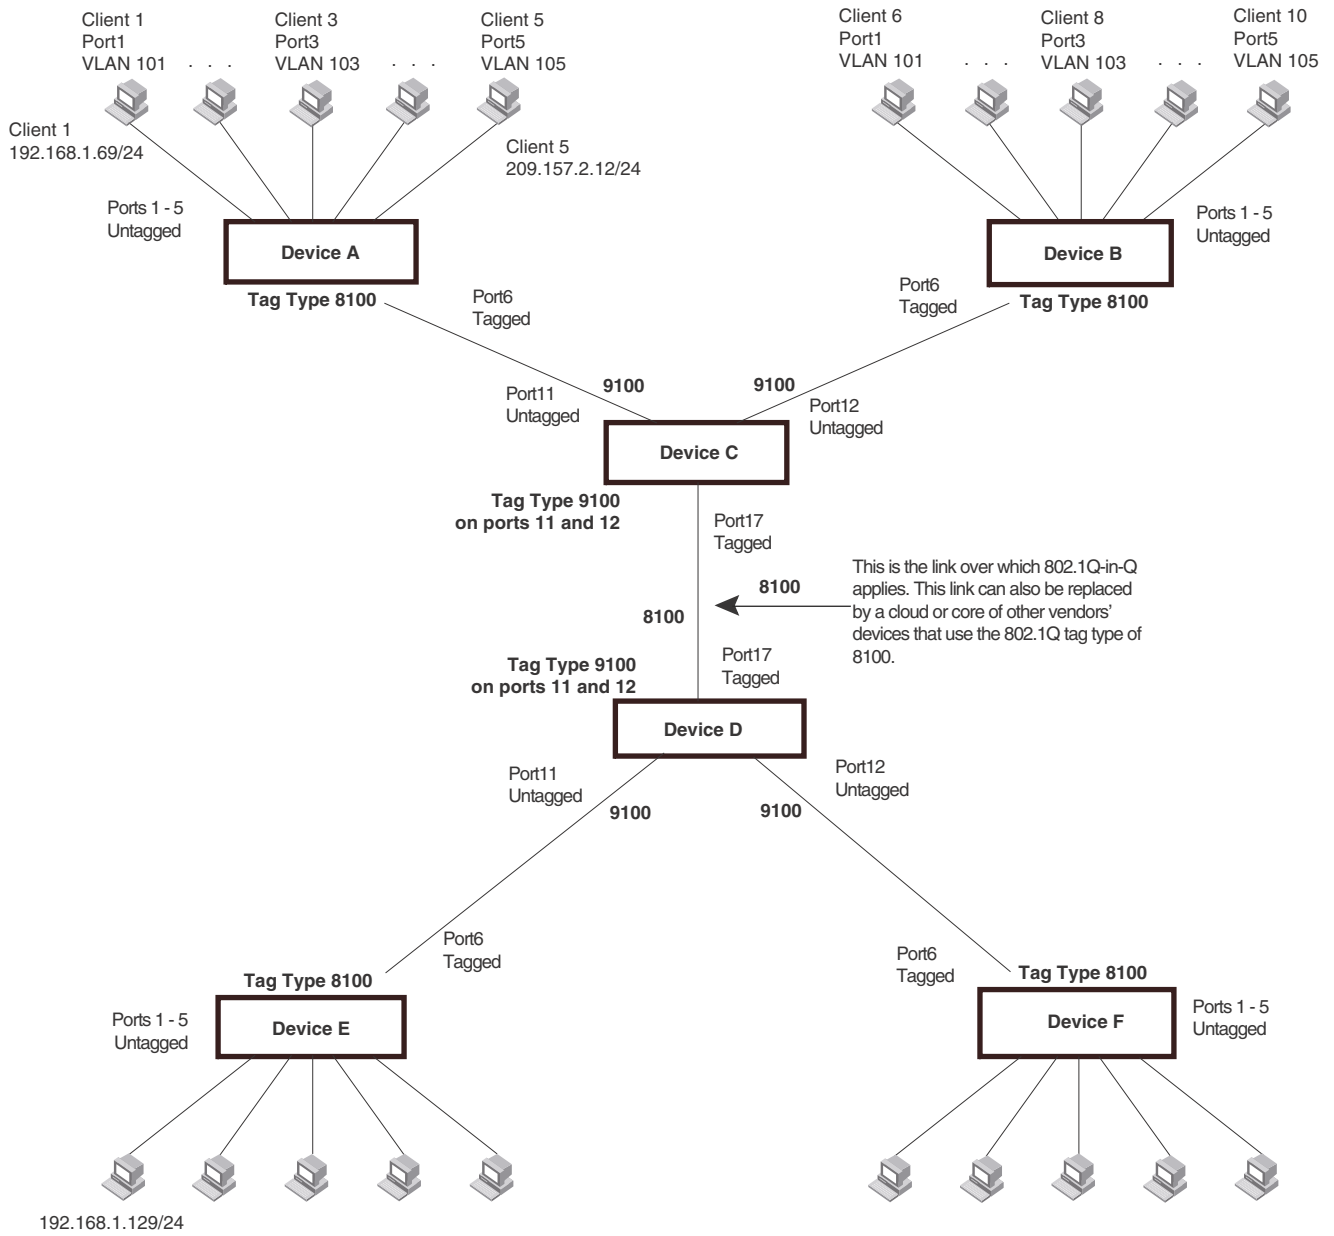
The **ethernet** <port number> to <port number> parameter specifies the ports that will use the defined 802.1Q tag. This parameter operates with the following rules:

- If you do not specify a port or range of ports, the 802.1Q tag applies to all Ethernet ports on the device.

Example configuration

[Figure 70](#) shows an example 802.1Q-in-Q configuration.

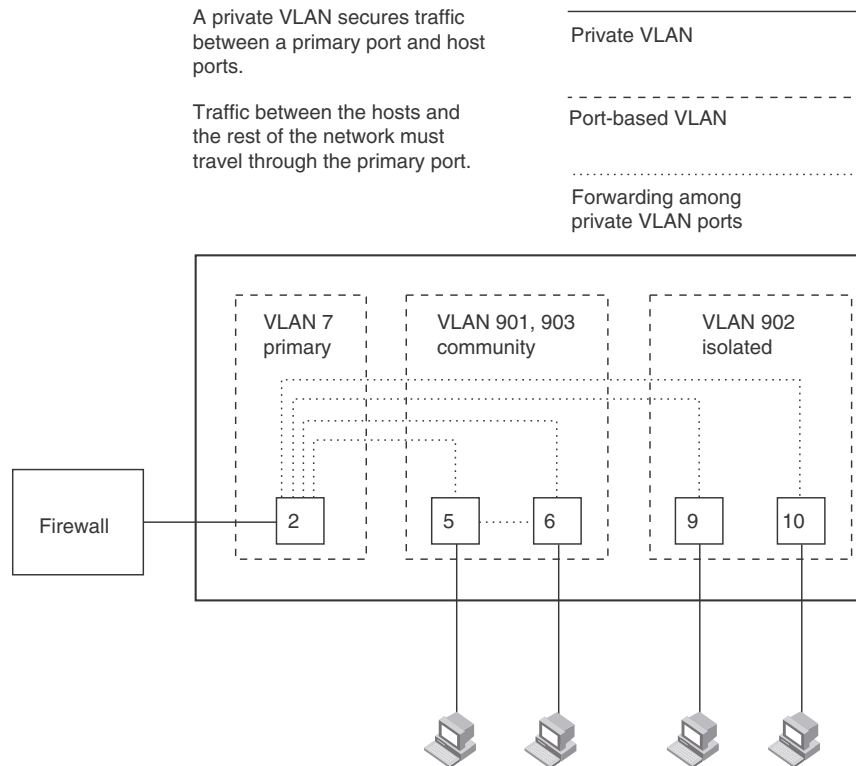
FIGURE 70 Example 802.1Q-in-Q configuration



Configuring private VLANs

A private VLAN is a VLAN that has the properties of standard Layer 2 port-based VLANs but also provides additional control over flooding packets on a VLAN. [Figure 71](#) shows an example of an application using a private VLAN.

FIGURE 71 Private VLAN used to secure communication between a workstation and servers



This example uses a private VLAN to secure traffic between hosts and the rest of the network through a firewall. Five ports in this example are members of a private VLAN. The first port (port 2) is attached to a firewall. The next four ports (ports 5, 6, 9, and 10) are attached to hosts that rely on the firewall to secure traffic between the hosts and the rest of the network. In this example, two of the hosts (on ports 5 and 6) are in a community private VLAN, and thus can communicate with one another as well as through the firewall. The other two hosts (on ports 9 and 10), are in an isolated VLAN and thus can communicate only through the firewall. The two hosts are secured from communicating with one another even though they are in the same VLAN.

By default, the private VLAN does not forward broadcast or unknown-unicast packets from outside sources into the private VLAN. If needed, you can override this behavior for broadcast packets, unknown-unicast packets.

You can configure a combination of the following types of private VLANs:

- Primary – The primary private VLAN ports are “promiscuous”. They can communicate with all the isolated private VLAN ports and community private VLAN ports in the isolated and community VLANs that are mapped to the promiscuous port.

- Secondary – The secondary private VLAN are secure VLANs that are separated from the rest of the network by the primary private VLAN. Every secondary private VLAN is associated with a primary private VLAN. The two types of secondary private VLANs are isolated private VLAN and community private VLAN.
 - Isolated – Broadcasts and unknown-unicasts packet received on isolated ports are sent only to the primary port. They are not flooded to other ports in the isolated VLAN.
On PowerConnect B-Series TI24X devices, the broadcasts, unknown-unicasts, and unregistered-multicast packets received on isolated ports are sent to the primary port. They are not flooded to other ports in the isolated VLAN.
 - Community – Broadcasts and unknown unicasts received on community ports are sent to the primary port and also are flooded to the other ports in the community VLAN.
On PowerConnect B-Series TI24X devices, the broadcasts, unknown unicasts, and unregistered multicast received on community ports are sent to the primary port and are flooded to the other ports in the community VLAN.

Each private VLAN must have a primary VLAN. The primary VLAN is the interface between the secured ports and the rest of the network. The private VLAN can have any combination of community and isolated VLANs. The community VLAN and isolated VLAN cannot forward traffic to each other. You cannot forwarding traffic between different private VLANs.

[Table 42](#) list the differences between private VLANs and standard VLANs.

TABLE 42 Comparison of private VLANs and standard port-based VLANs

Forwarding behavior	Private VLANs	Standard VLANs
All ports within a VLAN constitute a common Layer broadcast domain	No	Yes
Broadcasts and unknown unicasts are forwarded to all the VLAN ports by default	No (isolated VLAN) Yes (community VLAN)	Yes
Known unicasts	Yes (forwarding is done only between ports of the same community VLAN and the primary VLAN port)	Yes

Configuration notes

NOTE

PowerConnect B-Series TI24X devices support 802.1Q tagged ports on private VLAN. Private VLAN is a hardware-based feature. Private VLANs on the PowerConnect device forwards unknown-unicast, unregistered multicast, and broadcast in hardware.

- Normally, in any port-based VLAN, the device floods unknown unicast, unregistered multicast, and broadcast packets in hardware, although selective packets, such as IGMP, may be sent to only to the CPU for analysis, based on the IGMP snooping configuration. When Protocol or Subnet VLANs are enabled, or if private VLAN mappings are enabled, the device will flood unknown unicast, unregistered multicast, and broadcast packets in software.
- There is currently no support for IGMP snooping within private VLANs. In order for clients in private VLANs to receive multicast traffic, IGMP snooping must be disabled so that all multicast packets are treated as unregistered packets and are flooded in software to all the ports.

- PowerConnect B-Series TI24X forward all known unicast traffic in hardware. PowerConnect B-Series TI24X devices, multiple MAC entries do not appear in the MAC address table because the PowerConnect B-Series TI24X transparently manages multiple MAC entries in hardware.
- You can configure private VLANs and dual-mode VLAN ports on the same device. However, the dual-mode VLAN ports cannot be members of private VLANs.

Configuration notes and limitations for PowerConnect devices

Consider the following statements when configuring a private VLAN on a PowerConnect B-Series TI24X device:

- Each private VLAN can have multiple isolated VLANs or community VLANs. You can use any combination of isolated or community VLANs with the primary VLAN.
- You cannot configure a common port in any pair of primary, isolated, and community VLANs.
- You cannot configure the default VLAN (VLAN 1) as a private VLAN.
- You can configure Virtual interface (VE) on primary VLAN. The VE configuration in secondary private VLAN is not supported.
- A secondary VLAN can have only one primary VLAN associated with it.
- You can configure static trunk or dynamic trunk in both primary and secondary private VLAN.
- Spanning tree protocol is automatically disabled on private VLAN.
- You cannot enable Metro Ring Protocol (MRP), Virtual Switch Redundancy Protocol (VSRP), and Virtual Router Redundancy Protocol (VRRP) on private VLAN.
- You cannot enable single span on a private VLAN configured system and you cannot configure private VLAN on a single span enabled system.
- You cannot configure private VLAN through VLAN group.
- ICMP redirect is automatically disabled for private VLAN.
- To enhance private VLAN security, unique VLAN ID is used to identify the primary VLAN for different private VLANs. The known unicast among isolated VLAN and known unicast across different isolated or community VLAN is restricted.
- To enhance private VLAN flexibility, tagged, untagged, and dual mode ports are supported.
- All ports in the primary VLAN are promiscuous. You cannot configure an individual port in the primary to secondary private VLAN mapping. Traffic is forwarded to all ports in the primary private vlan when received from a secondary VLAN.

Command syntax

To configure a private VLAN, configure each of the component VLANs (isolated, community, and primary) as a separate port-based VLAN:

- Use standard VLAN configuration commands to create the VLAN and add ports.
- Identify the private VLAN type (isolated, community, or primary)
- For the primary VLAN, map the other private VLANs to the ports in the primary VLAN

Configuring an isolated or community private VLAN

To configure a community private VLAN, enter commands such as the following.

```
PowerConnect(config)# vlan 901
PowerConnect(config-vlan-901)# untagged ethernet 5 to 6
PowerConnect(config-vlan-901)# pvlan type community
```

These commands create port-based VLAN 901, add ports 5 and 6 to the VLAN as untagged ports, then specify that the VLAN is a community private VLAN.

Syntax: `untagged ethernet <portnum> [to [<portnum> | ethernet <portnum>]`

Syntax: `[no] pvlan type community | isolated | primary`

The **untagged** command adds the ports to the VLAN.

The **pvlan type** command specifies that this port-based VLAN is a private VLAN:

- **community** – Broadcasts and unknown unicasts received on community ports are sent to the primary port and also are flooded to the other ports in the community VLAN.
- **isolated** – Broadcasts and unknown unicasts received on isolated ports are sent only to the primary port. They are not flooded to other ports in the isolated VLAN.
- **primary** – The primary private VLAN ports are “promiscuous”. They can communicate with all the isolated private VLAN ports and community private VLAN ports in the isolated and community VLANs that are mapped to the promiscuous port.

To configure a community private VLAN on a PowerConnect B-Series TI24X device enter the following commands.

```
PowerConnect(config)# vlan 901
PowerConnect(config-vlan-901)# untagged ethernet 5 to 6
PowerConnect(config-vlan-901)# pvlan type community
```

These commands create port-based VLAN 901, add ports 5 and 6 to the VLAN as untagged ports, and then specify that the VLAN is a community private VLAN.

To configure an isolated private VLAN on a PowerConnect B-Series TI24X device, enter the following commands.

```
PowerConnect(config)# vlan 902
PowerConnect(config-vlan-902)# tagged ethernet 9 to 10
PowerConnect(config-vlan-902)# pvlan type isolated
```

Configuring the primary VLAN

NOTE

In PowerConnect B-Series TI24X devices, all the ports in the primary private VLAN are active.

Syntax: `untagged ethernet <portnum> [to <portnum> | ethernet<portnum>]`

Syntax: `[no] pvlan type community | isolated | primary`

Syntax: `[no] pvlan mapping <vlan-id> ethernet <portnum>`

The **untagged** command adds the ports to the VLAN.

The **pvlan type** command specifies that this port-based VLAN is a private VLAN. Specify **primary** as the type.

The **pvlan mapping** command identifies the other private VLANs for which this VLAN is the primary. The command also specifies the primary VLAN ports to which you are mapping the other private VLANs. The parameters of the **pvlan mapping** command are as follows:

- The **<vlan-id>** parameter specifies another private VLAN. The other private VLAN you want to specify must already be configured.
- The **ethernet <portnum>** parameter specifies the primary VLAN port to which you are mapping all the ports in the other private VLAN (the one specified by **<vlan-id>**).

To configure a primary private VLAN, on PowerConnect B-Series TI24X device enter the following commands:

```
PowerConnect(config)# vlan 7
PowerConnect(config-vlan-7)# untagged ethernet 2
PowerConnect(config-vlan-7)# pvlan type primary
PowerConnect(config-vlan-7)# pvlan mapping 901
```

These commands create port-based VLAN 7, add port 2 as an untagged port, identify the VLAN as the primary VLAN in a private VLAN, and map the secondary private VLANs to the primary VLAN.

NOTE

For PowerConnect devices, you do not have to specify the port in the primary vlan mapping.

CLI example for Figure 71

To configure the private VLANs shown in [Figure 71](#) on page 300, enter the following commands.

```
PowerConnect(config)# vlan 901
PowerConnect(config-vlan-901)# untagged ethernet 5 to 6
PowerConnect(config-vlan-901)# pvlan type community
PowerConnect(config-vlan-901)# exit
PowerConnect(config)# vlan 902
PowerConnect(config-vlan-902)# untagged ethernet 9 to 10
PowerConnect(config-vlan-902)# pvlan type isolated
PowerConnect(config-vlan-902)# exit
PowerConnect(config)# vlan 903
PowerConnect(config-vlan-903)# untagged ethernet 7 to 8
PowerConnect(config-vlan-903)# pvlan type community
PowerConnect(config-vlan-903)# exit
PowerConnect(config)# vlan 7
PowerConnect(config-vlan-7)# untagged ethernet 2
PowerConnect(config-vlan-7)# pvlan type primary
PowerConnect(config-vlan-7)# pvlan mapping 901
PowerConnect(config-vlan-7)# pvlan mapping 902
PowerConnect(config-vlan-7)# pvlan mapping 903
```

Enabling broadcast, unregistered multicast or unknown unicast traffic to the private VLAN on PowerConnect device

To configure the ports in primary VLAN to forward broadcast, unregistered multicast, or unknown unicast, enter the following commands:

```
PowerConnect(config-vlan-7)# pvlan preference unknown-unicast flood
PowerConnect(config-vlan-7)# pvlan preference broadcast flood
```

Syntax: [no] pvlan preference broadcast | unknown-unicast

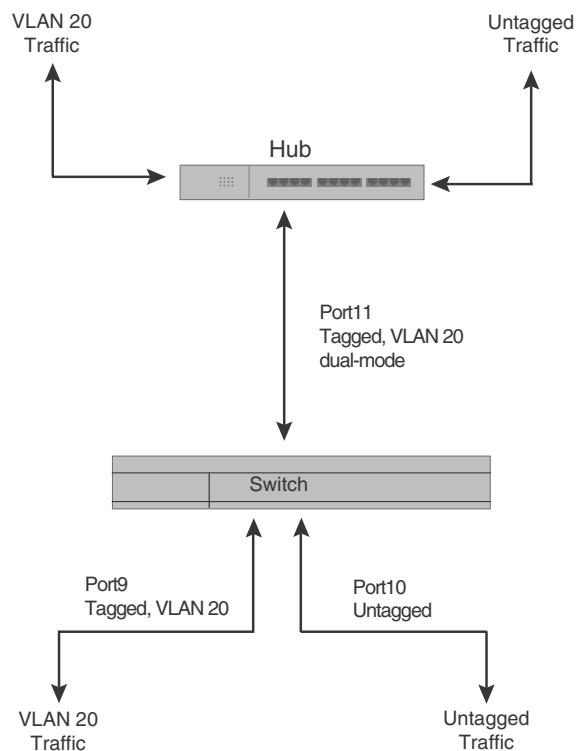
These commands enable forwarding of broadcast, unregistered multicast flood, and unknown-unicast packets to ports within the private VLAN.

Dual-mode VLAN ports

Configuring a tagged port as a **dual-mode** port allows it to accept and transmit both tagged traffic and untagged traffic at the same time. A dual-mode port accepts and transmits frames belonging to VLANs configured for the port, as well as frames belonging to the default VLAN (that is, untagged traffic).

For example, in [Figure 72](#), port 11 is a dual-mode port belonging to VLAN 20. Traffic for VLAN 20, as well as traffic for the default VLAN, flows from a hub to this port. The dual-mode feature allows traffic for VLAN 20 and untagged traffic to go through the port at the same time.

FIGURE 72 Dual-mode VLAN port example



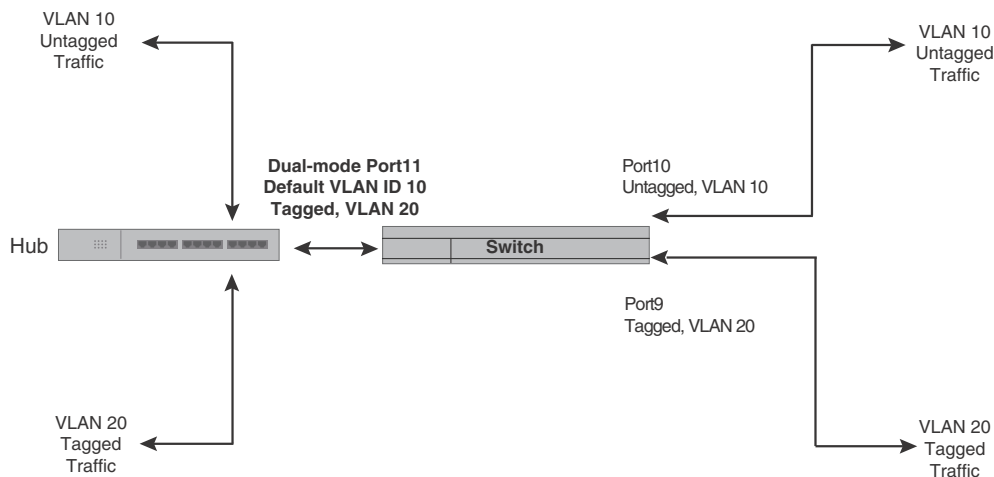
To enable the dual-mode feature on port 11 in [Figure 72](#), enter the following commands.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# tagged e 11
PowerConnect(config-vlan-20)# tagged e 9
PowerConnect(config-vlan-20)# int e 11
PowerConnect(config-if-e10000-11)# dual-mode
PowerConnect(config-if-e10000-11)# exit
```

Syntax: [no] dual-mode

You can configure a dual-mode port to transmit traffic for a specified VLAN (other than the DEFAULT-VLAN) as untagged, while transmitting traffic for other VLANs as tagged. [Figure 73](#) illustrates this enhancement.

FIGURE 73 Specifying a default VLAN ID for a dual-mode port



In [Figure 73](#), tagged port 11 is a dual-mode port belonging to VLANs 10 and 20. The default VLAN assigned to this dual-mode port is 10. This means that the port transmits tagged traffic on VLAN 20 (and all other VLANs to which the port belongs) and transmits untagged traffic on VLAN 10.

The dual-mode feature allows tagged traffic for VLAN 20 and untagged traffic for VLAN 10 to go through port 11 at the same time. A dual-mode port transmits only untagged traffic on its default VLAN (that is, either VLAN 1, or a user-specified VLAN ID), and only tagged traffic on all other VLANs.

The following commands configure VLANs 10 and 20 in [Figure 73](#). Tagged port 11 is added to VLANs 10 and 20, then designated a dual-mode port whose specified default VLAN is 10. In this configuration, port 11 transmits only untagged traffic on VLAN 10 and only tagged traffic on VLAN 20.

```
PowerConnect(config)# vlan 10 by port
PowerConnect(config-vlan-10)# untagged e 10
PowerConnect(config-vlan-10)# tagged e 11
PowerConnect(config-vlan-10)# exit
PowerConnect(config)# vlan 20 by port
PowerConnect(config-vlan-20)# tagged e 9
PowerConnect(config-vlan-20)# tagged e 11
PowerConnect(config-vlan-20)# exit
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# dual-mode 10
PowerConnect(config-if-e10000-11)# exit
```

Syntax: [no] dual-mode [<vlan-id>]

Notes:

- If you do not specify a <vlan-id> in the **dual mode** command, the port default VLAN is set to 1. The port transmits untagged traffic on the DEFAULT-VLAN.

- The dual-mode feature is disabled by default. Only tagged ports can be configured as dual-mode ports.
- In trunk group, either all of the ports must be dual-mode, or none of them can be.

The **show vlan** command displays a separate row for dual-mode ports on each VLAN.

Example

```
PowerConnect# show vlan
Total PORT-VLAN entries: 3
Maximum PORT-VLAN entries: 16

Legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S1) 1 2 3 4 5 6 7 8
  Untagged Ports: (S2) 1 2 3 4 5 6 7 8 12 13 14 15 16 17 18 19
  Untagged Ports: (S2) 20 21 22 23 24
  Tagged Ports: None
  Uplink Ports: None
DualMode Ports: None
PORT-VLAN 10, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: (S2) 10
  Tagged Ports: None
  Uplink Ports: None
DualMode Ports: (S2) 11
PORT-VLAN 20, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: None
  Tagged Ports: (S2) 9
  Uplink Ports: None
DualMode Ports: (S2) 11
```

Displaying VLAN information

After you configure the VLANs, you can verify the configuration using the **show** commands described in this section.

NOTE

If a VLAN name begins with "GVRP_VLAN_", the VLAN was created by the GARP VLAN Registration Protocol (GVRP). If a VLAN name begins with "STATIC_VLAN_", the VLAN was created by GVRP and then was converted into a statically configured VLAN.

Displaying VLANs in alphanumeric order

VLANs are displayed in alphanumeric order, as shown in the following example.

10 Displaying VLAN information

```
PowerConnect# show run
Current configuration:
!
ver 4.2.00b
!
!
global-stp
!
trunk ethe 1 to 2
trunk ethe 9 ethe 13
  config-trunk-ind
  disable ethe 13
!
vlan 202 by port
  tagged ethe 5 to 6
  router-interface ve 22
!
vlan 203 by port
  tagged ethe 5 to 6
  router-interface ve 23
!
vlan 204 by port
  tagged ethe 5 to 6
  router-interface ve 24
!
vlan 205 by port
  tagged ethe 5 to 6
  router-interface ve 25
!
```

Displaying system-wide VLAN information

Use the **show vlans** command to display VLAN information for all the VLANs configured on the device.

The following example shows the display output for a PowerConnect B-Series TI24X device.

```
PowerConnect# show vlans
Total PORT-VLAN entries: 2
Maximum PORT-VLAN entries: 8
legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S2) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
  Untagged Ports: (S2) 17 18 19 20 21 22 23 24
  Untagged Ports: (S4) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
  Untagged Ports: (S4) 17 18 19 20 21 22 23 24
  Tagged Ports: None

PORT-VLAN 10, Name IP_VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S1) 1 2 3 4 5 6
  Tagged Ports: None
```

Ports that are tagged but are not dual-mode ports are listed as tagged ports. In the following example display output, ports 7 and 8 are dual-mode ports in port-based VLAN 4. Ports 7 and 8 also belong to port-based VLAN 3, but they are tagged ports only in VLAN 3 and are not configured as dual-mode ports.

```

PowerConnect# show vlan 4
Total PORT-VLAN entries: 5
Maximum PORT-VLAN entries: 3210
PORT-VLAN 4, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: None
  Tagged Ports:   6   9  10  11
  Uplink Ports: None
  DualMode Ports:  7   8
PowerConnect# show vlan 3
Total PORT-VLAN entries: 5
Maximum PORT-VLAN entries: 3210
PORT-VLAN 3, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: None
  Tagged Ports:   6   7   8   9  10
  Uplink Ports: None
  DualMode Ports: None
    
```

Syntax: `show vlans [<vlan-id> | ethernet <portnum>]`

The *<vlan-id>* parameter specifies a VLAN for which you want to display the configuration information.

The *<portnum>* parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port.

Displaying VLAN information for specific ports

Use one of the following methods to display VLAN information for specific ports.

To display VLAN information for all the VLANs of which port 1 is a member, enter the following command.

```

PowerConnect# show vlans e 1
Total PORT-VLAN entries: 3
Maximum PORT-VLAN entries: 8

legend: [S=Slot]

PORT-VLAN 100, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: (S7) 1 2 3 4
  Tagged Ports: None
    
```

Syntax: `show vlans [<vlan-id> | ethernet []<portnum>]`

The *<vlan-id>* parameter specifies a VLAN for which you want to display the configuration information.

The *<portnum>* parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port.

10 Displaying VLAN information

Configuring Trunk Groups and Dynamic Link Aggregation 11

Trunk group overview

The Trunk group feature allows you to manually configure multiple high-speed load-sharing links between two Layer 2 Switches or Layer 3 Switches or between a Layer 2 Switch and Layer 3 Switch and a server.

In addition to enabling load sharing of traffic, trunk groups provide redundant, alternate paths for traffic if any of the segments fail.


- **Trunk groups** are manually-configured aggregate links containing multiple ports.
- **802.3ad link aggregation** is a protocol that dynamically creates and manages trunk groups.

NOTE

You can use both types of trunking on the same device. However, you can use only one type of trunking for a given port. For example, you can configure port 1 as a member of a static trunk group or you can enable 802.3ad link aggregation on the port, but you cannot do both.

Figure 74 shows an example of a configuration that uses trunk groups.

FIGURE 74 Trunk group application within a network



NOTE


The ports in a trunk group make a single logical link. Therefore, all the ports in a trunk group must be connected to the same device at the other end.

Trunk group connectivity to a server

To support termination of a trunk group, the server must have either multiple network interface cards (NICs) or either a dual or quad interface card installed. The trunk server is designated as a server with multiple adapters or a single adapter with multiple ports that share the same MAC and IP address.

Figure 75 shows an example of a trunk group between a server and a device.

FIGURE 75 Trunk group between a server and switches



Trunk group rules

Table 43 lists the maximum number of trunk groups you can configure on a device and the valid number of ports in a trunk group. The table applies to static and LACP trunk ports.

TABLE 43 Trunk group support

Model	Maximum number of Gbps trunk groups	Valid number of ports in a group
PowerConnect B-Series TI24X	31	2, 3, 4, 5, 6, 7, or 8

- You cannot configure a port as a member of a trunk group if 802.3ad link aggregation is enabled on the port.
- Unlike other devices, trunk groups on devices listed in Table 43 are not classified as switch trunk groups or server trunk groups.
- On PowerConnect B-Series TI24X devices, port assignment on a module need not be consecutive. The port range can contain gaps. For example, you can configure ports 1, 3, and 4 (excluding 2). Refer to “Flexible trunk group membership” on page 314.
- Port assignment on a module need not be consecutive. The port range can contain gaps. For example, you can configure ports 1, 3, and 4 (excluding 2). Refer to “Flexible trunk group membership” on page 314.
- You can select any port to be the primary port of the trunk group.
- You cannot combine Gbps and 10-Gbps ports in the same trunk group.

- Make sure the device on the other end of the trunk link can support the same number of ports in the link. For example, if you configure a three-port trunk group on the device and the other end is a different type of switch, make sure the other switch can support a three-port trunk group.
- All the ports must be connected to the same device at the other end.
- All trunk group member properties must match the lead port of the trunk group with respect to the following parameters:
 - port tag type (untagged or tagged port)
 - statically configured port speed and duplex
 - QoS priority

To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.

Trunk group configuration examples

Figure 76 shows some examples of valid 2-port trunk group links between devices. The trunk groups in this example are switch trunk groups between two devices. Ports in a valid 2-port trunk group on one device are connected to two ports in a valid 2-port trunk group on another device. The same rules apply to 3-port, 4-port, etc., trunk groups.




FIGURE 76 Examples of 2-port and 3-port trunk groups

Flexible trunk group membership

This section describes flexible trunk group membership on the PowerConnect B-Series TI24X devices.

PowerConnect B-Series TI24X devices

Trunking is supported on non-consecutive ports in a module. For example, you can configure ports e 4, 6, and 7 (excluding e 5) together on a module as a trunk group. In releases prior to the above, the ports in a trunk group must be consecutive. This feature is supported on static and LACP trunk ports, as well as GbE and 10-GbE ports.

Viewing the first and last ports in a trunk group

Output for many of the **show** commands will show the first and last port in a trunk as FirstPort-LastPort, if the ports are consecutive, and FirstPort*LastPort if the ports are not consecutive.

With the configuration above, output from the **show mac** command resembles the following, which shows the first and last ports.

```
PowerConnect#show mac
Total active entries from all ports = 1
MAC-Address Port Type Index
0007.e910.c201 7*21 Dynamic 2920
```

For a trunk group with members 7 to 9, the output from the show mac command resembles the following.

```
PowerConnect#show mac
Total active entries from all ports = 1
MAC-Address Port Type Index
0007.e910.c201 7-9 Dynamic 2920
```

Trunk group load sharing

Trunk groups on the PowerConnect B-Series TI24X devices are not classified as switch trunk groups or server trunk groups.

Devices load-share across the ports in the trunk group. The method used for the load sharing depends on the following:

- Device type – Chassis device or Stackable device
- Traffic type – Layer 2 or Layer 3
- Software version your device is running

NOTE

Layer 3 routed IP is not load balanced. This traffic types will however still be forwarded on the trunk ports.

Note regarding IPv6

Devices that support IPv6 take a the IPv6 address for a packet into account when sharing traffic across a trunk group. The load sharing is performed in the same way it is for IPv4 addresses; that is; trunk types for which traffic load is shared based on IPv4 address information can now use IPv6 addresses to make the load sharing decision.

Load sharing for unknown unicast, multicast, and broadcast traffic

Devices load balance unknown unicast, multicast, and broadcast traffic based on the source port and VLAN ID and not on any source or destination information in the packet.

For example, when the switch receives unknown unicast, multicast, and broadcast packets, and the packets are from the same source port, the packets are forwarded to the same port of the trunk group. Conversely, when the switch receives unknown unicast, multicast, and broadcast packets, and the packets are from different source ports, the packets are load-balanced across all the ports of the trunk group.

Note that this does not apply to known unicast traffic, which is always load balanced across all the ports of a trunk group based on the traffic's Layer 2 and Layer 3 source and destination parameters.

How trunk load sharing works

Load balancing procedures differ depending on the software version your device is running. In earlier releases, the device load balances all bridged traffic based on source and destination MAC addresses. In subsequent releases, the load balancing method for bridged traffic varies depending on the traffic type.

NOTE

There is no change in load balancing for routed traffic, which is always based on the source and destination IP addresses and protocol field .

NOTE

Load balancing on the PowerConnect B-Series TI24X is hardware-based.

[Table 44](#) describes how the PowerConnect B-Series TI24X devices load balance traffic.

TABLE 44 Trunk group load sharing on PowerConnect B-Series TI24X devices

Traffic type	Load balancing method
L2 Bridged Non-IP	Source MAC, Destination MAC, Module ID (0 for 24-port device), Ingress port, VLAN ID (default VLAN ID used for untagged packets), EtherType
L2 Bridged IPv4 TCP/UDP	Source IP, Destination IP, Source TCP/UDP Port, Destination TCP/UDP Port
L2 Bridged IPv4 Non-TCP/UDP	Source IP, Destination IP
L2 Bridged IPv6 TCP/UDP	Source IP, Destination IP, Source TCP/UDP Port, Destination TCP/UDP Port
L2 Bridged IPv6 Non-TCP/UDP	Source IP, Destination IP

Configuring a trunk group

Follow the steps given below to configure a trunk group.

1. Disconnect the cables from those ports on both systems that will be connected by the trunk group. Do not configure the trunk groups with the cables connected.

NOTE

If you connect the cables before configuring the trunk groups and then reboot, the traffic on the ports can create a spanning tree loop.

2. Configure the trunk group on one of the two Layer 2 Switches or Layer 3 Switches involved in the configuration.

NOTE

Downtime is incurred when adding a new port to a trunk group. It is suggested that you schedule the addition of ports to a trunk group to minimize downtime and its impact to the production network.

3. Save the configuration changes to the startup-config file.
4. Dynamically place the new trunk configuration into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.
5. If the device at the other end of the trunk group is another Layer 2 Switch or Layer 3 Switch, repeat Steps 2 – 4 for the other device.
6. When the trunk groups on both devices are operational, reconnect the cables to those ports that are now configured as trunk groups, starting with the first port (lead port) of each trunk group.
7. To verify the link is operational, use the **show trunk** command.

CLI syntax

The following is the CLI syntax for creating a trunk group. Configuration examples are shown in later sections of this chapter.

Syntax: `[no] trunk ethernet <primary-portnum> to <portnum> [ethernet <primary-portnum> to <portnum>]`

Syntax: `trunk deploy`

Each **ethernet** parameter introduces a port group. Enter the slot number (if applicable) and the port number of the Ethernet port.

The `<primary-portnum>` **to** `<portnum>` parameters specify a port group. Notice that each port group must begin with a primary port. After you enter this command, the primary port of the first port group specified (which must be the group with the lower port numbers) becomes the primary port for the entire trunk group.

To configure a trunk group consisting of two groups of two ports each, enter commands such as the following.

```
PowerConnect(config)#trunk ethernet 1 to 2 ethernet 3 to 4
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

Example 1: Configuring the trunk groups shown in Figure 75

To configure the trunk groups shown in [Figure 74](#), enter the following commands. Notice that the commands are entered on multiple devices.

To configure the trunk group link between device1 and the device, enter the following commands.

NOTE

The text shown in italics in the CLI example below shows messages echoed to the screen in answer to the CLI commands entered.

```
PowerConnect(config)#trunk e 5 to 8
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

To configure the trunk group link between device2 and the server, enter the following commands

```
PowerConnect(config)#trunk e 2 to 4
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

You then configure the trunk group on the device.

```
PowerConnect(config)#trunk ethernet 17 to 18
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

NOTE

The **trunk deploy** command dynamically places trunk configuration changes into effect, without a software reload.

Example 2: Configuring a trunk group that spans two Gbps Ethernet modules in a chassis device

This section shows how to configure a trunk group that spans two modules in a chassis device.

Multi-slot trunk groups are supported on 1-GbE ports, 10-GbE ports, as well as on static and LACP trunk ports.

To configure a trunk group consisting of two groups of ports, 1 – 2 on module 1 and 5 – 6 on module 4, enter the following commands.

```
PowerConnect(config)#trunk ethernet 1 to 2 ethernet 5 to 6
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

NOTE

The **trunk deploy** command dynamically places trunk configuration changes into effect, without a software reload.

Example 3: Configuring a multi-slot trunk group with one port per module

You can select one port per module in a multi-slot trunk group.

To configure a two-port multi-slot trunk group consisting of ports 1 on module 1 and 2 on module 2, enter the following commands.

```
PowerConnect(config)#trunk ethernet 1 to 1 ethernet 2 to 2
Trunk will be created in next trunk deploy
PowerConnect(config)#write memory
PowerConnect(config)#trunk deploy
```

Notice that the groups of ports meet the criteria for a multi-slot trunk group. Each group contains the same number of ports (two) and begins on a primary port (1 and 3).

Example 4: Configuring a trunk group of 10 Gbps Ethernet ports

You can configure 10 Gbps Ethernet ports together in a trunk group.

To configure a trunk group containing two 10 Gbps Ethernet ports, enter commands such as the following.

```
PowerConnect(config)#trunk ethernet 1 to 2
PowerConnect(config-trunk-1-2)# write memory
PowerConnect(config-trunk-1-2)# exit
PowerConnect(config)#trunk deploy
```

These commands configure a trunk group consisting of 10 Gbps Ethernet ports 1 and 2, then deploy the trunk group. The trunk configuration does not take effect until you deploy it.

Syntax: **[no] trunk ethernet** <primary-portnum> **to** <secondary-portnum>

Syntax: trunk deploy

The <primary-portnum> parameter specifies the trunk group primary port.

The <secondary-portnum> parameter specifies the secondary port in the trunk group.

NOTE

Two-port trunk groups are supported for 10 Gbps Ethernet. You cannot specify more than two ports.

To display configuration information and load-sharing statistics for the trunk group, enter the **show trunk** command. Refer to [“Displaying trunk group configuration information”](#) on page 323.

Additional trunking options

The following trunking options can be performed on ports in deployed trunks. Note that these options are *not* supported on LACP trunk ports on PowerConnect B-Series TI24X devices.

The additional trunking options are as follows:

- Naming a trunk port
- Disabling or re-enabling a trunk port
- Deleting a static trunk group (applies to static trunks only)
- Specifying the minimum number of ports in a trunk group (applies to static trunks only)
- Monitoring a trunk port
- Configuring outbound rate shaping on a trunk port
- Enabling sFlow forwarding on an individual port in a trunk
- Setting the sFlow sampling rate on an individual port in a trunk

NOTE

Depending on the operational state of LACP-enabled ports, at any time, these ports may join a trunk group, change trunk group membership, exit a trunk group, or possibly never join a trunk group. Therefore, before configuring trunking options on LACP-enabled ports (e.g., naming the port, disabling the port, etc.), verify the actual trunk group port membership using the **show trunk** command. To view the status of LACP, use the **show link-aggregate** command.

Naming a trunk port

To name an individual port in a trunk group, enter a command such as the following at the trunk group configuration level.

```
PowerConnect(config)#trunk e 1 to 4
PowerConnect(config-trunk-1-4)#port-name customer1 ethernet 2
```

This command assigns the name “customer1” to port 2 in the trunk group consisting of ports 1 – 4.

Syntax: [no] port-name <ASCII string> ethernet <portnum>

The <ASCII string> parameter specifies the port name. The name can be up to 49 characters long.

The <portnum> parameter is a valid port in the trunk group.

Disabling or re-enabling a trunk port

This feature is supported on individual ports of a static trunk group.

You can disable or re-enable individual ports in a trunk group. To disable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level.

```
PowerConnect(config)#trunk e 1 to 4
PowerConnect(config-trunk-1-4)#config-trunk-ind
PowerConnect(config-trunk-1-4)#disable ethernet 2
```

Syntax: [no] config-trunk-ind

Syntax: [no] disable ethernet <portnum>

The **config-trunk-ind** command enables configuration of individual ports in the trunk group. If you do not use this command, the **disable** and **enable** commands will be valid only for the primary port in the trunk group and will disable or enable all ports in the trunk group. You need to enter the **config-trunk-ind** command only once in a trunk group. After you enter the command, all applicable port configuration commands apply to individual ports only.

NOTE

If you enter **no config-trunk-ind**, all port configuration commands are removed from the individual ports and the configuration of the primary port is applied to all the ports. Also, once you enter the **no config-trunk-ind** command, the **enable**, **disable**, and **monitor** commands are valid only on the primary port and apply to the entire trunk group.

The **disable** command disables the port. The states of other ports in the trunk group are not affected.

If you have configured a name for the trunk port, you can specify the port name, as shown in the following example.

```
PowerConnect(config-trunk-1-4)#config-trunk-ind  
PowerConnect(config-trunk-1-4)#disable customer1
```

Syntax: **disable** <portname>

To enable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level.

```
PowerConnect(config-trunk-1-4)#config-trunk-ind  
PowerConnect(config-trunk-1-4)#enable ethernet 2
```

Syntax: **enable ethernet**<portnum>]

Syntax: **enable** <portname>

Disabling or re-enabling a range or list of trunk ports

To disable a range of ports in a trunk group, enter commands such as the following.

```
PowerConnect(config)#trunk ethernet 1 to 4  
PowerConnect(config-trunk-1-4)#config-trunk-ind  
PowerConnect(config-trunk-1-4)#disable ethernet 3 to 4
```

This command disables ports 3 – 4 in trunk group 1 – 4.

To disable a list of ports, enter a command such as the following.

```
PowerConnect(config-trunk-1-4)#disable ethernet 1 ethernet 3 ethernet 4
```

This command disables ports 1, 3, and 4 in the trunk group.

You can specify a range and a list on the same command line. For example, to re-enable some trunk ports, enter a command such as the following.

```
PowerConnect(config-trunk-1-4)#enable ethernet 1 to 2 ethernet 4
```

Syntax: [no] config-trunk-ind

Syntax: [no] **disable ethernet** <portnum> [to <portnum> | **ethernet** <portnum>]

Syntax: [no] **enable ethernet** <portnum> [to <portnum> | **ethernet** <portnum>]

The **to** <portnum> parameter indicates that you are specifying a range. Specify the lower port number in the range first, then **to**, then the higher port number in the range.

The <portnum> parameter specifies an individual port. You can enter this parameter multiple times to specify a list, as shown in the examples above.

Modifying Trunk Group Membership

You can change port membership by removing individual ports from the trunk group. To remove a port from a trunk group, use one of the following methods.

To remove ports 3 and 4 from the trunk group, enter the following command:

```
PowerConnect(config)# no trunk ethernet 3 to 4
```

Syntax: no trunk ethernet | pos <portnum> [to <portnum>]

The <portnum> parameter indicates the port you are removing.

NOTE

Make sure you enter the lower port in the range before the “to” and the higher port in the range after the “to”.

As a shortcut, you also can enter just the lower port in the range. The software automatically removes all higher ports in addition to the specified port. For example, to remove ports 3 and 4, you can enter the following command:

```
PowerConnect(config)# no trunk ethernet 3
```

Therefore, for trunk group 1 – 4, the following commands are not valid:

```
PowerConnect(config)# no trunk ethernet 2
```

These commands are invalid because the trunk group cannot contain only a single port. These commands, if the software allowed them, would result in a trunk group consisting only of port 1.

On most devices, trunk groups can contain two ports or four ports but cannot contain only three ports. Therefore, the following command also is invalid for trunk group 1 – 4:

```
PowerConnect(config)# no trunk ethernet 4
```

This command is invalid because it would result in a trunk group containing three ports, 1 – 3.

Deleting a static trunk group

Use the **no trunk ethernet** command to delete a static trunk group.

NOTE

To delete an LACP trunk group, use the CLI command **no link-aggregate active | passive**.

To delete a trunk group, use **no** in front of the command you used to create the trunk group. For example, to remove one of the trunk groups configured in the examples above, enter the following command.

```
PowerConnect(config)#no trunk ethernet 1 to 2 ethernet 3 to 4
```

Syntax: no trunk ethernet | pos <portnum> to <portnum>

Specifying the minimum number of ports in a static trunk group

You can configure devices to disable all of the ports in a trunk group when the number of active member ports drops below a specified threshold value. For example, if a trunk group has 4 ports, and the threshold for the trunk group is 3, then the trunk group is disabled if the number of available ports in the trunk group drops below 3. If the trunk group is disabled, then traffic is forwarded over a different link or trunk group.

11 Configuring a trunk group

For example, the following commands establish a trunk group consisting of 4 ports, then establish a threshold for this trunk group of three ports.

```
PowerConnect(config)#trunk e 31 to34
PowerConnect(config-trunk-31-34)#threshold 3
```

In this example, if the number of active ports drops below three, then all the ports in the trunk group are disabled.

Syntax: [no] threshold <number>

- <number> - Specify a threshold number from 2 (default) up to the number of ports in a trunk group. The total number of threshold ports must be greater than 1.

NOTE

When using the **no threshold** command, it is not necessary to enter a number.

Configuration notes:

- This feature is supported on static trunk groups only. It is not supported on LACP trunk groups.
- When UDLD is enabled on a trunk port, trunk threshold is not supported.
- The **disable module** command can be used to disable the ports on a module. However, on 10 Gbps modules, the **disable module** command does not cause the remote connection to be dropped. If a trunk group consists of 10 Gbps ports, and you use the **disable module** command to disable ports in the trunk group, which then causes the number of active ports in the trunk group to drop below the threshold value, the trunk group is not disabled.
- If you establish a threshold for a trunk used in conjunction with Metro Ring Protocol (MRP) on 10 Gbps interfaces, then you must also enable Link Fault Signaling (LFS).
- If you specify a threshold for a trunk group, the other end of the trunk group must also have the same threshold configuration.

Monitoring a trunk port

You can monitor the traffic on an individual port of a static trunk group. For configuration details, refer to [“Monitoring an individual trunk port”](#) on page 397.

Configuring outbound rate shaping for a trunk port

You can configure the maximum rate at which outbound traffic is sent out on a static trunk port. For configuration details, refer to [“Configuring outbound rate shaping for a trunk port”](#) on page 425.

Enabling sFlow forwarding on a trunk port

You can enable sFlow forwarding on individual ports of a static trunk group. For configuration details, refer to [“Enabling sFlow forwarding on individual interfaces”](#) on page 1087.

Setting the sFlow sampling rate on a trunk port

You can configure an individual trunk port to use a different sampling rate than the global default sampling rate. This feature is supported on static trunk ports. For configuration details, refer to [“Changing the sampling rate for a trunk port”](#) on page 1086.

Displaying trunk group configuration information

To display configuration information for the trunk groups, use the **show trunk** command. This command displays information for configured trunk groups and operational trunk groups. A configured trunk group is one that has been configured in the software but has not been placed into operation by a reset or reboot. An operational trunk group is one that has been placed into operation by a reset or reboot.

Enter the **show trunk** command at any CLI level.

Syntax: `show trunk [ethernet | pos <portnum> to <portnum>]`

NOTE

The **show trunk** command does not display any form of trunk when links are up.

```
PowerConnect#show trunk
```

```
Configured trunks:
```

```
Trunk ID: 1
HW Trunk ID: 1
Ports_Configured: 8
Primary Port Monitored: Jointly
Ports      1      2      3      4      5      6      7      8
Port Names none    none  none   none   none   longna test  none
Port_Status enable enable enable enable disable disable enable enable
Monitor    on     on     off    on     off    off   off   off
Mirror Port 3     4     N/A    5     N/A    N/A   N/A   N/A
Monitor Dir both  in    N/A    out   N/A    N/A   N/A   N/A
```

```
Operational trunks:
```

```
Trunk ID: 1
HW Trunk ID: 1
Duplex: Full
Speed: 1G
Tag: No
Priority: level0
Active Ports: 6
Ports      1      2      3      4      5      6      7      8
Link_Status active active active active down  down  active active
LACP_Status ready  ready  ready  expired down  down  ready  ready
Load Sharing
Mac Address 3      2      2      2      0      0      6      1
IP          0      0      0      0      0      0      0      0
Multicast  4      2      5      2      0      0      2      3
```

[Table 45](#) describes the information displayed by the **show trunk** command.

TABLE 45 CLI trunk group information

This field...	Displays...
Trunk ID	The trunk group number. The software numbers the groups in the display to make the display easy to use.
HW Trunk ID	The trunk ID.

11 Dynamic link aggregation

TABLE 45 CLI trunk group information (Continued)

This field...	Displays...
Duplex	The mode of the port, which can be one of the following: <ul style="list-style-type: none">• None – The link on the primary trunk port is down.• Full – The primary port is running in full-duplex.• Half – The primary port is running in half-duplex. NOTE: This field and the following fields apply only to operational trunk groups.
Speed	The speed set for the port. The value can be one of the following: <ul style="list-style-type: none">• None – The link on the primary trunk port is down.• 10 – The port speed is 10 Mbps.• 100 – The port speed is 100 Mbps.• 1G – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1Q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.
Active Ports	The number of ports in the trunk group that are currently active.
Ports	The ports in the trunk group.
Link_Status	The link status or each port in the trunk group.
LACP_Status	For more information about this feature, refer to the section “Displaying and determining the status of aggregate links” on page 335: <ul style="list-style-type: none">• Ready - The port is functioning normally in the trunk group and is able to transmit and receive LACP packets.• Expired - The time has expired (as determined by timeout values) and the port has shut down because the port on the other side of the link has stopped transmitting packets.• Down - The port physical link is down.
Load Sharing	The number of traffic flows currently being load balanced on the trunk ports. All traffic exchanged within the flow is forwarded on the same trunk port. For information about trunk load sharing, refer to “Trunk group load sharing” on page 314.

Dynamic link aggregation

The device software supports the IEEE 802.3ad standard for link aggregation. This standard describes the Link Aggregation Control Protocol (LACP), a mechanism for allowing ports on both sides of a redundant link to form a trunk link (aggregate link), without the need for manual configuration of the ports into trunk groups.

When you enable link aggregation on a group of ports, the ports can negotiate with the ports at the remote ends of the links to establish trunk groups.

The link aggregation feature automates trunk configuration but can coexist with the trunk group feature. Link aggregation parameters do not interfere with trunk group parameters.

NOTE

Use the link aggregation feature only if the device at the other end of the link you want to aggregate also supports IEEE 802.3ad link aggregation. Otherwise, you need to manually configure the trunk links.

Link aggregation support is disabled by default. You can enable the feature on an individual port basis, in active or passive mode:

- **Active mode** – When you enable a port for active link aggregation, the port can exchange standard LACP Protocol Data Unit (LACPDU) messages to negotiate trunk group configuration with the port on the other side of the link. In addition, the port actively sends LACPDU messages on the link to search for a link aggregation partner at the other end of the link, and can initiate an LACPDU exchange to negotiate link aggregation parameters with an appropriately configured remote port.
- **Passive mode** – When you enable a port for passive link aggregation, the port can exchange LACPDU messages with the port at the remote end of the link, but the port cannot search for a link aggregation port or initiate negotiation of an aggregate link. Thus, the port at the remote end of the link must initiate the LACPDU exchange.

NOTE

Dell recommends that you disable or remove the cables from the ports you plan to enable for dynamic link aggregation. Doing so prevents the possibility that LACP will use a partial configuration to talk to the other side of a link. A partial configuration does not cause errors, but does sometimes require LACP to be disabled and re-enabled on both sides of the link to ensure that a full configuration is used. It's easier to disable a port or remove its cable first. This applies both for active link aggregation and passive link aggregation.

Examples of valid LACP trunk groups

Ports follow the same configuration rules for dynamically created aggregate links as they do for statically configured trunk groups. Refer to [“Trunk group rules”](#) on page 312 and [“Trunk group load sharing”](#) on page 314.

[Figure 77](#) on page 325 shows some examples of valid aggregate links.

FIGURE 77 Examples of valid aggregate links

In this example, assume that link aggregation is enabled on all of the links between the device on the left and the device on the right (which can be either a device or another vendor device). The ports that are members of aggregate links in this example are following the configuration rules for trunk links on devices.

The rules apply to a device even if the device at the other end is from another vendor and uses different rules. Refer to [“Trunk group rules”](#) on page 312.

Configuration notes and limitations

This section lists the configuration considerations and limitations for dynamic link aggregation.

PowerConnect B-Series TI24X devices

The following notes and feature limitations apply to the PowerConnect B-Series TI24X devices:

- You cannot use 802.3ad link aggregation on a port configured as a member of a static trunk group.
- The dynamic link aggregation ((802.3ad) implementation on PowerConnect B-Series TI24X devices allows any number of ports up to eight to be aggregated into a link.

11 Dynamic link aggregation




- The default key assigned to an aggregate link is based on the port type (1 Gbps port or 10 Gbps port). The device assigns different keys to 10 Gbps ports than to 1 Gbps ports so that ports with different physical capabilities will not be able to form a trunk.

NOTE

The trunks that will be formed by link aggregation will strictly adhere to the static trunking rules on the device. Be careful in selecting keys if you are manually configuring link aggregation keys. Make sure that the possible trunks that you expect to be formed conform to the static trunking rules.

- When the feature dynamically adds or changes a trunk group, the **show trunk** command displays the trunk as both configured and active. However, the **show running-config** or **write terminal** command does not contain a trunk command defining the new or changed trunk group.

Ports enabled for link aggregation follow the same rules as ports configured for trunk groups.



- If the feature places a port into a trunk group as a secondary port, all configuration information except information related to link aggregation is removed from the port. For example, if port 3 has an IP interface, and the link aggregation feature places port 3 into a trunk group consisting of ports 1 – 4, the IP interface is removed from the port.

NOTE

Unique Key should be configured for each group of ports that have differing VLAN membership, regardless of being tagged or untagged.

- If you use this feature on a Layer 3 Switch that is running OSPF or BGP4, the feature causes these protocols to reset when a dynamic link change occurs. The reset includes ending and restarting neighbor sessions with OSPF and BGP4 peers, and clearing and relearning dynamic route entries and forwarding cache entries. Although the reset causes a brief interruption, the protocols automatically resume normal operation.
- You can enable link aggregation on 802.1Q tagged ports (ports that belong to more than one port-based VLAN).

NOTE

It is recommended to disable the link aggregation on ports before you disable link aggregation.

Adaptation to trunk disappearance

The device will tear down an aggregate link if the device at the other end of the link reboots or brings all the links down. Tearing the aggregate link down prevents a mismatch if the other device has a different trunk configuration following the reboot or re-establishment of the links. Once the other device recovers, 802.3 can renegotiate the link without a mismatch.

Flexible trunk eligibility

The criteria for trunk port eligibility in an aggregate link are flexible. A range of ports can contain down ports and still be eligible to become an aggregate link.

It also increases the tolerance for down ports during link negotiation. In previous releases, all the ports in a valid trunk configuration (2-port, 4-port, or 8-port trunk starting on a valid primary port number) need to be up.

The device places the ports into 2-port groups by default, consisting of an odd-numbered port and the next even-numbered port. For example, ports 1 and 2 are a two-port group, as are ports 3 and 4, and so on. If either of the ports in a two-port group is up, the device considers both ports to be eligible to be in an aggregate link.

[Figure 78](#) shows an example of 2-port groups in a range of four ports on which link aggregation is enabled. Based on the states of the ports, some or all of them will be eligible to be used in an aggregate link.

FIGURE 78 Two-port groups used to determine aggregation eligibility

11 Dynamic link aggregation




Table 46 shows examples of the ports from Figure 78 that will be eligible for an aggregate link based on individual port states.

TABLE 46 Port eligibility for link aggregation

	Port group 1		Port group 2		Trunk eligibility
	1	2	3	4	
Link State	Up	Up	Up	Up	4-port 1 – 4
	Up	Up	Up	Down	4-port 1 – 4
	Up	Down	Up	Down	4-port 1 – 4
	Up	Up	Down	Up	4-port 1 – 4
	Down	Down	Down	Up	2-port 3 – 4
	Up	Down	Down	Down	2-port 1 – 2

As shown in these examples, all or a subset of the ports within a port range will be eligible for formation into an aggregate link based on port states. Notice that the sets of ports that are eligible for the aggregate link must be valid static trunk configurations.

Enabling dynamic link aggregation

By default, link aggregation is disabled on all ports. To enable link aggregation on a set of ports, enter commands such as the following at the Interface configuration level of the CLI.

NOTE

Configuration commands for link aggregation differ depending on whether you are using the default link aggregation key automatically assigned by the software, or if you are assigning a different, unique key. Follow the commands below, according to the type of key you are using. For more information about keys, refer to “Key” on page 331.

Using the default key assigned by the software

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#link-aggregate active
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-e10000-2)#link-aggregate active
```

The commands in this example enable the active mode of link aggregation on ports 1 and 2. The ports can send and receive LACPDU messages. Note that these ports will use the default key, since one has not been explicitly configured.

NOTE

In conformance with the 802.3ad specification, the default key assigned to an aggregate link is based on the port type (1 Gbps port or 10 Gbps port). The device assigns different keys to 10 Gbps ports than 1 Gbps ports, so that ports with different physical capabilities will not be able to form a trunk.

Assigning a unique key

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#link-aggregate configure key 10000
PowerConnect(config-if-e10000-1)#link-aggregate active
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-e10000-2)#link-aggregate configure key 10000
PowerConnect(config-if-e10000-2)#link-aggregate active
```

The commands in this example assign the key 10000 and enable the active mode of link aggregation on ports 1 and 2. The ports can send and receive LACPDU messages.

NOTE

As shown in this example, when configuring a key, it is pertinent that you assign the key prior to enabling link aggregation.

The following commands enable passive link aggregation on ports 5 – 8.

```
PowerConnect(config)#interface ethernet 5 to 8
PowerConnect(config-mif-5-8)#link-aggregate passive
```

The commands in this example enable the passive mode of link aggregation on ports 5 – 8. These ports wait for the other end of the link to contact them. After this occurs, the ports can send and receive LACPDU messages.

To disable link aggregation on a port, enter a command such as the following.

```
PowerConnect(config-if-e10000-8)# link-aggregate
```

Syntax: [no] link-aggregate <active | passive >

Syntax: [no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>]

NOTE

For more information about keys, including details about the syntax shown above, refer to “Key” on page 331.

Dynamic Operation of Allocation Keys

The device dynamically changes a port key based on changes to the port VLAN membership.

11 Dynamic link aggregation

When you change a port VLAN membership, the device searches through existing key groups for a port with matching port properties. Specifically, it searches for a match on all three of the following properties:

- VLAN ID
- default key
- port tag type (tagged or untagged)

If it finds a match, the port (whose VLAN membership you are changing) gets the matching port key. If it does not find a match, the port gets a new key.

NOTE

For multi-slot trunk groups, you must manually configure the keys in the trunk group(s) to match. For instructions on configuring keys manually, see [“Configuring keys for ports with link aggregation enabled”](#) on page 334.

How changing the VLAN membership of a port affects trunk groups and dynamic keys

When you change a port VLAN membership and the port is currently a member of a trunk group, the following changes occur to the trunk group:

- The device tears down the existing trunk group.
- All ports in the trunk group get a new key.
- The new key group aggregates into a new trunk group.

When you change a port VLAN membership, and the port is not a member of a trunk group, the following changes occur:

- The port gets a new key depending on changes to the port VLAN tag type, as follows:
 - Tagged to Tagged VLAN – The primary port of the trunk group gets a new key.
 - Tagged to Untagged VLAN – The port gets the default key for untagged ports.
 - Untagged to Tagged VLAN – If the device finds a port with matching port properties, the port gets that port key. If it does not find one, the port gets a new key.
 - Untagged to Untagged VLAN – The port gets a new key depending on whether it is in the default VLAN or not. If there is a trunk group associated with the key, it is not affected.
- All other ports keep their existing key.
- The new key groups try to aggregate into trunk groups.

Link aggregation parameters

You can change the settings on individual ports for the following link aggregation parameters:

- System priority
- Port priority
- Timeout
- Key

System priority

The system priority parameter specifies the link aggregation priority on the device, relative to the devices at the other ends of the links on which link aggregation is enabled. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

System Priority does not take effect until you toggle the **link-aggregate <active | passive>** command.

NOTE

If you are connecting the device to another vendor device and the link aggregation feature is not working, set the system priority on the device to a lower priority (a higher priority value). In some cases, this change allows the link aggregation feature to operate successfully between the two devices.

Port priority

The port priority parameter determines the active and standby links. When a group of ports is negotiating with a group of ports on another device to establish a trunk group, the port with the highest priority becomes the default active port. The other ports (with lower priorities) become standby ports in the trunk group. You can specify a priority from 0 – 65535. A higher value indicates a lower priority. The default is 1.

NOTE

The primary port in the port group becomes the default active port. The primary port is the lowest-numbered port in a valid trunk-port group.

Link type

The link type parameter specifies whether the trunk is connecting to a server (server link) or to another networking device (switch link). The default link type is switch.

Timeout

You can specify a timeout mode, which determines how fast ports are removed from a trunk. You can specify a short timeout mode.

Key

Every port that is 802.3ad-enabled has a key. The key identifies the group of potential trunk ports to which the port belongs. Ports with the same key are called a key group and are eligible to be in the same trunk group.

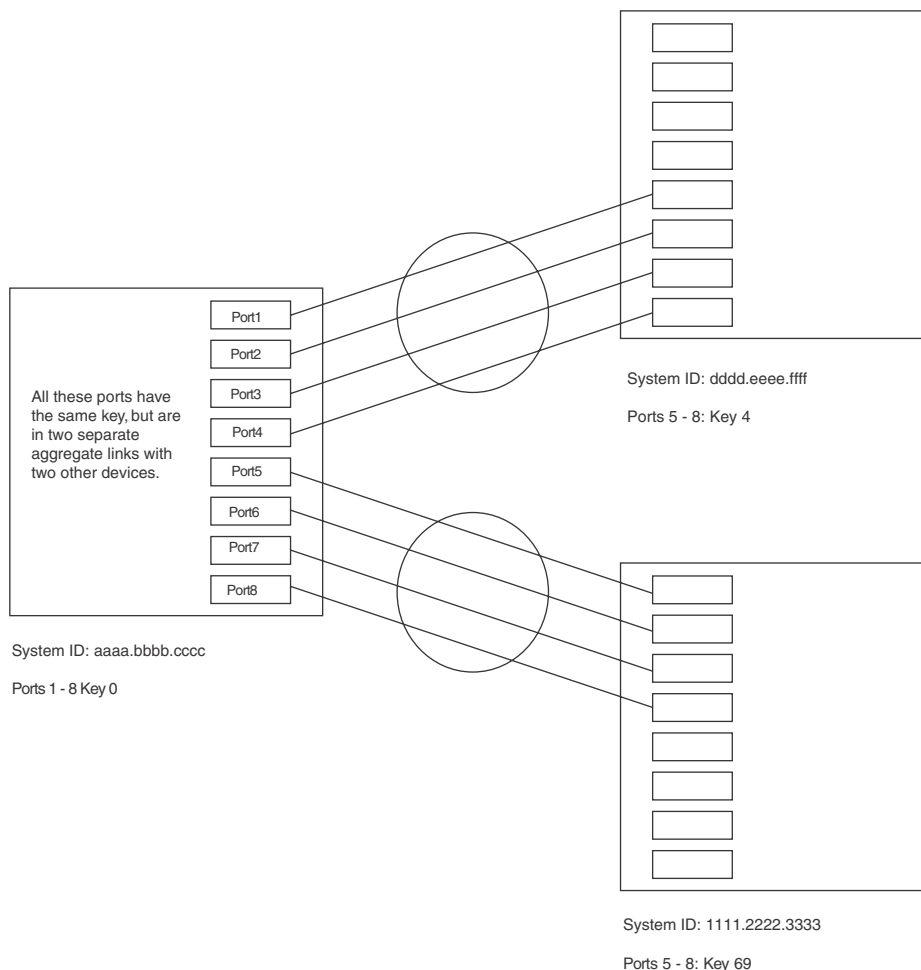
When you enable link-aggregation on an untagged port, the software assigns a default key to the port. For tagged ports, you must manually configure link-aggregation keys. Refer to [“Configuring keys for ports with link aggregation enabled”](#) on page 334.

All ports within an aggregate link must have the same key. However, if the device has ports that are connected to two different devices, and the port groups allow the ports to form into separate aggregate links with the two devices, then each group of ports can have the same key while belonging to separate aggregate links with different devices. [Figure 79](#) on page 332 shows an example.

NOTE

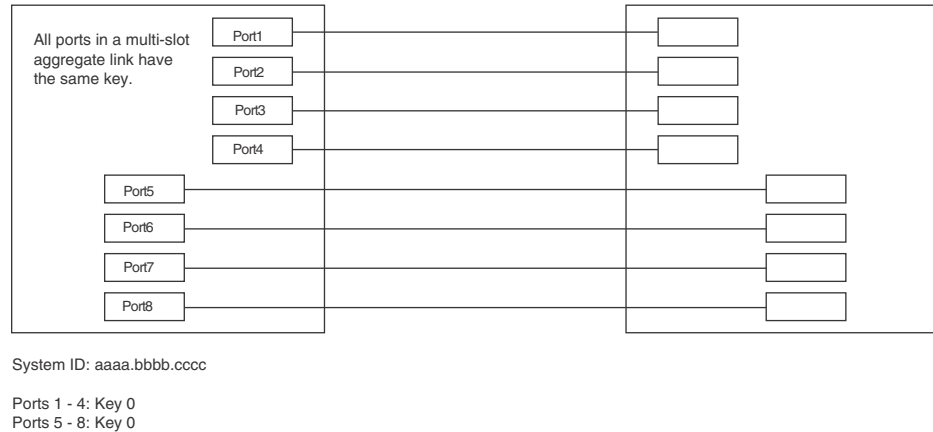
It is recommended to configure a unique key if ports are tagged or untagged in a VLAN.

FIGURE 79 Ports with the same key in different aggregate links



Notice that the keys between one device and another do not need to match. The only requirement for key matching is that all the ports within an aggregate link on a given device must have the same key.

Devices that support multi-slot trunk groups can form multi-slot aggregate links using link aggregation. However, the link aggregation keys for the groups of ports on each module must match. For example, if you want to allow link aggregation to form an aggregate link containing ports 1 – 4 and 5 – 8, you must change the link aggregation key on one or both groups of ports so that the key is the same on all eight ports. [Figure 80](#) on page 333 shows an example.

FIGURE 80 Multi-slot aggregate link

By default, the device ports are divided into 4-port groups. The software dynamically assigns a unique key to each 4-port group. If you need to divide a 4-port group into two 2-port groups, change the key in one of the groups so that the two 2-port groups have different keys. For example, if you plan to use ports 1 and 2 in VLAN 1, and ports 3 and 4 in VLAN 2, change the key for ports 3 and 4.

NOTE

If you change the key for a port group, recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

Viewing keys for tagged ports

To display link aggregation information, including the key for a specific port, enter a command such as the following at any level of the CLI.

```
PowerConnect# show link-aggregation ethernet 1
System ID: 00e0.52a9.bb00
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp]
1      0      0      0 No L No No No No No No
```

The command in this example shows the key and other link aggregation information for port 1.

To display link aggregation information, including the key for all ports on which link aggregation is enabled, enter the following command at any level of the CLI.

```
PowerConnect# show link-aggregation
System ID: 0004.8055.b200
Long timeout: 90, default: 90
Short timeout: 3, default: 3
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1      1      1  10000 Yes S Agg Syn Col Dis Def No Dwn
2      1      1  10000 Yes S Agg Syn Col Dis Def No Dwn
1      1      1  10000 Yes S Agg Syn Col Dis Def No Dwn
2      1      1  10000 Yes S Agg Syn Col Dis Def No Dwn
1      1      1    480 Yes S Agg Syn Col Dis Def No Dwn
2      1      1    480 Yes S Agg Syn Col Dis Def No Dwn
3      1      1    480 Yes S Agg Syn Col Dis Def No Dwn
4      1      1    480 Yes S Agg Syn Col Dis Def No Dwn
```

11 Dynamic link aggregation

17	1	1	481	Yes	S	Agg	Syn	Col	Dis	Def	No	Ope
18	1	1	481	Yes	S	Agg	Syn	Col	Dis	Def	No	Ope
19	1	1	481	Yes	S	Agg	Syn	Col	Dis	Def	No	Ope
20	1	1	481	Yes	S	Agg	Syn	Col	Dis	Def	No	Ope

Syntax: `show link-aggregation [ethernet [<portnum>]`

Possible values: N/A

Default value: N/A

Configuring link aggregation parameters

You can configure one or more parameters on the same command line, and in any order.

NOTE

For key configuration only, configuration commands differ depending on whether or not link aggregation is enabled on the ports. Follow the appropriate set of commands below, according to your system configuration.

Configuring a port group key if link aggregation is disabled

Use this command sequence to change the key for ports that do not have link aggregation enabled, and for all other link aggregation parameters (i.e., system priority, port priority).

For example, to change the software-assigned key for a port group to another value, enter commands similar to the following.

```
PowerConnect(config)#interface ethernet 1 to 4
PowerConnect(config-mif-1-4)#link-aggregate configure key 10000
PowerConnect(config-mif-1-4)#interface ethernet 5 to 8
PowerConnect(config-mif-5-8)#link-aggregate configure key 10000
```

Configuring keys for ports with link aggregation enabled

As shown in this command sequence, to change the key on ports that already have link aggregation enabled, you must first turn OFF link aggregation, configure the new key, then re-enable link aggregation.

```
PowerConnect(config)# interface ethernet 1 to 4
PowerConnect(config-mif-1-4)# link-aggregate
PowerConnect(config-mif-1-4)# link-aggregate configure key 10000
PowerConnect(config-mif-1-4)# link-aggregate active
PowerConnect(config-mif-1-4)# interface ethernet 5 to 8
PowerConnect(config-mif-5-8)# link-aggregate
PowerConnect(config-mif-5-8)# link-aggregate configure key 10000
PowerConnect(config-mif-5-8)# link-aggregate active
```

These commands change the key for ports 1 – 4 and 5 – 8 to 10000. Since all ports in an aggregate link must have the same key, the command in this example enables ports 1 – 4 and 5 – 8 to form a multi-slot aggregate link.

Syntax: `[no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>]`

The **system-priority** <num> parameter specifies the device link aggregation priority. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

The **port-priority** <num> parameter specifies an individual port priority within the port group. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

The **key** <num> parameter identifies the group of ports that are eligible to be aggregated into a trunk group. The software automatically assigns a key to each group of ports. The software assigns the keys in ascending numerical order, beginning with 0. You can change a port group key to a value from 10000 – 65535.

NOTE

If you change the key for a port group, Dell recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

Configuring port timeout

You can control the time it takes to remove ports from a trunk with link aggregation enabled by configuring the link aggregated port with a “short” timeout mode. Once a port is configured with a timeout mode, it will remain in that timeout mode whether it is up or down or whether or not it is part of a trunk.

All ports in a trunk should have the same timeout mode, which is checked when link aggregation is enabled on ports.

To configure a port with a short timeout mode, enter a command such as the following.

```
PowerConnect(config)# interface ethernet1
PowerConnect(config-if-e10000-1)# link-aggregation configure timeout short
```

Syntax: [no] link-aggregate configure timeout [short]

If the timeout mode is not configured for a port and link aggregation is enabled, the port starts with a short timeout mode. Once a trunk is formed, the timeout mode is changed to the long timeout mode. The value for “long” and “short” is displayed in the output for the **show link-aggregate** command.

Displaying and determining the status of aggregate links

The **show link-aggregation** command provides the ability to view the status of dynamic links. You can determine the status of ports that are members of an aggregate link, and whether LACPDU messages are being transmitted between the ports.

The following section provides details about the events that can affect the status of ports in an aggregate link and the status of LACP messages exchanged between the ports. Later sections provide instructions for viewing these status reports.

Events that affect the status of ports in an aggregate link

Devices can block traffic on a port or shut down a port that is part of a trunk group or aggregate link, when a port joins a trunk group and the port on the other end of the link shuts down or stops transmitting LACP packets. Depending on the timeout value set on the port, the link aggregation information expires. If this occurs, the device shuts down the port and notifies all the upper layer protocols that the port is down.

Devices can also block traffic on a port that is initially configured with link aggregation. The port is blocked until it joins a trunk group. In this case, traffic is blocked, but the port is still operational.

A port remains blocked until one of the following events occurs:

- Both ports in the aggregate link have the same key

11 Displaying and determining the status of aggregate links

- LACP brings the port back up
- The port joins a trunk group

Displaying link aggregation and port status information

Use the **show link-aggregation** command to determine the operational status of ports associated with aggregate links.

To display the link aggregation information for a specific port, enter a command such as the following at any level of the CLI.

```
PowerConnect# show link-aggregation ethernet 5
System ID: 00e0.52a9.bb00
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
5      1      1      20000 No L Agg Syn Col Dis No No Ope
```

The command in this example shows the link aggregation information for port 1.

To display the link aggregation information for all ports on which link aggregation is enabled, enter the following command at any level of the CLI.

```
PowerConnect#show link-agg
System ID: 0024.3817.50bb
Long timeout: 120, default: 120
Short timeout: 3, default: 3
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
5      1      1      20000 No L Agg Syn Col Dis No No Ope
6      1      1      20000 No L Agg Syn Col Dis No No Ope
25     1      1      10000 Yes L Agg Syn Col Dis No No Ope
26     1      1      10000 Yes L Agg Syn Col Dis No No Ope
```

Syntax: **show link-aggregation [ethernet <portnum>]**

Use **ethernet <portnum>** to display link-aggregation information for a specific port.

NOTE

Ports that are configured as part of an aggregate link must also have the same key. For more information about assigning keys, refer to the section [“Link aggregation parameters”](#) on page 330.

The **show link aggregation** command output is described in [“CLI display of link aggregation information”](#) on page 336.

TABLE 47 CLI display of link aggregation information

This field...	Displays...
System ID	Lists the base MAC address of the device. This is also the MAC address of port 1 .
Port	Lists the port number.
Sys P	Lists the system priority configured for this port.
Port P	Lists the port link aggregation priority.
Key	Lists the link aggregation key.

TABLE 47 CLI display of link aggregation information (Continued)

This field...	Displays...
Act	<p>Indicates the link aggregation mode, which can be one of the following:</p> <ul style="list-style-type: none"> • No – The mode is passive or link aggregation is disabled (off) on the port. <p>If link aggregation is enabled (and the mode is passive), the port can send and receive LACPDU messages to participate in negotiation of an aggregate link initiated by another port, but cannot search for a link aggregation port or initiate negotiation of an aggregate link.</p> <ul style="list-style-type: none"> • Yes – The mode is active. The port can send and receive LACPDU messages.
Tio	<p>Indicates the timeout value of the port. The timeout value can be one of the following:</p> <ul style="list-style-type: none"> • L – Long. The trunk group has already been formed and the port is therefore using a longer message timeout for the LACPDU messages exchanged with the remote port. Typically, these messages are used as confirmation of the health of the aggregate link. • S – Short. The port has just started the LACPDU message exchange process with the port at the other end of the link. The S timeout value also can mean that the link aggregation information received from the remote port has expired and the ports are starting a new information exchange.
Agg	<p>Indicates the link aggregation state of the port. The state can be one of the following:</p> <ul style="list-style-type: none"> • Agg – Link aggregation is enabled on the port. • No – Link aggregation is disabled on the port.
Syn	<p>Indicates the synchronization state of the port. The state can be one of the following:</p> <ul style="list-style-type: none"> • No – The port is out of sync with the remote port. The port does not understand the status of the LACPDU process and is not prepared to enter a trunk link. • Syn – The port is in sync with the remote port. The port understands the status of the LACPDU message exchange process, and therefore knows the trunk group to which it belongs, the link aggregation state of the remote port, and so on.
Col	<p>Indicates the collection state of the port, which determines whether the port is ready to send traffic over the trunk link.</p> <ul style="list-style-type: none"> • Col – The port is ready to send traffic over the trunk link. • No – The port is not ready to send traffic over the trunk link.
Dis	<p>Indicates the distribution state of the port, which determines whether the port is ready to receive traffic over the trunk link.</p> <ul style="list-style-type: none"> • Dis – The port is ready to receive traffic over the trunk link. • No – The port is not ready to receive traffic over the trunk link.
Def	<p>Indicates whether the port is using default link aggregation values. The port uses default values if it has not received link aggregation information through LACP from the port at the remote end of the link. This field can have one of the following values:</p> <ul style="list-style-type: none"> • Def – The port has not received link aggregation values from the port at the other end of the link and is therefore using its default link aggregation LACP settings. • No – The port has received link aggregation information from the port at the other end of the link and is using the settings negotiated with that port.

11 Clearing the negotiated aggregate links table

TABLE 47 CLI display of link aggregation information (Continued)

This field...	Displays...
Exp	Indicates whether the negotiated link aggregation settings have expired. The settings expire if the port does not receive an LACPDU message from the port at the other end of the link before the message timer expires. This field can have one of the following values: <ul style="list-style-type: none">• Exp – The link aggregation settings this port negotiated with the port at the other end of the link have expired. The port is now using its default link aggregation settings.• No – The link aggregation values that this port negotiated with the port at the other end of the link have not expired, so the port is still using the negotiated settings.
Ope	<ul style="list-style-type: none">• Ope (operational) - The port is operating normally.• Ina (inactive) - The port is inactive because the port on the other side of the link is down or has stopped transmitting LACP packets.• Blo (blocked) - The port is blocked because the adjacent port is not configured with link aggregation or because it is not able to join a trunk group. To unblock the port and bring it to an operational state, enable link aggregation on the adjacent port and ensure that the ports have the same key.• Dwn.

Displaying LACP status information

Use the **show trunk** command to determine the status of LACP. Refer to [“Displaying trunk group configuration information”](#) on page 323.

Clearing the negotiated aggregate links table

When a group of ports negotiates a trunk group configuration, the software stores the negotiated configuration in a table. You can clear the negotiated link aggregation configurations from the software. When you clear the information, the software does not remove link aggregation parameter settings you have configured. Only the configuration information negotiated using LACP is removed.

NOTE

The software automatically updates the link aggregation configuration based on LACPDU messages. However, clearing the link aggregation information can be useful if you are troubleshooting a configuration.

To clear the link aggregation information, enter the following command at the Privileged EXEC level of the CLI.

```
PowerConnect#clear link-aggregate
```

Syntax: clear link-aggregate

Configuring single link LACP

A single instance of link aggregation (or single link LACP) can be used for unidirectional link detection. Single link LACP is based on the 802.3ad LACP protocol; but allows you to form an aggregated link with only one Ethernet port. It is the preferred method for detecting unidirectional links across multi-vendor devices, instead of link-keepalive (UDLD), since it is based on a standard rather than on a proprietary solution.

Configuration notes

- This feature is supported on 1-GbE and 10-GbE ports.
- This feature is not supported on static trunk ports.
- This feature is not intended for the creation of trunk groups.
- The single link LACP timer is always short (3 seconds) and is not configurable. PDUs are sent out every three seconds.
- This feature is not supported on ports that have the **link-keepalive** command (UDLD) configured.

CLI syntax

To form a single link LACP, the port on both sides of the link must have LACP enabled. You can then define a single link LACP at the interface level of the device by entering the following command.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#link-aggregate configure singleton
Link-aggregation active
```

Syntax: [no] link-aggregate configure singleton

When single link LACP is configured, the **show link-aggregate** command displays the following information.

```
PowerConnect#show link-agg
System ID: 0024.3817.50bb
Long timeout: 120, default: 120
Short timeout: 3, default: 3
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
5      1      1      20000 No L Agg Syn Col Dis No No Ope
6      1      1      20000 No L Agg Syn Col Dis No No Ope
25     1      1      10000 Yes L Agg Syn Col Dis No No Ope
26     1      1      10000 Yes L Agg Syn Col Dis No No Ope
```

If **singleton** is configured on the port, the “Key” column displays “singleton”. Refer to [“CLI display of link aggregation information”](#) on page 336 to interpret the information on the displayed output.

Also, when ports are logically brought up or down while **singleton** is configured on the port, the following Syslog messages are generated:

- Logical link on interface ethernet <port#> is up.
- Logical link on interface ethernet <port#> is down.

11 Configuring single link LACP

Configuring GARP VLAN Registration Protocol

GVRP overview

GARP VLAN Registration Protocol (GVRP) is a Generic Attribute Registration Protocol (GARP) application that provides VLAN registration service by means of dynamic configuration (registration) and distribution of VLAN membership information.

A device enabled for GVRP can do the following:

- Learn about VLANs from other devices and configure those VLANs on the ports that learn about the VLANs. The device listens for GVRP Protocol Data Units (PDUs) from other devices, and implements the VLAN configuration information in the PDUs.
- Advertise VLANs configured on the device to other devices. The device sends GVRP PDUs advertising its VLANs to other devices. GVRP advertises statically configured VLANs and VLANs learned from other devices through GVRP.

GVRP enables a device to dynamically create 802.1Q-compliant VLANs on links with other devices that are running GVRP. GVRP reduces the chances for errors in VLAN configuration by automatically providing VLAN ID consistency across the network. You can use GVRP to propagate VLANs to other GVRP-aware devices automatically, without the need to manually configure the VLANs on each device. In addition, if the VLAN configuration on a device changes, GVRP automatically changes the VLAN configurations of the affected devices.

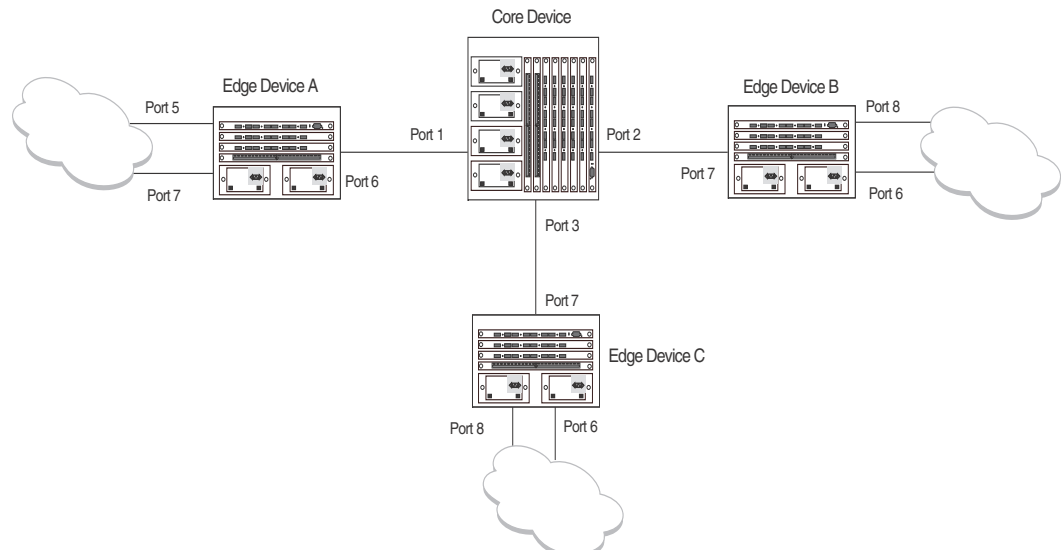
The Dell implementation of GARP and GVRP is based on the following standards:

- ANSI/IEEE standard 802.1D, 1998 edition
- IEEE standard 802.1Q, 1998 edition; approved December 8, 1998
- IEEE draft P802.1w/D10, March 26, 2001
- IEEE draft P802.1u/D9, November 23, 2000
- IEEE draft P802.1t/D10, November 20, 2000

Application examples

[Figure 81](#) shows an example of a network that uses GVRP. This section describes various ways you can use GVRP in a network such as this one. “[CLI examples](#)” on page 357 lists the CLI commands to implement the applications of GVRP described in this section.

FIGURE 81 Example of GVRP



In this example, a core device is attached to three edge devices. Each of the edge devices is attached to other edge devices or host stations (represented by the clouds).

The effects of GVRP in this network depend on which devices the feature is enabled on, and whether both learning and advertising are enabled. In this type of network (a core device and edge devices), you can have the following four combinations:

- Dynamic core and fixed edge
- Dynamic core and dynamic edge
- Fixed core and dynamic edge
- Fixed core and fixed edge

Dynamic core and fixed edge

In this configuration, all ports on the core device are enabled to learn and advertise VLAN information. The edge devices are configured to advertise their VLAN configurations on the ports connected to the core device. GVRP learning is disabled on the edge devices.

TABLE 48

Core device	Edge device A	Edge device B	Edge device C
<ul style="list-style-type: none"> • GVRP is enabled on all ports. • Both learning and advertising are enabled. <p>NOTE: Since learning is disabled on all the edge devices, advertising on the core device has no effect in this configuration.</p>	<ul style="list-style-type: none"> • GVRP is enabled on port 6. Learning is disabled. • VLAN 20 • Port 5 (untagged) • Port 6 (tagged) • VLAN 40 • Port 7 (untagged) • Port 6 (tagged) 	<ul style="list-style-type: none"> • GVRP is enabled on port 7. Learning is disabled. • VLAN 20 • Port 8 (untagged) • Port 7 (tagged) • VLAN 30 • Port 6 (untagged) • Port 7 (tagged) 	<ul style="list-style-type: none"> • GVRP is enabled on port 7. Learning is disabled. • VLAN 30 • Port 8 (untagged) • Port 7 (tagged) • VLAN 40 • Port 6 (untagged) • Port 7 (tagged)

In this configuration, the edge devices are statically (manually) configured with VLAN information. The core device dynamically configures itself to be a member of each of the edge device VLANs. The operation of GVRP on the core device results in the following VLAN configuration on the device:

- VLAN 20
 - 1 (tagged)
 - 2 (tagged)
- VLAN 30
 - 2 (tagged)
 - 3 (tagged)
- VLAN 40
 - 1 (tagged)
 - 3 (tagged)

VLAN 20 traffic can now travel through the core between edge devices A and B. Likewise, VLAN 30 traffic can travel between B and C and VLAN 40 traffic can travel between A and C. If an edge device is moved to a different core port or the VLAN configuration of an edge device is changed, the core device automatically reconfigures itself to accommodate the change.

Notice that each of the ports in the dynamically created VLANs is tagged. All GVRP VLAN ports configured by GVRP are tagged, to ensure that the port can be configured for additional VLANs.

NOTE

This example assumes that the core device has no static VLANs configured. However, you can have static VLANs on a device that is running GVRP. GVRP can dynamically add other ports to the statically configured VLANs but cannot delete statically configured ports from the VLANs.

Dynamic core and dynamic edge

GVRP is enabled on the core device and on the edge devices. This type of configuration is useful if the devices in the edge clouds are running GVRP and advertise their VLANs to the edge devices. The edge devices learn the VLANs and also advertise them to the core. In this configuration, you do not need to statically configure the VLANs on the edge or core devices, although you can have statically configured VLANs on the devices. The devices learn the VLANs from the devices in the edge clouds.

Fixed core and dynamic edge

GVRP learning is enabled on the edge devices. The VLANs on the core device are statically configured, and the core device is enabled to advertise its VLANs but not to learn VLANs. The edge devices learn the VLANs from the core.

Fixed core and fixed edge

The VLANs are statically configured on the core and edge devices. On each edge device, VLAN advertising is enabled but learning is disabled. GVRP is not enabled on the core device. This configuration enables the devices in the edge clouds to learn the VLANs configured on the edge devices.

VLAN names

The **show vlans** command lists VLANs created by GVRP as “GVRP_VLAN_<vlan-id>”. VLAN names for statically configured VLANs are not affected. To distinguish between statically-configured VLANs that you add to the device and VLANs that you convert from GVRP-configured VLANs into statically-configured VLANs, the **show vlans** command displays a converted VLAN name as “STATIC_VLAN_<vlan-id>”.

Configuration notes

- If you disable GVRP, all GVRP configuration information is lost if you save the configuration change (**write memory** command) and then reload the software. However, if you reload the software without first saving the configuration change, the GVRP configuration is restored following a software reload.
- The maximum number of VLANs supported on a device enabled for GVRP is the same as the maximum number on a device that is not enabled for GVRP.
 - To display the maximum number of VLANs allowed on your device, enter the **show default values** command. See the “vlan” row in the System Parameters section. Make sure you allow for the default VLAN (1), the GVRP base VLAN (4093), and the Single STP VLAN (4094). These VLANs are maintained as “Registration Forbidden” in the GVRP database. Registration Forbidden VLANs cannot be advertised or learned by GVRP.
 - To increase the maximum number of VLANs supported on the device, enter the **system-max vlan <num>** command at the global CONFIG level of the CLI, then save the configuration and reload the software. The maximum number you can specify is listed in the Maximum column of the **show default values** display.
- The default VLAN (VLAN 1) is not advertised by the Dell implementation of GVRP. The default VLAN contains all ports that are not members of statically configured VLANs or VLANs enabled for GVRP.

NOTE

The default VLAN has ID 1 by default. You can change the VLAN ID of the default VLAN, but only before GVRP is enabled. You cannot change the ID of the default VLAN after GVRP is enabled.

- Single STP must be enabled on the device. The Dell implementation of GVRP requires Single STP. If you do not have any statically configured VLANs on the device, you can enable Single STP as follows.

```
PowerConnect(config)# vlan 1
PowerConnect(config-vlan-1)# exit
PowerConnect(config)# span
PowerConnect(config)# span single
```

These commands enable configuration of the default VLAN (VLAN 1), which contains all the device ports, and enable STP and Single STP.

- All VLANs that are learned dynamically through GVRP are added to the single spanning tree.
- All ports that are enabled for GVRP become tagged members of the GVRP base VLAN (4093). If you need to use this VLAN ID for another VLAN, you can change the GVRP VLAN ID. Refer to [“Changing the GVRP base VLAN ID”](#) on page 345. The software adds the GVRP base VLAN to the single spanning tree.
- All VLAN ports added by GVRP are tagged.

- GVRP is supported only for tagged ports or for untagged ports that are members of the default VLAN. GVRP is not supported for ports that are untagged and are members of a VLAN other than the default VLAN.
- To configure GVRP on a trunk group, enable the protocol on the primary port in the trunk group. The GVRP configuration of the primary port is automatically applied to the other ports in the trunk group.
- You can use GVRP on a device even if the device has statically configured VLANs. GVRP does not remove any ports from the statically configured VLANs, although GVRP can add ports to the VLANs. GVRP advertises the statically configured VLANs. Ports added by GVRP do not appear in the running-config and will not appear in the startup-config file when save the configuration. You can manually add a port to make the port a permanent member of the VLAN. After you manually add the port, the port will appear in the running-config and be saved to the startup-config file when you save the configuration.
- VLANs created by GVRP do not support virtual routing interfaces or protocol-based VLANs. virtual routing interfaces and protocol-based VLANs are still supported on statically configured VLANs even if GVRP adds ports to those VLANs.
- You cannot manually configure any parameters on a VLAN that is created by GVRP. For example, you cannot change STP parameters for the VLAN.
- The GVRP timers (Join, Leave, and Leaveall) must be set to the same values on all the devices that are exchanging information using GVRP.
- If the network has a large number of VLANs, the GVRP traffic can use a lot of CPU resources. If you notice high CPU utilization after enabling GVRP, set the GVRP timers to longer values. In particular, set the Leaveall timer to a longer value. Refer to [“Changing the GVRP timers”](#) on page 347.
- The feature is supported only on Ethernet ports.

NOTE

If you plan to change the GVRP base VLAN ID (4093) or the maximum configurable value for the Leaveall timer (300000 ms by default), you must do so before you enable GVRP.

Configuring GVRP

To configure a device for GVRP, globally enable support for the feature, then enable the feature on specific ports. Optionally, you can disable VLAN learning or advertising on specific interfaces.

You can also change the protocol timers and the GVRP base VLAN ID.

Changing the GVRP base VLAN ID

By default, GVRP uses VLAN 4093 as a base VLAN for the protocol. All ports that are enabled for GVRP become tagged members of this VLAN. If you need to use VLAN ID 4093 for a statically configured VLAN, you can change the GVRP base VLAN ID.

NOTE

If you want to change the GVRP base VLAN ID, you must do so before enabling GVRP.

To change the GVRP base VLAN ID, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# gvrp-base-vlan-id 1001
```

This command changes the GVRP VLAN ID from 4093 to 1001.

Syntax: [no] **gvrp-base-vlan-id** <vlan-id>

The <vlan-id> parameter specifies the new VLAN ID. You can specify a VLAN ID from 2 – 4092 or 4095.

Increasing the maximum configurable value of the Leaveall timer

By default, the highest value you can specify for the Leaveall timer is 300000 ms. You can increase the maximum configurable value of the Leaveall timer to 1000000 ms.

NOTE

You must enter this command before enabling GVRP. Once GVRP is enabled, you cannot change the maximum Leaveall timer value.

NOTE

This command does not change the default value of the Leaveall timer itself. The command only changes the maximum value to which you can set the Leaveall timer.

To increase the maximum value you can specify for the Leaveall timer, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# gvrp-max-leaveall-timer 1000000
```

Syntax: [no] **gvrp-max-leaveall-timer** <ms>

The <ms> parameter specifies the maximum number of ms to which you can set the Leaveall timer. You can specify from 300000 – 1000000 (one million) ms. The value must be a multiple of 100 ms. The default is 300000 ms.

Enabling GVRP

To enable GVRP, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# gvrp-enable  
PowerConnect(config-gvrp)# enable all
```

The first command globally enables support for the feature and changes the CLI to the GVRP configuration level. The second command enables GVRP on all ports on the device.

The following command enables GVRP on ports 1,2 and 4.

```
PowerConnect(config-gvrp)# enable ethernet 1 ethernet 2 ethernet 4
```

Syntax: [no] **gvrp-enable**

Syntax: [no] **enable all | ethernet <portnum> [ethernet <portnum> | to <portnum>]**

The **all** parameter enables GVRP on all ports.

The **ethernet <portnum> [ethernet <portnum> | to <portnum>]** parameter enables GVRP on the specified list or range of Ethernet ports.

- To specify a list, enter each port as **ethernet <portnum>** followed by a space.

- To specify a range, enter the first port in the range as **ethernet <portnum>** followed by to followed by the last port in the range. For example, to add ports 1 – 8, enter the following command: **enable ethernet 1 to 8**.

You can combine lists and ranges in the same command. For example: **enable ethernet 1 to 8 ethernet 1 ethernet 2 ethernet 3**.

Disabling VLAN advertising

To disable VLAN advertising on a port enabled for GVRP, enter a command such as the following at the GVRP configuration level.

```
PowerConnect(config-gvrp)# block-applicant ethernet 1 ethernet 2 ethernet 3
```

This command disables advertising of VLAN information on ports 1, 2, and 3.

Syntax: [no] block-applicant all | ethernet <portnum> [ethernet <portnum> | to <portnum>]

NOTE

Leaveall messages are still sent on the GVRP ports.

Disabling VLAN learning

To disable VLAN learning on a port enabled for GVRP, enter a command such as the following at the GVRP configuration level.

```
PowerConnect(config-gvrp)# block-learning ethernet 2
```

This command disables learning of VLAN information on port 2.

NOTE

The port still advertises VLAN information unless you also disable VLAN advertising.

Syntax: [no] block-learning all | ethernet <portnum> [ethernet <portnum> | to <portnum>]

Changing the GVRP timers

GVRP uses the following timers:

- Join – The maximum number of milliseconds (ms) a device GVRP interfaces wait before sending VLAN advertisements on the interfaces. The actual interval between Join messages is randomly calculated to a value between 0 and the maximum number of milliseconds specified for Join messages. You can set the Join timer to a value from 200 – one third the value of the Leave timer. The default is 200 ms.
- Leave – The number of ms a GVRP interface waits after receiving a Leave message on the port to remove the port from the VLAN indicated in the Leave message. If the port receives a Join message before the Leave timer expires, GVRP keeps the port in the VLAN. Otherwise, the port is removed from the VLAN. When a port receives a Leave message, the port GVRP state is changed to Leaving. Once the Leave timer expires, the port GVRP state changes to Empty. You can set the Leave timer to a value from three times the Join timer – one fifth the value of the Leaveall timer. The default is 600 ms.

NOTE

When all ports in a dynamically created VLAN (one learned through GVRP) leave the VLAN, the VLAN is immediately deleted from the device's VLAN database. However, this empty VLAN is still maintained in the GVRP database for an amount of time equal to the following.

(number-of-GVRP-enabled-up-ports) * (2 * join-timer)

While the empty VLAN is in the GVRP database, the VLAN does not appear in the **show vlans** display but does still appear in the **show gvrp vlan all** display.

- **Leaveall** – The minimum interval at which GVRP sends Leaveall messages on all GVRP interfaces. Leaveall messages ensure that the GVRP VLAN membership information is current by aging out stale VLAN information and adding information for new VLAN memberships, if the information is missing. A Leaveall message instructs the port to change the GVRP state for all its VLANs to Leaving, and remove them unless a Join message is received before the Leave timer expires. By default, you can set the Leaveall timer to a value from five times the Leave timer – maximum value allowed by software (configurable from 300000 – 1000000 ms). The default is 10000.

NOTE

The actual interval is a random value between the Leaveall interval and 1.5 * the Leaveall time or the maximum Leaveall time, whichever is lower.

NOTE

You can increase the maximum configurable value of the Leaveall timer from 300000 ms up to 1000000 ms using the **gvrp-max-leaveall-timer** command. (Refer to [“Increasing the maximum configurable value of the Leaveall timer”](#) on page 346.)

Timer configuration requirements

- All timer values must be in multiples of 100 ms.
- The Leave timer must be $\geq 3 \times$ the Join timer.
- The Leaveall timer must be $\geq 5 \times$ the Leave timer.
- The GVRP timers must be set to the same values on all the devices that are exchanging information using GVRP.

Changing the Join, Leave, and Leaveall timers

The same CLI command controls changes to the Join, Leave, and Leaveall timers. To change values to the timers, enter a command such as the following.

```
PowerConnect(config-gvrp)# join-timer 1000 leave-timer 3000 leaveall-timer 15000
```

This command changes the Join timer to 1000 ms, the Leave timer to 3000 ms, and the Leaveall timer to 15000.

Syntax: [no] join-timer <ms> leave-timer <ms> leaveall-timer <ms>

NOTE

When you enter this command, all the running GVRP timers are canceled and restarted using the new times specified by the command.

Resetting the timers to their defaults

To reset the Join, Leave, and Leaveall timers to their default values, enter the following command.

```
PowerConnect(config-gvrp)# default-timers
```

Syntax: default-timers

This command resets the timers to the following values:

- Join – 200 ms
- Leave – 600 ms
- Leaveall – 10000 ms

Converting a VLAN created by GVRP into a statically-configured VLAN

You cannot configure VLAN parameters on VLANs created by GVRP. Moreover, VLANs and VLAN ports added by GVRP do not appear in the running-config and cannot be saved in the startup-config file.

To be able to configure and save VLANs or ports added by GVRP, you must convert the VLAN ports to statically-configured ports.

To convert a VLAN added by GVRP into a statically-configured VLAN, add the ports using commands such as the following.

```
PowerConnect(config)# vlan 22
PowerConnect(config-vlan-22)# tagged ethernet 1 to 8
```

These commands convert GVRP-created VLAN 22 containing ports 1 through 8 into statically-configured VLAN 22.

Syntax: [no] vlan <vlan-id>

Syntax: [no] tagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

Use the same commands to statically add ports that GVRP added to a VLAN.

NOTE

You cannot add the VLAN ports as untagged ports.

NOTE

After you convert the VLAN, the VLAN name changes from “GVRP_VLAN_<vlan-id>” to “STATIC_VLAN_<vlan-id>”.

Displaying GVRP information

You can display the following GVRP information:

- GVRP configuration information
- GVRP VLAN information
- GVRP statistics

- CPU utilization statistics
- GVRP diagnostic information

Displaying GVRP configuration information

To display GVRP configuration information, enter a command such as the following.

```
PowerConnect# show gvrp
GVRP is enabled on the system

GVRP BASE VLAN ID      : 4093
GVRP MAX Leaveall Timer : 300000 ms

GVRP Join Timer        : 200 ms
GVRP Leave Timer       : 600 ms
GVRP Leave-all Timer  : 10000 ms

=====
Configuration that is being used:

block-learning ethe 3
block-applicant ethe 7 ethe 11
enable ethe 1 to 7 ethe 1 ethe 7 ethe 11

=====

Spanning Tree: SINGLE SPANNING TREE
Dropped Packets Count: 0

=====

Number of VLANs in the GVRP Database: 15
Maximum Number of VLANs that can be present: 4095

=====
```

Syntax: `show gvrp [ethernet <port-num>]`

This display shows the following information.

TABLE 49 CLI display of summary GVRP information

This field...	Displays...
Protocol state	The state of GVRP. The display shows one of the following: <ul style="list-style-type: none"> • GVRP is disabled on the system • GVRP is enabled on the system
GVRP BASE VLAN ID	The ID of the base VLAN used by GVRP.
GVRP MAX Leaveall Timer	The maximum number of ms to which you can set the Leaveall timer. NOTE: To change the maximum value, refer to “Increasing the maximum configurable value of the Leaveall timer” on page 346.
GVRP Join Timer	The value of the Join timer. NOTE: For descriptions of the Join, Leave, and Leaveall timers or to change the timers, refer to “Changing the GVRP timers” on page 347.
GVRP Leave Timer	The value of the Leave timer.

TABLE 49 CLI display of summary GVRP information (Continued)

This field...	Displays...
GVRP Leave-all Timer	The value of the Leaveall timer.
Configuration that is being used	The configuration commands used to enable GVRP on individual ports. If GVRP learning or advertising is disabled on a port, this information also is displayed.
Spanning Tree	The type of STP enabled on the device. NOTE: The current release supports GVRP only with Single STP.
Dropped Packets Count	The number of GVRP packets that the device has dropped. A GVRP packet can be dropped for either of the following reasons: <ul style="list-style-type: none"> GVRP packets are received on a port on which GVRP is not enabled. NOTE: If GVRP support is not globally enabled, the device does not drop the GVRP packets but instead forwards them at Layer 2. <ul style="list-style-type: none"> GVRP packets are received with an invalid GARP Protocol ID. The protocol ID must always be 0x0001.
Number of VLANs in the GVRP Database	The number of VLANs in the GVRP database. NOTE: This number includes the default VLAN (1), the GVRP base VLAN (4093), and the single STP VLAN (4094). These VLANs are not advertised by GVRP but are maintained as “Registration Forbidden”.
Maximum Number of VLANs that can be present	The maximum number of VLANs that can be configured on the device. This number includes statically configured VLANs, VLANs learned through GVRP, and VLANs 1, 4093, and 4094. To change the maximum number of VLANs the device can have, use the system-max vlan <num> command. Refer to “Displaying and modifying system parameter default settings” on page 184.

To display detailed GVRP information for an individual port, enter a command such as the following.

```
PowerConnect# show gvrp ethernet 1
Port 1 -
  GVRP Enabled      : YES
  GVRP Learning    : ALLOWED
  GVRP Applicant   : ALLOWED
  Port State       : UP
  Forwarding       : YES

VLAN Membership:      [VLAN-ID]          [MODE]
                     1                FORBIDDEN
                     2                FIXED
                     1001             NORMAL
                     1003             NORMAL
                     1004             NORMAL
                     1007             NORMAL
                     1009             NORMAL
                     1501             NORMAL
                     2507             NORMAL
                     4001             NORMAL
                     4093             FORBIDDEN
                     4094             FORBIDDEN
```

This display shows the following information.

TABLE 50 CLI display of detailed GVRP information for a port

This field...	Displays...
Port number	The port for which information is being displayed.
GVRP Enabled	Whether GVRP is enabled on the port.
GVRP Learning	Whether the port can learn VLAN information from GVRP.
GVRP Applicant	Whether the port can advertise VLAN information into GVRP.
Port State	The port link state, which can be UP or DOWN.
Forwarding	Whether the port is in the GVRP Forwarding state: <ul style="list-style-type: none"> • NO – The port is in the Blocking state. • YES – The port is in the Forwarding state.
VLAN Membership	The VLANs of which the port is a member. For each VLAN, the following information is shown: <ul style="list-style-type: none"> • VLAN ID – The VLAN ID. • Mode – The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • FIXED – The port will always be a member of this VLAN and the VLAN will always be advertised on this port by GVRP. A port becomes FIXED when you configure the port as a tagged member of a statically configured VLAN. • FORBIDDEN – The VLAN is one of the special VLANs that is not advertised or learned by GVRP. In the current release, the following VLANs are forbidden: the default VLAN (1), the GVRP base VLAN (4093), or the Single STP VLAN (4094). • NORMAL – The port became a member of this VLAN after learning about the VLAN through GVRP. The port membership in the VLAN depends on GVRP. If the VLAN is removed from the ports that send GVRP advertisements to this device, then the port will stop being a member of the VLAN.

Displaying GVRP VLAN information

To display information about all the VLANs on the device, enter the following command.

```
PowerConnect# show gvrp vlan brief
```

```
Number of VLANs in the GVRP Database: 7
Maximum Number of VLANs that can be present: 4095
```

[VLAN-ID]	[MODE]	[VLAN-INDEX]
1	STATIC-DEFAULT	0
7	STATIC	2
11	STATIC	4
1001	DYNAMIC	7
1003	DYNAMIC	8
4093	STATIC-GVRP-BASE-VLAN	6
4094	STATIC-SINGLE-SPAN-VLAN	5

=====

Syntax: `show gvrp vlan all | brief | <vlan-id>`

This display shows the following information.

TABLE 51 CLI display of summary VLAN information for GVRP

This field...	Displays...
Number of VLANs in the GVRP Database	The number of VLANs in the GVRP database. NOTE: This number includes the default VLAN (1), the GVRP base VLAN (4093), and the single STP VLAN (4094). These VLANs are not advertised by GVRP but are included in the total count.
Maximum Number of VLANs that can be present	The maximum number of VLANs that can be configured on the device. This number includes statically configured VLANs, VLANs learned through GVRP, and VLANs 1, 4093, and 4094. To change the maximum number of VLANs the device can have, use the system-max vlan <num> command. Refer to “Displaying and modifying system parameter default settings” on page 184.
VLAN-ID	The VLAN ID.
MODE	The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • STATIC – The VLAN is statically configured and cannot be removed by GVRP. This includes VLANs you have configured as well as the default VLAN (1), base GVRP VLAN (4093), and Single STP VLAN (4094). • DYNAMIC – The VLAN was learned through GVRP.
VLAN-INDEX	A number used as an index into the internal database.

To display detailed information for a specific VLAN, enter a command such as the following.

```
PowerConnect# show gvrp vlan 1001

VLAN-ID: 1001, VLAN-INDEX: 7, STATIC: NO, DEFAULT: NO, BASE-VLAN: NO
Timer to Delete Entry Running: NO
Legend: [S=Slot]

Forbidden Members: None

Fixed Members: None

Normal(Dynamic) Members: (S2) 1
```

This display shows the following information.

TABLE 52 CLI display of summary VLAN information for GVRP

This field...	Displays...
VLAN-ID	The VLAN ID.
VLAN-INDEX	A number used as an index into the internal database.
STATIC	Whether the VLAN is a statically configured VLAN.
DEFAULT	Whether this is the default VLAN.
BASE-VLAN	Whether this is the base VLAN for GVRP.
Timer to Delete Entry Running	Whether all ports have left the VLAN and the timer to delete the VLAN itself is running. The timer is described in the note for the Leave timer in “Changing the GVRP timers” on page 347.
Legend	The meanings of the letter codes used in other parts of the display.
Forbidden Members	The ports that cannot become members of a VLAN advertised or learned by GVRP.

TABLE 52 CLI display of summary VLAN information for GVRP (Continued)

This field...	Displays...
Fixed Members	The ports that are statically configured members of the VLAN. GVRP cannot remove these ports.
Normal(Dynamic) Members	The ports that were added by GVRP. These ports also can be removed by GVRP.
MODE	The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • STATIC – The VLAN is statically configured and cannot be removed by GVRP. This includes VLANs you have configured as well as the default VLAN (1), base GVRP VLAN (4093), and Single STP VLAN (4094). • DYNAMIC – The VLAN was learned through GVRP.

To display detailed information for all VLANs, enter the **show gvrp vlan all** command.

Displaying GVRP statistics

To display GVRP statistics for a port, enter a command such as the following.

```
PowerConnect# show gvrp statistics ethernet 1
PORT 1 Statistics:
  Leave All Received           : 147
  Join Empty Received         : 4193
  Join In Received            : 599
  Leave Empty Received        : 0
  Leave In Received           : 0
  Empty Received              : 588
  Leave All Transmitted       : 157
  Join Empty Transmitted      : 1794
  Join In Transmitted         : 598
  Leave Empty Transmitted     : 0
  Leave In Transmitted        : 0
  Empty Transmitted           : 1248
  Invalid Messages/Attributes Skipped : 0
  Failed Registrations        : 0
```

Syntax: **show gvrp statistics all | ethernet <port-num>**

This display shows the following information for the port.

TABLE 53 CLI display of GVRP statistics

This field...	Displays...
Leave All Received	The number of Leaveall messages received.
Join Empty Received	The number of Join Empty messages received.
Join In Received	The number of Join In messages received.
Leave Empty Received	The number of Leave Empty messages received.
Leave In Received	The number of Leave In messages received.
Empty Received	The number of Empty messages received.
Leave All Transmitted	The number of Leaveall messages sent.
Join Empty Transmitted	The number of Join Empty messages sent.
Join In Transmitted	The number of Join In messages sent.

TABLE 53 CLI display of GVRP statistics (Continued)

This field...	Displays...
Leave Empty Transmitted	The number of Leave Empty messages sent.
Leave In Transmitted	The number of Leave In messages sent.
Empty Transmitted	The number of Empty messages sent.
Invalid Messages/Attributes Skipped	The number of invalid messages or attributes received or skipped. This can occur in the following cases: <ul style="list-style-type: none"> • The incoming GVRP PDU has an incorrect length. • "End of PDU" was reached before the complete attribute could be parsed. • The Attribute Type of the attribute that was being parsed was not the GVRP VID Attribute Type (0x01). • The attribute that was being parsed had an invalid attribute length. • The attribute that was being parsed had an invalid GARP event. • The attribute that was being parsed had an invalid VLAN ID. The valid range is 1 – 4095.
Failed Registrations	The number of failed registrations that have occurred. A failed registration can occur for the following reasons: <ul style="list-style-type: none"> • Join requests were received on a port that was blocked from learning dynamic VLANs (GVRP Blocking state). • An entry for a new GVRP VLAN could not be created in the GVRP database.

To display GVRP statistics for all ports, enter the **show gvrp statistics all** command.

Displaying CPU utilization statistics

You can display CPU utilization statistics for GVRP.

To display CPU utilization statistics for GVRP for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.

```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP              0.01        0.03        0.09        0.22         9
BGP              0.00        0.00        0.00        0.00         0
GVRP           0.00       0.03       0.04       0.07        4
ICMP            0.00        0.00        0.00        0.00         0
IP              0.00        0.00        0.00        0.00         0
OSPF            0.00        0.00        0.00        0.00         0
RIP             0.00        0.00        0.00        0.00         0
STP             0.00        0.00        0.00        0.00         0
VRRP           0.00        0.00        0.00        0.00         0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. An example is given below.

12 Displaying GVRP information

```
PowerConnect# show process cpu
The system has only been up for 6 seconds.
Process Name   5Sec(%)   1Min(%)   5Min(%)   15Min(%)   Runtime(ms)
ARP            0.01      0.00      0.00      0.00        0
BGP            0.00      0.00      0.00      0.00        0
GVRP           0.00      0.00      0.00      0.00        0
ICMP           0.01      0.00      0.00      0.00        1
IP             0.00      0.00      0.00      0.00        0
OSPF           0.00      0.00      0.00      0.00        0
RIP            0.00      0.00      0.00      0.00        0
STP            0.00      0.00      0.00      0.00        0
VRRP           0.00      0.00      0.00      0.00        0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.

```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name   Sec(%)   Time(ms)
ARP            0.00      0
BGP            0.00      0
GVRP           0.01      1
ICMP           0.00      0
IP             0.00      0
OSPF           0.00      0
RIP            0.00      0
STP            0.01      1
VRRP           0.00      0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The `<num>` parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Displaying GVRP diagnostic information

To display diagnostic information, enter the following command.

```

PowerConnect# debug gvrp packets
      GVRP:  Packets debugging is on
GVRP: 0x2095ced4: 01 80 c2 00 00 21 00 e0 52 ab 87 40 00 3a 42 42
GVRP: 0x2095cee4: 03 00 01 01 02 00 04 05 00 02 04 05 00 07 04 05
GVRP: 0x2095cef4: 00 09 04 05 00 0b 04 02 03 e9 04 01 03 eb 04 01
GVRP: 0x2095cf04: 03 ec 04 01 03 ef 04 01 03 f1 04 01 05 dd 04 01
GVRP: 0x2095cf14: 09 cb 04 01 0f a1 00 00
GVRP: Port 1 RCV
GVRP: 0x2095ced4: 01 80 c2 00 00 21 00 e0 52 ab 87 40 00 28 42 42
GVRP: 0x2095cee4: 03 00 01 01 04 02 03 e9 04 01 03 eb 04 01 03 ec
GVRP: 0x2095cef4: 04 01 03 ef 04 01 03 f1 04 01 05 dd 04 01 09 cb
GVRP: 0x2095cf04: 04 01 0f a1 00 00
GVRP: Port 1 TX
GVRP: 0x207651b8: 01 80 c2 00 00 21 00 04 80 2c 0e 20 00 3a 42 42
GVRP: 0x207651c8: 03 00 01 01 02 00 04 05 03 e9 04 05 03 eb 04 05
GVRP: 0x207651d8: 03 ec 04 05 03 ef 04 05 03 f1 04 05 05 dd 04 05
GVRP: 0x207651e8: 09 cb 04 05 0f a1 04 02 00 02 04 01 00 07 04 01
GVRP: 0x207651f8: 00 09 04 01 00 0b 00 00
GVRP: Port 1 TX
GVRP: 0x207651b8: 01 80 c2 00 00 21 00 04 80 2c 0e 20 00 18 42 42
GVRP: 0x207651c8: 03 00 01 01 04 02 00 02 04 01 00 07 04 01 00 09
GVRP: 0x207651d8: 04 01 00 0b 00 00

```

Syntax: debug gvrp packets

Clearing GVRP statistics

To clear the GVRP statistics counters, enter a command such as the following.

```
PowerConnect# clear gvrp statistics all
```

This command clears the counters for all ports. To clear the counters for a specific port only, enter a command such as the following.

```
PowerConnect# clear gvrp statistics ethernet 1
```

Syntax: clear gvrp statistics all | ethernet <portnum>

CLI examples

The following sections show the CLI commands for implementing the applications of GVRP described in [“Application examples”](#) on page 341.

NOTE

Although some of the devices in these configuration examples do not have statically configured VLANs, this is not a requirement. You always can have statically configured VLANs on a device that is running GVRP.

Dynamic core and fixed edge

In this configuration, the edge devices advertise their statically configured VLANs to the core device. The core device does not have any statically configured VLANs but learns the VLANs from the edge devices.

Enter the following commands on the core device.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable all
```

These commands globally enable GVRP support and enable the protocol on all ports.

Enter the following commands on edge device A.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# untag ethernet 5
PowerConnect(config-vlan-20)# tag ethernet 6
PowerConnect(config-vlan-20)# vlan 40
PowerConnect(config-vlan-40)# untag ethernet 5
PowerConnect(config-vlan-40)# tag ethernet 6
PowerConnect(config-vlan-40)# exit
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable ethernet 6
PowerConnect(config-gvrp)# block-learning ethernet 6
```

These commands statically configure two port-based VLANs, enable GVRP on port 6, and block GVRP learning on the port. The device will advertise the VLANs but will not learn VLANs from other devices.

Enter the following commands on edge device B.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# untag ethernet 8
PowerConnect(config-vlan-20)# tag ethernet 7
PowerConnect(config-vlan-20)# vlan 30
PowerConnect(config-vlan-30)# untag ethernet 6
PowerConnect(config-vlan-30)# tag ethernet 7
PowerConnect(config-vlan-30)# exit
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable ethernet 7
PowerConnect(config-gvrp)# block-learning ethernet 7
```

Enter the following commands on edge device C.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# vlan 30
PowerConnect(config-vlan-30)# untag ethernet 8
PowerConnect(config-vlan-30)# tag ethernet 7
PowerConnect(config-vlan-20)# vlan 40
PowerConnect(config-vlan-40)# untag ethernet 6
PowerConnect(config-vlan-40)# tag ethernet 7
PowerConnect(config-vlan-40)# exit
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable ethernet 7
PowerConnect(config-gvrp)# block-learning ethernet 7
```

Dynamic core and dynamic edge

In this configuration, the core and edge devices have no statically configured VLANs and are enabled to learn and advertise VLANs. The edge and core devices learn the VLANs configured on the devices in the edge clouds. To enable GVRP on all the ports, enter the following command on each edge device **and** on the core device.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable all
```

Fixed core and dynamic edge

In this configuration, GVRP learning is enabled on the edge devices. The VLANs on the core device are statically configured, and the core device is enabled to advertise its VLANs but not to learn VLANs. The edge devices learn the VLANs from the core.

Enter the following commands on the core device.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# tag ethernet 1
PowerConnect(config-vlan-20)# tag ethernet 2
PowerConnect(config-vlan-20)# vlan 30
PowerConnect(config-vlan-30)# tag ethernet 2
PowerConnect(config-vlan-30)# tag ethernet 3
PowerConnect(config-vlan-30)# vlan 40
PowerConnect(config-vlan-40)# tag ethernet 4
PowerConnect(config-vlan-40)# tag ethernet 3
PowerConnect(config-vlan-40)# vlan 50
PowerConnect(config-vlan-50)# untag ethernet 5
PowerConnect(config-vlan-50)# tag ethernet 6
PowerConnect(config-vlan-50)# exit
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable ethernet 1 ethernet 2 ethernet 3
PowerConnect(config-gvrp)# block-learning ethernet 1 ethernet 2 ethernet 3
```

These VLAN commands configure VLANs 20, 30, 40, and 50. The GVRP commands enable the protocol on the ports that are connected to the edge devices, and disable VLAN learning on those ports. All the VLANs are advertised by GVRP.

Enter the following commands on edge devices A, B, and C.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# gvrp-enable
PowerConnect(config-gvrp)# enable all
PowerConnect(config-gvrp)# block-applicant all
```

Fixed core and fixed edge

The VLANs are statically configured on the core and edge devices. On each edge device, VLAN advertising is enabled but learning is disabled. GVRP is not configured on the core device. This configuration enables the devices in the edge clouds to learn the VLANs configured on the edge devices.

12 CLI examples

This configuration does not use any GVRP configuration on the core device.

The configuration on the edge device is the same as in [“Dynamic core and fixed edge”](#) on page 357.

Configuring Rule-Based IP Access Control Lists

ACL overview

This chapter describes how Access Control Lists (ACLs) are implemented and configured in the PowerConnect B-Series T124X devices.

Devices support **rule-based ACLs** (sometimes called hardware-based ACLs), where the decisions to permit or deny packets are processed in hardware and all permitted packets are switched or routed in hardware. All denied packets are also dropped in hardware. In addition, PowerConnect B-Series T124X devices support inbound ACLs only. Outbound ACLs are not supported.

NOTE

PowerConnect B-Series T124X devices support hardware-based ACLs only. These devices do not support flow-based ACLs.

Rule-based ACLs program the ACL entries you assign to an interface into Content Addressable Memory (CAM) space allocated for the ports. The ACLs are programmed into hardware at startup (or as new ACLs are entered and bound to ports). Devices that use rule-based ACLs program the ACLs into the CAM entries and use these entries to permit or deny packets in the hardware, without sending the packets to the CPU for processing.

Rule-based ACLs are supported on the following interface types:

- Gbps Ethernet ports
- 10 Gbps Ethernet ports
- Trunk groups
- Virtual routing interfaces

Types of IP ACLs

You can configure the following types of IP ACLs:

- **Standard** – Permits or denies packets based on source IP address. Valid standard ACL IDs are 1 – 99 or a character string.
- **Extended** – Permits or denies packets based on source and destination IP address and also based on IP protocol information. Valid extended ACL IDs are a number from 100 – 199 or a character string.

ACL IDs and entries

ACLs consist of ACL IDs and ACL entries:

- **ACL ID** – An ACL ID is a number from 1 – 99 (for a standard ACL) or 100 – 199 (for an extended ACL) or a character string. The ACL ID identifies a collection of individual ACL entries. When you apply ACL entries to an interface, you do so by applying the ACL ID that contains the ACL entries to the interface, instead of applying the individual entries to the interface. This makes applying large groups of access filters (ACL entries) to interfaces simple. Refer to “[Numbered and named ACLs](#)” on page 362.

NOTE

This is different from IP access policies. If you use IP access policies, you apply the individual policies to interfaces.

- **ACL entry** – Also called an **ACL rule**, this is a filter command associated with an ACL ID. The maximum number of ACL rules you can configure is a system-wide parameter and depends on the device you are configuring. You can configure up to the maximum number of entries in any combination in different ACLs. The total number of entries in all ACLs cannot exceed the system maximum, listed in [Table 54](#).

TABLE 54 Maximum number of ACL entries

System	Maximum ACL rules per port region	Maximum ACL entries per system
PowerConnect B-Series TI24X Layer 2 or Layer 3 Switch	1534	1534

The PowerConnect B-Series TI24X supports a maximum of 1015 ACL entries per ACL.

You configure ACLs on a global basis, then apply them to the incoming traffic on specific ports. The software applies the entries within an ACL in the order they appear in the ACL configuration. As soon as a match is found, the software takes the action specified in the ACL entry (permit or deny the packet) and stops further comparison for that packet.

Numbered and named ACLs

When you configure an ACL, you can refer to the ACL by a numeric ID or by an alphanumeric name. The commands to configure numbered ACLs are different from the commands for named ACLs.

- **Numbered ACL** – If you refer to the ACL by a numeric ID, you can use 1 – 99 for a standard ACL or 100 – 199 for an extended ACL.
- **Named ACL** – If you refer to the ACL by a name, you specify whether the ACL is a standard ACL or an extended ACL, then specify the name.

You can configure up to 99 standard numbered IP ACLs and 100 extended numbered IP ACLs. You also can configure up to 99 standard named ACLs and 100 extended named ACLs by number.

Default ACL action

The default action when no ACLs are configured on a device is to permit all traffic. However, once you configure an ACL and apply it to a port, the default action for that port is to deny all traffic that is not explicitly permitted on the port:

- If you want to tightly control access, configure ACLs consisting of permit entries for the access you want to permit. The ACLs implicitly deny all other access.

- If you want to secure access in environments with many users, you might want to configure ACLs that consist of explicit deny entries, then add an entry to permit all access to the end of each ACL. The software permits packets that are not denied by the deny entries.

How hardware-based ACLs work

When you bind an ACL to inbound traffic on an interface, the device programs the Layer 4 CAM with the ACL. Permit and deny rules are programmed. Most ACL rules require one Layer 4 CAM entry. However, ACL rules that match on more than one TCP or UDP application port may require several CAM entries. The Layer 4 CAM entries for ACLs do not age out. They remain in the CAM until you remove the ACL:

- If a packet received on the interface matches an ACL rule in the Layer 4 CAM, the device permits or denies the packet according to the ACL.
- If a packet does not match an ACL rule, the packet is dropped, since the default action on an interface that has ACLs is to deny the packet.

How fragmented packets are processed

The descriptions above apply to non-fragmented packets. The default processing of fragments by hardware-based ACLs is as follows:

- The first fragment of a packet is permitted or denied using the ACLs. The first fragment is handled the same way as non-fragmented packets, since the first fragment contains the Layer 4 source and destination application port numbers. The device uses the Layer 4 CAM entry if one is programmed, or applies the interface's ACL entries to the packet and permits or denies the packet according to the first matching ACL.
- For other fragments of the same packet, they are subject to a rule only if there is no Layer 4 information in the rule or in any preceding rules.

The fragments are forwarded even if the first fragment, which contains the Layer 4 information, was denied. Generally, denying the first fragment of a packet is sufficient, since a transaction cannot be completed without the entire packet.

For tighter control, you can configure the port to drop all packet fragments. Refer to [“Enabling strict control of ACL filtering of fragmented packets”](#) on page 383.

Hardware aging of Layer 4 CAM entries

Rule-based ACLs use Layer 4 CAM entries. The device permanently programs rule-based ACLs into the CAM. The entries never age out.

Configuration considerations

- PowerConnect B-Series TI24X devices support inbound ACLs. Outbound ACL are not supported.
- Hardware-based ACLs are supported on:
 - Gbps Ethernet ports
 - 10 Gbps Ethernet ports

13 Configuring standard numbered ACLs

- Trunk groups
- Virtual routing interfaces
- ACLs on the PowerConnect B-Series TI24X devices apply to all traffic, including management traffic.
- ACL logging is supported for denied packets and packets that are sent to the CPU to generate the log if logging is enabled on the port and the ACL that is applied to that port. ACL logging is not supported for packets that are processed in hardware (permitted packets).
- The number of ACL rules supported per device is listed in [Table 54](#).
- Hardware-based ACLs support only one ACL per port. The ACL of course can contain multiple entries (rules). For example, hardware-based ACLs do not support ACLs 101 and 102 on port 1, but hardware-based ACLs do support ACL 101 containing multiple entries.
- By default, the first fragment of a fragmented packet received by the device is permitted or denied using the ACLs, but subsequent fragments of the same packet are forwarded in hardware. Generally, denying the first fragment of a packet is sufficient, since a transaction cannot be completed without the entire packet.
- The following ACL features and options are not supported on the PowerConnect B-Series TI24X devices:
 - Applying an ACL on a device that has Super Aggregated VLANs (SAVs) enabled.
 - ACL logging – ACL logging is supported for packets that are sent to the CPU for processing (denied packets). ACL logging is not supported for packets that are processed in hardware (permitted packets).
 - Flow-based ACLs
- PowerConnect B-Series TI24X devices support MAC filters instead of Layer 2 ACLs.
- You can apply an ACL to a port that has TCP SYN protection or ICMP smurf protection, or both, enabled.

NOTE

PowerConnect B-Series TI24X does not support ACLs on Group VEs, even though the CLI contains commands for this action.

Configuring standard numbered ACLs

This section describes how to configure standard numbered ACLs with numeric IDs and provides configuration examples.

Standard ACLs permit or deny packets based on source IP address. You can configure up to 99 standard numbered ACLs. There is no limit to the number of ACL entries an ACL can contain except for the system-wide limitation. For the number of ACL entries supported on a device, refer to [“ACL IDs and entries”](#) on page 361.

Standard numbered ACL syntax

Syntax: `[no] access-list <ACL-num> deny | permit <source-ip> | <hostname> <wildcard> [log]`

or

Syntax: `[no] access-list <ACL-num> deny | permit <source-ip>/<mask-bits> | <hostname> [log]`

Syntax: [no] access-list <ACL-num> deny | permit host <source-ip> | <hostname> [log]

Syntax: [no] access-list <ACL-num> deny | permit any [log]

Syntax: [no] ip access-group <ACL-num> in

The <ACL-num> parameter is the access list number from 1 – 99.

The **deny | permit** parameter indicates whether packets that match a policy in the access list are denied (dropped) or permitted (forwarded).

The <source-ip> parameter specifies the source IP address. Alternatively, you can specify the host name.

NOTE

To specify the host name instead of the IP address, the host name must be configured using the DNS resolver on the device. To configure the DNS resolver name, use the **ip dns server-address...** command at the global CONFIG level of the CLI.

The <wildcard> parameter specifies the mask value to compare against the host address specified by the <source-ip> parameter. The <wildcard> is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the <source-ip>. Ones mean any value matches. For example, the <source-ip> and <wildcard> values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in CIDR format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “209.157.22.26 0.0.0.255” as “209.157.22.26/24”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into ones. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of subnet lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/<mask-bits>” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** command.

The **host** <source-ip> | <hostname> parameter lets you specify a host IP address or name. When you use this parameter, you do not need to specify the mask. A mask of all zeros (0.0.0.0) is implied.

The **any** parameter configures the policy to match on all host addresses.

The **log** argument configures the device to generate Syslog entries and SNMP traps for packets that are denied by the access policy.

The **in** parameter applies the ACL to incoming traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or virtual interface.

NOTE

If the ACL is for a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface.

Configuration example for standard numbered ACLs

To configure a standard ACL and apply it to incoming traffic on port 1, enter the following commands.

```
PowerConnect(config)# access-list 1 deny host 209.157.22.26 log
PowerConnect(config)# access-list 1 deny 209.157.29.12 log
PowerConnect(config)# access-list 1 deny host IPHost1 log
PowerConnect(config)# access-list 1 permit any
PowerConnect(config)# int eth 1
PowerConnect(config-if-1)# ip access-group 1 in
PowerConnect(config)# write memory
```

The commands in this example configure an ACL to deny packets from three source IP addresses from being received on port 1. The last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries.

Configuring standard named ACLs

This section describes how to configure standard named ACLs with alphanumeric IDs. This section also provides configuration examples.

Standard ACLs permit or deny packets based on source IP address. You can configure up to 99 standard named ACLs. There is no limit to the number of ACL entries an ACL can contain except for the system-wide limitation. For the number of ACL entries supported on a device, refer to [“ACL IDs and entries”](#) on page 361.

The commands for configuring named ACL entries are different from the commands for configuring numbered ACL entries. The command to configure a numbered ACL is **access-list**. The command for configuring a named ACL is **ip access-list**. In addition, when you configure a numbered ACL entry, you specify all the command parameters on the same command. When you configure a named ACL, you specify the ACL type (standard or extended) and the ACL name with one command, which places you in the configuration level for that ACL. Once you enter the configuration level for the ACL, the command syntax is the same as the syntax for numbered ACLs.

Standard named ACL syntax

Syntax: [no] ip access-list standard <ACL-name> | <ACL-num>

Syntax: deny | permit <source-ip> | <hostname> <wildcard> [log]

or

Syntax: deny | permit <source-ip>/<mask-bits> | <hostname> [log]

Syntax: deny | permit host <source-ip> | <hostname> [log]

Syntax: deny | permit any [log]

Syntax: [no] ip access-group <ACL-name> in

The `<ACL-name>` parameter is the access list name. You can specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, "ACL for Net1").

The `<ACL-num>` parameter allows you to specify an ACL number if you prefer. If you specify a number, you can specify from 1 – 99 for standard ACLs.

NOTE

For convenience, the software allows you to configure numbered ACLs using the syntax for named ACLs. The software also still supports the older syntax for numbered ACLs. Although the software allows both methods for configuring numbered ACLs, numbered ACLs are always formatted in the startup-config and running-config files in using the older syntax, as follows.

```
access-list 1 deny host 209.157.22.26 log
access-list 1 deny 209.157.22.0 0.0.0.255 log
access-list 1 permit any
access-list 101 deny tcp any any eq http log
```

The **deny** | **permit** parameter indicates whether packets that match a policy in the access list are denied (dropped) or permitted (forwarded).

The `<source-ip>` parameter specifies the source IP address. Alternatively, you can specify the host name.

NOTE

To specify the host name instead of the IP address, the host name must be configured using the DNS resolver on the device. To configure the DNS resolver name, use the **ip dns server-address...** command at the global CONFIG level of the CLI.

The `<wildcard>` parameter specifies the mask value to compare against the host address specified by the `<source-ip>` parameter. The `<wildcard>` is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the `<source-ip>`. Ones mean any value matches. For example, the `<source-ip>` and `<wildcard>` values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in CIDR format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of "209.157.22.26 0.0.0.255" as "209.157.22.26/24". The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into ones. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of subnet lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in `"/<mask-bits>"` format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** command.

The **host** <source-ip> | <hostname> parameter lets you specify a host IP address or name. When you use this parameter, you do not need to specify the mask. A mask of all zeros (0.0.0.0) is implied.

The **any** parameter configures the policy to match on all host addresses.

The **log** argument configures the device to generate Syslog entries and SNMP traps for packets that are denied by the access policy.

The **in** parameter applies the ACL to incoming traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or virtual interface.

NOTE

If the ACL is bound to a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface.

Configuration example for standard named ACLs

To configure a standard named ACL, enter commands such as the following.

```
PowerConnect(config)#ip access-list standard Net1
PowerConnect(config-std-nACL)# deny host 209.157.22.26 log
PowerConnect(config-std-nACL)# deny 209.157.29.12 log
PowerConnect(config-std-nACL)# deny host IPHost1 log
PowerConnect(config-std-nACL)# permit any
PowerConnect(config-std-nACL)# exit
PowerConnect(config)#int eth 1
PowerConnect(config-if-e10000-1)# ip access-group Net1 in
```

The commands in this example configure a standard ACL named “Net1”. The entries in this ACL deny packets from three source IP addresses from being forwarded on port 1. Since the implicit action for an ACL is “deny”, the last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries. For an example of how to configure the same entries in a numbered ACL, refer to [“Configuring standard numbered ACLs”](#) on page 364.

Notice that the command prompt changes after you enter the ACL type and name. The “std” in the command prompt indicates that you are configuring entries for a standard ACL. For an extended ACL, this part of the command prompt is “ext”. The “nACL” indicates that you are configuring a named ACL.

Configuring extended numbered ACLs

This section describes how to configure extended numbered ACLs.

Extended ACLs let you permit or deny packets based on the following information:

- IP protocol
- Source IP address or host name

- Destination IP address or host name
- Source TCP or UDP port (if the IP protocol is TCP or UDP)
- Destination TCP or UDP port (if the IP protocol is TCP or UDP)

The IP protocol can be one of the following well-known names or any IP protocol number from 0 – 255:

- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- Internet Gateway Routing Protocol (IGRP)
- Internet Protocol (IP)
- Open Shortest Path First (OSPF)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

For TCP and UDP, you also can specify a comparison operator and port name or number. For example, you can configure a policy to block web access to a specific website by denying all TCP port 80 (HTTP) packets from a specified source IP address to the website IP address.

NOTE

PowerConnect support extended ACLs.

Extended numbered ACL syntax

Syntax: `[no] access-list <ACL-num> deny | permit <ip-protocol> <source-ip> | <hostname> <wildcard> [<operator> <source-tcp/udp-port>] <destination-ip> | <hostname> [<icmp-num> | <icmp-type>] <wildcard> [<tcp/udp comparison operator> <destination-tcp/udp-port>] [dscp-marking <0-63>] [802.1p-priority-marking <0 - 7>... | [802.1p-and-internal-marking] [internal-priority-marking] [dscp-matching <0-63>] [log] [precedence <name> | <0 - 7>] [tos <0 - 63> | <name>] [traffic policy <name>]`

Syntax: `[no] access-list <ACL-num> deny | permit host <ip-protocol> any any`

Syntax: `[no] ip access-group <ACL-num> in`

The `<ACL-num>` parameter is the extended access list number. Specify a number from 100 – 199.

The `deny | permit` parameter indicates whether packets that match the policy are dropped or forwarded.

The `<ip-protocol>` parameter indicates the type of IP packet you are filtering. You can specify a well-known name for any protocol whose number is less than 255. For other protocols, you must enter the number. Enter “?” instead of a protocol to list the well-known names recognized by the CLI.

The `<source-ip> | <hostname>` parameter specifies the source IP host for the policy. If you want the policy to match on all source addresses, enter `any`.

The *<wildcard>* parameter specifies the portion of the source IP host address to match against. The *<wildcard>* is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the *<source-ip>*. Ones mean any value matches. For example, the *<source-ip>* and *<wildcard>* values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “209.157.22.26 0.0.0.255” as “209.157.22.26/24”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of subnet lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/*<mask-bits>*” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** command.

The *<destination-ip>* | *<hostname>* parameter specifies the destination IP host for the policy. If you want the policy to match on all destination addresses, enter **any**.

The *<icmp-type>* | *<icmp-num>* parameter specifies the ICMP protocol type:

- This parameter applies only if you specified **icmp** as the *<ip-protocol>* value.
- If you use this parameter, the ACL entry is sent to the CPU for processing.
- If you do not specify a message type, the ACL applies to all types of ICMP messages.

The *<icmp-num>* parameter can be a value from 0 – 255.

The *<icmp-type>* parameter can have one of the following values, depending on the software version the device is running:

- any-icmp-type
- echo
- echo-reply
- information-request
- log
- mask-reply
- mask-request
- parameter-problem
- redirect
- source-quench

- time-exceeded
- timestamp-reply
- timestamp-request
- traffic policy
- unreachable
- *<num>*

The *<tcp/udp comparison operator>* parameter specifies a comparison operator for the TCP or UDP port number. This parameter applies only when you specify **tcp** or **udp** as the IP protocol. For example, if you are configuring an entry for HTTP, specify **tcp eq http**. You can enter one of the following operators:

- **eq** – The policy applies to the TCP or UDP port name or number you enter after **eq**.
- **established** – This operator applies only to TCP packets. If you use this operator, the policy applies to TCP packets that have the ACK (Acknowledgment) or RST (Reset) bits set on (set to “1”) in the Control Bits field of the TCP packet header. Thus, the policy applies only to established TCP sessions, not to new sessions. Refer to Section 3.1, “Header Format”, in RFC 793 for information about this field.

NOTE

This operator applies only to destination TCP ports, not source TCP ports.

- **gt** – The policy applies to TCP or UDP port numbers greater than the port number or the numeric equivalent of the port name you enter after **gt**.
- **lt** – The policy applies to TCP or UDP port numbers that are less than the port number or the numeric equivalent of the port name you enter after **lt**.
- **neq** – The policy applies to all TCP or UDP port numbers except the port number or port name you enter after **neq**.
- **range** – The policy applies to all TCP or UDP port numbers that are between the first TCP or UDP port name or number and the second one you enter following the **range** parameter. The **range** includes the port names or numbers you enter. For example, to apply the policy to all ports between and including 23 (Telnet) and 53 (DNS), enter the following: **range 23 53**. The first port number in the **range** must be lower than the last number in the **range**.

The *<tcp/udp-port>* parameter specifies the TCP or UDP port number or well-known name. You can specify a well-known name for any application port whose number is less than 1024. For other application ports, you must enter the number. Enter “?” instead of a port to list the well-known names recognized by the CLI.

The **in** parameter specifies that the ACL applies to incoming traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or a virtual interface.

NOTE

If the ACL is for a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface. Refer to [“Configuring standard numbered ACLs”](#) on page 364.

The **precedence** *<name>* | *<num>* parameter of the **ip access-list** command specifies the IP precedence. The precedence option for of an IP packet is set in a three-bit field following the four-bit header-length field of the packet header. You can specify one of the following:

- **critical** or **5** – The ACL matches packets that have the critical precedence. If you specify the option number instead of the name, specify number 5.

- **flash** or **3** – The ACL matches packets that have the flash precedence. If you specify the option number instead of the name, specify number 3.
- **flash-override** or **4** – The ACL matches packets that have the flash override precedence. If you specify the option number instead of the name, specify number 4.
- **immediate** or **2** – The ACL matches packets that have the immediate precedence. If you specify the option number instead of the name, specify number 2.
- **internet** or **6** – The ACL matches packets that have the internetwork control precedence. If you specify the option number instead of the name, specify number 6.
- **network** or **7** – The ACL matches packets that have the network control precedence. If you specify the option number instead of the name, specify number 7.
- **priority** or **1** – The ACL matches packets that have the priority precedence. If you specify the option number instead of the name, specify number 1.
- **routine** or **0** – The ACL matches packets that have the routine precedence. If you specify the option number instead of the name, specify number 0.

The **tos** *<name>* | *<num>* parameter of the **ip access-list** command specifies the IP ToS. You can specify one of the following:

- **max-reliability** or **2** – The ACL matches packets that have the maximum reliability ToS. The decimal value for this option is 2.
- **max-throughput** or **4** – The ACL matches packets that have the maximum throughput ToS. The decimal value for this option is 4.
- **min-delay** or **8** – The ACL matches packets that have the minimum delay ToS. The decimal value for this option is 8.
- **min-monetary-cost** or **1** – The ACL matches packets that have the minimum monetary cost ToS. The decimal value for this option is 1.

NOTE

This value is not supported on 10 Gbps Ethernet modules.

- **normal** or **0** – The ACL matches packets that have the normal ToS. The decimal value for this option is 0.
- *<num>* – A number from 0 – 15 that is the sum of the numeric values of the options you want. The ToS field is a four-bit field following the Precedence field in the IP header. You can specify one or more of the following. To select more than one option, enter the decimal value that is equivalent to the sum of the numeric values of all the ToS options you want to select. For example, to select the **max-reliability** and **min-delay** options, enter number 10. To select all options, select 15.

The **dscp-marking** option enables you to configure an ACL that marks matching packets with a specified DSCP value. Enter a value from 0 – 63. Refer to [“Using an IP ACL to mark DSCP values \(DSCP marking\)”](#) on page 387.

The **dscp-matching** option matches on the packet DSCP value. Enter a value from 0 – 63. This option does not change the packet forwarding priority through the device or mark the packet. Refer to [“DSCP matching”](#) on page 389.

The **log** parameter enables SNMP traps and Syslog messages for packets denied by the ACL.

You can enable logging on ACLs and filters that support logging even when the ACLs and filters are already in use. To do so, re-enter the ACL or filter command and add the **log** parameter to the end of the ACL or filter. The software replaces the ACL or filter command with the new one. The new ACL or filter, with logging enabled, takes effect immediately.

The **traffic-policy** option enables the device to rate limit inbound traffic and to count the packets and bytes per packet to which ACL permit or deny clauses are applied. For configuration procedures and examples, refer to [Chapter 17, “Configuring Traffic Policies”](#).

Configuration examples for extended numbered ACLs

To configure an extended access list that blocks all Telnet traffic received on port 1 from IP host 209.157.22.26, enter the following commands.

```
PowerConnect(config)#access-list 101 deny tcp host 209.157.22.26 any eq telnet
log
PowerConnect(config)#access-list 101 permit ip any any
PowerConnect(config)#int eth 1
PowerConnect(config-if-e10000-1)#ip access-group 101 in
```

Here is another example of commands for configuring an extended ACL and applying it to an interface. These examples show many of the syntax choices. Notice that some of the entries are configured to generate log entries while other entries are not thus configured.

```
PowerConnect(config)# access-list 102 perm icmp 209.157.22.0/24 209.157.21.0/24
PowerConnect(config)# access-list 102 deny igmp host rkwong 209.157.21.0/24 log
PowerConnect(config)# access-list 102 deny igmp 209.157.21.0/24 host rkwong log
PowerConnect(config)# access-list 102 deny ip host 209.157.21.100 host
209.157.22.1 log
PowerConnect(config)# access-list 102 deny ospf any any log
PowerConnect(config)# access-list 102 permit ip any any
```

The first entry permits ICMP traffic from hosts in the 209.157.22.x network to hosts in the 209.157.21.x network.

The second entry denies IGMP traffic from the host device named “rkwong” to the 209.157.21.x network.

The third entry denies IGMP traffic from the 209.157.21.x network to the host device named “rkwong”.

The fourth entry denies all IP traffic from host 209.157.21.100 to host 209.157.22.1 and generates Syslog entries for packets that are denied by this entry.

The fifth entry denies all OSPF traffic and generates Syslog entries for denied traffic.

The sixth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 102 to the incoming traffic on port 2 and to the incoming traffic on port 3.

```
PowerConnect(config)# int eth 2
PowerConnect(config-if-2)# ip access-group 102 in
PowerConnect(config-if-2)# exit
PowerConnect(config)# int eth 3
PowerConnect(config-if-3)# ip access-group 102 in
PowerConnect(config)# write memory
```

Here is another example of an extended ACL.

13 Configuring extended named ACLs

```
PowerConnect(config)# access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24
PowerConnect(config)# access-list 103 deny tcp 209.157.21.0/24 eq ftp
209.157.22.0/24
PowerConnect(config)# access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24
lt telnet neq 5
PowerConnect(config)# access-list 103 deny udp any range 5 6 209.157.22.0/24
```

The first entry in this ACL denies TCP traffic from the 209.157.21.x network to the 209.157.22.x network.

The second entry denies all FTP traffic from the 209.157.21.x network to the 209.157.22.x network.

The third entry denies TCP traffic from the 209.157.21.x network to the 209.157.22.x network, if the TCP port number of the traffic is less than the well-known TCP port number for Telnet (23), and if the TCP port is not equal to 5. Thus, TCP packets whose TCP port numbers are 5 or are greater than 23 are allowed.

The fourth entry denies UDP packets from any source to the 209.157.22.x network, if the UDP port number from the source network is 5 or 6 and the destination UDP port is 7 or 8.

The fifth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 103 to the incoming traffic on ports 1 and 2.

```
PowerConnect(config)# int eth 1
PowerConnect(config-if-1)# ip access-group 103 in
PowerConnect(config-if-1)# exit
PowerConnect(config)# int eth 2
PowerConnect(config-if-2)# ip access-group 103 in
PowerConnect(config)# write memory
```

Configuring extended named ACLs

The commands for configuring named ACL entries are different from the commands for configuring numbered ACL entries. The command to configure a numbered ACL is **access-list**. The command for configuring a named ACL is **ip access-list**. In addition, when you configure a numbered ACL entry, you specify all the command parameters on the same command. When you configure a named ACL, you specify the ACL type (standard or extended) and the ACL number with one command, which places you in the configuration level for that ACL. Once you enter the configuration level for the ACL, the command syntax is the same as the syntax for numbered ACLs.

Extended ACLs let you permit or deny packets based on the following information:

- IP protocol
- Source IP address or host name
- Destination IP address or host name
- Source TCP or UDP port (if the IP protocol is TCP or UDP)
- Destination TCP or UDP port (if the IP protocol is TCP or UDP)

The IP protocol can be one of the following well-known names or any IP protocol number from 0 – 255:

- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- Internet Gateway Routing Protocol (IGRP)
- Internet Protocol (IP)
- Open Shortest Path First (OSPF)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

For TCP and UDP, you also can specify a comparison operator and port name or number. For example, you can configure a policy to block web access to a specific website by denying all TCP port 80 (HTTP) packets from a specified source IP address to the website's IP address.

Extended named ACL syntax

Syntax: `[no] ip access-list extended <ACL-name> deny | permit <ip-protocol> <source-ip> | <hostname> <wildcard> [<operator> <source-tcp/udp-port>] <destination-ip> | <hostname> [<icmp-num> | <icmp-type>] <wildcard> [<tcp/udp comparison operator> <destination-tcp/udp-port>] [dscp-marking <0-63> [802.1p-priority-marking <0 -7>... | [802.1p-and-internal-marking] [internal-priority-marking] [dscp-matching <0-63>] [log] [precedence <name> | <0 - 7>] [tos <0 - 63> | <name>] [traffic policy <name>]`

Syntax: `[no] access-list <num> deny | permit host <ip-protocol> any any`

Syntax: `[no] ip access-group <num> in`

The `<ACL-name>` parameter is the access list name. You can specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, "ACL for Net1").

The `deny` | `permit` parameter indicates whether packets that match the policy are dropped or forwarded.

The `<ip-protocol>` parameter indicates the type of IP packet you are filtering. You can specify a well-known name for any protocol whose number is less than 255. For other protocols, you must enter the number. Enter "?" instead of a protocol to list the well-known names recognized by the CLI.

The `<source-ip>` | `<hostname>` parameter specifies the source IP host for the policy. If you want the policy to match on all source addresses, enter `any`.

The `<wildcard>` parameter specifies the portion of the source IP host address to match against. The `<wildcard>` is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet's source address must match the `<source-ip>`. Ones mean any value matches. For example, the `<source-ip>` and `<wildcard>` values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of "209.157.22.26 0.0.0.255" as "209.157.22.26/24". The CLI automatically converts the CIDR number into the appropriate ACL

mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of subnet lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/<mask-bits>” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** command.

The <destination-ip> | <hostname> parameter specifies the destination IP host for the policy. If you want the policy to match on all destination addresses, enter **any**.

The <icmp-type> | <icmp-num> parameter specifies the ICMP protocol type:

- This parameter applies only if you specified **icmp** as the <ip-protocol> value.
- If you use this parameter, the ACL entry is sent to the CPU for processing.
- If you do not specify a message type, the ACL applies to all types of ICMP messages.

The <icmp-num> parameter can be a value from 0 – 255.

The <icmp-type> parameter can have one of the following values, depending on the software version the device is running:

- any-icmp-type
- echo
- echo-reply
- information-request
- log
- mask-reply
- mask-request
- parameter-problem
- redirect
- source-quench
- time-exceeded
- timestamp-reply
- timestamp-request
- traffic policy
- unreachable
- <num>

The <tcp/udp comparison operator> parameter specifies a comparison operator for the TCP or UDP port number. This parameter applies only when you specify **tcp** or **udp** as the IP protocol. For example, if you are configuring an entry for HTTP, specify **tcp eq http**. You can enter one of the following operators:

- **eq** – The policy applies to the TCP or UDP port name or number you enter after **eq**.
- **established** – This operator applies only to TCP packets. If you use this operator, the policy applies to TCP packets that have the ACK (Acknowledgment) or RST (Reset) bits set on (set to “1”) in the Control Bits field of the TCP packet header. Thus, the policy applies only to established TCP sessions, not to new sessions. Refer to Section 3.1, “Header Format”, in RFC 793 for information about this field.

NOTE

This operator applies only to destination TCP ports, not source TCP ports.

- **gt** – The policy applies to TCP or UDP port numbers greater than the port number or the numeric equivalent of the port name you enter after **gt**.
- **lt** – The policy applies to TCP or UDP port numbers that are less than the port number or the numeric equivalent of the port name you enter after **lt**.
- **neq** – The policy applies to all TCP or UDP port numbers except the port number or port name you enter after **neq**.
- **range** – The policy applies to all TCP or UDP port numbers that are between the first TCP or UDP port name or number and the second one you enter following the range parameter. The range includes the port names or numbers you enter. For example, to apply the policy to all ports between and including 23 (Telnet) and 53 (DNS), enter the following: **range 23 53**. The first port number in the range must be lower than the last number in the range.

The `<tcp/udp-port>` parameter specifies the TCP or UDP port number or well-known name. You can specify a well-known name for any application port whose number is less than 1024. For other application ports, you must enter the number. Enter “?” instead of a port to list the well-known names recognized by the CLI.

The `in` parameter specifies that the ACL applies to incoming traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or a virtual interface.

NOTE

If the ACL is for a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface. Refer to [“Configuring standard numbered ACLs”](#) on page 364.

The **precedence** `<name> | <num>` parameter of the `ip access-list` command specifies the IP precedence. The precedence option for of an IP packet is set in a three-bit field following the four-bit header-length field of the packet’s header. You can specify one of the following:

- **critical** or **5** – The ACL matches packets that have the critical precedence. If you specify the option number instead of the name, specify number 5.
- **flash** or **3** – The ACL matches packets that have the flash precedence. If you specify the option number instead of the name, specify number 3.
- **flash-override** or **4** – The ACL matches packets that have the flash override precedence. If you specify the option number instead of the name, specify number 4.
- **immediate** or **2** – The ACL matches packets that have the immediate precedence. If you specify the option number instead of the name, specify number 2.
- **internet** or **6** – The ACL matches packets that have the internetwork control precedence. If you specify the option number instead of the name, specify number 6.
- **network** or **7** – The ACL matches packets that have the network control precedence. If you specify the option number instead of the name, specify number 7.

- **priority** or **1** – The ACL matches packets that have the priority precedence. If you specify the option number instead of the name, specify number 1.
- **routine** or **0** – The ACL matches packets that have the routine precedence. If you specify the option number instead of the name, specify number 0.

The **tos** *<name>* | *<num>* parameter of the **ip access-list** command specifies the IP ToS. You can specify one of the following:

- **max-reliability** or **2** – The ACL matches packets that have the maximum reliability ToS. The decimal value for this option is 2.
- **max-throughput** or **4** – The ACL matches packets that have the maximum throughput ToS. The decimal value for this option is 4.
- **min-delay** or **8** – The ACL matches packets that have the minimum delay ToS. The decimal value for this option is 8.
- **min-monetary-cost** or **1** – The ACL matches packets that have the minimum monetary cost ToS. The decimal value for this option is 1.

NOTE

This value is not supported on 10 Gigabit Ethernet modules.

- **normal** or **0** – The ACL matches packets that have the normal ToS. The decimal value for this option is 0.
- *<num>* – A number from 0 – 15 that is the sum of the numeric values of the options you want. The ToS field is a four-bit field following the Precedence field in the IP header. You can specify one or more of the following. To select more than one option, enter the decimal value that is equivalent to the sum of the numeric values of all the ToS options you want to select. For example, to select the max-reliability and min-delay options, enter number 10. To select all options, select 15.
- The **dscp-cos-mapping** option maps the DSCP value in incoming packets to a hardware table that provides mapping of each of the 0 – 63 DSCP values, and distributes them among eight traffic classes (internal priorities) and eight 802.1p priorities.

NOTE

The **dscp-cos-mapping** option overrides port-based priority settings.

The **dscp-marking** option enables you to configure an ACL that marks matching packets with a specified DSCP value. Enter a value from 0 – 63. Refer to [“Using an IP ACL to mark DSCP values \(DSCP marking\)”](#) on page 387.

The **dscp-matching** option matches on the packet’s DSCP value. Enter a value from 0 – 63. This option does not change the packet’s forwarding priority through the device or mark the packet. Refer to [“DSCP matching”](#) on page 389.

The **log** parameter enables SNMP traps and Syslog messages for packets denied by the ACL:

- You can enable logging on ACLs and filters that support logging even when the ACLs and filters are already in use. To do so, re-enter the **ACL** or **filter** command and add the **log** parameter to the end of the ACL or filter. The software replaces the **ACL** or **filter** command with the new one. The new ACL or filter, with logging enabled, takes effect immediately.

The **traffic-policy** option enables the device to rate limit inbound traffic and to count the packets and bytes per packet to which ACL permit or deny clauses are applied. For configuration procedures and examples, refer to the chapter [“About traffic policies”](#) on page 427.

Configuration example for extended named ACLs

To configure an extended named ACL, enter commands such as the following.

```
PowerConnect(config)#ip access-list extended "block Telnet"
PowerConnect(config-ext-nACL)#deny tcp host 209.157.22.26 any eq telnet log
PowerConnect(config-ext-nACL)#permit ip any any
PowerConnect(config-ext-nACL)#exit
PowerConnect(config)#int eth 1
PowerConnect(config-if-1)#ip access-group "block Telnet" in
```

The options at the ACL configuration level and the syntax for the **ip access-group** command are the same for numbered and named ACLs and are described in ["Configuring extended numbered ACLs"](#) on page 368 and ["Configuring extended numbered ACLs"](#) on page 368.

Preserving user input for ACL TCP/UDP port numbers

ACL implementations automatically display the TCP/UDP port name instead of the port number, regardless of user preference. This feature preserves the user input (name or number) and now displays either the port name or the number.

A new command has been added to enable this feature.

```
PowerConnect(config)# ip preserve-ACL-user-input-format
```

Syntax: ip preserve-ACL-user-input-format

The following example shows how this feature works for a TCP port (this feature works the same way for UDP ports). In this example, the user identifies the TCP port by number (80) when configuring ACL group 140. However, **show ip access-list 140** reverts back to the port name for the TCP port (http in this example). After the user issues the new **ip preserve-ACL-user-input-format** command, **show ip access-list 140** displays either the TCP port number or name, depending on how it was configured by the user.

```
PowerConnect(config)# access-list 140 permit tcp any any eq 80
PowerConnect(config)# access-list 140 permit tcp any any eq ftp
PowerConnect# show ip access-lists 140
Extended IP access list 140
permit tcp any any eq http
permit tcp any any eq ftp
PowerConnect(config)# ip preserve-ACL-user-input-format
PowerConnect# show ip access-lists 140
Extended IP access list 140
permit tcp any any eq 80
permit tcp any any eq ftp
```

Managing ACL comment text

ACL comment text describes entries in an ACL. The comment text appears in the output of **show** commands that display ACL information.

This section describes how to add ACL comments.

Adding a comment to an entry in a numbered ACL

To add comments to entries in a numbered ACL, enter commands such as the following.

```
PowerConnect(config)#access-list 100 remark The following line permits TCP
packets
PowerConnect(config)#access-list 100 permit tcp 192.168.4.40/24 2.2.2.2/24
PowerConnect(config)#access-list 100 remark The following permits UDP packets
PowerConnect(config)#access-list 100 permit udp 192.168.2.52/24 2.2.2.2/24
PowerConnect(config)#access-list 100 deny ip any any
```

For example, using the same example configuration above, you could instead enter the following commands.

```
PowerConnect(config)#ip access-list extended 100
PowerConnect(config-ext-nACL)#remark The following line permits TCP packets
PowerConnect(config-ext-nACL)#permit tcp 192.168.4.40/24 2.2.2.2/24
PowerConnect(config-ext-nACL)#remark The following permits UDP packets
PowerConnect(config-ext-nACL)#permit udp 192.168.2.52/24 2.2.2.2/24
PowerConnect(config-ext-nACL)#deny ip any any
```

Syntax: [no] access-list <ACL-num> remark <comment-text>

or

Syntax: [no] ip access-list standard | extended <ACL-num>

Syntax: remark <comment-text>

For <ACL-num>, enter the number of the ACL.

The <comment-text> can be up to 128 characters in length. The comment must be entered separately from the actual ACL entry; that is, you cannot enter the ACL entry and the ACL comment with the same **access-list** or **ip access-list** command. Also, in order for the remark to be displayed correctly in the output of **show** commands, the comment must be entered immediately before the ACL entry it describes. Note that an ACL comment is tied to the ACL entry immediately following the comment. Therefore, if the ACL entry is removed, the ACL comment is also removed.

The **standard** | **extended** parameter indicates the ACL type.

Applying an ACL to a virtual interface in a protocol- or subnet-based VLAN

By default, when you apply an ACL to a virtual interface in a protocol-based or subnet-based VLAN, the ACL takes effect on all protocol or subnet VLANs to which the untagged port belongs. To prevent the device from denying packets on other virtual interfaces that do not have an ACL applied, configure an ACL that permits packets in the IP subnet of the virtual interface in all protocol-based or subnet-based VLANs to which the untagged port belongs. The following is an example configuration.

```
PowerConnect# conf t
PowerConnect(config)# vlan 1 name DEFAULT-VLAN by port
PowerConnect(config-vlan-1)# ip-subnet 192.168.10.0 255.255.255.0
PowerConnect(config-vlan-ip-subnet)# static ethe 1
PowerConnect(config-vlan-ip-subnet)# router-interface ve 10
PowerConnect(config-vlan-ip-subnet)# ip-subnet 10.15.1.0 255.255.255.0
PowerConnect(config-vlan-ip-subnet)# static ethe 1
```

```
PowerConnect(config-vlan-ip-subnet)# router-interface ve 20
PowerConnect(config-vlan-ip-subnet)# logging console
PowerConnect(config-vlan-ip-subnet)# exit
PowerConnect(config-vlan-1)# no vlan-dynamic-discovery
    Vlan dynamic discovery is disabled
PowerConnect(config-vlan-1)# int e 2
PowerConnect(config-if-e10000-2)# disable
PowerConnect(config-if-e10000-2)# interface ve 10
PowerConnect(config-vif-10)# ip address 192.168.10.254 255.255.255.0
PowerConnect(config-vif-10)# int ve 20
PowerConnect(config-vif-20)# ip access-group test1 in
PowerConnect(config-vif-20)# ip address 10.15.1.10 255.255.255.0
PowerConnect(config-vif-20)# exit
PowerConnect(config)# ip access-list extended test1
PowerConnect(config-ext-nACL)# permit ip 10.15.1.0 0.0.0.255 any log
PowerConnect(config-ext-nACL)# permit ip 192.168.10.0 0.0.0.255 any log
PowerConnect(config-ext-nACL)# end
PowerConnect#
```

Enabling ACL logging

You may want the software to log entries in the Syslog for packets that are denied by ACL filters. ACL logging is disabled by default; it must be explicitly enabled on a port.

When you enable logging for ACL entries, statistics for packets that match the deny conditions of the ACL entries are logged. For example, if you configure a standard ACL entry to deny all packets from source address 209.157.22.26, statistics for packets that are explicitly denied by the ACL entry are logged in the Syslog buffer and in SNMP traps sent by the device.

The first time an ACL entry denies a packet, the software immediately generates a Syslog entry and an SNMP trap. The software also starts a five-minute timer. The timer keeps track of all packets explicitly denied by the ACL entries. After five minutes, the software generates a single Syslog entry for each ACL entry that denied a packet. The Syslog entry (message) indicates the number of packets denied by the ACL entry during the previous five minutes. Note however that packet count may be inaccurate if the packet rate is high and exceeds the CPU processing rate.

If no ACL entries explicitly deny packets during an entire five-minute timer interval, the timer stops. The timer restarts when an ACL entry explicitly denies a packet.

NOTE

The timer for logging packets denied by Layer 2 filters is a different timer than the ACL logging timer.

Configuration notes

Note the following before configuring ACL logging:

- You can enable ACL logging on physical and virtual interfaces.
- ACL logging logs denied packets only.
- When ACL logging is disabled, packets that match the ACL rule are forwarded or dropped in hardware. When ACL logging is enabled, all packets that match the ACL deny rule are sent to the CPU. When ACL logging is enabled, Dell recommends that you configure a traffic conditioner, then link the ACL to the traffic conditioner to prevent CPU overload. For example:

13 Enabling ACL logging

```
PowerConnect(config)# traffic-policy TPD1 rate-limit fixed 100 exceed-action drop
PowerConnect(config)# access-list 101 deny ip host 210.10.12.2 any
traffic-policy TPD1 log
```

- ACL logging is intended for debugging purpose. Dell recommends that you disable ACL logging after the debug session is over.

Configuration Tasks

To enable ACL logging, complete the following steps:

1. Create ACL entries with the log option
2. Enable ACL logging on individual ports

NOTE

The command syntax for enabling ACL logging is different on IPv4 devices than on IPv6 devices. See the configuration examples in the next section.

3. Bind the ACLs to the ports on which ACL logging is enabled

Example Configuration

The following shows an example configuration on an IPv4 device.

```
PowerConnect(config)# access-list 1 deny host 209.157.22.26 log
PowerConnect(config)# access-list 1 deny 209.157.29.12 log
PowerConnect(config)# access-list 1 deny host IPHost1 log
PowerConnect(config)# access-list 1 permit any
PowerConnect(config)# interface e 4
PowerConnect(config-if-e10000-4)# ACL-logging
PowerConnect(config-if-e10000-4)# ip access-group 1 in
```

The above commands create ACL entries that include the log option, enable ACL logging on interface e 4, then bind the ACL to interface e 4. Statistics for packets that match the deny statements will be logged.

Syntax: ACL-logging

The **ACL-logging** command applies to IPv4 devices only.

Displaying ACL Log Entries

The first time an entry in an ACL permits or denies a packet and logging is enabled for that entry, the software generates a Syslog message and an SNMP trap. Messages for packets permitted or denied by ACLs are at the warning level of the Syslog.

When the first Syslog entry for a packet permitted or denied by an ACL is generated, the software starts an ACL timer. After this, the software sends Syslog messages every five minutes. If an ACL entry does not permit or deny any packets during the timer interval, the software does not generate a Syslog entry for that ACL entry.

NOTE

For an ACL entry to be eligible to generate a Syslog entry for denied packets, logging must be enabled for the entry. The Syslog contains entries only for the ACL entries that deny packets and have logging enabled.

To display Syslog entries, enter the following command from any CLI prompt:

```
PowerConnect#show log
Syslog logging: enabled (0 messages dropped, 2 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 9 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Dynamic Log Buffer (50 lines):
0d00h12m18s:W:ACL: ACL: List 122 denied tcp 20.20.15.6(0)(Ethernet 4 0000.0804.01
20.20.18.6(0), 1 event(s)
0d00h12m18s:W:ACL: ACL: List 122 denied tcp 20.20.15.2(0)(Ethernet 4 0000.0804.01
20.20.18.2(0), 1 event(s)
0d00h12m18s:W:ACL: ACL: List 122 denied tcp 20.20.15.4(0)(Ethernet 4 0000.0804.01
20.20.18.4(0), 1 event(s)
0d00h12m18s:W:ACL: ACL: List 122 denied tcp 20.20.15.3(0)(Ethernet 4 0000.0804.01
20.20.18.3(0), 1 event(s)
0d00h12m18s:W:ACL: ACL: List 122 denied tcp 20.20.15.5(0)(Ethernet 4 0000.0804.01
20.20.18.5(0), 1 event(s)
0d00h12m18s:I:ACL: 122 applied to port 4 by  from console session
0d00h10m12s:I:ACL: 122 removed from port 4 by  from console session
0d00h09m56s:I:ACL: 122 removed from port 4 by  from console session
0d00h09m38s:I:ACL: 122 removed from port 4 by  from console session
```

Syntax: show log

Enabling strict control of ACL filtering of fragmented packets

The default processing of fragments by hardware-based ACLs is as follows:

- The first fragment of a packet is permitted or denied using the ACLs. The first fragment is handled the same way as non-fragmented packets, since the first fragment contains the Layer 4 source and destination application port numbers. The device uses the Layer 4 CAM entry if one is programmed, or applies the interface's ACL entries to the packet and permits or denies the packet according to the first matching ACL.
- For other fragments of the same packet, they are subject to a rule only if there is no Layer 4 information in the rule or in any preceding rules.

The fragments are forwarded even if the first fragment, which contains the Layer 4 information, was denied. Generally, denying the first fragment of a packet is sufficient, since a transaction cannot be completed without the entire packet.

For tighter control, you can configure the port to drop all packet fragments. To do so, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# ip access-group frag deny
```

This option begins dropping all fragments received by the port as soon as you enter the command. This option is especially useful if the port is receiving an unusually high rate of fragments, which can indicate a hacker attack.

Syntax: [no] ip access-group frag deny

Enabling ACL support for switched traffic in the router image

By default, when an ACL is applied to a physical or virtual routing interface, the Layer 3 device filters routed traffic only. It does not filter traffic that is switched from one port to another within the same VLAN or virtual routing interface, even if an ACL is applied to the interface.

You can enable the device to filter switched traffic within a VLAN or virtual routing interface. When filtering is enabled, the device uses the ACLs applied to inbound traffic to filter traffic received by a port from another port in the same virtual routing interface.

To enable this feature, enter a command such as the following.

```
PowerConnect(config)# access-list 101 bridged-routed
```

Applying the ACL rule above to an interface, enables filtering of traffic switched within a VLAN or virtual routing interface.

Syntax: [no] ip access-list <ACL-ID> bridged-routed

The <ACL-ID> parameter specifies a standard or extended numbered or named ACL.

You can use this feature in conjunction with **enable ACL-per-port-per-vlan**, to assign an ACL to a single port within a virtual interface. In this case, all of the Layer 3 traffic (bridged and routed) are filtered by the ACL.

```
PowerConnect(config)# enable ACL-per-port-per-vlan
PowerConnect(config)# write memory
PowerConnect(config)# exit
PowerConnect# reload
```

Enabling ACL filtering based on VLAN membership or VE port membership

NOTE

This section applies to IPv4 ACLs only. IPv6 ACLs do not support ACL filtering based on VLAN membership or VE port membership.

You can apply an inbound IPv4 ACL to specific VLAN members on a port (Layer 2 devices only) or to specific ports on a virtual interface (VE) (Layer 3 Devices only).

By default, this feature support is disabled. To enable it, enter the following commands at the Global CONFIG level of the CLI.

```
PowerConnect (config)# enable ACL-per-port-per-vlan
PowerConnect (config)# write memory
PowerConnect (config)# exit
PowerConnect# reload
```

After entering the above commands, you can do the following:

- Apply an IPv4 ACL to specific VLAN members on a port – refer to [“Applying an IPv4 ACL to specific VLAN members on a port \(Layer 2 devices only\)”](#) on page 385
- Apply an IPv4 ACL to a subset of ports on a VE – refer to [“Applying an IPv4 ACL to a subset of ports on a virtual interface \(Layer 3 devices only\)”](#) on page 385

Syntax: [no] enable ACL-per-port-per-vlan

Enter the no form of the command to disable this feature.

Applying an IPv4 ACL to specific VLAN members on a port (Layer 2 devices only)

When you bind an IPv4 ACL to a port, the port filters all inbound traffic on the port. However, on a tagged port, there may be a need to treat packets for one VLAN differently from packets for another VLAN. You can configure a tagged port on a Layer 2 device to filter packets based on the packets' VLAN membership.

NOTE

Before you can bind an ACL to specific VLAN members on a port, you must first enable support for this feature. If this feature is not already enabled on your device, enable it as instructed in the section [“Enabling ACL filtering based on VLAN membership or VE port membership”](#) on page 384.

To apply an IPv4 ACL to a specific VLAN on a port, enter commands such as the following on a tagged port.

```
PowerConnect(config)# vlan 12 name vlan12
PowerConnect(config-vlan-12)# untag ethernet 5 to 8
PowerConnect(config-vlan-12)# tag ethernet 23 to 24
PowerConnect(config-vlan-12)# exit
PowerConnect(config)# access-list 10 deny host 209.157.22.26 log
PowerConnect(config)# access-list 10 deny 209.157.29.12 log
PowerConnect(config)# access-list 10 deny host IPHost1 log
PowerConnect(config)# access-list 10 permit
PowerConnect(config)# int e 23
PowerConnect(config-if-e10000-23)# per-vlan 12
PowerConnect(config-if-e10000-23-vlan-12)# ip access-group 10 in
```

The commands in this example configure port-based VLAN 12, and add ports e 5 – 8 as untagged ports and ports e 23 – 24 as tagged ports to the VLAN. The commands following the VLAN configuration commands configure ACL 10. Finally, the last three commands apply ACL 10 on VLAN 12 for which port e 23 is a member.

Syntax: `per-vlan <VLAN ID>`

Syntax: `[no] ip access-group <ACL ID>`

The `<VLAN ID>` parameter specifies the VLAN name or number to which you will bind the ACL.

The `<ACL ID>` parameter is the access list name or number.

Applying an IPv4 ACL to a subset of ports on a virtual interface (Layer 3 devices only)

You can apply an IPv4 ACL to a virtual routing interface. The virtual interface is used for routing between VLANs and contains all the ports within the VLAN. The IPv4 ACL applies to all the ports on the virtual routing interface. You also can specify a subset of ports within the VLAN containing a specified virtual interface when assigning an ACL to that virtual interface.

Use this feature when you do not want the IPv4 ACLs to apply to all the ports in the virtual interface VLAN or when you want to streamline IPv4 ACL performance for the VLAN.

NOTE

Before you can bind an IPv4 ACL to specific ports on a virtual interface, you must first enable support for this feature. If this feature is not already enabled on your device, enable it as instructed in the section [“Enabling ACL filtering based on VLAN membership or VE port membership”](#) on page 384.

To apply an ACL to a subset of ports within a virtual interface, enter commands such as the following.

```
PowerConnect(config)# vlan 10 name IP-subnet-vlan
PowerConnect(config-vlan-10)# untag ethernet 1 to 12
PowerConnect(config-vlan-10)# router-interface ve 1
PowerConnect(config-vlan-10)# exit
PowerConnect(config)# access-list 1 deny host 209.157.22.26 log
PowerConnect(config)# access-list 1 deny 209.157.29.12 log
PowerConnect(config)# access-list 1 deny host IPHost1 log
PowerConnect(config)# access-list 1 permit any
PowerConnect(config)# interface ve 1
PowerConnect(config-vif-1)# ip access-group 1 in ethernet 1 ethernet 3 ethernet 4
to 5
```

The commands in this example configure port-based VLAN 10, add ports 1 – 12 to the VLAN, and add virtual routing interface 1 to the VLAN. The commands following the VLAN configuration commands configure ACL 1. Finally, the last two commands apply ACL 1 to a subset of the ports associated with virtual interface 1.

Syntax: [no] ip access-group <ACL ID> in ethernet <portnum> [to<portnum>]

The <ACL ID> parameter is the access list name or number.

Filtering on IP precedence and ToS values

To configure an extended IP ACL that matches based on IP precedence, enter commands such as the following.

```
PowerConnect(config)#access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24
precedence internet
PowerConnect(config)#access-list 103 deny tcp 209.157.21.0/24 eq ftp
209.157.22.0/24 precedence 6
PowerConnect(config)#access-list 103 permit ip any any
```

The first entry in this ACL denies TCP traffic from the 209.157.21.x network to the 209.157.22.x network, if the traffic has the IP precedence option “internet” (equivalent to “6”).

The second entry denies all FTP traffic from the 209.157.21.x network to the 209.157.22.x network, if the traffic has the IP precedence value “6” (equivalent to “internet”).

The third entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

To configure an IP ACL that matches based on ToS, enter commands such as the following.


```
PowerConnect(config)#access-list 104 deny tcp 209.157.21.0/24 209.157.22.0/24
tos normal
PowerConnect(config)#access-list 104 deny tcp 209.157.21.0/24 eq ftp
209.157.22.0/24 tos 13
PowerConnect(config)#access-list 104 permit ip any any
```

The first entry in this IP ACL denies TCP traffic from the 209.157.21.x network to the 209.157.22.x network, if the traffic has the IP ToS option “normal” (equivalent to “0”).

The second entry denies all FTP traffic from the 209.157.21.x network to the 209.157.22.x network, if the traffic has the IP ToS value “13” (equivalent to “max-throughput”, “min-delay”, and “min-monetary-cost”).

The third entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

QoS options for IP ACLs

Quality of Service (QoS) options enable you to perform QoS for packets that match the ACLs. Using an ACL to perform QoS is an alternative to directly setting the internal forwarding priority based on incoming port, VLAN membership, and so on. (This method is described in [“Assigning QoS priorities to traffic”](#) on page 408.)

The following QoS ACL options are supported:

- **dscp-marking** – Marks the DSCP value in the outgoing packet with the value you specify.
- **802.1p-and internal-marking** – Supported on PowerConnect devices only with the DSCP marking option, this command assigns traffic that matches the ACL to a hardware forwarding queue and re-marks the packets that match the ACL with the 802.1p priority.
- **dscp-matching** – Matches on the packet DSCP value. This option does not change the packet forwarding priority through the device or mark the packet.

Using an IP ACL to mark DSCP values (DSCP marking)

The **dscp-marking** option for extended ACLs allows you to configure an ACL that marks matching packets with a specified DSCP value. You also can use DSCP marking to assign traffic to a specific hardware forwarding queue (refer to [“Using an ACL to change the forwarding queue for PowerConnect B-Series TI24X devices”](#) on page 388).

For example, the following commands configure an ACL that marks all IP packets with DSCP value 5. The ACL is then applied to incoming packets on interface 7. Consequently, all inbound packets on interface 7 are marked with the specified DSCP value.

```
PowerConnect(config)# access-list 120 permit ip any any dscp-marking 5
dscp-cos-mapping
PowerConnect(config)# interface 7
PowerConnect(config-if-e10000-7)# ip access-group 120 in
```

Syntax: ...**dscp-marking** <dscp-value>

The **dscp-marking** <dscp-value> parameter maps a DSCP value to an internal forwarding priority. The DSCP value can be from 0 – 63.

Combined ACL for 802.1p marking on PowerConnect B-Series TI24X**NOTE**

For feature support on PowerConnect B-Series TI24X devices, refer to [“Combined ACL for 802.1p marking on PowerConnect B-Series TI24X devices”](#) on page 388.

Using an ACL to change the forwarding queue for PowerConnect B-Series TI24X devices**NOTE**

For feature support on PowerConnect B-Series TI24X devices, refer to [“Using an ACL to change the forwarding queue for PowerConnect B-Series TI24X devices”](#) on page 388.

Combined ACL for 802.1p marking on PowerConnect B-Series TI24X devices

PowerConnect B-Series TI24X devices support a simple method for assigning an 802.1p priority and internal marking priority to packets without affecting the actual packet or the DSCP marking. PowerConnect B-Series TI24X devices use the same configured value for both the internal marking priority and the 802.1p priority marking value. This feature is enabled through the use of an ACL option. The option applies to IP, TCP, and UDP traffic.

For IP

```
PowerConnect(config)# acc 104 per ip any any 802.1p-and-internal-marking 1
```

Syntax: `access-list <num(100-199)> permit ip any any 802.1p-and-internal-marking <priority value (0-7)>`

For TCP

```
PowerConnect(config)# acc 105 per tcp any any 802.1p-and-internal-marking 1
```

Syntax: `access-list <num(100-199)> permit tcp any any 802.1p-and-internal-marking <priority value (0-7)>`

For UDP

```
PowerConnect(config)# acc 105 per udp any any 802.1p-and-internal-marking 1
```

Syntax: `access-list <num(100-199)> permit udp any any 802.1p-and-internal-marking <priority value (0-7)>`

In each of these examples, the 802.1p value and internal priority value were set to 1.

Using an ACL to change the forwarding queue for PowerConnect B-Series TI24X devices**NOTE**

This section applies to PowerConnect B-Series TI24X devices.

The **802.1p-and-internal-marking <0 – 7>** parameter does the following:

- Re-marks the packets of the 802.1Q traffic that match the ACL with this new priority, or marks the packets of the non-802.1Q traffic that match the ACL with this priority, later at the outgoing 802.1Q interface.

- Assigns traffic that matches the ACL to the specific hardware forwarding queue (qosp0 – qosp7>.

NOTE

The **802.1p-and-internal-marking** option overrides port-based priority settings.

In addition to changing the internal forwarding priority, if the outgoing interface is an 802.1Q interface, this parameter maps the specified priority to its equivalent 802.1p (CoS) priority and marks the packet with the new 802.1p priority. [Table 55](#) lists the default mappings of hardware forwarding queues to 802.1p priorities.

TABLE 55 Default mapping of forwarding queues to 802.1p priorities

Forwarding Queue	qosp0	qosp1	qosp2	qosp3	qosp4	qosp5	qosp6	qosp7
802.1p	0	1	2	3	4	5	6	7

The complete CLI syntax for 802.1p priority marking and internal priority marking is shown in [“Configuring extended numbered ACLs”](#) on page 368 and [“Configuring extended named ACLs”](#) on page 374. Also refer to [“Combined ACL for 802.1p marking on PowerConnect B-Series TI24X devices”](#) on page 388. The following shows the syntax specific to this feature on PowerConnect B-Series TI24X devices.

Syntax: ... **802.1p-and-internal-marking** <0 – 7>

DSCP matching

The **dscp-matching** option matches on the packet DSCP value. This option does not change the packet forwarding priority through the device or mark the packet.

To configure an ACL that matches on a packet with DSCP value 29, enter a command such as the following.

```
PowerConnect(config)# access-list 112 permit ip 1.1.1.0 0.0.0.255 2.2.2.x
0.0.0.255 dscp-matching 29
```

The complete CLI syntax for this feature is shown in [“Configuring extended numbered ACLs”](#) on page 368 and [“Configuring extended named ACLs”](#) on page 374. The following shows the syntax specific to this feature.

Syntax: ...**dscp-matching** <0 – 63>

NOTE

For complete syntax information, refer to [“Extended numbered ACL syntax”](#) on page 369.

ACL-based rate limiting

ACL-based rate limiting provides the facility to limit the rate for IP traffic that matches the permit conditions in extended IP ACLs. This feature is available in the Layer 2 and Layer 3 code.

For more details, including configuration procedures, refer to [Chapter 17, “Configuring Traffic Policies”](#).

Using ACLs to control multicast features

You can use ACLs to control the following multicast features:

- Limit the number of multicast groups that are covered by a static rendezvous point (RP)
- Control which multicast groups for which candidate RPs sends advertisement messages to bootstrap routers
- Identify which multicast group packets will be forwarded or blocked on an interface

For configuration procedures, refer to [Chapter 19, “Configuring IP Multicast Protocols”](#).

Enabling and viewing hardware usage statistics for an ACL

The number of configured ACL rules can affect the rate at which hardware resources are used. You can use the **show access-list hw-usage on** command to enable hardware usage statistics, followed by the **show access-list <access-list-id>** command to determine the hardware usage for an ACL. To gain more hardware resources, you can modify the ACL rules so that it uses less hardware resource.

NOTES:

- When an ACL is not attached to any port, the **show access-list** command displays an estimate of the current TCAM usage, assuming it is attached to one port and one VLAN. Once the ACL is attached to a port, the **show access-list <access-list-id>** command shows the exact current TCAM usage by the ACL.
- If the ACL contains filters with Layer 4 source or destination port ranges and the ACL is not attached to any port or VLAN, then the minimum and maximum number of estimated TCAM usage per filter is displayed in 'x or y' format where 'x' is the minimum number and 'y' is the maximum number of estimated TCAM entries.
- Whenever the ACL is attached to a different VLAN (on the same or another port), the TCAM usage count is incremented to reflect the current usage. The following shows an example of the **show access-list** command output before an ACL is attached to a port.

```
PowerConnect(config-if-e10000-2-vlan-2)#show acc 111
Extended IP access list 111 (hw usage : 3 or up to 13)
permit tcp any range 10 40 any (hw usage : 1 or 5)
permit tcp any range 10 60 any (hw usage : 1 or 7)
```

The following shows an example **show access-list** command output after an ACL is attached to a port.

```
PowerConnect(config-if-e10000-2)#show access-list 111
Extended IP access list 111 (hw usage : 3)
permit tcp any range 10 40 any (hw usage : 1)
permit tcp any range 10 60 any (hw usage : 1)
```

Displaying ACL information

To display the number of Layer 4 CAM entries used by each ACL, enter the following command.

```
PowerConnect#show access-list all

Extended IP access list 100 (Total flows: N/A, Total packets: N/A, Total rule cam
use: 3)
permit udp host 192.168.2.169 any (Flows: N/A, Packets: N/A, Rule cam use: 1)
permit icmp any any (Flows: N/A, Packets: N/A, Rule cam use: 1)
deny ip any any (Flows: N/A, Packets: N/A, Rule cam use: 1)
```

Syntax: **show access-list <ACL-num> | <ACL-name> | all**

The Rule cam use field lists the number of CAM entries used by the ACL or entry. The number of CAM entries listed for the ACL itself is the total of the CAM entries used by the ACL entries.

For flow-based ACLs, the Total flows and Flows fields list the number of Layer 4 session table flows in use for the ACL.

Enabling and viewing hardware usage statistics for an ACL

The number of configured ACL rules can affect the rate at which hardware resources are used. You can use the **show access-list hw-usage on** command to enable hardware usage statistics, followed by the **show access-list <access-list-id>** command to determine the hardware usage for an ACL. To gain more hardware resources, you can modify the ACL rules so that it uses less hardware resource.

To enable and view hardware usage statistics, enter commands such as the following:

```
PowerConnect# show access-list hw-usage on
PowerConnect# show access-list 100
Extended IP access list 100 (hw usage : 2)
deny ip any any (hw usage : 1
```

Syntax: show access-list hw-usage on

Syntax: show access-list <access-list-id>

where <access-list-id> is a valid ACL name or number.

Displaying ACL information

To display the number of Layer 4 CAM entries used by each ACL, enter the following command.

```
PowerConnect#show access-list all

Extended IP access list 100 (Total flows: N/A, Total packets: N/A, Total rule cam
use: 3)
permit udp host 192.168.2.169 any (Flows: N/A, Packets: N/A, Rule cam use: 1)
permit icmp any any (Flows: N/A, Packets: N/A, Rule cam use: 1)
deny ip any any (Flows: N/A, Packets: N/A, Rule cam use: 1)
```

Syntax: show access-list <ACL-num> | <ACL-name> | all

The Rule cam use field lists the number of CAM entries used by the ACL or entry. The number of CAM entries listed for the ACL itself is the total of the CAM entries used by the ACL entries.

For flow-based ACLs, the Total flows and Flows fields list the number of Layer 4 session table flows in use for the ACL.

The Total packets and Packets fields apply only to flow-based ACLs.

Troubleshooting ACLs

Use the following methods to troubleshoot ACLs:

- To display the number of Layer 4 CAM entries being used by each ACL, enter the **show access-list <ACL-num> | <ACL-name> | all** command. Refer to [“Displaying ACL information”](#) on page 392.

- To determine whether the issue is specific to fragmentation, remove the Layer 4 information (TCP or UDP application ports) from the ACL, then reapply the ACL.

If you are using another feature that requires ACLs, either use the same ACL entries for filtering and for the other feature, or change to flow-based ACLs

13 Troubleshooting ACLs

Configuring Port Mirroring and Monitoring

Mirroring support by platform

The procedures in this chapter describe how to configure port mirroring on devices.

Port mirroring is a method of monitoring network traffic that forwards a copy of each incoming or outgoing packet from one port on a network switch to another port where the packet can be analyzed. Port mirroring may be used as a diagnostic tool or debugging feature, especially for preventing attacks. Port mirroring can be managed locally or remotely.

Configure port mirroring by assigning a port from which to copy all packets, and a “mirror” port where the copies of the packets are sent (also known as the monitor port). A packet received on, or issued from, the first port is forwarded to the second port as well. Attach a protocol analyzer on the mirror port to monitor each segment separately. The analyzer captures and evaluates the data without affecting the client on the original port.

The mirror port may be a port on the same switch with an attached RMON probe, a port on a different switch in the same hub, or the switch processor.

[Table 56](#) lists which devices support the mirroring features discussed in this chapter.

TABLE 56 Mirroring support fo PowerConnect B-Series TI24X devices

Mirroring Feature	PowerConnect B-Series TI24X Series
Basic Port Mirroring and Monitoring	X
ACL-Based Inbound Mirroring	X
MAC Filter-Based Mirroring	X

Configuring port mirroring and monitoring

PowerConnect B-Series TI24X devices support one mirror port per device. To configure port monitoring, first specify the **mirror port**, then enable monitoring on the **monitored port**.

The **mirror port** is the port to which the monitored traffic is copied. Attach your protocol analyzer to the mirror port. The **monitored port** is the port whose traffic you want to monitor.

Configuration notes

Refer to the following rules when configuring port mirroring and monitoring:

- Port monitoring and sFlow support:
 - PowerConnect B-Series TI24X devices support sFlow and port monitoring together on the same port.

- If you configure both ACL mirroring and ACL based rate limiting on the same port, then all packets that match are mirrored, including the packets that exceed the rate limit.
- [Table 57](#) lists the number of mirror and monitor ports supported on the devices.

TABLE 57 Number of mirror and monitored ports supported

Port type	Maximum ports supported
	PowerConnect B-Series T124X
Ingress mirror ports	1
Egress mirror ports	1
Ingress monitored ports	no limit
Egress monitored ports	no limit

- You can configure a mirror port specifically as an ingress port, an egress port, or both.
- Mirror ports can run at any speed and are not related to the speed of the ingress or egress monitored ports.
- The same port cannot be both a monitored port and the mirror port.
- The same port can be monitored by one mirror port for ingress traffic and another mirror port for egress traffic.
- The mirror port cannot be a trunk port.
- The monitored port and its mirror port do not need to belong to the same port-based VLAN:
 - If the mirror port is in a *different* VLAN from the monitored port, the packets are tagged with the monitor port VLAN ID.
 - If the mirror port is in the *same* VLAN as the monitored port, the packets are tagged or untagged, depending on the mirror port configuration.
- More than one monitored port can be assigned to the same mirror port.
- If the primary interface of a trunk is enabled for monitoring, the entire trunk will be monitored. You can also enable an individual trunk port for monitoring using the **config-trunk-ind** command.
- For ingress ACL mirroring, the previous ingress rule also applies. The analyzer port setting command **acl-mirror-port** must be specified for each port, even though the hardware only supports one port per device. This applies whether the analyzer port is on the local device or on a remote device. For example, when port mirroring is set to a remote device, any mirroring (ACL, MAC filter, or VLAN) enabled ports are globally set to a single analyzer port, as shown in the following example.

```
PowerConnect(config)# mirror ethernet 24
PowerConnect(config)# mirror ethernet 48
PowerConnect(config)# interface ethernet 1
PowerConnect(interface ethernet 1)# monitor ethernet 48 both
```

The analyzer port (48) is set to all devices in the system

```
PowerConnect# interface ethernet 2
PowerConnect(interface ethernet 2)# ip access-group 101 in
PowerConnect# exit
PowerConnect# interface ethernet 1
PowerConnect(interface ethernet 1)# acl-mirror-port ethernet 48
```

The previous command is required even though the analyzer port is already set globally by the port mirroring command.

```
PowerConnect#interface ethernet 3
PowerConnect(interface ethernet 3)# ip access-group 101
PowerConnect(interface ethernet 3)# acl-mirror-port ethernet 48
PowerConnect(interface ethernet 3)# permit ip any any mirror
PowerConnect(interface ethernet 3)# ip access-group 102
PowerConnect(interface ethernet 3)# deny ip any any log
```

Monitoring a port

This section describes how to configure port mirroring and monitoring.

To configure port monitoring on an individual port on a device, enter commands similar to the following.

```
PowerConnect(config)#mirror-port ethernet 4
PowerConnect(config)#interface ethernet 11
PowerConnect(config-if-e10000-11)#monitor ethernet 4 both
```

Traffic on port e 11 will be monitored, and the monitored traffic will be copied to port e 4, the mirror port.

Syntax: [no] mirror-port ethernet <portnum> [input | output]

Syntax: [no] monitor ethernet<portnum> both | in | out

The <portnum> parameter for **mirror-port ethernet** specifies the port to which the monitored traffic will be copied. The <portnum> parameter for **monitor ethernet** specifies the port on which traffic will be monitored.

The **input** and **output** parameters configure the mirror port exclusively for ingress or egress traffic. If you do not specify one, both types of traffic apply.

The **both**, **in**, and **out** parameters specify the traffic direction you want to monitor on the mirror port. There is no default.

To display the port monitoring configuration, enter the **show monitor** and **show mirror** commands.

Monitoring an individual trunk port

You can monitor the traffic on an individual port of a static trunk group. Note that monitoring an LACP trunk port is not supported on PowerConnect B-Series TI24X devices.

By default, when you monitor the primary port in a trunk group, aggregated traffic for all the ports in the trunk group is copied to the mirror port. You can configure the device to monitor individual ports in a trunk group. You can monitor the primary port or a secondary port individually.

To configure port monitoring on an individual port in a trunk group, enter commands such as the following.

```
PowerConnect(config)#mirror-port ethernet 6
PowerConnect(config)#trunk e 2 to 5
PowerConnect(config-trunk-2-5)#config-trunk-ind
PowerConnect(config-trunk-2-5)#monitor ethernet 4 ethernet 6 in
```

Traffic on trunk port e 4 will be monitored, and the monitored traffic will be copied to port e 6, the mirror port.

Syntax: [no] mirror-port ethernet [<portnum>] [input | output]

Syntax: [no] config-trunk-ind

Syntax: [no] monitor ethernet <portnum> both | in | out

The <portnum> parameter for **mirror-port ethernet** specifies the port to which the monitored traffic will be copied. The <portnum> parameter for **monitor ethernet** specifies the port on which traffic will be monitored.

The **input** or **output** parameters configure the mirror port exclusively for ingress or egress traffic. If you do not specify one, both types of traffic apply.

The **config-trunk-ind** command enables configuration of individual ports in the trunk group. You enter the **config-trunk-ind** command only once in a trunk group. After you enter the command, all applicable port configuration commands apply to individual ports only.

NOTE

If you enter **no config-trunk-ind**, all port configuration commands are removed from the individual ports and the configuration of the primary port is applied to all the ports. Also, once you enter the **no config-trunk-ind** command, the enable, disable, and monitor commands are valid only on the primary port and apply to the entire trunk group.

The **both**, **in**, and **out** parameters specify the traffic direction you want to monitor on the mirror port. There is no default.

To display the port monitoring configuration, enter the **show monitor** and **show mirror** commands

ACL-based inbound mirroring

This section describes how to configure ACL-based inbound mirroring for PowerConnect devices.

Creating an ACL-based inbound mirror clause for PowerConnect B-Series TI24X devices

The following example shows how to configure an ACL-based inbound mirror clause for PowerConnect B-Series TI24X devices.

1. Configure the mirror port.

```
PowerConnect(config)#mirror-port ethernet 2
```

2. Configure the ACL inbound mirror clause.

```
PowerConnect(config)#access-list 101 permit ip any any mirror
```

At this point not all IP traffic will be mirrored to port 2, since the ACL has not yet been applied to any port.

3. Apply the ACL inbound clause to the monitor port.

```
PowerConnect(config)#int e 5
PowerConnect(config-if-e10000-5)#ip access-group 101 in
```

4. Configure the monitor port to use the mirror port.

```
PowerConnect(config-if-e10000-5)#acl-mirror-port ethernet 2
```

To display ACL mirror settings, enter the **show access-list all** command.

```
PowerConnect#show access-list all
Extended IP access list 101
permit ip any any mirror
```

Specifying the destination mirror port

You can specify physical ports or a trunk to mirror traffic from. If you complete the rest of the configuration but do not specify a destination mirror port, the port-mirroring ACL will be non-operational. This can be useful if you want to be able to mirror traffic by a set criteria on-demand. With this configuration, you just configure a destination mirror port whenever you want the port-mirroring ACL to become operational.

The following sections describe how to specify a destination port for a port or a trunk as well as the special considerations required when mirroring traffic from a virtual interface.

Specifying the destination mirror port for physical ports

When you want traffic that has been selected by ACL-based Inbound Mirroring to be mirrored, you must configure a destination mirror port. This configuration is performed at the Interface Configuration of the port whose traffic you are mirroring. The destination port must be the same for all ports in a port region as described in [“Ports from a port region must be mirrored to the same destination mirror port”](#) on page 399.

In the following example, ACL mirroring traffic from port 1 is mirrored to port 3.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 3
```

Syntax: [no] ACL-mirror-port ethernet<portnum>

The <portnum> variable specifies port to which ACL-mirror traffic from the configured interface will be mirrored.

The <portnum> parameter specifies the mirror port to which the monitored port traffic will be copied.

Ports from a port region must be mirrored to the same destination mirror port

Port regions as described in [“Enabling or disabling the Spanning Tree Protocol \(STP\)”](#) on page 175 are important when defining a destination mirror port. This is because all traffic mirrored from any single port in a port region will be mirrored to the same destination mirror port as traffic mirrored from any other port in the same port region. For example, ports 1 to 12 are in the same port region. If you configure ports 1 and 2 to mirror their traffic, they should use the same destination mirror port as shown in the following configuration.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 3
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-e10000-2)#ACL-mirror-port ethernet 3
```

If ports within the same port region are mirrored to different destination ports, an error message will be generated as shown in the following example, and the configuration will be disallowed.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 3
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-e10000-2)#ACL-mirror-port ethernet 7
Error - Inbound Mirror port 3 already configured for port region 1 - 12
```

When a destination port is configured for any port within a port region, traffic from any ACL with a mirroring clause assigned to any port in that port region will be mirrored to that destination port. This will occur even if a destination port is not explicitly configured for the port with the ACL configured. In the following example, an ACL with a mirroring clause (101) is applied to a port (1). Another port in the same region (2) has a destination port set (3). In this example, traffic generated from operation of ACL 101 is mirrored to port 3 even though a destination port has not explicitly been defined for traffic from port 1.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ip access-group 101 in
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-e10000-2)#ACL-mirror-port ethernet 3
```

NOTE

If a destination mirror port is not configured for any ports within the port region where the port-mirroring ACL is configured, the ACL will not mirror the traffic but the ACL will be applied to traffic on the port.

Specifying the destination mirror port for trunk ports

You can mirror the traffic that has been selected by ACL-based Inbound Mirroring from a trunk by configuring a destination port for the primary port within the trunk configuration as shown.

```
PowerConnect(config)#trunk ethernet 1 to 4
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 8
```

Using this configuration, all trunk traffic will be mirrored to port 8.

Limitations when configuring ACL-based mirroring with trunks

The **config-trunk-ind** option as described in [“Disabling or re-enabling a trunk port”](#) on page 319 cannot operate with ACL-Based Mirroring as described in the following:

- If a trunk is configured with the **config-trunk-ind** option, ACL-Based Mirroring will not be allowed.
- If the **config-trunk-ind** option is added to a trunk, any ports that are configured for ACL-based Mirroring will have monitoring removed and the following message will be displayed.

Trunk port monitoring, if any, has been removed.

If an individual port is configured for ACL-Based Mirroring, you cannot add it to a trunk. If you try to add a port that is configured for ACL-Based Mirroring to a trunk, the following message appears.

Note - ACL-mirror-port configuration is removed from port 2 in new trunk.

NOTE

If you want to add a port configured for ACL-Based Mirroring to a trunk, you must first remove the **ACL-mirror-port** from the port configuration. You can then add the port to a trunk that can then be configured for ACL-Based Trunk Mirroring.

Behavior of ACL-based mirroring when deleting trunks

If you delete a trunk that has ACL-Based Mirroring configured, the ACL-Based Mirroring configuration will be configured on the individual ports that made up the trunk.

For example, if a trunk is configured as shown in the following example and is then deleted from the configuration as shown, each of the ports that previously were contained in the trunk will be configured for ACL-Based Mirroring.

```
PowerConnect(config)#trunk ethernet 1 to 2
PowerConnect(config)#trunk deploy
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000)#ACL-mirror-port ethernet 3
```

To delete the trunk, enter the following command.

```
PowerConnect(config)#no trunk ethernet 1 to 2
```

Configuration for ACL-Based Mirroring on ports 1 and 2 that results from the trunk being deleted.

```
interface ethernet 1
  ACL-mirror-port ethernet 3
interface ethernet 2
  ACL-mirror-port ethernet 3
```

Configuring ACL-based mirroring for ACLs bound to virtual interfaces

For configurations that have an ACL configured for ACL-Based Mirroring bound to a virtual interface, you must configure the **ACL-mirror-port** command on a physical port that is a member of the same VLAN as the virtual interface. Additionally, only traffic that arrives at ports that belong to the same port group as the physical port where the **ACL-mirror-port** command is configured will be mirrored. This follows the same rules described in [“Ports from a port region must be mirrored to the same destination mirror port”](#) on page 399.

For example, in the following configuration ports 1,2 and 3 are in VLAN 10 with ve 10. Ports 1 and 2 belong to the same port group while port 3 belongs to another port group.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#tagged ethernet 1 to 2
PowerConnect(config-vlan-10)#tagged ethernet 3
PowerConnect(config-vlan-10)#router-interface ve 10

PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 5
PowerConnect(config)#interface ve 10
PowerConnect(config-vif-10)#ip address 10.10.10.254/24
PowerConnect(config-vif-10)#ip access-group 102 in
PowerConnect(config)#access-list 102 permit ip any any mirror
```

In this configuration, the **ACL-mirror-port** command is configured on port 1 which is a member of ve 10. Because of this, ACL-Based Mirroring will only apply to VLAN 10 traffic that arrives on ports 1 and 2. It will not apply to VLAN 10 traffic that arrives on port 3 because that port belongs to a different port group than ports 1 and 2. This is because if you apply ACL-Based Mirroring on an entire VE, and enable mirroring in only one port region, traffic that is in the same VE but on a port in a different port region will not be mirrored.

To make the configuration apply ACL-Based Mirroring to VLAN 10 traffic arriving on port 3, you must add the following command to the configuration.

```
PowerConnect(config)#interface ethernet 3
PowerConnect(config-if-e10000-3)#ACL-mirror-port ethernet 5
```

If a port is in both mirrored and non-mirrored VLANs, only traffic on the port from the mirrored VLAN will be mirrored. For example, the following configuration adds VLAN 20 to the previous configuration. In this example, ports 1 and 2 are in both VLAN 10 and VLAN 20. ACL-Based Mirroring is only applied to VLAN 10. Consequently, traffic that is on ports 1 and 2 that belongs to VLAN 20 will not be mirrored.

```
PowerConnect(config)#vlan 10
PowerConnect(config-vlan-10)#tagged ethernet 1 to 2
PowerConnect(config-vlan-10)#tagged ethernet 3
PowerConnect(config-vlan-10)#router-interface ve 10

PowerConnect(config)#vlan 20
PowerConnect(config-vlan-20)#tagged ethernet 1 to 2

PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#ACL-mirror-port ethernet 5
PowerConnect(config)#interface ve 10
PowerConnect(config-vif-10)#ip address 10.10.10.254/24
PowerConnect(config-vif-10)#ip access-group 102 in
PowerConnect(config)#access-list 102 permit ip any any mirror
```

MAC filter-based mirroring

This feature allows traffic entering an ingress port to be monitored from a mirror port connected to a data analyzer, based on specific source and destination MAC addresses. This feature supports mirroring of inbound traffic only. Outbound mirroring is not supported.

MAC-Filter-Based Mirroring allows a user to specify a particular stream of data for mirroring using a filter. This eliminates the need to analyze all incoming data to the monitored port. To configure MAC-Filter-Based Mirroring, the user must perform three steps:

- Define a mirror port
- Create a MAC filter with a mirroring clause
- Apply the MAC filter to an interface

The following sections describe these steps.

Configuring MAC filter-based mirroring on PowerConnect B-Series TI24X devices

The following example shows how to configure MAC filter-based mirroring on PowerConnect B-Series TI24X devices.

1. Configure the mirror port.

```
PowerConnect(config)#mirror-port ethernet 2
```

2. Configure the MAC filter inbound mirror clause.

```
PowerConnect(config)#mac filter 1 permit 0000.0000.0010 ffff.ffff.ffff any mirror
```

3. Apply the MAC filter inbound mirror clause to the monitor port.


```
PowerConnect(config)#int e 5  
PowerConnect(config-if-e10000-5)#mac filter-group 1
```

4. Configure the monitor port to use the mirror port.

```
PowerConnect(config-if-e10000-5)#acl-mirror-port ethernet 2
```

To display ACL mirror settings, enter the **show access-list all** command.

```
PowerConnect#show access-list all  
Extended IP access list 101  
permit ip any any mirror
```

14 MAC filter-based mirroring

Configuring Quality of Service

Classification

Quality of Service (QoS) features are used to prioritize the use of bandwidth in a switch. When QoS features are enabled, traffic is classified as it arrives at the switch, and processed through on the basis of configured priorities. Traffic can be dropped, prioritized for guaranteed delivery, or subject to limited delivery options as configured by a number of different mechanisms.

This chapter describes how QoS is implemented and configured in PowerConnect devices.

Classification is the process of selecting packets on which to perform QoS, reading the QoS information and assigning a priority to the packets. The classification process assigns a priority to packets as they enter the switch. These priorities can be determined on the basis of information contained within the packet or assigned to the packet as it arrives at the switch. Once a packet or traffic flow is classified, it is mapped to a forwarding priority queue.

Packets on devices are classified in up to eight traffic classes with values between 0 and 7. Packets with higher priority classifications are given a precedence for forwarding.

Processing of classified traffic

The **trust level** in effect on an interface determines the type of QoS information the device uses for performing QoS. The device establishes the trust level based on the configuration of various features and if the traffic is switched or routed. The trust level can be one of the following:

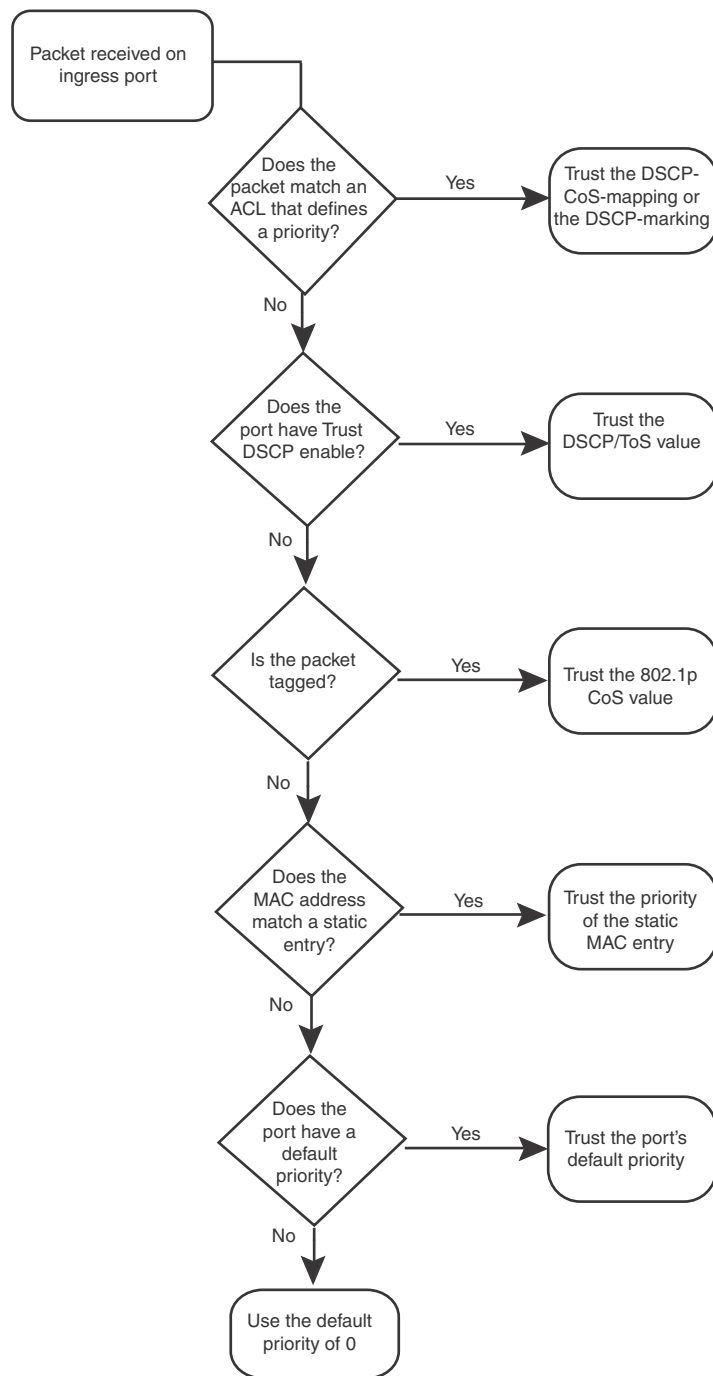
- Ingress port default priority
- Static MAC address
- Layer 2 Class of Service (CoS) value – This is the 802.1p priority value in the Ethernet frame. It can be a value from 0 – 7. The 802.1p priority is also called the Class of Service.
- Layer 3 Differentiated Service codepoint (DSCP) – This is the value in the six most significant bits of the IP packet header 8-bit DSCP field. It can be a value from 0 – 63. These values are described in RFCs 2472 and 2475. The DSCP value is sometimes called the DiffServ value. The device automatically maps a packet's DSCP value to a hardware forwarding queue. Refer to [“Viewing QoS settings”](#) on page 418.
- ACL keyword – An ACL can also prioritize traffic and mark it before sending it along to the next hop. This is described in the ACL chapter in the section [“QoS options for IP ACLs”](#) on page 387.

Given the variety of different criteria, there are multiple possibilities for traffic classification within a stream of network traffic. For this reason, the priority of packets must be resolved based on which criteria takes precedence. Precedence follows the scheme illustrated in [Figure 82](#)

Determining the trust level of a packet'

[Figure 82](#) illustrates how the device determines the trust level of a packet.

FIGURE 82 Determining a packet trust level



As shown in the figure, the first criteria considered is whether the packet matches on an ACL that defines a priority. If this is not the case and the packet is tagged, the packet is classified with the 802.1p CoS value. If neither of these are true, the packet is next classified based on the static MAC address, ingress port default priority, or the default priority of zero (0).

Once a packet is classified by one of the procedures mentioned, it is mapped to an internal forwarding queue. There are eight queues designated as 0 to 7. The internal forwarding priority maps to one of these eight queues as shown in [Table 58](#) through [Table 61](#). The mapping between the internal priority and the forwarding queue cannot be changed.

[Table 58](#) through [Table 61](#) show the default QoS mappings that are used if the trust level for CoS or DSCP is enabled.

TABLE 58 Default QoS mappings, columns 0 to 15

DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
802.1p (COS) Value	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
DSCP value	0	1	2	3	4	5	6	7	8	9	10	11	12	12	14	15
Internal Forwarding Priority	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
Forwarding Queue	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1

TABLE 59 Default QoS mappings, columns 16 to 31

DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
802.1p (COS) Value	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
DSCP value	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Internal Forwarding Priority	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
Forwarding Queue	2	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3

TABLE 60 Default QoS mappings, columns 32 to 47

DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
802.1p (COS) Value	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
DSCP value	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
Internal Forwarding Priority	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5
Forwarding Queue	4	4	4	4	4	4	4	4	5	5	5	5	5	5	5	5

TABLE 61 Default QoS mappings, columns 48 to 63

DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
802.1p (COS) Value	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7
DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63

TABLE 61 Default QoS mappings, columns 48 to 63 (Continued)

DSCP value	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Internal Forwarding Priority	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7
Forwarding Queue	6	6	6	6	6	6	6	6	7	7	7	7	7	7	7	7

Mapping between DSCP value and Forwarding Queue cannot be changed. However, mapping between DSCP values and the other properties can be changed as follows:

- **DSCP to internal forwarding priority mapping** – You can change the mapping between the DSCP value and the Internal Forwarding priority value from the default values shown in [Table 58](#) through [Table 61](#). This mapping is used for COS marking and determining the internal priority when the trust level is DSCP. Refer to “[Changing the DSCP -> internal forwarding priority mappings](#)” on page 412.
- **Internal forwarding priority to forwarding queue** – You can reassign an internal forwarding priority to a different hardware forwarding queue. Refer to “[Changing the internal forwarding priority -> hardware forwarding queue mappings](#)” on page 413.

QoS queues

Devices support the eight QoS queues (qosp0 – qosp7) listed in [Table 62](#).

TABLE 62 QoS queues

QoS Priority Level	QoS Queue
0	qosp0 (lowest priority queue)
1	qosp1
2	qosp2
3	qosp3
4	qosp4
5	qosp5
6	qosp6
7	qosp7 (highest priority queue)

The queue names listed above are the default names. If desired, you can rename the queues as instructed in “[Renaming the queues](#)” on page 416.

Packets are classified and assigned to specific queues based on the criteria shown in [Figure 82](#).

Assigning QoS priorities to traffic

By default, all traffic is in the best-effort queue (qosp0) and is honored on tagged ports on all PowerConnect switches. You can assign traffic to a higher queue based on the following:

- Incoming port (sometimes called the ingress port)
- Static MAC entry

The following sections describe how to change the priority for each of the items listed above.

Although it is possible for a packet to qualify for an adjusted QoS priority based on more than one of the criteria listed in the section above, the system always gives a packet the highest priority for which it qualifies. Thus, if a packet is entitled to the premium queue because of its IP source and destination addresses, but is entitled only to the high queue because of its incoming port, the system places the packet in the premium queue on the outgoing port.

When you apply a QoS priority to one of the items listed above, you specify a number from 0 – 7. The priority number specifies the IEEE 802.1 equivalent to one of the eight QoS queues on devices. The numbers correspond to the queues as shown in [Table 58](#).

Changing a port priority

To change the QoS priority of port 1 to the premium queue (qosp7), enter the following commands.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-e10000-1)#priority 7
```

The device will assign priority 7 to untagged switched traffic received on port 1.

Syntax: [no] priority <num>

The <num> parameter can be from 0 – 7 and specifies the IEEE 802.1 equivalent to one of eight QoS queues listed in [Table 58](#).

Assigning static MAC entries to priority queues

By default, all MAC entries are in the best effort queue. When you configure a static MAC entry, you can assign the entry to a higher QoS level.

To configure a static MAC entry and assign the entry to the premium queue, enter commands such as the following.

```
PowerConnect(config)#vlan 9
PowerConnect(config-vlan-9)#static-mac-address 1145.1163.67FF ethernet 1 priority
7
PowerConnect(config-vlan-9)#write memory
```

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <num>]
[host-type | router-type | fixed-host]

The **priority** <num> parameter can be from 0 – 7 and specifies the IEEE 802.1 equivalent to one of the eight QoS queues.

NOTE

The location of the **static-mac-address** command in the CLI depends on whether you configure port-based VLANs on the device. If the device does not have more than one port-based VLAN (VLAN 1, which is the default VLAN containing all ports), the **static-mac-address** command is at the global CONFIG level of the CLI. If the device has more than one port-based VLAN, then the **static-mac-address** command is not available at the global CONFIG level. In this case, the command is available at the configuration level for each port-based VLAN.

Buffer allocation/threshold for QoS queues

By default, Ironware software allocates a certain number of buffers to the outbound transport queue for each port based on QoS priority. The buffers control the total number of packets permitted in the outbound queue for the port. If desired, you can increase or decrease the maximum number of outbound transmit buffers allocated to all QoS queues, or to specific QoS queues on a port or group of ports. For more information, refer to [“Egress buffer thresholds for QoS priorities”](#) on page 187. On PowerConnect B-Series T124X devices, this feature is called **egress buffer threshold**. For more information, refer to [“Egress buffer thresholds for QoS priorities”](#) on page 187.

Marking

Marking is the process of changing the packet QoS information (the 802.1p and DSCP information in a packet) for the next hop. For example, for traffic coming from a device that does not support DiffServ, you can change the packet IP Precedence value into a DSCP value before forwarding the packet.

You can mark a packet Layer 2 CoS value, its Layer 3 DSCP value, or both values. The Layer 2 CoS or DSCP value the device marks in the packet is the same value that results from mapping the packet QoS value into a Layer 2 CoS or DSCP value.

Marking is optional and is disabled by default. Marking is performed using ACLs. When marking is not used, the device still performs the mappings listed in [“Classification”](#) on page 405 for scheduling the packet, but leaves the packet QoS values unchanged when the device forwards the packet.

For configuration syntax, rules, and examples of QoS marking, refer to [“QoS options for IP ACLs”](#) on page 387.

Configuring DSCP-based QoS

IronWare releases support basic DSCP-based QoS (also called Type of Service (ToS) based QoS) as described in this chapter. However, the PowerConnect B-Series T124X family of switches do not support other advanced DSCP-based QoS features as described.

IronWare releases also support marking of the DSCP value. The software can read Layer 3 Quality of Service (QoS) information in an IP packet and select a forwarding queue for the packet based on the information. The software interprets the value in the six most significant bits of the IP packet header 8-bit ToS field as a Diffserv Control Point (DSCP) value, and maps that value to an internal forwarding priority.

The internal forwarding priorities are mapped to one of the eight forwarding queues (qosp0 – qosp7) on the device. During a forwarding cycle, the device gives more preference to the higher numbered queues, so that more packets are forwarded from these queues. So for example, queue qosp7 receives the highest preference while queue qosp0, the best-effort queue, receives the lowest preference.

Application notes

- DSCP-based QoS is not automatically honored for routed and switched traffic. The default is 802.1p to CoS mapping. To honor DSCP-based QoS, you must change the priority mapping to DSCP to CoS mapping. Refer to [“Using ACLs to honor DSCP-based QoS”](#) on page 411.
- When DSCP marking is enabled, the device changes the contents of the inbound packet ToS field to match the DSCP-based QoS value. This differs from the BigIron, which marks the outbound packet ToS field.

Using ACLs to honor DSCP-based QoS

This section shows how to configure devices to honor DSCP-based QoS for routed and switched traffic.

PowerConnect B-Series TI24X Devices

These devices support DSCP-based QoS on a per-port basis. DSCP-based QoS is not automatically honored for switched traffic. The default is 802.1p to CoS mapping. To honor DSCP-based QoS, enter the following command at the interface level of the CLI.

```
PowerConnect(config-if-e10000-11)trust dscp
```

When **trust dscp** is enabled, the interface honors the Layer 3 DSCP value. By default, the interface honors the Layer 2 CoS value.

NOTE

Use the **bridged-routed** keyword in the ACL to honor DSCP for switched traffic in the Layer 3 image. Refer to [“Enabling ACL support for switched traffic in the router image”](#) on page 384.

Configuring the QoS mappings

You can optionally change the following QoS mappings:

- DSCP -> internal forwarding priority
- Internal forwarding priority -> hardware forwarding queue

The mappings are globally configurable and apply to all interfaces.

Default DSCP -> Internal forwarding priority mappings

The DSCP values are described in RFCs 2474 and 2475. [Table 63](#) list the default mappings of DSCP values to internal forwarding priority values.

TABLE 63 Default DSCP to internal forwarding priority mappings

Internal forwarding priority	DSCP value
0 (lowest priority queue)	0 - 7
1	8 - 15
2	16 - 23
3	24 - 31

TABLE 63 Default DSCP to internal forwarding priority mappings (Continued)

Internal forwarding priority	DSCP value
4	32 - 39
5	40 - 47
6	48 - 55
7 (highest priority queue)	56 - 63

Notice that DSCP values range from 0 - 63, whereas the internal forwarding priority values range from 0 - 7. Any DSCP value within a given range is mapped to the same internal forwarding priority value. For example, any DSCP value from 8 - 15 maps to priority 1.

After performing this mapping, the device maps the internal forwarding priority value to one of the hardware forwarding queues.

Table 64 list the default mappings of internal forwarding priority values to the hardware forwarding queues.

TABLE 64 Default mappings of internal forwarding priority values

Internal forwarding priority	Forwarding queues
0 (lowest priority queue)	qosp0
1 ¹	qosp1
2	qosp2
3	qosp3
4	qosp4
5	qosp5
6	qosp6
7 (highest priority queue)	qosp7

1. PowerConnect B-Series TI24X devices supports seven priorities instead of eight when sFlow is enabled. QoS queue 1 is reserved for sFlow and not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

You can change the DSCP -> internal forwarding mappings. You also can change the internal forwarding priority-> hardware forwarding queue mappings.

Changing the DSCP -> internal forwarding priority mappings

To change the DSCP -> internal forwarding priority mappings for all the DSCP ranges, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#qos-tos map dscp-priority 0 2 3 4 to 1
PowerConnect(config)#qos-tos map dscp-priority 8 to 5
PowerConnect(config)#qos-tos map dscp-priority 16 to 4
PowerConnect(config)#qos-tos map dscp-priority 24 to 2
PowerConnect(config)#qos-tos map dscp-priority 32 to 0
```

```
PowerConnect(config)#qos-tos map dscp-priority 40 to 7
PowerConnect(config)#qos-tos map dscp-priority 48 to 3
PowerConnect(config)#qos-tos map dscp-priority 56 to 6
PowerConnect(config)#ip rebind-ACL all
```

The first command in the above example maps priority 1 to DSCP values 0, 2, 3, and 4.

These commands configure the mappings displayed in the DSCP to forwarding priority portion of the QoS information display. To read this part of the display, select the first part of the DSCP value from the d1 column and select the second part of the DSCP value from the d2 row. For example, to read the DSCP to forwarding priority mapping for DSCP value 24, select 2 from the d1 column and select 4 from the d2 row. The mappings that are changed by the command above are shown below in bold type.

```
PowerConnect#show qos-tos
```

...portions of table omitted for simplicity...

```
DSCP-Priority map: (dscp = d1d2)
```

d2	0	1	2	3	4	5	6	7	8	9
d1										
0	1	0	1	1	1	0	0	0	5	1
1	6	1	1	1	1	1	4	2	2	2
2	2	2	2	2	2	3	3	3	3	3
3	3	3	0	4	4	4	4	4	4	4
4	7	5	5	5	5	5	5	5	3	6
5	6	6	6	6	6	6	6	7	7	7
6	7	7	7	7						

Syntax: [no] qos-tos map dscp-priority <dscp-value> [<dscp-value> ...] to <priority>

The <dscp-value> [<dscp-value> ...] parameter specifies the DSCP value ranges you are remapping.

PowerConnect B-Series T124X devices, you can specify up to **eight** DSCP values in the same command, to map to the same forwarding priority. For example

```
PowerConnect(config)#qos-tos map dscp-priority 1 2 3 4 5 6 7 8 to 6.
```

The <priority> parameter specifies the internal forwarding priority.

Changing the internal forwarding priority -> hardware forwarding queue mappings

To reassign an internal forwarding priority to a different hardware forwarding queue, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#qos tagged-priority 2 qosp0
```

Syntax: [no] qos tagged-priority <num> <queue>

The <num> parameter can be from 0 – 7 and specifies the internal forwarding priority.

The <queue> parameter specifies the hardware forwarding queue to which you are reassigning the priority. The default queue names are as follows:

- qosp7

- qos6
- qos5
- qos4
- qos3
- qos2
- qos1
- qos0

Scheduling

Scheduling is the process of mapping a packet to an internal forwarding queue based on its QoS information, and servicing the queues according to a mechanism.

This section describes the scheduling methods used on PowerConnect B-Series T124X devices.

QoS Queuing methods

The following QoS queuing methods are supported in all IronWare releases for the PowerConnect devices:

- **Weighted round robin (WRR)** – WRR ensures that all queues are serviced during each cycle. A weighted fair queuing algorithm is used to rotate service among the eight queues on the PowerConnect devices. The rotation is based on the weights you assign to each queue. This method rotates service among the queues, forwarding a specific number of packets in one queue before moving on to the next one.

WRR is the default queuing method and uses a default set of queue weights.

The number of packets serviced during each visit to a queue depends on the percentages you configure for the queues. The software automatically converts the percentages you specify into weights for the queues.

NOTE

Queue cycles on the PowerConnect devices are based on bytes. These devices service a given number of bytes (based on weight) in each queue cycle.

- **Strict priority(SP)** – SP ensures service for high priority traffic. The software assigns the maximum weights to each queue, to cause the queuing mechanism to serve as many packets in one queue as possible before moving to a lower queue. This method biases the queuing mechanism to favor the higher queues over the lower queues.

For example, strict queuing processes as many packets as possible in qos3 before processing any packets in qos2, then processes as many packets as possible in qos2 before processing any packets in qos1, and so on.

- **Hybrid WRR and SP** – An additional configurable queuing mechanism combines both the strict priority and weighted round robin mechanisms. The combined method enables the device to give strict priority to delay-sensitive traffic such as VOIP traffic, and weighted round robin priority to other traffic types.

By default, when you select the combined SP and WRR queuing method, the device assigns strict priority to traffic in qosp7 and qosp6, and weighted round robin priority to traffic in qosp0 through qosp5. Thus, the device schedules traffic in queue 7 and queue 6 first, based on the strict priority queuing method. When there is no traffic in queue 7 and queue 6, the device schedules the other queues in round-robin fashion from the highest priority queue to the lowest priority queue.

By default, when you specify the combined SP and WRR queuing method, the system balances the traffic among the queues as shown in [Table 65](#). If desired, you can change the default bandwidth values as instructed in the section “[Changing the bandwidth allocations of the hybrid WRR and SP queues](#)” on page 417.

TABLE 65 Default bandwidth for combined SP and WRR queuing methods

Queue	Default bandwidth
qosp7	Strict priority (highest priority)
qosp6	Strict priority
qosp5	25%
qosp4	15%
qosp3	15%
qosp2	15%
qosp1	15%
qosp0	15% (lowest priority)

Selecting the QoS queuing method

By default, devices use the weighted fair queuing method of packet prioritization. To change the method to strict priority or back to weighted fair queuing, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)#qos mechanism strict
```

To change the method back to weighted round robin, enter the following command.

```
PowerConnect(config)#qos mechanism weighted
```

Syntax: [no] qos mechanism strict | weighted

To change the queuing mechanism to the combined SP and WRR method, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)#qos mechanism mixed-sp-wrr
```

Syntax: mechanism mixed-sp-wrr

Configuring the QoS queues

Each of the queues has the following configurable parameters:

- The queue name
- The minimum percentage of a port outbound bandwidth guaranteed to the queue

Renaming the queues

The default queue names are qos7, qos6, qos5, qos4, qos3, qos2, qos1, and qos0. You can change one or more of the names if desired.

To rename queue “qosp3” to “92-octane”, enter the following command.

```
PowerConnect(config)#qos name qosp3 92-octane
```

Syntax: qos name <old-name> <new-name>

The <old-name> parameter specifies the name of the queue before the change.

The <new-name> parameter specifies the new name of the queue. You can specify an alphanumeric string up to 32 characters long.

Changing the minimum bandwidth percentages of the WRR queues

If you are using the weighted round robin mechanism instead of the strict mechanism, you can change the weights for each queue by changing the minimum percentage of bandwidth you want each queue to guarantee for its traffic.

By default, the eight QoS queues on PowerConnect devices receive the following minimum guaranteed percentages of a port total bandwidth. Note that the defaults differ when jumbo frames are enabled.

TABLE 66 Default minimum bandwidth percentages on devices

Queue	Default minimum percentage of bandwidth	
	Without jumbo frames	With jumbo frames
qosp7	75%	44%
qosp6	7%	8%
qosp5	3%	8%
qosp4	3%	8%
qosp3	3%	8%
qosp2	3%	8%
qosp1	3%	8%
qosp0	3%	8%

When the queuing method is weighted round robin, the software internally translates the percentages into weights. The weight associated with each queue controls how many packets are processed for the queue at a given stage of a cycle through the weighted round robin algorithm.

NOTE

Queue cycles on the PowerConnect devices are based on bytes. These devices service a given number of bytes (based on the weight) in each queue cycle.

The bandwidth allocated to each queue is based on the relative weights of the queues. You can change the bandwidth percentages allocated to the queues by changing the queue weights.

There is no minimum bandwidth requirement for a given queue. For example, queue qos3 is not required to have at least 50% of the bandwidth.

Command syntax

To change the bandwidth percentages for the queues, enter commands such as the following. Note that this example uses the default queue names.

```
PowerConnect(config)#qos profile qosp7 25 qosp6 15 qosp5 12 qosp4 12 qosp3 10
qosp2 10 qosp1 10 qosp0 6
Profile qosp7      : Priority7    bandwidth requested  25% calculated  25%
Profile qosp6      : Priority6    bandwidth requested  15% calculated  15%
Profile qosp5      : Priority5    bandwidth requested  12% calculated  12%
Profile qosp4      : Priority4    bandwidth requested  12% calculated  12%
Profile qosp3      : Priority3    bandwidth requested  10% calculated  10%
Profile qosp2      : Priority2    bandwidth requested  10% calculated  10%
Profile qosp1      : Priority1    bandwidth requested  10% calculated  10%
Profile qosp0      : Priority0    bandwidth requested   6% calculated   6%
```

Syntax: [no] qos profile <queue> <percentage> <queue> <percentage> <queue> <percentage> <queue> <percentage> <queue> <percentage> <queue> <percentage> <queue> <percentage>

Each <queue> parameter specifies the name of a queue. You can specify the queues in any order on the command line, but you must specify each queue.

The <percentage> parameter specifies a number for the percentage of the device outbound bandwidth that is allocated to the queue. QoS queues require a minimum bandwidth percentage of 3% for each priority. When jumbo frames are enabled, the minimum bandwidth requirement is 8%. If these minimum values are not met, QoS may not be accurate.

Configuration notes

- The total of the percentages you enter must equal 100.
- PowerConnect devices do not adjust the bandwidth percentages you enter.

PowerConnect B-Series TI24X devices supports seven priorities instead of eight when sFlow is enabled. QoS queue 1 is reserved for sFlow and not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

Changing the bandwidth allocations of the hybrid WRR and SP queues

To change the default bandwidth percentages for the queues when the device is configured to use the combined SP and WRR queuing mechanism, enter commands such as the following. Note that this example uses the default queue names.

```
PowerConnect(config)#qos profile qosp7 sp qosp6 sp qosp5 20 qosp4 16 qosp3 16
qosp2 16 qosp1 16 qosp0 16
```

Syntax: [no] qos profile <queue 7> sp | <queue 6> sp | <percentage> <queue 5> <percentage> <queue 4> <percentage> <queue 3> <percentage> <queue 2> <percentage> <queue 1> <percentage> <queue 0> <percentage>]

Each <queue x> parameter specifies the name of a queue. You can specify the queues in any order on the command line, but you must specify each queue. Note that queue 7 supports strict priority only, queue 6 supports both strict priority and WRR queuing mechanisms, and queues 0 – 5 support the WRR queuing mechanism only.

The **sp** parameter configures strict priority as the queuing mechanism. Note that only queue 7 and queue 6 support this method.

The *<percentage>* parameter configures WRR as the queuing mechanism and specifies the percentage of the device outbound bandwidth allocated to the queue. The queues require a minimum bandwidth percentage of 3% for each priority. When jumbo frames are enabled, the minimum bandwidth requirement is 8%. If these minimum values are not met, QoS may not be accurate.

NOTE

The percentages must add up to 100. The device does not adjust the bandwidth percentages you enter. In contrast, the BigIron QoS does adjust the bandwidth percentages to ensure that each queue has at least its required minimum bandwidth percentage.

Viewing QoS settings

To display the QoS settings for all of the queues, enter the **show qos-profiles** command.

```
PowerConnect#show qos-profiles all
bandwidth scheduling mechanism: weighted priority
Profile qosp7      : Priority7   bandwidth requested 25% calculated 25%
Profile qosp6      : Priority6   bandwidth requested 15% calculated 15%
Profile qosp5      : Priority5   bandwidth requested 12% calculated 12%
Profile qosp4      : Priority4   bandwidth requested 12% calculated 12%
Profile qosp3      : Priority3   bandwidth requested 10% calculated 10%
Profile qosp2      : Priority2   bandwidth requested 10% calculated 10%
Profile qosp1      : Priority1   bandwidth requested 10% calculated 10%
Profile qosp0      : Priority0   bandwidth requested 6%  calculated 6%
```

Syntax: **show qos-profiles all | <name>**

The **all** parameter displays the settings for all eight queues.

The *<name>* parameter displays the settings for the specified queue.

Viewing DSCP-based QoS settings

To display configuration information for DSCP-based QoS, enter the following command at any level of the CLI


```
PowerConnect#show qos-tos
DSCP-->Traffic-Class map: (DSCP = d1d2: 00, 01...63)
  d2|  0  1  2  3  4  5  6  7  8  9
d1  |
-----+-----
  0  |  0  0  0  0  0  0  0  0  1  1
  1  |  1  1  1  1  1  1  2  2  2  2
  2  |  2  2  2  2  3  3  3  3  3  3
  3  |  3  3  4  4  4  4  4  4  4  4
  4  |  5  5  5  5  5  5  5  5  6  6
  5  |  6  6  6  6  6  6  7  7  7  7
  6  |  7  7  7  7
-----+-----

Traffic-Class-->802.1p-Priority map (use to derive DSCP--802.1p-Priority):
Traffic | 802.1p
Class   | Priority
-----+-----
  0     |     0
  1     |     1
  2     |     2
  3     |     3
  4     |     4
  5     |     5
  6     |     6
  7     |     7
-----+-----
```

Syntax: show qos-tos

This command shows the following information.

TABLE 67 DSCP-based QoS configuration information

This field...	Displays...
DSCP-priority map	
d1 and d2	The DSCP to forwarding priority mappings that are currently in effect. NOTE: The example above shows the default mappings. If you change the mappings, the command displays the changed mappings
Traffic class -> 802.1p priority map	
Traffic Class and 802.1p Priority	The traffic class to 802.1p Priority mappings that are currently in effect. NOTE: The example above shows the default mappings. If you change the mappings, the command displays the changed mappings.

15 Viewing DSCP-based QoS settings

Configuring Rate Limiting and Rate Shaping on the PowerConnect B-Series TI24X

Rate limiting overview

This chapter describes how to configure rate limiting and rate shaping on PowerConnect B-Series TI24X devices.

Rate limiting applies to inbound ports and rate shaping applies to outbound ports.

Port-based fixed rate limiting is supported on inbound ports. This feature allows you to specify the maximum number of *kilobits* a given port on a PowerConnect device can receive. The port drops bytes or kilobits that exceed the limit you specify. You can configure a Fixed Rate Limiting policy on a port inbound direction only. Fixed rate limiting applies to all traffic on the rate limited port.

Fixed rate limiting is at line rate and occurs in hardware. Refer to [“Rate limiting in hardware”](#) on page 421.

When you specify the maximum number of kilobits, you specify it in bits per second (bps). The Fixed Rate Limiting policy applies to one-second intervals and allows the port to receive the number of kilobits you specify in the policy, but drops additional bytes or kilobits. Unused bandwidth is not carried over from one interval to the next.

NOTE

Dell recommends that you do not use Fixed Rate Limiting on ports that receive route control traffic or Spanning Tree Protocol (STP) control traffic. If the port drops control packets due to the Fixed Rate Limiting policy, routing or STP can be disrupted.

Rate limiting in hardware

Each device supports line-rate rate limiting in hardware. The device creates entries in Content Addressable Memory (CAM) for the rate limiting policies. The CAM entries enable the device to perform the rate limiting in hardware instead of sending the traffic to the CPU. The device sends the first packet in a given traffic flow to the CPU, which creates a CAM entry for the traffic flow. A CAM entry consists of the source and destination addresses of the traffic. The device uses the CAM entry for rate limiting all the traffic within the same flow. A rate limiting CAM entry remains in the CAM for two minutes before aging out.

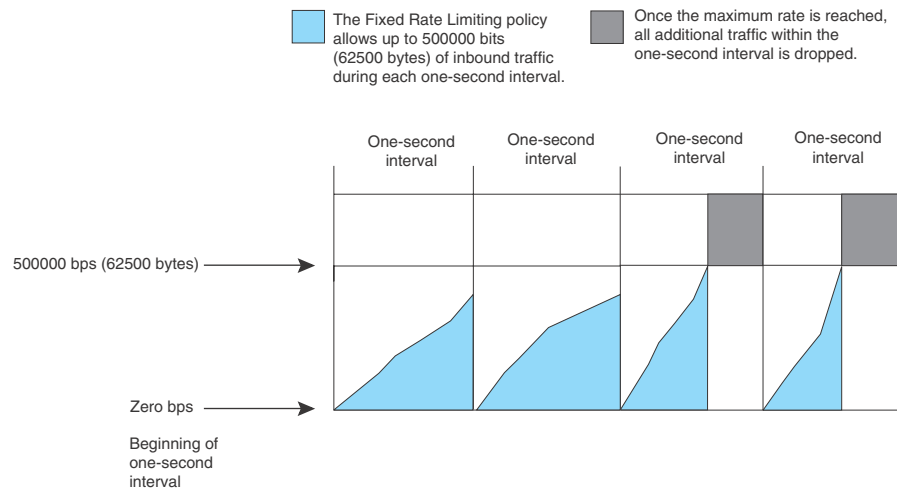
How Fixed Rate Limiting works

Fixed Rate Limiting counts the number of kilobits (PowerConnect devices) that a port receives, in one second intervals. If the number exceeds the maximum number you specify when you configure the rate, the port drops all further inbound packets for the duration of the one-second interval.

Once the one-second interval is complete, the port clears the counter and re-enables traffic.

Figure 83 shows an example of how Fixed Rate Limiting works. In this example, a Fixed Rate Limiting policy is applied to a port to limit the inbound traffic to 500000 bits (62500 bytes) a second. During the first two one-second intervals, the port receives less than 500000 bits in each interval. However, the port receives more than 500000 bits during the third and fourth one-second intervals, and consequently drops the excess traffic.

FIGURE 83 Fixed Rate Limiting



NOTE

The software counts the kilobits by polling statistics counters for the port every 100 milliseconds, which provides 10 readings each second. Due to the polling interval, the Fixed Rate Limiting policy has an accuracy of within 10% of the port's line rate. It is therefore possible for the policy to sometimes allow more traffic than the limit you specify, but the extra traffic is never more than 10% of the port's line rate.

Configuration notes

- Rate limiting is available only on inbound ports.
- Fixed rate limiting is supported on GbE and 10 GbE ports on PowerConnect switches.
- Fixed rate limiting is not supported on tagged ports in the base Layer 3 and full Layer 3 images
- The rate limit on IPv6 hardware takes several seconds to take effect at higher configured rate limit values. For example, if the configured rate limit is 750 Mbps, line-rate limiting could take up to 43 seconds to take effect.

Configuring a port-based rate limiting policy

To configure rate limiting on a PowerConnect port, enter commands such as the following.

```
PowerConnect(config)#interface ethernet 24
PowerConnect(config-if-e10000-24)#rate input fixed 64
```

These commands configure a fixed rate limiting policy that allows port 24 to receive a maximum of 64 kilobits per second (65536 bytes per second). If the port receives additional bits during a given one-second interval, the port drops all inbound packets on the port until the next one-second interval starts.

Syntax: `[no] rate-limit input fixed <average-rate>`

For PowerConnect devices, the `<average-rate>` parameter specifies the maximum number of *kilobits* per second (kbps) the port can receive. The minimum rate that can be configured is 64,000 bytes per second.

Configuring an ACL-based rate limiting policy

IP ACL-based rate limiting of inbound traffic provides the facility to limit the rate for IP traffic that matches the permit conditions in extended IP ACLs. This feature is available in the Layer 2 and Layer 3 code.

To configure ACL-based rate limiting on a device, you create individual **traffic policies**, then reference the traffic policies in one or more ACL entries (also called clauses or statements). The traffic policies become effective on ports to which the ACLs are bound.

For configuration procedures for ACL-based rate limiting, refer to [Chapter 17, “Configuring Traffic Policies”](#).

Displaying the fixed rate limiting configuration

To display the fixed rate limiting configuration on the device, enter the following command.

```
PowerConnect#show rate-limit fixed
Total rate-limited interface count: 11.
Port      Configured Input Rate      Actual Input Rate
  1         1000000                    1000000
  3         10000000                   10005000
  7         10000000                   10000000
  9         7500000                    7502000
 11         8000000                    7999000
 12         8000000                    7999000
 13         8000000                    7999000
 14         8000000                    7999000
 15         8000000                    7999000
 21         8000000                    8000000
 25         7500000                    7502000
```

Syntax: `show rate-limit fixed`

The command lists the ports on which fixed rate limiting is configured, and provides the information listed in [Table 68](#) for each of the ports.

TABLE 68 CLI display of Fixed Rate Limiting information

This field...	Displays...
Total rate-limited interface count	The total number of ports that are configured for Fixed Rate Limiting.
Port	The port number.

TABLE 68 CLI display of Fixed Rate Limiting information (Continued)

This field...	Displays...
Configured Input Rate	The maximum rate requested for inbound traffic. The rate is measured in or kilobits per second (kbps) for PowerConnect devices.
Actual Input Rate	The actual maximum rate provided by the hardware. The rate is measured in kilobits per second (kbps) for PowerConnect devices.

Rate shaping overview

Outbound Rate Shaping is a port-level feature that is used to shape the rate and control the bandwidth of outbound traffic on a port. This feature smooths out excess and bursty traffic to the configured maximum limit before it is sent out on a port. Packets are stored in available buffers and then forwarded at a rate no greater than the configured limit. This process provides for better control over the inbound traffic of neighboring devices.

The device has one global rate shaper for a port and one rate shaper for each port priority queue. Rate shaping is done on a single-token basis, where each token is defined to be 1 byte.

Configuration notes

The following rules apply when configuring outbound rate shapers:

- Outbound rate shapers can be configured *only* on physical ports, not on virtual or loopback ports.
- For trunk ports, the rate shaper must be configured on individual ports of a trunk using the **config-trunk-ind** command (trunk configuration level); you cannot configure a rate shaper for a trunk.
- When outbound rate shaping is enabled on a port on an IPv4 device, the port QoS queuing method (**qos mechanism**) will be strict mode. This applies to IPv4 devices only. On IPv6 devices, the QoS mechanism is whatever method is configured on the port, even when outbound rate shaping is enabled.
- You can configure a rate shaper for a port and for the individual priority queues of that port. However, if a port rate shaper is configured, that value overrides the rate shaper value of a priority queue if the priority queue rate shaper is greater than the rate shaper for the port.
- On PowerConnect B-Series T124X devices, configured rate shaper values are rounded up to the nearest values programmable by the hardware.

Configuring outbound rate shaping for a port

To configure the maximum rate at which outbound traffic is sent out on a port, enter commands such as the following.

```
PowerConnect(config)#interface e 2
PowerConnect(config-if-e10000-2)#rate-limit output shaping 1300
```

- On PowerConnect B-Series T124X devices, the configured 1300 Kbps outbound rate shaping on port 2 is rounded up to the nearest value programmable by the hardware, which is 1344 Kbps. This value is the actual limit on the port for outbound traffic.

Syntax: [no] **rate-limit output shaping** <value>

On PowerConnect B-Series TI24X devices, you can specify a value up to the port line rate for `<value>`.

Configuring outbound rate shaping for a specific priority

To configure the maximum rate at which outbound traffic is sent out on a port priority queue, enter commands such as the following.

```
PowerConnect(config)#interface e 2
PowerConnect(config-if-e10000-2)#rate-limit output shaping 500 priority 7
```

- On PowerConnect B-Series TI24X devices, the configured 500 Kbps limit for outbound traffic on Priority queue 7 on port 2 is rounded up to the nearest multiple of 504 Kbps, which is 504 Kbps.

Syntax: `[no] rate-limit output shaping <value> priority <priority-queue>`

On PowerConnect B-Series TI24X devices, you can specify a value up to the port line rate for `<value>`.

Specify 0-7 for `<priority-queue>`

Configuring outbound rate shaping for a trunk port

This feature is supported on individual ports of a static trunk group. However, it is not supported on LACP trunk ports on PowerConnect B-Series TI24X devices.

To configure the maximum rate at which outbound traffic is sent out on a trunk port, enter the following on each trunk port where outbound traffic will be shaped.

```
PowerConnect(config)#trunk e 13 to 16
PowerConnect(config-trunk-13-16)#config-trunk-ind
PowerConnect(config-trunk-13-16)#rate-limit output shaping ethe 15 651
PowerConnect(config-trunk-13-16)#rate-limit output shaping ethe 14 1300
```

The above commands configure an outbound rate shaper on port 14 and port 15.

- On PowerConnect B-Series TI24X devices, the configured outbound rate shaper of 651 Kbps on port 15 is rounded up to 704 Kbps. The configured outbound rate shaper of 1300 on port 14 is rounded up to 1344 Kbps.

Syntax: `[no] rate-limit output shaping ethernet <portnum> <value>`

On PowerConnect B-Series TI24X devices, you can specify a `<value>` up to the port line rate.

Displaying rate shaping configurations

To display the configured outbound rate shaper on a device, enter the following command.

```
PowerConnect#show rate-limit output-shaping
Outbound Rate Shaping Limits in Kbps:
  Port  PortMax  Prio0  Prio1  Prio2  Prio3  Prio4  Prio5  Prio6  Prio7
    1         -      -      -      -      -      -      -      -      651
    2      1302      -      -      -      -      -      -      -      -
   15       651      -      -      -      -      -      -      -      -
```

The display lists the ports on a device, the configured outbound rate shaper on a port and for a priority for a port.

16 Rate shaping overview

Configuring Traffic Policies

About traffic policies

This chapter describes how traffic policies are implemented and configured in the PowerConnect B-Series TI24X devices.

Devices use **traffic policies** for the following:

- to rate limit inbound traffic
- to count the packets and bytes per packet to which ACL permit or deny clauses are applied

Traffic policies consist of policy names and policy definitions:

- **Traffic policy name** – This is a string of up to 8 alphanumeric characters that identifies individual traffic policy definitions.
- **Traffic policy definition (TPD)** – This is the command filter associated with a traffic policy name. A TPD can define any one of the following:
 - Rate limiting policy
 - ACL counting policy
 - Combined rate limiting and ACL counting policy

The maximum number of supported active TPDs is a system-wide parameter and depends on the device you are configuring. The total number of active TPDs cannot exceed the system maximum. Refer to [“Maximum number of traffic policies supported on a device”](#) on page 428.

When you apply a traffic policy to an interface, you do so by adding a reference to the traffic policy in an ACL entry, instead of applying the individual traffic policy to the interface. The traffic policy becomes an **active traffic policy** or **active TPD** when you bind its associated ACL to an interface.

To configure traffic policies for ACL-based rate limiting, refer to [“Configuring ACL-based fixed rate limiting”](#) on page 430 and [“Configuring ACL-based adaptive rate limiting”](#) on page 431.

To configure traffic policies for ACL counting, refer to [“Enabling ACL statistics”](#) on page 434.

Configuration notes and feature limitations

Note the following when configuring traffic policies:

- Traffic policies are supported on all PowerConnect B-Series TI24X.
- This feature is supported in the Layer 2 and Layer 3 code.
- This feature applies to IP ACLs only.
- Traffic policies are not supported on 10 Gbps Ethernet interfaces.
- The maximum number of supported active TPDs is a system-wide parameter and depends on the device you are configuring. The total number of active TPDs cannot exceed the system maximum. Refer to [“Maximum number of traffic policies supported on a device”](#) on page 428.

17 Maximum number of traffic policies supported on a device

- You can reference the same traffic policy in more than one ACL entry within an access list. For example, two or more ACL statements in ACL 101 can reference a TPD named TPD1.
- You can reference the same traffic policy in more than one access list. For example, ACLs 101 and 102 could both reference a TPD named TPD1.
- To modify or delete an active traffic policy, you must first unbind the ACL that references the traffic policy.
- When you define a TPD (when you enter the CLI command **traffic-policy**), explicit marking of CoS parameters, such as traffic class and 802.1p priority, are not available on the device. In the case of a TPD defining rate limiting, the device re-marks CoS parameters based on the DSCP value in the packet header and the determined conformance level of the rate limited traffic, as shown in [Table 69](#).

TABLE 69 CoS parameters for packets that use rate limiting traffic policies

If the packet conformance level is...	and the packet DSCP value is...	the device sets the traffic class and 802.1p priority to...
0 (Green) or 1 (Yellow)	0 – 7	0 (lowest priority queue)
	8 – 15	1
	16 – 23	2
	24 – 31	3
	32 – 39	4
	40 – 47	5
	48 – 55	6
	56 – 63	7 (highest priority queue)
2 (Red)	N/A	0 (lowest priority queue)

- When you define a TPD, reference the TPD in an ACL entry, then apply the ACL to a VE in the Layer 3 router code, the rate limit policy is accumulative for all of the ports in the port region. If the VE/VLAN contains ports that are in different port regions, the rate limit policy is applied per port region.

For example, TPD1 has a rate limit policy of 600M and is referenced in ACL 101. ACL 101 is applied to VE 1, which contains ports e 11 to e 14. Because ports e 11 and 12 are in a different port region than e 13 and 14, the rate limit policy will be 600M for ports e 11 and 12, and 600M for ports e 13 and 14.

Maximum number of traffic policies supported on a device

The maximum number of supported active traffic policies is a system-wide parameter and depends on the device you are configuring, as follows:

- By default, up to 1024 active traffic policies are supported on Layer 2 and on Layer 3 switches. This value is fixed on Layer 2 switches and cannot be modified.
- The number of active traffic policies supported on Layer 3 switches varies depending on the configuration and the available system memory. The default value and also the maximum number of traffic policies supported on Layer 3 switches is 1024.

Setting the maximum number of traffic policies supported on a Layer 3 device

If desired you can adjust the maximum number of active traffic policies that a Layer 3 device will support. To do so, enter commands such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)# system-max hw-traffic-conditioner 25
PowerConnect(config)# write memory
PowerConnect(config)# reload
```

NOTE

You must save the configuration and reload the software to place the change into effect.

Syntax: [no] **system-max hw-traffic-conditioner** <num>

<num> is a value from 0 to *n*, where 0 disables hardware resources for traffic policies, and *n* is a number up to 1024. The maximum number you can configure depends on the configuration and available memory on your device. If the configuration you enter causes the device to exceed the available memory, the device will reject the configuration and display a warning message on the console.

NOTE

Dell does not recommend setting the system-max for traffic policies to 0 (zero), since this renders traffic policies ineffective.

ACL-based rate limiting using traffic policies

ACL-based rate limiting provides the facility to limit the rate for IP traffic that matches the permit conditions in extended IP ACLs. This feature is available in the Layer 2 and Layer 3 code.

To configure ACL-based rate limiting, you create individual **traffic policies**, then reference the traffic policies in one or more ACL entries (also called clauses or statements). The traffic policies become effective on ports to which the ACLs are bound. Refer to [“About traffic policies”](#) on page 427.

When you configure a traffic policy for rate limiting, the device automatically enables **rate limit counting**, similar to the two-rate three-color marker (trTCM) mechanism described in RFC 2698 for adaptive rate limiting, and the single-rate three-color marker (srTCM) mechanism described in RFC 2697 for fixed rate limiting. This feature counts the number of bytes and trTCM or srTCM conformance level per packet to which rate limiting traffic policies are applied. Refer to [“ACL and rate limit counting”](#) on page 434.

You can configure ACL-based rate limiting on the following interface types:

- physical Ethernet interfaces
- virtual interfaces
- trunk ports
- specific VLAN members on a port (New in 02.3.03 – refer to [“Applying an IPv4 ACL to specific VLAN members on a port \(Layer 2 devices only\)”](#) on page 385)
- a subset of ports on a virtual interface (New in 02.3.03 – refer to [“Applying an IPv4 ACL to a subset of ports on a virtual interface \(Layer 3 devices only\)”](#) on page 385.)

Support for fixed rate limiting and adaptive rate limiting

PowerConnect B-Series TI24X devices support the following types of ACL-based rate limiting:

- **Fixed rate limiting** – Enforces a strict bandwidth limit. The device forwards traffic that is within the limit but either drops all traffic that exceeds the limit, or forwards all traffic that exceeds the limit at the lowest priority level, according to the action specified in the traffic policy.
- **Adaptive rate limiting** – Enforces a flexible bandwidth limit that allows for bursts above the limit. You can configure Adaptive Rate Limiting to forward, modify the IP precedence of and forward, or drop traffic based on whether the traffic is within the limit or exceeds the limit.

Configuring ACL-based fixed rate limiting

Use the procedures in this section to configure ACL-based fixed rate limiting. Before configuring this feature, see what to consider in [“Configuration notes and feature limitations”](#) on page 427.

Fixed rate limiting enforces a strict bandwidth limit. The port forwards traffic that is within the limit. If the port receives more than the specified number of fragments in a one-second interval, the device either drops or forwards subsequent fragments in hardware, depending on the action you specify.

To implement the ACL-based fixed rate limiting feature, first create a traffic policy, then reference the policy in an extended ACL statement. Lastly, bind the ACL to an interface. Follow the steps below.

1. Create a traffic policy. Enter a command such as the following.

```
PowerConnect(config)# traffic-policy TPD1 rate-limit fixed 100 exceed-action drop
```

2. Create an extended ACL entry or modify an existing extended ACL entry that references the traffic policy.

Example

```
PowerConnect(config)# access-list 101 permit ip host 210.10.12.2 any traffic-policy TPD1
```

3. Bind the ACL to an interface.

```
PowerConnect(config)# int e 5
PowerConnect(config-if-e5)# ip access-group 101 in
PowerConnect(config-if-e5)# exit
```

The above commands configure a fixed rate limiting policy that allows port e5 to receive a maximum traffic rate of 100 kbps. If the port receives additional bits during a given one-second interval, the port drops the additional inbound packets that are received within that one-second interval.

Syntax: [no] traffic-policy <TPD name> rate-limit fixed <cir value> exceed-action <action> [count]

Syntax: access-list <num> permit | deny.... traffic policy <TPD name>

Syntax: [no] ip access-group <num> in

NOTE

For brevity, some parameters were omitted from the above access-list syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the **no** form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

<TPD name> is the name of the traffic policy definition. This value can be 8 or fewer alphanumeric characters.

rate-limit fixed specifies that the traffic policy will enforce a strict bandwidth.

<cir value> is the committed information rate in kbps. This value can be from 64 – 1000000 Kbps.

exceed-action <action> specifies the action to be taken when packets exceed the configured cir value. Refer to [“Specifying the action to be taken for packets that are over the limit”](#) on page 433.

The **count** parameter is optional and enables ACL counting. Refer to [“ACL and rate limit counting”](#) on page 434.

Configuring ACL-based adaptive rate limiting

Use the procedures in this section to configure ACL-based adaptive rate limiting. Before configuring this feature, see what to consider in [“Configuration notes and feature limitations”](#) on page 427.

[Table 70](#) lists the configurable parameters for ACL-based adaptive rate limiting.

TABLE 70 ACL-Based adaptive rate limiting parameters

Parameter	Definition
Committed Information Rate (CIR)	The guaranteed kilobit rate of inbound traffic that is allowed on a port.
Committed Burst Size (CBS)	The number of bytes per second allowed in a burst before some packets will exceed the committed information rate. Larger bursts are more likely to exceed the rate limit. The CBS must be a value greater than zero (0). Dell recommends that this value be equal to or greater than the size of the largest possible IP packet in a stream. For PowerConnect B-Series TI24X devices, the CBS value is specified in kilobits.
Peak Information Rate (PIR)	The peak maximum kilobit rate for inbound traffic on a port. The PIR must be equal to or greater than the CIR.
Peak Burst Size (PBS)	The number of bytes per second allowed in a burst before all packets will exceed the peak information rate. The PBS must be a value greater than zero (0). Dell recommends that this value be equal to or greater than the size of the largest possible IP packet in the stream. For PowerConnect B-Series TI24X devices, the PBS value is specified in kilobits.

If a port receives more than the configured bit or byte rate in a one-second interval, the port will either drop or forward subsequent data in hardware, depending on the action you specify.

To implement the ACL-based adaptive rate limiting feature, first create a traffic policy then reference the policy in an extended ACL statement. Lastly, bind the ACL to an interface. Follow the steps below.

1. Create a traffic policy. Enter a command such as the following.

```
PowerConnect(config)# traffic-policy TPDAfour rate-limit adaptive cir 10000
cbs 1600 pir 20000 pbs 4000 exceed-action drop
```

2. Create a new extended ACL entry or modify an existing extended ACL entry that references the traffic policy.

Example

```
PowerConnect(config)# access-list 104 permit ip host 210.10.12.2 any
traffic-policy TPDAfour
```

3. Bind the ACL to an interface.

```
PowerConnect(config)# int e 7
PowerConnect(config-if-e7)# ip access-group 104 in
PowerConnect(config-if-e7)# exit
```

The above commands configure an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps on port e7 and allows bursts of up to 1600 bytes. It also enforces a peak rate of 20000 kbps and allows bursts of 4000 bytes above the PIR limit. If the port receives additional bits during a given one-second interval, the port drops all packets on the port until the next one-second interval starts.

Syntax: `[no] traffic-policy <TPD name> rate-limit adaptive cir <cir value> cbs <cbs value> pir <pir value> pbs <pbs value> exceed-action <action> [count]`

Syntax: `access-list <num> permit | deny.... traffic policy <TPD name>`

Syntax: `[no] ip access-group <num> in`

NOTE

For brevity, some parameters were omitted from the above access-list syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the **no** form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

<TPD name> is the name of the traffic policy definition. This value can be 8 or fewer alphanumeric characters.

rate-limit adaptive specifies that the policy will enforce a flexible bandwidth limit that allows for bursts above the limit.

<cir value> is the committed information rate in kbps. Refer to [Table 70](#).

<cbs value> is the committed burst size in bytes. Refer to [Table 70](#).

<pir value> is the peak information rate in kbps. Refer to [Table 70](#).

<pbs value> is the peak burst size in bytes. Refer to [Table 70](#).

exceed-action <action> specifies the action to be taken when packets exceed the configured values. Refer to [“Specifying the action to be taken for packets that are over the limit”](#) on page 433.

The **count** parameter is optional and enables ACL counting. Refer to [“ACL and rate limit counting”](#) on page 434.

Specifying the action to be taken for packets that are over the limit

You can specify the action to be taken when packets exceed the configured cir value for fixed rate limiting, or the cir, cbs, pir, and pbs values for adaptive rate limiting. You can specify one of the following actions:

- Drop packets that exceed the limit
- Permit packets that exceed the limit and forward them at the lowest priority level

Dropping packets that exceed the limit

This section shows some example configurations and provides the CLI syntax for configuring a port to drop packets that exceed the configured limits for rate limiting.

Example

The following shows an example fixed rate limiting configuration.

```
PowerConnect(config)# traffic-policy TPD1 rate-limit fixed 10000 exceed-action drop
```

The above command sets the fragment threshold at 10,000 per second. If the port receives more than 10,000 packet fragments in a one-second interval, the device drops the excess fragments.

Syntax: `traffic-policy <TPD name> rate-limit fixed <cir value> exceed-action drop`

Example

The following shows an example adaptive rate limiting configuration.

```
PowerConnect(config)# traffic-policy TPDfour rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000 exceed-action drop
```

The above commands configure an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps on port e7 and allows bursts of up to 1600 bytes. It also enforces a peak rate of 20000 kbps and allows bursts of 4000 bytes above the PIR limit. If the port receives additional bits during a given one-second interval, the port drops all packets on the port until the next one-second interval starts.

Syntax: `traffic-policy rate-limit adaptive cir <cir value> cbs <cbs value> pir <pir value> pbs <pbs value> exceed-action drop`

Permitting packets that exceed the limit

This section shows some example configurations and provides the CLI syntax for configuring a port to permit packets that exceed the configured limit for rate limiting.

Example

The following shows an example fixed rate limiting configuration.

```
PowerConnect(config)# traffic-policy TPD1 rate-limit fixed 10000 exceed-action permit-at-low-pri
```

The above command sets the fragment threshold at 10,000 per second. If the port receives more than 10,000 packet fragments in a one-second interval, the device takes the specified action. The action specified with this command is to permit excess fragments and forward them at the lowest priority level.

Syntax: [no] traffic-policy <TPD name> rate-limit fixed <cir value> exceed-action permit-at-low-pri

Example

The following shows an example adaptive rate limiting configuration.

```
PowerConnect(config)# traffic-policy TPDAfour rate-limit adaptive cir 10000 cbs
1600 pir 20000 pbs 4000 exceed-action permit-at-low-pri
```

The above commands configure an adaptive rate limiting policy that enforces a guaranteed committed rate of 10000 kbps on port e7 and allows bursts of up to 1600 bytes. It also enforces a peak rate of 20000 kbps and allows bursts of 4000 bytes above the PIR limit. If the port receives additional bits during a given one-second interval, the port permits all packets on the port and forwards the packets at the lowest priority level.

Syntax: traffic-policy rate-limit adaptive cir <cir value> cbs <cbs value> pir <pir value> pbs <pbs value> exceed-action permit-at-low-pri

ACL and rate limit counting

ACL counting, also called **ACL statistics**, enables the device to count the number of packets and the number of bytes per packet to which ACL filters are applied.

Rate limit counting counts the number of bytes and conformance level per packet to which rate limiting traffic policies are applied. The device uses the counting method similar to the two-rate three-color marker (trTCM) mechanism described in RFC 2698 for adaptive rate limiting, and the single-rate three-color marker (srTCM) mechanism described in RFC 2697 for fixed rate limiting. Rate limit counting is automatically enabled when a traffic policy is enforced (active). You can view these counters using the show commands listed in [“Viewing traffic policies”](#) on page 437.

For more information about traffic policies, refer to [“About traffic policies”](#) on page 427.

Enabling ACL statistics

NOTE

ACL statistics and **ACL counting** are used interchangeably throughout this chapter and mean the same thing.

Use the procedures in this section to configure ACL statistics. Before configuring this feature, see what to consider in [“Configuration notes and feature limitations”](#) on page 427.

You also can enable ACL statistics when you create a traffic policy for rate limiting. Refer to [“Enabling ACL statistics with rate limiting traffic policies”](#) on page 435.

Follow the steps given below to implement the ACL counting feature.

1. Create a traffic policy. Enter a command such as the following

```
PowerConnect(config)# traffic-policy TPD5 count
```

2. Create an extended ACL entry or modify an existing extended ACL entry that references the traffic policy definition.

Example

```
PowerConnect(config)# access-list 101 permit ip host 210.10.12.2 any
traffic-policy TPD5
```


3. Bind the ACL to an interface.

```
PowerConnect(config)# int e 4
PowerConnect(config-if-e4)# ip access-group 101 in
PowerConnect(config-if-e4)# exit
```

The above commands configure an ACL counting policy and apply it to port e4. Port e4 counts the number of packets and the number of bytes on the port that were permitted or denied by ACL filters.

Syntax: [no] traffic-policy <TPD name> count

Syntax: access-list <num> permit | deny.... traffic policy <TPD name>

Syntax: [no] ip access-group <num> in

NOTE

For brevity, some parameters were omitted from the above access-list syntax.

The software allows you to add a reference to a non-existent TPD in an ACL statement and to bind that ACL to an interface. The software does not issue a warning or error message for non-existent TPDs.

Use the **no** form of the command to delete a traffic policy definition. Note that you cannot delete a traffic policy definition if it is currently in use on a port. To delete a traffic policy, first unbind the associated ACL.

<TPD name> is the name of the traffic policy definition. This value can be 8 alphanumeric characters or less.

Enabling ACL statistics with rate limiting traffic policies

The configuration example in the section [“Enabling ACL statistics”](#) on page 434 shows how to enable ACL counting without having to configure parameters for rate limiting. You also can enable ACL counting while defining a rate limiting traffic policy, as illustrated in the following configuration examples.

Example

To enable ACL counting while defining traffic policies for fixed rate limiting, enter commands such as the following at the Global CONFIG Level of the CLI.

```
PowerConnect(config)# traffic-policy TPD1 rate-limit fixed 1000 count
exceed-action drop
PowerConnect(config)# traffic-policy TPD2 rate-limit fixed 10000 exceed-action
drop count
```

Syntax: [no] traffic-policy <TPD name> rate-limit fixed <cir value> exceed-action <action> count

Example

To enable ACL counting while defining traffic policies for adaptive rate limiting, enter commands such as the following at the Global CONFIG Level of the CLI.

```
traffic-policy TPDA4 rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000
count exceed-action drop
traffic-policy TPDA5 rate-limit adaptive cir 10000 cbs 1600 pir 20000 pbs 4000
exceed-action permit-at-low-pri count
```

Syntax: traffic-policy rate-limit adaptive cir <cir value> cbs <cbs value> pir <pir value> pbs <pbs value> exceed-action <action> count

Viewing ACL and rate limit counters

When ACL counting is enabled on the device, you can use **show** commands to display the total packet count and byte count of the traffic filtered by ACL statements. The output of the show commands also display the rate limiting traffic counters, which are automatically enabled for active rate limiting traffic policies.

Use either the **show access-list accounting** command or the **show statistics traffic-policy** command to display ACL and traffic policy counters. The output of these commands are identical. The following shows an example output.

```
PowerConnect# show access-list accounting traffic-policy g_voip
Traffic Policy - g_voip:
General Counters:
Port Region#                Byte Count                Packet Count
-----
7 (1 - 12)                   85367040                  776064
All port regions              84367040                  776064
Rate Limiting Counters:
Port Region#                Green Conformance         Yellow Conformance         Red Conformance
-----
7 (1 - 12)                   329114195612139520      37533986897781760        0
All port regions              329114195612139520      37533986897781760        0
```

Syntax: show access-list accounting traffic-policy [<TPD name>]

or

Syntax: show statistics traffic-policy [<TPD name>]

[Table 71](#) explains the output of the **show access-list accounting** and **show statistics traffic-policy** commands.

TABLE 71 ACL and rate limit counting statistics

This line...	Displays...
Traffic Policy	The name of the traffic policy.
General Counters	
Port Region #	The port region to which the active traffic policy applies.
Byte Count	The number of bytes that were filtered (matched ACL clauses).
Packet Count	The number of packets that were filtered (matched ACL clauses).
Rate Limiting Counters	
Port Region#	The port region to which the active traffic policy applies.
Green Conformance	The number of bytes that did not exceed the CIR packet rate.
Yellow Conformance	The number of bytes that exceeded the CIR packet rate.
Red Conformance	The number of bytes that exceeded the PIR packet rate.

Clearing ACL and rate limit counters

The device keeps a running tally of the number of packets and the number of bytes per packet that are filtered by ACL statements and rate limiting traffic policies. You can clear these accumulated counters, essentially resetting them to zero. To do so, use either the **clear access-list accounting traffic-policy** or the **clear statistics traffic-policy** command.

To clear the counters for ACL counting and rate limit counting, enter commands such as the following.

```
PowerConnect(config)# clear access-list accounting traffic-policy CountOne
PowerConnect(config)# clear statistics traffic-policy CountTwo
```

Syntax: **clear access-list accounting traffic-policy** <TPD name>

or

Syntax: **clear statistics traffic-policy** <TPD name>

where <TPD name> is the name of the traffic policy definition for which you want to clear traffic policy counters.

Viewing traffic policies

To view traffic policies that are currently defined on the device, enter the **show traffic-policy** command. An example display output is shown below. [Table 72](#) defines the output.

```
PowerConnect# show traffic-policy t_voip
Traffic Policy - t_voip:
Metering Enabled, Parameters:
    Mode: Adaptive Rate-Limiting
    cir: 100 kbps, cbs: 2000 bytes, pir: 200 kbps, pbs: 4000
bytes
Counting Not Enabled
Number of References/Bindings:1
```

Syntax: **show traffic-policy** [<TPD name>]

To display all traffic policies, enter the **show traffic-policy** command without entering a TPD name.

TABLE 72 Traffic policy information

This line...	Displays...
Traffic Policy	The name of the traffic policy.
Metering	Shows whether or not rate limiting was configured as part of the traffic policy: <ul style="list-style-type: none"> • Enabled – The traffic policy includes a rate limiting configuration. • Disabled – The traffic policy does not include a rate limiting configuration
Mode	If rate limiting is enabled, this field shows the type of metering enabled on the port: <ul style="list-style-type: none"> • Fixed Rate-Limiting • Adaptive Rate-Limiting
cir	The committed information rate, in kbps, for the adaptive rate-limiting policy.
cbs	The committed burst size, in bytes per second, for the adaptive rate-limiting policy.
pir	The peak information rate, in kbps, for the adaptive rate-limiting policy.
pbs	The peak burst size, in bytes per second, for the adaptive rate-limiting policy.

17 Viewing traffic policies

TABLE 72 Traffic policy information (Continued)

This line...	Displays...
Counting	Shows whether or not ACL counting was configured as part of the traffic policy: <ul style="list-style-type: none"><li data-bbox="626 359 1276 380">• Enabled – Traffic policy includes an ACL counting configuration.<li data-bbox="626 390 1422 411">• Disabled – Traffic policy does not include an ACL traffic counting configuration.
Number of References/Bindings	The number of port regions to which this traffic policy applies. For example, if the traffic policy is applied to a trunk group that includes ports e 9, 10, 11, and 12, the value in this field would be 2, because these four trunk ports are in two different port regions.

Configuring IP Multicast Traffic Reduction for PowerConnect B-Series TI24X Switches

IGMP snooping overview

When a device processes a multicast packet, by default, it broadcasts the packets to all ports except the incoming port of a VLAN. Packets are flooded by hardware without going to the CPU. This behavior causes some clients to receive unwanted traffic.

IGMP snooping provides multicast containment by forwarding traffic to only the ports that have IGMP receivers for a specific multicast group (destination address). A device maintains the IGMP group membership information by processing the IGMP reports and leave messages, so traffic can be forwarded to ports receiving IGMP reports.

An IPv4 multicast address is a destination address in the range of 224.0.0.0 to 239.255.255.255. Addresses of 224.0.0.X are reserved. Because packets destined for these addresses may require VLAN flooding, devices do not snoop in the reserved range. Data packets destined to addresses in the reserved range are flooded to the entire VLAN by hardware, and mirrored to the CPU. Multicast data packets destined for the non-reserved range of addresses are snooped. A client must send IGMP reports in order to receive traffic. If an application outside the reserved range requires VLAN flooding, the user must configure a static group that applies to the entire VLAN.

An IGMP device's responsibility is to broadcast general queries periodically, and to send group queries when receiving a leave message, to confirm that none of the clients on the port still want specific traffic before removing the traffic from the port. IGMP V2 lets clients specify what group (destination address) will receive the traffic but not to specify the source of the traffic. IGMP V3 is for source-specific multicast traffic, adding the capability for clients to INCLUDE or EXCLUDE specific traffic sources. An IGMP V3 device port state could be INCLUDE or EXCLUDE, and there are different types of group records for client reports.

The interfaces respond to general or group queries by sending a membership report that contains one or more of the following records associated with a specific group:

- Current-state record that indicates from which sources the interface wants to receive and not receive traffic. This record contains the source address of interfaces and whether or not traffic will be included (IS_IN) or not excluded (IS_EX) from this source.
- Filter-mode-change record. If the interface state changes from IS_IN to IS_EX, a TO_EX record is included in the membership report. Likewise, if the interface state changes from IS_EX to IS_IN, a TO_IN record appears in the membership report.
- An IGMP V2 leave report is equivalent to a TO_IN (empty) record in IGMP V3. This record means that no traffic from this group will be received regardless of the source.
- An IGMP V2 group report is equivalent to an IS_EX (empty) record in IGMP V3. This record means that all traffic from this group will be received regardless of source.
- Source-list-change record. If the interface wants to add or remove traffic sources from its membership report, the report can contain an ALLOW record, which includes a list of new sources from which the interface wishes to receive traffic. It can also contain a BLOCK record, which lists the current traffic sources from which the interface wants to stop receiving traffic.

IGMP protocols provide a method for clients and a device to exchange messages, and let the device build a database indicating which port wants what traffic. The protocols do not specify forwarding methods. They require IGMP snooping or multicast protocols such as PIM to handle packet forwarding. PIM can route multicast packets within and outside a VLAN, while IGMP snooping can switch packets only within a VLAN.

If a VLAN is not IGMP snooping-enabled, it floods multicast data and control packets to the entire VLAN in hardware. When snooping is enabled, IGMP packets are trapped to the CPU. Data packets are mirrored to the CPU in addition to being VLAN flooded. The CPU then installs hardware resources, so that subsequent data packets can be switched to desired ports in hardware without going to the CPU. If there is no client report or port to queriers for a data stream, the hardware resource drops it.

On PowerConnect B-Series T124X devices, the hardware can either match the group address only (* G), or both the source and group (S G) of the data stream. If any IGMPv3 is configured in any port of a VLAN, this VLAN uses (S G) match; otherwise, it uses (* G). This is 32-bit IP address matching, not 23-bit multicast MAC address 01-00-5e-xx-xx-xx matching.

PowerConnect B-Series T124X devices support up to 2K of IGMP groups, which are produced by client membership reports.

IGMP V1, V2, and V3 snooping support

Table 73 shows IGMP snooping version support by software release on PowerConnect B-Series T124X devices.

TABLE 73 IGMP snooping support

Software release	IGMP version support	Supported in software code...
PowerConnect B-Series T124X devices	IGMP V1 snooping IGMP V2 snooping IGMP V3 snooping	L2

Queriers and non-queriers

An IGMP snooping-enabled device can be configured as a querier (active) or non-querier (passive). An IGMP querier sends queries; a non-querier listens for IGMP queries and forwards them to the entire VLAN. VLANs can be independently configured to be queriers or non-queriers. If a VLAN has a connection to a PIM-enabled port on another router, the VLAN should be configured as a non-querier. When multiple IGMP snooping devices are connected together, and there is no connection to a PIM-enabled port, one of the devices should be configured as a querier. If multiple devices are configured as queriers, after these devices exchange queries, then all except the winner stop sending queries. The device with the lowest address becomes the querier. Although the system will work when multiple devices are configured as queriers, Dell recommends that only one device (preferably the one with the traffic source) is configured as a querier.

The non-queriers always forward multicast data traffic and IGMP messages to router ports which receive IGMP queries or PIM hellos. Dell recommends that you configure the device with the data traffic source (server) as a querier. If a server is attached to a non-querier, the non-querier always forwards traffic to the querier regardless of whether there are any clients on the querier.

NOTE

In a topology of one or more connecting devices, at least one device must be running PIM, or configured as active. Otherwise, none of the devices can send out queries, and traffic cannot be forwarded to clients.

IGMP snooping enhancements

This section describes the enhancements to IGMP snooping . These features are also supported on PowerConnect B-Series TI24X devices, except where noted.

Support for IGMP V3 snooping

Refer to [“IGMP snooping overview”](#) on page 439.

VLAN-specific configuration

IGMP snooping can be enabled on some VLANs or on all VLANs. Each VLAN can be independently configured to be a querier or non-querier and can be configured for IGMP V2 or IGMP V3. In general, the **ip multicast** commands apply globally to all VLANs except those configured with VLAN-specific **multicast** commands. The VLAN-specific multicast commands supersede the global **ip multicast** commands.

IGMP snooping can be configured for IGMP V2 or IGMP V3 on individual ports of a VLAN. An interface or router sends the queries and reports that include its IGMP version specified on it. The version configuration only applies to sending queries. The snooping device recognizes and processes IGMP V2 and IGMP V3 packets regardless of the version configuration.

To avoid version deadlock, an interface retains its version configuration even when it receives a report with a lower version.

Tracking and fast leave

Devices support fast leave for IGMP V2, and tracking and fast leave for IGMP V3. Fast leave stops the traffic immediately when the port receives a leave message. Tracking traces all IGMP V3 clients. Refer to [“IGMP V3 membership tracking and fast leave”](#) on page 452 and [“Fast leave for IGMP V2”](#) on page 452.

Configuration notes and feature limitations for PowerConnect B-Series TI24X devices

The following details apply to PowerConnect B-Series TI24X devices:

- Servers (traffic sources) are not required to send IGMP memberships.
- The default IGMP version is V2.
- Hardware resource is installed only when there is data traffic. If a VLAN is configured for IGMPv3, the hardware matches (S G), otherwise it matches (* G).
- A user can configure the maximum numbers of groups and hardware switched data streams.
- The device supports static groups that apply to the entire VLAN, or to just a few ports. The device acts as a proxy to send IGMP reports for the static groups when receiving queries.

- A user can configure static router ports to force all multicast traffic to these specific ports.
- Fast leave for IGMPv2 is supported. Fast leave stops traffic immediately when the port receives a leave message.
- Tracking and fast leave for IGMPv3 is supported. If the only client on a port leaves, traffic is stopped immediately.
- An IGMP device can be configured as a querier (active) or non-querier (passive). Queriers send queries. Non-queriers listen for queries and forward them to the entire VLAN.
- Every VLAN can be independently configured to be a querier or a non-querier.
- If a VLAN has a connection to a PIM-enabled port on another router, this VLAN should be configured as a non-querier (passive). When multiple snooping devices connect together and there is no connection to PIM ports, one device should be configured as a querier (active). If multiple devices are configured as active (queriers), only one will keep sending queries after exchanging queries.
- An IGMP device can be configured to rate-limit the forwarding IGMPv2 membership reports to queriers.
- The querier must configure an IP address to send out queries.

The implementation allows snooping on some VLANs or all VLANs. Each VLAN can independently enable or disable IGMP, or configure V2 or V3. In general, global configuration commands **ip multicast** apply to every VLAN except those that have local **multicast** configurations (which supersede the global configuration). IGMP also allows independent configuration of individual ports in a VLAN for either IGMPv2 or IGMPv3. Configuring a specific version on a port or a VLAN only applies to the queries sent by the device. The device always processes client reports of any version regardless of the configured version.

IGMP snooping requires hardware resources. If resources are inadequate, the data stream without a resource is mirrored to CPU in addition to being VLAN flooded, which can cause high CPU usage. Dell recommends that you avoid global enabling of snooping unless necessary.

When any port in a VLAN is configured for IGMPv3, the VLAN matches both source and group (S G) in hardware switching. If no ports are configured for IGMPv3, the VLAN matches group only (* G). Matching (S G) requires more hardware resources than matching (* G) when there are multiple servers sharing the same group. For example, two data streams from different sources to the same group require two (S G) entries in IGMPv3, but only one (* G) in IGMPv2. To conserve resources, IGMPv3 should be used only in source-specific applications. When VLANs are independently configured for versions, some VLANs can match (* G) while others match (S G).

IGMP snooping requires clients to send membership reports in order to receive data traffic. If a client application does not send reports, you must configure static groups to force traffic to client ports. A static group can apply to only some ports or to the entire VLAN.

PIM SM traffic snooping overview

When multiple PIM sparse routers connect through a snooping-enabled device, the device always forwards multicast traffic to these routers. For example, PIM sparse routers R1, R2 and R3 connect through a device. Assume R2 needs traffic, and R1 sends it to the device, which forwards it to both R2 and R3, even though R3 does not need it. A PIM SM snooping-enabled device listens to join and prune messages exchanged by PIM sparse routers, and stops traffic to the router that sends prune messages. This allows the device to forward the data stream to R2 only.

PIM SM traffic snooping requires IGMP snooping to be enabled on the device. IGMP snooping configures the device to listen for IGMP messages. PIM SM traffic snooping provides a finer level of multicast traffic control by configuring the device to listen specifically for PIM SM join and prune messages sent from one PIM SM router to another through the device.

PIM SM snooping support

Table 74 shows PIM SM snooping version support by PowerConnect B-Series T124X devices.

TABLE 74 PIM SM snooping support

Version support	Supported in software code...
PIM SM V2 snooping	L2, L3

Application examples

Figure 84 shows an example application of the PIM SM traffic snooping feature. In this example, a device is connected through an IP router to a PIM SM group source that is sending traffic for two PIM SM groups. The device also is connected to a receiver for each of the groups.

When PIM SM traffic snooping is enabled, the device starts listening for PIM SM join and prune messages and IGMP group membership reports. Until the device receives a PIM SM join message or an IGMP group membership report, the device forwards IP multicast traffic out all ports. Once the device receives a join message or group membership report for a group, the device forwards subsequent traffic for that group only on the ports from which the join messages or IGMP reports were received.

In this example, the router connected to the receiver for group 239.255.162.1 sends a join message toward the group source. Since PIM SM traffic snooping is enabled on the device, the device examines the join message to learn the group ID, then makes a forwarding entry for the group ID and the port connected to the receiver router. The next time the device receives traffic for 239.255.162.1 from the group source, the device forwards the traffic only on port 1, since that is the only port connected to a receiver for the group.

Notice that the receiver for group 239.255.162.69 is directly connected to the device. As result, the device does not see a join message on behalf of the client. However, since IGMP snooping also is enabled, the device uses the IGMP group membership report from the client to select the port for forwarding traffic to group 239.255.162.69 receivers.

The IGMP snooping feature and the PIM SM traffic snooping feature together build a list of groups and forwarding ports for the VLAN. The list includes PIM SM groups learned through join messages as well as MAC addresses learned through IGMP group membership reports. In this case, even though the device never sees a join message for the receiver for group 239.255.162.69, the device nonetheless learns about the receiver and forwards group traffic to the receiver.

The device stops forwarding IP multicast traffic on a port for a group if the port receives a prune message for the group.

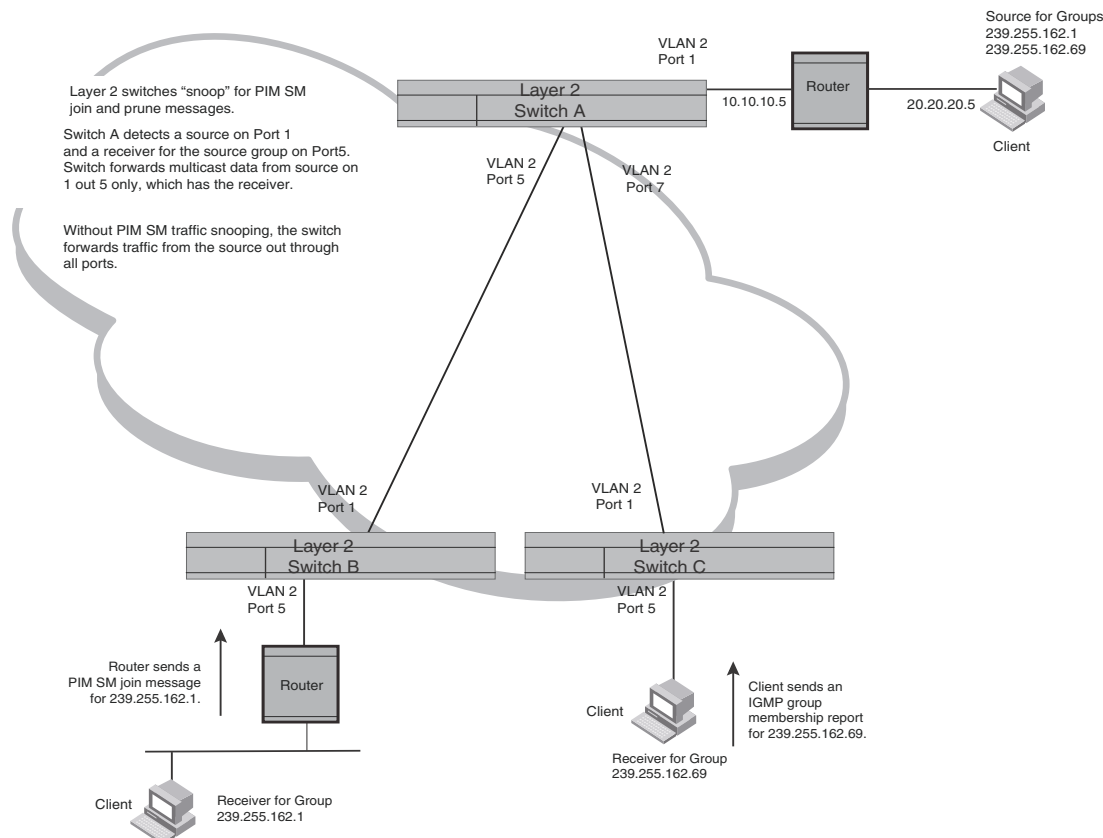
Notice that the ports connected to the source and the receivers are all in the same port-based VLAN on the device. This is required for the PIM SM snooping feature. The devices on the edge of the Global Ethernet cloud are configured for IGMP snooping and PIM SM traffic snooping. Although this application uses multiple devices, the feature has the same requirements and works the same way as it does on a single device.

The following figure shows another example application for PIM SM traffic snooping. This example shows devices on the edge of a Global Ethernet cloud (a Layer 2 Packet over SONET cloud). Assume that each device is attached to numerous other devices such as other Layer 2 Switches and Layer 3 Switches (routers).

NOTE

This example assumes that the devices are actually devices running Layer 2 Switch software.

FIGURE 84 PIM SM traffic reduction in Global Ethernet environment



The devices on the edge of the Global Ethernet cloud are configured for IGMP snooping and PIM SM traffic snooping. Although this application uses multiple devices, the feature has the same requirements and works the same way as it does on a single device.

Configuration notes and limitations

- PIM SM snooping applies only to PIM SM version 2 (PIM SM V2).
- PowerConnect B-Series TI24X devices support PIM SM traffic snooping in the Layer 2 code.
- IGMP snooping must be enabled on the device that will be running PIM SM snooping. The PIM SM traffic snooping feature requires IGMP snooping.

NOTE

Use the passive mode of IGMP snooping instead of the active mode. The passive mode assumes that a router is sending group membership queries as well as join and prune messages on behalf of receivers. The active mode configures the device to send group membership queries.

- All the device ports connected to the source and receivers or routers must be in the same port-based VLAN.
- The PIM SM snooping feature assumes that the group source and the device are in different subnets and communicate through a router. The source must be in a different IP subnet than the receivers. A PIM SM router sends PIM join and prune messages on behalf of a multicast group receiver only when the router and the source are in different subnet. When the receiver and source are in the same subnet, they do not need the router in order to find one another. They find one another directly within the subnet.

The device forwards all IP multicast traffic by default. Once you enable IGMP snooping and PIM SM traffic snooping, the device initially blocks all PIM SM traffic instead of forwarding it. The device forwards PIM SM traffic to a receiver only when the device receives a join message from the receiver. Consequently, if the source and the downstream router are in the same subnet, and PIM SM traffic snooping is enabled, the device blocks the PIM SM traffic and never starts forwarding the traffic. This is because the device never receives a join message from the downstream router for the group. The downstream router and group find each other without a join message because they are in the same subnet.

NOTE

If the “route-only” feature is enabled on a Layer 3 Switch, PIM SM traffic snooping will not be supported.

Configuring IGMP snooping

Configuring IGMP snooping on a device consists of the following global, VLAN-specific, and port-specific tasks:

Global tasks

Perform the following global tasks:

- [“Configuring the IGMP V3 snooping software resource limits”](#)
- [“Enabling IGMP snooping globally on the device”](#)
- [“Configuring the global IGMP mode”](#)
- [“Configuring the global IGMP version”](#)
- [“Modifying the age interval for group membership entries”](#)
- [“Modifying the query interval \(active IGMP snooping mode only\)”](#)
- [“Modifying the maximum response time”](#)
- [“Configuring report control” \(rate limiting\)](#)
- [“Modifying the wait time before stopping traffic when receiving a leave message”](#)
- [“Modifying the multicast cache age time”](#)
- [“Enabling or disabling error and warning messages”](#)

VLAN-specific tasks

Perform the following VLAN-specific tasks:

- “Configuring the IGMP mode for a VLAN” (active or passive)
- “Disabling IGMP snooping on a VLAN”
- “Configuring the IGMP version for a VLAN”
- “Configuring static router ports.”
- “Turning off static group proxy”
- “IGMP V3 membership tracking and fast leave”
- “Fast leave for IGMP V2”
- “Fast convergence”

Port-specific tasks

Perform the following port-specific tasks:

- “Disabling transmission and receipt of IGMP packets on a port”
- “Configuring the IGMP version for individual ports in a VLAN”

Configuring the IGMP V3 snooping software resource limits

By default, PowerConnect B-Series TI24X devices by default support up to 512 IGMP snooping multicast cache (mcache) entries and a maximum of 1K IGMP group addresses. If necessary, you can change the default values using the procedures in this section.

About IGMP snooping mcache entries and group addresses

An IGMP snooping group address entry is created when an IGMP join message is received for a group. An IGMP snooping mcache entry is created when data traffic is received for that group. Each mcache entry represents one data stream. The egress port list for the mcache entry is obtained from the IGMP group address entry. If there is an existing IGMP group address entry when an mcache is created, data traffic for that multicast group is switched in hardware.

Changing the maximum number of supported IGMP snooping mcache entries

When IGMP snooping is enabled, by default, the system supports up to 512 IGMP snooping mcache entries. If necessary, you can change the maximum number of IGMP snooping cache entries supported on the device. To do so, enter a command such as the following.

```
PowerConnect(config)# system-max igmp-snoop-mcache 2000
```

Syntax: [no] system-max igmp-snoop-mcache <num>

where <num> is a value between 512 and 2048 on PowerConnect B-Series TI24X devices. The default value is 512.

Setting the maximum number of IGMP group addresses

When IGMP snooping is enabled, by default, PowerConnect B-Series TI24X devices support up to 4K of IGMP group addresses by default, and the configurable range is from 4096 to 8192. The configured number is the upper limit of an expandable database. Client memberships exceeding the group limit are not processed. Enter a command such as the following to define the maximum number of IGMP group addresses.

```
PowerConnect(config)# system-max igmp-max-group-addr 1600
```

Syntax: [no] system-max igmp-max-group-addr <num>

On PowerConnect B-Series TI24X devices, <num> is a value between 4096 and 8192 and the default is 4096.

Enabling IGMP snooping globally on the device

Use the procedures in this section to enable IGMP snooping on a global basis.

Configuration notes for layer devices

- If Layer 3 multicast routing is enabled on your system, do not attempt to enable Layer 2 IGMP snooping. Layer 2 IGMP snooping is automatically enabled with Layer 3 multicast routing.
- If the "route-only" feature is enabled on the Layer 3 Switch, then IP multicast traffic reduction will not be supported.
- IGMP snooping is not supported on the default VLAN of Layer 3 Switches.

Configuring the IGMP mode

You can configure active or passive IGMP modes on the device. The default mode is passive. If you specify an IGMP mode for a VLAN, it overrides the global setting.

- **Active** - When active IGMP mode is enabled, a device actively sends out IGMP queries to identify multicast groups on the network, and makes entries in the IGMP table based on the group membership reports it receives.

NOTE

Routers in the network generally handle this operation. Use the active IGMP mode only when the device is in a stand-alone Layer 2 Switched network with no external IP multicast router attachments. In this case, enable the active IGMP mode on only one of the devices and leave the other devices configured for passive IGMP mode.

- **Passive** - When passive IGMP mode is enabled, it forwards reports to the router ports which receive queries. IGMP snooping in the passive mode does not send queries. However, it forwards queries to the entire VLAN.

Configuring the global IGMP mode

To globally set the IGMP mode to active, enter the following command.

```
PowerConnect(config)# ip multicast active
```

Syntax: [no] ip multicast [active | passive]

If you do not enter either *active* or *passive*, the passive mode is assumed.

Configuring the IGMP mode for a VLAN

If you specify an IGMP mode for a VLAN, it overrides the global setting.

To set the IGMP mode for VLAN 20 to active, enter the following commands.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast active
```

Syntax: [no] multicast active | passive

Configuring the IGMP version

Configuring the global IGMP version

When you globally enable IGMP snooping, you can specify IGMP V2 or IGMP V3 for the device. The following command enables IGMP V3.

```
PowerConnect(config)# ip multicast version 3
```

Syntax: [no] ip multicast version 2|3

If you do not specify a version number, IGMP V2 is assumed.

Configuring the IGMP version for a VLAN

You can specify the IGMP version for a VLAN. For example, the following commands configure VLAN 20 to use IGMP V3.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast version 3
```

Syntax: [no] multicast version 2 | 3

If no IGMP version is specified, then the globally-configured IGMP version is used. If an IGMP version is specified for individual ports, those ports use that version, instead of the VLAN version.

Configuring the IGMP version for individual ports in a VLAN

You can specify the IGMP version for individual ports in a VLAN. For example, the following commands configure ports 4, 5, and 6 to use IGMP V3. The other ports either use the IGMP version specified with the multicast version command, or the globally-configured IGMP version.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast port-version 3 ethe 4 to 6
```

Syntax: [no] multicast port-version 2 | 3 <port-numbers>

Disabling IGMP snooping on a VLAN

When IGMP snooping is enabled globally, you can still disable it for a specific VLAN. For example, the following commands cause IGMP snooping to be disabled for VLAN 20. This setting overrides the global setting.

```
PowerConnect(config)# vlan 20
```

```
PowerConnect(config-vlan-20)# multicast disable-multicast-snoop
```

Syntax: [no] multicast disable-multicast-snoop

Disabling transmission and receipt of IGMP packets on a port

When a VLAN is snooping-enabled, all IGMP packets are trapped to the CPU without hardware VLAN flooding. The CPU can block IGMP packets to and from a multicast-disabled port, and does not add it to the output interfaces of hardware resources. This prevents the disabled port from receiving multicast traffic. However, if static groups to the entire VLAN are defined, the traffic from these groups is VLAN flooded, including to disabled ports. Traffic from disabled ports cannot be blocked in hardware, and is switched in the same way as traffic from enabled ports.

This command has no effect on a VLAN that is not snooping-enabled because all multicast traffic is VLAN flooded.

To disable transmission and receipt of IGMP packets on a port, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 3
PowerConnect(config-if-e10000-3)# ip-multicast-disable
```

The above commands disable IGMP snooping on port 3 but does not affect the state of IGMP on other ports.

Syntax: [no] ip-multicast-disable

Modifying the age interval for group membership entries

When the device receives a group membership report, it makes an entry for that group in the IGMP group table. The age interval specifies how long the entry can remain in the table before the device receives another group membership report. When multiple devices connect together, all devices should be configured for the same age interval, which should be at least twice the length of the query interval, so that missing one report won't stop traffic. Non-querier age intervals should be the same as the age interval of the querier.

To modify the age interval, enter a command such as the following.

```
PowerConnect(config)# ip multicast age-interval 280
```

Syntax: [no] ip multicast age-interval <interval>

The <interval> parameter specifies the aging time. You can specify a value from 20 - 7200 seconds. The default is 260 seconds.

Modifying the query interval (active IGMP snooping mode only)

If IP multicast traffic reduction is set to active mode, you can modify the query interval to specify how often the device sends group membership queries. When multiple queriers connect together, they should all be configured with the same query interval.

To modify the query interval, enter a command such as the following.

```
PowerConnect(config)# ip multicast query-interval 120
```

Syntax: [no] ip multicast query-interval <interval>

The <interval> parameter specifies the time between queries. You can specify a value from 10 - 3600 seconds. The default is 125 seconds.

Modifying the maximum response time

The maximum response time is the number of seconds that a client can wait before responding to a query sent by the switch. The default response time is 10 seconds maximum.

To change the maximum response time, enter a command such as the following

```
PowerConnect(config)# ip multicast max-response-time 5
```

Syntax: [no] ip multicast max-response-time <interval>

For <interval>, enter a value from 1 - 10 seconds. The default is 10 seconds.

Configuring report control

A device in passive mode forwards reports and leave messages from clients to the upstream router ports that are receiving queries.

You can configure report control to rate-limit report forwarding within the same group to no more than once every 10 seconds. This rate-limiting does not apply to the first report answering a group-specific query.

NOTE

This feature applies to IGMP V2 only. The leave messages are not rate limited.

IGMP V2 membership reports of the same group from different clients are considered to be the same and are rate-limited.

Use the following command to alleviate report storms from many clients answering the upstream router query.

```
PowerConnect(config)# ip multicast report-control
```

Syntax: [no] ip multicast report-control

The original command, **ip igmp-report-control**, has been renamed to **ip multicast report-control**. The original command is still accepted; however, it is renamed when you issue a **show configuration** command.

Modifying the wait time before stopping traffic when receiving a leave message

You can define the wait time before stopping traffic to a port when a leave message is received. The device sends group-specific queries once per second to ask if any client in the same port still needs this group. The value range is from 1 to 5, and the default is 2. Due to internal timer granularity, the actual wait time is between n and (n+1) seconds (n is the configured value).

```
PowerConnect(config)# ip multicast leave-wait-time 1
```

Syntax: [no] ip multicast leave-wait-time <num>

<num> is the number of seconds from 1 to 5. The default is 2 seconds.

Modifying the multicast cache age time

You can set the time for an mcache to age out when it does not receive traffic. The traffic is hardware switched. One minute before aging out an mcache, the device mirrors a packet of this mcache to CPU to reset the age. If no data traffic arrives within one minute, this mcache is deleted. A lower value quickly removes resources consumed by idle streams, but it mirrors packets to CPU often. A higher value is recommended only data streams are continually arriving.

```
PowerConnect(config)# ip multicast mcache-age 180
```

Syntax: [no] ip multicast mcache-age <num>

<num> is the number of seconds from 60 to 3600. The default is 120 seconds.

Enabling or disabling error and warning messages

The device prints error or warning messages when it runs out of software resources or when it receives packets with the wrong checksum or groups. These messages are rate-limited. You can turn off these messages by entering a command such as the following.

```
PowerConnect(config)# ip multicast verbose-off
```

Syntax: [no] ip multicast verbose-off

Configuring static router ports

The device forwards all multicast control and data packets to router ports which receive queries. Although router ports are learned, you can force multicast traffic to specified ports even though these ports never receive queries. To configure static router ports, enter commands such as the following.

```
PowerConnect(config)# vlan 70
PowerConnect(config-vlan-70)# multicast router-port e 4 to 5 e 8
```

Syntax: [no] multicast router-port <port-numbers>

Turning off static group proxy

If a device has been configured for static groups, it acts as a proxy and sends membership reports for the static groups when it receives general or group-specific queries. When a static group configuration is removed, it is deleted from the active group table immediately. However, leave messages are not sent to the querier, and the querier should age out the group. Proxy activity can be turned off. The default is on. To turn proxy activity off for VLAN 20, enter commands similar to the following.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast proxy-off
```

Syntax: [no] multicast proxy-off

IGMP V3 membership tracking and fast leave

IGMP V3 gives clients membership tracking and fast leave capability. In IGMP V2, only one client on an interface needs to respond to a router's queries. This can leave some clients invisible to the router, making it impossible to track the membership of all clients in a group. When a client leaves the group, the device sends group-specific queries to the interface to see if other clients on that interface need the data stream of the client who is leaving. If no client responds, the device waits a few seconds before it stops the traffic. You can configure the wait time using the **ip multicast leave-wait-time** command.

IGMP V3 requires every client to respond to queries, allowing the device to track all clients. When tracking is enabled, and an IGMP V3 client sends a leave message and there is no other client, the device immediately stops forwarding traffic to the interface. This feature requires the entire VLAN be configured for IGMP V3 with no IGMP V2 clients. If a client does not send a report during the specified group membership time (the default is 260 seconds), that client is removed from the tracking list.

Every group on a physical port keeps its own tracking record. However, it can only track group membership; it cannot track by (source, group). For example, Client A and Client B belong to group1 but each receives traffic streams from different sources. Client A receives a stream from (source_1, group1) and Client B receives a stream from (source_2, group1). The device still waits for the configured leave-wait-time before it stops the traffic because these two clients are in the same group. If the clients are in different groups, then the waiting period is not applied and traffic is stopped immediately.

Enabling IGMP V3 membership tracking and fast leave for the VLAN

To enable the tracking and fast leave feature for VLAN 20, enter the following commands.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast tracking
```

Syntax: [no] multicast tracking

The membership tracking and fast leave features are supported for IGMP V3 only. If any port or any client is not configured for IGMP V3, then the multicast tracking command is ignored.

Fast leave for IGMP V2

When a device receives an IGMP V2 leave message, it sends out multiple group-specific queries. If no other client replies within the waiting period, the device stops forwarding traffic. When fast-leave-v2 is configured, when the device receives a leave message, it immediately stops forwarding to that port. The device does not send group specific-queries. You must ensure that no snooping-enabled ports have multiple clients. When two devices connect together, the querier should not be configured for fast-leave-v2, since the port might have multiple clients through the non-querier. The number of queries, and the waiting period (in seconds) can be configured using the **ip multicast leave-wait-time** command. The default is 2 seconds.

Enabling fast leave for IGMP V2

To configure fast leave for IGMP V2, enter the following commands.

```
PowerConnect(config)# vlan 20
PowerConnect(config-vlan-20)# multicast fast-leave-v2
```

Syntax: [no] multicast fast-leave-v2

Fast convergence

In addition to sending periodic general queries, an active device sends general queries when it detects a new port. However, because the device does not recognize the other device's port up event, multicast traffic might still require up to the query-interval time to resume after a topology change. Fast convergence allows the device to listen to topology change events in Layer 2 protocols such as spanning tree, and then send general queries to shorten the convergence time.

If the Layer 2 protocol cannot detect a topology change, fast convergence may not work in some cases. For example, if the direct connection between two devices switches from one interface to another, the rapid spanning tree protocol (802.1w) considers this optimization, rather than a topology change. In this example, other devices will not receive topology change notifications, and will be unable to send queries to speed up the convergence. Fast convergence works well with the regular spanning tree protocol in this case.

Enabling Fast convergence

To enable fast-convergence, enter the following commands.

```
PowerConnect(config)# vlan 70
PowerConnect(config-vlan-70)# multicast fast-convergence
```

Syntax: multicast fast-convergence

Configuring PIM SM snooping

Configuring PIM SM snooping on a device consists of the following global and VLAN-specific tasks.

Global task

Perform the following global tasks:

- [“Enabling or disabling PIM SM snooping”](#)

VLAN-specific tasks

Perform the following VLAN-specific tasks:

- [“Enabling PIM SM snooping on a VLAN”](#)
- [“Disabling PIM SM snooping on a VLAN”](#)

Enabling or disabling PIM SM snooping

PIM SM snooping should be used only in topologies where multiple PIM sparse routers connect through a device. PIM SM snooping does not work on a PIM dense mode router which does not send join messages, and traffic to PIM dense ports is stopped. A PIM SM snooping-enabled device displays a warning if it receives PIM dense join/prune messages.

To enable PIM sparse snooping globally, enter a command such as the following.

```
PowerConnect(config)# ip pimsm-snooping
This command enables PIM SM traffic snooping. The PIM SM traffic snooping feature
assumes that the network has routers that are running PIM SM.
```

NOTE

The device must be in passive mode before it can be configured for PIM SM snooping.

To disable the feature, enter the following command.

```
PowerConnect(config)# no ip pimsm-snooping
```

If you also want to disable IP multicast traffic reduction, enter the following command.

```
PowerConnect(config)# no ip multicast
```

Syntax: [no] ip pimsm-snooping

Enabling PIM SM snooping on a VLAN

You can enable PIM SM snooping for a specific VLAN. For example, the following commands enable PIM SM snooping on VLAN 20.

```
PowerConnect(config)# vlan 20
```

```
PowerConnect(config-vlan-20)# multicast pimsm-snooping
```

Syntax: [no] multicast pimsm-snooping

Disabling PIM SM snooping on a VLAN

When PIM SM snooping is enabled globally, you can still disable it for a specific VLAN. For example, the following commands disable PIM SM snooping for VLAN 20. This setting overrides the global setting.

```
PowerConnect(config)# vlan 20
```

```
PowerConnect(config-vlan-20)# multicast disable-pimsm-snoop
```

Syntax: [no] multicast disable-pimsm-snoop

IGMP snooping show commands

This section shows how to display information about IGMP snooping, including:

- [“Displaying the IGMP snooping configuration”](#)
- [“Displaying IGMP snooping errors”](#)
- [“Displaying IGMP group information”](#)
- [“Displaying IGMP snooping mcache information”](#)
- [“Displaying software resource usage for VLANs”](#)
- [“Displaying the status of IGMP snooping traffic”](#)

Displaying the IGMP snooping configuration

To display the global IGMP snooping configuration, enter the following command at any level of the CLI.

```
PowerConnect# show ip multicast
Summary of all vlans. Please use "sh ip mu vlan <vlan-id>" for details
Version=2, Intervals: Query=125, Group Age=260, Max Resp=10, Other Qr=260
VL10: cfg V3, vlan cfg passive, , pimsm (vlan cfg), 1 grp, 0 (SG) cache, no rtr
port
```

To display the IGMP snooping information for a specific VLAN , enter a command such as the following.

```
PowerConnect# show ip multicast vlan 10
Version=3, Intervals: Query=10, Group Age=260, Max Resp=10, Other Qr=30
VL3:dft V2,glb cfg passive,port down,, pimsm(glb cfg), 0 (*G) cache, no rtr port
e2      has      3 groups, non-QR (passive), default V3
**** Warning! has V2 client (life=240),
      group: 239.0.0.3, life = 240
      group: 224.1.1.2, life = 240
      group: 224.1.1.1, life = 240

e4      has      0 groups, non-QR (passive), default V3
```

Syntax: `show ip multicast vlan [<vlan-id>]`

If you do not specify a <vlan-id>, information for all VLANs is displayed.

This display shows the following information.

Table 0.2:

This field...	Displays...
Version	The global IGMP version. In this example, the device is configured for IGMP version 2.
Query	How often a querier sends a general query on the interface. In this example, the general queries are sent every 125 seconds.
Group Age	The number of seconds membership groups can be members of this group before aging out.
Max Resp	The maximum number of seconds a client waits before replying to a query.
Other Qr	How long it took a switch with a lower IP address to become a new querier. This value is 2 x Query + Max Resp.
cfg	The IGMP version for the specified VLAN. In this example, VL10: cfg V3 indicates that VLAN 10 is configured for IGMP V3.
vlan cfg	The IGMP configuration mode, which is either passive or active.
dft	The default config on this VLAN
glb cfg	The global configuration mode. This mode is selected when the VLAN does not have any specific config for PIM SM.
pimsm	Indicates that PIM SM is enabled on the VLAN.
rtr port	The router ports, which are the ports receiving queries.

Displaying IGMP snooping errors

To display information about possible IGMP errors, enter the following command.

```
PowerConnect# show ip multicast error
snoop SW processed pkt: 173, up-time 160 sec
```

Syntax: show ip multicast error

The following table describes the output from the **show ip multicast error** command.

Table 0.3:

This field	Displays
SW processed pkt	The number of multicast packets processed by IGMP snooping.
up-time	The time since the IGMP snooping is enabled.

Displaying IGMP group information

To display information about IGMP groups, enter the following command.

```
PowerConnect# show ip multicast group
p-:physical, ST:static, QR:querier, EX:exclude, IN:include, Y:yes, N:no
VL70 : 3 groups, 4 group-port, tracking_enabled
   group      p-port  ST    QR    life mode  source
1    224.1.1.2    33   no   yes   120 EX    0
2    224.1.1.1    33   no   yes   120 EX    0
3    226.1.1.1    35   yes  yes   100 EX    0
4    226.1.1.1    33   yes  yes   100 EX    0
```

In this example, an IGMP V2 group is in EXCLUDE mode with a source of 0. The group only excludes traffic from the 0 (zero) source list, which actually means that all traffic sources are included.

To display detailed IGMP group information for a specific group, enter the following command.

```
PowerConnect# show ip multicast group 226.1.1.1 detail
Display group 226.1.1.1 in all interfaces in details.
p-:physical, ST:static, QR:querier, EX:exclude, IN:include, Y:yes, N:no
VL70 : 1 groups, 2 group-port, tracking_enabled
   group      p-port  ST    QR    life mode  source
1    226.1.1.1    35   yes  yes   120 EX    0
   group: 226.1.1.1, EX, permit 0 (source, life):
   life=120, deny 0:
   group      p-port  ST    QR    life mode  source
2    226.1.1.1    33   yes  yes   120 EX    0
   group: 226.1.1.1, EX, permit 0 (source, life):
   life=120, deny 0:
```

If the tracking and fast leave features are enabled, you can display the list of clients that belong to a particular group by entering the following command.

```
PowerConnect# show ip multicast group 224.1.1.1 tracking
Display group 224.1.1.1 in all interfaces with tracking enabled.
p-:physical, ST:static, QR:querier, EX:exclude, IN:include, Y:yes, N:no
VL70 : 1 groups, 1 group-port, tracking_enabled
   group      p-port  ST    QR    life mode  source
*** Note: has 1 static groups to the entire vlan, not displayed here
1    224.1.1.1    33   no   yes   100 EX    0
   receive reports from 1 clients: (age)
   (2.2.100.2 60)
```

Syntax: show ip multicast group [*<group-address>*] [*[detail]*] [*[tracking]*]

If you want a report for a specific multicast group, enter that group's address for `<group-address>`.

Enter detail to display the source list of a specific VLAN.

Enter tracking for information on interfaces that have tracking enabled.

The following table describes the information displayed by the **show ip multicast group** command.

Table 0.4:

This field...	Displays...
group	The address of the group (destination address in this case, 224.1.1.1)
p-port	The physical port on which the group membership was received.
ST	Yes indicates that the IGMP group was configured as a static group; No means the address was learned from reports.
QR	Yes means the port is a querier port; No means it is not. A port becomes a non-querier port when it receives a query from a source with a lower source IP address than the device.
life	The number of seconds the group can remain in EXCLUDE mode. An EXCLUDE mode changes to INCLUDE mode if it does not receive an "IS_EX" or "TO_EX" message during a certain period of time. The default is 260 seconds. There is no life displayed in INCLUDE mode.
mode	Indicates current mode of the interface: INCLUDE or EXCLUDE. If the interface is in INCLUDE mode, it admits traffic only from the source list. If an interface is in EXCLUDE mode, it denies traffic from the source list and accepts the rest.
source	Identifies the source list that will be included or excluded on the interface. For example, if an IGMP V2 group is in EXCLUDE mode with a source of 0, the group excludes traffic from the 0 (zero) source list, which actually means that all traffic sources are included.

Displaying IGMP snooping mcache information

The IGMP snooping mcache contains multicast forwarding information for VLANs. To display information in the multicast forwarding mcache, enter the following command.

```
PowerConnect# show ip multicast mcache
Example: (S G) cnt=: cnt is number of SW processed packets
        OIF: e22 TR(32,33), TR is trunk, e32 primary, e33 output
vlan 10, 1 caches. use 1 VIDX
1      (10.10.10.2 239.0.0.3) cnt=0
        OIF: tag e2
        age=2s up-time=2s change=2s vidx=8191 (ref-cnt=1)
```

Syntax: show ip multicast mcache

The following table describes the output of the **show ip multicast mcache** command.

Table 0.5:

This field...	Displays...
(source group)	Source and group addresses of this data stream. (* group) means match group only; (source group) means match both.
cnt	The number of packets processed in software. Packets are switched in hardware, which increases this number slowly.
OIF	The output interfaces. If <code>entire vlan</code> is displayed, this indicates that static groups apply to the entire VLAN.

Table 0.5:

This field...	Displays...
age	The mcache age. The mcache will be reset to 0 if traffic continues to arrive, otherwise the mcache will be aged out when it reaches the time defined by the ip multicast mcache-age command.
uptime	The up time of this mcache in seconds.
vidx	Vidx specifies output port list index. Range is from 4096 to 8191
ref-cnt	The vidx is shared among mcaches having the same output interfaces. Ref-cnt indicates the number of mcaches using this vidx.

Displaying software resource usage for VLANs

To display information about the software resources used, enter the following command.

```
PowerConnect# show ip multicast resource
          alloc in-use  avail get-fail   limit  get-mem  size init
igmp group          256    1   255      0   32000    1   16  256
igmp phy port      1024    1  1023      0  200000    1   22 1024
... entries deleted ...
snoop mcache entry   128    2   126      0    8192    3   56  128
total pool memory 109056 bytes
has total 2 forwarding hash
VIDX sharing hash   : size=2      anchor=997  2nd-hash=no  fast-trav=no
Available vidx: 4060. IGMP/MLD use 2
```

Syntax: show ip multicast resource

The following table describes the output from the **show ip multicast resource** command.

Table 0.6:

This field...	Displays...
alloc	The allocated number of units.
in-use	The number of units which are currently being used.
avail	The number of available units.
get-fail	This displays the number of resource failures. NOTE: It is important to pay attention to this field.
limit	The upper limit of this expandable field. The limit of mcast group is configured by the system-max igmp-max-group-addr command. The limit of snoop mcache entry is configured by the system-max multicast-snoop-mcache command.
get-mem	The number of memory allocation. This number should continue to increase.
size	The size of a unit (in bytes).
init	The initial allocated amount of memory. More memory may be allocated if resources run out.
Available vidx	The output interface (OIF) port mask used by mcache. The entire device has a maximum of 4096 vidx. Different mcaches with the same OIF share the same vidx. If vidx is not available, the stream cannot be hardware-switched.

Displaying the status of IGMP snooping traffic

To display status information for IGMP snooping traffic, enter the following command.

```
PowerConnect# show ip multicast traffic
IGMP snooping: Total Recv: 22, Xmit: 26
Q: query, Qry: general Q, G-Qry: group Q, GSQry: group-source Q, Mbr: member
Recv      QryV2      QryV3      G-Qry      GSQry      MbrV2      MbrV3      Leave
VL1        0          0          0          0          4          0          0
VL70       18         0          0          0          0          0          0
Recv      IsIN       IsEX       ToIN       ToEX       ALLOW      BLOCK      Pkt-Err
VL1        0          4          0          0          0          0          0
VL70       0          0          0          0          0          0          0

Send      QryV2      QryV3      G-Qry      GSQry      MbrV2      MbrV3
VL1        0          0          8          0          0          0
VL70       0          0          0          0          0          18
VL70      pimsm-snooping, Hello: 12, Join/Prune: 9
```

Syntax: show ip multicast traffic

The following table describes the information displayed by the **show ip multicast traffic** command.

Table 0.7:

This field...	Displays...
Q	Query
Qry	General Query
QryV2	Number of general IGMP V2 queries received or sent.
QryV3	Number of general IGMP V3 queries received or sent.
G-Qry	Number of group-specific queries received or sent.
GSQry	Number of group source-specific queries received or sent.
Mbr	The membership report.
MbrV2	The IGMP V2 membership report.
MbrV3	The IGMP V3 membership report.
IsIN	Number of source addresses that were included in the traffic.
IsEX	Number of source addresses that were excluded in the traffic.
ToIN	Number of times the interface mode changed from EXCLUDE to INCLUDE.
ToEX	Number of times the interface mode changed from INCLUDE to EXCLUDE.
ALLO	Number of times that additional source addresses were allowed on the interface.
BLK	Number of times that sources were removed from an interface.
Pkt-Err	Number of packets having errors, such as checksum.
Pimsm-snooping hello, join, prune	Number of PIM sparse hello, join, and prune packets

PIM SM snooping show commands

This section shows how to display information about PIM SM snooping, including:

- “Displaying PIM SM snooping information”
- “Displaying PIM SM snooping information on a Layer 2 switch”
- “Displaying PIM SM snooping information for a specific group or source group pair”

Displaying PIM SM snooping information

To display PIM SM snooping information, enter the following command.

```
PowerConnect# show ip multicast pimsm-snooping
vlan 1, has 2 caches.
1    (* 230.1.1.1) has 1 pim join ports out of 1 OIF
    1 (age=60)
    1 has 1 src: 20.20.20.66(60)
2    (* 230.2.2.2) has 1 pim join ports out of 1 OIF
    1 (age=60)
    1 has 1 src: 20.20.20.66(60)
```

This output shows the number of PIM join outgoing interfaces (OIF) out of the total OIF. The join/prune messages are source-specific. In this case, if the mcache is in (* G), the display function will also print the traffic source information.

Syntax: `show ip multicast pimsm-snooping [<vlan-id>]`

Use the <vlan-id> parameter to display PIM SM snooping information for a specific VLAN.

Displaying PIM SM snooping information on a Layer 2 switch

You can display PIM SM snooping information for all groups by entering the following command at any level of the CLI on a Layer 2 Switch.

```
PowerConnect# show ip multicast pimsm-snooping vlan 100
VLAN ID 100, total 3 entries
PIMSM Neighbor list:
    1.100.100.12      : 3 expire 120 s
    1.100.100.10     : 3 expire 170 s
    1.100.100.7      : 3 expire 160 s
1    Group: 224.0.1.22, fid 08ac, NO cam
    Forwarding Port: 3
    PIMv2 Group Port: 3
    (Source, Port) list: 1 entries
2    Group: 239.255.162.2, fid 08aa, cam 8
    Forwarding Port: 1 2
    PIMv2 Group Port: 1 2
    (Source, Port) list: 3 entries
3    Group: 239.255.163.2, fid 08a9, cam 10
    Forwarding Port: 1 2
    PIMv2 Group Port: 1 2
    (Source, Port) list: 3 entries
VLAN ID 4008, total 0 entries
PIMSM Neighbor list:
```

Syntax: `show ip pimsm-snooping vlan <vlan-id>`

Enter the ID of the VLAN for the `vlan <vlan-id>` parameter.

If you want to display PIM SM snooping information for one source or one group, enter a command as in the following example. The command also displays the (source, port) list of the group.

```
PowerConnect# show ip pimsm-snooping 239.255.163.2
Show pimsm snooping group 239.255.163.2 in all vlan
VLAN ID 100
Group: 239.255.163.2, fid 08a9, cam 10
  Forwarding Port: 1 2
  PIMv2 Group Port: 1 2
  (Source, Port) list: 3 entries
    1   192.168.176.44, age=0, port: 2
    2   158.158.158.158, age=0, port: 1
    3   1.1.7.1, age=0, port: 2
```

Syntax: `show ip pimsm-snooping <group-address> | <source-address>`

If the address you entered is within the range of source addresses, then the router treats it as the source address. Likewise, if the address falls in the range of group addresses, then the router assumes that you are requesting a report for that group.

This display shows the following information.

Table 0.8:

This field...	Displays...
VLAN ID	The port-based VLAN to which the information listed below apply and the number of members in the VLAN.
PIM SM Neighbor list	The PIM SM routers that are attached to the Layer 2 Switch ports. The value following “expires” indicates how many seconds the Layer 2 Switch will wait for a hello message from the neighbor before determining that the neighbor is no longer present and removing the neighbor from the list.
Multicast Group	The IP address of the multicast group. NOTE: The fid and camindex values are used by Dell Technical Support for troubleshooting.
Forwarding Port	The ports attached to the group receivers. A port is listed here when it receives a join message for the group, an IGMP membership report for the group, or both.
PIMv2 Group Port	The ports on which the Layer 2 Switch has received PIM SM join messages for the group.
Source, Port list	The IP address of each PIM SM source and the Layer 2 Switch ports connected to the receivers of the source.

Displaying PIM SM snooping information for a specific group or source group pair

To display PIM SM snooping information for a specific group, enter a command such as the following at any level of the CLI.

18 Clear commands for IGMP snooping

```
PowerConnect# show ip multicast pimsm-snooping 230.1.1.1
Show pimsm snooping group 230.1.1.1 in all vlans
vlan 10,has 2 caches.
1 (*230.1.1.1) has 1 pim join ports out of 1 OIF
  1(age=120)
  1 has 1 src:20.20.20.66(120)
```

To display PIM SM snooping information for a specific (source, group) pair, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip multicast pimsm-snooping 230.2.2.2 20.20.20.66
Show pimsm snooping source 20.20.20.66, group 230.2.2.2 in all vlans
vlan 10:(*230.2.2.2) has 1 pim join ports out of 2 OIF
  1(age=0)
  1 has 1 src:20.20.20.66(0)
```

Syntax: `show ip multicast pimsm-snooping <group-address> [<source-ip-address>]`

The device determines which address is the group address and which one is the source address based on the ranges that the address fall into. If the address is within the range of source addresses, then the router treats it as the source address. Likewise, if the address falls in the range of group addresses, then the router assumes it is a group address.

The output shows the following information.

Table 0.9:

This field...	Displays...
vlan	The VLAN membership ID of the source.
port	The port on which the source is sending traffic. In this example, the port number is 1.
age	The age of the port, in seconds.
src	The source address and age. The age (number of seconds) is indicated in brackets immediately following the source.

Clear commands for IGMP snooping

The clear IGMP snooping commands should be used only in troubleshooting conditions, or to recover from errors.

Clearing the IGMP mcache

To clear the mcache on all VLANs, enter the following command.

```
PowerConnect# clear ip multicast mcache
```

Syntax: `clear ip multicast mcache`

Clearing the mcache on a specific VLAN

To clear the mcache on a specific VLAN, enter the following command.

```
PowerConnect# clear ip multicast vlan 10 mcache
```

Syntax: `clear ip multicast vlan <vlan-id> mcache`

The `<vlan-id>` parameter specifies the specific VLAN to clear the cache.

Clearing traffic on a specific VLAN

To clear the traffic counters on a specific VLAN, enter the following command.

```
PowerConnect# clear ip multicast vlan 10 traffic
```

Syntax: `clear ip multicast vlan <vlan-id> traffic`

The `<vlan-id>` parameter specifies the specific VLAN on which to clear the traffic counters.

Clearing IGMP counters on VLANs

To clear IGMP snooping on error and traffic counters for all VLANs, enter the following command.

```
PowerConnect# clear ip multicast counters
```

Syntax: `clear ip multicast counters`

18 Clear commands for IGMP snooping

Configuring IP Multicast Protocols

This chapter describes how to configure Layer 3 Switches for Protocol Independent Multicast (PIM). Layer 3 Switches support the following IP multicast versions:

- Internet Group Management Protocol (IGMP) V1 and V2
- Internet Group Management Protocol (IGMP) V3
- PIM Dense mode (PIM DM) V1 (draft-ietf-pim-dm-05) and V2 (draft-ietf-pim-v2-dm-03)
- PIM Sparse mode (PIM SM) V2 (RFC 2362)

NOTE

Each multicast protocol uses IGMP. IGMP is automatically enabled on an interface when you configure PIM and is disabled on the interface if you disable PIM .

NOTE

This chapter applies only to IP multicast routing. To configure Layer 2 multicast features, refer to [Chapter 18, "Configuring IP Multicast Traffic Reduction for PowerConnect B-Series TI24X Switches"](#).

Overview of IP multicasting

Multicast protocols allow a group or channel to be accessed over different networks by multiple stations (clients) for the receipt and transmit of multicast data.

Distribution of stock quotes, video transmissions such as news services and remote classrooms, and video conferencing are all examples of applications that use multicast routing.

IPv4 multicast group addresses

In IPv4 Multicast, host groups are identified by Class D addresses, i.e., those with "1110" as their higher-order four bits. In Internet standard "dotted decimal" notation, these group addresses range from 224.0.0.0 to 239.255.255.255. However, the IANA IPv4 Multicast Address Registry (referencing RFC 3171) stipulates that the range 224.0.0.0 through 224.0.0.255 should not be used for regular multicasting applications.

"The range of addresses between 224.0.0.0 and 224.0.0.255, inclusive, is reserved for the use of routing protocols and other low-level topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Multicast routers should not forward any multicast datagram with destination addresses in this range, regardless of its TTL."

Mapping of IPv4 Multicast group addresses to Ethernet MAC addresses

The IANA owns a block of Ethernet MAC addresses for Multicast usage that are in the range 0100.5e00.0000 through 0100.5e7f.ffff. For a given IPv4 Multicast group, there is a simple way of obtaining the appropriate Ethernet Destination MAC address that must be used in Layer 2 encapsulation. This is defined in RFC 1112, as follows:

“An IP host group address is mapped to an Ethernet multicast address by placing the low-order 23-bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01-00-5E-00-00-00 (hex). Because there are 28 significant bits in an IP host group address, more than one host group address may map to the same Ethernet multicast address.”

NOTE

Since there are 5 bits in the IPv4 Group address that are not used in the mapping, there is a possibility for up to 32 IPv4 Multicast Groups to use the same Ethernet Destination MAC address. Taking this into account along with the reserved IPv4 Group address range, it is discouraged for applications to use IPv4 Multicast Group Addresses that may conflict with the reserved addresses at the Layer 2 level. This is because some devices may use just the Ethernet Destination MAC address to take actions on the packet.

Supported Layer 3 multicast routing protocols

Layer 3 Switches support two different multicast routing protocols— Protocol-Independent Multicast (PIM) protocol along with the Internet Group Membership Protocol (IGMP).

PIM is broadcast and pruning multicast protocols that deliver IP multicast datagrams. The protocols employ reverse path lookup check and pruning to allow source-specific multicast delivery trees to reach all group members. PIM build a different multicast tree for each source and destination host group.

NOTE

PIM can concurrently operate on different ports of a Layer 3 Switch.

Multicast terms

The following are commonly used terms in discussing multicast-capable routers. These terms are used throughout this chapter:

- **Node:** Refers to a router or Layer 3 Switch.
- **Root Node:** The node that initiates the tree building process. It is also the router that sends the multicast packets down the multicast delivery tree.
- **Upstream:** Represents the direction from which a router receives multicast data packets. An upstream router is a node that sends multicast packets.
- **Downstream:** Represents the direction to which a router forwards multicast data packets. A downstream router is a node that receives multicast packets from upstream transmissions.
- **Group Presence:** Means that a multicast group has been learned from one of the directly connected interfaces. Members of the multicast group are present on the router.
- **Intermediate nodes:** Routers that are in the path between source routers and leaf routers.

- **Leaf nodes:** Routers that do not have any downstream routers.
- **Multicast Tree:** A unique tree is built for each source group (S,G) pair. A multicast tree is comprised of a root node and one or more nodes that are leaf or intermediate nodes.

Changing global IP multicast parameters

The following configurable parameters apply to PIM-DM, PIM-SM:

- Maximum number of PIM groups – You can change the maximum number of groups of each type for which the software will allocate memory. By default, Layer 3 Switches support up to 1024 PIM groups.
- Internet Group Membership Protocol (IGMP) V1 and V2 parameters – You can change the query interval, group membership time, and maximum response time.
- Hardware forwarding of fragmented IP multicast packets – You can enable the Layer 3 Switch to forward all fragments of fragmented IP multicast packets in hardware.

Changing dynamic memory allocation for IP multicast groups

Layer 3 Switches support up to 1024 PIM groups by default. Memory for the groups is allocated dynamically as needed. For each protocol, previous releases support a maximum of 255 groups and 255 IGMP memberships.

NOTE

The number of interface groups you can configure for PIM is unlimited; therefore, the **system-max pim-max-int-group** commands that define their maximum table sizes have been removed.

The software allocates memory globally for each group, and also allocates memory separately for each interface IGMP membership in a multicast group. An interface becomes a member of a multicast group when the interface receives an IGMP group membership report. For example, if the Layer 3 Switch learns about one multicast group, global memory for one group is used. In addition, if three interfaces on the device receive IGMP group membership reports for the group, interface memory for three IGMP memberships also is used.

Since the same group can use multiple allocations of memory (one for the group itself and one for each interface membership in the group), you can increase the maximum number of IGMP memberships, up to 8192.

NOTE

The total for IGMP memberships applies to the device, not to individual interfaces. You can have up to 8192 IGMP memberships on all the individual interfaces, not up to 8192 IGMP memberships on each interface.

Defining the maximum number of PIM cache entries

The PIM cache system parameter defines the maximum number of repeated PIM traffic being sent from the same source address and being received by the same destination address. To define this maximum, enter a command such as the following.

```
PowerConnect(config)#system-max pim-mcache 999
```

Syntax: system-max pim-mcache <num>

The <num> parameter specifies the maximum number of multicast cache entries for PIM. Enter a number from 256 – 4096. The default is 1024.

Changing IGMP V1 and V2 parameters

IGMP allows devices to limit the multicast of IGMP packets to only those ports on the router that are identified as IP Multicast members. This section applies to devices that support IGMP versions 1 and 2.

The router actively sends out host queries to identify IP Multicast groups on the network, inserts the group information in an IGMP packet, and forwards the packet to IP Multicast neighbors.

The following IGMP V1 and V2 parameters apply to PIM :

- IGMP query interval – Specifies how often the Layer 3 Switch queries an interface for group membership. Possible values are 1 – 3600. The default is 60.
- IGMP group membership time – Specifies how many seconds an IP Multicast group can remain on a Layer 3 Switch interface in the absence of a group report. Possible values are 1 – 7200. The default is 60.
- IGMP maximum response time – Specifies how many seconds the Layer 3 Switch will wait for an IGMP response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 – 10. The default is 5.

To change these parameters, you must first enable IP multicast routing by entering the following CLI command at the global CLI level.

```
PowerConnect(config)#ip multicast-routing
```

Syntax: [no] ip multicast-routing

NOTE

You must enter the **ip multicast-routing** command before changing the global IP Multicast parameters. Otherwise, the changes do not take effect and the software uses the default values.

Modifying IGMP (V1 and V2) query interval period

The IGMP query interval period defines how often a router will query an interface for group membership. Possible values are 1-3, 600 seconds and the default value is 60 seconds.

To modify the default value for the IGMP (V1 and V2) query interval, enter the following.

```
PowerConnect(config)#ip igmp query 120
```

Syntax: ip igmp query-interval <1-3600>

Modifying IGMP (V1 and V2) membership time

Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 1 – 7200 seconds and the default value is 140 seconds.

To define an IGMP (V1 and V2) membership time of 240 seconds, enter the following.

```
PowerConnect(config)#ip igmp group-membership-time 240
```

Syntax: ip igmp group-membership-time <1-7200>

Modifying IGMP (V1 and V2) maximum response time

Maximum response time defines how long the Layer 3 Switch will wait for an IGMP (V1 and V2) response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 – 10. The default is 10.

To change the IGMP (V1 and V2) maximum response time, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#ip igmp max-response-time 8
```

Syntax: [no] ip igmp max-response-time <num>

The <num> parameter specifies the number of seconds and can be a value from 1 – 10. The default is 10.

NOTE

Adding an interface to a multicast group

You can manually add an interface to a multicast group. This is useful in the following cases:

- Hosts attached to the interface are unable to add themselves as members of the group using IGMP.
- There are no members for the group attached to the interface.

When you manually add an interface to a multicast group, the device forwards multicast packets for the group but does not itself accept packets for the group.

You can manually add a multicast group to individual ports only. If the port is a member of a virtual routing interface, you must add the ports to the group individually.

To manually add a port to a multicast group, enter a command such as the following at the configuration level for the port.

```
PowerConnect(config-if-1)#ip igmp static-group 224.2.2.2
```

This command adds port 1 to multicast group 224.2.2.2.

To add a port that is a member of a virtual routing interface to a multicast group, enter a command such as the following at the configuration level for the virtual routing interface.

```
PowerConnect(config-vif-1)#ip igmp static-group 224.2.2.2 ethernet 2
```

This command adds port 2 in virtual routing interface 1 to multicast group 224.2.2.2.

Syntax: [no] ip igmp static-group <ip-addr> [ethernet <portnum>]

The <ip-addr> parameter specifies the group number.

The **ethernet** <portnum> parameter specifies the port number. Use this parameter if the port is a member of a virtual routing interface, and you are entering this command at the configuration level for the virtual routing interface.

Manually added groups are included in the group information displayed by the following commands:

- **show ip igmp group**
- **show ip pim group**

PIM Dense

NOTE

This section describes the “dense” mode of PIM, described in RFC 1075. Refer to “[PIM Sparse](#)” on page 478 for information about PIM Sparse.

NOTE

On PowerConnect B-Series T124X devices, this feature is supported.

PIM was introduced to simplify some of the complexity of the routing protocol at the cost of additional overhead tied with a greater replication of forwarded multicast packets. PIM builds source-routed multicast delivery trees and employs reverse path check when forwarding multicast packets.

There are two modes in which PIM operates: Dense and Sparse. The Dense Mode is suitable for densely populated multicast groups, primarily in the LAN environment. The Sparse Mode is suitable for sparsely populated multicast groups with the focus on WAN.

PIM primarily using the IP routing table instead of maintaining its own, thereby being routing protocol independent.

Initiating PIM multicasts on a network

Once PIM is enabled on each router, a network user can begin a video conference multicast from the server on R1 as shown in [Figure 85](#). When a multicast packet is received on a PIM-capable router interface, the interface checks its IP routing table to determine whether the interface that received the message provides the shortest path back to the source. If the interface does provide the shortest path back to the source, the multicast packet is then forwarded to all neighboring PIM routers. Otherwise, the multicast packet is discarded and a prune message is sent back upstream.

In [Figure 85](#), the root node (R1) is forwarding multicast packets for group 229.225.0.1, which it receives from the server, to its downstream nodes, R2, R3, and R4. Router R4 is an intermediate router with R5 and R6 as its downstream routers. Because R5 and R6 have no downstream interfaces, they are leaf nodes. The receivers in this example are those workstations that are resident on routers R2, R3, and R6.

Pruning a multicast tree

As multicast packets reach these leaf routers, the routers check their IGMP databases for the group. If the group is not in a router IGMP database, the router discards the packet and sends a prune message to the upstream router. The router that discarded the packet also maintains the prune state for the source, group (S,G) pair. The branch is then pruned (removed) from the multicast tree. No further multicast packets for that specific (S,G) pair will be received from that upstream router until the prune state expires. You can configure the PIM Prune Timer (the length of time that a prune state is considered valid).

For example, in [Figure 85](#) the sender with address 207.95.5.1 is sending multicast packets to the group 229.225.0.1. If a PIM switch receives any groups other than that group, the switch discards the group and sends a prune message to the upstream PIM switch.

In Figure 86, switch S5 is a leaf node with no group members in its IGMP database. Therefore, the switch must be pruned from the multicast tree. S5 sends a prune message upstream to its neighbor switch S4 to remove itself from the multicast delivery tree and install a prune state, as seen in Figure 86. Switch S5 will not receive any further multicast traffic until the prune age interval expires.

When a node on the multicast delivery tree has all of its downstream branches (downstream interfaces) in the prune state, a prune message is sent upstream. In the case of S4, if both S5 and S6 are in a prune state at the same time, S4 becomes a leaf node with no downstream interfaces and sends a prune message to S1. With S4 in a prune state, the resulting multicast delivery tree would consist only of leaf nodes S2 and S3.

FIGURE 85 Transmission of multicast packets from the source to host group members

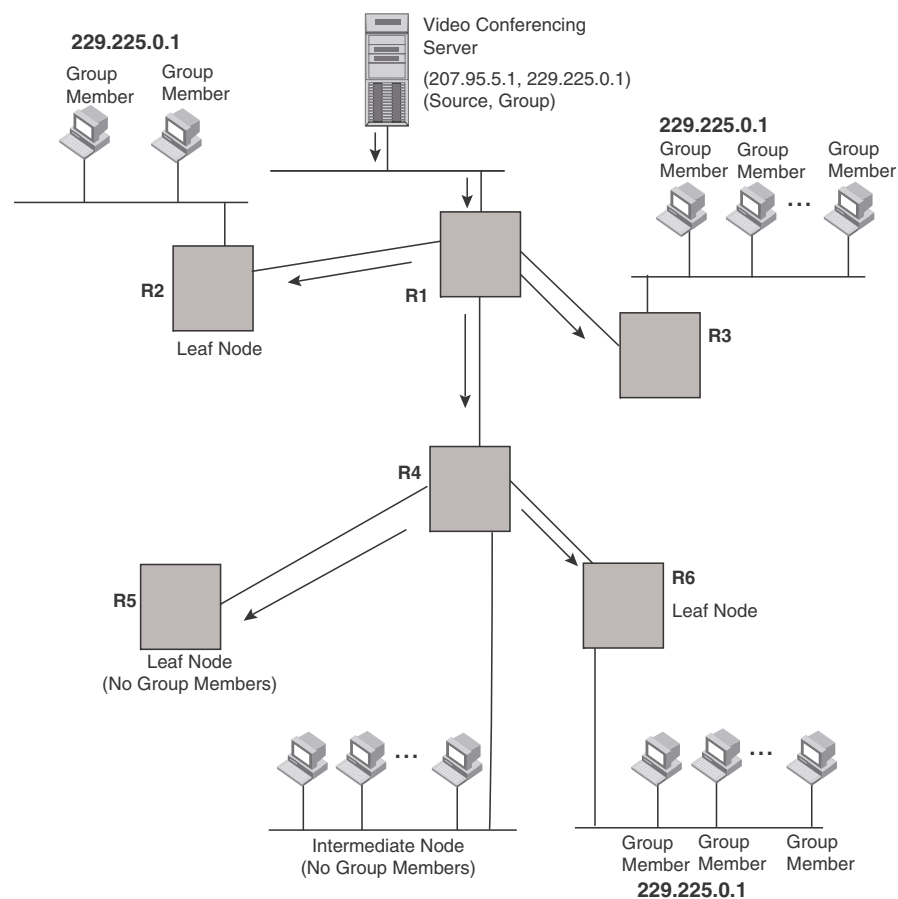
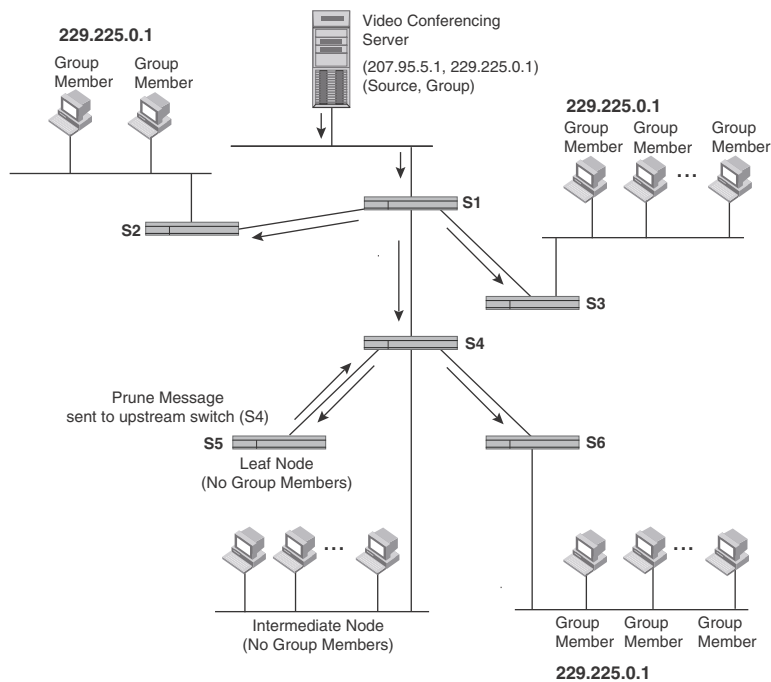


FIGURE 86 Pruning leaf nodes from a multicast tree



Grafts to a multicast Tree

A PIM switch restores pruned branches to a multicast tree by sending graft messages towards the upstream switch. Graft messages start at the leaf node and travel up the tree, first sending the message to its neighbor upstream switch.

In the example above, if a new 229.225.0.1 group member joins on switch S6, which was previously pruned, a graft is sent upstream to S4. Since the forwarding state for this entry is in a prune state, S4 sends a graft to S1. Once S4 has joined the tree, S4 and S6 once again receive multicast packets.

Prune and graft messages are continuously used to maintain the multicast delivery tree. No configuration is required on your part.

PIM DM versions

Devices support PIM DM V1 and V2. The default is V2. You can specify the version on an individual interface basis.

The primary difference between PIM DM V1 and V2 is the methods the protocols use for messaging:

- PIM DM V1 – uses the Internet Group Management Protocol (IGMP) to send messages
- PIM DM V2 – sends messages to the multicast address 224.0.0.13 (ALL-PIM-ROUTERS) with protocol number 103

The CLI commands for configuring and managing PIM DM are the same for V1 and V2. The only difference is the command you use to enable the protocol on an interface.

NOTE

Version 2 is the default PIM DM version. The only difference between version 1 and version 2 is the way the protocol sends messages. The change is not apparent in most configurations. You can use version 2 instead of version 1 with no impact to your network. However, if you want to continue to use PIM DM V1 on an interface, you must change the version, then save the configuration.

NOTE

The note above does not mean you can run different PIM versions on devices that are connected to each other. The devices must run the same version of PIM. If you want to connect a Layer 3 Switch running PIM to a device that is running PIM V1, you must change the version on the Layer 3 Switch to V1 (or change the version on the device to V2, if supported).

Configuring PIM DM

NOTE

This section describes how to configure the “dense” mode of PIM, described in RFC 1075. Refer to [“Configuring PIM Sparse”](#) on page 479 for information about configuring PIM Sparse.

Enabling PIM on the router and an interface

By default, PIM is disabled. To enable PIM, perform the following:

- Enable the feature globally.
- Configure the IP interfaces that will use PIM.
- Enable PIM locally on the ports that have the IP interfaces you configured for PIM.

Suppose you want to initiate the use of desktop video for fellow users on a sprawling campus network. All destination workstations have the appropriate hardware and software but the devices that connect the various buildings need to be configured to support PIM multicasts from the designated video conference server as shown in [Figure 85](#) on page 471.

PIM is enabled on each of the devices shown in [Figure 85](#), on which multicasts are expected. You can enable PIM on each router independently or remotely from one of the routers with a Telnet connection. Follow the same steps for each router. A reset of the router is required when PIM is first enabled. Thereafter, all changes are dynamic.

Globally enabling and disabling PIM

To globally enable PIM, enter the following command.

```
PowerConnect(config)#router pim
```

Syntax: [no] router pim

The behavior of the **[no] router pim** command is as follows:

- Entering **router pim** command to enable PIM does not require a software reload.
- Entering a **no router pim** command removes all configuration for PIM multicast on a Layer 3 Switch (**router pim** level) only.

Globally Enabling and Disabling PIM without Deleting Multicast Configuration

As stated above entering a **no router pim** command deletes the PIM configuration. If you want to disable PIM without deleting any PIM configuration, enter the following command.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#disable-pim
```

Syntax: [no] disable-pim

Use the [no] version of the command to re-enable PIM.

Enabling a PIM version

Using the CLI

To enable PIM on an interface, globally enable PIM, then enable PIM on interface 3, enter the following commands.

```
PowerConnect(config)#router pim
PowerConnect(config)#int e 3
PowerConnect(config-if-e10000-3)#ip address 207.95.5.1/24
PowerConnect(config-if-e1000-3)#ip pim
```

Syntax: [no] ip pim [version 1 | 2]

The **version 1 | 2** parameter specifies the PIM DM version. The default version is 2.

If you have enabled PIM version 1 but need to enable version 2 instead, enter either of the following commands at the configuration level for the interface.

```
PowerConnect(config-if-1)#ip pim version 2
PowerConnect(config-if-1)#no ip pim version 1
```

To disable PIM DM on the interface, enter the following command.

```
PowerConnect(config-if-1)#no ip pim
```

Modifying PIM global parameters

PIM global parameters come with preset values. The defaults work well in most networks, but you can modify the following parameters if you need to:

- Neighbor timeout
- Hello timer
- Prune timer
- Prune wait timer
- Graft retransmit timer
- Inactivity timer

Modifying neighbor timeout

Neighbor timeout is the interval after which a PIM router will consider a neighbor to be absent. Absence of PIM hello messages from a neighboring router indicates that a neighbor is not present.

The default value is 180 seconds.

To apply a PIM neighbor timeout value of 360 seconds to all ports on the router operating with PIM, enter the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#nbr-timeout 360
```

Syntax: nbr-timeout <60-8000>

The default is 180 seconds.

Modifying hello timer

This parameter defines the interval at which periodic hellos are sent out PIM interfaces. Routers use hello messages to inform neighboring routers of their presence. The default rate is 60 seconds.

To apply a PIM hello timer of 120 seconds to all ports on the router operating with PIM, enter the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#hello-timer 120
```

Syntax: `hello-timer <10-3600>`

The default is 60 seconds.

Modifying prune timer

This parameter defines how long a PIM router will maintain a prune state for a forwarding entry.

The first received multicast interface is forwarded to all other PIM interfaces on the router. If there is no presence of groups on that interface, the leaf node sends a prune message upstream and stores a prune state. This prune state travels up the tree and installs a prune state.

A prune state is maintained until the prune timer expires or a graft message is received for the forwarding entry. The default value is 180 seconds.

To set the PIM prune timer to 90, enter the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)##prune-timer 90
```

Syntax: `prune-timer <10-3600>`

The default is 180 seconds.

Modifying the prune wait timer

The CLI command **prune-wait** allows you to configure the amount of time a PIM router will wait before stopping traffic to neighbor routers that do not want the traffic. The value can be from zero to three seconds. The default is three seconds. A smaller prune wait value reduces flooding of unwanted traffic.

A prune wait value of zero causes the PIM router to stop traffic immediately upon receiving a prune message. If there are two or more neighbors on the physical port, then the **prune-wait** command should not be used because one neighbor may send a prune message while the other sends a join message at the during time or in less than three seconds.

To set the prune wait time to zero, enter the following commands.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#prune-wait 0
```

Syntax: `prune-wait <time>`

where `<time>` can be 0 - 3 seconds. A value of 0 causes the PIM router to stop traffic immediately upon receiving a prune message. The default is 3 seconds.

Viewing the prune wait time

To view the prune wait time, enter the **show ip pim dense** command at any level of the CLI.

```
PowerConnect# show ip pim dense
Global PIM Dense Mode Settings
Hello interval: 60, Neighbor timeout: 180
Graft Retransmit interval: 10, Inactivity interval: 180
Route Expire interval: 200, Route Discard interval: 340
Prune age: 180, Prune wait: 3
```

Modifying graft retransmit timer

The Graft Retransmit Timer defines the interval between the transmission of graft messages.

A graft message is sent by a router to cancel a prune state. When a router receives a graft message, the router responds with a Graft Ack (acknowledge) message. If this Graft Ack message is lost, the router that sent the graft message will resend it.

To change the graft retransmit timer from the default of 180 to 90 seconds, enter the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#graft-retransmit-timer 10
```

Syntax: `graft-retransmit-timer <2-10>`

The default is 3 seconds.

Modifying inactivity timer

The router deletes a forwarding entry if the entry is not used to send multicast packets. The PIM inactivity timer defines how long a forwarding entry can remain unused before the router deletes it.

To apply a PIM inactivity timer of 90 seconds to all PIM interfaces, enter the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#inactivity-timer 90
```

Syntax: `inactivity-timer <10-3600>`

The default is 180 seconds.

Selection of shortest path back to source

By default, when a multicast packet is received on a PIM-capable router interface in a multi-path topology, the interface checks its IP routing table to determine the shortest path back to the source. If the alternate paths have the same cost, the first alternate path in the table is picked as the path back to the source. For example, in the table below, the first four routes have the same cost back to the source. However, 137.80.127.3 will be chosen as the path to the source since it is the first one on the list. The router rejects traffic from any port other than Port V11 on which 137.80.127.3 resides.

```
Total number of IP routes: 19
B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
Destination      NetMask          Gateway          Port          Cost Type
..
9      172.17.41.4      255.255.255.252*137.80.127.3      v11           2      O
      172.17.41.4      255.255.255.252 137.80.126.3      v10           2      O
      172.17.41.4      255.255.255.252 137.80.129.1      v13           2      O
      172.17.41.4      255.255.255.252 137.80.128.3      v12           2      O
10     172.17.41.8      255.255.255.252 0.0.0.0          2            1      D
```

When the Highest IP RPF feature is enabled, the selection of the shortest path back to the source is based on which Reverse Path Forwarding (RPF) neighbor in the IP routing table has the highest IP address, if the cost of the routes are the same. For example, in the table above, Gateway 137.80.129.1 will be chosen as the shortest path to the source because it is the RPF neighbor with the highest IP address.

When choosing the RPF, the router first checks the Multicast Routing Table. If the table is not available, it chooses an RPF from the IP Routing Table. Multicast route is configured using the **ip mroute** command.

To enable the Highest IP RPF feature, enter commands such as the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#highest-ip-rpf
```

The command immediately enables the Highest IP RPF feature; there is no need to reboot the device.

Syntax: [no] highest-ip-rpf

Entering the **no** version of the command disables the feature; the shortest path back to the source will be based on the first entry in the IP routing table. If some PIM traffic paths were selected based on the highest IP RPF, these paths are changed immediately to use the first RPF in the routing table.

Failover time in a multi-path topology

When a port in a multi-path topology fails, and the failed port is the input port of the downstream router, a new path is re-established within a few seconds, depending on the routing protocol being used.

No configuration is required for this feature.

Modifying the TTL

The TTL defines the minimum value required in a packet for it to be forwarded out of the interface.

For example, if the TTL for an interface is set at 10, it means that only those packets with a TTL value of 10 or more will be forwarded. Likewise, if an interface is configured with a TTL Threshold value of 1, all packets received on that interface will be forwarded. Possible TTL values are 1 to 31. The default TTL value is 1.

Configuration notes

- If the TTL for an interface is greater than 1, PIM packets received on the interface are always forwarded in software because each packet TTL must be examined. Therefore, Dell does not recommend modifying the TTL under normal operating conditions.
- Multicast packets with a TTL value of 1 are switched within the same VLAN. These packets cannot be routed between different VLANs.

Configuration syntax

To configure a TTL of 24, enter the following.

```
PowerConnect(config-if-24)#ip pim ttl 24
```

Syntax: ip pim ttl <1-31>

Dropping PIM traffic in hardware

Unwanted PIM Dense or PIM Sparse multicast traffic can be dropped in hardware on Layer 3 Switches. . Refer to [“Passive multicast route insertion”](#) on page 501.

PIM Sparse

Devices support Protocol Independent Multicast (PIM) Sparse version 2. PIM Sparse provides multicasting that is especially suitable for widely distributed multicast environments. The Dell implementation is based on RFC 2362.

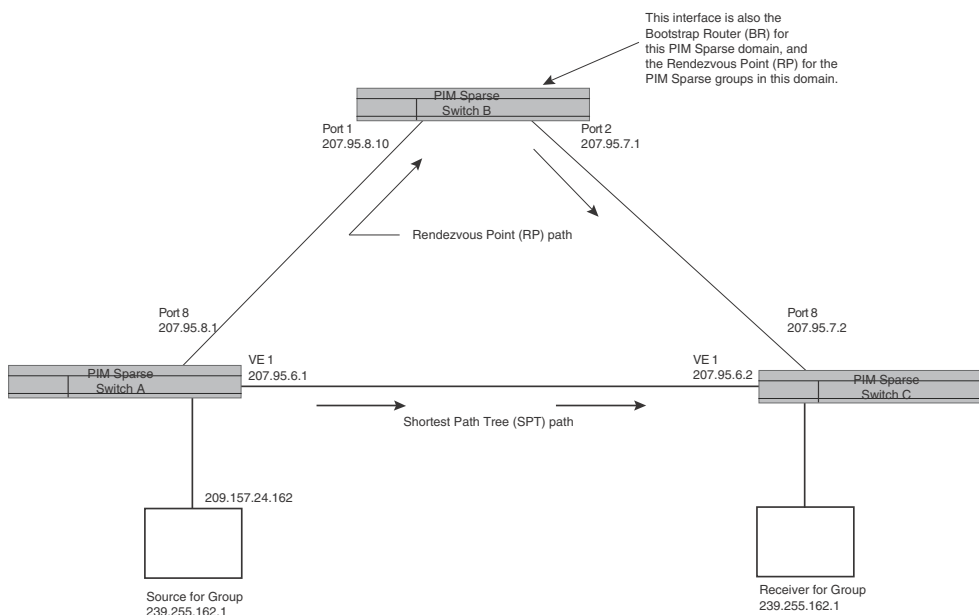
NOTE

On PowerConnect B-Series TI24X devices, this feature is supported.

In a PIM Sparse network, a PIM Sparse router that is connected to a host that wants to receive information for a multicast group must explicitly send a join request on behalf of the receiver (host).

PIM Sparse routers are organized into domains. A PIM Sparse domain is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary. [Figure 87](#) shows a simple example of a PIM Sparse domain. This example shows three Layer 3 Switches configured as PIM Sparse routers. The configuration is described in detail following the figure.

FIGURE 87 Example of a PIM Sparse domain



PIM Sparse switch types

Switches that are configured with PIM Sparse interfaces also can be configured to fill one or more of the following roles:

- **PMBR** – A PIM switch that has some interfaces within the PIM domain and other interface outside the PIM domain. PBMRs connect the PIM domain to the Internet.

NOTE

You cannot configure a routing interface as a PMBR interface for PIM Sparse in the current software release.

- **BSR** – The Bootstrap Router (BSR) distributes RP information to the other PIM Sparse switches within the domain. Each PIM Sparse domain has one active BSR. For redundancy, you can configure ports on multiple switches as candidate BSRs. The PIM Sparse protocol uses an election process to select one of the candidate BSRs as the BSR for the domain. The BSR with the highest BSR priority (a user-configurable parameter) is elected. If the priorities result in a tie, then the candidate BSR interface with the highest IP address is elected. In the example in [Figure 87](#), PIM Sparse switch B is the BSR. Port 2 is configured as a candidate BSR.
- **RP** – The RP is the meeting point for PIM Sparse sources and receivers. A PIM Sparse domain can have multiple RPs, but each PIM Sparse multicast group address can have only one active RP. PIM Sparse switches learn the addresses of RPs and the groups for which they are responsible from messages that the BSR sends to each of the PIM Sparse switches. In the example in [Figure 87](#), PIM Sparse Switch B is the RP. Port 2 is configured as a candidate Rendezvous Point (RP).

To enhance overall network performance, Layer 3 Switches use the RP to forward only the first packet from a group source to the group receivers. After the first packet, the Layer 3 Switch calculates the shortest path between the receiver and source (the Shortest Path Tree, or SPT) and uses the SPT for subsequent packets from the source to the receiver. The Layer 3 Switch calculates a separate SPT for each source-receiver pair.

NOTE

Dell recommends that you configure the same ports as candidate BSRs and RPs.

RP paths and SPT paths

[Figure 87](#) shows two paths for packets from the source for group 239.255.162.1 and a receiver for the group. The source is attached to PIM Sparse Switch A and the recipient is attached to PIM Sparse Switch C. PIM Sparse Switch B in is the RP for this multicast group. As a result, the default path for packets from the source to the receiver is through the RP. However, the path through the RP sometimes is not the shortest path. In this case, the shortest path between the source and the receiver is over the direct link between Switch A and Switch C, which bypasses the RP (Switch B).

To optimize PIM traffic, the protocol contains a mechanism for calculating the Shortest Path Tree (SPT) between a given source and receiver. PIM Sparse switches can use the SPT as an alternative to using the RP for forwarding traffic from a source to a receiver. By default, Layer 3 Switches forward the first packet they receive from a given source to a given receiver using the RP path, but forward subsequent packets from that source to that receiver through the SPT. In [Figure 87](#), Switch A forwards the first packet from group 239.255.162.1 source to the destination by sending the packet to Switch B, which is the RP. Switch B then sends the packet to Switch C. For the second and all future packets that Switch A receives from the source for the receiver, Switch A forwards them directly to Switch C using the SPT path.

Configuring PIM Sparse

To configure a Layer 3 Switch for PIM Sparse, perform the following tasks:

- Configure the following global parameter:
 - Enable the PIM Sparse mode of multicast routing.
- Configure the following interface parameters:
 - Configure an IP address on the interface
 - Enable PIM Sparse.
 - Identify the interface as a PIM Sparse border, if applicable.

NOTE

You cannot configure a routing interface as a PMBR interface for PIM Sparse in the current software release.

- Configure the following PIM Sparse global parameters:
 - Identify the Layer 3 Switch as a candidate PIM Sparse Bootstrap Router (BSR), if applicable.
 - Identify the Layer 3 Switch as a candidate PIM Sparse Rendezvous Point (RP), if applicable.
 - Specify the IP address of the RP (if you want to statically select the RP).

NOTE

Dell recommends that you configure the same Layer 3 Switch as both the BSR and the RP.

Limitations in this release

The implementation of PIM Sparse in the current software release has the following limitations:

- PIM Border Routers (PMBRs) are not supported. Thus, you cannot configure a routing interface as a PMBR interface for PIM Sparse.
- PIM Sparse and regular PIM (dense mode) cannot be used on the same interface.

Configuring Global PIM Sparse parameters

To configure the PIM Sparse global parameters, use either of the following methods.

To configure basic global PIM Sparse parameters, enter commands such as the following on each Layer 3 Switch within the PIM Sparse domain.

```
PowerConnect(config)#router pim
```

Syntax: [no] router pim

NOTE

You do not need to globally enable IP multicast routing when configuring PIM Sparse.

The command in this example enables IP multicast routing, and enables the PIM Sparse mode of IP multicast routing. The command does not configure the Layer 3 Switch as a candidate PIM Sparse Bootstrap Router (BSR) and candidate Rendezvous Point (RP). You can configure a Layer 3 Switch as a PIM Sparse switch without configuring the it as a candidate BSR and RP. However, if you do configure the Layer 3 Switch as one of these, Dell recommends that you configure it as both. Refer to [“Configuring BSRs”](#) on page 481.

The behavior of the **[no] router pim** command is as follows:

- Entering **no router pim** command to disable PIM does not require a software reload.
- Entering a **no router pim** command removes all configuration for PIM multicast on a Layer 3 Switch (**router pim** level) only.

Globally enabling and disabling PIM without deleting the multicast configuration

As stated above entering a **no router pim** command deletes the PIM configuration. If you want to disable PIM without deleting any PIM configuration, enter the following command.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#disable-pim
```

Syntax: [no] disable-pim

Use the [no] version of the command to re-enable PIM.

Configuring PIM interface parameters

After you enable IP multicast routing and PIM Sparse at the global level, you must enable it on the individual interfaces connected to the PIM Sparse network. To do so, use the following CLI method.

To enable PIM Sparse mode on an interface, enter commands such as the following.

```
PowerConnect(config)#interface ethernet 2
PowerConnect(config-if-2)#ip address 207.95.7.1 255.255.255.0
PowerConnect(config-if-2)#ip pim-sparse
```

Syntax: [no] ip pim-sparse

The commands in this example add an IP interface to port 2, then enable PIM Sparse on the interface.

If the interface is on the border of the PIM Sparse domain, you also must enter the following command.

```
PowerConnect(config-if-2)#ip pim border
```

Syntax: [no] ip pim border

NOTE

You cannot configure a routing interface as a PMBR interface for PIM Sparse in the current software release.

Configuring BSRs

In addition to the global and interface parameters in the sections above, you need to identify an interface on at least one Layer 3 Switch as a candidate PIM Sparse Bootstrap router (BSR) and candidate PIM Sparse Rendezvous Point (RP).

NOTE

It is possible to configure the Layer 3 Switch as only a candidate BSR or RP, but Dell recommends that you configure the same interface on the same Layer 3 Switch as both a BSR and an RP.

This section presents how to configure BSRs. Refer to [“Configuring RPs”](#) on page 482 for instructions on how to configure RPs.

To configure the Layer 3 Switch as a candidate BSR and RP, enter commands such as the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#bsr-candidate ethernet 2 30 255
BSR address: 207.95.7.1, hash mask length: 30, priority: 255
```

This command configures the PIM Sparse interface on port 2 as a BSR candidate, with a hash mask length of 30 and a priority of 255. The information shown in italics above is displayed by the CLI after you enter the candidate BSR configuration command.

Syntax: **[no] bsr-candidate ethernet** [*<portnum>* | **loopback** *<num>* | **ve** *<num>* *<hash-mask-length>* [*<priority>*]

The *<portnum>* | **loopback** *<num>* | **ve** *<num>* parameter specifies the interface. The Layer 3 Switch will advertise the specified interface IP address as a candidate BSR:

- Enter **ethernet** *<portnum>* for a physical interface (port).
- Enter **ve** *<num>* for a virtual interface.
- Enter **loopback** *<num>* for a loopback interface.

The *<hash-mask-length>* parameter specifies the number of bits in a group address that are significant when calculating the group-to-RP mapping. You can specify a value from 1 – 32.

NOTE

Dell recommends you specify 30 for IP version 4 (IPv4) networks.

The *<priority>* specifies the BSR priority. You can specify a value from 0 – 255. When the election process for BSR takes place, the candidate BSR with the highest priority becomes the BSR. The default is 0.

Configuring RPs

Enter a command such as the following to configure the Layer 3 Switch as a candidate RP.

```
PowerConnect(config-pim-router)#rp-candidate ethernet 2
```

Syntax: **[no] rp-candidate ethernet***<portnum>* | **loopback** *<num>* | **ve** *<num>*

The *<portnum>* | **loopback** *<num>* | **ve** *<num>* parameter specifies the interface. The Layer 3 Switch will advertise the specified interface IP address as a candidate RP:

- Enter **ethernet** *<portnum>* for a physical interface (port).
- Enter **ve** *<num>* for a virtual interface.
- Enter **loopback** *<num>* for a loopback interface.

By default, this command configures the Layer 3 Switch as a candidate RP for all group numbers beginning with 224. As a result, the Layer 3 Switch is a candidate RP for all valid PIM Sparse group numbers. You can change this by adding or deleting specific address ranges. The following example narrows the group number range for which the Layer 3 Switch is a candidate RP by explicitly adding a range.

```
PowerConnect(config-pim-router)#rp-candidate add 224.126.0.0 16
```

Syntax: **[no] rp-candidate add** *<group-addr>* *<mask-bits>*

The *<group-addr>* *<mask-bits>* specifies the group address and the number of significant bits in the subnet mask. In this example, the Layer 3 Switch is a candidate RP for all groups that begin with 224.126. When you add a range, you override the default. The Layer 3 Switch then becomes a candidate RP only for the group address ranges you add.

You also can change the group numbers for which the Layer 3 Switch is a candidate RP by deleting address ranges. For example, to delete all addresses from 224.126.22.0 – 224.126.22.255, enter the following command.

```
PowerConnect(config-pim-router)#rp-candidate delete 224.126.22.0 24
```

Syntax: `rp-candidate delete <group-addr> <mask-bits>`

The usage of the `<group-addr> <mask-bits>` parameter is the same as for the `rp-candidate add` command.

If you enter both commands shown in the example above, the net effect is that the Layer 3 Switch becomes a candidate RP for groups 224.126.0.0 – 224.126.21.255 and groups 224.126.23.0 – 224.126.255.255.

Updating PIM-Sparse forwarding entries with new RP configuration

If you make changes to your static RP configuration, the entries in the PIM-Sparse multicast forwarding table continue to use the old RP configuration until they are aged out.

The `clear pim rp-map` command allows you to update the entries in the static multicast forwarding table immediately after making RP configuration changes. This command is meant to be used with `rp-address` command.

To update the entries in a PIM sparse static multicast forwarding table with new RP configuration, enter the following command at the privileged EXEC level of the CLI.

```
PowerConnect#clear pim rp-map
```

Syntax: `clear pim rp-map`

Statically specifying the RP

Dell recommends that you use the PIM Sparse protocol RP election process so that a backup RP can automatically take over if the active RP router becomes unavailable. However, if you do not want the RP to be selected by the RP election process but instead you want to explicitly identify the RP by its IP address, you can do using the following CLI method.

If you explicitly specify the RP, the Layer 3 Switch uses the specified RP for all group-to-RP mappings and overrides the set of candidate RPs supplied by the BSR.

NOTE

Specify the same IP address as the RP on all PIM Sparse routers within the PIM Sparse domain. Make sure the router is on the backbone or is otherwise well connected to the rest of the network.

To specify the IP address of the RP, enter commands such as the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#rp-address 207.95.7.1
```

Syntax: `[no] rp-address <ip-addr>`

The `<ip-addr>` parameter specifies the IP address of the RP.

The command in the example above identifies the router interface at IP address 207.95.7.1 as the RP for the PIM Sparse domain. The Layer 3 Switch will use the specified RP and ignore group-to-RP mappings received from the BSR.

Changing the Shortest Path Tree (SPT) threshold

In a typical PIM Sparse domain, there may be two or more paths from a DR (designated router) for a multicast source to a PIM group receiver:

- **Path through the RP** – This is the path the Layer 3 Switch uses the first time it receives traffic for a PIM group. However, the path through the RP may not be the shortest path from the Layer 3 Switch to the receiver.
- **Shortest Path** – Each PIM Sparse router that is a DR for a multicast source calculates a shortest path tree (SPT) to all the PIM Sparse group receivers within the domain, with the Layer 3 Switch itself as the root of the tree. The first time a Layer 3 Switch configured as a PIM router receives a packet for a PIM receiver, the Layer 3 Switch sends the packet to the RP for the group. The Layer 3 Switch also calculates the SPT from itself to the receiver. The next time the Layer 3 Switch receives a PIM Sparse packet for the receiver, the Layer 3 Switch sends the packet toward the receiver using the shortest route, which may not pass through the RP.

By default, the device switches from the RP to the SPT after receiving the first packet for a given PIM Sparse group. The Layer 3 Switch maintains a separate counter for each PIM Sparse source-group pair.

After the Layer 3 Switch receives a packet for a given source-group pair, the Layer 3 Switch starts a PIM data timer for that source-group pair. If the Layer 3 Switch does not receive another packet for the source-group pair before the timer expires, it reverts to using the RP for the next packet received for the source-group pair. In accordance with the PIM Sparse RFC recommendation, the timer is 210 seconds and is not configurable. The counter is reset to zero each time the Layer 3 Switch receives a packet for the source-group pair.

You can change the number of packets that the Layer 3 Switch sends using the RP before switching to using the SPT. To do so, use the following CLI method.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#spt-threshold 1000
```

Syntax: [no] spt-threshold infinity | <num>

The **infinity** | <num> parameter specifies the number of packets. If you specify infinity, the Layer 3 Switch sends packets using the RP indefinitely and does not switch over to the SPT. If you enter a specific number of packets, the Layer 3 Switch does not switch over to using the SPT until it has sent the number of packets you specify using the RP.

Changing the PIM join and prune message interval

By default, the Layer 3 Switch sends PIM Sparse Join/Prune messages every 60 seconds. These messages inform other PIM Sparse routers about clients who want to become receivers (Join) or stop being receivers (Prune) for PIM Sparse groups.

You can change the Join/Prune message interval using the following CLI method.

NOTE

Use the same Join/Prune message interval on all the PIM Sparse routers in the PIM Sparse domain. If the routers do not all use the same timer interval, the performance of PIM Sparse can be adversely affected.

To change the Join/Prune interval, enter commands such as the following.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#message-interval 30
```

Syntax: [no] message-interval <num>

The <num> parameter specifies the number of seconds and can range from 1 – 65535. The default is 60.

Dropping PIM traffic in hardware

Unwanted PIM Dense or PIM Sparse multicast traffic can be dropped in hardware on Layer 3 Switches. Refer to [“Passive multicast route insertion”](#) on page 501.

On PowerConnect B-Series TI24X devices, anycast RP is supported on a fully-meshed topology. Configure MSDP peers and configure mesh-groups that use a fully meshed peer topology. MSDP peering can be established between physical interfaces or loopback interfaces. MSDP is used between all intra-domain RPs in a full-mesh configuration to provide redundancy and backup of PIM-SM network.

```
RP1(config)#router ospf
RP1(config-ospf-router)# area 0
RP1(config-ospf-router)# exit
RP1(config)# interface loopback 1
RP1(config-lbif-1)# ip ospf area 0
RP1(config-lbif-1)# ip ospf passive
RP1(config-lbif-1)# ip address 10.0.0.1/32
RP1(config-lbif-1)# ip pim-sparse
RP1(config-lbif-1)# exit
RP1(config)# interface loopback 2
RP1(config-lbif-2)# ip ospf area 0
RP1(config-lbif-2)# ip ospf passive
RP1(config-lbif-2)# ip address 10.1.1.1/32
RP1(config-lbif-2)# exit
RP1(config)# interface ethernet 1
RP1(config-if-e10000-1)# ip ospf area 0
RP1(config-if-e10000-1)# ip address 192.1.1.1/24
RP1(config-if-e10000-1)# ip pim-sparse
RP1(config)# interface ethernet 2
RP1(config-if-e10000-2)# ip ospf area 0
RP1(config-if-e10000-2)# ip ospf cost 5
RP1(config-if-e10000-2)# ip address 192.2.1.1/24
RP1(config-if-e10000-2)# ip pim-sparse
RP1(config)# interface ethernet 3
RP1(config-if-e10000-3)# ip ospf area 0
RP1(config-if-e10000-3)# ip ospf cost 10
RP1(config-if-e10000-3)# ip address 192.3.1.1/24
RP1(config-if-e10000-3)# ip pim-sparse
RP1(config-if-e10000-3)# exit
RP1(config)# router pim
RP1(config-pim-router)# rp-candidate loopback 1
RP1(config-pim-router)# exit
RP1(config)# router msdp
RP1(config-msdp-router)# msdp-peer 10.1.1.2 connect-source loopback 2
RP1(config-msdp-router)# originator-id loopback 2
RP1(config)#ip router-id 10.1.1.1RP2(config)#router ospf
RP2(config-ospf-router)# area 0
RP2(config-ospf-router)# exit
RP2(config)# interface loopback 1
RP2(config-lbif-1)# ip ospf area 0
RP2(config-lbif-1)# ip ospf passive
RP2(config-lbif-1)# ip address 10.0.0.1/32
RP2(config-lbif-1)# ip pim-sparse
```

```

RP2(config-lbif-1)# exit
RP2(config)# interface loopback 2
RP2(config-lbif-2)# ip ospf area 0
RP2(config-lbif-2)# ip ospf passive
RP2(config-lbif-2)# ip address 10.1.1.2/32
RP2(config-lbif-2)# exit
RP2(config)# interface ethernet 1
RP2(config-if-e10000-1)# ip ospf area 0
RP2(config-if-e10000-1)# ip address 192.1.1.2/24
RP2(config-if-e10000-1)# ip pim-sparse
RP2(config)# interface ethernet 2
RP2(config-if-e10000-2)# ip ospf area 0
RP2(config-if-e10000-2)# ip ospf cost 5
RP2(config-if-e10000-2)# ip address 192.5.2.1/24
RP2(config-if-e10000-2)# ip pim-sparse
RP2(config)# interface ethernet 3
RP2(config-if-e10000-3)# ip ospf area 0
RP2(config-if-e10000-3)# ip ospf cost 10
RP2(config-if-e10000-3)# ip address 192.6.1.2/24
RP2(config-if-e10000-3)# ip pim-sparse
RP2(config-if-e10000-3)# exit
RP2(config)# router pim
RP2(config-pim-router)# rp-candidate loopback 1
RP2(config-pim-router)# exit
RP2(config)# router msdp
RP2(config-msdp-router)# msdp-peer 10.1.1.1 connect-source loopback 2
RP2(config-msdp-router)# originator-id loopback 2
RP2(config)#ip router-id 10.1.1.2

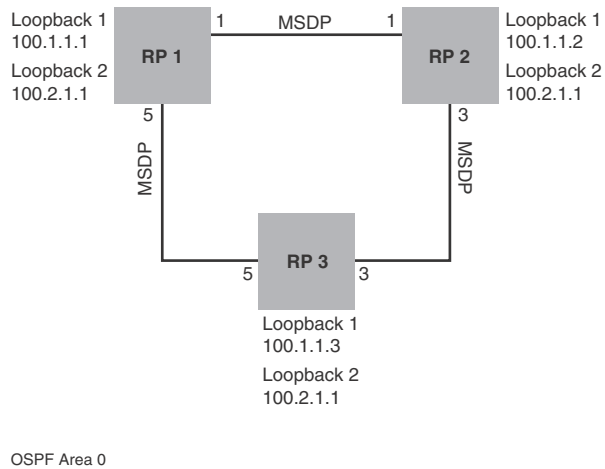
```

The example shown in Figure 88 is an anycast-enabled network with three RPs connected in a triangular mesh topology. Loopback 2 in RP 1, RP 2, and RP 3 have the same IP address, which is the anycast RP address. Loopback 1 in RP1, RP2, and RP3 have different IP addresses and are configured as MSDP peering IP addresses in a triangular mesh configuration.

OSPF is configured as the IGP for the network and all of the devices are in OSPF area 0. This example demonstrates only anycast RP configurations. Assuming a PIM-SM network with three anycast RPs configured, the RP address is configured to be the anycast RP address that was configured on Loopback interface 2 of RP1, RP2, and RP3. Based on the IGP routes, all routers in the network are registered to the shortest path anycast RP. This shares the load between all the RPs, and provides hot backup.

The configuration examples demonstrate the commands required to enable this application.

FIGURE 88 Anycast enabled network in a triangular mesh topology



RP 1 Configuration

The following commands provide the configuration for the RP 1 router in Figure 88.

```
RP1(config)#router ospf
RP1(config-ospf-router)# area 0
RP1(config-ospf-router)# exit
RP1(config)# interface loopback 1
RP1(config-lbif-1)# ip ospf area 0
RP1(config-lbif-1)# ip ospf passive
RP1(config-lbif-1)# ip address 100.1.1.1/32
RP1(config-lbif-1)# exit
RP1(config)# interface loopback 2
RP1(config-lbif-2)# ip ospf area 0
RP1(config-lbif-2)# ip ospf passive
RP1(config-lbif-2)# ip pim-sparse
RP1(config-lbif-2)# ip address 100.2.1.1/32
RP1(config-lbif-2)# exit
RP1(config)# interface ethernet 1
RP1(config-if-e10000-1)# ip ospf area 0
RP1(config-if-e10000-1)# ip ospf cost 2
RP1(config-if-e10000-1)# ip address 192.1.1.1/24
RP1(config-if-e10000-1)# ip pim-sparse
RP1(config)# interface ethernet 5
RP1(config-if-e10000-5)# ip ospf area 0
RP1(config-if-e10000-5)# ip address 192.3.1.1/24
RP1(config-if-e10000-5)# ip pim-sparse
RP1(config)# router pim
RP1(config-pim-router)# rp-address 100.2.1.1
RP1(config-pim-router)# exit
RP1(config)# router msdp
RP1(config-msdp-router)# msdp-peer 100.1.1.2 connect-source loopback 1
RP1(config-msdp-router)# msdp-peer 100.1.1.3 connect-source loopback 1
RP1(config-msdp-router)# mesh-group mesh1 100.1.1.2
RP1(config-msdp-router)# mesh-group mesh1 100.1.1.3
RP1(config-msdp-router)# originator-id loopback 1
RP1(config-msdp-router)# exit
RP1(config)#ip router-id 100.1.1.1
```

RP 2 Configuration

The following commands provide the configuration for the RP 2 router in .Figure 88.

```
RP2(config)#router ospf
RP2(config-ospf-router)# area 0
RP2(config-ospf-router)# exit
RP2(config)# interface loopback 1
RP2(config-lbif-1)# ip ospf area 0
RP2(config-lbif-1)# ip ospf passive
RP2(config-lbif-1)# ip address 100.1.1.2/32
RP2(config-lbif-1)# exit
RP2(config)# interface loopback 2
RP2(config-lbif-2)# ip ospf area 0
RP2(config-lbif-2)# ip ospf passive
RP2(config-lbif-2)# ip address 100.2.1.1/32
RP2(config-lbif-2)# ip pim-sparse
RP2(config-lbif-2)# exit
RP2(config)# interface ethernet 1
RP2(config-if-e10000-1)# ip ospf area 0
RP2(config-if-e10000-1)# ip ospf cost 2
RP2(config-if-e10000-1)# ip address 192.1.1.2/24
RP2(config-if-e10000-1)# ip pim-sparse
RP2(config)# interface ethernet 3
RP2(config-if-e10000-3)# ip ospf area 0
RP2(config-if-e10000-3)# ip ospf cost 2
RP2(config-if-e10000-3)# ip address 192.2.1.2/24
RP2(config-if-e10000-3)# ip pim-sparse
RP2(config-if-e10000-3)# exit
RP2(config)# router pim
RP2(config-pim-router)# rp-address 100.2.1.1
RP2(config-pim-router)# exit
RP2(config)# router msdp
RP2(config-msdp-router)# msdp-peer 100.1.1.1 connect-source loopback 1
RP2(config-msdp-router)# msdp-peer 100.1.1.3 connect-source loopback 1
RP2(config-msdp-router)# mesh-group mesh1 100.1.1.1
RP2(config-msdp-router)# mesh-group mesh1 100.1.1.3
RP2(config-msdp-router)# originator-id loopback 1
RP2(config-msdp-router)# exit
RP2(config)#ip router-id 100.1.1.2
```

RP 3 Configuration

The following commands provide the configuration for the RP 3 router in Figure 88.

```
RP3(config)#router ospf
RP3(config-ospf-router)# area 0
RP3(config-ospf-router)# exit
RP3(config)# interface loopback 1
RP3(config-lbif-1)# ip ospf area 0
RP3(config-lbif-1)# ip ospf passive
RP3(config-lbif-1)# ip address 100.1.1.3/32
RP3(config-lbif-1)# exit
RP3(config)# interface loopback 2
RP3(config-lbif-2)# ip ospf area 0
RP3(config-lbif-2)# ip ospf passive
RP3(config-lbif-2)# ip address 100.2.1.1/32
RP3(config-lbif-2)# ip pim-sparse
RP3(config-lbif-2)# exit
RP3(config)# interface ethernet 3
RP3(config-if-e10000-3)# ip ospf area 0
RP3(config-if-e10000-3)# ip ospf cost 2
```

```
RP3(config-if-e10000-3)# ip address 192.2.1.3/24
RP3(config-if-e10000-3)# ip pim-sparse
RP3(config)# interface ethernet 5
RP3(config-if-e10000-5)# ip ospf area 0
RP3(config-if-e10000-5)# ip ospf cost 2
RP3(config-if-e10000-5)# ip address 192.3.1.3/24
RP3(config-if-e10000-5)# ip pim-sparse
RP3(config-if-e10000-5)# exit
RP3(config)# router pim
RP3(config-pim-router)# rp-address 100.2.1.1
RP3(config-pim-router)# exit
RP3(config)# router msdp
RP3(config-msdp-router)# msdp-peer 100.1.1.1 connect-source loopback 1
RP3(config-msdp-router)# msdp-peer 100.1.1.2 connect-source loopback 1
RP3(config-msdp-router)# mesh-group mesh1 100.1.1.1
RP3(config-msdp-router)# mesh-group mesh1 100.1.1.2
RP3(config-msdp-router)# originator-id loopback 1
RP3(config-msdp-router)# exit
RP3(config)#ip router-id 100.1.1.3
```

Displaying PIM Sparse configuration information and statistics

You can display the following PIM Sparse information:

- Basic PIM Sparse configuration information
- Group information
- BSR information
- Candidate RP information
- RP-to-group mappings
- RP information for a PIM Sparse group
- RP set list
- PIM Neighbor information
- The PIM flow cache
- The PIM multicast cache
- PIM traffic statistics

Displaying basic PIM Sparse configuration information

To display basic configuration information for PIM Sparse, enter the following command at any CLI level.

```
PowerConnect# show ip pim sparse

Global PIM Sparse Mode Settings
  Hello interval: 60, Neighbor timeout: 180
  Bootstrap Msg interval: 130, Candidate-RP Advertisement interval: 60
  Join/Prune interval: 60, SPT Threshold: 1

Interface Ethernet e8
TTL Threshold: 1, Enabled
Local Address: 207.95.8.1

Interface Ve 1
TTL Threshold: 1, Enabled
Local Address: 207.95.6.1
```

Syntax: show ip pim sparse

This example shows the PIM Sparse configuration information on PIM Sparse router A in [Figure 87](#).

This display shows the following information.

TABLE 75 Output of show ip pim sparse

This field...	Displays...
Global PIM Sparse mode settings	
Hello interval	How frequently the Layer 3 Switch sends PIM Sparse hello messages to its PIM Sparse neighbors. This field show the number of seconds between hello messages. PIM Sparse routers use hello messages to discover one another.
Neighbor timeout	How many seconds the Layer 3 Switch will wait for a hello message from a neighbor before determining that the neighbor is no longer present and removing cached PIM Sparse forwarding entries for the neighbor.
Bootstrap Msg interval	How frequently the BSR configured on the Layer 3 Switch sends the RP set to the RPs within the PIM Sparse domain. The RP set is a list of candidate RPs and their group prefixes. A candidate RP group prefix indicates the range of PIM Sparse group numbers for which it can be an RP. NOTE: This field contains a value only if an interface on the Layer 3 Switch is elected to be the BSR. Otherwise, the field is blank.
Candidate-RP Advertisement interval	How frequently the candidate PR configured on the Layer 3 Switch sends candidate RP advertisement messages to the BSR. NOTE: This field contains a value only if an interface on the Layer 3 Switch is configured as a candidate RP. Otherwise, the field is blank.
Join/Prune interval	How frequently the Layer 3 Switch sends PIM Sparse Join/Prune messages for the multicast groups it is forwarding. This field show the number of seconds between Join/Prune messages. The Layer 3 Switch sends Join/Prune messages on behalf of multicast receivers who want to join or leave a PIM Sparse group. When forwarding packets from PIM Sparse sources, the Layer 3 Switch sends the packets only on the interfaces on which it has received join requests in Join/Prune messages for the source group. You can change the Join/Prune interval if needed. Refer to “Changing the PIM join and prune message interval” on page 484.
SPT Threshold	The number of packets the Layer 3 Switch sends using the path through the RP before switching to using the SPT path.

TABLE 75 Output of show ip pim sparse (Continued)

This field...	Displays...
PIM Sparse Interface Information	
<p>NOTE: You also can display IP multicast interface information using the show ip pim interface command. However, this command lists all IP multicast interfaces, including regular PIM (dense mode) interfaces. The show ip pim sparse command lists only the PIM Sparse interfaces.</p>	
Interface	<p>The type of interface and the interface number. The interface type can be one of the following:</p> <ul style="list-style-type: none"> • Ethernet • VE <p>The number is either a port number or the virtual interface (VE) number.</p>
TTL Threshold	<p>Following the TTL threshold value, the interface state is listed. The interface state can be one of the following:</p> <ul style="list-style-type: none"> • Disabled • Enabled
Local Address	Indicates the IP address configured on the port or virtual interface.

Displaying a list of multicast groups

To display a list of the IP multicast groups the Layer 3 Switch is forwarding, enter the following command at any CLI level.

```
PowerConnect# show ip pim group
Total number of groups: 1
Index 1      Group 228.1.0.88
      Group member at e2: v33 EX 0,
      Group member at e13: v55 EX 0,
      Group member at e1: v5 EX 0,
```

Syntax: **show ip pim group**

This display shows the following information.

TABLE 76 Output of show ip pim group

This field...	Displays...
Total number of Groups	<p>Lists the total number of IP multicast groups the Layer 3 Switch is forwarding.</p> <p>NOTE: This list can include groups that are not PIM Sparse groups. If interfaces on the Layer 3 Switch are configured for regular PIM (dense mode) , these groups are listed too.</p>
Index	The index number of the table entry in the display.
Group	The multicast group address
Group member	
Ports	The Layer 3 Switch ports connected to the receivers of the groups.

Displaying BSR information

To display BSR information, enter the following command at any CLI level.

```
PowerConnect# show ip pim bsr

PIMv2 Bootstrap information

This system is the elected Bootstrap Router (BSR)
  BSR address: 207.95.7.1
  Uptime: 00:33:52, BSR priority: 5, Hash mask length: 32
  Next bootstrap message in 00:00:20

Next Candidate-RP-advertisement in 00:00:10
  RP: 207.95.7.1
    group prefixes:
      224.0.0.0 / 4
Candidate-RP-advertisement period: 60
```

This example shows information displayed on a Layer 3 Switch that has been elected as the BSR. The following example shows information displayed on a Layer 3 Switch that is not the BSR. Notice that some fields shown in the example above do not appear in the example below.

```
PowerConnect# show ip pim bsr

PIMv2 Bootstrap information
  local BSR address = 207.95.7.1
  local BSR priority = 5
```

Syntax: show ip pim bsr

This display shows the following information.

TABLE 77 Output of show ip pim bsr

This field...	Displays...
BSR address or local BSR address	The IP address of the interface configured as the PIM Sparse Bootstrap Router (BSR). NOTE: If the word “local” does not appear in the field, this Layer 3 Switch is the BSR. If the word “local” does appear, this Layer 3 Switch is not the BSR.
Uptime	The amount of time the BSR has been running. NOTE: This field appears only if this Layer 3 Switch is the BSR.
BSR priority or local BSR priority	The priority assigned to the interface for use during the BSR election process. During BSR election, the priorities of the candidate BSRs are compared and the interface with the highest BSR priority becomes the BSR. NOTE: If the word “local” does not appear in the field, this Layer 3 Switch is the BSR. If the word “local” does appear, this Layer 3 Switch is not the BSR.
Hash mask length	The number of significant bits in the IP multicast group comparison mask. This mask determines the IP multicast group numbers for which the Layer 3 Switch can be a BSR. The default is 32 bits, which allows the Layer 3 Switch to be a BSR for any valid IP multicast group number. NOTE: This field appears only if this Layer 3 Switch is the BSR.
Next bootstrap message in	Indicates how many seconds will pass before the BSR sends its next Bootstrap message. NOTE: This field appears only if this Layer 3 Switch is the BSR.
Next Candidate-PR-advertisement message in	Indicates how many seconds will pass before the BSR sends its next candidate PR advertisement message. NOTE: This field appears only if this Layer 3 Switch is the BSR.
RP	Indicates the IP address of the Rendezvous Point (RP). NOTE: This field appears only if this Layer 3 Switch is the BSR.

TABLE 77 Output of show ip pim bsr (Continued)

This field...	Displays...
group prefixes	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP. NOTE: This field appears only if this Layer 3 Switch is the BSR.
Candidate-RP-advertisement period	Indicates how frequently the BSR sends candidate RP advertisement messages. NOTE: This field appears only if this Layer 3 Switch is the BSR.

Displaying Pim resources

To display the hardware resource information such as hardware allocation, availability, and limit for software data structure, enter the following command.

```
PowerConnect# show ip pim resource
```

	alloc	in-use	avail	allo-fail	up-limit	get-mem
NBR list	64	0	64	0	512	0
timer	256	0	256	0	4096	0
pimsm J/P elem	0	0	0	0	48960	0
pimsm group2rp	0	0	0	0	4096	0
pimsm L2 reg xmt	64	0	64	0	no-limit	0
mcache	256	0	256	0	1024	0
mcache hash link	997	0	997	0	no-limit	0
mcache 2nd hash	9	0	9	0	997	0
graft if no mcache	197	0	197	0	no-limit	0
pim/dvm global group	256	0	256	0	no-limit	0
pim prune	128	0	128	0	40960	0
Output intf-vlan	2000	0	2000	0	no-limit	0
group hash link	97	0	97	0	no-limit	0
2D vlan for nbr, glb	2000	0	2000	0	no-limit	0
Output intf.	1024	0	1024	0	no-limit	0
2D for glb grp	1024	0	1024	0	no-limit	0
pim/dvm config. intf	128	2	126	0	no-limit	2
Prune rate limit	256	0	256	0	no-limit	0
Distributed add cpu	128	0	128	0	no-limit	0
L2 VIDX	256	0	256	0	4096	0
L2 VIDX hash	997	0	997	0	no-limit	0
igmp group	256	0	256	0	4096	0
igmp phy port	1024	0	1024	0	no-limit	0
igmp exist phy port	1024	4	1020	0	no-limit	4
igmp G/GS query	128	0	128	0	no-limit	0
igmp v3 source	2000	0	2000	0	500000	0
igmp v3 tracking	0	0	0	0	no-limit	0
igmp glb sorted list	2000	0	2000	0	500000	0

total pool memory 286918 bytes

#of PIM ports: physical 2, VEs 0 (max: 512), loopback 0, tunnels 0
 Total Mlls in pool: 943 Allocated MLL: 0 Available MLL: 943
 SW processed pkts 0

Syntax: show ip pim resource

For the PowerConnect devices, the number of Hardware (Hw) resources available are 2048.

For each software data structure listed in the output, the following information is shown.

TABLE 78 Output of show ip pim resource

This field...	Displays...
alloc	Number of nodes of that data that are currently allocated in memory.
in-use	Number of allocated nodes in use
avail	Number of allocated nodes are not in use
allo-fail	Number of allocated notes that failed
up-limit	Maximum number of nodes that can be allocated for a data structure. This may or may not be configurable, depending on the data structure
get-mem	Number of attempts made to use allocated nodes
#of PIM ports	Total number of PIM ports, by port type, on the device
Total, allocated, and available Mils	In Layer 3 multicast, this refers to the Multicast Linked List that contains information on where (S,G) gets forwarded. Each (S,G) entry requires a single MLL entry to forward traffic to all physical, untagged ports. Also, one MLL entry is required per VLAN that has tagged outbound ports. There can be up to 1024 MLL entries.

Displaying candidate RP information

To display candidate RP information, enter the following command at any CLI level.

```
PowerConnect# show ip pim rp-candidate

Next Candidate-RP-advertisement in 00:00:10
  RP: 207.95.7.1
    group prefixes:
      224.0.0.0 / 4

Candidate-RP-advertisement period: 60
```

This example show information displayed on a Layer 3 Switch that is a candidate RP. The following example shows the message displayed on a Layer 3 Switch that is not a candidate RP.

```
PowerConnect# show ip pim rp-candidate
```

This system is not a Candidate-RP.

Syntax: show ip pim rp-candidate

This display shows the following information.

TABLE 79 Output of show ip pim rp-candidate

This field...	Displays...
Candidate-RP-advertisement in	Indicates how many seconds will pass before the BSR sends its next RP message. NOTE: This field appears only if this Layer 3 Switch is a candidate RP.
RP	Indicates the IP address of the Rendezvous Point (RP). NOTE: This field appears only if this Layer 3 Switch is a candidate RP.

TABLE 79 Output of show ip pim rp-candidate (Continued)

This field...	Displays...
group prefixes	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP. NOTE: This field appears only if this Layer 3 Switch is a candidate RP.
Candidate-RP-advertisement period	Indicates how frequently the BSR sends candidate RP advertisement messages. NOTE: This field appears only if this Layer 3 Switch is a candidate RP.

Displaying RP-to-group mappings

To display RP-to-group-mappings, enter the following command at any CLI level.

```
PowerConnect# show ip pim rp-map
Number of group-to-RP mappings: 6
```

```
Group address      RP address
-----
1 239.255.163.1    99.99.99.5
2 239.255.163.2    99.99.99.5
3 239.255.163.3    99.99.99.5
4 239.255.162.1    99.99.99.5
5 239.255.162.2    43.43.43.1
6 239.255.162.3    99.99.99.5
```

Syntax: show ip pim rp-map

This display shows the following information.

TABLE 80 Output of show ip pim rp-map

This field...	Displays...
Group address	Indicates the PIM Sparse multicast group address using the listed RP.
RP address	Indicates the IP address of the Rendezvous Point (RP) for the listed PIM Sparse group.

Displaying RP information for a PIM Sparse group

To display RP information for a PIM Sparse group, enter the following command at any CLI level.

```
PowerConnect# show ip pim rp-hash 239.255.162.1

RP: 207.95.7.1, v2
Info source: 207.95.7.1, through bootstrap
```

Syntax: show ip pim rp-hash <group-addr>

The <group-addr> parameter is the address of a PIM Sparse IP multicast group.

This display shows the following information.

TABLE 81 Output of show ip pim rp-hash

This field...	Displays...
RP	Indicates the IP address of the Rendezvous Point (RP) for the specified PIM Sparse group. Following the IP address is the port or virtual interface through which this Layer 3 Switch learned the identity of the RP.
Info source	Indicates the IP address on which the RP information was received. Following the IP address is the method through which this Layer 3 Switch learned the identity of the RP.

Displaying the RP set list

To display the RP set list, enter the following command at any CLI level.

```
PowerConnect# show ip pim rp-set
Group address Static-RP-address Override
-----
Access-List 44 99.99.99.5 On
Number of group prefixes Learnt from BSR: 1
Group prefix = 239.255.162.0/24 #RPs expected: 1
#RPs received: 1
RP 1: 43.43.43.1 priority=0 age=0
```

Syntax: show ip pim rp-set

This display shows the following information.

TABLE 82 Output of show ip pim rp-set

This field...	Displays...
Number of group prefixes	The number of PIM Sparse group prefixes for which the RP is responsible.
Group prefix	Indicates the multicast groups for which the RP listed by the previous field is a candidate RP.
RPs expected/received	Indicates how many RPs were expected and received in the latest Bootstrap message.
RP <num>	Indicates the RP number. If there are multiple RPs in the PIM Sparse domain, a line of information for each of them is listed, and they are numbered in ascending numerical order.
priority	The RP priority of the candidate RP. During the election process, the candidate RP with the highest priority is elected as the RP.
age	The age (in seconds) of this RP-set. NOTE: If this Layer 3 Switch is not a BSR, this field contains zero. Only the BSR ages the RP-set.

Displaying multicast neighbor information

To display information about the Layer 3 Switch PIM neighbors, enter the following command at any CLI level.

```
PowerConnect# show ip pim nbr
```

```
Port Neighbor      Holdtime Age    UpTime
          sec      sec    sec
e8  207.95.8.10    180    60    900
Port Neighbor      Holdtime Age    UpTime
          sec      sec    sec
v1  207.95.6.2     180    60    900
```

Syntax: show ip pim nbr

This display shows the following information.

TABLE 83 Output of show ip pim nbr

This field...	Displays...
Port	The interface through which the Layer 3 Switch is connected to the neighbor.
Neighbor	The IP interface of the PIM neighbor interface.
Holdtime sec	Indicates how many seconds the neighbor wants this Layer 3 Switch to hold the entry for this neighbor in memory. The neighbor sends the Hold Time in its Hello packets: <ul style="list-style-type: none"> If the Layer 3 Switch receives a new Hello packet before the Hold Time received in the previous packet expires, the Layer 3 Switch updates its table entry for the neighbor. If the Layer 3 Switch does not receive a new Hello packet from the neighbor before the Hold time expires, the Layer 3 Switch assumes the neighbor is no longer available and removes the entry for the neighbor.
Age sec	The number of seconds since the Layer 3 Switch received the last hello message from the neighbor.
UpTime sec	The number of seconds the PIM neighbor has been up. This timer starts when the Layer 3 Switch receives the first Hello messages from the neighbor.

Displaying information about an upstream neighbor device

You can view information about the upstream neighbor device for a given source IP address for IP Protocol Independent Multicast (PIM) packets. For PIM, the software uses the IP route table or multicast route table to lookup the upstream neighbor device.

Enter the following command at the Privileged EXEC level of the CLI.

```
PowerConnect# show ip pim rpf 1.1.20.2
directly connected or through an L2 neighbor
```

Syntax: show ip pim rpf <IP address>

where <IP address> is a valid source IP address

NOTE

If there are multiple equal cost paths to the source, the **show ip pim rpf** command output may not be accurate. If your system has multiple equal cost paths, use the command **show ip pim mcache** to view information about the upstream neighbor.

Displaying the PIM flow cache

To display the PIM flow cache for, enter the following command at any CLI level.

```
PowerConnect# show ip pim flowcache 238.0.0.1

Multicast flow (24.1.1.100 238.0.0.1):
Vidx for source vlan forwarding: 2080
[mcastPrintMll]: vrId 0, ipGrp 238.0.0.1, grpPrefix 32, ipSrc 24.1.1.100, srcPre
32, nextHopIdx 0

1 flow printed
```

Syntax: show ip pim flowcache

This display shows the following information.

TABLE 84 Output of show ip pim flowcache

This field...	Displays...
Source	Indicates the source of the PIM Sparse group.
Group	Indicates the PIM Sparse group.
Parent	Indicates the port or virtual interface from which the Layer 3 Switch receives packets from the group source.
CamFlags	This field is used by Dell technical support for troubleshooting.
CamIndex	This field is used by Dell technical support for troubleshooting.
Fid	This field is used by Dell technical support for troubleshooting.
Flags	This field is used by Dell technical support for troubleshooting.

Displaying the PIM multicast cache

To display the PIM multicast cache, enter the following command at any CLI level.

PowerConnect

```
(* 228.1.0.88) RP 1.1.1.1, in e21 (e21), cnt=0
  Source is directly connected
  Sparse Mode, RPT=1 SPT=0 REG=0 MSDP Adv=0 MSDP Create=0
  L3 (SW) 3: e2(VL33), e1(VL5), TR(e13,e14)(VL55)
  fast=1 slow=0 pru=1 graft
  age=0s up-time=4123m HW=0
```

Syntax:

This display shows the following information.

TABLE 85 Output of show ip pim mcache

This field...	Displays...
(<source>, <group>)	The comma-separated values in parentheses is a source-group pair. The <source> is the PIM source for the multicast <group>. For example, the following entry means source 209.157.24.162 for group 239.255.162.1: (209.157.24.162,239.255.162.1) If the <source> value is * (asterisk), this cache entry uses the RP path. The * value means "all sources". If the <source> is a specific source address, this cache entry uses the SPT path.
RP<ip-addr>	Indicates the RP for the group for this cache entry. NOTE: The RP address appears only if the RPT flag is set to 1 and the SPT flag is set to 0 (see below).

TABLE 85 Output of show ip pim mcache (Continued)

This field...	Displays...
forward port	The port through which the Layer 3 Switch reaches the source.
Count	The number of packets forwarded using this cache entry.
Sparse Mode	Indicates whether the cache entry is for regular PIM (dense mode) or PIM Sparse. This flag can have one of the following values: <ul style="list-style-type: none"> • 0 – The entry is not for PIM Sparse (and is therefore for the dense mode of PIM). • 1– The entry is for PIM Sparse.
RPT	Indicates whether the cache entry uses the RP path or the SPT path. The RPT flag can have one of the following values: <ul style="list-style-type: none"> • 0 – The SPT path is used instead of the RP path. • 1– The RP path is used instead of the SPT path. <p>NOTE: The values of the RP and SPT flags are always opposite (one is set to 0 and the other is set to 1).</p>
SPT	Indicates whether the cache entry uses the RP path or the SPT path. The SP flag can have one of the following values: <ul style="list-style-type: none"> • 0 – The RP path is used instead of the SPT path. • 1– The SPT path is used instead of the RP path. <p>NOTE: The values of the RP and SPT flags are always opposite (one is set to 0 and the other is set to 1).</p>
Register Suppress	Indicates whether the Register Suppress timer is running. This field can have one of the following values: <ul style="list-style-type: none"> • 0 – The timer is not running. • 1 – The timer is running.
member ports	Indicates the Layer 3 Switch physical ports to which the receivers for the source and group are attached. The receivers can be directly attached or indirectly attached through other PIM Sparse routers.
virtual ports	Indicates the virtual interfaces to which the receivers for the source and group are attached. The receivers can be directly attached or indirectly attached through other PIM Sparse routers.
prune ports	Indicates the physical ports on which the Layer 3 Switch has received a prune notification (in a Join/Prune message) to remove the receiver from the list of recipients for the group.
virtual prune ports	Indicates the virtual interfaces ports on which the Layer 3 Switch has received a prune notification (in a Join/Prune message) to remove the receiver from the list of recipients for the group.

Displaying PIM traffic statistics

To display PIM traffic statistics, use the following CLI method.

```
PowerConnect# show ip pim traffic
```

```

Port      Hello           J/P           Register      RegStop       Assert
      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]
e8      19      19      32      0      0      0      37      0      0      0

Port      Hello           J/P           Register      RegStop       Assert
      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]
v1      18      19      0      20      0      0      0      0      0      0

Port      Hello           J/P           Register      RegStop       Assert
      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]      [Rx      Tx]
v2      0      19      0      0      0      16      0      0      0      0

Total 37      57      32      0      0      0      0      0      0      0
IGMP Statistics:
  Total Recv/Xmit 85/110
  Total Discard/chksum 0/0

```

Syntax: show ip pim traffic

NOTE

If you have configured interfaces for standard PIM (dense mode) on the Layer 3 Switch, statistics for these interfaces are listed first by the display.

This display shows the following information.

TABLE 86 Output of show ip pim traffic

This field...	Displays...
Port	The port or virtual interface on which the PIM interface is configured.
Hello	The number of PIM Hello messages sent or received on the interface.
J/P	The number of Join/Prune messages sent or received on the interface. NOTE: Unlike PIM dense, PIM Sparse uses the same messages for Joins and Prunes.
Register	The number of Register messages sent or received on the interface.
RegStop	The number of Register Stop messages sent or received on the interface.
Assert	The number of Assert messages sent or received on the interface.
Total Recv/Xmit	The total number of IGMP messages sent and received by the Layer 3 Switch.
Total Discard/chksum	The total number of IGMP messages discarded, including a separate counter for those that failed the checksum comparison.

Displaying and clearing PIM errors

If you want to determine how many PIM errors there are on the device, enter the following command.

```

PowerConnect# show ip pim error
**** Warning counter pim route change = 1
HW tagged replication enabled, SW processed pkts 0

```

Syntax: show ip pim error

This command displays the number of warnings and non-zero PIM errors on the device. This count can increase during transition periods such as reboots and topology changes; however, if the device is stable, the number of errors should not increase. If warnings keep increasing in a stable topology, then there may be a configuration error or problems on the device.

To clear the counter for PIM errors, enter the following command.

```
PowerConnect# clear pim counters
```

```
clear pim counters
```

For PowerConnect B-Series TI24X devices, you must configure a fully meshed topology between MSDP peers. This is mandated for this release because of lack of any EGP that provides a peer RPF check for SA messages that are forwarded between MSDP peers. This limitation is not applicable to PowerConnect device because BGP is not supported on the device. This limitation is not applicable to PowerConnect device because BGP is not supported on the device.

Passive multicast route insertion

Passive Multicast Route Insertion (PMRI) enables a Layer 3 switch running PIM Sparse to create an entry for a multicast route (e.g., (S,G)), with no directly attached clients or when connected to another PIM router (transit network).

PMRI is critical for Service Providers wanting to deliver IP-TV services or multicast-based video services. Service Providers, who have transit networks, distribute multicast-based video services to other Service Providers, regardless of whether a client subscribes to a video service.

To configure PMRI, enter the following command at the **router pim** level of the CLI.

```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#hardware-drop
```

Syntax: [no] hardware-drop

When you enable PMRI, the **show ip pim mcache** command output displays the multicast cache entry along with a drop flag, indicating that the device is dropping packets in hardware. If the HW flag is set to 1 (HW=1), it implies that the packets are being dropped in hardware. If the HW flag is set to 0, (HW=0), it indicates that the packets are being processed in software. The following shows an example display output.

```
PowerConnect# show ip pim mcache
1 (10.10.10.18 226.0.1.56) in v10 (e1), cnt=2
Source is directly connected
Sparse Mode, RPT=0 SPT=1 REG=1 MSDP Adv=0 MSDP Create=0
fast=0 slow=0 pru=1 graft age drop
age=0s up-time=2m HW=1 L2-vidx=8191
```

Multicast Source Discovery Protocol (MSDP)

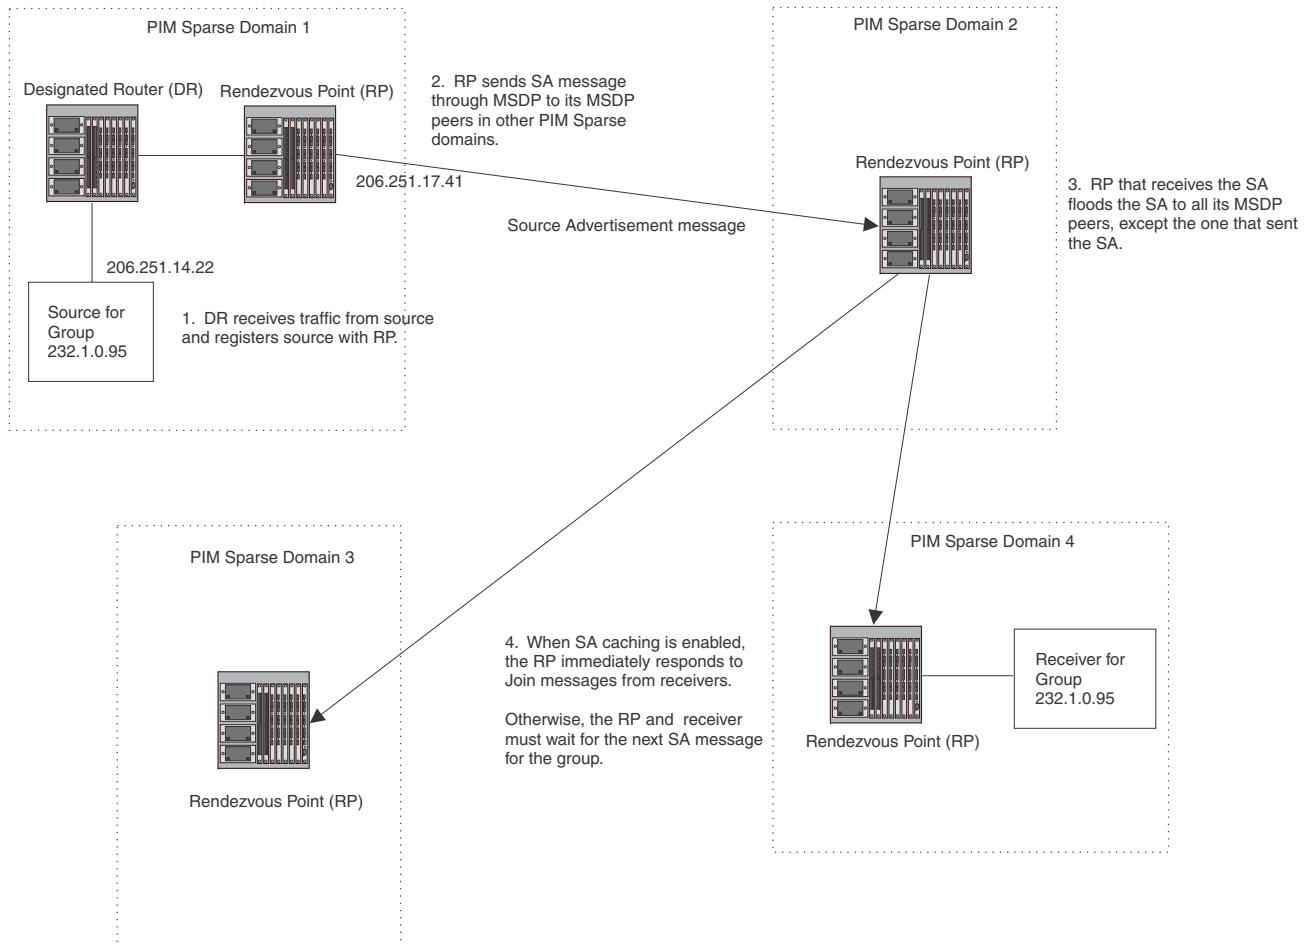
The Multicast Source Discovery Protocol (MSDP) is used by Protocol Independent Multicast (PIM) Sparse routers to exchange routing information for PIM Sparse multicast groups across PIM Sparse domains. Routers running MSDP can discover PIM Sparse sources that are in other PIM Sparse domains.

For PowerConnect B-Series TI24X devices, you must configure a fully meshed topology between MSDP peers. This is mandated for this release because of lack of any EGP that provides a peer RPF check for SA messages that are forwarded between MSDP peers.

PIM Sparse routers use MSDP to register PIM Sparse multicast sources in a domain with the Rendezvous Point (RP) for that domain.

Figure 89 shows an example of some PIM Sparse domains. For simplicity, this example shows only one Designated Router (DR), one group source, and one receiver for the group. Only one PIM Sparse router within each domain needs to run MSDP.

FIGURE 89 PIM Sparse domains joined by MSDP routers



In this example, the source for PIM Sparse multicast group 232.0.1.95 is in PIM Sparse domain 1. The source sends a packet for the group to its directly attached DR. The DR sends a Group Advertisement message for the group to the domain RP. The RP is configured for MSDP, which enables the RP to exchange source information with other PIM Sparse domains by communicating with RPs in other domains that are running MSDP.

The RP sends the source information to each of its peers by sending a Source Active message. The message contains the IP address of the source, the group address to which the source is sending, and the IP address of the RP interface with its peer. By default, the IP address included in the RP address field of the SA message is the IP address of the originating RP. However, if MSDP is instructed to use a specific address as the IP address of the RP in a Source Address message (CLI command **originator-id** <type> <number>), the Source Active message can be the IP address of any interface on the originating RP. The interface is usually a loopback interface.

In this example, the Source Active message contains the following information:

- Source address: 206.251.14.22
- Group address: 232.1.0.95
- RP address: 206.251.17.41

[Figure 89](#) shows only one peer for the MSDP router (which is also the RP here) in domain 1, so the Source Active message goes to only that peer. When an MSDP router has multiple peers, it sends a Source Active message to each of those peers. Each peer sends the Source Advertisement to its other MSDP peers. The RP that receives the Source Active message also sends a Join message for the group if the RP that received the message has receivers for the group.

Peer Reverse Path Forwarding (RPF) flooding

When the MSDP router (also the RP) in domain 2 receives the Source Active message from its peer in domain 1, the MSDP router in domain 2 forwards the message to all its other peers. The propagation process is sometimes called “peer Reverse Path Forwarding (RPF) flooding”. This term refers to the fact that the MSDP router uses its PIM Sparse RPF tree to send the message to its peers within the tree. In [Figure 89](#), the MSDP router floods the Source Active message it receives from its peer in domain 1 to its other peers, in domains 3 and 4.

Note that the MSDP router in domain 2 does not forward the Source Active back to its peer in domain 1, because that is the peer from which the router received the message. An MSDP router never sends a Source Active message back to the peer that sent it. The peer that sent the message is sometimes called the “RPF peer”. The MSDP router uses the unicast routing table for its Exterior Gateway Protocol (EGP) to identify the RPF peer by looking for the route entry that is the next hop toward the source. Often, the EGP protocol is Border Gateway Protocol (BGP) version 4.

NOTE

MSDP depends on BGP for interdomain operations.

The MSDP routers in domains 3 and 4 also forward the Source Active message to all their peers except the ones that sent them the message. [Figure 89](#) does not show additional peers.

Source active caching

When an MSDP router that is also an RP receives a Source Active message, the RP checks its PIM Sparse multicast group table for receivers for the group. If the DR has a receiver for the group being advertised in the Source Active message, the DR sends a Join message for that receiver back to the DR in the domain from which the Source Active message came. Usually, the DR is also the MSDP router that sent the Source Active message.

In [Figure 89](#), if the MSDP router and RP in domain 4 has a table entry for the receiver, the RP sends a Join message on behalf of the receiver back through the RPF tree to the RP for the source, in this case the RP in domain 1.

Some MSDP routers that are also RPs can cache Source Active messages. If the RP is not caching Source Active messages, the RP does not send a Join message unless it already has a receiver that wants to join the group. Otherwise, the RP does not send a Join message and does not remember the information in the Source Active message after forwarding it. If the RP receives a request from a receiver for the group, the RP and receiver must wait for the next Source Active message for that group before the RP can send a Join message for the receiver.

However, if Source Active caching is enabled on the MSDP and RP router, the RP caches the Source Active messages it receives. In this case, even if the RP does not have a receiver for a group when the RP receives the Source Active message for the group, the RP can immediately send a Join for a new receiver that wants to join the group, without waiting for the next Source Active message from the RP in the source domain.

The maximum size of the cache used to store MSDP Source Active messages is 8K and the default size is 4K.

Configuring MSDP

To configure MSDP on a Layer 3 Switch, perform the following tasks:

- Enable MSDP
- Configure the MSDP peers

NOTE

The PIM Sparse Rendezvous Point (RP) is also an MSDP peer.

NOTE

Routers that run MSDP must also run BGP. Also, the source address used by the MSDP router must be the same source address used by BGP. This limitation is not applicable to PowerConnect device because BGP is not supported on the device.

Enabling MSDP

To enable MSDP, enter the following command.

```
PowerConnect(config)# router msdp
```

NOTE

When enabling and disabling MSDP, you do not need to save the configuration and reload the software. The configuration change takes effect immediately, as soon as you enter the CLI command.

Syntax: [no] router msdp

Configuring MSDP peers

On a device, you can configure a maximum of 15 MSDP peers. To configure an MSDP peer, enter a command such as the following at the MSDP configuration level.

```
PowerConnect(config-msdp-router)# msdp-peer 205.216.162.1
```

Syntax: [no] msdp-peer <ip-addr> [connect-source loopback <num>]

The <ip-addr> parameter specifies the IP address of the neighbor.

The **connect-source loopback** <num> parameter specifies the loopback interface you want to use as the source for sessions with the neighbor.

NOTE

It is strongly recommended that you use the **connect-source loopback** <num> parameter when issuing the **msdp-peer** command. If you do not use this parameter, the Layer 3 Switch uses the subnet interface configured on the port.

Also, make sure the IP address of the connect-source loopback is the same as the source IP address used by the MSDP router, the PIM-RP, and the BGP router. This limitation is not applicable to PowerConnect device because BGP is not supported on the device.

The commands in the following example add an MSDP neighbor and specify a loopback interface as the source interface for sessions with the neighbor. By default, the Layer 3 Switch uses the subnet address configured on the physical interface where you configure the neighbor as the source address for sessions with the neighbor.

```
PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)# ip address 9.9.9.9/32
PowerConnect(config-lbif-1)# exit
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.99 connect-source loopback 1
```

Designating an interface IP address as the RP IP address

When an RP receives a Source Active message, it checks its PIM Sparse multicast group table for receivers for the group. If it finds a receiver, the RP sends a Join message for that receiver back to the RP that originated the Source Active message. The originator RP is identified by its RP address.

By default, the IP address included in the RP address field of the SA message is the IP address of the originating RP. Beginning with this release, an SA message can use the IP address of any interface on the originating RP. (The interface is usually a loopback interface.)

To designate an interface IP address to be the IP address of the RP, enter commands such as the following:

```
PowerConnect(config)# interface loopback 2
PowerConnect(config-lbif-2)# ip address 2.2.1.99/32
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# originator-id loopback 2
PowerConnect(config-msdp-router)# exit
```

Syntax: [no] **originator-id** <type> <number>

The **originator-id** parameter instructs MSDP to use the specified address as the IP address of the RP in an SA message. This address must be the address of the interface used to connect the RP to the source. There are no default originator-ids.

The <type> parameter indicates the type of interface used by the RP. Ethernet, loopback and virtual routing interfaces (ve) can be used.

The <number> parameter specifies the interface number (for example: loopback number, port number or virtual routing interface number.)

Filtering MSDP source-group pairs

The following commands allow you to filter individual source-group pairs in MSDP Source-Active messages:

- **sa-filter in** – Filters source-group pairs received in Source-Active messages from an MSDP neighbor
- **sa-filter originate** – Filters source-group pairs in Source-Active messages in advertisements to an MSDP neighbor

Filtering incoming source-active messages

The following example configures filters for incoming Source-Active messages from three MSDP neighbors:

- For peer 2.2.2.99, all source-group pairs in Source-Active messages from the neighbor are filtered out (dropped).
- For peer 2.2.2.97, all source-group pairs except those with 10.x.x.x as the source are permitted.
- For peer 2.2.2.96, all source-group pairs except those associated with RP 2.2.42.3 are permitted.

Example

The following commands configure an IP address on port 3/1. This is the port on which the MSDP neighbors will be configured.

```
PowerConnect(config)# interface ethernet 3/1
PowerConnect(config-if-3/1)# ip address 2.2.2.98/24
PowerConnect(config-if-3/1)# exit
```

The following commands configure a loopback interface. The Layer 3 Switch will use this interface as the source address for communicating with the MSDP neighbors.

```
PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)# ip address 9.9.9.8/32
PowerConnect(config-lbif-1)# exit
```

The following commands configure extended ACLs. The ACLs will be used in route maps, which will be used by the Source-Active filters.

```
PowerConnect(config)# access-list 123 permit 10.0.0.0 0.255.255.255 any
PowerConnect(config)# access-list 124 permit 2.2.42.3 0.0.0.0 any
PowerConnect(config)# access-list 125 permit any any
```

The following commands configure the route maps.

```
PowerConnect(config)# route-map msdp_map deny 1
PowerConnect(config-routemap msdp_map)# match ip address 123
PowerConnect(config-routemap msdp_map)# exit
PowerConnect(config)# route-map msdp2_map permit 1
PowerConnect(config-routemap msdp2_map)# match ip address 125
PowerConnect(config-routemap msdp2_map)# exit
PowerConnect(config)# route-map msdp2_rp_map deny 1
PowerConnect(config-routemap msdp2_rp_map)# match ip route-source 124
PowerConnect(config-routemap msdp2_rp_map)# exit
```

The following commands enable MSDP and configure the MSDP neighbors on port 3/1.


```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.99 connect-source loopback
1
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.97 connect-source loopback
1
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.96 connect-source loopback
1
```

The following commands configure the Source-Active filters.

```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# sa-filter in 2.2.2.99
PowerConnect(config-msdp-router)# sa-filter in 2.2.2.97 route-map msdp_map
PowerConnect(config-msdp-router)# sa-filter in 2.2.2.96 route-map msdp2_map
rp-route-map msdp2_rp_map
```

The **sa-filter** commands configure the following filters:

- **sa-filter in 2.2.2.99** – This command drops all source-group pairs received from neighbor 2.2.2.99.

NOTE

The default action is to deny all source-group pairs from the specified neighbor. If you want to permit some pairs, use route maps.

- **sa-filter in 2.2.2.97 route-map msdp_map** – This command drops source-group pairs received from neighbor 2.2.2.97 if the pairs have source address 10.x.x.x and any group address.
- **sa-filter in 2.2.2.96 route-map msdp2_map rp-route-map msdp2_rp_map** – This command accepts all source-group pairs except those associated with RP 2.2.42.3.

CLI syntax

Syntax: [no] **sa-filter in** <ip-addr> [**route-map** <map-tag>] [**rp-route-map** <rp-map-tag>]

The <ip-addr> parameter specifies the IP address of the MSDP neighbor. The filter applies to Active-Source messages received from this neighbor.

The **route-map** <map-tag> parameter specifies a route map. The Layer 3 Switch applies the filter to source-group pairs that match the route map. Use the **match ip address** <acl-id> command in the route map to specify an extended ACL that contains the source and group addresses.

The **rp-route-map** <rp-map-tag> parameter specifies a route map to use for filtering based on Rendezvous Point (RP) address. Use this parameter if you want to filter Source-Active messages based on their origin. If you use the **route-map** parameter instead, messages are filtered based on source-group pairs but not based on origin. Use the **match ip route-source** <acl-id> command in the route map to specify the RP address.

NOTE

The default filter action is deny. If you want to permit some source-group pairs, use a route map. A permit action in the route map allows the Layer 3 Switch to receive the matching source-group pairs. A deny action in the route map drops the matching source-group pairs.

Filtering advertised source-active messages

The following example configures the Layer 3 Switch to advertise all source-group pairs except the ones that have source address 10.x.x.x.

Example

The following commands configure an IP address on port 3/1. This is the port on which the MSDP neighbors will be configured.

```
PowerConnect(config)# interface ethernet 3/1
PowerConnect(config-if-3/1)# ip address 2.2.2.98/24
PowerConnect(config-if-3/1)# exit
```

The following commands configure a loopback interface. The Layer 3 Switch will use this interface as the source address for communicating with the MSDP neighbors.

```
PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)# ip address 9.9.9.8/32
PowerConnect(config-lbif-1)# exit
```

The following command configures an extended ACL to specify the source and group addresses you want to filter.

```
PowerConnect(config)# access-list 123 permit 10.0.0.0 0.255.255.255 any
```

The following commands configure a route map. The map matches on source address 10.x.x.x and any group address. Since the action is deny, the Source-Active filter that uses this route map will remove the source-group pairs that match this route map from the Source-Active messages to the neighbor.

```
PowerConnect(config)# route-map msdp_map deny 1
PowerConnect(config-routemap msdp_map)# match ip address 123
PowerConnect(config-routemap msdp_map)# exit
```

The following commands enable MSDP and configure MSDP neighbors.

```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.99 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 2.2.2.97 connect-source loopback 1
PowerConnect(config-msdp-router)# exit
```

The following commands configure the Source-Active filter.

```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# sa-filter originate route-map msdp_map
```

This filter removes source-group pairs that match route map msdp_map from Source-Active messages before sending them to MSDP neighbors.

CLI syntax

Syntax: [no] **sa-filter originate** [**route-map** <map-tag>]

The **route-map** <map-tag> parameter specifies a route map. The Layer 3 Switch applies the filter to source-group pairs that match the route map. Use the **match ip address** <acl-id> command in the route map to specify an extended ACL that contains the source and group addresses.

NOTE

The default filter action is deny. If you want to permit some source-group pairs, use a route map. A permit action in the route map allows the Layer 3 Switch to receive the matching source-group pairs. A deny action in the route map drops the matching source-group pairs.

MSDP mesh groups

A PIM Sparse domain can have several RPs that are connected to each other to form an MSDP mesh group. To qualify as a mesh group, the RPs have to be fully meshed; that is, each RP must be connected to all peer RPs in a domain. (See Figure 90.)

A mesh group reduces the forwarding of SA messages within a domain. Instead of having every RP in a domain forward SA messages to all the RPs within that domain, only one RP forwards the SA message. Since an MSDP mesh group is fully meshed, peers do not forward SA messages received in a domain from one member to every member of the group. The RP that originated the SA or the first RP in a domain that receives the SA message is the only one that can forward the message to the members of a mesh group. An RP can forward an SA message to any MSRP router as long as that peer is farther away from the originating RP than the current MSRP router.

Figure 90 shows an example of an MSDP mesh group. In a PIM-SM mesh group the RPs are configured to be peers of each other. They can also be peers of RPs in other domains.

Configuring an MSDP mesh group

To configure an MSDP mesh group, enter commands such as the following on each device that will be included in the mesh group:

```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 1.1.3.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.4.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.2.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 17.17.17.7
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.4.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.3.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.2.1
PowerConnect(config-msdp-router)# exit
```

Syntax: [no] **mesh-group** <group-name> <peer-address>

The example configuration above reflects the configuration in Figure 90. On RP 1.1.1.1, you specify its peers within the same domain (1.1.3.1, 1.1.4.1, and 1.1.2.1).

You first configure the MSDP peers using the **msdp-peer** command to assign their IP addresses and the loopback interfaces. This information will be used as the source for sessions with the neighbor.

Next, place the MSDP peers within a domain into a mesh group. Use the **mesh-group** command. There are no default mesh groups.

The **group-name** parameter identifies the group. Enter up to 31 characters for group-name. You can have up to 4 mesh groups within a multicast network. Each mesh group can include up to 32 peers.

The **peer-address** parameter specifies the IP address of the MSDP peer that is being placed in the group.

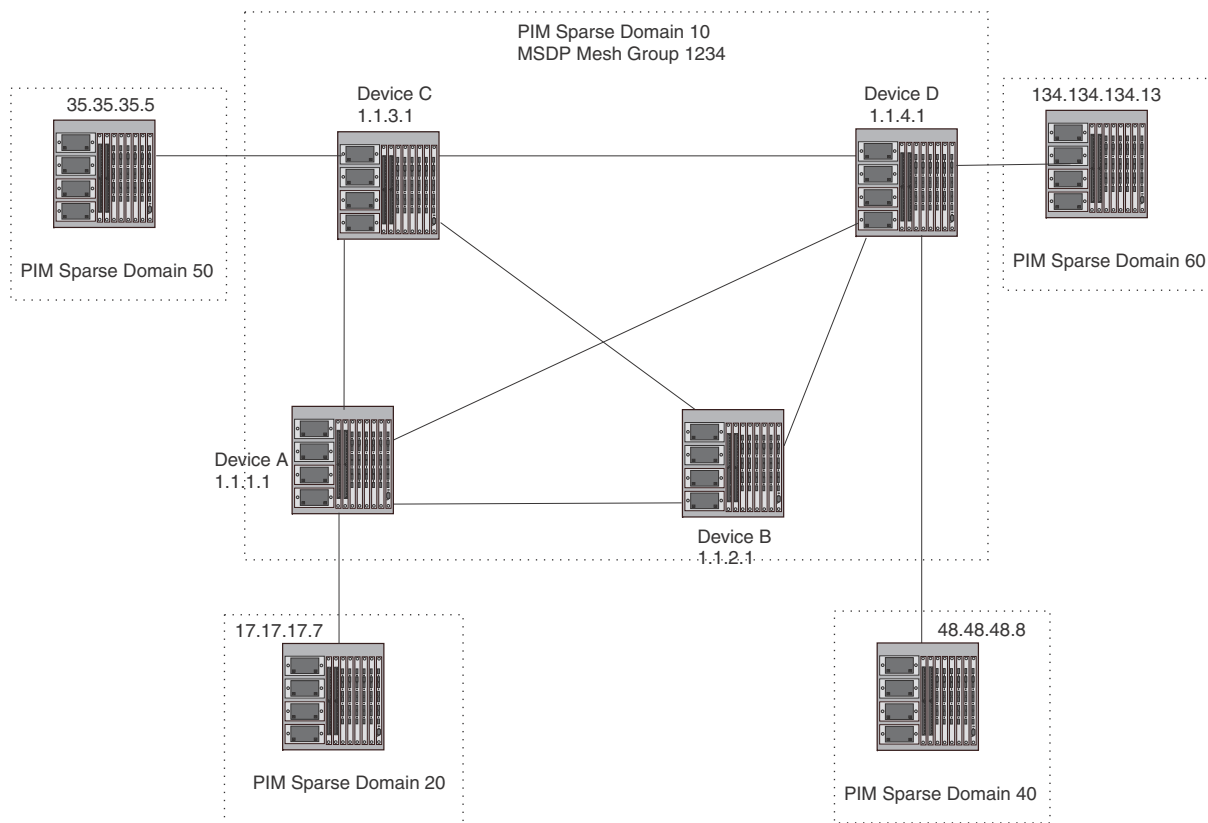
NOTE

On each of the device that will be part of the mesh-group, there must be a mesh-group definition for all the peers in the mesh-group.

Up to 32 MSDP peers can be configured per mesh group.

Example configuration

In Figure 90, devices A, B, C, and D are in Mesh Group 1234. The example configuration following the figure shows how the devices are configured to be part of the MSDP mesh group. The example also shows the features that need to be enabled for the MSDP mesh group to work.

FIGURE 90 MSDP mesh group 1234**Configuration for Device A**

The following set of commands configure the MSDP peers of Device A (1.1.1.1) that are inside and outside MSDP mesh group 1234. Device A peers inside the mesh group 1234 are 1.1.2.1, 1.1.3.1, and 1.1.4.1. Device 17.17.17.7 is a peer of Device A, but is outside mesh group 1234. Multicast is enabled on Device A interfaces. PIM and BGP are also enabled.

```
PowerConnect(config)# router pim

PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 1.1.3.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.4.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.2.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 17.17.17.7
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.4.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.3.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.2.1
PowerConnect(config-msdp-router)# exit
```

```

PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)#ip address 1.1.1.1 255.255.255.0
PowerConnect(config-lbif-1)# ip pim-sparse
PowerConnect(config-lbif-1)# exit

PowerConnect(config)# interface ethernet 1/1
PowerConnect(config-if-1/1)# ip address 14.14.14.1 255.255.255.0
PowerConnect(config-if-1/1)# ip pim-sparse
PowerConnect(config-if-1/1)# exit

PowerConnect(config)# interface ethernet 2/1
PowerConnect(config-if-2/1)# ip address 12.12.12.1 255.255.255.0
PowerConnect(config-if-2/1)# ip pim-sparse
PowerConnect(config-if-2/1)# exit

PowerConnect(config)# interface ethernet 2/20
PowerConnect(config-if-2/20)# ip address 159.159.159.1 255.255.255.0
PowerConnect(config-if-2/20)# ip pim-sparse
PowerConnect(config-if-2/20)# exit

PowerConnect(config)# interface ethernet 4/1
PowerConnect(config-if-4/1)# ip address 31.31.31.1 255.255.255.0
PowerConnect(config-if-4/1)# ip pim-sparse
PowerConnect(config-if-4/1)# exit

PowerConnect(config)# interface ethernet 4/8
PowerConnect(config-if-4/8)# ip address 17.17.17.1 255.255.255.0
PowerConnect(config-if-4/8)# ip pim-sparse
PowerConnect(config-if-4/8)# ip pim border
PowerConnect(config-if-4/8)# exit

PowerConnect(config)# router pim
PowerConnect(config-router-pim)# bsr-candidate loopback 1 1 31
PowerConnect(config-router-pim)# rp-candidate loopback 1
PowerConnect(config-router-pim)# exit

PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# local-as 111
PowerConnect(config-bgp-router)# neighbor 31.31.31.3 remote-as 333
PowerConnect(config-bgp-router)# neighbor 31.31.31.3 next-hop-self
PowerConnect(config-bgp-router)# neighbor 12.12.12.2 remote-as 222
PowerConnect(config-bgp-router)# neighbor 12.12.12.2 next-hop-self
PowerConnect(config-bgp-router)# neighbor 14.14.14.4 remote-as 444
PowerConnect(config-bgp-router)# neighbor 14.14.14.4 next-hop-self
PowerConnect(config-bgp-router)# neighbor 17.17.17.7 remote-as 777
PowerConnect(config-bgp-router)# neighbor 17.17.17.7 next-hop-self
PowerConnect(config-bgp-router)# redistribute connected
PowerConnect(config-bgp-router)# write memory

```

Configuration for Device B

The following set of commands configure the MSDP peers of Device B. All Device B peers (1.1.1.1, 1.1.3.1, and 1.1.4.1) are in the MSDP mesh group 1234. Multicast is enabled on Device B interfaces. PIM and BGP are also enabled.

```
PowerConnect(config)# router pim
```

19 Multicast Source Discovery Protocol (MSDP)

```
PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 1.1.3.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.1.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.4.1 connect-source loopback 1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.1.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.3.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.4.1
PowerConnect(config-msdp-router)# exit

PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)# ip address 1.1.2.1 255.255.255.0
PowerConnect(config-lbif-1)# ip pim-sparse
PowerConnect(config-lbif-1)# exit

PowerConnect(config)# interface ethernet 1/1
PowerConnect(config-if-1/1)# ip address 12.12.12.2 255.255.255.0
PowerConnect(config-if-1/1)# ip pim-sparse
PowerConnect(config-if-1/1)# exit

PowerConnect(config)# interface ethernet 1/12
PowerConnect(config-if-1/12)# ip address 165.165.165.1 255.255.255.0
PowerConnect(config-if-1/12)# ip pim-sparse
PowerConnect(config-if-1/12)# exit

PowerConnect(config)# interface ethernet 1/24
PowerConnect(config-if-1/24)# ip address 168.72.2.2 255.255.255.0
PowerConnect(config-if-1/24)# exit

PowerConnect(config)# interface ethernet 1/25
PowerConnect(config-if-1/25)# ip address 24.24.24.2 255.255.255.0
PowerConnect(config-if-1/25)# ip pim-sparse
PowerConnect(config-if-1/24)# exit

PowerConnect(config)# interface ethernet 8/1
PowerConnect(config-if-8/1)# ip address 32.32.32.2 255.255.255.0
PowerConnect(config-if-8/1)# ip pim-sparse
PowerConnect(config-if-1/24)# exit

PowerConnect(config)# router pim
PowerConnect(config-router-pim)# bsr-candidate loopback 1 2 32
PowerConnect(config-router-pim)# rp-candidate loopback 1
PowerConnect(config-router-pim)# exit

PowerConnect(config)# router bgp
PowerConnect(config-router-bgp)# local-as 222
PowerConnect(config-router-bgp)# neighbor 32.32.32.3 remote-as 333
PowerConnect(config-router-bgp)# neighbor 32.32.32.3 next-hop-self
PowerConnect(config-router-bgp)# neighbor 24.24.24.4 remote-as 444
PowerConnect(config-router-bgp)# neighbor 24.24.24.4 next-hop-self
PowerConnect(config-router-bgp)# neighbor 12.12.12.1 remote-as 111
PowerConnect(config-router-bgp)# neighbor 12.12.12.1 next-hop-self
PowerConnect(config-router-bgp)# redistribute connected
PowerConnect(config-router-bgp)# write memory
```

Configuration for Device C

The following set of commands configure the MSDP peers of Device C (1.1.3.1) that are inside and outside MSDP mesh group 1234. Device C peers inside the mesh group 1234 are 1.1.1.1, 1.1.2.1, and 1.1.4.1. Device 35.35.35.5 is a peer of Device C, but is outside mesh group 1234. Multicast is enabled on Device C interfaces. PIM and BGP are also enabled. This configuration is not applicable to PowerConnect device because BGP is not supported on the device.

```
PowerConnect(config)# router pim

PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 35.35.35.5
PowerConnect(config-msdp-router)# msdp-peer 1.1.2.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.4.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.1.1 connect-source loopback 1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.2.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.1.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.4.1
PowerConnect(config-msdp-router)# exit

PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-1)# ip address 1.1.3.1 255.255.255.0
PowerConnect(config-lbif-1)# ip pim-sparse
PowerConnect(config-lbif-1)# exit

PowerConnect(config)# interface ethernet 3/1
PowerConnect(config-if-3/1)# ip address 32.32.32.3 255.255.255.0
PowerConnect(config-if-3/1)# ip pim-sparse
PowerConnect(config-if-3/1)# exit

PowerConnect(config)# interface ethernet 10/1
PowerConnect(config-if-10/1)# ip address 31.31.31.3 255.255.255.0
PowerConnect(config-if-10/1)# ip pim-sparse
PowerConnect(config-if-10/1)# exit

PowerConnect(config)# interface ethernet 10/8
PowerConnect(config-if-10/8)# ip address 35.35.35.3 255.255.255.0
PowerConnect(config-if-10/8)# ip pim-sparse
PowerConnect(config-if-10/8)# ip pim border
PowerConnect(config-if-10/8)# exit

PowerConnect(config)# interface ethernet 12/2
PowerConnect(config-if-12/1)# ip address 34.34.34.3 255.255.255.0
PowerConnect(config-if-12/1)# ip pim-sparse
PowerConnect(config-if-12/1)# exit

PowerConnect(config)# interface ethernet 14/4
PowerConnect(config-if-14/4)# ip address 154.154.154.1 255.255.255.0
PowerConnect(config-if-12/1)# ip pim-sparse
PowerConnect(config-if-12/1)# exit

PowerConnect(config)# router pim
PowerConnect(config-router-pim)# bsr-candidate loopback 1 1 3
PowerConnect(config-router-pim)# rp-candidate loopback 1
PowerConnect(config-router-pim)# exit

PowerConnect(config)# router bgp
PowerConnect(config-router-bsr)# local-as 333
PowerConnect(config-router-bsr)# neighbor 35.35.35.5 remote-as 555
PowerConnect(config-router-bsr)# neighbor 35.35.35.5 next-hop-self
PowerConnect(config-router-bsr)# neighbor 32.32.32.2 remote-as 222
```

```

PowerConnect(config-router-bsr)# neighbor 32.32.32.2 next-hop-self
PowerConnect(config-router-bsr)# neighbor 34.34.34.4 remote-as 444
PowerConnect(config-router-bsr)# neighbor 34.34.34.4 next-hop-self
PowerConnect(config-router-bsr)# neighbor 31.31.31.1 remote-as 111
PowerConnect(config-router-bsr)# neighbor 31.31.31.1 next-hop-self
PowerConnect(config-router-bsr)# redistribute connected
PowerConnect(config-router-bsr)# write memory

```

Configuration for Device D

The following set of commands configure the MSDP peers of Device D (1.1.4.1) that are inside and outside MSDP mesh group 1234. Device D peers inside the mesh group 1234 are 1.1.1.1, 1.1.2.1, and 1.1.3.1. Device 48.48.48.8 and 134.134.134.13 are also peers of Device D, but are outside mesh group 1234. Multicast is enabled on Device D interfaces. PIM and BGP are also enabled. This configuration is not applicable to PowerConnect device because BGP is not supported on the device.

```

PowerConnect(config)# router pim

PowerConnect(config)# router msdp
PowerConnect(config-msdp-router)# msdp-peer 1.1.3.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.1.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 1.1.2.1 connect-source loopback 1
PowerConnect(config-msdp-router)# msdp-peer 48.48.48.8
PowerConnect(config-msdp-router)# msdp-peer 134.134.134.13
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.1.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.3.1
PowerConnect(config-msdp-router)# mesh-group 1234 1.1.2.1
PowerConnect(config-msdp-router)# exit
PowerConnect(config)# interface loopback 1
PowerConnect(config-lbif-)# ip address 1.1.4.1 255.255.255.0
PowerConnect(config-lbif-)# ip pim-sparse
PowerConnect(config-lbif-)# exit
PowerConnect(config)# interface ethernet 1/1
PowerConnect(config-if-)# ip address 24.24.24.4 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# exit
PowerConnect(config)# interface ethernet 2/6
PowerConnect(config-if-)# ip address 156.156.156.1 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# exit
PowerConnect(config)# interface ethernet 5/1
PowerConnect(config-if-)# ip address 34.34.34.4 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# exit
PowerConnect(config)# interface ethernet 7/1
PowerConnect(config-if-)# ip address 14.14.14.4 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# exit
PowerConnect(config)# interface ethernet 7/7
PowerConnect(config-if-)# ip address 48.48.48.4 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# ip pim border
PowerConnect(config-if-)# exit
PowerConnect(config)# interface ethernet 7/8
PowerConnect(config-if-)# ip address 134.134.134.4 255.255.255.0
PowerConnect(config-if-)# ip pim-sparse
PowerConnect(config-if-)# ip pim border
PowerConnect(config-if-)# exit

```



```

PowerConnect(config)# router pim
PowerConnect(config-router-pim)# bsr-candidate loopback 1 14 34
PowerConnect(config-router-pim)# rp-candidate loopback 1
PowerConnect(config-router-pim)# exit
PowerConnect(config)# router bgp
PowerConnect(config-router-bsr)# local-as 444
PowerConnect(config-router-bsr)# neighbor 34.34.34.3 remote-as 333
PowerConnect(config-router-bsr)# neighbor 34.34.34.3 next-hop-self
PowerConnect(config-router-bsr)# neighbor 14.14.14.1 remote-as 111
PowerConnect(config-router-bsr)# neighbor 14.14.14.1 next-hop-self
PowerConnect(config-router-bsr)# neighbor 24.24.24.2 remote-as 222
PowerConnect(config-router-bsr)# neighbor 24.24.24.2 next-hop-self
PowerConnect(config-router-bsr)# neighbor 48.48.48.8 remote-as 888
PowerConnect(config-router-bsr)# neighbor 48.48.48.8 next-hop-self
PowerConnect(config-router-bsr)# neighbor 134.134.134.13 remote-as 1313
PowerConnect(config-router-bsr)# neighbor 134.134.134.13 next-hop-self
PowerConnect(config-router-bsr)# redistribute connected
PowerConnect(config-router-bsr)# write memory

```

Displaying MSDP information

You can display the following MSDP information:

- Summary information – the IP addresses of the peers, the state of the Layer 3 Switch MSDP session with each peer, and statistics for Keepalive, Source Active, and Notification messages sent to and received from each of the peers
- Peer information – the IP address of the peer, along with detailed MSDP and TCP statistics
- Source Active cache entries – the Source Active messages cached by the Layer 3 Switch

Displaying summary information

To display summary MSDP information, enter the following command at any level of the CLI:

```
PowerConnect(config-msdp-router)# show ip msdp summary
```

```

MSDP Peer Status Summary
KA: Keepalive SA:Source-Active NOT: Notification
Peer Address      State           In   Out   In   Out   In   Out
206.251.17.30    ESTABLISH      3    3     0   640   0    0
206.251.17.41    ESTABLISH      0    3    651  0     0    0

```

Syntax: show ip msdp summary

This display shows the following information.

TABLE 87 MSDP summary information

This field...	Displays...
Peer Address	The IP address of the peer interface with the Layer 3 Switch
State	The state of the MSDP router connection with the peer. The state can be one of the following: <ul style="list-style-type: none"> • CONNECT – The session is in the active open state. • ESTABLISH – The MSDP session is fully up. • IDLE – The session is idle or inactive. • LISTEN – The session is in the passive open state.
KA In	The number of MSDP Keepalive messages the MSDP router has received from the peer
KA Out	The number of MSDP Keepalive messages the MSDP router has sent to the peer
SA In	The number of Source Active messages the MSDP router has received from the peer
SA Out	The number of Source Active messages the MSDP router has sent to the peer
NOT In	The number of Notification messages the MSDP router has received from the peer
NOT Out	The number of Notification messages the MSDP router has sent to the peer

Displaying peer information

To display summary MSDP peer information, use the following CLI method.

```
PowerConnect(config-msdp-router)# show ip msdp peer

Total number of MSDP Peers: 2

  IP Address      State
  1 206.251.17.30 ESTABLISHED
  Keep Alive Time Hold Time
  60              90

                Message Sent      Message Received
Keep Alive      2              3
Notifications   0              0
Source-Active   0              640
Last Connection Reset Reason:Reason Unknown
Notification Message Error Code Received:Unspecified
Notification Message Error SubCode Received:Not Applicable
Notification Message Error Code Transmitted:Unspecified
Notification Message Error SubCode Transmitted:Not Applicable
TCP Connection state: ESTABLISHED
Local host: 206.251.17.29, Local Port: 8270
Remote host: 206.251.17.30, Remote Port: 639
ISentSeq:      16927  SendNext:      685654  TotUnAck:      0
SendWnd:       16384  TotSent:       668727  ReTrans:       1
IRcvSeq:      45252428  RcvNext:      45252438  RcvWnd:       16384
TotalRcv:      10    RcvQue:        0    SendQue:       0
```

Syntax: show ip msdp peer

This display shows the following information.

TABLE 88 MSDP peer information

This field...	Displays...
Total number of MSDP peers	The number of MSDP peers configured on the Layer 3 Switch
IP Address	The IP address of the peer interface with the Layer 3 Switch
State	The state of the MSDP router connection with the peer. The state can be one of the following: <ul style="list-style-type: none"> • CONNECT – The session is in the active open state. • ESTABLISH – The MSDP session is fully up. • IDLE – The session is idle or inactive. • LISTEN – The session is in the passive open state.
Keep Alive Time	The keep alive time, which specifies how often this MSDP router sends keep alive messages to the neighbor. The keep alive time is 60 seconds and is not configurable.
Hold Time	The hold time, which specifies how many seconds the MSDP router will wait for a KEEPALIVE or UPDATE message from an MSDP neighbor before deciding that the neighbor is dead. The hold time is 90 seconds and is not configurable.
Keep Alive Message Sent	The number of Keep Alive messages the MSDP router has sent to the peer.
Keep Alive Message Received	The number of Keep Alive messages the MSDP router has received from the peer.
Notifications Sent	The number of Notification messages the MSDP router has sent to the peer.
Notifications Received	The number of Notification messages the MSDP router has received from the peer.
Source-Active Sent	The number of Source Active messages the MSDP router has sent to the peer.
Source-Active Received	The number of Source Active messages the MSDP router has received from the peer.
Last Connection Reset Reason	The reason the previous session with this neighbor ended.
Notification Message Error Code Received	If the MSDP router receives a NOTIFICATION messages from the neighbor, the message contains an error code corresponding to one of the following errors. Some errors have subcodes that clarify the reason for the error. Where applicable, the subcode messages are listed underneath the error code messages: <ul style="list-style-type: none"> • 1 – Message Header Error • 2 – SA-Request Error • 3 – SA-Message/SA-Response Error • 4 – Hold Timer Expired • 5 – Finite State Machine Error • 6 – Notification • 7 – Cease For information about these error codes, see section 17 in the Internet draft describing MSDP, “draft-ietf-msdp-spec”.
Notification Message Error SubCode Received	See above.
Notification Message Error Code Transmitted	The error message corresponding to the error code in the NOTIFICATION message this MSDP router sent to the neighbor. See the description for the Notification Message Error Code Received field for a list of possible codes.

TABLE 88 MSDP peer information (Continued)

This field...	Displays...
Notification Message Error SubCode Transmitted	See above.
TCP Statistics	
TCP connection state	<p>The state of the connection with the neighbor. The connection can have one of the following states:</p> <ul style="list-style-type: none"> • LISTEN – Waiting for a connection request. • SYN-SENT – Waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED – Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED – Data can be sent and received over the connection. This is the normal operational state of the connection. • FIN-WAIT-1 – Waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 – Waiting for a connection termination request from the remote TCP. • CLOSE-WAIT – Waiting for a connection termination request from the local user. • CLOSING – Waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK – Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request). • TIME-WAIT – Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request. • CLOSED – There is no connection state.
Local host	The IP address of the MSDP router interface with the peer.
Local port	The TCP port the MSDP router is using for the BGP4 TCP session with the neighbor.
Remote host	The IP address of the neighbor.
Remote port	The TCP port number of the peer end of the connection.
ISentSeq	The initial send sequence number for the session.
SendNext	The next sequence number to be sent.
TotUnAck	The number of sequence numbers sent by the MSDP router that have not been acknowledged by the neighbor.
SendWnd	The size of the send window.
TotSent	The number of sequence numbers sent to the neighbor.
ReTrans	The number of sequence numbers that the MSDP router retransmitted because they were not acknowledged.
IRcvSeq	The initial receive sequence number for the session.
RcvNext	The next sequence number expected from the neighbor.
RcvWnd	The size of the receive window.
TotalRcv	The number of sequence numbers received from the neighbor.

TABLE 88 MSDP peer information (Continued)

This field...	Displays...
RcvQue	The number of sequence numbers in the receive queue.
SendQue	The number of sequence numbers in the send queue.

Displaying source active cache information

To display the Source Actives in the MSDP cache, use the following CLI method.

```
PowerConnect(config-msdp-router)# show ip msdp sa-cache
```

```
Total Entry 4096, Used 1800 Free 2296
Index  SourceAddr  GroupAddr  Age
1      (100.100.1.254, 232.1.0.95), RP:206.251.17.41, Age:0
2      (100.100.1.254, 237.1.0.98), RP:206.251.17.41, Age:30
3      (100.100.1.254, 234.1.0.48), RP:206.251.17.41, Age:30
4      (100.100.1.254, 239.1.0.51), RP:206.251.17.41, Age:30
5      (100.100.1.254, 234.1.0.154), RP:206.251.17.41, Age:30
6      (100.100.1.254, 236.1.0.1), RP:206.251.17.41, Age:30
7      (100.100.1.254, 231.1.0.104), RP:206.251.17.41, Age:90
8      (100.100.1.254, 239.1.0.157), RP:206.251.17.41, Age:30
9      (100.100.1.254, 236.1.0.107), RP:206.251.17.41, Age:30
10     (100.100.1.254, 233.1.0.57), RP:206.251.17.41, Age:90
```

Syntax: show ip msdp sa-cache

This display shows the following information.

TABLE 89 MSDP source active cache

This field...	Displays...
Total Entry	The total number of entries the cache can hold.
Used	The number of entries the cache currently contains.
Free	The number of additional entries for which the cache has room.
Index	The cache entry number.
SourceAddr	The IP address of the multicast source.
GroupAddr	The IP multicast group to which the source is sending information.
RP	The RP through which receivers can access the group traffic from the source
Age	The number of seconds the entry has been in the cache

Clearing MSDP information

You can clear the following MSDP information:

- Peer information
- Source Active cache
- MSDP statistics

Clearing peer information

To clear MSDP peer information, enter the following command at the Privileged EXEC level of the CLI:

```
PowerConnect# clear ip msdp peer 205.216.162.1
Remote connection closed
```

Syntax: `clear ip msdp peer <ip-addr>`

The command in this example clears the MSDP peer connection with MSDP router 205.216.162.1. The CLI displays a message to indicate when the connection has been successfully closed.

Clearing the source active cache

To clear the entries from the Source Active cache, enter the following command at the Privileged EXEC level of the CLI:

```
PowerConnect# clear ip msdp sa-cache
```

Syntax: `clear ip msdp sa-cache [<source-addr> | <group-addr>]`

The command in this example clears all the cache entries. Use the `<source-addr>` parameter to clear only the entries for a specified source. Use the `<group-addr>` parameter to clear only the entries for a specific group.

Clearing MSDP statistics

To clear MSDP statistics, enter the following command at the Privileged EXEC level of the CLI:

```
PowerConnect# clear ip msdp statistics
```

Syntax: `clear ip msdp statistics [<ip-addr>]`

The command in this example clears statistics for all the peers. To clear statistics for a specific peer, enter the peer IP address.

Using ACLs to control multicast features

You can use ACLs to control the following multicast features:

- Limit the number of multicast groups that are covered by a static rendezvous point (RP)
- Control which multicast groups for which candidate RPs sends advertisement messages to bootstrap routers
- Identify which multicast group packets will be forwarded or blocked on an interface

Using ACLs to limit static RP groups

You can limit the number of multicast groups covered by a static RP using standard ACLs. In the ACL, you specify the group to which the RP address applies. The following examples set the RP address to be applied to multicast groups with some minor variations.

To configure an RP that covers multicast groups in 239.255.162.x, enter commands such as the following.

```
PowerConnect(config)#access-list 2 permit 239.255.162.0 0.0.0.255
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#rp-address 43.43.43.1 2
```

To configure an RP that covers multicast groups in the 239.255.162.x range, except the 239.255.162.2 group, enter commands such as the following.

```
PowerConnect(config)#access-list 5 deny host 239.255.162.2
PowerConnect(config)#access-list 5 permit 239.255.0.0 0.0.255.255
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#bsr-candidate ve 43 32 100
PowerConnect(config-pim-router)#rp-candidate ve 43
PowerConnect(config-pim-router)#rp-address 99.99.99.5 5
```

To configure an RP for multicast groups using the override switch, enter commands such as the following.

```
PowerConnect(config)#access-list 44 permit 239.255.162.0 0.0.0.255
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#rp-address 43.43.43.1
PowerConnect(config-pim-router)#rp-address 99.99.99.5 44 override
```

Syntax: [no] **rp-address** <ip-address> [<access-list-num>] [override]

The access-list-num parameter is the number of the standard ACL that will filter the multicast group.

NOTE

Extended ACLs cannot be used to limit static RP groups.

The **override** parameter directs the Layer 3 Switch to ignore the information learned by a BSR if there is a conflict between the RP configured in this command and the information that is learned by the BSR. In previous releases, static RP configuration precedes the RP address learned from the PIM Bootstrap protocol. With this enhancement, an RP address learned dynamically from PIM Bootstrap protocol takes precedence over static RP configuration unless the override parameter is used.

You can use the **show ip pim rp-set** command to display the ACLs used to filter the static RP groups.

Example

```
PowerConnect#show ip pim rp-set
Group address      Static-RP-address  Override
-----
Access-List 44    99.99.99.5         On
Number of group prefixes Learnt from BSR: 1
Group prefix = 224.0.0.0/4 #RPs: 1
  RP 1: 43.43.43.1 priority=0 age=0
```

In the example above, the display shows the following information:

- The Group Address table shows the static RP address that is covered by the access list, and whether or not the override parameter has been enabled.
- The Group prefix line shows the multicast group prefix for the static RP.
- The RP #line shows the configured IP address of the RP candidate.

The **show ip pim rp-map** to show the group-to-RP mapping.

```
PowerConnect#show ip pim rp-map
Number of group-to-RP mappings: 6
  Group address  RP address
-----
1 239.255.163.1  43.43.43.1
2 239.255.163.2  43.43.43.1
3 239.255.163.3  43.43.43.1
4 239.255.162.1  99.99.99.5
5 239.255.162.2  99.99.99.5
6 239.255.162.3  99.99.99.5
```

The display shows the multicast group addresses covered by the RP candidate and the IP address of the RP for the listed multicast group. In the example above, you see the following:

- The first three lines show the multicast group addresses that are covered by the RP candidate.
- The last three lines show the multicast group addresses covered by the static RP.

Using ACLs to limit PIM RP candidate advertisement

You can use standard ACLs to control the groups for which the candidate RP will send advertisement messages to the bootstrap router. For example, ACL 5 can be configured to be applied to the multicast groups within the IP address 239.x.x.x range. You can configure the Layer 3 Switch to advertise itself as a candidate RP to the bootstrap router only for groups in the range of 239.x.x.x. Enter commands such as the following.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#ip address 99.99.99.5 255.255.255.0
PowerConnect(config-if-1)#ip pim-sparse
PowerConnect(config-if-1)#exit
PowerConnect(config)#access-list 5 deny host 239.255.162.2
PowerConnect(config)#access-list 5 permit 239.0.0.0 0.0.255.255
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#bsr-candidate ethernet 1 32 100
PowerConnect(config-pim-router)#rp-candidate ethernet 1 group-list 5
```

The example above shows a configuration for an Ethernet interface. To configure ACLs that are applied to a virtual routing interface, enter commands such as the following.

```
PowerConnect(config)#interface ve 16
PowerConnect(config-vif-16)#ip address 16.16.16.1 255.255.255.0
PowerConnect(config-vif-16)#ip pim-sparse
PowerConnect(config-vif-16)#exit
PowerConnect(config)#access-list 5 deny host 239.255.162.2
PowerConnect(config)#access-list 5 permit 239.255.0.0 0.0.255.255
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#bsr-candidate ve 16 32 100
PowerConnect(config-pim-router)#rp-candidate ve 16 group-list 5
```

To configure ACLs that are applied to a loopback interface, enter commands such as the following.

```
PowerConnect(config)#interface loopback 1
PowerConnect(config-lbif-1)#ip address 88.88.88.8 255.255.255.0
PowerConnect(config-lbif-1)#ip pim-sparse
PowerConnect(config-lbif-1)#exit
PowerConnect(config)#access-list 5 deny host 239.255.162.2
PowerConnect(config)#access-list 5 permit 239.255.0.0 0.0.255.255
```



```
PowerConnect(config)#router pim
PowerConnect(config-pim-router)#bsr-candidate loopback 1 32 100
PowerConnect(config-pim-router)#rp-candidate loopback 1 group-list 5
```

Syntax: **[no] rp-candidate ethernet** <portnum> | **loopback** <num> | **ve** <num> [**group-list** <access-list-num>]

The <portnum> | **loopback** <num> | **ve** <num> parameter specifies the interface. The Layer 3 Switch will advertise the specified interface IP address as a candidate RP:

- Enter **ethernet** <portnum> for a physical interface (port).
- Enter **ve** <num> for a virtual interface.
- Enter **loopback** <num> for a loopback interface.

The **group-list** <access-list-num> indicates that a standard ACL is used to filter for which multicast group the advertisement will be made.

NOTE

Extended ACLs cannot be used for group-list.

Configuring a static multicast route

Static multicast routes allow you to control the network path used by multicast traffic. Static multicast routes are especially useful when the unicast and multicast topologies of a network are different. You can avoid the need to make the topologies similar by instead configuring static multicast routes.

You can configure more than one static multicast route. The Layer 3 Switch always uses the most specific route that matches a multicast source address. Thus, if you want to configure a multicast static route for a specific multicast source and also configure another multicast static route for all other sources, you can configure two static routes as shown in the examples below.

To add static routes to multicast router A (refer to [Figure 91](#)), enter commands such as the following.

```
PIMRouterA(config)#ip mroute 1 207.95.10.0 255.255.255.0 interface ethernet 2
distance 1
PIMRouterA(config)#ip mroute 2 0.0.0.0 0.0.0.0 interface ethernet 3 distance 1
PIMRouterA(config)#write memory
```

Syntax: **mroute** <route-num> <ip-addr> **interface ethernet** <portnum> | **ve** <num> [**distance** <num>]

or

Syntax: **mroute** <route-num> <ip-addr> **rpf_address** <rpf-num>

The <route-num> parameter specifies the route number.

The <ip-addr> command specifies the PIM source for the route.

NOTE

In IP multicasting, a route is handled in terms of its source, rather than its destination.

You can use the **ethernet** <portnum> parameter to specify a physical port or the **ve** <num> parameter to specify a virtual interface.

NOTE

The **ethernet** *<portnum>* parameter does not apply to PIM SM.

The **distance** *<num>* parameter sets the administrative distance for the route. When comparing multiple paths for a route, the Layer 3 Switch prefers the path with the lower administrative distance.

NOTE

Regardless of the administrative distances, the Layer 3 Switch always prefers directly connected routes over other routes.


The **rpf_address** *<rpf-num>* parameter specifies an RPF number.

The example above configures two static multicast routes. The first route is for a specific source network, 207.95.10.0/24. If the Layer 3 Switch receives multicast traffic for network 207.95.10.0/24, the traffic must arrive on port 2. The second route is for all other multicast traffic. Traffic from multicast sources other than 207.95.10.0/24 must arrive on port 3.

Figure 91 shows an example of an IP Multicast network. The two static routes configured in the example above apply to this network. The commands in the example above configure PIM router A to accept PIM packets from 207.95.10.0/24 when they use the path that arrives at port 2, and accept all other PIM packets only when they use the path that arrives at port 3.

The distance parameter sets the administrative distance. This parameter is used by the software to determine the best path for the route. Thus, to ensure that the Layer 3 Switch uses the default static route, assign a low administrative distance value. When comparing multiple paths for a route, the Layer 3 Switch prefers the path with the lower administrative distance.

FIGURE 91 Example of multicast static routes



To add a static route to a virtual interface, enter commands such as the following.

```
PowerConnect(config)#mroute 3 0.0.0.0 0.0.0.0 int ve 1 distance 1
PowerConnect(config)#write memory
```

Tracing a multicast route

The Dell implementation of Mtrace is based on “A ‘traceroute’ facility for IP Multicast”, an Internet draft by S. Casner and B. Fenner. To trace a PIM route, use the following CLI method..

To trace a PIM route to PIM source 209.157.24.62 in group 239.255.162.1, enter a command such as the following.

```
PowerConnect#mtrace source 209.157.24.62 group 239.255.162.1
```

```
Type Control-c to abort
```

```
Tracing the route for tree 209.157.23.188
```

```
0 207.95.7.2
0 207.95.7.2 Thresh 0
1 207.95.7.1 Thresh 0
2 207.95.8.1 Thresh 0
3 207.157.24.62
```

Syntax: `mtrace source <ip-addr> group <multicast-group>`

The **source <ip-addr>** parameter specifies the address of the route source.


NOTE

In IP multicasting, a route is handled in terms of its source, rather than its destination. When you trace an IP route, you specify its destination, but when you trace a PIM route, you specify its source.

The **group <multicast-group>** parameter specifies the PIM group the source IP address is in.

Figure 92 shows an example of an IP multicast group. The command example shown above is entered on PIM router A.

FIGURE 92 Example of a PIM group



The command example above indicates that the source address 209.157.24.62 is three hops (three PIM switches) away from PIM Switch A. In PIM terms, each of the three switches has a forwarding state for the specified source address and multicast group. The value following “Thresh” in some of the lines indicates the TTL threshold. The threshold 0 means that all multicast packets are forwarded on the interface. If an administrator has set the TTL threshold to a higher value, only packets whose TTL is higher than the threshold are forwarded on the interface. The threshold is listed only for the PIM switch hops between the source and destination.

Displaying the multicast configuration for another multicast router

To display another PIM router PIM configuration, enter a command such as the following.

```
PowerConnect#mrintfo 207.95.8.1
207.95.8.1 -> 207.95.8.10 [PIM/0 /1]
207.95.10.2 -> 0.0.0.0 [PIM/0 /1 /leaf]
209.157.25.1 -> 0.0.0.0 [PIM/0 /1 /leaf]
209.157.24.1 -> 0.0.0.0 [PIM/0 /1 /leaf]
207.95.6.1 -> 0.0.0.0 [PIM/0 /1 /leaf]
128.2.0.1 -> 0.0.0.0 [PIM/0 /1 /leaf]
```

Syntax: mrintfo <ip-addr>

The <ip-addr> parameter specifies the IP address of the PIM router.

The output in this example is based on the PIM group shown in [Figure 92](#) on page 525. The output shows the PIM interfaces configured on PIM router C (207.95.8.1). In this example, the PIM router has six PIM interfaces. One of the interfaces goes to PIM router B. The other interfaces go to leaf nodes, which are multicast end nodes attached to the router PIM interfaces. (For simplicity, the figure shows only one leaf node.)

When the arrow following an interface in the display points to a router address, this is the address of the next hop PIM router on that interface. In this example, PIM interface 207.95.8.1 on PIM router 207.95.8.1 is connected to PIM router 207.95.8.10. The connection can be a direct one or can take place through non-PIM routers. In this example, the PIM routers are directly connected.

When the arrow following an interface address points to zeros (0.0.0.0), the interface is not connected to a PIM router. The interface is instead connected to a leaf node.

NOTE

This display shows the PIM interface configuration information, but does not show the link states for the interfaces.

The information in brackets indicates the following:

- The multicast interface type .
- The Time-to-Live (TTL) for the interface.
- The metric for the interface
- Whether the interface is connected to a leaf node (“leaf” indicates a leaf node and blank indicates another PIM router)

For example, the information for the first interface listed in the display is “PIM/0 /1”. This information indicates that the interface is a PIM interface, has a TTL of 0, and a metric of 1. The interface is not a leaf node interface and thus is an interface to another PIM router.

The information for the second interface in the display is “PIM/0 /1/leaf”. This information indicates that the interface is a PIM interface, has a TTL of 0 and a metric of 1, and is connected to a leaf node.

IGMP V3

The Internet Group Management Protocol (IGMP) allows an IPV4 interface to communicate IP Multicast group membership information to its neighboring routers. The routers in turn limit the multicast of IP packets with multicast destination addresses to only those interfaces on the router that are identified as IP Multicast group members. This release introduces the support of IGMP version 3 (IGMP V3) on Layer 3 Switches.

In IGMP V2, when a router sent a query to the interfaces, the clients on the interfaces respond with a membership report of multicast groups to the router. The router can then send traffic to these groups, regardless of the traffic source. When an interface no longer needs to receive traffic from a group, it sends a leave message to the router which in turn sends a group-specific query to that interface to see if any other clients on the same interface is still active.

In contrast, IGMP V3 provides selective filtering of traffic based on traffic source. A router running IGMP V3 sends queries to every multicast enabled interface at the specified interval. These queries determine if any interface wants to receive traffic from the router. The queries include the IP address of the traffic source (S) or the ID of the multicast group (G, or both).

The interfaces respond to these queries by sending a membership report that contains one or more of the following records that are associated with a specific group:

- Current-State Record that indicates from which sources the interface wants to receive and not receive traffic. The record contains source address of interfaces and whether or not traffic will be received or included (IS_IN) or not received or excluded (IS_EX) from that source.
- Filter-mode-change record. If the interface changes its current state from IS_IN to IS_EX, a TO_EX record is included in the membership report. Likewise, if an interface current state changes from IS_EX to IS_IN, a TO_IN record appears in the membership report.

IGMP V2 Leave report is equivalent to a TO_IN(empty) record in IGMP V3. This record means that no traffic from this group will be received regardless of the source.

An IGMP V2 group report is equivalent to an IS_EX(empty) record in IGMP V3. This record means that all traffic from this group will be received regardless of source.

- Source-List-Change Record. If the interface wants to add or remove traffic sources from its membership report, the membership report can have an ALLOW record, which contains a list of new sources from which the interface wishes to receive traffic. It can also contain a BLOCK record, which lists current traffic sources from which the interfaces wants to stop receiving traffic.

In response to membership reports from the interfaces, the router sends a Group-Specific or a Group-and-Source Specific query to the multicast interfaces. Each query is sent three times with a one-second interval in between each transmission to ensure the interfaces receive the query. For example, a router receives a membership report with a Source-List-Change record to block old sources from an interface. The router sends Group-and-Source Specific Queries to the source and group (S,G) identified in the record. If none of the interfaces is interested in the (S,G), it is removed from (S,G) list for that interface on the router.

Each IGMP V3-enabled router maintains a record of the state of each group and each physical port within a virtual routing interface. This record contains the group, group-timer, filter mode, and source records information for the group or interface. Source records contain information on the source address of the packet and source timer. If the source timer expires when the state of the group or interface is in Include mode, the record is removed.

Default IGMP version

IGMP V3 is available on devices ; however, devices are shipped with IGMP V2 enabled. You must enable IGMP V3 globally or per interface.

Also, you must specify what version of IGMP you want to run on a device globally, on each interface (physical port or virtual routing interface), and on each physical port within a virtual routing interface. If you do not specify an IGMP version, IGMP V2 will be used.

Compatibility with IGMP V1 and V2

Different multicast groups, interfaces, and routers can run their own version of IGMP. Their version of IGMP is reflected in the membership reports that the interfaces send to the router. Routers and interfaces must be configured to recognize the version of IGMP you want them to process.

An interface or router sends the queries and reports that include its IGMP version specified on it. It may recognize a query or report that has a different version, but it may not process them. For example, an interface running IGMP V2 can recognize IGMP V3 packets, but cannot process them. Also, a router running IGMP V3 can recognize and process IGMP V2 packet, but when that router sends queries to an IGMP V2 interface, the host on that interface may not recognize the IGMP V3 queries. The interface or router does not automatically downgrade the IGMP version running on them to avoid version deadlock.

If an interface continuously receives queries from routers that are running versions of IGMP that are different from what is on the interface, the interface logs warning messages in the syslog every five minutes. Reports sent by interfaces to routers that contain different versions of IGMP do not trigger warning messages; however, you can see the versions of the packets using the **show ip igmp traffic** command.

The version of IGMP can be specified globally, per interface (physical port or virtual routing interface), and per physical port within a virtual routing interface. The IGMP version set on a physical port within a virtual routing interface supersedes the version set on a physical or virtual routing interface. Likewise, the version on a physical or virtual routing interface supersedes the version set globally on the device. The sections below present how to set the version of IGMP.

Globally enabling the IGMP version

Using the CLI

To globally identify the IGMP version on a device, enter the following command.

```
PowerConnect(config)# ip igmp version 3
```

Syntax: `ip igmp version <version-number>`

Enter 1, 2, or 3 for <version-number>. Version 2 is the default version.

Enabling the IGMP version per interface setting

To specify the IGMP version for a physical port, enter a command such as the following.

```
PowerConnect(config)# interface eth 5
PowerConnect(config-if-5)# ip igmp version 3
```

To specify the IGMP version for a virtual routing interface on a physical port, enter a command such as the following.

```
PowerConnect(config)# interface ve 3
PowerConnect(config-vif-3)# ip igmp version 3
```

Syntax: [no] ip igmp version <version-number>

Enter 1, 2, or 3 for <version-number>. Version 2 is the default version.

Enabling the IGMP version on a physical port within a virtual routing interface

To specify the IGMP version recognized by a physical port that is a member of a virtual routing interface, enter a command such as the following.

```
PowerConnect(config)# interface ve 3
PowerConnect(config-vif-3)# ip igmp version 2
PowerConnect(config-vif-3)# ip igmp port-version 3 e3-e7 e9
```

In this example, the second line sets IGMP V2 on virtual routing interface 3. However, the third line set IGMP V3 on ports 3 through 7 and port e9. All other ports in this virtual routing interface are configured with IGMP V2.

Syntax: ip igmp port-version <version-number> ethernet <port-number>

Enter 1, 2, or 3 for <version-number>. IGMP V2 is the default version.

The **ethernet** <port-number> parameter specifies which physical port within a virtual routing interface is being configured.

Enabling membership tracking and fast leave

IGMP V3 provides membership tracking and fast leave to clients. In IGMP V2, only one client on an interface needs to respond to a router queries; therefore, some of the clients may be invisible to the router, making it impossible for the router to track the membership of all clients in a group. Also, when a client leaves the group, the router sends group specific queries to the interface to see if other clients on that interface need the data stream of the client who is leaving. If no client responds, the router waits three seconds before it stops the traffic.

IGMP V3 contains the tracking and fast leave feature that you enable on virtual routing interfaces. Once enabled, all physical ports on that virtual routing interface will have the feature enabled. IGMP V3 requires all clients to respond to general and group specific queries so that all clients on an interface can be *tracked*. *Fast leave* allows clients to leave the group without the three second waiting period, if the following conditions are met:

- If the interface, to which the client belongs, has IGMP V3 clients only. Therefore, all physical ports on a virtual routing interface must have IGMP V3 enabled and no IGMP V1 or V2 clients can be on the interface. (Although IGMP V3 can handle V1 and V2 clients, these two clients cannot be on the interface in order for fast leave to take effect.)
- No other client on the interface is receiving traffic from the group to which the client belongs. Every group on the physical interface of a virtual routing interface keeps its own tracking record. However, it can track group membership only; it cannot track by (source, group).

For example, two clients (Client A and Client B) belong to group1 but each is receiving traffic streams from different sources. Client A receives a stream from (source_1, group1) and Client B receives it from (source_2, group1). The router still waits for three seconds before it stops the traffic because the two clients are in the same group. If the clients are in different groups, then the three second waiting period is not applied and traffic is stopped immediately. The **show ip igmp group tracking** command displays that clients in a group that are being tracked.

If a client sends a leave message, the client is immediately removed from the group. If a client does not send a report during the the specified group membership time (the default is 140 seconds), that client is removed from the tracking list.

Using the CLI

To enable the tracking and fast leave feature, enter commands such as the following.

```
PowerConnect(config)# interface ve 13
PowerConnect(config-vif-13)# ip igmp tracking
```

Syntax: ip igmp tracking

Setting the query interval

The IGMP query interval period defines how often a router will query an interface for group membership. Possible values are 10 – 3,600 seconds and the default value is 60 seconds, but the value you enter must be a little more than twice the group membership time.

To modify the default value for the IGMP query interval, enter the following.

```
PowerConnect(config)# ip igmp query-interval 120
```

Syntax: ip igmp query-interval <10-3600>

The interval must be a little more than two times the group membership time.

Setting the group membership time

Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 20 – 7200 seconds and the default value is 140 seconds.

To define an IGMP membership time of 240 seconds, enter the following.

```
PowerConnect(config)# ip igmp group-membership-time 240
```

Syntax: ip igmp group-membership-time <20-7200>

Setting the maximum response time

The maximum response time defines the maximum number of seconds that a client can wait before it replies to the query sent by the router. Possible values are 1 – 10. The default is 10.

To change the IGMP maximum response time, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# ip igmp max-response-time 8
```

Syntax: [no] ip igmp max-response-time <num>

The *<num>* parameter specifies the maximum number of seconds for the response time. Enter a value from 1 – 10. The default is 10.

IGMP V3 and source specific multicast protocols

Enabling IGMP V3 enables source specific multicast (SSM) filtering for PIM Dense (PIM-DM) for multicast group addresses in the 224.0.1.0 through 239.255.255.255 address range. However, if PIM Sparse is used as the multicast protocol, the SSM protocol should be enabled if you want to filter unwanted traffic before the Shortest Path Tree protocol switchover occurs for groups in the 232/8 range. Not configuring the SSM protocol in PIM Sparse may cause the switch or router to leak unwanted packets with the same group, but containing undesired sources, to clients. After SPT switch over, the leak stops and source specific multicast works correctly even without configuring the SSM protocol.

If the SSM protocol is not enabled and before the SPT switchover, the multicast router creates one (*, G) entry for the entire multicast group, which can have many sources. If the SSM protocol is enabled, one (S,G) entry is created for every member of the multicast group, even for members with non-existent traffic. For example, if there are 1,000 members in the group, 1,000 (S,G) entries will be created. Therefore, enabling the SSM protocol for PIM-SM requires more resources than leaving the protocol disabled.

Enabling SSM

To enable the SSM protocol on a device running PIM-SM, enter a command such as the following.

```
PowerConnect(config)# router pim
PowerConnect(config-pim-router)# ssm-enable
```

Syntax: [no] ssm-enable

Enter the ssm-enable command under the router pim level to globally enable the SSM protocol on a Layer 3 Switch.

Displaying IGMP V3 information on Layer 3 Switches

The sections below present the show commands available for IGMP V3 on Layer 3 Switches. For show commands on Layer 2 Switches, use the **show ip multicast** commands which are discussed in the section “[IGMP snooping show commands](#)” on page 454.

Displaying IGMP group status

NOTE

This report is available on Layer 3 Switches.

To display the status of all IGMP multicast groups on a device, enter the following command.

```
PowerConnect#show ip igmp group
Interface v18 : 1 groups
  group          phy-port static querier life mode  #_src
1  239.0.0.1     e20   no    yes           include 19
Interface v110 : 3 groups
  group          phy-port static querier life mode  #_src
2  239.0.0.1     e5    no    yes           include 10
3  239.0.0.1     e6    no    yes          100  exclude 13
4  224.1.10.1    e5    no    yes           include 1
```

To display the status of one IGMP multicast group, enter a command such as the following.

```
PowerConnect#show ip igmp group 239.0.0.1 detail
Display group 239.0.0.1 in all interfaces.
Interface v18 : 1 groups
  group          phy-port static querier life mode  #_src
1  239.0.0.1     e20   no    yes           include 19
  group: 239.0.0.1, include, permit 19 (source, life):
    (3.3.3.1 40) (3.3.3.2 40) (3.3.3.3 40) (3.3.3.4 40) (3.3.3.5 40)
    (3.3.3.6 40) (3.3.3.7 40) (3.3.3.8 40) (3.3.3.9 40) (3.3.3.10 40)
    (3.3.3.11 40) (3.3.3.12 40) (3.3.3.13 40) (3.3.3.14 40) (3.3.3.15 40)
    (3.3.3.16 40) (3.3.3.17 40) (3.3.3.18 40) (3.3.3.19 40)
Interface v110 : 1 groups
  group          phy-port static querier life mode  #_src
2  239.0.0.1     e5    no    yes           include 10
  group: 239.0.0.1, include, permit 10 (source, life):
    (2.2.3.0 80) (2.2.3.1 80) (2.2.3.2 80) (2.2.3.3 80) (2.2.3.4 80)
    (2.2.3.5 80) (2.2.3.6 80) (2.2.3.7 80) (2.2.3.8 80) (2.2.3.9 80)
```

If the tracking and fast leave feature is enabled, you can display the list of clients that belong to a particular group by entering commands such as the following.

```
PowerConnect#show ip igmp group 224.1.10.1 tracking
Display group 224.1.10.1 in all interfaces with tracking enabled.
Interface v13 : 1 groups, tracking_enabled
  group          phy-port static querier life mode  #_src
1  224.1.10.1    e15   no    yes           include 3
  receive reports from 3 clients:
    110.110.110.7 110.110.110.8 110.110.110.9
```

Syntax: show ip igmp group [<group-address>] [detail | tracking]

If you want a report for a specific multicast group, enter that group address for <group-address>. Omit the <group-address> if you want a report for all multicast groups.

Enter **detail** if you want to display the source list of the multicast group.

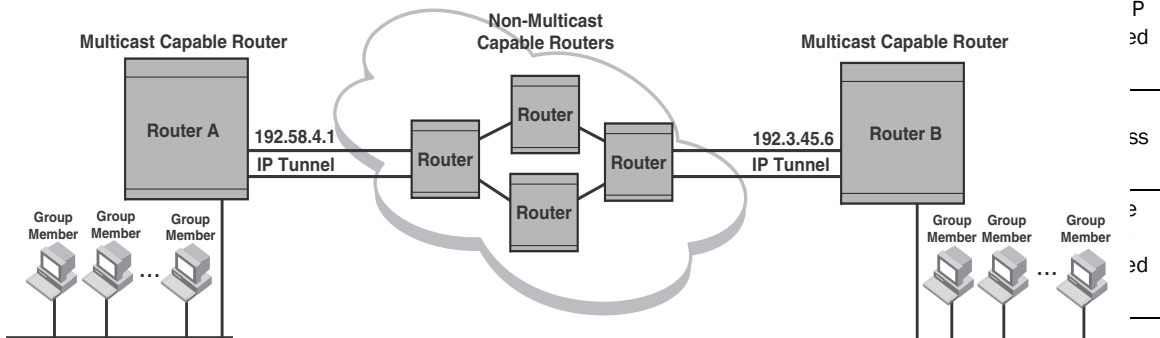
Enter **tracking** if you want information on interfaces that have tracking enabled.

The following table defines the statistics for the **show ip igmp group** command output.

TABLE 90 Output of show ip igmp group

This field	Displays
Group	The address of the multicast group
Phy-port	The physical port on which the multicast group was received.

TABLE 90 Output of show ip igmp group (Continued)

This field	Displays
Static	A “yes” entry in this column indicates that the multicast group was configured as a
	
#_src	<p>Identifies the source list that will be included or excluded on the interface. If IGMP V2 group is in Exclude mode with a #_src of 0, the group excludes traffic from 0 (zero) source list, which means that all traffic sources are included.</p>
Group:	<p>If you requested a <i>detailed</i> report, the following information is displayed:</p> <ul style="list-style-type: none"> • The multicast group address • The mode of the group • A list of sources from which traffic will be admitted (include) or denied (exclude) on the interface is listed. • The life of each source list. <p>If you requested a <i>tracking</i> report, the clients from which reports were received are identified.</p>

Displaying the IGMP status of an interface

You can display the status of a multicast enabled port by entering a command such as the following.

NOTE

This report is available on Layer 3 Switches.

```
PowerConnect#show ip igmp interface
query interval = 60, max response time= 3, group membership time=140
v5: default V2, PIM dense, addr=1.1.1.2
  e412 has 0 groups, non-Querier (age=40), default V2
v20: configured V3, PIM dense (port down), addr=1.1.20.1
v110: configured V3, PIM dense, addr=110.110.110.1
  e6 has 2 groups, Querier, default V3
    group: 239.0.0.1, exclude, life=100, deny 13
    group: 224.1.10.1, include, permit 2
  e5 has 3 groups, Querier, default V3
    group: 224.2.2.2, include, permit 100
    group: 239.0.0.1, include, permit 10
    group: 224.1.10.1, include, permit 1
```

Syntax: show ip igmp interface [ve | ethernet <number> <group-address>]

Enter **ve** and its <number> or **ethernet** and its <number> to display information for a specific virtual routing interface or ethernet interface.

Entering an address for <group-address> displays information for a specified group on the specified interface.

The report shows the following information.

TABLE 91 Output of show ip igmp interface

This field	Displays
Query interval	Displays how often a querier sends a general query on the interface.
Max response	The maximum number of seconds a client can wait before it replies to the query.
Group membership time	The number of seconds multicast groups can be members of this group before aging out.
(details)	<p>The following is displayed for each interface:</p> <ul style="list-style-type: none"> • The ID of the interface • The IGMP version that it is running (default IGMP V2 or configured IGMP V3) • The multicast protocol it is running: PIM-DM, PIM-SM • Address of the multicast group on the interface • If the interface is a virtual routing interface, the physical port to which that interface belongs, the number of groups on that physical port, whether or not the port is a querier or a non-querier port, the age of the port, and other multicast information for the port are displayed.

Displaying IGMP traffic status

To display the traffic status on each virtual routing interface, enter the following command.

NOTE

This report is available on Layer 3 Switches.

```
PowerConnect#show ip igmp traffic
Recv  QryV2  QryV3  G-Qry  GSQry  MbrV2  MbrV3  Leave  IsIN  IsEX  ToIN  ToEX  ALLOW  BLK
v5      29      0      0      0      0      0      0      0      0      0      0      0      0
v18     15      0      0      0      0      30     0      60     0      0      0      0      0
v110    0      0      0      0      0      97     0     142    37     2      2      3      2
Send  QryV1  QryV2  QryV3  G-Qry  GSQry
v5      0      2      0      0      0
v18     0      0      30     30     0
v110    0      0      30     44     11
```

Syntax: show ip igmp traffic

The report shows the following information.

TABLE 92 Output of show ip igmp traffic

This field	Displays
QryV2	Number of general IGMP V2 query received or sent by the virtual routing interface.
QryV3	Number of general IGMP V3 query received or sent by the virtual routing interface.
G-Qry	Number of group specific query received or sent by the virtual routing interface.
GSQry	Number of source specific query received or sent by the virtual routing interface.

TABLE 92 Output of show ip igmp traffic (Continued)

This field	Displays
MbrV2	The IGMP V2 membership report.
MbrV3	The IGMP V3 membership report.
Leave	Number of IGMP V2 "leave" messages on the interface. (See ToEx for IGMP V3.)
IsIN	Number of source addresses that were included in the traffic.
IsEX	Number of source addresses that were excluded in the traffic.
ToIN	Number of times the interface mode changed from exclude to include.
ToEX	Number of times the interface mode changed from include to exclude.
ALLOW	Number of times that additional source addresses were allowed or denied on the interface.
BLK	Number of times that sources were removed from an interface.

Clearing IGMP statistics

To clear statistics for IGMP traffic, enter the following command.

```
PowerConnect# clear igmp traffic
```

Syntax: clear igmp traffic

This command clears all the multicast traffic information on all interfaces on the device.

IGMP Proxy

IGMP Proxy provides a means for the PowerConnect B-Series TI24X routers to receive any or all multicast traffic from an upstream device if the router is not able to run PIM.

IGMP Proxy enables the router to issue IGMP host messages on behalf of hosts that the router discovered through standard PIM interfaces. The router acts as a proxy for its hosts and performs the host portion of the IGMP task on the upstream interface as follows:

- When queried, the router sends group membership reports for the groups learned
- When one of its hosts joins a multicast address group to which none of its other hosts belong, the router sends unsolicited membership reports to that group.
- When the last of its hosts in a particular multicast group leaves the group, the PowerConnect B-Series TI24X router sends an unsolicited leave group membership report to group for all routers (multicast IP address 244.0.0.2)

Configuration notes

When using IGMP Proxy, you must do the following.

1. Configure PIM on all multicast client ports to build the group membership table. The group membership table will be reported by the proxy interface. Refer to [“Globally enabling and disabling PIM”](#) on page 473.
2. Enable IP multicast on an interface to an upstream PowerConnect B-Series TI24X router that will be the IGMP proxy interface and configure IGMP Proxy on that interface

Also note the following limitations:

- IGMP Proxy cannot be enabled on the same interface on which PIM SM, PIM DM, is enabled.
- IGMP Proxy is only supported in a PIM Dense environment where there are IGMP clients connected to the Dell device. The Dell device will not send IGMP reports on an IGMP proxy interface for remote clients connected to a PIM neighbor, as it will not be aware of groups that the remote clients are interested in.

Configuring IGMP Proxy

Follow the steps given below to configure IGMP Proxy.

1. Configure router PIM globally.

```
PowerConnect(config)#router pim
```

2. Configure an IP address on the interface (physical or virtual routing interface) that will serve as the IGMP proxy for an upstream device by entering commands such as the following.

```
PowerConnect(config)#int e 1/3
PowerConnect(config-if-e1000-1/3)#ip address 207.95.5.1/24
```

3. Enable IGMP Proxy on the interface.

```
PowerConnect(config-if-e1000-1/3)#ip igmp proxy
```

Syntax: [no] ip igmp proxy

Once IGMP Proxy is configured and the PowerConnect B-Series T124X router receives a query on an IGMP Proxy interface, the router sends a report in response to the query before the IGMP maximum response time expires.

Displaying IGMP Proxy traffic

Use the **show ip igmp traffic** command to see traffic for IGMP Proxy.

```
PowerConnect#show ip igmp traffic
Recv  QryV2  QryV3  G-Qry  GSQry  MbrV2  MbrV3  Leave  IsIN  IsEX  ToIN  ToEX  ALLO  BLK
e1/14    0      0      0      0  27251    0      12     0  27251  12    0     0     0
v10     250    0      0      0   244     0      0     0   244    0     0     0     0
Send  QryV1  QryV2  QryV3  G-Qry  GSQry  MbrV1 Mbrv2 Leave
e1/14    0  1365    0      48     0      0      0     0
v10      0     1     0      0     0      0  25602    1
```

Syntax: show ip igmp traffic

Refer to “[Displaying IGMP traffic status](#)” on page 534 to interpret the information in the output. The fields in bold show information for IGMP Proxy.

Configuring LLDP

This chapter describes how to configure the LLDP protocol:

Link layer discovery protocol (LLDP) – The Layer 2 network discovery protocol described in the IEEE 802.1AB standard, *Station and Media Access Control Connectivity Discovery*. This protocol enables a station to advertise its capabilities to, and to discover, other LLDP-enabled stations in the same 802 LAN segments.

LLDP enables network discovery between Network Connectivity devices (such as switches).

The information generated by LLDP can be used to diagnose and troubleshoot misconfigurations on both sides of a link. For example, the information generated can be used to discover devices with misconfigured or unreachable IP addresses, and to detect port speed and duplex mismatches.

LLDP facilitate interoperability across multiple vendor devices. devices running LLDP can interoperate with third-party devices running LLDP.

The LLDP implementation adheres to the IEEE 802.1AB and TIA-1057 standards.

Terms used in this chapter

LLDP agent – The protocol entity that implements LLDP for a particular IEEE 802 device. Depending on the configured LLDP operating mode, an LLDP agent can send and receive LLDP advertisements (frames), or send LLDP advertisements only, or receive LLDP advertisements only.

LLDPDU (LLDP Data Unit) – A unit of information in an LLDP packet that consists of a sequence of short variable length information elements, known as **TLVs**. LLDP pass-through is not supported in conformance to IEEE standard.

MIB (Management Information Base) – A virtual database that identifies each manageable object by its name, syntax, accessibility, and status, along with a text description and unique object identifier (OID). The database is accessible by a Network Management Station (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP).

Network connectivity device – A forwarding 802 LAN device, such as a router, switch, or wireless access point.

Station – A node in a network.

TLV (Type-Length-Value) – An information element in an LLDPDU that describes the type of information being sent, the length of the information string, and the value (actual information) that will be transmitted.

TTL (Time-to-Live) – Specifies the length of time that the receiving device should maintain the information acquired through LLDP in its MIB.


LLDP overview

LLDP enables a station attached to an IEEE 802 LAN/MAN to advertise its capabilities to, and to discover, other stations in the same 802 LAN segments.

The information distributed by LLDP (the advertisement) is stored by the receiving device in a standard Management Information Base (MIB), accessible by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP). The information also can be viewed from the CLI, using **show LLDP** commands.

Figure 93 illustrates LLDP connectivity

FIGURE 93 LLDP connectivity



Benefits of LLDP

LLDP provides the following benefits:

- Network Management:
 - Simplifies the use of and enhances the ability of network management tools in multi-vendor environments
 - Enables discovery of accurate physical network topologies such as which devices are neighbors and through which ports they connect

- Enables discovery of stations in multi-vendor environments
- Network Inventory Data:
 - Supports optional system name, system description, system capabilities and management address
 - System description can contain the device product name or model number, version of hardware type, and operating system
 - Provides device capability, such as switch, router, or WLAN access port
- Network troubleshooting:
 - Information generated by LLDP can be used to detect speed and duplex mismatches
 - Accurate topologies simplify troubleshooting within enterprise networks
 - Can discover devices with misconfigured or unreachable IP addresses

General operating principles

LLDP use the services of the Data Link sublayers, Logical Link Control and Media Access Control, to transmit and receive information to and from other **LLDP Agents** (protocol entities that implement LLDP).

LLDP is a one-way protocol. An LLDP agent can transmit and receive information to and from another LLDP agent located on an adjacent device, but it cannot solicit information from another LLDP agent, nor can it acknowledge information received from another LLDP agent.

Operating modes

When LLDP is enabled on a global basis, by default, each port on the Dell device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

Transmit mode

An LLDP agent sends LLDP packets to adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

An LLDP agent initiates the transmission of LLDP packets whenever the transmit countdown timing counter expires, or whenever LLDP information has changed. When a transmit cycle is initiated, the LLDP manager extracts the MIB objects and formats this information into TLVs. The TLVs are inserted into an LLDPDU, addressing parameters are prepended to the LLDPDU, and the information is sent out LLDP-enabled ports to adjacent LLDP-enabled devices.

Receive mode

An LLDP agent receives LLDP packets from adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.


When an LLDP agent receives LLDP packets, it checks to ensure that the LLDPDUs contain the correct sequence of mandatory TLVs, then validates optional TLVs. If the LLDP agent detects any errors in the LLDPDUs and TLVs, it drops them in software. TLVs that are not recognized but do not contain basic formatting errors, are assumed to be valid and are assigned a temporary identification index and stored for future possible alter retrieval by network management. All validated TLVs are stored in the neighbor database.

LLDP packets

LLDP agents transmit information about a sending device/port in packets called LLDP Data Units (LLDPDUs). All the LLDP information to be communicated by a device is contained within a single 1500 byte packet. A device receiving LLDP packets is not permitted to combine information from multiple packets.

As shown in [Figure 94](#), each LLDPDU has three mandatory TLVs, an End of LLDPDU TLV, plus optional TLVs as selected by network management.

FIGURE 94 LLDPDU packet format



M = mandatory TLV (required for all LLDPDUs)

Each LLDPDU consists of an untagged Ethernet header and a sequence of short, variable length information elements known as TLVs.

TLVs have Type, Length, and Value fields, where:

- **Type** identifies the kind of information being sent
- **Length** indicates the length (in octets) of the information string
- **Value** is the actual information being sent (for example, a binary bit map or an alpha-numeric string containing one or more fields).

TLV support

This section lists the LLDP TLV support.

LLDP TLVs

There are two types of LLDP TLVs, as specified in the IEEE 802.3AB standard:

- **Basic management TLVs** consist of both optional general system information TLVs as well as mandatory TLVs.

Mandatory TLVs cannot be manually configured. They are always the first three TLVs in the LLDPDU, and are part of the packet header.

General system information TLVs are optional in LLDP implementations and are defined by the Network Administrator.

Devices support the following Basic Management TLVs:

- Chassis ID (mandatory)
- Port ID (mandatory)
- Time to Live (mandatory)
- Port description
- System name
- System description
- System capabilities
- Management address
- End of LLDPDU
- **Organizationally-specific TLVs** are optional in LLDP implementations and are defined and encoded by individual organizations or vendors. These TLVs include support for, but are not limited to, the IEEE 802.1 and 802.3 standards and the TIA-1057 standard.

Devices support the following Organizationally-specific TLVs:

- **802.1 organizationally-specific TLVs**
 - Port VLAN ID
 - VLAN name TLV
- **802.3 organizationally-specific TLVs**
 - MAC/PHY configuration/status
 - Power through MDI
 - Link aggregation
 - Maximum frame size

Mandatory TLVs

When an LLDP agent transmits LLDP packets to other agents in the same 802 LAN segments, the following mandatory TLVs are always included:

- Chassis ID
- Port ID
- Time to Live (TTL)

This section describes the above TLVs in detail.

Chassis ID

The Chassis ID identifies the device that sent the LLDP packets.

There are several ways in which a device may be identified. A chassis ID subtype, included in the TLV and shown in [Table 93](#), indicates how the device is being referenced in the Chassis ID field.

TABLE 93 Chassis ID subtypes

ID subtype	Description
0	Reserved
1	Chassis component
2	Interface alias

TABLE 93 Chassis ID subtypes

ID subtype	Description
3	Port component
4	MAC address
5	Network address
6	Interface name
7	Locally assigned
8 - 255	Reserved

Devices use chassis ID subtype 4, the base MAC address of the device. Other third party devices may use a chassis ID subtype other than 4. The chassis ID will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
Chassis ID (MAC address): 0012.f233.e2c0
```

The chassis ID TLV is always the first TLV in the LLDPDU.

Port ID

The Port ID identifies the port from which LLDP packets were sent.

There are several ways in which a port may be identified, as shown in [Figure 94](#). A port ID subtype, included in the TLV, indicates how the port is being referenced in the Port ID field.

TABLE 94 Port ID subtypes

ID subtype	Description
0	Reserved
1	Interface alias
2	Port component
3	MAC address
4	Network address
5	Interface name
6	Agent circuit ID
7	Locally assigned
8 - 255	Reserved


Devices use port ID subtype 3, the permanent MAC address associated with the port. Other third party devices may use a port ID subtype other than 3. The port ID appears similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
Port ID (MAC address): 0012.f233.e2d3
```

The LLDPDU format is shown in [“LLDPDU packet format”](#) on page 540.

The Port ID TLV format is shown below.

FIGURE 95 Port ID TLV packet format



TTL value

The Time to Live (TTL) Value is the length of time the receiving device should maintain the information acquired by LLDP in its MIB.

The TTL value is automatically computed based on the LLDP configuration settings. The TTL value will appear similar to the following on the remote device, and in the CLI display output on the device (show lldp local-info).

```
Time to live: 40 seconds
```


If the TTL field has a value other than zero, the receiving LLDP agent is notified to completely replace all information associated with the LLDP agent/port with the information in the received LLDPDU.

If the TTL field value is zero, the receiving LLDP agent is notified that all system information associated with the LLDP agent/port is to be deleted. This TLV may be used, for example, to signal that the sending port has initiated a port shutdown procedure.

The LLDPDU format is shown in [“LLDPDU packet format”](#) on page 540.

The TTL TLV format is shown below.

FIGURE 96 TTL TLV packet format



MIB support

Devices support the following standard MIB modules:

- LLDP-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB

Syslog messages

Syslog messages for LLDP provide management applications with information related to MIB data consistency and general status. These Syslog messages correspond to the lldpRemTablesChange SNMP notifications. Refer to [“Enabling LLDP SNMP notifications and syslog messages”](#) on page 547.

Configuring LLDP

This section describes how to enable and configure LLDP.

[Table 95](#) lists the LLDP global-level tasks and the default behavior/value for each task.

TABLE 95 LLDP global configuration tasks and default behavior /value

Global task	Default behavior / value when LLDP is enabled
Enabling LLDP on a global basis	Disabled
Specifying the maximum number of LLDP neighbors per device	Automatically set to 392 neighbors per device
Specifying the maximum number of LLDP neighbors per port	Automatically set to 4 neighbors per port
Enabling SNMP notifications and Syslog messages	Disabled
Changing the minimum time between SNMP traps and Syslog messages	Automatically set to 2 seconds when SNMP notifications and Syslog messages for LLDP are enabled
Enabling and disabling TLV advertisements	When LLDP transmit is enabled, by default, the device will automatically advertise LLDP capabilities, except for the system description, VLAN name, and power-via-MDI information, which may be configured by the system administrator. Also, if desired, you can disable the advertisement of individual TLVs.
Changing the minimum time between LLDP transmissions	Automatically set to 2 seconds
Changing the interval between regular LLDP transmissions	Automatically set to 30 seconds
Changing the holdtime multiplier for transmit TTL	Automatically set to 4
Changing the minimum time between port reinitializations	Automatically set to 2 seconds

Configuration notes and considerations

- LLDP is supported on Ethernet interfaces only.
- If a port is 802.1X-enabled, the transmission and reception of LLDP packets will only take place while the port is authorized.
- Cisco Discovery Protocol (CDP) and Brocade Discovery Protocol (FDP) run independently of LLDP. Therefore, these discovery protocols can run simultaneously on the same device.
- By default, the device limits the number of neighbors per port to four, and staggers the transmission of LLDP packets on different ports, in order to minimize any high-usage spikes to the CPU.
- Ports that are in blocking mode (spanning tree) can still receive LLDP packets from a forwarding port.
- Auto-negotiation status indicates what is being advertised by the port for 802.3 auto-negotiation.

Enabling and disabling LLDP

LLDP is enabled by default on individual ports. However, to run LLDP, you must first enable it on a global basis (on the entire device).

To enable LLDP globally, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)#lldp run
```

Syntax: [no] lldp run

Changing a port LLDP operating mode

LLDP packets are not exchanged until LLDP is enabled on a global basis. When LLDP is enabled on a global basis, by default, each port on the Dell device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

You can configure a different operating mode for each port on the device. For example, you could disable the receipt and transmission of LLDP packets on port e 1, configure port e 3 to only receive LLDP packets, and configure port e 5 to only transmit LLDP packets.

The following sections show how to change the operating mode.

Enabling and disabling receive and transmit mode

To disable the receipt and transmission of LLDP packets on individual ports, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#no lldp enable ports e 4 e 5
```

The above command disables LLDP on ports 4 and 5. These ports will not transmit nor receive LLDP packets.

To enable LLDP on a port after it has been disabled, enter the following command.

```
PowerConnect(config)#lldp enable ports e 4
```

Syntax: [no] lldp enable ports ethernet <port-list> | all

Use the [no] form of the command to disable the receipt and transmission of LLDP packets on a port.

Enabling and disabling receive only mode

When LLDP is enabled on a global basis, by default, each port on the device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode from receive and transmit mode to receive only mode, simply disable the transmit mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#no lldp enable transmit ports e 4 e 5 e 6
```

The above command changes the LLDP operating mode on ports 4, 5, and 6 from transmit and receive mode to receive only mode.

To change a port LLDP operating mode from transmit only to receive only, first disable the transmit only mode, then enable the receive only mode. Enter commands such as the following.

```
PowerConnect(config)#no lldp enable transmit ports e 7 e 8 e 9
PowerConnect(config)#lldp enable receive ports e 7 e 8 e 9
```

The above commands change the LLDP operating mode on ports 7, 8, and 9, from transmit only to receive only. Note that if you do not disable the transmit only mode, you will configure the port to both transmit and receive LLDP packets.

Syntax: `[no] lldp enable receive ports ethernet <port-list> | all`

Use the [no] form of the command to disable the receive only mode.

Enabling and Disabling Transmit Only Mode

When LLDP is enabled on a global basis, by default, each port on the device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode to transmit only mode, simply disable the receive mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#no lldp enable receive ports e 4 e 5 e 6
```

The above command changes the LLDP operating mode on ports 4, 5, and 6 from transmit and receive mode to transmit only mode. Any incoming LLDP packets will be dropped in software.

To change a port LLDP operating mode from receive only to transmit only, first disable the receive only mode, then enable the transmit only mode. For example, enter commands such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#no lldp enable receive ports e 7 e 8
PowerConnect(config)#lldp enable transmit ports e 7 e 8
```

The above commands change the LLDP operating mode on ports 7 and 8 from receive only mode to transmit only mode. Any incoming LLDP packets will be dropped in software. Note that if you do not disable receive only mode, you will configure the port to both receive and transmit LLDP packets.

Syntax: `[no] lldp enable transmit ports ethernet <port-list> | all`

Use the [no] form of the command to disable the *transmit only* mode.

Specifying the maximum number of LLDP neighbors

You can change the limit of the number of LLDP neighbors for which LLDP data will be retained, per device as well as per port.

Per device

You can change the maximum number of neighbors for which LLDP data will be retained for the entire system.

For example, to change the maximum number of LLDP neighbors for the entire device to 26, enter the following command.

```
PowerConnect(config)#lldp max-total-neighbors 26
```

Syntax: `[no] lldp max-total-neighbors <value>`

Use the [no] form of the command to remove the static configuration and revert to the default value of 392.

where <value> is a number between 16 and 65536. The default number of LLDP neighbors per device is 392.

Use the **show lldp** command to view the configuration.

Per port

You can change the maximum number of LLDP neighbors for which LLDP data will be retained for each port. By default, the maximum number is four and you can change this to a value between one and 64.

For example, to change the maximum number of LLDP neighbors to six, enter the following command.

```
PowerConnect(config)#lldp max-neighbors-per-port 6
```

Syntax: [no] lldp max-neighbors-per-port <value>

Use the [no] form of the command to remove the static configuration and revert to the default value of four.

where <value> is a number from 1 to 64. The default is number of LLDP neighbors per port is four.

Use the **show lldp** command to view the configuration.

Enabling LLDP SNMP notifications and syslog messages

SNMP notifications and Syslog messages for LLDP provide management applications with information related to MIB data updates and general status.

When you enable LLDP SNMP notifications, corresponding Syslog messages are enabled as well. When you enable LLDP SNMP notifications, the device will send traps and corresponding Syslog messages whenever there are changes to the LLDP data received from neighboring devices.

LLDP SNMP notifications and corresponding Syslog messages are disabled by default. To enable them, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#lldp enable snmp notifications ports e 2 to 6
```

The above command enables SNMP notifications and corresponding Syslog messages on ports 2 and 6. By default, the device will send no more than one SNMP notification and Syslog message within a five second period. If desired, you can change this interval. Refer to [“Specifying the minimum time between SNMP traps and syslog messages”](#) on page 547.

Syntax: [no] lldp enable snmp notifications ports ethernet <port-list> | all

Specifying the minimum time between SNMP traps and syslog messages

When SNMP notifications and Syslog messages for LLDP are enabled, the device will send no more than one SNMP notification and corresponding Syslog message within a five second period. If desired, you can throttle the amount of time between transmission of SNMP traps (lldpRemTablesChange) and Syslog messages from five seconds up to a value equal to one hour (3600 seconds).

NOTE

Because LLDP Syslog messages are rate limited, some LLDP information given by the system will not match the current LLDP statistics (as shown in the **show lldp statistics** command output).

To change the minimum time interval between traps and Syslog messages, enter a command such as the following.

```
PowerConnect(config)#lldp snmp-notification-interval 60
```

When the above command is applied, the LLDP agent will send no more than one SNMP notification and Syslog message every 60 seconds.

Syntax: `[no] lldp snmp-notification-interval <seconds>`

where <seconds> is a value between 5 and 3600. The default is 5 seconds.

Changing the minimum time between LLDP transmissions

The LLDP transmit delay timer limits the number of LLDP frames an LLDP agent can send within a specified time frame. When you enable LLDP, the system automatically sets the LLDP transmit delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between 1 and 8192 seconds.

NOTE

The LLDP transmit delay timer must not be greater than one quarter of the LLDP transmission interval (CLI command `lldp transmit-interval`).

The LLDP transmit delay timer prevents an LLDP agent from transmitting a series of successive LLDP frames during a short time period, when rapid changes occur in LLDP. It also increases the probability that multiple changes, rather than single changes, will be reported in each LLDP frame.

To change the LLDP transmit delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#lldp transmit-delay 7
```

The above command causes the LLDP agent to wait a minimum of seven seconds after transmitting an LLDP frame and before sending another LLDP frame.

Syntax: `[no] lldp transmit-delay <seconds>`

where <seconds> is a value between 1 and 8192. The default is two seconds. Note that this value must not be greater than one quarter of the LLDP transmission interval (CLI command `lldp transmit-interval`).

Changing the interval between regular LLDP transmissions

The LLDP transmit interval specifies the number of seconds between regular LLDP packet transmissions. When you enable LLDP, by default, the device will wait 30 seconds between regular LLDP packet transmissions. If desired, you can change the default behavior from 30 seconds to a value between 5 and 32768 seconds.

To change the LLDP transmission interval, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#lldp transmit-interval 40
```

The above command causes the LLDP agent to transmit LLDP frames every 40 seconds.

Syntax: `[no] lldp transmit-interval <seconds>`

where <seconds> is a value from 5 to 32768. The default is 30 seconds.

NOTE

Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the holdtime multiplier for transmit TTL

The holdtime multiplier for transmit TTL is used to compute the actual time-to-live (TTL) value used in an LLDP frame. The TTL value is the length of time the receiving device should maintain the information in its MIB. When you enable LLDP, the device automatically sets the holdtime multiplier for TTL to four. If desired, you can change the default behavior from four to a value between two and ten.

To compute the TTL value, the system multiplies the LLDP transmit interval by the holdtime multiplier. For example, if the LLDP transmit interval is 30 and the holdtime multiplier for TTL is 4, then the value 120 is encoded in the TTL field in the LLDP header.

To change the holdtime multiplier, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#lldp transmit-hold 6
```

Syntax: [no] lldp transmit-hold <value>

where <value> is a number from 2 to 10. The default value is 4.

NOTE

Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the minimum time between port reinitializations

The LLDP re-initialization delay timer specifies the minimum number of seconds the device will wait from when LLDP is disabled on a port, until it will honor a request to re-enable LLDP on that port. When you enable LLDP, the system sets the re-initialization delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between one and ten seconds.

To set the re-initialization delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
PowerConnect(config)#lldp reinit-delay 5
```

The above command causes the device to wait five seconds after LLDP is disabled, before attempting to honor a request to re-enable it.

Syntax: [no] lldp reinit-delay <seconds>

where <seconds> is a value from 1 - 10. The default is two seconds.

LLDP TLVs advertised by the device

When LLDP is enabled on a global basis, the device will automatically advertise the following information, except for the features noted:

General system information:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

802.1 capabilities:

- VLAN name (not automatically advertised)
- Untagged VLAN ID

802.3 capabilities:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

The above TLVs are described in detail in the following sections.

NOTE

The system description, VLAN name, and power-via-MDI information TLVs are not automatically enabled. The following sections show how to enable these advertisements.

General system information

Except for the system description, the device will advertise the following system information when LLDP is enabled on a global basis:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

Management Address

A management address is an IPv4 address that can be used to manage the device. If no management address is explicitly configured to be advertised, the device will use the first available IPv4 address. A Layer 3 switch will select the first available address of each type from those configured on the following types of interfaces, in the following order of preference:

- Physical port on which LLDP will be transmitting the packet
- Virtual router interface (VE) on a VLAN that the port is a member of
- Loopback interface
- Virtual router interface (VE) on any other VLAN
- Other physical port

If no IP address is configured on any of the above, the port's current MAC address will be advertised.

To advertise a IPv4 management address, enter a command such as the following:

```
PowerConnect(config)#lldp advertise management-address ipv4 209.157.2.1 ports e 4
```

The management address will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**):

```
Management address (IPv4): 209.157.2.1
```

Syntax: [no] lldp advertise management-address ipv4 <ipv4 address> ports ethernet <port list> | all

<ipv4 address> or <ipv6 address> or both are the addresses that may be used to reach higher layer entities to assist discovery by network management. In addition to management addresses, the advertisement will include the system interface number associated with the management address.

For <port list>, specify the port(s) in the format <portnum>. You can list all of the ports individually; use the keyword to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Port description

The port description TLV identifies the port from which the LLDP agent transmitted the advertisement. The port description is taken from the ifDescr MIB object from MIB-II.

By default, the port description is automatically advertised when LLDP is enabled on a global basis. To disable advertisement of the port description, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise port-description ports e 4 to 12
```

The port description will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
Port description: "GigabitEthernet20"
```

Syntax: [no] lldp advertise port-description ports ethernet <port-list> | all

System capabilities

The system capabilities TLV identifies the primary functions of the device and indicates whether these primary functions are enabled. The primary functions can be one or more of the following (more than one for example, if the device is both a bridge and a router):

- Repeater
- Bridge
- WLAN access point
- Router
- Telephone
- DOCSIS cable device
- Station only (devices that implement end station capability)
- Other

System capabilities for devices are based on the type of software image in use (e.g., Layer 2 switch or Layer 3 router). The enabled capabilities will be the same as the available capabilities, except that when using a router image (base or full Layer 3), if the global route-only feature is turned on, the bridge capability will not be included, since no bridging takes place.

By default, the system capabilities are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise system-capabilities ports e 4 to 12
```

The system capabilities will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
System capabilities : bridge
Enabled capabilities: bridge
```

Syntax: [no] lldp advertise system-capabilities ports ethernet <port-list> | all

System description

The system description is the network entity, which can include information such as the product name or model number, the version of the system hardware type, the software operating system level, and the networking software version. The information corresponds to the sysDescr MIB object in MIB-II.

To advertise the system description, enter a command such as the following.

```
PowerConnect(config)#lldp advertise system-description ports e 4 to 12
```

The system description will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
+ System description : "Brocade Communications, Inc.,TI24X, IronWare Version
04.0.00b256T3e1 Compiled on Sep 04 2007 at 0\
3:54:29 labeled as SXS04000b256"
```

NOTE

The contents of the show command output will vary depending on which TLVs are configured to be advertised.

Syntax: [no] lldp advertise system-description ports ethernet <port-list> | all

System name

The system name is the system administratively assigned name, taken from the sysName MIB object in MIB-II. The sysName MIB object corresponds to the name defined with the CLI command **hostname**.

By default, the system name is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise system-name ports e 4 to 12
```

The system name will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
System name: "PowerConnect"
```

Syntax: [no] lldp advertise system-name ports ethernet <port-list> | all

802.1 capabilities

Except for the VLAN name, the device will advertise the following 802.1 attributes when LLDP is enabled on a global basis:

- VLAN name (not automatically advertised)

- Untagged VLAN ID

VLAN name

The VLAN name TLV contains the name and VLAN ID of a VLAN configured on a port. An LLDPDU may include multiple instances of this TLV, each for a different VLAN.

To advertise the VLAN name, enter a command such as the following.

```
PowerConnect(config)#lldp advertise vlan-name vlan 99 ports e 4 to 12
```

The VLAN name will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
VLAN name (VLAN 99): "Voice-VLAN-99"
```

Syntax: [no] lldp advertise vlan-name vlan <vlan ID> ports ethernet <port-list> | all

For <vlan ID>, enter the VLAN ID to advertise.

Untagged VLAN id

The port VLAN ID TLV advertises the Port VLAN Identifier (PVID) that will be associated with untagged or priority-tagged frames. If the port is not an untagged member of any VLAN (i.e., the port is strictly a tagged port), the value zero will indicate that.

By default, the port VLAN ID is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise port-vlan-id ports e 4 to 12
```

The untagged VLAN ID will appear similar to the following on the remote device, and in the CLI display output on the device (Refer to **show lldp local-info command** on page 625).

```
Port VLAN ID: 99
```

Syntax: [no] lldp advertise port-vlan-id ports ethernet <port-list> | all

802.3 capabilities

Except for Power-via-MDI information, the device will advertise the following 802.3 attributes when LLDP is enabled on a global basis:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

Link aggregation

The **link-aggregation** TLV indicates the following:

- Whether the link is capable of being aggregated
- Whether the link is currently aggregated
- The primary trunk port

Devices advertise link aggregation information about standard link aggregation (LACP) as well as static trunk configuration.

By default, link-aggregation information is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise link-aggregation ports e 12
```

Syntax: [no] lldp advertise link-aggregation ports ethernet <port-list> | all

The link aggregation advertisement will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
Link aggregation: not capable
```

MAC/PHY configuration status

The MAC/PHY configuration and status TLV includes the following information:

- Auto-negotiation capability and status
- Speed and duplex mode
- Flow control capabilities for auto-negotiation
- Port speed down-shift and maximum port speed advertisement
- If applicable, indicates if the above settings are the result of auto-negotiation during link initiation or of a manual set override action

The advertisement reflects the effects of the following CLI commands:

- speed-duplex
- flow-control
- gig-default
- link-config

By default, the MAC/PHY configuration and status information are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise mac-phy-config-status ports e 4 to 12
```

The MAC/PHY configuration advertisement will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

```
+ 802.3 MAC/PHY          : auto-negotiation enabled
  Advertised capabilities: 10baseT-HD, 10baseT-FD, 100baseTX-HD,
  100baseTX-FD,
  fdxSPause, fdxBPause, 1000baseT-HD, 1000baseT-FD
  Operational MAU type: 100BaseTX-FD
```

Syntax: [no] lldp advertise mac-phy-config-status ports ethernet <port-list> | all

Maximum frame size

The maximum frame size TLV provides the maximum 802.3 frame size capability of the port. This value is expressed in octets and includes the four-octet Frame Check Sequence (FCS). The default maximum frame size is 1522. The advertised value may change depending on whether the **aggregated-vlan** or **jumbo** CLI commands are in effect.

By default, the maximum frame size is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
PowerConnect(config)#no lldp advertise max-frame-size ports e 4 to 12
```

The maximum frame size advertisement will appear similar to the following on the remote device, and in the CLI display output on the device (**show lldp local-info**).

Maximum frame size: 1522 octets

Syntax: [no] lldp advertise max-frame-size ports ethernet <port-list> | all

Displaying LLDP statistics and configuration settings

You can use the following CLI **show** commands to display information about LLDP settings and statistics:

- **show lldp** – Displays a summary of the LLDP configuration settings.
- **show lldp statistics** – Displays LLDP global and per-port statistics.
- **show lldp neighbors** – Displays a list of the current LLDP neighbors.
- **show lldp neighbors detail** – Displays the details of the latest advertisements received from LLDP neighbors.
- **show lldp local-info** – Displays the details of the LLDP advertisements that will be transmitted on each port.

This above **show** commands are described in this section.

LLDP configuration summary

To display a summary of the LLDP configuration settings on the device, enter the **show lldp** command at any level of the CLI.

The following shows an example report.

```
PowerConnect#show lldp
LLDP transmit interval      : 10 seconds
LLDP transmit hold multiplier : 4 (transmit TTL: 40 seconds)
LLDP transmit delay        : 1 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay    : 1 seconds
LLDP-MED fast start repeat count : 3

LLDP maximum neighbors      : 392
LLDP maximum neighbors per port : 4
```

Syntax: **show lldp**

The following table describes the information displayed by the **show lldp statistics** command.

Table 2:

This field...	Displays...
LLDP transmit interval	The number of seconds between regular LLDP packet transmissions.
LLDP transmit hold multiplier	The multiplier used to compute the actual time-to-live (TTL) value of an LLDP advertisement. The TTL value is the transmit interval multiplied by the transmit hold multiplier.
LLDP transmit delay	The number of seconds the LLDP agent will wait after transmitting an LLDP frame and before transmitting another LLDP frame.
LLDP SNMP notification interval	The number of seconds between transmission of SNMP LLDP traps (lldpRemTablesChange) and SNMP LLDP-MED traps (lldpXMedTopologyChangeDetected).

Table 2:

This field...	Displays...
LLDP reinitialize delay	The minimum number of seconds the device will wait from when LLDP is disabled on a port, until a request to re-enable LLDP on that port will be honored.
LLDP maximum neighbors	The maximum number of LLDP neighbors for which LLDP data will be retained, per device.
LLDP maximum neighbors per port	The maximum number of LLDP neighbors for which LLDP data will be retained, per port.

LLDP statistics

The **show lldp statistics** command displays an overview of LLDP neighbor detection on the device, as well as packet counters and protocol statistics. The statistics are displayed on a global basis.

The following shows an example report.

```
PowerConnect#show lldp statistics
Last neighbor change time: 23 hours 50 minutes 40 seconds ago

Neighbor entries added          : 14
Neighbor entries deleted        : 5
Neighbor entries aged out       : 4
Neighbor advertisements dropped : 0
```

Port	Tx Pkts Total	Rx Pkts Total	Rx Pkts w/Errors	Rx Pkts Discarded	Rx TLVs Unrecognz	Rx TLVs Discarded	Neighbors Aged Out
1	60963	75179	0	0	0	0	4
2	0	0	0	0	0	0	0
3	60963	60963	0	0	0	0	0
4	60963	121925	0	0	0	0	0
5	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0
10	60974	0	0	0	0	0	0
11	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0
14	0	0	0	0	0	0	0

Syntax: show lldp statistics

NOTE

You can reset LLDP statistics using the CLI command **clear LLDP statistics**. Refer to [“The contents of the show output will vary depending on which TLVs are configured to be advertised.”](#) on page 561.

The following table describes the information displayed by the **show lldp statistics** command.

Table 3:

This field...	Displays...
Last neighbor change time	The elapsed time (in hours, minutes, and seconds) since a neighbor last advertised information. For example, the elapsed time since a neighbor was last added, deleted, or its advertised information changed.
Neighbor entries added	The number of new LLDP neighbors detected since the last reboot or since the last time the clear lldp statistics all command was issued.
Neighbor entries deleted	The number of LLDP neighbors deleted since the last reboot or since the last time the clear lldp statistics all command was issued.
Neighbor entries aged out	The number of LLDP neighbors dropped on all ports after the time-to-live expired. Note that LLDP entries age out naturally when a port cable or module is disconnected or when a port becomes disabled. However, if a disabled port is re-enabled, the system will delete the old LLDP entries.
Neighbor advertisements dropped	The number of valid LLDP neighbors the device detected, but could not add. This can occur, for example, when a new neighbor is detected and the device is already supporting the maximum number of neighbors possible. This can also occur when an LLDPDU is missing a mandatory TLV or is not formatted correctly.
Port	The local port number.
Tx Pkts Total	The number of LLDP packets the port transmitted.
Rx Pkts Total	The number of LLDP packets the port received.
Rx Pkts w/Errors	The number of LLDP packets the port received that have one or more detectable errors.
Rx Pkts Discarded	The number of LLDP packets the port received then discarded.
Rx TLVs Unrecognz	The number of TLVs the port received that were not recognized by the LLDP local agent. Unrecognized TLVs are retained by the system and can be viewed in the output of the show LLDP neighbors detail command or retrieved through SNMP.
Rx TLVs Discarded	The number of TLVs the port received then discarded.
Neighbors Aged Out	The number of times a neighbor information was deleted because its TTL timer expired.

LLDP neighbors

The **show lldp neighbors** command displays a list of the current LLDP neighbors per port.

The following shows an example report.

```
PowerConnect#show lldp neighbors
Lcl Port Chassis ID      Port ID      Port Description      System Name
5         0024.3876.29c0  0024.3878.04b4  10GigabitEthernet3/2/4  FCX648 Switch
6         0024.3876.29c0  0024.3878.04b4  10GigabitEthernet2/2/4  FCX648 Switch
9         0024.3823.7900  0024.3823.79d8  10GigabitEthernet10/1   FastIron SX 8~
25        000c.dbf1.8580  000c.dbf1.8583  GigabitEthernet4        FESX424 Switc~
26        000c.dbf1.8580  000c.dbf1.8582  GigabitEthernet3        FESX424 Switc~
```

Syntax: show lldp neighbors

The following table describes the information displayed by the **show lldp neighbors** command.

Table 4:

This field...	Displays...
Lcl Port	The local LLDP port number.
Chassis ID	The identifier for the chassis. Devices use the base MAC address of the device as the Chassis ID.
Port ID	The identifier for the port. Devices use the permanent MAC address associated with the port as the port ID.
Port Description	The description for the port. Devices use the ifDescr MIB object from MIB-II as the port description.
System Name	The administratively-assigned name for the system. Devices use the sysName MIB object from MIB-II, which corresponds to the CLI hostname command setting. NOTE: A tilde (~) at the end of a line indicates that the value in the field is too long to display in full and is truncated.

LLDP neighbors detail

The **show lldp neighbors detail** command displays the LLDP advertisements received from LLDP neighbors.

The following shows an example **show lldp neighbors detail** report.

NOTE

The **show lldp neighbors detail** output will vary depending on the data received. Also, values that are not recognized or do not have a recognizable format, may be displayed in hexadecimal binary form.

```
PowerConnect#show lldp neighbors detail ports e 9
Local port: 9
Neighbor: 0800.0f18.cc03, TTL 101 seconds
+ Chassis ID (network address): 10.43.39.151
+ Port ID (MAC address): 0800.0f18.cc03
+ Time to live: 120 seconds
+ Port description      : "LAN port"
+ System name          : "regDN 1015,MITEL 5235 DM"
+ System description   : "regDN 1015,MITEL 5235 DM,h/w rev 2,ASIC rev 1,f/w\
                        Boot 02.01.00.11,f/w Main 02.01.00.11"
+ System capabilities  : bridge, telephone
  Enabled capabilities: bridge, telephone
+ Management address (IPv4): 10.43.39.151
+ 802.3 MAC/PHY       : auto-negotiation enabled
  Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                        100BaseTX-FD
  Operational MAU type  : 100BaseTX-FD
+ MED capabilities: capabilities, networkPolicy, extendedPD
  MED device type      : Endpoint Class III
+ MED Network Policy
  Application Type     : Voice
  Policy Flags        : Known Policy, Tagged
  VLAN ID             : 300
  L2 Priority          : 7
  DSCP Value          : 7
+ MED Extended Power via MDI
  Power Type          : PD device
  Power Source        : Unknown Power Source
  Power Priority       : High (2)
  Power Value         : 6.2 watts (PSE equivalent: 6656 mWatts)
+ MED Hardware revision : "PCB Version: 2"
+ MED Firmware revision : "Boot 02.01.00.11"
+ MED Software revision : "Main 02.01.00.11"
+ MED Serial number     : ""
+ MED Manufacturer      : "Mitel Corporation"
+ MED Model name        : "MITEL 5235 DM"
+ MED Asset ID          : ""
```

A backslash (\) at the end of a line indicates that the text continues on the next line.

Except for the following field, the fields in the above output are described in the individual TLV advertisement sections in this chapter.

Table 5:

This field...	Displays...
Neighbor	The source MAC address from which the packet was received, and the remaining TTL for the neighbor entry.

Syntax: `show lldp neighbors detail [ports ethernet <port-list> | all]`

If you do not specify any ports or use the keyword **all**, by default, the report will show the LLDP neighbor details for all ports.

LLDP configuration details

The **show lldp local-info** command displays the local information advertisements (TLVs) that will be transmitted by the LLDP agent.

NOTE

The **show lldp local-info** output will vary based on LLDP configuration settings

```
PowerConnect#show lldp local-info
Local port: 5
+ Chassis ID (MAC address): 0024.3817.50bb
+ Port ID (MAC address): 0024.3817.50bf
+ Time to live: 120 seconds
+ System name          : "TX24 Router"
+ Port description     : "10GigabitEthernet5"
+ System capabilities : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY       : auto-negotiation supported, but disabled
  Operational MAU type : 10GigBaseSR
+ Link aggregation: aggregated (aggregated port ifIndex: 5)
+ Maximum frame size: 1522 octets
+ Port VLAN ID: none
+ Management address (IPv4): 2.2.2.1

Local port: 6
+ Chassis ID (MAC address): 0024.3817.50bb
+ Port ID (MAC address): 0024.3817.50c0
+ Time to live: 120 seconds
+ System name          : "TX24 Router"
+ Port description     : "10GigabitEthernet6"
+ System capabilities : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY       : auto-negotiation supported, but disabled
  Operational MAU type : 10GigBaseSR
+ Link aggregation: aggregated (aggregated port ifIndex: 5)
+ Maximum frame size: 1522 octets
+ Port VLAN ID: none
+ Management address (IPv4): 2.2.2.1
```

Syntax: `show lldp local-info [ports ethernet <port-list> | all]`

If you do not specify any ports or use the keyword `all`, by default, the report will show the local information advertisements for all ports.

```
PowerConnect#show lldp local-info ports ethernet 28
Local port: 28
+ Chassis ID (MAC address): 0024.3817.50bb
+ Port ID (MAC address): 0024.3817.50d6
+ Time to live: 120 seconds
+ System name          : "TX24 Router"
+ Port description     : "GigabitEthernet28"
+ System capabilities  : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY        : auto-negotiation enabled
  Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                           100BaseTX-FD, fdxSPause, fdxBPause, 1000BaseT-HD,
                           1000BaseT-FD
  Operational MAU type   : 1000BaseT-FD
+ Link aggregation: not capable
+ Maximum frame size: 1522 octets
+ Port VLAN ID: 4000
+ Management address (IPv4): 10.20.64.246
```

NOTE

The contents of the **show** output will vary depending on which TLVs are configured to be advertised.

A backslash (\) at the end of a line indicates that the text continues on the next line.

The fields in the below outputs are described in the individual TLV advertisement sections in this chapter.

Resetting LLDP statistics

To reset LLDP statistics, enter the **clear lldp statistics** command at the Global CONFIG level of the CLI. The device will clear the global and per-port LLDP neighbor statistics on the device (refer to [“LLDP statistics”](#) on page 556).

```
PowerConnect# clear lldp statistics
```

Syntax: clear lldp statistics [ports ethernet <port-list> | all]

If you do not specify any ports or use the keyword **all**, by default, the system will clear lldp statistics on all ports.

Clearing cached LLDP neighbor information

The device clears cached LLDP neighbor information after a port becomes disabled and the LLDP neighbor information ages out. However, if a port is disabled then re-enabled before the neighbor information ages out, the device will clear the cached LLDP neighbor information when the port is re-enabled.

If desired, you can manually clear the cache. For example, to clear the cached LLDP neighbor information for port e 20, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect#clear lldp neighbors ports e 20
```

Syntax: clear lldp neighbors [ports ethernet <port-list> | all]

20 Clearing cached LLDP neighbor information

If you do not specify any ports or use the keyword **all**, by default, the system will clear the cached LLDP neighbor information for all ports.

Configuring IP

Basic configuration

NOTE

The terms *Layer 3 Switch* and *router* are used interchangeably in this chapter and mean the same thing.

IP is enabled by default. Basic configuration consists of adding IP addresses and, for Layer 3 Switches, enabling a route exchange protocol, such as Routing Information Protocol (RIP).

If you are configuring a Layer 3 Switch, refer to [“Configuring IP addresses”](#) on page 579 to add IP addresses, then refer to one or more of the following to enable and configure the route exchange protocols:

- [Chapter 22, “Configuring RIP”](#)
- [Chapter 23, “Configuring OSPF Version 2 \(IPv4\)”](#)

If you are configuring a Layer 2 Switch, refer to [“Configuring the management IP address and specifying the default gateway”](#) on page 616 to add an IP address for management access through the network and to specify the default gateway.

The rest of this chapter describes IP and how to configure it in more detail. Use the information in this chapter if you need to change some of the IP parameters from their default values or you want to view configuration information or statistics.

Overview

Layer 2 Switches and Layer 3 Switches support Internet Protocol (IP) version 4. IP support on Layer 2 Switches consists of basic services to support management access and access to a default gateway. IP support on Layer 3 Switches includes all of the following, in addition to a highly configurable implementation of basic IP services including Address Resolution Protocol (ARP), ICMP Router Discovery Protocol (IRDP), and Reverse ARP (RARP):

- Route exchange protocols:
 - Routing Information Protocol (RIP)
 - Open Shortest Path First (OSPF)
- Multicast protocols:
 - Internet Group Membership Protocol (IGMP)
 - Protocol Independent Multicast Dense (PIM-DM)
 - Protocol Independent Multicast Sparse (PIM-SM)
- Router redundancy protocols:
 - Virtual Router Redundancy Protocol Extended (VRRPE)
 - Virtual Router Redundancy Protocol (VRRP)

IP interfaces

Layer 3 Switches and Layer 2 Switches allow you to configure IP addresses. On Layer 3 Switches, IP addresses are associated with individual interfaces. On Layer 2 Switches, a single IP address serves as the management access address for the entire device.

All Layer 3 Switches and Layer 2 Switches support configuration and display of IP address in classical subnet format (example: 192.168.1.1 255.255.255.0) and Classless Interdomain Routing (CIDR) format (example: 192.168.1.1/24). You can use either format when configuring IP address information. IP addresses are displayed in classical subnet format by default but you can change the display format to CIDR. Refer to [“Changing the network mask display to prefix format”](#) on page 623.

Layer 3 Switches

Layer 3 Switches allow you to configure IP addresses on the following types of interfaces:

- Ethernet ports
- Virtual routing interfaces (used by VLANs to route among one another)
- Loopback interfaces

Each IP address on a Layer 3 Switch must be in a different subnet. You can have only one interface that is in a given subnet. For example, you can configure IP addresses 192.168.1.1/24 and 192.168.2.1/24 on the same Layer 3 Switch, but you cannot configure 192.168.1.1/24 and 192.168.1.2/24 on the same Layer 3 Switch.

You can configure multiple IP addresses on the same interface.

The number of IP addresses you can configure on an individual interface depends on the Layer 3 Switch model. To display the maximum number of IP addresses and other system parameters you can configure on a Layer 3 Switch, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184.

You can use any of the IP addresses you configure on the Layer 3 Switch for Telnet, or SNMP access.

Layer 2 Switches

You can configure an IP address on a Layer 2 Switch for management access to the Layer 2 Switch. An IP address is required for Telnet access, and SNMP access.

You also can specify the default gateway for forwarding traffic to other subnets.

IP packet flow through a Layer 3 Switch

[Figure 97](#) shows how an IP packet moves through a Layer 3 Switch.

FIGURE 97 IP Packet Flow through a Layer 3 Switch

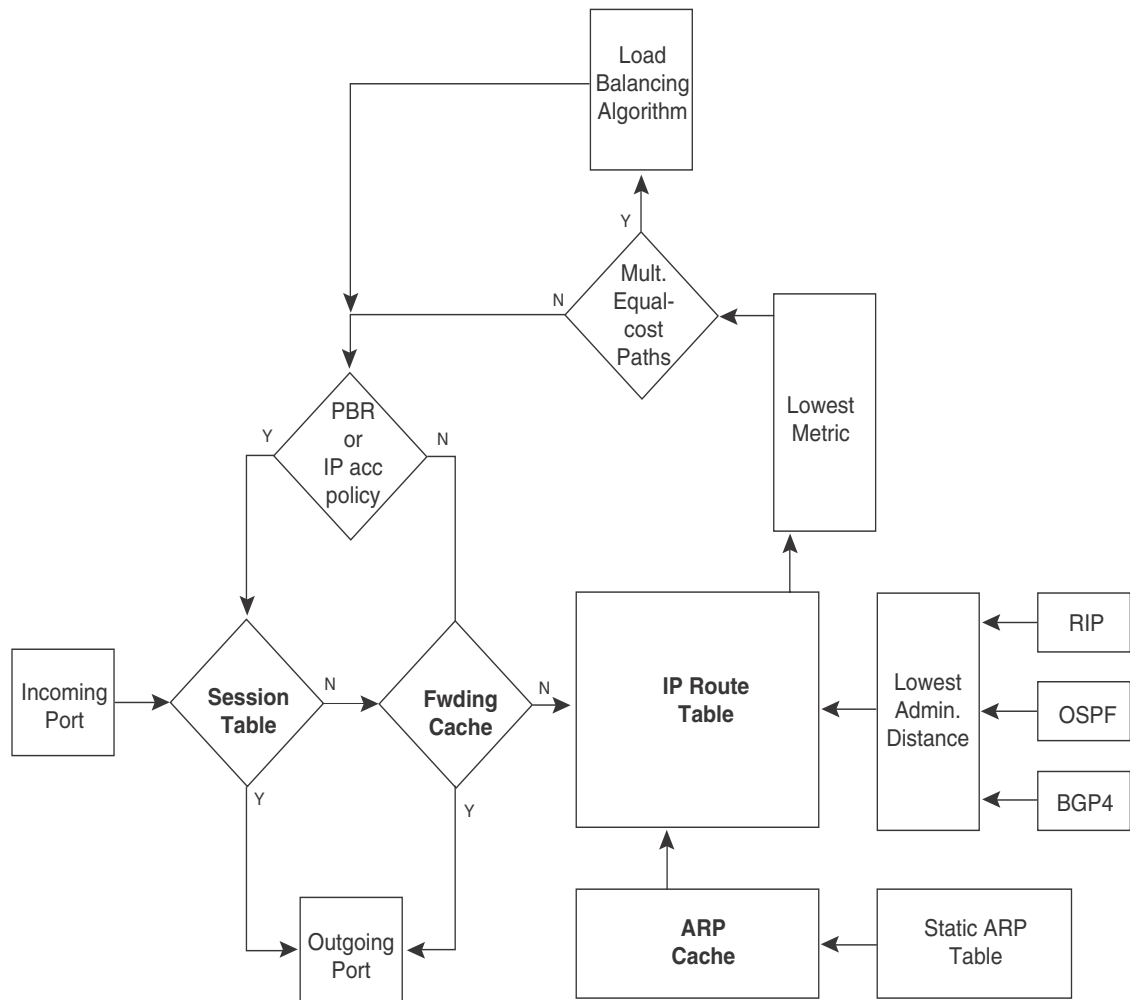


Figure 97 shows the following packet flow:

1. When the Layer 3 Switch receives an IP packet, the Layer 3 Switch checks for filters on the receiving interface.¹ If a deny filter on the interface denies the packet, the Layer 3 Switch discards the packet and performs no further processing, except generating a Syslog entry and SNMP message, if logging is enabled for the filter.
 2. If the packet is not denied at the incoming interface, the Layer 3 Switch looks in the session table for an entry that has the same source IP address and TCP or UDP port as the packet. If the session table contains a matching entry, the Layer 3 Switch immediately forwards the packet, by addressing it to the destination IP address and TCP or UDP port listed in the session table entry and sending the packet to a queue on the outgoing ports listed in the session table. The Layer 3 Switch selects the queue based on the Quality of Service (QoS) level associated with the session table entry.
 3. If the session table does not contain an entry that matches the packet source address and TCP or UDP port, the Layer 3 Switch looks in the IP forwarding cache for an entry that matches the packet destination IP address. If the forwarding cache contains a matching entry, the Layer 3 Switch forwards the packet to the IP address in the entry. The Layer 3 Switch sends the packet to a queue on the outgoing ports listed in the forwarding cache. The Layer 3 Switch selects the queue based on the Quality of Service (QoS) level associated with the forwarding cache entry.
1. The filter can be an Access Control List (ACL) or an IP access policy.

4. If the IP forwarding cache does not have an entry for the packet, the Layer 3 Switch checks the IP route table for a route to the packet destination. If the IP route table has a route, the Layer 3 Switch makes an entry in the session table or the forwarding cache, and sends the route to a queue on the outgoing ports:
 - If the running-config contains an IP access policy for the packet, the software makes an entry in the session table. The Layer 3 Switch uses the new session table entry to forward subsequent packets from the same source to the same destination.
 - If the running-config does not contain an IP access policy for the packet, the software creates a new entry in the forwarding cache. The Layer 3 Switch uses the new cache entry to forward subsequent packets to the same destination.

The following sections describe the IP tables and caches:

- ARP cache and static ARP table
- IP route table
- IP forwarding cache
- IP session table

The software enables you to display these tables. You also can change the capacity of the tables on an individual basis if needed by changing the memory allocation for the table.

ARP cache and static ARP table

The ARP cache contains entries that map IP addresses to MAC addresses. Generally, the entries are for devices that are directly attached to the Layer 3 Switch.

An exception is an ARP entry for an interface-based static IP route that goes to a destination that is one or more router hops away. For this type of entry, the MAC address is either the destination device MAC address or the MAC address of the router interface that answered an ARP request on behalf of the device, using proxy ARP.

ARP cache

The ARP cache can contain dynamic (learned) entries and static (user-configured) entries. The software places a dynamic entry in the ARP cache when the Layer 3 Switch learns a device MAC address from an ARP request or ARP reply from the device.

The software can learn an entry when the Layer 2 Switch or Layer 3 Switch receives an ARP request from another IP forwarding device or an ARP reply. Here is an example of a dynamic entry:

	IP Address	MAC Address	Type	Age	Port
1	207.95.6.102	0800.5afc.ea21	Dynamic	0	6

Each entry contains the destination device IP address and MAC address.

Static ARP table

In addition to the ARP cache, Layer 3 Switches have a static ARP table. Entries in the static ARP table are user-configured. You can add entries to the static ARP table regardless of whether the device the entry is for is connected to the Layer 3 Switch.

NOTE

Layer 3 Switches have a static ARP table. Layer 2 Switches do not.

The software places an entry from the static ARP table into the ARP cache when the entry interface comes up.

Here is an example of a static ARP entry.

Index	IP Address	MAC Address	Port
1	207.95.6.111	0800.093b.d210	1

Each entry lists the information you specified when you created the entry.

To display ARP entries, refer to the following:

- [“Displaying the ARP cache”](#) on page 628 – Layer 3 Switch
- [“Displaying the static ARP table”](#) on page 629 – Layer 3 Switch only
- [“Displaying ARP entries”](#) on page 638 – Layer 2 Switch

To configure other ARP parameters, refer to the following:

- [“Configuring ARP parameters”](#) on page 587 – Layer 3 Switch only

To increase the size of the ARP cache and static ARP table, refer to the following:

- For dynamic entries, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184. The ip-arp parameter controls the ARP cache size.
- Static entries, [“Changing the maximum number of entries the static ARP table can hold”](#) on page 591 – Layer 3 Switches only. The ip-static-arp parameter controls the static ARP table size.

IP route table

The IP route table contains paths to IP destinations.

NOTE

Layer 2 Switches do not have an IP route table. A Layer 2 Switch sends all packets addressed to another subnet to the default gateway, which you specify when you configure the basic IP information on the Layer 2 Switch.

The IP route table can receive the paths from the following sources:

- A directly-connected destination, which means there are no router hops to the destination
- A static IP route, which is a user-configured route
- A route learned through RIP
- A route learned through OSPF
- A route learned through BGP4

The IP route table contains the best path to a destination:

- When the software receives paths from more than one of the sources listed above, the software compares the administrative distance of each path and selects the path with the lowest administrative distance. The administrative distance is a protocol-independent value from 1 – 255.
- When the software receives two or more best paths from the same source and the paths have the same metric (cost), the software can load share traffic among the paths based on destination host or network address (based on the configuration and the Layer 3 Switch model).

Here is an example of an entry in the IP route table.

Destination	NetMask	Gateway	Port	Cost	Type
1.1.0.0	255.255.0.0	99.1.1.2	1	2	R

Each IP route table entry contains the destination IP address and subnet mask and the IP address of the next-hop router interface to the destination. Each entry also indicates the port attached to the destination or the next-hop to the destination, the route IP metric (cost), and the type. The type indicates how the IP route table received the route.

To display the IP route table, refer to [“Displaying the IP route table”](#) on page 632 (Layer 3 Switch only)

To configure a static IP route, refer to [“Configuring static routes”](#) on page 596 (Layer 3 Switch only)

To clear a route from the IP route table, refer to [“Clearing IP routes”](#) on page 634 (Layer 3 Switch only)

To increase the size of the IP route table for learned and static routes, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184:

- For learned routes, modify the ip-route parameter.
- For static routes, modify the ip-static-route parameter.

IP forwarding cache

The IP forwarding cache provides a fast-path mechanism for forwarding IP packets. The cache contains entries for IP destinations. When a Layer 3 Switch has completed processing and addressing for a packet and is ready to forward the packet, the device checks the IP forwarding cache for an entry to the packet destination:

- If the cache contains an entry with the destination IP address, the device uses the information in the entry to forward the packet out the ports listed in the entry. The destination IP address is the address of the packet final destination. The port numbers are the ports through which the destination can be reached.
- If the cache does not contain an entry and the traffic does not qualify for an entry in the session table instead, the software can create an entry in the forwarding cache.

Each entry in the IP forwarding cache has an age timer. If the entry remains unused for ten minutes, the software removes the entry. The age timer is not configurable.

Here is an example of an entry in the IP forwarding cache.

	IP Address	Next Hop	MAC	Type	Port	Vlan	Pri
1	192.168.1.11	DIRECT	0000.0000.0000	PU	n/a		0

Each IP forwarding cache entry contains the IP address of the destination, and the IP address and MAC address of the next-hop router interface to the destination. If the destination is actually an interface configured on the Layer 3 Switch itself, as shown here, then next-hop information indicates this. The port through which the destination is reached is also listed, as well as the VLAN and Layer 4 QoS priority associated with the destination if applicable.

To display the IP forwarding cache, refer to [“Displaying the forwarding cache”](#) on page 631.

NOTE

You cannot add static entries to the IP forwarding cache, although you can increase the number of entries the cache can contain. Refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184.

To increase the size of the IP forwarding cache, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184.

Layer 4 session table

The Layer 4 session provides a fast path for forwarding packets. A **session** is an entry that contains complete Layer 3 and Layer 4 information for a flow of traffic. Layer 3 information includes the source and destination IP addresses. Layer 4 information includes the source and destination TCP and UDP ports. For comparison, the IP forwarding cache contains the Layer 3 destination address but does not contain the other source and destination address information of a Layer 4 session table entry.

The Layer 2 Switch or Layer 3 Switch selects the session table instead of the IP forwarding table for fast-path forwarding for the following features:

- Layer 4 Quality-of-Service (QoS) policies
- IP access policies

To increase the size of the session table, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184. The ip-qos-session parameter controls the size of the session table.

IP route exchange protocols

Layer 3 Switches support the following IP route exchange protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

All these protocols provide routes to the IP route table. You can use one or more of these protocols, in any combination. The protocols are disabled by default. For configuration information, refer to the following:

- [Chapter 22, “Configuring RIP”](#)
- [Chapter 23, “Configuring OSPF Version 2 \(IPv4\)”](#)

IP multicast protocols

Layer 3 Switches also support the following Internet Group Membership Protocol (IGMP) based IP multicast protocols:

- Protocol Independent Multicast – Dense mode (PIM-DM)
- Protocol Independent Multicast – Sparse mode (PIM-SM)

For configuration information, refer to [Chapter 19, “Configuring IP Multicast Protocols”](#).

NOTE

Layer 2 Switches support IGMP and can forward IP multicast packets. Refer to [Chapter 18, “Configuring IP Multicast Traffic Reduction for PowerConnect B-Series TI24X Switches”](#).

IP interface redundancy protocols

You can configure a Layer 3 Switch to back up an IP interface configured on another Layer 3 Switch. If the link for the backed up interface becomes unavailable, the other Layer 3 Switch can continue service for the interface. This feature is especially useful for providing a backup to a network default gateway.

Layer 3 Switches support the following IP interface redundancy protocols:

- Virtual Router Redundancy Protocol (VRRP) – A standard router redundancy protocol based on RFC 2338. You can use VRRP to configure Layer 3 Switches and third-party routers to back up IP interfaces on other Layer 3 Switches or third-party routers.
- Virtual Router Redundancy Protocol Extended (VRRPE) – An extension to standard VRRP that adds additional features and overcomes limitations in standard VRRP. You can use VRRPE only on Layer 3 Switches.

For configuration information, refer to the following:

- Virtual Router Redundancy Protocol Extended (VRRPE) – refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).
- Virtual Router Redundancy Protocol (VRRP) – refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).

Access Control Lists and IP access policies

Layer 3 Switches provide two mechanisms for filtering IP traffic:

- Access Control Lists (ACLs)
- IP access policies

Both methods allow you to filter packets based on Layer 3 and Layer 4 source and destination information.

ACLs also provide great flexibility by providing the input to various other filtering mechanisms such as route maps, which are used by BGP4.

IP access policies allow you to configure QoS based on sessions (Layer 4 traffic flows).

Only one of these filtering mechanisms can be enabled on a device at a time. Devices can store forwarding information for both methods of filtering in the session table.

For configuration information, [Chapter 13, “Configuring Rule-Based IP Access Control Lists”](#)

Basic IP parameters and defaults – Layer 3 Switches

IP is enabled by default. The following IP-based protocols are all disabled by default:

- Routing protocols:
 - Routing Information Protocol (RIP) – refer to [Chapter 22, “Configuring RIP”](#)
 - Open Shortest Path First (OSPF) – refer to [Chapter 23, “Configuring OSPF Version 2 \(IPv4\)”](#)
- Multicast protocols:
 - Internet Group Membership Protocol (IGMP) – refer to [“Changing global IP multicast parameters”](#) on page 467

- Protocol Independent Multicast Dense (PIM-DM) – refer to “[PIM Dense](#)” on page 470
- Protocol Independent Multicast Sparse (PIM-SM) – refer to “[PIM Sparse](#)” on page 478
- Router redundancy protocols:
 - Virtual Router Redundancy Protocol Extended (VRRPE) – refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).
 - Virtual Router Redundancy Protocol (VRRP) – refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).

The following tables list the Layer 3 Switch IP parameters, their default values, and where to find configuration information.

NOTE

For information about parameters in other protocols based on IP, such as RIP, OSPF, and so on, refer to the configuration chapters for those protocols.

When parameter changes take effect

Most IP parameters described in this chapter are dynamic. They take effect immediately, as soon as you enter the CLI command. You can verify that a dynamic change has taken effect by displaying the running-config. To display the running-config, enter the **show running-config** or **write terminal** command at any CLI prompt.

To save a configuration change permanently so that the change remains in effect following a system reset or software reload, save the change to the startup-config file:

- To save configuration changes to the startup-config file, enter the **write memory** command from the Privileged EXEC level of any configuration level of the CLI.

Changes to memory allocation require you to reload the software after you save the changes to the startup-config file. When reloading the software is required to complete a configuration change described in this chapter, the procedure that describes the configuration change includes a step for reloading the software.

IP global parameters – Layer 3 Switches

[Table 96](#) lists the IP global parameters for Layer 3 Switches.

TABLE 96 IP global parameters – Layer 3 Switches

Parameter	Description	Default	See page...
IP state	The Internet Protocol, version 4	Enabled NOTE: You cannot disable IP.	n/a
IP address and mask notation	Format for displaying an IP address and its network mask information. You can enable one of the following: <ul style="list-style-type: none"> Class-based format; example: 192.168.1.1 255.255.255.0 Classless Interdomain Routing (CIDR) format; example: 192.168.1.1/24 	Class-based NOTE: Changing this parameter affects the display of IP addresses, but you can enter addresses in either format regardless of the display setting.	page 623
Router ID	The value that routers use to identify themselves to other routers when exchanging route information. OSPF and BGP4 use router IDs to identify routers. RIP does not use the router ID.	The IP address configured on the lowest-numbered loopback interface. If no loopback interface is configured, then the lowest-numbered IP address configured on the device.	page 584
Maximum Transmission Unit (MTU)	The maximum length an Ethernet packet can be without being fragmented.	1500 bytes for Ethernet II encapsulation 1492 bytes for SNAP encapsulation	page 582
Address Resolution Protocol (ARP)	A standard IP mechanism that routers use to learn the Media Access Control (MAC) address of a device on the network. The router sends the IP address of a device in the ARP request and receives the device MAC address in an ARP reply.	Enabled	page 587
ARP rate limiting	Lets you specify a maximum number of ARP packets the device will accept each second. If the device receives more ARP packets than you specify, the device drops additional ARP packets for the remainder of the one-second interval.	Disabled	page 588
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age. NOTE: You also can change the ARP age on an individual interface basis. Refer to Table 97 on page 575.	Ten minutes	page 589
Proxy ARP	An IP mechanism a router can use to answer an ARP request on behalf of a host, by replying with the router own MAC address instead of the host.	Disabled	page 589
Static ARP entries	An ARP entry you place in the static ARP table. Static entries do not age out.	No entries	page 590

TABLE 96 IP global parameters – Layer 3 Switches (Continued)

Parameter	Description	Default	See page...
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	page 592
Directed broadcast forwarding	A directed broadcast is a packet containing all ones (or in some cases, all zeros) in the host portion of the destination IP address. When a router forwards such a broadcast, it sends a copy of the packet out each of its enabled IP interfaces. NOTE: You also can enable or disable this parameter on an individual interface basis. Refer to Table 97 on page 575.	Disabled	page 592
Directed broadcast mode	The packet format the router treats as a directed broadcast. The following formats can be directed broadcast: <ul style="list-style-type: none"> All ones in the host portion of the packet destination address. All zeroes in the host portion of the packet destination address. 	All ones NOTE: If you enable all-zeroes directed broadcasts, all-ones directed broadcasts remain enabled.	page 593
Source-routed packet forwarding	A source-routed packet contains a list of IP addresses through which the packet must pass to reach its destination.	Enabled	page 593
Internet Control Message Protocol (ICMP) messages	The Layer 3 Switch can send the following types of ICMP messages: <ul style="list-style-type: none"> Echo messages (ping messages) Destination Unreachable messages 	Enabled	page 594
ICMP Router Discovery Protocol (IRDP)	An IP protocol a router can use to advertise the IP addresses of its router interfaces to directly attached hosts. You can enable or disable the protocol, and change the following protocol parameters: <ul style="list-style-type: none"> Forwarding method (broadcast or multicast) Hold time Maximum advertisement interval Minimum advertisement interval Router preference level NOTE: You also can enable or disable IRDP and configure the parameters on an individual interface basis. Refer to Table 97 on page 575.	Disabled	page 608
Reverse ARP (RARP)	An IP mechanism a host can use to request an IP address from a directly attached router when the host boots.	Enabled	page 610

TABLE 96 IP global parameters – Layer 3 Switches (Continued)

Parameter	Description	Default	See page...
Static RARP entries	An IP address you place in the RARP table for RARP requests from hosts. NOTE: You must enter the RARP entries manually. The Layer 3 Switch does not have a mechanism for learning or dynamically generating RARP entries.	No entries	page 611
Maximum BootP relay hops	The maximum number of hops away a BootP server can be located from a router and still be used by the router clients for network booting.	Four	page 616
Domain name for Domain Name Server (DNS) resolver	A domain name (example: brocade.router.com) you can use in place of an IP address for certain operations such as IP pings, trace routes, and Telnet management connections to the router.	None configured	page 581
DNS default gateway addresses	A list of gateways attached to the router through which clients attached to the router can reach DNSs.	None configured	page 581
IP load sharing	A feature that enables the router to balance traffic to a specific destination across multiple equal-cost paths. NOTE: Load sharing is sometimes called Equal Cost Multi Path (ECMP).	Enabled	page 605
Maximum IP load sharing paths	The maximum number of equal-cost paths across which the Layer 3 Switch is allowed to distribute traffic.	Four	page 608
Origination of default routes	You can enable a router to originate default routes for the following route exchange protocols, on an individual protocol basis: <ul style="list-style-type: none"> • RIP • OSPF • BGP4 	Disabled	page 650 page 690 page 773
Default network route	The router uses the default network route if the IP route table does not contain a route to the destination and also does not contain an explicit default route (0.0.0.0 0.0.0.0 or 0.0.0.0/0).	None configured	page 604
Static route	An IP route you place in the IP route table.	No entries	page 596
Source interface	The IP address the router uses as the source address for Telnet, RADIUS, or TACACS/TACACS+ packets originated by the router. The router can select the source address based on either of the following: <ul style="list-style-type: none"> • The lowest-numbered IP address on the interface the packet is sent on. • The lowest-numbered IP address on a specific interface. The address is used as the source for all packets of the specified type regardless of interface the packet is sent on. 	The lowest-numbered IP address on the interface the packet is sent on.	page 585

IP interface parameters – Layer 3 Switches

Table 97 lists the interface-level IP parameters for Layer 3 Switches.

TABLE 97 IP interface parameters – Layer 3 Switches

Parameter	Description	Default	See page...
IP state	The Internet Protocol, version 4	Enabled NOTE: You cannot disable IP.	n/a
IP address	A Layer 3 network interface address NOTE: Layer 2 Switches have a single IP address used for management access to the entire device. Layer 3 Switches have separate IP addresses on individual interfaces.	None configured ¹	page 579
Encapsulation type	The format of the packets in which the router encapsulates IP datagrams. The encapsulation format can be one of the following: <ul style="list-style-type: none"> Ethernet II SNAP 	Ethernet II	page 582
Maximum Transmission Unit (MTU)	The maximum length (number of bytes) of an encapsulated IP datagram the router can forward.	1500 for Ethernet II encapsulated packets 1492 for SNAP encapsulated packets	page 583
ARP age	Locally overrides the global setting. Refer to Table 96 on page 572.	Ten minutes	page 589
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	page 646
Directed broadcast forwarding	Locally overrides the global setting. Refer to Table 96 on page 572.	Disabled	page 592
ICMP Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings. Refer to Table 96 on page 572.	Disabled	page 609

TABLE 97 IP interface parameters – Layer 3 Switches (Continued)

Parameter	Description	Default	See page...
UDP broadcast forwarding	<p>The router can forward UDP broadcast packets for UDP applications such as BootP. By forwarding the UDP broadcasts, the router enables clients on one subnet to find servers attached to other subnets.</p> <p>NOTE: To completely enable a client UDP application request to find a server on another subnet, you must configure an IP helper address consisting of the server IP address or the directed broadcast address for the subnet that contains the server. See the next row.</p>	<p>The router helps forward broadcasts for the following UDP application protocols:</p> <ul style="list-style-type: none"> • bootps • dns • netbios-dgm • netbios-ns • tacacs • tftp • time 	page 613
IP helper address	<p>The IP address of a UDP application server (such as a BootP or DHCP server) or a directed broadcast address. IP helper addresses allow the router to forward requests for certain UDP applications from a client on one subnet to a server on another subnet.</p>	None configured	page 614

1. Some devices have a factory default, such as 209.157.22.154, used for troubleshooting during installation. For Layer 3 Switches, the address is on module 1 port 1.

Basic IP parameters and defaults – Layer 2 Switches

IP is enabled by default. The following tables list the Layer 2 Switch IP parameters, their default values, and where to find configuration information.

NOTE

Layer 2 Switches also provide IP multicast forwarding, which is enabled by default. For information about this feature, refer to [Chapter 18, “Configuring IP Multicast Traffic Reduction for PowerConnect B-Series TI24X Switches”](#).

IP global parameters – Layer 2 Switches

[Table 98](#) lists the IP global parameters for Layer 2 Switches.

TABLE 98 IP global parameters – Layer 2 Switches

Parameter	Description	Default	See page...
IP address and mask notation	<p>Format for displaying an IP address and its network mask information. You can enable one of the following:</p> <ul style="list-style-type: none"> • Class-based format; example: 192.168.1.1 255.255.255.0 • Classless Interdomain Routing (CIDR) format; example: 192.168.1.1/24 	<p>Class-based</p> <p>NOTE: Changing this parameter affects the display of IP addresses, but you can enter addresses in either format regardless of the display setting.</p>	page 623
IP address	<p>A Layer 3 network interface address</p> <p>NOTE: Layer 2 Switches have a single IP address used for management access to the entire device. Layer 3 Switches have separate IP addresses on individual interfaces.</p>	None configured ¹	page 616
Default gateway	The IP address of a locally attached router (or a router attached to the Layer 2 Switch by bridges or other Layer 2 Switches). The Layer 2 Switch and clients attached to it use the default gateway to communicate with devices on other subnets.	None configured	page 616
Address Resolution Protocol (ARP)	A standard IP mechanism that networking devices use to learn the Media Access Control (MAC) address of another device on the network. The Layer 2 Switch sends the IP address of a device in the ARP request and receives the device MAC address in an ARP reply.	<p>Enabled</p> <p>NOTE: You cannot disable ARP.</p>	n/a
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age.	<p>Ten minutes</p> <p>NOTE: You cannot change the ARP age on Layer 2 Switches.</p>	n/a
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	page 619
Domain name for Domain Name Server (DNS) resolver	A domain name (example: brocade.router.com) you can use in place of an IP address for certain operations such as IP pings, trace routes, and Telnet management connections to the router.	None configured	page 617
DNS default gateway addresses	A list of gateways attached to the router through which clients attached to the router can reach DNSs.	None configured	page 617

TABLE 98 IP global parameters – Layer 2 Switches (Continued)

Parameter	Description	Default	See page...
Source interface	The IP address the Layer 2 Switch uses as the source address for Telnet, RADIUS, or TACACS/TACACS+ packets originated by the router. The Layer 2 Switch uses its management IP address as the source address for these packets.	The management IP address of the Layer 2 Switch. NOTE: This parameter is not configurable on Layer 2 Switches.	n/a
DHCP gateway stamp	The device can assist DHCP/BootP Discovery packets from one subnet to reach DHCP/BootP servers on a different subnet by placing the IP address of the router interface that forwards the packet in the packet Gateway field. You can specify up to 32 gateway lists. A gateway list contains up to eight gateway IP addresses. You activate DHCP assistance by associating a gateway list with a port. When you configure multiple IP addresses in a gateway list, the Layer 2 Switch inserts the addresses into the DHCP Discovery packets in a round robin fashion.	None configured	page 622

1. Some devices have a factory default, such as 209.157.22.154, used for troubleshooting during installation. For Layer 3 Switches, the address is on port 1.

Interface IP parameters – Layer 2 Switches

[Table 99](#) lists the interface-level IP parameters for Layer 2 Switches.

TABLE 99 Interface IP parameters – Layer 2 Switches

Parameter	Description	Default	See page...
DHCP gateway stamp	You can configure a list of DHCP stamp addresses for a port. When the port receives a DHCP/BootP Discovery packet from a client, the port places the IP addresses in the gateway list into the packet Gateway field.	None configured	page 622

Configuring IP parameters – Layer 3 Switches

The following sections describe how to configure IP parameters. Some parameters can be configured globally while others can be configured on individual interfaces. Some parameters can be configured globally and overridden for individual interfaces.

NOTE

This section describes how to configure IP parameters for Layer 3 Switches. For IP configuration information for Layer 2 Switches, refer to [“Configuring IP parameters – Layer 2 Switches”](#) on page 616.

Configuring IP addresses

You can configure an IP address on the following types of Layer 3 Switch interfaces:

- Ethernet port
- Virtual routing interface (also called a Virtual Ethernet or “VE”)
- Loopback interface

By default, you can configure up to 24 IP addresses on each interface.

On Compact Layer 3 Switches, you can increase this amount to up to 64 IP subnet addresses per port by increasing the size of the subnet-per-interface table.

Refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184.

NOTE

Once you configure a virtual routing interface on a VLAN, you cannot configure Layer 3 interface parameters on individual ports. Instead, you must configure the parameters on the virtual routing interface itself.

Devices support both classical IP network masks (Class A, B, and C subnet masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks:

- To enter a classical network mask, enter the mask in IP address format. For example, enter “209.157.22.99 255.255.255.0” for an IP address with a Class-C subnet mask.
- To enter a prefix network mask, enter a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter “209.157.22.99/24” for an IP address that has a network mask with 24 significant bits (ones).

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the display to prefix format. Refer to [“Changing the network mask display to prefix format”](#) on page 623.

Assigning an IP address to an Ethernet port

To assign an IP address to port 1, enter the following commands.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#ip address 192.45.6.1 255.255.255.0
```

You also can enter the IP address and mask in CIDR format, as follows.

```
PowerConnect(config-if-1)#ip address 192.45.6.1/24
```

Syntax: [no] ip address <ip-addr> <ip-mask> [ospf-ignore | ospf-passive | secondary]

or

Syntax: [no] ip address <ip-addr>/<mask-bits> [ospf-ignore | ospf-passive | secondary]

The **ospf-ignore** | **ospf-passive** parameters modify the Layer 3 Switch defaults for adjacency formation and interface advertisement. Use one of these parameters if you are configuring multiple IP subnet addresses on the interface but you want to prevent OSPF from running on some of the subnets:

- **ospf-passive** – This option disables adjacency formation with OSPF neighbors. By default, when OSPF is enabled on an interface, the software forms OSPF router adjacencies between each primary IP address on the interface and the OSPF neighbor attached to the interface.

- **ospf-ignore** – This option disables OSPF adjacency formation and also disables advertisement of the interface into OSPF. The subnet is completely ignored by OSPF.

NOTE

The **ospf-passive** option disables adjacency formation but does not disable advertisement of the interface into OSPF. To disable advertisement in addition to disabling adjacency formation, you must use the **ospf-ignore** option.

Use the **secondary** parameter if you have already configured an IP address within the same subnet on the interface.

NOTE

When you configure more than one address in the same subnet, all but the first address are secondary addresses and do not form OSPF adjacencies..

Assigning an IP address to a loopback interface

Loopback interfaces are always up, regardless of the states of physical interfaces. They can add stability to the network because they are not subject to route flap problems that can occur due to unstable links between a Layer 3 Switch and other devices. You can configure up to eight loopback interfaces on a Chassis Layer 3 Switch. You can configure up to four loopback interfaces on a Compact Layer 3 Switch.

You can add up to 24 IP addresses to each loopback interface.

NOTE

If you configure the Layer 3 Switch to use a loopback interface to communicate with a BGP4 neighbor, you also must configure a loopback interface on the neighbor and configure the neighbor to use that loopback interface to communicate with the Layer 3 Switch. Refer to [“Adding a loopback interface”](#) on page 756.

To add a loopback interface, enter commands such as those shown in the following example.

```
PowerConnect(config-bgp-router)#exit
PowerConnect(config)#int loopback 1
PowerConnect(config-lbif-1)#ip address 10.0.0.1/24
```

Syntax: `interface loopback <num>`

The `<num>` parameter specifies the virtual interface number. You can specify from 1 to the maximum number of virtual interfaces supported on the device. To display the maximum number of virtual interfaces supported on the device, enter the **show default values** command. The maximum is listed in the System Parameters section, in the Current column of the virtual-interface row.

Refer to the syntax description in [“Assigning an IP address to an Ethernet port”](#) on page 579.

Assigning an IP address to a virtual interface

A virtual interface is a logical port associated with a Layer 3 Virtual LAN (VLAN) configured on a Layer 3 Switch. You can configure routing parameters on the virtual interface to enable the Layer 3 Switch to route protocol traffic from one Layer 3 VLAN to the other, without using an external router.¹

1.

You can configure IP routing interface parameters on a virtual interface. This section describes how to configure an IP address on a virtual interface. Other sections in this chapter that describe how to configure interface parameters also apply to virtual interfaces.

NOTE

The Layer 3 Switch uses the lowest MAC address on the device (the MAC address of port 1) as the MAC address for all ports within all virtual interfaces you configure on the device.

To add a virtual interface to a VLAN and configure an IP address on the interface, enter commands such as the following.

```
PowerConnect(config)#vlan 2 name IP-Subnet_1.1.2.0/24
PowerConnect(config-vlan-2)#untag e1 to 4
PowerConnect(config-vlan-2)#router-interface ve1
PowerConnect(config-vlan-2)#interface ve1
PowerConnect(config-vif-1)#ip address 1.1.2.1/24
```

The first two commands in this example create a Layer 3 protocol-based VLAN name “IP-Subnet_1.1.2.0/24” and add a range of untagged ports to the VLAN. The **router-interface** command creates virtual interface 1 as the routing interface for the VLAN. The last two commands change to the interface configuration level for the virtual interface and assign an IP address to the interface.

Syntax: `router-interface ve <num>`

Syntax: `interface ve <num>`

Refer to the syntax description in [“Assigning an IP address to an Ethernet port”](#) on page 579.

Deleting an IP address

To delete an IP address, enter a command such as the following.

```
PowerConnect(config-if-e10000-1)#no ip address 1.1.2.1
```

This command deletes IP address 1.1.2.1. You do not need to enter the subnet mask.

To delete all IP addresses from an interface, enter the following command.

```
PowerConnect(config-if-e10000-1)#no ip address *
```

Syntax: `no ip address <ip-addr> | *`

Configuring packet parameters

You can configure the following packet parameters on Layer 3 Switches. These parameters control how the Layer 3 Switch sends IP packets to other devices on an Ethernet network. The Layer 3 Switch always places IP packets into Ethernet packets to forward them on an Ethernet port.

- **Encapsulation type** – The format for the Layer 2 packets within which the Layer 3 Switch sends IP packets.
- **Maximum Transmission Unit (MTU)** – The maximum length of IP packet that a Layer 2 packet can contain. IP packets that are longer than the MTU are fragmented and sent in multiple Layer 2 packets. You can change the MTU globally or on individual ports:
 - **Global MTU** – The default MTU value depends on the encapsulation type on a port and is 1500 bytes for Ethernet II encapsulation and 1492 bytes for SNAP encapsulation.
 - **Port MTU** – A port default MTU depends on the encapsulation type enabled on the port.

Changing the encapsulation type

The Layer 3 Switch encapsulates IP packets into Layer 2 packets, to send the IP packets on the network. (A Layer 2 packet is also called a MAC layer packet or an Ethernet frame.) The source address of a Layer 2 packet is the MAC address of the Layer 3 Switch interface sending the packet. The destination address can be one of the following:

- The MAC address of the IP packet destination. In this case, the destination device is directly connected to the Layer 3 Switch.
- The MAC address of the next-hop gateway toward the packet destination.
- An Ethernet broadcast address.

The entire IP packet, including the source and destination address and other control information and the data, is placed in the data portion of the Layer 2 packet. Typically, an Ethernet network uses one of two different formats of Layer 2 packet:

- Ethernet II
- Ethernet SNAP (also called IEEE 802.3)

The control portions of these packets differ slightly. All IP devices on an Ethernet network must use the same format. Layer 3 Switches use Ethernet II by default. You can change the IP encapsulation to Ethernet SNAP on individual ports if needed.

NOTE

All devices connected to the Layer 3 Switch port must use the same encapsulation type.

To change the IP encapsulation type on interface 5 to Ethernet SNAP, enter the following commands.

```
PowerConnect(config)#int e 5  
PowerConnect(config-if-e10000-5)#ip encapsulation snap
```

Syntax: ip encapsulation snap | ethernet_ii

Changing the Maximum Transmission Unit (MTU)

The Maximum Transmission Unit (MTU) is the maximum length of IP packet that a Layer 2 packet can contain. IP packets that are longer than the MTU are fragmented and sent in multiple Layer 2 packets. You can change the MTU globally or on individual ports.

The default MTU is 1500 bytes for Ethernet II packets and 1492 for Ethernet SNAP packets.

MTU enhancements

Devices contain the following enhancements to jumbo packet support:

- Hardware forwarding of Layer 3 jumbo packets – Layer 3 IP unicast jumbo packets received on a port that supports the frame's MTU size and forwarded to another port that also supports the frame's MTU size are forwarded in hardware. Previous releases support hardware forwarding of Layer 2 jumbo frames only.
- ICMP unreachable message if a frame is too large to be forwarded – If a jumbo packet has the Do not Fragment (DF) bit set, and the outbound interface does not support the packet's MTU size, the device sends an ICMP unreachable message to the device that sent the packet.

NOTE

These enhancements apply only to transit traffic forwarded through the device.

Configuration considerations for increasing the MTU

- When you increase the MTU size of a port, the increase uses system resources. Increase the MTU size only on the ports that need it. For example, if you have one port connected to a server that uses jumbo frames and two other ports connected to clients that can support the jumbo frames, increase the MTU only on those three ports. Leave the MTU size on the other ports at the default value (1500 bytes). Globally increase the MTU size only if needed.
- Use the same MTU size on all ports that will be supporting jumbo frames. If the device needs to fragment a jumbo frame (and the frame does not have the DF bit set), the device fragments the frame into 1500-byte fragments, even if the outbound port has a larger MTU. For example, if a port has an MTU setting of 8000 and receives an 8000-byte frame, then must forward the frame onto a port with an MTU of 4000, the device does not fragment the 8000-byte frame into two 4000-byte frames. Instead, the device fragments the 8000-byte frame into six fragments (five 1500-byte fragments and a final, smaller fragment.)

Globally changing the Maximum Transmission Unit (MTU)

The Maximum Transmission Unit (MTU) is the maximum size an IP packet can be when encapsulated in a Layer 2 packet. If an IP packet is larger than the MTU allowed by the Layer 2 packet, the Layer 3 Switch fragments the IP packet into multiple parts that will fit into the Layer 2 packets, and sends the parts of the fragmented IP packet separately, in different Layer 2 packets. The device that receives the multiple fragments of the IP packet reassembles the fragments into the original packet.

You can increase the MTU size to accommodate jumbo packet sizes up to 9198 bytes in a PowerConnect B-Series TI24X device.

To globally enable jumbo support on all ports, enter commands such as the following.

```
PowerConnect(config)#jumbo
PowerConnect(config)#write memory
PowerConnect(config)#end
PowerConnect#reload
```

Syntax: [no] jumbo

NOTE

You must save the configuration change and then reload the software to enable jumbo support.

Changing the Maximum Transmission Unit on an individual port

By default, the maximum Ethernet MTU sizes are as follows:

- 1500 bytes – The maximum for Ethernet II encapsulation
- 1492 bytes – The maximum for SNAP encapsulation

When jumbo mode is enabled, the maximum Ethernet MTU sizes are as follows:

- 9198 bytes – The maximum for Ethernet II encapsulation
- 9190 bytes – The maximum for SNAP encapsulation

To change the MTU for interface 5 to 1000, enter the following commands.

```
PowerConnect(config)#int e 5
PowerConnect(config-if-5)#ip-port-mtu 1000
PowerConnect(config-if-5)#write memory
PowerConnect(config-if-5)#end
```

Syntax: [no] ip-port-mtu <num>

NOTE

The new command **ip-port-mtu** replace the command **ip mtu**. On the PowerConnect the IP MTU check on egress is validated based on the physical port instead of the ip interface. Therefore, the command **ip-port-mtu** can be set only on a physical port. In the case of a VE, we can set the **ip-port-mtu** on a port member of a VE. In contrast with the **ip mtu** command, the multiple physical ports in a VE can have a different IP MTU. However, on PowerConnect, all VLANs of a port would have the same IP MTU size.

The `<num>` parameter specifies the MTU. Ethernet II packets can hold IP packets from 576 – 1500 bytes long. If jumbo mode is enabled, Ethernet II packets can hold IP packets up to 9198 bytes long. Ethernet SNAP packets can hold IP packets from 576 – 1492 bytes long. If jumbo mode is enabled, SNAP packets can hold IP packets up to 9190 bytes long. The default MTU for Ethernet II packets is 1500. The default MTU for SNAP packets is 1492.

Path MTU discovery (RFC 1191) support

Most devices support the path MTU discovery method described in RFC 1191. When the device receives an IP packet that has its Do not Fragment (DF) bit set, and the packet size is greater than the MTU value of the outbound interface, then the device returns an ICMP Destination Unreachable message to the source of the packet, with the Code indicating "fragmentation needed and DF set". The ICMP Destination Unreachable message includes the MTU of the outbound interface. The source host can use this information to help determine the maximum MTU of a path to a destination.

RFC 1191 is supported on all interfaces.

Changing the router ID

In most configurations, a Layer 3 Switch has multiple IP addresses, usually configured on different interfaces. As a result, a Layer 3 Switch identity to other devices varies depending on the interface to which the other device is attached. Some routing protocols, including Open Shortest Path First (OSPF) and Border Gateway Protocol version 4 (BGP4), identify a Layer 3 Switch by just one of the IP addresses configured on the Layer 3 Switch, regardless of the interfaces that connect the Layer 3 Switches. This IP address is the router ID.

NOTE

Routing Information Protocol (RIP) does not use the router ID.

NOTE

If you change the router ID, all current BGP4 sessions are cleared.

By default, the router ID on a Layer 3 Switch is one of the following:

- If the router has loopback interfaces, the default router ID is the IP address configured on the lowest numbered loopback interface configured on the Layer 3 Switch. For example, if you configure loopback interfaces 1, 2, and 3 as follows, the default router ID is 9.9.9.9/24:
 - Loopback interface 1, 9.9.9.9/24
 - Loopback interface 2, 4.4.4.4/24
 - Loopback interface 3, 1.1.1.1/24
- If the device does not have any loopback interfaces, the default router ID is the lowest numbered IP interface configured on the device.

If you prefer, you can explicitly set the router ID to any valid IP address. The IP address cannot be in use on another device in the network.

NOTE

Layer 3 Switches use the same router ID for both OSPF and BGP4. If the router is already configured for OSPF, you may want to use the router ID that is already in use on the router rather than set a new one. To display the router ID, enter the **show ip CLI** command at any CLI level.

To change the router ID, enter a command such as the following.

```
PowerConnect(config)#ip router-id 209.157.22.26
```

Syntax: `ip router-id <ip-addr>`

The `<ip-addr>` can be any valid, unique IP address.

NOTE

You can specify an IP address used for an interface on the Layer 3 Switch, but do not specify an IP address in use by another device.

Specifying a single source interface for Telnet, TACACS/TACACS+, or RADIUS Packets

When the Layer 3 Switch originates a Telnet, TACACS/TACACS+, or RADIUS packet, the source address of the packet is the lowest-numbered IP address on the interface that sends the packet. You can configure the Layer 3 Switch to always use the lowest-numbered IP address on a specific interface as the source addresses for these types of packets. When you configure the Layer 3 Switch to use a single source interface for all Telnet, TACACS/TACACS+, or RADIUS packets, the Layer 3 Switch uses the same IP address as the source for all packets of the specified type, regardless of the ports that actually send the packets.

Identifying a single source IP address for Telnet, TACACS/TACACS+, or RADIUS packets provides the following benefits:

- If your Telnet, TACACS/TACACS+, or RADIUS server is configured to accept packets only from specific IP addresses, you can use this feature to simplify configuration of the server by configuring the device to always send the packets from the same link or source address.
- If you specify a loopback interface as the single source for Telnet, TACACS/TACACS+, or RADIUS packets, servers can receive the packets regardless of the states of individual links. Thus, if a link to the server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS/TACACS+, or RADIUS packets. You can configure a source interface for one or more of these types of packets separately.

To specify an Ethernet or a loopback or virtual interface as the source for all TACACS/TACACS+ packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the port or interface as the source IP address for TACACS/TACACS+ packets originated by the device.

The following sections show the syntax for specifying a single source IP address for Telnet, TACACS/TACACS+, and RADIUS packets.

Telnet packets

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all Telnet packets, enter commands such as the following.

```
PowerConnect(config)#int loopback 2
PowerConnect(config-lbif-2)#ip address 10.0.0.2/24
PowerConnect(config-lbif-2)#exit
PowerConnect(config)#ip telnet source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all Telnet packets from the Layer 3 Switch.

Syntax: `ip telnet source-interface ethernet [<portnum> | loopback <num> | ve <num>]`

The <num> parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the <portnum> is the port number.

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the Layer 3 Switch.

```
PowerConnect(config)#interface ethernet 4
PowerConnect(config-if-4)#ip address 209.157.22.110/24
PowerConnect(config-if-4)#exit
PowerConnect(config)#ip telnet source-interface ethernet 4
```

TACACS/TACACS+ packets

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all TACACS/TACACS+ packets, enter commands such as the following.

```
PowerConnect(config)#int ve 1
PowerConnect(config-vif-1)#ip address 10.0.0.3/24
PowerConnect(config-vif-1)#exit
PowerConnect(config)#ip tacacs source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all TACACS/TACACS+ packets from the Layer 3 Switch.

Syntax: `ip tacacs source-interface ethernet [<portnum> | loopback <num> | ve <num>]`

The <num> parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the <portnum> is the port number.

RADIUS packets

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all RADIUS packets, enter commands such as the following.

```
PowerConnect(config)#int ve 1
PowerConnect(config-vif-1)#ip address 10.0.0.3/24
PowerConnect(config-vif-1)#exit
PowerConnect(config)#ip radius source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all RADIUS packets from the Layer 3 Switch.

Syntax: `ip radius source-interface ethernet<portnum> | loopback <num> | ve <num>`

The *<num>* parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the *<portnum>* is the port number.

Configuring ARP parameters

Address Resolution Protocol (ARP) is a standard IP protocol that enables an IP Layer 3 Switch to obtain the MAC address of another device interface when the Layer 3 Switch knows the IP address of the interface. ARP is enabled by default and cannot be disabled.

NOTE

Layer 2 Switches also support ARP. The description in “[How ARP works](#)” also applies to ARP on Layer 2 Switches. However, the configuration options described later in this section apply only to Layer 3 Switches, not to Layer 2 Switches.

How ARP works

A Layer 3 Switch needs to know a destination MAC address when forwarding traffic, because the Layer 3 Switch encapsulates the IP packet in a Layer 2 packet (MAC layer packet) and sends the Layer 2 packet to a MAC interface on a device directly attached to the Layer 3 Switch. The device can be the packet final destination or the next-hop router toward the destination.

The Layer 3 Switch encapsulates IP packets in Layer 2 packets regardless of whether the ultimate destination is locally attached or is multiple router hops away. Since the Layer 3 Switch IP route table and IP forwarding cache contain IP address information but not MAC address information, the Layer 3 Switch cannot forward IP packets based solely on the information in the route table or forwarding cache. The Layer 3 Switch needs to know the MAC address that corresponds with the IP address of either the packet locally attached destination or the next-hop router that leads to the destination.

For example, to forward a packet whose destination is multiple router hops away, the Layer 3 Switch must send the packet to the next-hop router toward its destination, or to a default route or default network route if the IP route table does not contain a route to the packet destination. In each case, the Layer 3 Switch must encapsulate the packet and address it to the MAC address of a locally attached device, the next-hop router toward the IP packet destination.

To obtain the MAC address required for forwarding a datagram, the Layer 3 Switch does the following:

- First, the Layer 3 Switch looks in the ARP cache (not the static ARP table) for an entry that lists the MAC address for the IP address. The ARP cache maps IP addresses to MAC addresses. The cache also lists the port attached to the device and, if the entry is dynamic, the age of the entry. A dynamic ARP entry enters the cache when the Layer 3 Switch receives an ARP reply or receives an ARP request (which contains the sender IP address and MAC address). A static entry enters the ARP cache from the static ARP table (which is a separate table) when the interface for the entry comes up.

To ensure the accuracy of the ARP cache, each dynamic entry has its own age timer. The timer is reset to zero each time the Layer 3 Switch receives an ARP reply or ARP request containing the IP address and MAC address of the entry. If a dynamic entry reaches its maximum allowable age, the entry times out and the software removes the entry from the table. Static entries do not age out and can be removed only by you.

- If the ARP cache does not contain an entry for the destination IP address, the Layer 3 Switch broadcasts an ARP request out all its IP interfaces. The ARP request contains the IP address of the destination. If the device with the IP address is directly attached to the Layer 3 Switch, the device sends an ARP response containing its MAC address. The response is a unicast packet addressed directly to the Layer 3 Switch. The Layer 3 Switch places the information from the ARP response into the ARP cache.

ARP requests contain the IP address and MAC address of the sender, so all devices that receive the request learn the MAC address and IP address of the sender and can update their own ARP caches accordingly.

NOTE

The ARP request broadcast is a MAC broadcast, which means the broadcast goes only to devices that are directly attached to the Layer 3 Switch. A MAC broadcast is not routed to other networks. However, some routers, including Layer 3 Switches, can be configured to reply to ARP requests from one network on behalf of devices on another network. Refer to [“Enabling proxy ARP”](#) on page 589.

NOTE

If the router receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (the Layer 3 Switch knows of no route to the destination address), the router sends an ICMP Host Unreachable message to the source.

Rate limiting ARP packets

You can limit the number of ARP packets the device accepts during each second. By default, the software does not limit the number of ARP packets the device can receive. Since the device sends ARP packets to the CPU for processing, if a device in a busy network receives a high number of ARP packets in a short period of time, some CPU processing might be deferred while the CPU processes the ARP packets.

To prevent the CPU from becoming flooded by ARP packets in a busy network, you can restrict the number of ARP packets the device will accept each second. When you configure an ARP rate limit, the device accepts up to the maximum number of packets you specify, but drops additional ARP packets received during the one-second interval. When a new one-second interval starts, the counter restarts at zero, so the device again accepts up to the maximum number of ARP packets you specified, but drops additional packets received within the interval.

To limit the number of ARP packets the device will accept each second, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#rate-limit-arp 100
```

This command configures the device to accept up to 100 ARP packets each second. If the device receives more than 100 ARP packets during a one-second interval, the device drops the additional ARP packets during the remainder of that one-second interval.

Syntax: `[no] rate-limit-arp <num>`

The `<num>` parameter specifies the number of ARP packets and can be from 0 – 100. If you specify 0, the device will not accept any ARP packets.

NOTE

If you want to change a previously configured the ARP rate limiting policy, you must remove the previously configured policy using the **no rate-limit-arp** *<num>* command before entering the new policy.

Changing the ARP aging period

When the Layer 3 Switch places an entry in the ARP cache, the Layer 3 Switch also starts an aging timer for the entry. The aging timer ensures that the ARP cache does not retain learned entries that are no longer valid. An entry can become invalid when the device with the MAC address of the entry is no longer on the network.

The ARP age affects dynamic (learned) entries only, not static entries. The default ARP age is ten minutes. On Layer 3 Switches, you can change the ARP age to a value from 0 – 240 minutes. You cannot change the ARP age on Layer 2 Switches. If you set the ARP age to zero, aging is disabled and entries do not age out.

On PowerConnect devices, ARP aging period is performed on route entry.

To globally change the ARP aging parameter to 20 minutes, enter the following command.

```
PowerConnect(config)#ip arp-age 20
```

Syntax: **ip arp-age** *<num>*

The *<num>* parameter specifies the number of minutes and can be from 0 – 240. The default is 10. If you specify 0, aging is disabled.

To override the globally configured IP ARP age on an individual interface, enter a command such as the following at the interface configuration level.

```
PowerConnect(config-if-e10000-1)#ip arp-age 30
```

Syntax: **[no] ip arp-age** *<num>*

The *<num>* parameter specifies the number of minutes and can be from 0 – 240. The default is the globally configured value, which is 10 minutes by default. If you specify 0, aging is disabled.

Enabling proxy ARP

Proxy ARP allows a Layer 3 Switch to answer ARP requests from devices on one network on behalf of devices in another network. Since ARP requests are MAC-layer broadcasts, they reach only the devices that are directly connected to the sender of the ARP request. Thus, ARP requests do not cross routers.

For example, if Proxy ARP is enabled on a Layer 3 Switch connected to two subnets, 10.10.10.0/24 and 20.20.20.0/24, the Layer 3 Switch can respond to an ARP request from 10.10.10.69 for the MAC address of the device with IP address 20.20.20.69. In standard ARP, a request from a device in the 10.10.10.0/24 subnet cannot reach a device in the 20.20.20.0 subnet if the subnets are on different network cables, and thus is not answered.

NOTE

An ARP request from one subnet can reach another subnet when both subnets are on the same physical segment (Ethernet cable), since MAC-layer broadcasts reach all the devices on the segment.

Proxy ARP is disabled by default on Layer 3 Switches. This feature is not supported on Layer 2 Switches.

You can enable proxy ARP at the Interface level, as well as at the Global CONFIG level, of the CLI.

NOTE

Configuring proxy ARP at the Interface level overrides the global configuration.

Enabling proxy ARP globally

To enable IP proxy ARP on a global basis, enter the following command.

```
PowerConnect(config)#ip proxy-arp
```

To again disable IP proxy ARP on a global basis, enter the following command.

```
PowerConnect(config)#no ip proxy-arp
```

Syntax: [no] ip proxy-arp

- Layer 3 devices only – BL3 and L3

Enabling local proxy ARP

Devices support Proxy Address Resolution Protocol (**Proxy ARP**), a feature that enables router ports to respond to ARP requests for subnets it can reach. However, router ports will not respond to ARP requests for IP addresses in the same subnet as the incoming ports. Resolves this issue with the introduction of Local Proxy ARP per IP interface. **Local Proxy ARP** enables router ports to reply to ARP requests for IP addresses within the same subnet and to forward all traffic between hosts in the subnet.

When Local Proxy ARP is enabled on a router port, the port will respond to ARP requests for IP addresses within the same subnet, if it has ARP entries for the destination IP addresses in the ARP cache. If it does not have ARP entries for the IP addresses, the port will attempt to resolve them by broadcasting its own ARP requests.

Local Proxy ARP is disabled by default. To use Local Proxy ARP, Proxy ARP (CLI command **ip proxy-arp**) must be enabled globally on the device. You can enter the CLI command to enable Local Proxy ARP even though Proxy ARP is not enabled, however, the configuration will not take effect until you enable Proxy ARP.

Use the **show run** command to view the ports on which Local Proxy ARP is enabled.

To enable Local Proxy ARP, enter commands such as the following.

```
PowerConnect(config)#int e 4  
PowerConnect(config-if-e10000-4)#ip local-proxy-arp
```

Syntax: [no] ip local-proxy-arp

Use the **no** form of the command to disable Local Proxy ARP.

Creating static ARP entries

Layer 3 Switches have a static ARP table, in addition to the regular ARP cache. The static ARP table contains entries that you configure.

Static entries are useful in cases where you want to pre-configure an entry for a device that is not connected to the Layer 3 Switch, or you want to prevent a particular entry from aging out. The software removes a dynamic entry from the ARP cache if the ARP aging interval expires before the entry is refreshed. Static entries do not age out, regardless of whether the device receives an ARP request from the device that has the entry address.

NOTE

You cannot create static ARP entries on a Layer 2 Switch.

The maximum number of static ARP entries you can configure depends on the software version running on the device. Refer to [“Changing the maximum number of entries the static ARP table can hold”](#) on page 591.

To display the ARP cache and static ARP table, refer to the following:

- To display the ARP table, refer to [“Displaying the ARP cache”](#) on page 628.
- To display the static ARP table, refer to [“Displaying the static ARP table”](#) on page 629.

To create a static ARP entry, enter a command such as the following.

```
PowerConnect(config)#arp 1 192.53.4.2 1245.7654.2348 e 2
```

Syntax: arp <num> <ip-addr> <mac-addr> ethernet <portnum>

The <num> parameter specifies the entry number. You can specify a number from 1 up to the maximum number of static entries allowed on the device.

The <ip-addr> command specifies the IP address of the device that has the MAC address of the entry.

The <mac-addr> parameter specifies the MAC address of the entry.

The <portnum> command specifies the port number attached to the device that has the MAC address of the entry.

Changing the maximum number of entries the static ARP table can hold

NOTE

The basic procedure for changing the static ARP table size is the same as the procedure for changing other configurable cache or table sizes. Refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184

To increase the maximum number of static ARP table entries you can configure on a Layer 3 Switch, enter commands such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#system-max ip-static-arp 1000
PowerConnect(config)#write memory
PowerConnect(config)#end
PowerConnect#reload
```

NOTE

You must save the configuration to the startup-config file and reload the software after changing the static ARP table size to place the change into effect.

Syntax: system-max ip-static-arp <num>

Configuring forwarding parameters

The following configurable parameters control the forwarding behavior of Layer 3 Switches:

- Time-To-Live (TTL) threshold
- Forwarding of directed broadcasts
- Forwarding of source-routed packets
- Ones-based and zero-based broadcasts

All these parameters are global and thus affect all IP interfaces configured on the Layer 3 Switch.

To configure these parameters, use the procedures in the following sections.

Changing the TTL threshold

The TTL threshold prevents routing loops by specifying the maximum number of router hops an IP packet originated by the Layer 3 Switch can travel through. Each device capable of forwarding IP that receives the packet decrements (decreases) the packet TTL by one. If a device receives a packet with a TTL of 1 and reduces the TTL to zero, the device drops the packet.

The default TTL is 64. You can change the TTL to a value from 1– 255.

To modify the TTL threshold to 25, enter the following commands.

```
PowerConnect(config)#ip ttl 25
```

Syntax: ip ttl <1-255>

Enabling forwarding of directed broadcasts

A directed broadcast is an IP broadcast to all devices within a single directly-attached network or subnet. A net-directed broadcast goes to all devices on a given network. A subnet-directed broadcast goes to all devices within a given subnet.

NOTE

A less common type, the all-subnets broadcast, goes to all directly-attached subnets. Forwarding for this broadcast type also is supported, but most networks use IP multicasting instead of all-subnet broadcasting.

Forwarding for all types of IP directed broadcasts is disabled by default. You can enable forwarding for all types if needed. You cannot enable forwarding for specific broadcast types.

To enable forwarding of IP directed broadcasts, enter the following command.

```
PowerConnect(config)#ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

NOTE

This command is not supported on the PowerConnect B-Series TI24X devices.

Device software makes the forwarding decision based on the router's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last hop router.

To disable the directed broadcasts, enter the following command in the CONFIG mode.

```
PowerConnect(config)#no ip directed-broadcast
```

To enable directed broadcasts on an individual interface instead of globally for all interfaces, enter commands such as the following.

```
PowerConnect(config)#interface ethernet 1
PowerConnect(config-if-1)#ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

Disabling forwarding of IP source-routed packets

A source-routed packet specifies the exact router path for the packet. The packet specifies the path by listing the IP addresses of the router interfaces through which the packet must pass on its way to the destination. The Layer 3 Switch supports both types of IP source routing:

- Strict source routing – requires the packet to pass through only the listed routers. If the Layer 3 Switch receives a strict source-routed packet but cannot reach the next hop interface specified by the packet, the Layer 3 Switch discards the packet and sends an ICMP Source-Route-Failure message to the sender.

NOTE

The Layer 3 Switch allows you to disable sending of the Source-Route-Failure messages. Refer to [“Disabling ICMP messages”](#) on page 594.

- Loose source routing – requires that the packet pass through all of the listed routers but also allows the packet to travel through other routers, which are not listed in the packet.

The Layer 3 Switch forwards both types of source-routed packets by default. To disable the feature, use either of the following methods. You cannot enable or disable strict or loose source routing separately.

To disable forwarding of IP source-routed packets, enter the following command.

```
PowerConnect(config)#no ip source-route
```

Syntax: [no] ip source-route

To re-enable forwarding of source-routed packets, enter the following command.

```
PowerConnect(config)#ip source-route
```

Enabling support for zero-based IP subnet broadcasts

By default, the Layer 3 Switch treats IP packets with all ones in the host portion of the address as IP broadcast packets. For example, the Layer 3 Switch treats IP packets with 209.157.22.255/24 as the destination IP address as IP broadcast packets and forwards the packets to all IP hosts within the 209.157.22.x subnet (except the host that sent the broadcast packet to the Layer 3 Switch).

Most IP hosts are configured to receive IP subnet broadcast packets with all ones in the host portion of the address. However, some older IP hosts instead expect IP subnet broadcast packets that have all zeros instead of all ones in the host portion of the address. To accommodate this type of host, you can enable the Layer 3 Switch to treat IP packets with all zeros in the host portion of the destination IP address as broadcast packets.

NOTE

When you enable the Layer 3 Switch for zero-based subnet broadcasts, the Layer 3 Switch still treats IP packets with all ones the host portion as IP subnet broadcasts too. Thus, the Layer 3 Switch can be configured to support all ones only (the default) or all ones **and** all zeroes.

NOTE

This feature applies only to IP subnet broadcasts, not to local network broadcasts. The local network broadcast address is still expected to be all ones.

To enable the Layer 3 Switch for zero-based IP subnet broadcasts in addition to ones-based IP subnet broadcasts, enter the following command.

```
PowerConnect(config)#ip broadcast-zero
PowerConnect(config)#write memory
PowerConnect(config)#end
PowerConnect#reload
```

NOTE

You must save the configuration and reload the software to place this configuration change into effect.

Syntax: [no] ip broadcast-zero

NOTE

This command is not supported on the PowerConnect B-Series TI24X devices.

Disabling ICMP messages

Devices are enabled to reply to ICMP echo messages and send ICMP Destination Unreachable messages by default.

You can selectively disable the following types of Internet Control Message Protocol (ICMP) messages:

- Echo messages (ping messages) – The Layer 3 Switch replies to IP pings from other IP devices.
- Destination Unreachable messages – If the Layer 3 Switch receives an IP packet that it cannot deliver to its destination, the Layer 3 Switch discards the packet and sends a message back to the device that sent the packet to the Layer 3 Switch. The message informs the device that the destination cannot be reached by the Layer 3 Switch.

Disabling replies to broadcast ping requests

By default, devices are enabled to respond to broadcast ICMP echo packets, which are ping requests.

To disable response to broadcast ICMP echo packets (ping requests), enter the following command.

```
PowerConnect(config)#no ip icmp echo broadcast-request
```

Syntax: [no] ip icmp echo broadcast-request

If you need to re-enable response to ping requests, enter the following command.

```
PowerConnect(config)#ip icmp echo broadcast-request
```


Disabling ICMP destination unreachable messages

By default, when a device receives an IP packet that the device cannot deliver, the device sends an ICMP Unreachable message back to the host that sent the packet. You can selectively disable a device response to the following types of ICMP Unreachable messages:

- **Administration** – The packet was dropped by the device due to a filter or ACL configured on the device.
- **Fragmentation-needed** – The packet has the Do not Fragment bit set in the IP Flag field, but the device cannot forward the packet without fragmenting it.
- **Host** – The destination network or subnet of the packet is directly connected to the device, but the host specified in the destination IP address of the packet is not on the network.
- **Port** – The destination host does not have the destination TCP or UDP port specified in the packet. In this case, the host sends the ICMP Port Unreachable message to the device, which in turn sends the message to the host that sent the packet.
- **Protocol** – The TCP or UDP protocol on the destination host is not running. This message is different from the Port Unreachable message, which indicates that the protocol is running on the host but the requested protocol port is unavailable.
- **Source-route-failure** – The device received a source-routed packet but cannot locate the next-hop IP address indicated in the packet Source-Route option.

You can disable the device from sending these types of ICMP messages on an individual basis. To do so, use the following CLI method.

NOTE

Disabling an ICMP Unreachable message type does not change the device ability to forward packets. Disabling ICMP Unreachable messages prevents the device from generating or forwarding the Unreachable messages.

To disable all ICMP Unreachable messages, enter the following command.

```
PowerConnect(config)#no ip icmp unreachable
```

Syntax: [no] ip icmp unreachable [host | protocol | administration | fragmentation-needed | port | source-route-fail]

- If you enter the command without specifying a message type (as in the example above), all types of ICMP Unreachable messages listed above are disabled. If you want to disable only specific types of ICMP Unreachable messages, you can specify the message type. To disable more than one type of ICMP message, enter the **no ip icmp unreachable** command for each messages type.
- The **administration** parameter disables ICMP Unreachable (caused by Administration action) messages.
- The **fragmentation-needed** parameter disables ICMP Fragmentation-Needed But Do not-Fragment Bit Set messages.
- The **host** parameter disables ICMP Host Unreachable messages.
- The **port** parameter disables ICMP Port Unreachable messages.
- The **protocol** parameter disables ICMP Protocol Unreachable messages.
- The **source-route-fail** parameter disables ICMP Unreachable (caused by Source-Route-Failure) messages.

To disable ICMP Host Unreachable messages but leave the other types of ICMP Unreachable messages enabled, enter the following commands instead of the command shown above.

```
PowerConnect(config)#no ip icmp unreachable host
```

If you have disabled all ICMP Unreachable message types but you want to re-enable certain types, for example ICMP Host Unreachable messages, you can do so by entering the following command.

```
PowerConnect(config)#ip icmp unreachable host
```

Configuring static routes

The IP route table can receive routes from the following sources:

- Directly-connected networks – When you add an IP interface, the Layer 3 Switch automatically creates a route for the network the interface is in.
- RIP – If RIP is enabled, the Layer 3 Switch can learn about routes from the advertisements other RIP routers send to the Layer 3 Switch. If the route has a lower administrative distance than any other routes from different sources to the same destination, the Layer 3 Switch places the route in the IP route table.
- OSPF – Refer to RIP, but substitute “OSPF” for “RIP”.
- Default network route – A statically configured default route that the Layer 3 Switch uses if other default routes to the destination are not available. Refer to [“Configuring a default network route”](#) on page 604.
- Statically configured route – You can add routes directly to the route table. When you add a route to the IP route table, you are creating a static IP route. This section describes how to add static routes to the IP route table.

NOTE

On PowerConnect B-Series TI24X devices, this feature is supported.

Static route types

You can configure the following types of static IP routes:

- Standard – the static route consists of the destination network address and network mask, and the IP address of the next-hop gateway. You can configure multiple standard static routes with the same metric for load sharing or with different metrics to provide a primary route and backup routes.
- Interface-based – the static route consists of the destination network address and network mask, and the Layer 3 Switch interface through which you want the Layer 3 Switch to send traffic for the route. Typically, this type of static route is for directly attached destination networks.
- Null – the static route consists of the destination network address and network mask, and the “null0” parameter. Typically, the null route is configured as a backup route for discarding traffic if the primary route is unavailable.

Static IP route parameters

When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route destination network.

- The route path, which can be one of the following:
 - The IP address of a next-hop gateway
 - An Ethernet port
 - A virtual interface (a routing interface used by VLANs for routing Layer 3 protocol traffic among one another)
 - A “null” interface. The Layer 3 Switch drops traffic forwarded to the null interface.

You also can specify the following optional parameters:

- The metric for the route – The value the Layer 3 Switch uses when comparing this route to other routes in the IP route table to the same destination. The metric applies only to routes that the Layer 3 Switch has already placed in the IP route table. The default metric for static IP routes is 1.
- The administrative distance for the route – The value that the Layer 3 Switch uses to compare this route with routes from other route sources to the same destination before placing a route in the IP route table. This parameter does not apply to routes that are already in the IP route table. The default administrative distance for static IP routes is 1.

The default metric and administrative distance values ensure that the Layer 3 Switch always prefers static IP routes over routes from other sources to the same destination.

Multiple static routes to the same destination provide load sharing and redundancy

You can add multiple static routes for the same destination network to provide one or more of the following benefits:

- IP load balancing – When you add multiple IP static routes for the same destination to different next-hop gateways, and the routes each have the same metric and administrative distance, the Layer 3 Switch can load balance traffic to the routes’ destination. For information about IP load balancing, refer to [“Configuring IP load sharing”](#) on page 605.
- Path redundancy – When you add multiple static IP routes for the same destination, but give the routes different metrics or administrative distances, the Layer 3 Switch uses the route with the lowest administrative distance by default, but uses another route to the same destination if the first route becomes unavailable.

Refer to the following sections for examples and configuration information:

- [“Configuring load balancing and redundancy using multiple static routes to the same destination”](#) on page 600
- [“Configuring standard static IP routes and interface or null static routes to the same destination”](#) on page 601


Static route states follow port states

IP static routes remain in the IP route table only so long as the port or virtual interface used by the route is available. If the port or virtual routing interface becomes unavailable, the software removes the static route from the IP route table. If the port or virtual routing interface becomes available again later, the software adds the route back to the route table.

This feature allows the Layer 3 Switch to adjust to changes in network topology. The Layer 3 Switch does not continue trying to use routes on unavailable paths but instead uses routes only when their paths are available.

Figure 98 shows an example of a network containing a static route. The static route is configured on Switch A, as shown in the CLI example following the figure.

FIGURE 98 Example of a static route



The following command configures a static route to 207.95.7.0, using 207.95.6.157 as the next-hop gateway.

```
PowerConnect(config)# ip route 207.95.7.0/24 207.95.6.157
```

When you configure a static IP route, you specify the destination address for the route and the next-hop gateway or Layer 3 Switch interface through which the Layer 3 Switch can reach the route. The Layer 3 Switch adds the route to the IP route table. In this case, Switch A knows that 207.95.6.157 is reachable through port 2, and also assumes that local interfaces within that subnet are on the same port. Switch A deduces that IP interface 207.95.7.188 is also on port 2.

The software automatically removes a static IP route from the IP route table if the port used by that route becomes unavailable. When the port becomes available again, the software automatically re-adds the route to the IP route table.

Configuring a static IP route

To configure an IP static route with a destination address of 192.0.0.0 255.0.0.0 and a next-hop router IP address of 195.1.1.1, enter the following commands.

```
PowerConnect(config)# ip route 192.0.0.0 255.0.0.0 195.1.1.1
```

To configure a static IP route with an Ethernet port instead of a next-hop address, enter a command such as the following.

```
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 ethernet 1
```

The command in the example above configures a static IP route for destination network 192.128.2.69/24. Since an Ethernet port is specified instead of a gateway IP address as the next hop, the Layer 3 Switch always forwards traffic for the 192.128.2.69/24 network to port 1. The command in the following example configures an IP static route that uses virtual interface 3 as its next hop.

```
PowerConnect(config)# ip route 192.128.2.71 255.255.255.0 ve 3
```

The command in the following example configures an IP static route that uses port 2 as its next hop.

```
PowerConnect(config)# ip route 192.128.2.73 255.255.255.0 ethernet 2
```

Syntax: `ip route <dest-ip-addr> <dest-mask>`
`<next-hop-ip-addr> |`
`ethernet <portnum> | ve <num>`
`[<metric>] [distance <num>]`

or

Syntax: `ip route <dest-ip-addr>/<mask-bits>`
`<next-hop-ip-addr> |`
`ethernet <portnum> | ve <num>`
`[<metric>] [distance <num>]`

The `<dest-ip-addr>` is the route destination. The `<dest-mask>` is the network mask for the route destination IP address. Alternatively, you can specify the network mask information by entering a forward slash followed by the number of bits in the network mask. For example, you can enter 192.0.0.0 255.255.255.0 as 192.0.0.0/.24.

The `<next-hop-ip-addr>` is the IP address of the next-hop router (gateway) for the route.

If you do not want to specify a next-hop IP address, you can instead specify a port or interface number on the Layer 3 Switch. The `<num>` parameter is a virtual interface number. If you instead specify an Ethernet port, the `<portnum>` is the port number. In this case, the Layer 3 Switch forwards packets destined for the static route destination network to the specified interface. Conceptually, this feature makes the destination network like a directly connected network, associated with a specific Layer 3 Switch interface.

NOTE

The port or virtual interface you use for the static route next hop must have at least one IP address configured on it. The address does not need to be in the same subnet as the destination network.

The `<metric>` parameter can be a number from 1 – 16. The default is 1.

NOTE

If you specify 16, RIP considers the metric to be infinite and thus also considers the route to be unreachable.

The **distance** `<num>` parameter specifies the administrative distance of the route. When comparing otherwise equal routes to a destination, the Layer 3 Switch prefers lower administrative distances over higher ones, so make sure you use a low value for your default route. The default is 1.

NOTE

The Layer 3 Switch will replace the static route if the it receives a route with a lower administrative distance. Refer to [“Changing administrative distances”](#) on page 777 for a list of the default administrative distances for all types of routes.

NOTE

You can also assign the default router as the destination by entering 0.0.0.0 0.0.0.0 xxx.xxx.xxx.xxx.

Configuring a “Null” route

You can configure the Layer 3 Switch to drop IP packets to a specific network or host address by configuring a “null” (sometimes called “null0”) static route for the address. When the Layer 3 Switch receives a packet destined for the address, the Layer 3 Switch drops the packet instead of forwarding it.

To configure a null static route, use the following CLI method.

To configure a null static route to drop packets destined for network 209.157.22.x, enter the following commands.

```
PowerConnect(config)# ip route 209.157.22.0 255.255.255.0 null0
PowerConnect(config)# write memory
```

Syntax: `ip route <ip-addr> <ip-mask> null0 [<metric>] [distance <num>]`

or

Syntax: `ip route <ip-addr>/<mask-bits> null0 [<metric>] [distance <num>]`

To display the maximum value for your device, enter the **show default values** command. The maximum number of static IP routes the system can hold is listed in the ip-static-route row in the System Parameters section of the display. To change the maximum value, use the **system-max ip-static-route <num>** command at the global CONFIG level.

The `<ip-addr>` parameter specifies the network or host address. The Layer 3 Switch will drop packets that contain this address in the destination field instead of forwarding them.

The `<ip-mask>` parameter specifies the network mask. Ones are significant bits and zeros allow any value. For example, the mask 255.255.255.0 matches on all hosts within the Class C subnet address specified by `<ip-addr>`. Alternatively, you can specify the number of bits in the network mask. For example, you can enter 209.157.22.0/24 instead of 209.157.22.0 255.255.255.0.

The **null0** parameter indicates that this is a null route. You must specify this parameter to make this a null route.

The `<metric>` parameter adds a cost to the route. You can specify from 1 – 16. The default is 1.

The distance `<num>` parameter configures the administrative distance for the route. You can specify a value from 1 – 255. The default is 1. The value 255 makes the route unusable.

NOTE

The last two parameters are optional and do not affect the null route, unless you configure the administrative distance to be 255. In this case, the route is not used and the traffic might be forwarded instead of dropped.

Configuring load balancing and redundancy using multiple static routes to the same destination

You can configure multiple static IP routes to the same destination, for the following benefits:

- **IP load sharing** – If you configure more than one static route to the same destination, and the routes have different next-hop gateways but have the same metrics, the Layer 3 Switch load balances among the routes using basic round-robin. For example, if you configure two static routes with the same metrics but to different gateways, the Layer 3 Switch alternates between the two routes. For information about IP load balancing, refer to [“Configuring IP load sharing”](#) on page 605.
- **Backup Routes** – If you configure multiple static IP routes to the same destination, but give the routes different next-hop gateways and different metrics, the Layer 3 Switch will always use the route with the lowest metric. If this route becomes unavailable, the Layer 3 Switch will fail over to the static route with the next-lowest metric, and so on.

NOTE

You also can bias the Layer 3 Switch to select one of the routes by configuring them with different administrative distances. However, make sure you do not give a static route a higher administrative distance than other types of routes, unless you want those other types to be preferred over the static route. For a list of the default administrative distances, refer to [“Changing administrative distances”](#) on page 777.

The steps for configuring the static routes are the same as described in the previous section. The following sections provide examples.

To configure multiple static IP routes, enter commands such as the following.

```
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 209.157.22.1
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 192.111.10.1
```

The commands in the example above configure two static IP routes. The routes go to different next-hop gateways but have the same metrics. These commands use the default metric value (1), so the metric is not specified. These static routes are used for load sharing among the next-hop gateways.

The following commands configure static IP routes to the same destination, but with different metrics. The route with the lowest metric is used by default. The other routes are backups in case the first route becomes unavailable. The Layer 3 Switch uses the route with the lowest metric if the route is available.

```
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 209.157.22.1
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 192.111.10.1 2
PowerConnect(config)# ip route 192.128.2.69 255.255.255.0 201.1.1.1 3
```

In this example, each static route has a different metric. The metric is not specified for the first route, so the default (1) is used. A metric is specified for the second and third static IP routes. The second route has a metric of two and the third route has a metric of 3. Thus, the second route is used only if the first route (which has a metric of 1) becomes unavailable. Likewise, the third route is used only if the first and second routes (which have lower metrics) are both unavailable.

For complete syntax information, refer to [“Configuring a static IP route”](#) on page 598.

Configuring standard static IP routes and interface or null static routes to the same destination

You can configure a null0 or interface-based static route to a destination and also configure a normal static route to the same destination, so long as the route metrics are different.

When the Layer 3 Switch has multiple routes to the same destination, the Layer 3 Switch always prefers the route with the lowest metric. Generally, when you configure a static route to a destination network, you assign the route a low metric so that the Layer 3 Switch prefers the static route over other routes to the destination.

This feature is especially useful for the following configurations. These are not the only allowed configurations but they are typical uses of this enhancement:

- When you want to ensure that if a given destination network is unavailable, the Layer 3 Switch drops (forwards to the null interface) traffic for that network instead of using alternate paths to route the traffic. In this case, assign the normal static route to the destination network a lower metric than the null route.

- When you want to use a specific interface by default to route traffic to a given destination network, but want to allow the Layer 3 Switch to use other interfaces to reach the destination network if the path that uses the default interface becomes unavailable. In this case, give the interface route a lower metric than the normal static route.

NOTE

You cannot add a null or interface-based static route to a network if there is already a static route of any type with the same metric you specify for the null or interface-based route.

Figure 99 shows an example of two static routes configured for the same destination network. In this example, one of the routes is a standard static route and has a metric of 1. The other static route is a null route and has a higher metric than the standard static route. The Layer 3 Switch always prefers the static route with the lower metric. In this example, the Layer 3 Switch always uses the standard static route for traffic to destination network 192.168.7.0/24, unless that route becomes unavailable, in which case the Layer 3 Switch sends traffic to the null route instead.

FIGURE 99 Standard and null static routes to the same destination network

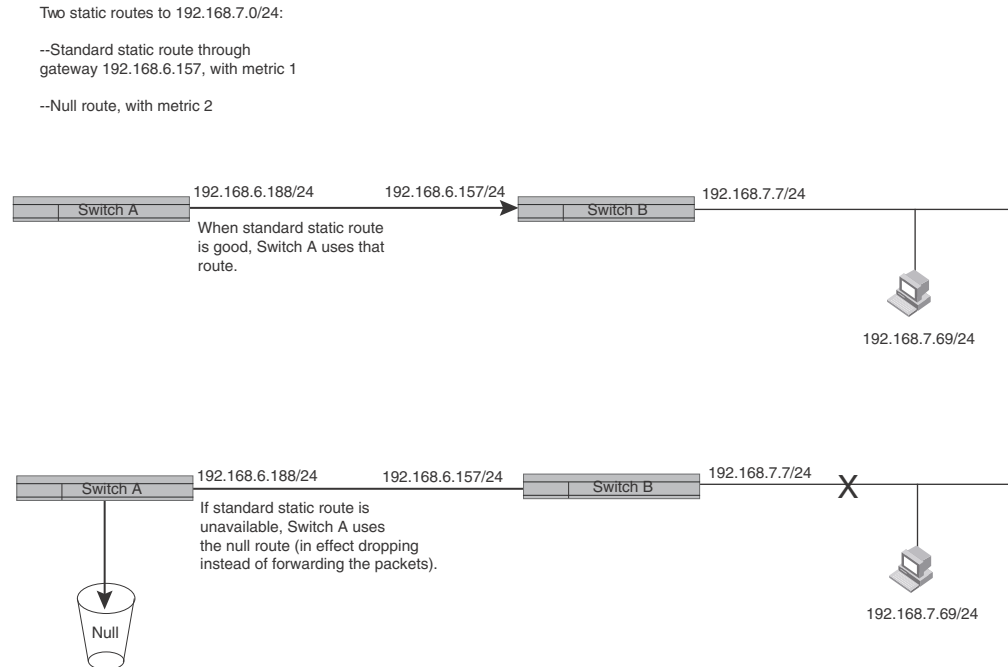


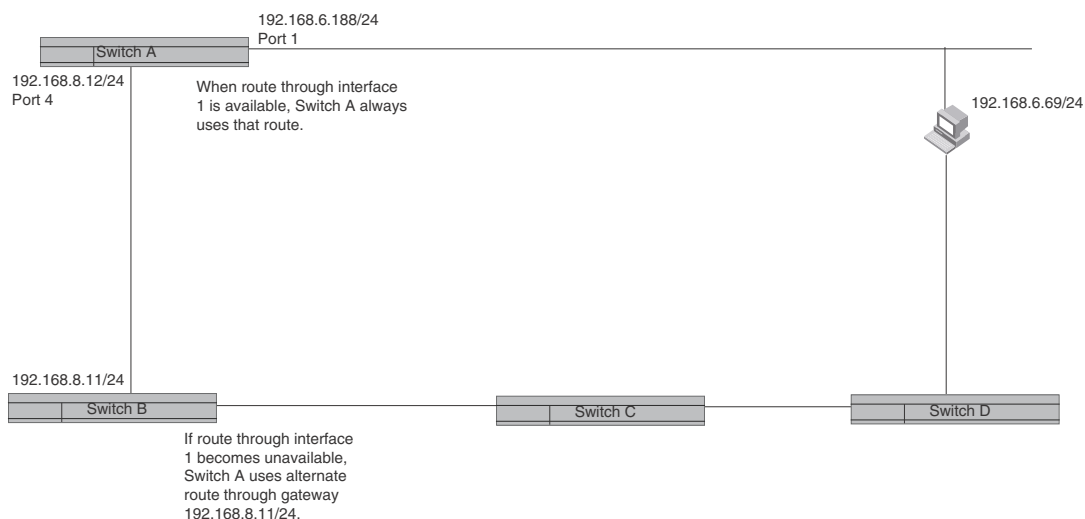
Figure 100 shows another example of two static routes. In this example, a standard static route and an interface-based static route are configured for destination network 192.168.6.0/24. The interface-based static route has a lower metric than the standard static route. As a result, the Layer 3 Switch always prefers the interface-based route when the route is available. However, if the interface-based route becomes unavailable, the Layer 3 Switch still forwards the traffic toward the destination using an alternate route through gateway 192.168.8.11/24.

FIGURE 100 Standard and interface routes to the same destination network

Two static routes to 192.168.7.0/24:

--Interface-based route through Port 1, with metric 1.

--Standard static route through gateway 192.168.8.11, with metric 3.



To configure a standard static IP route and a null route to the same network as shown in [Figure 99](#) on page 602, enter commands such as the following.

```
PowerConnect(config)# ip route 192.168.7.0/24 192.168.6.157/24 1
PowerConnect(config)# ip route 192.168.7.0/24 null0 3
```

The first command configures a standard static route, which includes specification of the next-hop gateway. The command also gives the standard static route a metric of 1, which causes the Layer 3 Switch to always prefer this route when the route is available.

The second command configures another static route for the same destination network, but the second route is a null route. The metric for the null route is 3, which is higher than the metric for the standard static route. If the standard static route is unavailable, the software uses the null route.

For complete syntax information, refer to [“Configuring a static IP route”](#) on page 598.

To configure a standard static route and an interface-based route to the same destination, enter commands such as the following.

```
PowerConnect(config)# ip route 192.168.6.0/24 ethernet 1 1
PowerConnect(config)# ip route 192.168.6.0/24 192.168.8.11/24 3
```

The first command configured an interface-based static route through Ethernet port 1. The command assigns a metric of 1 to this route, causing the Layer 3 Switch to always prefer this route when it is available. If the route becomes unavailable, the Layer 3 Switch uses an alternate route through the next-hop gateway 192.168.8.11/24.

Configuring a default network route

The Layer 3 Switch enables you to specify a candidate default route without the need to specify the next hop gateway. If the IP route table does not contain an explicit default route (for example, 0.0.0.0/0) or propagate an explicit default route through routing protocols, the software can use the default network route as a default route instead.

When the software uses the default network route, it also uses the default network route's next hop gateway as the gateway of last resort.

This feature is especially useful in environments where network topology changes can make the next hop gateway unreachable. This feature allows the Layer 3 Switch to perform default routing even if the default network route's default gateway changes.

The feature thus differs from standard default routes. When you configure a standard default route, you also specify the next hop gateway. If a topology change makes the gateway unreachable, the default route becomes unusable.

For example, if you configure 10.10.10.0/24 as a candidate default network route, if the IP route table does not contain an explicit default route (0.0.0.0/0), the software uses the default network route and automatically uses that route's next hop gateway as the default gateway. If a topology change occurs and as a result the default network route's next hop gateway changes, the software can still use the default network route. To configure a default network route, use the following CLI method.

If you configure more than one default network route, the Layer 3 Switch uses the following algorithm to select one of the routes.

1. Use the route with the lowest administrative distance.
2. If the administrative distances are equal:
 - Are the routes from different routing protocols (RIP, OSPF, or BGP4)? If so, use the route with the lowest IP address.
 - If the routes are from the same routing protocol, use the route with the best metric. The meaning of “best” metric depends on the routing protocol:
 - RIP – The metric is the number of hops (additional routers) to the destination. The best route is the route with the fewest hops.
 - OSPF – The metric is the path cost associated with the route. The path cost does not indicate the number of hops but is instead a numeric value associated with each route. The best route is the route with the lowest path cost.
 - BGP4 – The metric is the Multi-exit Discriminator (MED) associated with the route. The MED applies to routes that have multiple paths through the same AS. The best route is the route with the lowest MED.

Configuring a default network route

You can configure up to four default network routes.

To configure a default network route, enter commands such as the following.

```
PowerConnect(config)# ip default-network 209.157.22.0
PowerConnect(config)# write memory
```

Syntax: `ip default-network <ip-addr>`

The `<ip-addr>` parameter specifies the network address.

To verify that the route is in the route table, enter the following command at any level of the CLI.

```
PowerConnect# show ip route
Total number of IP routes: 2
Start index: 1  B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
      Destination      NetMask      Gateway      Port  Cost  Type
1      209.157.20.0      255.255.255.0  0.0.0.0      lb1   1    D
2      209.157.22.0      255.255.255.0  0.0.0.0      11   1    *D
```

This example shows two routes. Both of the routes are directly attached, as indicated in the Type column. However, one of the routes is shown as type “*D”, with an asterisk (*). The asterisk indicates that this route is a candidate default network route.

Configuring IP load sharing

The IP route table can contain more than one path to a given destination. When this occurs, the Layer 3 Switch selects the path with the lowest cost as the path for forwarding traffic to the destination. If the IP route table contains more than one path to a destination and the paths each have the lowest cost, then the Layer 3 Switch uses **IP load sharing** to select a path to the destination.¹You can enable a Layer 3 Switch to load balance across up to six equal-cost paths.

NOTE

IP load sharing is based on next-hop routing, and not on source routing.

NOTE

The term “path” refers to the next-hop router to a destination, not to the entire route to a destination. Thus, when the software compares multiple equal-cost paths, the software is comparing paths that use different next-hop routers, with equal costs, to the same destination.

In many contexts, the terms “route” and “path” mean the same thing. Most of the user documentation uses the term “route” throughout. The term “path” is used in this section to refer to an individual next-hop router to a destination, while the term “route” refers collectively to the multiple paths to the destination. Load sharing applies when the IP route table contains multiple, equal-cost paths to a destination.

NOTE

Devices also perform load sharing among the ports in aggregate links. Refer to [“Trunk group load sharing”](#) on page 314.

How multiple equal-cost paths enter the IP route table

IP load sharing applies to equal-cost paths in the IP route table. Routes that are eligible for load sharing can enter the table from any of the following sources:

- IP static routes
- Routes learned through RIP
- Routes learned through OSPF
- Routes learned through BGP4

1. IP load sharing is also called “Equal-Cost Multi-Path (ECMP)” load sharing or just “ECMP”

Administrative distance

The administrative distance is a unique value associated with each type (source) of IP route. Each path has an administrative distance. The administrative distance is not used when performing IP load sharing, but the administrative distance is used when evaluating multiple equal-cost paths to the same destination from different sources, such as RIP, OSPF and so on.

The value of the administrative distance is determined by the source of the route. The Layer 3 Switch is configured with a unique administrative distance value for each IP route source.

When the software receives multiple paths to the same destination and the paths are from different sources, the software compares the administrative distances of the paths and selects the path with the lowest distance. The software then places the path with the lowest administrative distance in the IP route table. For example, if the Layer 3 Switch has a path learned from OSPF and a path learned from RIP for a given destination, only the path with the lower administrative distance enters the IP route table.

Here are the default administrative distances on the Layer 3 Switch:

- Directly connected – 0 (this value is not configurable)
- Static IP route – 1 (applies to all static routes, including default routes and default network routes)
- Exterior Border Gateway Protocol (EBGP) – 20
- OSPF – 110
- RIP – 120
- Interior Gateway Protocol (IBGP) – 200
- Local BGP – 200
- Unknown – 255 (the router will not use this route)

Lower administrative distances are preferred over higher distances. For example, if the router receives routes for the same network from OSPF and from RIP, the router will prefer the OSPF route by default.

NOTE

You can change the administrative distances individually. Refer to the configuration chapter for the route source for information.

Since the software selects only the path with the lowest administrative distance, and the administrative distance is determined by the path source, IP load sharing does not apply to paths from different route sources. IP load sharing applies only when the IP route table contains multiple paths to the same destination, from the same IP route source.

IP load sharing does not apply to paths that come from different sources.

Path cost

The cost parameter provides a common basis of comparison for selecting from among multiple paths to a given destination. Each path in the IP route table has a cost. When the IP route table contains multiple paths to a destination, the Layer 3 Switch chooses the path with the lowest cost. When the IP route table contains more than one path with the lowest cost to a destination, the Layer 3 Switch uses IP load sharing to select one of the lowest-cost paths.

The source of a path cost value depends on the source of the path:

- IP static route – The value you assign to the metric parameter when you configure the route. The default metric is 1. Refer to “[Configuring load balancing and redundancy using multiple static routes to the same destination](#)” on page 600.
- RIP – The number of next-hop routers to the destination.
- OSPF – The Path Cost associated with the path. The paths can come from any combination of inter-area, intra-area, and external Link State Advertisements (LSAs).

NOTE

If the path is redistributed between two or more of the above sources before entering the IP route table, the cost can increase during the redistribution due to settings in redistribution filters.

Static route, OSPF load sharing

IP load sharing and load sharing for static routes, OSPF routes, routes are individually configured. Multiple equal-cost paths for a destination can enter the IP route table only if the source of the paths is configured to support multiple equal-cost paths.

[Table 100](#) lists the default and configurable maximum numbers of paths for each IP route source that can provide equal-cost paths to the IP route table. The table also lists where to find configuration information for the route source load sharing parameters.

The load sharing state for all the route sources is based on the state of IP load sharing. Since IP load sharing is enabled by default on all Layer 3 Switches, load sharing for static IP routes, RIP routes, OSPF routes, and BGP4 routes also is enabled by default.

TABLE 100 Default load sharing parameters for route sources

Route source	Default maximum number of paths	Maximum Number of paths	See...
PowerConnect B-Series T124X			
Static IP route	4 ¹	8	page 608
RIP	4 ¹	8	page 608
OSPF	4	8	page 608

1. This value depends on the value for IP load sharing, and is not separately configurable.

How IP load sharing works

When the Layer 3 Switch receives traffic for a destination and the IP route table contains multiple, equal-cost paths to that destination, the device checks the IP forwarding cache for a forwarding entry for the destination. The IP forwarding cache provides a fast path for forwarding IP traffic, including load-balanced traffic. The cache contains entries that associate a destination host or network with a path (next-hop router).

- If the IP forwarding sharing cache contains a forwarding entry for the destination, the device uses the entry to forward the traffic.

- If the IP load forwarding cache does not contain a forwarding entry for the destination, the software selects a path from among the available equal-cost paths to the destination, then creates a forwarding entry in the cache based on the calculation. Subsequent traffic for the same destination uses the forwarding entry.

Response to path state changes

If one of the load-balanced paths to a cached destination becomes unavailable, or the IP route table receives a new equal-cost path to a cached destination, the software removes the unavailable path from the IP route table. Then the software selects a new path.

Disabling or re-enabling load sharing

To disable IP load sharing, enter the following commands.

```
PowerConnect(config)# no ip load-sharing
```

Syntax: [no] ip load-sharing

Changing the maximum number of load sharing paths

By default, IP load sharing allows IP traffic to be balanced across up to four equal paths. You can change the maximum number of paths the Layer 3 Switch supports to a value from 2 – 8.

For optimal results, set the maximum number of paths to a value at least as high as the maximum number of equal-cost paths your network typically contains. For example, if the Layer 3 Switch you are configuring for IP load sharing has six next-hop routers, set the maximum paths value to six.

NOTE

If the setting for the maximum number of paths is lower than the actual number of equal-cost paths, the software does not use all the paths for load sharing.

To change the number of IP load sharing paths, enter a command such as the following.

```
PowerConnect(config)# ip load-sharing 6
```

Syntax: [no] ip load-sharing [<num>]

The <num> parameter specifies the number of paths and can be from 2 – 8.

Configuring IRDP

The ICMP Router Discovery Protocol (IRDP) is used by Layer 3 Switches to advertise the IP addresses of its router interfaces to directly attached hosts. IRDP is disabled by default. You can enable the feature on a global basis or on an individual port basis:

- If you enable the feature globally, all ports use the default values for the IRDP parameters.
- If you leave the feature disabled globally but enable it on individual ports, you also can configure the IRDP parameters on an individual port basis.

NOTE

You can configure IRDP parameters only on an individual port basis. To do so, IRDP must be disabled globally and enabled only on individual ports. You cannot configure IRDP parameters if the feature is globally enabled.

When IRDP is enabled, the Layer 3 Switch periodically sends Router Advertisement messages out the IP interfaces on which the feature is enabled. The messages advertise the Layer 3 Switch IP addresses to directly attached hosts who listen for the messages. In addition, hosts can be configured to query the Layer 3 Switch for the information by sending Router Solicitation messages.

Some types of hosts use the Router Solicitation messages to discover their default gateway. When IRDP is enabled on the Layer 3 Switch, the Layer 3 Switch responds to the Router Solicitation messages. Some clients interpret this response to mean that the Layer 3 Switch is the default gateway. If another router is actually the default gateway for these clients, leave IRDP disabled on the Layer 3 Switch.

IRDP uses the following parameters. If you enable IRDP on individual ports instead of enabling the feature globally, you can configure these parameters on an individual port basis:

- **Packet type** – The Layer 3 Switch can send Router Advertisement messages as IP broadcasts or as IP multicasts addressed to IP multicast group 224.0.0.1. The packet type is IP broadcast.
- **Maximum message interval and minimum message interval** – When IRDP is enabled, the Layer 3 Switch sends the Router Advertisement messages every 450 – 600 seconds by default. The time within this interval that the Layer 3 Switch selects is random for each message and is not affected by traffic loads or other network factors. The random interval minimizes the probability that a host will receive Router Advertisement messages from other routers at the same time. The interval on each IRDP-enabled Layer 3 Switch interface is independent of the interval on other IRDP-enabled interfaces. The default maximum message interval is 600 seconds. The default minimum message interval is 450 seconds.
- **Hold time** – Each Router Advertisement message contains a hold time value. This value specifies the maximum amount of time the host should consider an advertisement to be valid until a newer advertisement arrives. When a new advertisement arrives, the hold time is reset. The hold time is always longer than the maximum advertisement interval. Therefore, if the hold time for an advertisement expires, the host can reasonably conclude that the router interface that sent the advertisement is no longer available. The default hold time is three times the maximum message interval.
- **Preference** – If a host receives multiple Router Advertisement messages from different routers, the host selects the router that sent the message with the highest preference as the default gateway. The preference can be a number from 0-4294967296 to 0-4294967295. The default is 0.

Enabling IRDP globally

To globally enable IRDP, enter the following command.

```
PowerConnect(config)# ip irdp
```

This command enables IRDP on the IP interfaces on all ports. Each port uses the default values for the IRDP parameters. The parameters are not configurable when IRDP is globally enabled.

Enabling IRDP on an individual port

To enable IRDP on an individual interface and change IRDP parameters, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 3  
PowerConnect(config-if-3)# ip irdp maxadvertinterval 400
```

This example shows how to enable IRDP on a specific port and change the maximum advertisement interval for Router Advertisement messages to 400 seconds.

NOTE

To enable IRDP on individual ports, you must leave the feature globally disabled.

Syntax: `[no] ip irdp [broadcast | multicast] [holdtime <seconds>] [maxadvertinterval <seconds>] [minadvertinterval <seconds>] [preference <number>]`

The **broadcast | multicast** parameter specifies the packet type the Layer 3 Switch uses to send Router Advertisement:

- **broadcast** – The Layer 3 Switch sends Router Advertisement as IP broadcasts. This is the default.
- **multicast** – The Layer 3 Switch sends Router Advertisement as multicast packets addressed to IP multicast group 224.0.0.1.

The **holdtime <seconds>** parameter specifies how long a host that receives a Router Advertisement from the Layer 3 Switch should consider the advertisement to be valid. When a host receives a new Router Advertisement message from the Layer 3 Switch, the host resets the hold time for the Layer 3 Switch to the hold time specified in the new advertisement. If the hold time of an advertisement expires, the host discards the advertisement, concluding that the router interface that sent the advertisement is no longer available. The value must be greater than the value of the **maxadvertinterval** parameter and cannot be greater than 9000. The default is three times the value of the **maxadvertinterval** parameter.

The **maxadvertinterval** parameter specifies the maximum amount of time the Layer 3 Switch waits between sending Router Advertisements. You can specify a value from 1 to the current value of the **holdtime** parameter. The default is 600 seconds.

The **minadvertinterval** parameter specifies the minimum amount of time the Layer 3 Switch can wait between sending Router Advertisements. The default is three-fourths (0.75) the value of the **maxadvertinterval** parameter. If you change the **maxadvertinterval** parameter, the software automatically adjusts the **minadvertinterval** parameter to be three-fourths the new value of the **maxadvertinterval** parameter. If you want to override the automatically configured value, you can specify an interval from 1 to the current value of the **maxadvertinterval** parameter.

The **preference <number>** parameter specifies the IRDP preference level of this Layer 3 Switch. If a host receives Router Advertisements from multiple routers, the host selects the router interface that sent the message with the highest interval as the host default gateway. The valid range is 0-4294967296 to 0-4294967295. The default is 0.

Configuring RARP

The Reverse Address Resolution Protocol (RARP) provides a simple mechanism for directly-attached IP hosts to boot over the network. RARP allows an IP host that does not have a means of storing its IP address across power cycles or software reloads to query a directly-attached router for an IP address.

RARP is enabled by default. However, you must create a RARP entry for each host that will use the Layer 3 Switch for booting. A RARP entry consists of the following information:

- The entry number – the entry sequence number in the RARP table.
- The MAC address of the boot client.
- The IP address you want the Layer 3 Switch to give to the client.

When a client sends a RARP broadcast requesting an IP address, the Layer 3 Switch responds to the request by looking in the RARP table for an entry that contains the client MAC address:

- If the RARP table contains an entry for the client, the Layer 3 Switch sends a unicast response to the client that contains the IP address associated with the client MAC address in the RARP table.
- If the RARP table does not contain an entry for the client, the Layer 3 Switch silently discards the RARP request and does not reply to the client.

How RARP Differs from BootP/DHCP

RARP and BootP/DHCP are different methods for providing IP addresses to IP hosts when they boot. These methods differ in the following ways:

- Location of configured host addresses:
 - RARP requires static configuration of the host IP addresses on the Layer 3 Switch. The Layer 3 Switch replies directly to a host request by sending an IP address you have configured in the RARP table.
 - The Layer 3 Switch forwards BootP and DHCP requests to a third-party BootP/DHCP server that contains the IP addresses and other host configuration information.
- Connection of host to boot source (Layer 3 Switch or BootP/DHCP server):
 - RARP requires the IP host to be directly attached to the Layer 3 Switch.
 - An IP host and the BootP/DHCP server can be on different networks and on different routers, so long as the routers are configured to forward (“help”) the host boot request to the boot server.
 - You can centrally configure other host parameters on the BootP/DHCP server, in addition to the IP address, and supply those parameters to the host along with its IP address.

To configure the Layer 3 Switch to forward BootP/DHCP requests when boot clients and the boot servers are on different subnets on different Layer 3 Switch interfaces, refer to [“Configuring BootP/DHCP relay parameters”](#) on page 615.

Disabling RARP

RARP is enabled by default. To disable RARP, enter the following command at the global CONFIG level.

```
PowerConnect(config)# no ip rarp
```

Syntax: [no] ip rarp

To re-enable RARP, enter the following command.

```
PowerConnect(config)# ip rarp
```

Creating static RARP entries

You must configure the RARP entries for the RARP table. The Layer 3 Switch can send an IP address in reply to a client RARP request only if create a RARP entry for that client.

To assign a static IP RARP entry for static routes on a router, enter a command such as the following.

```
PowerConnect(config)# rarp 1 1245.7654.2348 192.53.4.2
```

This command creates a RARP entry for a client with MAC address 1245.7654.2348. When the Layer 3 Switch receives a RARP request from this client, the Layer 3 Switch replies to the request by sending IP address 192.53.4.2 to the client.

Syntax: `rarp <number> <mac-addr>.<ip-addr>`

The `<number>` parameter identifies the RARP entry number. You can specify an unused number from 1 to the maximum number of RARP entries supported on the device. To determine the maximum number of entries supported on the device, refer to the section [“Displaying and modifying system parameter default settings”](#) on page 184.

The `<mac-addr>` parameter specifies the MAC address of the RARP client.

The `<ip-addr>` parameter specifies the IP address the Layer 3 Switch will give the client in response to the client RARP request.

Changing the maximum number of static RARP entries supported

The number of RARP entries the Layer 3 Switch supports depends on how much memory the Layer 3 Switch has. To determine how many RARP entries your Layer 3 Switch can have, display the system default information using the procedure in the section [“Displaying and modifying system parameter default settings”](#) on page 184.

If your Layer 3 Switch allows you to increase the maximum number of RARP entries, you can use a procedure in the same section to do so.

NOTE

You must save the configuration to the startup-config file and reload the software after changing the RARP cache size to place the change into effect.

Configuring UDP broadcast and IP helper parameters

Some applications rely on client requests sent as limited IP broadcasts addressed to the UDP application port. If a server for the application receives such a broadcast, the server can reply to the client. Routers do not forward subnet directed broadcasts, so the client and server must be on the same network for the broadcast to reach the server. If the client and server are on different networks (on opposite sides of a router), the client request cannot reach the server.

You can configure the Layer 3 Switch to forward clients' requests to UDP application servers. To do so:

- Enable forwarding support for the UDP application port, if forwarding support is not already enabled.
- Configure a helper address on the interface connected to the clients. Specify the helper address to be the IP address of the application server or the subnet directed broadcast address for the IP subnet the server is in. A helper address is associated with a specific interface and applies only to client requests received on that interface. The Layer 3 Switch forwards client requests for any of the application ports the Layer 3 Switch is enabled to forward to the helper address.

Forwarding support for the following application ports is enabled by default:

- bootps (port 67)
- dns (port 53)
- tftp (port 69)

- time (port 37)
- netbios-ns (port 137)
- netbios-dgm (port 138)
- tacacs (port 65)

NOTE

The application names are the names for these applications that the Layer 3 Switch software recognizes, and might not match the names for these applications on some third-party devices. The numbers listed in parentheses are the UDP port numbers for the applications. The numbers come from RFC 1340.

NOTE

Forwarding support for BootP/DHCP is enabled by default. If you are configuring the Layer 3 Switch to forward BootP/DHCP requests, refer to [“Configuring BootP/DHCP relay parameters”](#) on page 615.

You can enable forwarding for other applications by specifying the application port number.

You also can disable forwarding for an application.

NOTE

If you disable forwarding for a UDP application, forwarding of client requests received as broadcasts to helper addresses is disabled. Disabling forwarding of an application does not disable other support for the application. For example, if you disable forwarding of Telnet requests to helper addresses, other Telnet support on the Layer 3 Switch is not also disabled.

Enabling forwarding for a UDP application

If you want the Layer 3 Switch to forward client requests for UDP applications that the Layer 3 Switch does not forward by default, you can enable forwarding support for the port. To enable forwarding support for a UDP application, use the following method. You also can disable forwarding for an application using this method.

NOTE

You also must configure a helper address on the interface that is connected to the clients for the application. The Layer 3 Switch cannot forward the requests unless you configure the helper address. Refer to [“Configuring an IP helper address”](#) on page 615.

To enable the forwarding of SNMP trap broadcasts, enter the following command.

```
PowerConnect(config)# ip forward-protocol udp snmp-trap
```

Syntax: [no] ip forward-protocol udp <udp-port-name> | <udp-port-num>

The <udp-port-name> parameter can have one of the following values. For reference, the corresponding port numbers from RFC 1340 are shown in parentheses. If you specify an application name, enter the name only, not the parentheses or the port number shown here:

- bootpc (port 68)
- bootps (port 67)
- discard (port 9)
- dns (port 53)
- dnsix (port 90)

- echo (port 7)
- mobile-ip (port 434)
- netbios-dgm (port 138)
- netbios-ns (port 137)
- ntp (port 123)
- tacacs (port 65)
- talk (port 517)
- time (port 37)
- tftp (port 69)

In addition, you can specify any UDP application by using the application UDP port number.

The `<udp-port-num>` parameter specifies the UDP application port number. If the application you want to enable is not listed above, enter the application port number. You also can list the port number for any of the applications listed above.

To disable forwarding for an application, enter a command such as the following.

```
PowerConnect(config)# no ip forward-protocol udp snmp
```

This command disables forwarding of SNMP requests to the helper addresses configured on Layer 3 Switch interfaces.

Configuring an IP helper address

To forward a client broadcast request for a UDP application when the client and server are on different networks, you must configure a helper address on the interface connected to the client. Specify the server IP address or the subnet directed broadcast address of the IP subnet the server is in as the helper address.

You can configure up to 16 helper addresses on each interface. You can configure a helper address on an Ethernet port or a virtual interface.

To configure a helper address on interface 2 on chassis module 1, enter the following commands.

```
PowerConnect(config)# interface e 2  
PowerConnect(config-if-2)# ip helper-address 1 207.95.7.6
```

The commands in this example change the CLI to the configuration level for port 2, then add a helper address for server 207.95.7.6 to the port. If the port receives a client request for any of the applications that the Layer 3 Switch is enabled to forward, the Layer 3 Switch forwards the client request to the server.

Syntax: `ip helper-address <num> <ip-addr>`

The `<num>` parameter specifies the helper address number and can be from 1 – 16.

The `<ip-addr>` command specifies the server IP address or the subnet directed broadcast address of the IP subnet the server is in.

Configuring BootP/DHCP relay parameters

A host on an IP network can use BootP/DHCP to obtain its IP address from a BootP/DHCP server. To obtain the address, the client sends a BootP/DHCP request. The request is a subnet directed broadcast and is addressed to UDP port 67. A limited IP broadcast is addressed to IP address 255.255.255.255 and is not forwarded by the Layer 3 Switch or other IP routers.

When the BootP/DHCP client and server are on the same network, the server receives the broadcast request and replies to the client. However, when the client and server are on different networks, the server does not receive the client request, because the Layer 3 Switch does not forward the request.

You can configure the Layer 3 Switch to forward BootP/DHCP requests. To do so, configure a helper address on the interface that receives the client requests, and specify the BootP/DHCP server IP address as the address you are helping the BootP/DHCP requests to reach. Instead of the server IP address, you can specify the subnet directed broadcast address of the IP subnet the server is in.

BootP/DHCP relay parameters

The following parameters control the Layer 3 Switch forwarding of BootP/DHCP requests:

- **Helper address** – The BootP/DHCP server IP address. You must configure the helper address on the interface that receives the BootP/DHCP requests from the client. The Layer 3 Switch cannot forward a request to the server unless you configure a helper address for the server.
- **Gateway address** – The Layer 3 Switch places the IP address of the interface that received the BootP/DHCP request in the request packet Gateway Address field (sometimes called the Router ID field). When the server responds to the request, the server sends the response as a unicast packet to the IP address in the Gateway Address field. (If the client and server are directly attached, the Gateway ID field is empty and the server replies to the client using a unicast or broadcast packet, depending on the server.)

By default, the Layer 3 Switch uses the lowest-numbered IP address on the interface that receives the request as the Gateway address. You can override the default by specifying the IP address you want the Layer 3 Switch to use.

- **Hop count** – Each router that forwards a BootP/DHCP packet increments the hop count by 1. Routers also discard a forwarded BootP/DHCP request instead of forwarding the request if the hop count is greater than the maximum number of BootP/DHCP hops allowed by the router. By default, a Layer 3 Switch forwards a BootP/DHCP request if its hop count is four or less, but discards the request if the hop count is greater than four. You can change the maximum number of hops the Layer 3 Switch will allow to a value from 1 – 15.

NOTE

The BootP/DHCP hop count is not the TTL parameter.

Configuring an IP helper address

The procedure for configuring a helper address for BootP/DHCP requests is the same as the procedure for configuring a helper address for other types of UDP broadcasts. Refer to [“Configuring an IP helper address”](#) on page 614.

Configuring the BOOTP/DHCP reply source address

You can configure the device so that a BOOTP/DHCP reply to a client contains the server IP address as the source address instead of the router IP address. To do so, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)# ip helper-use-responder-ip
```

Syntax: [no] ip helper-use-responder-ip

Changing the maximum number of hops to a BootP relay server

Each BootP/DHCP request includes a field Hop Count field. The Hop Count field indicates how many routers the request has passed through. When the Layer 3 Switch receives a BootP/DHCP request, the Layer 3 Switch looks at the value in the Hop Count field:

- If the hop count value is equal to or less than the maximum hop count the Layer 3 Switch allows, the Layer 3 Switch increments the hop count by one and forwards the request.
- If the hop count is greater than the maximum hop count the Layer 3 Switch allows, the Layer 3 Switch discards the request.

To change the maximum number of hops the Layer 3 Switch allows for forwarded BootP/DHCP requests, use either of the following methods.

NOTE

The BootP/DHCP hop count is not the TTL parameter.

To modify the maximum number of BootP/DHCP hops, enter the following command.

```
PowerConnect(config)# bootp-relay-max-hops 10
```

This command allows the Layer 3 Switch to forward BootP/DHCP requests that have passed through ten previous hops before reaching the Layer 3 Switch. Requests that have traversed 11 hops before reaching the switch are dropped. Since the hop count value initializes at zero, the hop count value of an ingressing DHCP Request packet is the number of Layer 3 routers that the packet has already traversed.

Syntax: bootp-relay-max-hops <1 through 15>

Configuring IP parameters – Layer 2 Switches

The following sections describe how to configure IP parameters on a Layer 2 Switch.

NOTE

This section describes how to configure IP parameters for Layer 2 Switches. For IP configuration information for Layer 3 Switches, refer to [“Configuring IP parameters – Layer 3 Switches”](#) on page 578.

Configuring the management IP address and specifying the default gateway

To manage a Layer 2 Switch using Telnet or Secure Shell (SSH) CLI connections, you must configure an IP address for the Layer 2 Switch. Optionally, you also can specify the default gateway.

Devices support both classical IP network masks (Class A, B, and C subnet masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks:

- To enter a classical network mask, enter the mask in IP address format. For example, enter “209.157.22.99 255.255.255.0” for an IP address with a Class-C subnet mask.
- To enter a prefix network mask, enter a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter “209.157.22.99/24” for an IP address that has a network mask with 24 significant bits (ones).

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the display to prefix format. Refer to [“Changing the network mask display to prefix format”](#) on page 623.

To assign an IP address to a Layer 2 Switch, enter a command such as the following at the global CONFIG level.

```
PowerConnect(config)# ip address 192.45.6.110 255.255.255.0
```

Syntax: `ip address <ip-addr> <ip-mask>`

or

Syntax: `ip address <ip-addr>/<mask-bits>`

You also can enter the IP address and mask in CIDR format, as follows.

```
PowerConnect(config)# ip address 192.45.6.1/24
```

To specify the Layer 2 Switch default gateway, enter a command such as the following.

```
PowerConnect(config)# ip default-gateway 192.45.6.1 255.255.255.0
```

```
ip default-gateway <ip-addr>
```

or

```
ip default-gateway <ip-addr>/<mask-bits>
```

NOTE

When configuring an IP address on a Layer 2 switch that has multiple VLANs, make sure the configuration includes a designated management VLAN that identifies the VLAN to which the global IP address belongs. Refer to [“Designated VLAN for Telnet management sessions to a Layer 2 Switch”](#) on page 863.

Configuring Domain Name Server (DNS) resolver

The Domain Name Server (DNS) resolver feature lets you use a host name to perform Telnet, ping, and traceroute commands. You can also define a DNS domain on a Layer 2 Switch or Layer 3 Switch and thereby recognize all hosts within that domain. After you define a domain name, the Layer 2 Switch or Layer 3 Switch automatically appends the appropriate domain to the host and forwards it to the domain name server.

For example, if the domain “newyork.com” is defined on a Layer 2 Switch or Layer 3 Switch and you want to initiate a ping to host “NYC01” on that domain, you need to reference only the host name in the command instead of the host name and its domain name. For example, you could enter either of the following commands to initiate the ping.

```
PowerConnect# ping nyc01
PowerConnect# ping nyc01.newyork.com
```

Defining a DNS entry

You can define up to four DNS servers for each DNS entry. The first entry serves as the primary default address. If a query to the primary address fails to be resolved after three attempts, the next gateway address is queried (also up to three times). This process continues for each defined gateway address until the query is resolved. The order in which the default gateway addresses are polled is the same as the order in which you enter them.

Suppose you want to define the domain name of newyork.com on a Layer 2 Switch and then define four possible default DNS gateway addresses. To do so, enter the following commands.

```
PowerConnect(config)# ip dns domain-name newyork.com
PowerConnect(config)# ip dns server-address 209.157.22.199 205.96.7.15
208.95.7.25 201.98.7.15
```

Syntax: `ip dns server-address <ip-addr> [<ip-addr>] [<ip-addr>] [<ip-addr>]`

In this example, the first IP address in the `ip dns server-address...` command becomes the primary gateway address and all others are secondary addresses. Because IP address 201.98.7.15 is the last address listed, it is also the last address consulted to resolve a query.

Using a DNS name To initiate a trace route

Example

Suppose you want to trace the route from a Layer 2 Switch to a remote server identified as NYC02 on domain newyork.com. Because the newyork.com domain is already defined on the Layer 2 Switch, you need to enter only the host name, NYC02, as noted below.

```
PowerConnect# traceroute nyc02
```

Syntax: `traceroute <host-ip-addr> [maxttl <value>] [minttl <value>] [numeric] [timeout <value>] [source-ip <ip addr>]`

The only required parameter is the IP address of the host at the other end of the route.

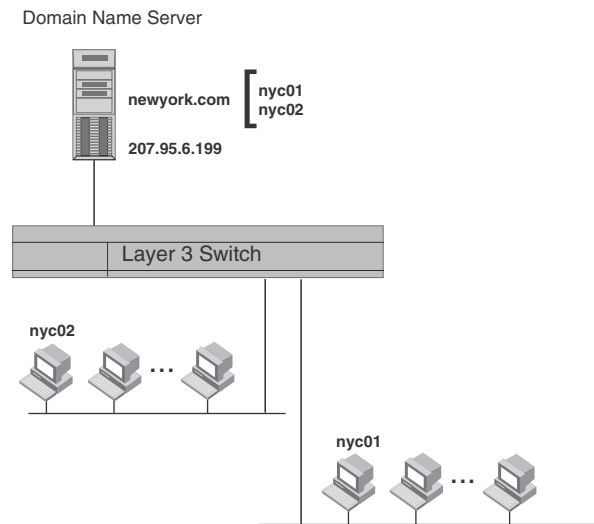
After you enter the command, a message indicating that the DNS query is in process and the current gateway address (IP address of the domain name server) being queried appear on the screen.

```
Type Control-c to abort
Sending DNS Query to 209.157.22.199
Tracing Route to IP node 209.157.22.80
To ABORT Trace Route, Please use stop-traceroute command.
Traced route to target IP node 209.157.22.80:
  IP Address          Round Trip Time1    Round Trip Time2
  207.95.6.30         93 msec             121 msec
```

NOTE

In the above example, 209.157.22.199 is the IP address of the domain name server (default DNS gateway address), and 209.157.22.80 represents the IP address of the NYC02 host.

FIGURE 101 Querying a Host on the newyork.com Domain



Changing the TTL threshold

The TTL threshold prevents routing loops by specifying the maximum number of router hops an IP packet originated by the Layer 2 Switch can travel through. Each device capable of forwarding IP that receives the packet decrements (decreases) the packet TTL by one. If a router receives a packet with a TTL of 1 and reduces the TTL to zero, the router drops the packet.

The default TTL is 64. You can change the TTL to a value from 1 – 255.

To modify the TTL threshold to 25, enter the following commands.

```
PowerConnect(config)# ip ttl 25
PowerConnect(config)# exit
```

Syntax: `ip ttl <1-255>`

Configuring DHCP Assist

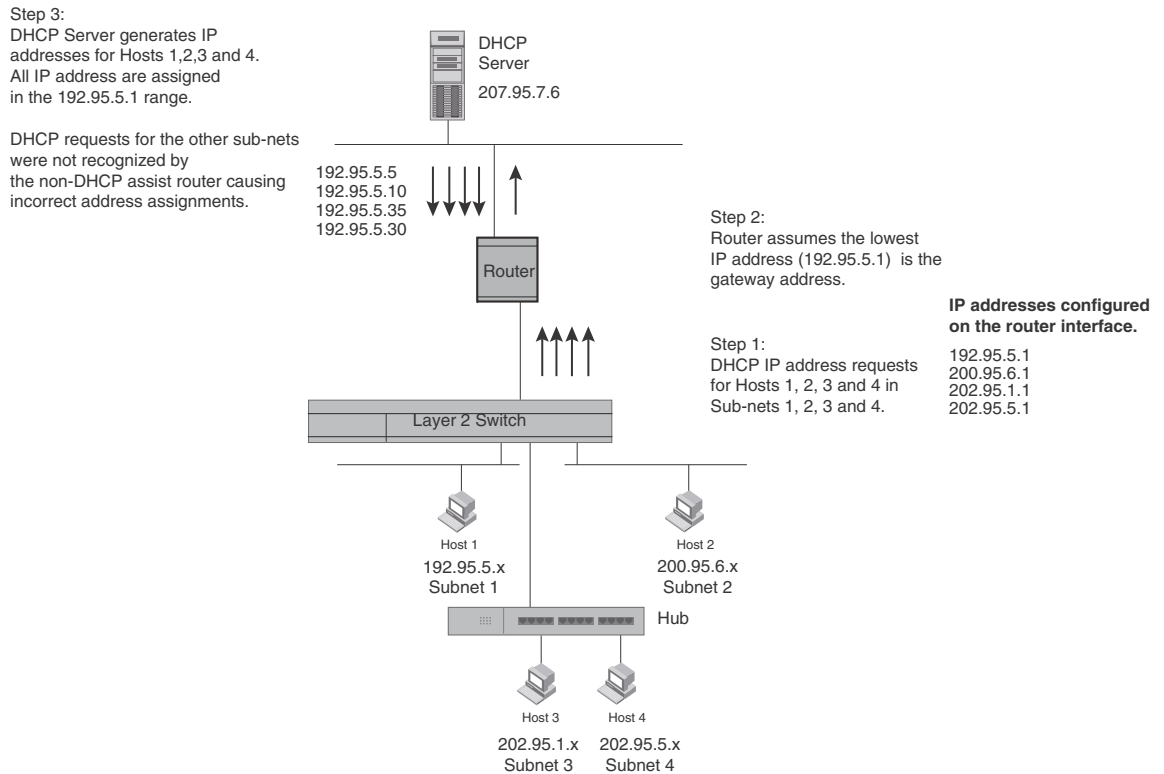
DHCP Assist allows a Layer 2 Switch to assist a router that is performing multi-netting on its interfaces as part of its DHCP relay function.

DHCP Assist ensures that a DHCP server that manages multiple IP subnets can readily recognize the requester IP subnet, even when that server is not on the client local LAN segment. The Layer 2 Switch does so by stamping each request with its IP gateway address in the DHCP discovery packet.

By allowing multiple subnet DHCP requests to be sent on the same wire, you can reduce the number of router ports required to support secondary addressing as well as reduce the number of DHCP servers required, by allowing a server to manage multiple subnet address assignments.

FIGURE 102 DHCP requests in a network without DHCP Assist on the Layer 2 Switch

21 Configuring IP parameters – Layer 2 Switches



In a network operating without DHCP Assist, hosts can be assigned IP addresses from the wrong subnet range because a router with multiple subnets configured on an interface cannot distinguish among DHCP discovery packets received from different subnets.

For example, in [Figure 102](#), a host from each of the four subnets supported on a Layer 2 Switch requests an IP address from the DHCP server. These requests are sent transparently to the router. Because the router is unable to determine the origin of each packet by subnet, it assumes the lowest IP address or the 'primary address' is the gateway for all ports on the Layer 2 Switch and stamps the request with that address.

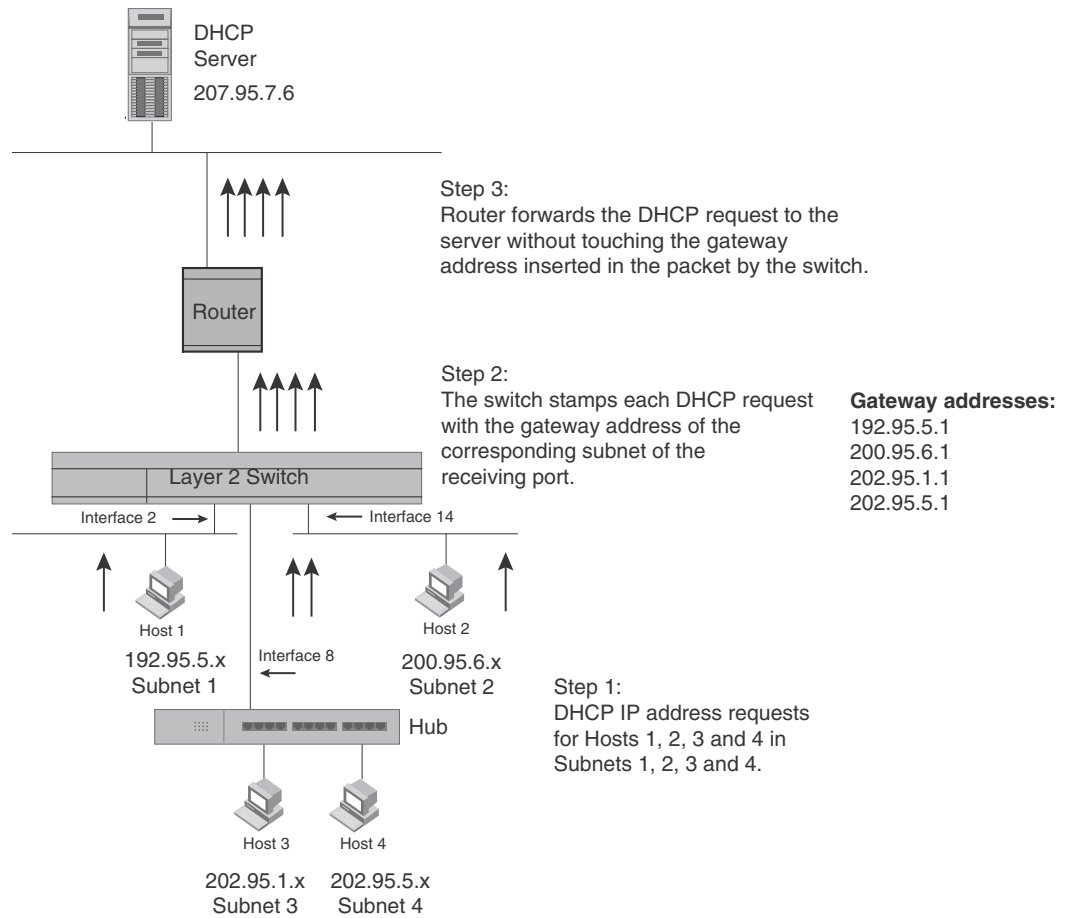
When the DHCP request is received at the server, it assigns all IP addresses within that range only.

With DHCP Assist enabled on a Layer 2 Switch, correct assignments are made because the Layer 2 Switch provides the stamping service.

How DHCP Assist works

Upon initiation of a DHCP session, the client sends out a DHCP discovery packet for an address from the DHCP server as seen in [Figure 103](#). When the DHCP discovery packet is received at a Layer 2 Switch with the DHCP Assist feature enabled, the gateway address configured on the receiving interface is inserted into the packet. This address insertion is also referred to as stamping.

FIGURE 103 DHCP requests in a network with DHCP Assist operating on a Switch



When the stamped DHCP discovery packet is then received at the router, it is forwarded to the DHCP server. The DHCP server then extracts the gateway address from each request and assigns an available IP address within the corresponding IP subnet (Figure 104). The IP address is then forwarded back to the workstation that originated the request.

NOTE

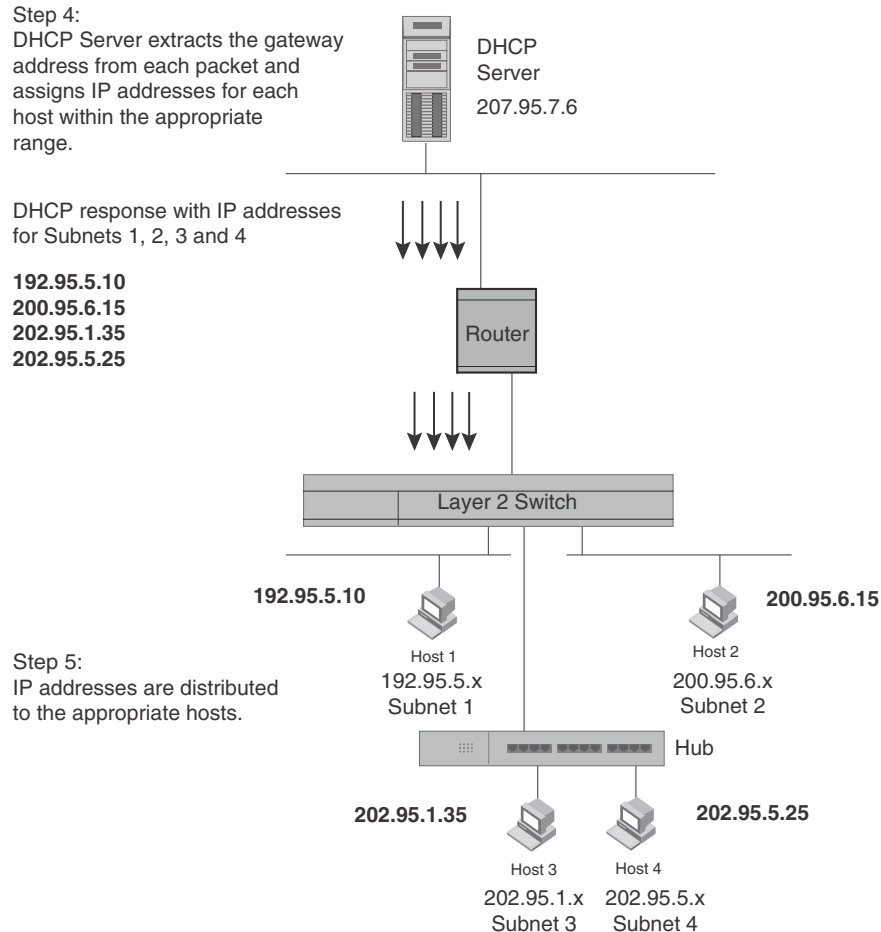
When DHCP Assist is enabled on any port, Layer 2 broadcast packets are forwarded by the CPU. Unknown unicast and multicast packets are still forwarded in hardware, although selective packets such as IGMP, are sent to the CPU for analysis. When DHCP Assist is not enabled, Layer 2 broadcast packets are forwarded in hardware.

NOTE

The DHCP relay function of the connecting router must be turned on.

FIGURE 104 DHCP Offers are forwarded back toward the requestors

21 Configuring IP parameters – Layer 2 Switches



NOTE

When DHCP Assist is enabled on any port, Layer 2 broadcast packets are forwarded by the CPU. Unknown unicast and multicast packets are still forwarded in hardware, although selective packets such as IGMP are sent to the CPU for analysis. When DHCP Assist is not enabled, Layer 2 broadcast packets are forwarded in hardware.

Configuring DHCP Assist

You can associate a gateway list with a port. You must configure a gateway list when DHCP Assist is enabled on a Layer 2 Switch. The gateway list contains a gateway address for each subnet that will be requesting addresses from a DHCP server. The list allows the stamping process to occur. Each gateway address defined on the Layer 2 Switch corresponds to an IP address of the router interface or other routers involved.

Up to eight addresses can be defined for each gateway list in support of ports that are multi-homed. When multiple IP addresses are configured for a gateway list, the Layer 2 Switch inserts the addresses into the discovery packet in a round robin fashion.

Up to 32 gateway lists can be defined for each Layer 2 Switch.

Example

To create the configuration indicated in [Figure 103](#) and [Figure 104](#), enter commands such as the following.

```
PowerConnect(config)# dhcp-gateway-list 1 192.95.5.1
PowerConnect(config)# dhcp-gateway-list 2 200.95.6.1
PowerConnect(config)# dhcp-gateway-list 3 202.95.1.1 202.95.5.1
PowerConnect(config)# int e 2
PowerConnect(config-if-e10000-2)# dhcp-gateway-list 1
PowerConnect(config-if-e10000-2)# int e8
PowerConnect(config-if-e10000-8)# dhcp-gateway-list 3
PowerConnect(config-if-e10000-8)# int e 14
PowerConnect(config-if-e10000-14)# dhcp-gateway-list 2
```

Syntax: `dhcp-gateway-list <num> <ip-addr>`

Displaying IP configuration information and statistics

The following sections describe IP display options for Layer 3 Switches and Layer 2 Switches:

- To display IP information on a Layer 3 Switch, refer to [“Displaying IP information – Layer 3 Switches”](#) on page 623.
- To display IP information on a Layer 2 Switch, refer to [“Displaying IP information – Layer 2 Switches”](#) on page 637.

Changing the network mask display to prefix format

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the displays to prefix format (example: /18) on a Layer 3 Switch or Layer 2 Switch using the following CLI method.

To enable CIDR format for displaying network masks, entering the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# ip show-subnet-length
```

Syntax: `[no] ip show-subnet-length`

Displaying IP information – Layer 3 Switches

You can display the following IP configuration information statistics on Layer 3 Switches:

- Global IP parameter settings and IP access policies – refer to [“Displaying global IP configuration information”](#) on page 624.
- CPU utilization statistics – refer to [“Displaying CPU utilization statistics”](#) on page 625.
- IP interfaces – refer to [“Displaying IP interface information”](#) on page 627.
- ARP entries – refer to [“Displaying ARP entries”](#) on page 628.
- Static ARP entries – refer to [“Displaying ARP entries”](#) on page 628.
- IP forwarding cache – refer to [“Displaying the forwarding cache”](#) on page 631.
- IP route table – refer to [“Displaying the IP route table”](#) on page 632.
- IP traffic statistics – refer to [“Displaying IP traffic statistics”](#) on page 634.

The sections below describe how to display this information.

In addition to the information described below, you can display the following IP information. This information is described in other parts of this guide:

- RIP
- OSPF
- BGP4
- PIM
- VRRP or VRRPE

Displaying global IP configuration information

To display IP configuration information, enter the following command at any CLI level.

```
PowerConnect# show ip
```

Global Settings

```
ttl: 64, arp-age: 10, bootp-relay-max-hops: 4
router-id : 207.95.11.128
enabled : UDP-Broadcast-Forwarding IRDP Proxy-ARP RARP OSPF
disabled: BGP4 Load-Sharing RIP FSRP VRRP
```

Static Routes

Index	IP Address	Subnet Mask	Next Hop Router	Metric	Distance
1	0.0.0.0	0.0.0.0	209.157.23.2	1	1

Policies

Index	Action	Source	Destination	Protocol	Port	Operator
1	deny	209.157.22.34	209.157.22.26	tcp	http	=
64	permit	any	any			

Syntax: show ip

NOTE

This command has additional options, which are explained in other sections in this guide, including the sections below this one.

This display shows the following information.

TABLE 101 CLI Display of global IP configuration information – Layer 3 Switch

This field...	Displays...
Global settings	
ttl	The Time-To-Live (TTL) for IP packets. The TTL specifies the maximum number of router hops a packet can travel before reaching the router. If the packet TTL value is higher than the value specified in this field, the router drops the packet. To change the maximum TTL, refer to “Changing the TTL threshold” on page 592.
arp-age	The ARP aging period. This parameter specifies how many minutes an inactive ARP entry remains in the ARP cache before the router ages out the entry. To change the ARP aging period, refer to “Changing the ARP aging period” on page 589.
bootp-relay-max-hops	The maximum number of hops away a BootP server can be located from the router and still be used by the router clients for network booting. To change this value, refer to “Changing the maximum number of hops to a BootP relay server” on page 616.

TABLE 101 CLI Display of global IP configuration information – Layer 3 Switch (Continued)

This field...	Displays...
router-id	The 32-bit number that uniquely identifies the router. By default, the router ID is the numerically lowest IP interface configured on the router. To change the router ID, refer to “Changing the router ID” on page 584.
enabled	The IP-related protocols that are enabled on the router.
disabled	The IP-related protocols that are disabled on the router.
Static routes	
Index	The row number of this entry in the IP route table.
IP Address	The IP address of the route destination.
Subnet Mask	The network mask for the IP address.
Next Hop Router	The IP address of the router interface to which the router sends packets for the route.
Metric	The cost of the route. Usually, the metric represents the number of hops to the destination.
Distance	The administrative distance of the route. The default administrative distance for static IP routes in routers is 1. To list the default administrative distances for all types of routes or to change the administrative distance of a static route, refer to “Changing administrative distances” on page 777.
Policies	
Index	The policy number. This is the number you assigned the policy when you configured it.
Action	The action the router takes if a packet matches the comparison values in the policy. The action can be one of the following: <ul style="list-style-type: none"> deny – The router drops packets that match this policy. permit – The router forwards packets that match this policy.
Source	The source IP address the policy matches.
Destination	The destination IP address the policy matches.
Protocol	The IP protocol the policy matches. The protocol can be one of the following: <ul style="list-style-type: none"> ICMP IGMP IGRP OSPF TCP UDP
Port	The Layer 4 TCP or UDP port the policy checks for in packets. The port can be displayed by its number or, for port types the router recognizes, by the well-known name. For example, TCP port 80 can be displayed as HTTP. NOTE: This field applies only if the IP protocol is TCP or UDP.
Operator	The comparison operator for TCP or UDP port names or numbers. NOTE: This field applies only if the IP protocol is TCP or UDP.

Displaying CPU utilization statistics

You can display CPU utilization statistics for IP protocols using the **show process cpu** command.

21 Displaying IP configuration information and statistics

The **show process cpu** command includes CPU utilization statistics for ACL, 802.1x, and L2VLAN. L2VLAN contains any packet transmitted to a VLAN by the CPU, including unknown unicast, multicast, broadcast, and CPU forwarded Layer 2 traffic.

To display CPU utilization statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.

```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ACL           0.00       0.00       0.00       0.00        0
ARP             0.01       0.01       0.01       0.01       714
BGP            0.00       0.00       0.00       0.00        0
DOT1X        0.00       0.00       0.00       0.00        0
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.00       0.00       0.00       0.00       161
IP             0.00       0.00       0.00       0.00       229
L2VLAN       0.01       0.00       0.00       0.01       673
OSPF           0.00       0.00       0.00       0.00        0
RIP            0.00       0.00       0.00       0.00        9
STP            0.00       0.00       0.00       0.00        7
VRRP           0.00       0.00       0.00       0.00        0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example.

```
PowerConnect# show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ACL           0.00       0.00       0.00       0.00        0
ARP             0.01       0.01       0.01       0.01       714
BGP            0.00       0.00       0.00       0.00        0
DOT1X        0.00       0.00       0.00       0.00        0
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.00       0.00       0.00       0.00       161
IP             0.00       0.00       0.00       0.00       229
L2VLAN       0.01       0.00       0.00       0.01       673
OSPF           0.00       0.00       0.00       0.00        0
RIP            0.00       0.00       0.00       0.00        9
STP            0.00       0.00       0.00       0.00        7
VRRP           0.00       0.00       0.00       0.00        0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.


```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name   Sec(%)   Time(ms)
ACL           0         0.00
ARP            1         0.01
BGP            0         0.00
DOT1X        0         0.00
GVRP          0         0.00
ICMP           0         0.00
IP             0         0.00
L2VLAN      1         0.01
OSPF           0         0.00
RIP            0         0.00
STP            0         0.00
VRRP           0         0.00
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The <num> parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Displaying IP interface information

To display IP interface information, enter the following command at any CLI level.

```
PowerConnect# show ip interface

Interface      IP-Address      OK?  Method   Status      Protocol
Ethernet 1    207.95.6.173   YES  NVRAM    up          up
Ethernet 2    3.3.3.3        YES  manual   up          up
Loopback 1    1.2.3.4        YES  NVRAM    down       down
```

Syntax: `show ip interface [ethernet [<portnum>] | [loopback <num>] | [ve <num>]`

This display shows the following information.

TABLE 102 CLI display of interface IP configuration information

This field...	Displays...
Interface	The type and the port number of the interface.
IP-Address	The IP address of the interface. NOTE: If an “s” is listed following the address, this is a secondary address. When the address was configured, the interface already had an IP address in the same subnet, so the software required the “secondary” option before the software could add the interface.
OK?	Whether the IP address has been configured on the interface.
Method	Whether the IP address has been saved in NVRAM. If you have set the IP address for the interface in the CLI , but have not saved the configuration, the entry for the interface in the Method field is “manual”.

TABLE 102 CLI display of interface IP configuration information (Continued)

This field...	Displays...
Status	The link status of the interface. If you have disabled the interface with the disable command, the entry in the Status field will be “administratively down”. Otherwise, the entry in the Status field will be either “up” or “down”.
Protocol	Whether the interface can provide two-way communication. If the IP address is configured, and the link status of the interface is up, the entry in the protocol field will be “up”. Otherwise the entry in the protocol field will be “down”.

To display detailed IP information for a specific interface, enter a command such as the following.

```
PowerConnect# show ip interface ethernet 1
Interface Ethernet 1
  port state: UP
  ip address: 192.168.9.51      subnet mask: 255.255.255.0
  encapsulation: ETHERNET, mtu: 1500, metric: 1
  directed-broadcast-forwarding: disabled
  proxy-arp: disabled
  ip arp-age: 10 minutes
  Ip Flow switching is disabled
  No Helper Addresses are configured.
  No inbound ip access-list is set
  No outgoing ip access-list is set
```

Displaying ARP entries

You can display the ARP cache and the static ARP table. The ARP cache contains entries for devices attached to the Layer 3 Switch. The static ARP table contains the user-configured ARP entries. An entry in the static ARP table enters the ARP cache when the entry interface comes up.

The tables require separate display commands.

Displaying the ARP cache

To display the contents of the ARP cache, enter the following command at any CLI level.

```
PowerConnect# show arp

Total number of ARP entries: 5
  IP Address      MAC Address      Type      Age      Port      Status
1   207.95.6.102   0800.5afc.ea21   Dynamic   0        6        Valid
2   207.95.6.18    00a0.24d2.04ed   Dynamic   3        6        Pend
3   207.95.6.54    00a0.24ab.cd2b   Dynamic   0        6        Pend
4   207.95.6.101   0800.207c.a7fa   Dynamic   0        6        Valid
5   207.95.6.211   00c0.2638.ac9c   Dynamic   0        6        Valid
```

Syntax: `show arp [ethernet <portnum> | mac-address <xxxx.xxx.xxx> [<mask>] | <ip-addr> [<ip-mask>]] [<num>]`

The `<portnum>` parameter lets you restrict the display to entries for a specific port.

The `mac-address <xxxx.xxx.xxx>` parameter lets you restrict the display to entries for a specific MAC address.

The `<mask>` parameter lets you specify a mask for the `mac-address <xxxx.xxxx.xxxx>` parameter, to display entries for multiple MAC addresses. Specify the MAC address mask as “f”s and “0”s, where “f”s are significant bits.

The `<ip-addr>` and `<ip-mask>` parameters let you restrict the display to entries for a specific IP address and network mask. Specify the IP address masks in standard decimal mask format (for example, 255.255.0.0).

NOTE

The `<ip-mask>` parameter and `<mask>` parameter perform different operations. The `<ip-mask>` parameter specifies the network mask for a specific IP address, whereas the `<mask>` parameter provides a filter for displaying multiple MAC addresses that have specific values in common.

The `<num>` parameter lets you display the table beginning with a specific entry number.

NOTE

The entry numbers in the ARP cache are not related to the entry numbers for static ARP table entries.

This display shows the following information. The number in the left column of the CLI display is the row number of the entry in the ARP cache. This number is not related to the number you assign to static MAC entries in the static ARP table.

TABLE 103 CLI display of ARP cache

This field...	Displays...
IP Address	The IP address of the device.
MAC Address	The MAC address of the device.
Type	<p>The ARP entry type, which can be one of the following:</p> <ul style="list-style-type: none"> Dynamic – The Layer 3 Switch learned the entry from an incoming packet. Static – The Layer 3 Switch loaded the entry from the static ARP table when the device for the entry was connected to the Layer 3 Switch. DHCP – The Layer 3 Switch learned the entry from the DHCP binding address table. <p>NOTE: If the type is DHCP, the port number will not be available until the entry gets resolved through ARP.</p>
Age	<p>The number of minutes the entry has remained unused. If this value reaches the ARP aging period, the entry is removed from the table.</p> <p>To display the ARP aging period, refer to “Displaying global IP configuration information” on page 624. To change the ARP aging interval, refer to “Changing the ARP aging period” on page 589.</p> <p>NOTE: Static entries do not age out.</p>
Port	<p>The port on which the entry was learned.</p> <p>NOTE: If the ARP entry type is DHCP, the port number will not be available until the entry gets resolved through ARP.</p>
Status	<p>The status of the entry, which can be one of the following:</p> <ul style="list-style-type: none"> Valid – This a valid ARP entry. Pend – The ARP entry is not yet resolved.

Displaying the static ARP table

To display the static ARP table instead of the ARP cache, enter the following command at any CLI level.

21 Displaying IP configuration information and statistics

```
PowerConnect# show ip static-arp
```

```
Static ARP table size: 512, configurable from 512 to 1024
```

Index	IP Address	MAC Address	Port
1	207.95.6.111	0800.093b.d210	1
3	207.95.6.123	0800.093b.d211	1

This example shows two static entries. Note that since you specify an entry index number when you create the entry, it is possible for the range of index numbers to have gaps, as shown in this example.

NOTE

The entry number you assign to a static ARP entry is not related to the entry numbers in the ARP cache.

Syntax: `show ip static-arp [ethernet<portnum> | mac-address <xxxx.xxxx.xxxx> [<mask>] | <ip-addr> [<ip-mask>]] [<num>]`

The `<portnum>` parameter lets you restrict the display to entries for a specific port.

The `mac-address <xxxx.xxxx.xxxx>` parameter lets you restrict the display to entries for a specific MAC address.

The `<mask>` parameter lets you specify a mask for the `mac-address <xxxx.xxxx.xxxx>` parameter, to display entries for multiple MAC addresses. Specify the MAC address mask as “f”s and “0”s, where “f”s are significant bits.

The `<ip-addr>` and `<ip-mask>` parameters let you restrict the display to entries for a specific IP address and network mask. Specify the IP address masks in standard decimal mask format (for example, 255.255.0.0).

NOTE

The `<ip-mask>` parameter and `<mask>` parameter perform different operations. The `<ip-mask>` parameter specifies the network mask for a specific IP address, whereas the `<mask>` parameter provides a filter for displaying multiple MAC addresses that have specific values in common.

The `<num>` parameter lets you display the table beginning with a specific entry number.

TABLE 104 CLI display of static ARP table

This field...	Displays...
Static ARP table size	The maximum number of static entries that can be configured on the device using the current memory allocation. The range of valid memory allocations for static ARP entries is listed after the current allocation. To change the memory allocation for static ARP entries, refer to “Changing the maximum number of entries the static ARP table can hold” on page 591.
Index	The number of this entry in the table. You specify the entry number when you create the entry.
IP Address	The IP address of the device.
MAC Address	The MAC address of the device.
Port	The port attached to the device the entry is for.

Displaying the forwarding cache

To display the IP forwarding cache, enter the following command at any CLI level.

```
PowerConnect# show ip cache
```

```
Total number of cache entries: 3
D:Dynamic P:Permanent F:Forward U:Us C:Complex Filter
W:Wait ARP I:ICMP Deny K:Drop R:Fragment S:Snap Encap
      IP Address      Next Hop      MAC              Type  Port  Vlan  Pri
1      192.168.1.11     DIRECT        0000.0000.0000  PU   n/a   0
2      192.168.1.255     DIRECT        0000.0000.0000  PU   n/a   0
3      255.255.255.255   DIRECT        0000.0000.0000  PU   n/a   0
```

Syntax: `show ip cache [<ip-addr>] | [<num>]`

The `<ip-addr>` parameter displays the cache entry for the specified IP address.

The `<num>` parameter displays the cache beginning with the row following the number you enter. For example, to begin displaying the cache at row 10, enter the following command.

show ip cache 9

The `show ip cache` command displays the following information.

TABLE 105 CLI display of IP forwarding cache – Layer 3 Switch

This field...	Displays...
IP Address	The IP address of the destination.
Next Hop	The IP address of the next-hop router to the destination. This field contains either an IP address or the value DIRECT. DIRECT means the destination is either directly attached or the destination is an address on this device. For example, the next hop for loopback addresses and broadcast addresses is shown as DIRECT.
MAC	The MAC address of the destination. NOTE: If the entry is type U (indicating that the destination is this device), the address consists of zeroes.
Type	The type of host entry, which can be one or more of the following: <ul style="list-style-type: none"> • D – Dynamic • P – Permanent • F – Forward • U – Us • C – Complex Filter • W – Wait ARP • I – ICMP Deny • K – Drop • R – Fragment • S – Snap Encap
Port	The port through which this device reaches the destination. For destinations that are located on this device, the port number is shown as “n/a”.
VLAN	Indicates the VLANs the listed port is in.
Pri	The QoS priority of the port or VLAN.

Displaying the IP route table

To display the IP route table, enter the following command at any CLI level.

```
PowerConnect# show ip route
Total number of IP routes: 514
Start index: 1 B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
Destination      NetMask          Gateway          Port Cost Type
1.1.0.0          255.255.0.0     99.1.1.2        1     2    R
1.2.0.0          255.255.0.0     99.1.1.2        1     2    R
1.3.0.0          255.255.0.0     99.1.1.2        1     2    R
1.4.0.0          255.255.0.0     99.1.1.2        1     2    R
1.5.0.0          255.255.0.0     99.1.1.2        1     2    R
1.6.0.0          255.255.0.0     99.1.1.2        1     2    R
1.7.0.0          255.255.0.0     99.1.1.2        1     2    R
1.8.0.0          255.255.0.0     99.1.1.2        1     2    R
1.9.0.0          255.255.0.0     99.1.1.2        1     2    R
1.10.0.0         255.255.0.0     99.1.1.2        1     2    S
```

Syntax: `show ip route [<ip-addr> [<ip-mask>] [longer] [none-bgp]] | <num> | bgp | direct | ospf | rip | static`

The `<ip-addr>` parameter displays the route to the specified IP address.

The `<ip-mask>` parameter lets you specify a network mask or, if you prefer CIDR format, the number of bits in the network mask. If you use CIDR format, enter a forward slash immediately after the IP address, then enter the number of mask bits (for example: 209.157.22.0/24 for 209.157.22.0 255.255.255.0).

The **longer** parameter applies only when you specify an IP address and mask. This option displays only the routes for the specified IP address and mask. See the example below.

The **none-bgp** parameter displays only the routes that did not come from BGP4.

The `<num>` option displays the route table entry whose row number corresponds to the number you specify. For example, if you want to display the tenth row in the table, enter “10”.

The **bgp** option displays the BGP4 routes.

The **direct** option displays only the IP routes that are directly attached to the Layer 3 Switch.

The **ospf** option displays the OSPF routes.

The **rip** option displays the RIP routes.

The **static** option displays only the static IP routes.

The **default** routes are displayed first.

Here is an example of how to use the **direct** option. To display only the IP routes that go to devices directly attached to the Layer 3 Switch, enter the following command.

```
PowerConnect# show ip route direct
Start index: 1 B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
Destination      NetMask          Gateway          Port Cost Type
209.157.22.0     255.255.255.0   0.0.0.0          11     1    D
```

Notice that the route displayed in this example has “D” in the Type field, indicating the route is to a directly connected device.

Here is an example of how to use the **static** option. To display only the static IP routes, enter the following command.

```
PowerConnect# show ip route static
Start index: 1 B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
Destination NetMask Gateway Port Cost Type
192.144.33.11 255.255.255.0 209.157.22.12 1 2 S
```

Notice that the route displayed in this example has “S” in the Type field, indicating the route is static.

Here is an example of how to use the **longer** option. To display only the routes for a specified IP address and mask, enter a command such as the following.

```
PowerConnect# show ip route 209.159.0.0/16 longer
Starting index: 1 B:BGP D:Directly-Connected R:RIP S:Static O:OSPF
Destination NetMask Gateway Port Cost Type

52 209.159.38.0 255.255.255.0 207.95.6.101 1 1 S
53 209.159.39.0 255.255.255.0 207.95.6.101 1 1 S
54 209.159.40.0 255.255.255.0 207.95.6.101 1 1 S
55 209.159.41.0 255.255.255.0 207.95.6.101 1 1 S
56 209.159.42.0 255.255.255.0 207.95.6.101 1 1 S
57 209.159.43.0 255.255.255.0 207.95.6.101 1 1 S
58 209.159.44.0 255.255.255.0 207.95.6.101 1 1 S
59 209.159.45.0 255.255.255.0 207.95.6.101 1 1 S
60 209.159.46.0 255.255.255.0 207.95.6.101 1 1 S
```

This example shows all the routes for networks beginning with 209.159. The mask value and **longer** parameter specify the range of network addresses to be displayed. In this example, all routes within the range 209.159.0.0 – 209.159.255.255 are listed.

The **summary** option displays a summary of the information in the IP route table. The following is an example of the output from this command.

Example

```
PowerConnect# show ip route summary

IP Routing Table - 35 entries:
 6 connected, 28 static, 0 RIP, 1 OSPF, 0 BGP, 0 ISIS, 0 MPLS
Number of prefixes:
 /0: 1 /16: 27 /22: 1 /24: 5 /32: 1
```

Syntax: show ip route summary

In this example, the IP route table contains 35 entries. Of these entries, 6 are directly connected devices, 28 are static routes, and 1 route was calculated through OSPF. One of the routes has a zero-bit mask (this is the default route), 27 have a 22-bit mask, 5 have a 24-bit mask, and 1 has a 32-bit mask.

The following table lists the information displayed by the **show ip route** command.

TABLE 106 CLI display of IP route table

This field...	Displays...
Destination	The destination network of the route.
NetMask	The network mask of the destination address.
Gateway	The next-hop router.
Port	The port through which this router sends packets to reach the route's destination.
Cost	The route's cost.
Type	<p>The route type, which can be one of the following:</p> <ul style="list-style-type: none"> • B – The route was learned from BGP. • D – The destination is directly connected to this Layer 3 Switch. • R – The route was learned from RIP. • S – The route is a static route. • * – The route is a candidate default route. • O – The route is an OSPF route. Unless you use the ospf option to display the route table, "O" is used for all OSPF routes. If you do use the ospf option, the following type codes are used: <ul style="list-style-type: none"> • O – OSPF intra area route (within the same area). • IA – The route is an OSPF inter area route (a route that passes from one area into another). • E1 – The route is an OSPF external type 1 route. • E2 – The route is an OSPF external type 2 route.

Clearing IP routes

If needed, you can clear the entire route table or specific individual routes.

To clear all routes from the IP route table, enter the following command.

```
PowerConnect# clear ip route
```

To clear route 209.157.22.0/24 from the IP routing table, enter the following command.

```
PowerConnect# clear ip route 209.157.22.0/24
```

Syntax: `clear ip route [<ip-addr> <ip-mask>]`

or

Syntax: `clear ip route [<ip-addr>/<mask-bits>]`

Displaying IP traffic statistics

To display IP traffic statistics, enter the following command at any CLI level.


```
PowerConnect# show ip traffic
IP Statistics
  139 received, 145 sent, 0 forwarded
  0 filtered, 0 fragmented, 0 reassembled, 0 bad header
  0 no route, 0 unknown proto, 0 no buffer, 0 other errors

ICMP Statistics
Received:
  0 total, 0 errors, 0 unreachable, 0 time exceed
  0 parameter, 0 source quench, 0 redirect, 0 echo,
  0 echo reply, 0 timestamp, 0 timestamp reply, 0 addr mask
  0 addr mask reply, 0 irdp advertisement, 0 irdp solicitation
Sent:
  0 total, 0 errors, 0 unreachable, 0 time exceed
  0 parameter, 0 source quench, 0 redirect, 0 echo,
  0 echo reply, 0 timestamp, 0 timestamp reply, 0 addr mask
  0 addr mask reply, 0 irdp advertisement, 0 irdp solicitation

UDP Statistics
  1 received, 0 sent, 1 no port, 0 input errors

TCP Statistics
  0 active opens, 0 passive opens, 0 failed attempts
  0 active resets, 0 passive resets, 0 input errors
  138 in segments, 141 out segments, 4 retransmission

RIP Statistics
  0 requests sent, 0 requests received
  0 responses sent, 0 responses received
  0 unrecognized, 0 bad version, 0 bad addr family, 0 bad req format
  0 bad metrics, 0 bad resp format, 0 resp not from rip port
  0 resp from loopback, 0 packets rejected
```

The **show ip traffic** command displays the following information.

TABLE 107 CLI display of IP traffic statistics – Layer 3 Switch

This field...	Displays...
IP statistics	
received	The total number of IP packets received by the device.
sent	The total number of IP packets originated and sent by the device.
forwarded	The total number of IP packets received by the device and forwarded to other devices.
filtered	The total number of IP packets filtered by the device.
fragmented	The total number of IP packets fragmented by this device to accommodate the MTU of this device or of another device.
reassembled	The total number of fragmented IP packets that this device re-assembled.
bad header	The number of IP packets dropped by the device due to a bad packet header.
no route	The number of packets dropped by the device because there was no route.
unknown proto	The number of packets dropped by the device because the value in the Protocol field of the packet header is unrecognized by this device.
no buffer	This information is used by Dell customer support.
other errors	The number of packets dropped due to error types other than those listed above.

TABLE 107 CLI display of IP traffic statistics – Layer 3 Switch (Continued)

This field...	Displays...
ICMP statistics	
The ICMP statistics are derived from RFC 792, "Internet Control Message Protocol", RFC 950, "Internet Standard Subnetting Procedure", and RFC 1256, "ICMP Router Discovery Messages". Statistics are organized into Sent and Received. The field descriptions below apply to each.	
total	The total number of ICMP messages sent or received by the device.
errors	This information is used by Dell customer support.
unreachable	The number of Destination Unreachable messages sent or received by the device.
time exceed	The number of Time Exceeded messages sent or received by the device.
parameter	The number of Parameter Problem messages sent or received by the device.
source quench	The number of Source Quench messages sent or received by the device.
redirect	The number of Redirect messages sent or received by the device.
echo	The number of Echo messages sent or received by the device.
echo reply	The number of Echo Reply messages sent or received by the device.
timestamp	The number of Timestamp messages sent or received by the device.
timestamp reply	The number of Timestamp Reply messages sent or received by the device.
addr mask	The number of Address Mask Request messages sent or received by the device.
addr mask reply	The number of Address Mask Replies messages sent or received by the device.
irdp advertisement	The number of ICMP Router Discovery Protocol (IRDP) Advertisement messages sent or received by the device.
irdp solicitation	The number of IRDP Solicitation messages sent or received by the device.
UDP statistics	
received	The number of UDP packets received by the device.
sent	The number of UDP packets sent by the device.
no port	The number of UDP packets dropped because they did not have a valid UDP port number.
input errors	This information is used by Dell customer support.
TCP statistics	
The TCP statistics are derived from RFC 793, "Transmission Control Protocol".	
active opens	The number of TCP connections opened by sending a TCP SYN to another device.
passive opens	The number of TCP connections opened by this device in response to connection requests (TCP SYNs) received from other devices.
failed attempts	This information is used by Dell customer support.
active resets	The number of TCP connections this device reset by sending a TCP RESET message to the device at the other end of the connection.
passive resets	The number of TCP connections this device reset because the device at the other end of the connection sent a TCP RESET message.
input errors	This information is used by Dell customer support.
in segments	The number of TCP segments received by the device.

TABLE 107 CLI display of IP traffic statistics – Layer 3 Switch (Continued)

This field...	Displays...
out segments	The number of TCP segments sent by the device.
retransmission	The number of segments that this device retransmitted because the retransmission timer for the segment had expired before the device at the other end of the connection had acknowledged receipt of the segment.
RIP statistics	
The RIP statistics are derived from RFC 1058, "Routing Information Protocol".	
requests sent	The number of requests this device has sent to another RIP router for all or part of its RIP routing table.
requests received	The number of requests this device has received from another RIP router for all or part of this device RIP routing table.
responses sent	The number of responses this device has sent to another RIP router request for all or part of this device RIP routing table.
responses received	The number of responses this device has received to requests for all or part of another RIP router routing table.
unrecognized	This information is used by Dell customer support.
bad version	The number of RIP packets dropped by the device because the RIP version was either invalid or is not supported by this device.
bad addr family	The number of RIP packets dropped because the value in the Address Family Identifier field of the packet header was invalid.
bad req format	The number of RIP request packets this router dropped because the format was bad.
bad metrics	This information is used by Dell customer support.
bad resp format	The number of responses to RIP request packets dropped because the format was bad.
resp not from rip port	This information is used by Dell customer support.
resp from loopback	The number of RIP responses received from loopback interfaces.
packets rejected	This information is used by Dell customer support.

Displaying IP information – Layer 2 Switches

You can display the following IP configuration information statistics on Layer 2 Switches:

- Global IP settings – refer to [“Displaying global IP configuration information”](#) on page 637.
- ARP entries – refer to [“Displaying ARP entries”](#) on page 638.
- IP traffic statistics – refer to [“Displaying IP traffic statistics”](#) on page 639.

Displaying global IP configuration information

To display the Layer 2 Switch IP address and default gateway, enter the following command.

21 Displaying IP configuration information and statistics

```
PowerConnect# show ip

Switch IP address: 192.168.1.2

Subnet mask: 255.255.255.0

Default router address: 192.168.1.1
TFTP server address: None
Configuration filename: None
Image filename: None
```

Syntax: show ip

This display shows the following information.

TABLE 108 CLI display of global IP configuration information – Layer 2 Switch

This field...	Displays...
IP configuration	
Switch IP address	The management IP address configured on the Layer 2 Switch. Specify this address for Telnet.
Subnet mask	The subnet mask for the management IP address.
Default router address	The address of the default gateway, if you specified one.
Most recent TFTP access	
TFTP server address	The IP address of the most-recently contacted TFTP server, if the switch has contacted a TFTP server since the last time the software was reloaded or the switch was rebooted.
Configuration filename	The name under which the Layer 2 Switch startup-config file was uploaded or downloaded during the most recent TFTP access.
Image filename	The name of the Layer 2 Switch flash image (system software file) that was uploaded or downloaded during the most recent TFTP access.

Displaying ARP entries

To display the entries the Layer 2 Switch has placed in its ARP cache, enter the following command from any level of the CLI.

```
PowerConnect# show arp

IP           Mac           Port Age VlanId
192.168.1.170 0010.5a11.d042 7 0 1
Total Arp Entries : 1
```

Syntax: show arp

This display shows the following information.

TABLE 109 CLI display of ARP cache

This field...	Displays...
IP	The IP address of the device.
Mac	The MAC address of the device. NOTE: If the MAC address is all zeros, the entry is for the default gateway, but the Layer 2 Switch does not have a link to the gateway.
Port	The port on which the entry was learned.
Age	The number of minutes the entry has remained unused. If this value reaches the ARP aging period, the entry is removed from the cache.
VlanId	The VLAN the port that learned the entry is in. NOTE: If the MAC address is all zeros, this field shows a random VLAN ID, since the Layer 2 Switch does not yet know which port the device for this entry is attached to.
Total ARP Entries	The number of entries in the ARP cache.

Displaying IP traffic statistics

To display IP traffic statistics on a Layer 2 Switch, enter the following command at any CLI level.

```
PowerConnect# show ip traffic

IP Statistics
 27 received, 24 sent
 0 fragmented, 0 reassembled, 0 bad header
 0 no route, 0 unknown proto, 0 no buffer, 0 other errors

ICMP Statistics
Received:
 0 total, 0 errors, 0 unreachable, 0 time exceed
 0 parameter, 0 source quench, 0 redirect, 0 echo,
 0 echo reply, 0 timestamp, 0 timestamp rely, 0 addr mask
 0 addr mask reply, 0 irdp advertisement, 0 irdp solicitation
Sent:
 0 total, 0 errors, 0 unreachable, 0 time exceed
 0 parameter, 0 source quench, 0 redirect, 0 echo,
 0 echo reply, 0 timestamp, 0 timestamp rely, 0 addr mask
 0 addr mask reply, 0 irdp advertisement, 0 irdp solicitation

UDP Statistics
 0 received, 0 sent, 0 no port, 0 input errors

TCP Statistics
 1 current active tcbs, 4 tcbs allocated, 0 tcbs freed 0 tcbs protected
 0 active opens, 0 passive opens, 0 failed attempts
 0 active resets, 0 passive resets, 0 input errors
 27 in segments, 24 out segments, 0 retransmission
```

Syntax: show ip traffic

The **show ip traffic** command displays the following information.

TABLE 110 CLI display of IP traffic statistics – Layer 2 Switch

This field...	Displays...
IP statistics	
received	The total number of IP packets received by the device.
sent	The total number of IP packets originated and sent by the device.
fragmented	The total number of IP packets fragmented by this device to accommodate the MTU of this device or of another device.
reassembled	The total number of fragmented IP packets that this device re-assembled.
bad header	The number of IP packets dropped by the device due to a bad packet header.
no route	The number of packets dropped by the device because there was no route.
unknown proto	The number of packets dropped by the device because the value in the Protocol field of the packet header is unrecognized by this device.
no buffer	This information is used by Dell customer support.
other errors	The number of packets that this device dropped due to error types other than the types listed above.
ICMP statistics	
The ICMP statistics are derived from RFC 792, "Internet Control Message Protocol", RFC 950, "Internet Standard Subnetting Procedure", and RFC 1256, "ICMP Router Discovery Messages". Statistics are organized into Sent and Received. The field descriptions below apply to each.	
total	The total number of ICMP messages sent or received by the device.
errors	This information is used by Dell customer support.
unreachable	The number of Destination Unreachable messages sent or received by the device.
time exceed	The number of Time Exceeded messages sent or received by the device.
parameter	The number of Parameter Problem messages sent or received by the device.
source quench	The number of Source Quench messages sent or received by the device.
redirect	The number of Redirect messages sent or received by the device.
echo	The number of Echo messages sent or received by the device.
echo reply	The number of Echo Reply messages sent or received by the device.
timestamp	The number of Timestamp messages sent or received by the device.
timestamp reply	The number of Timestamp Reply messages sent or received by the device.
addr mask	The number of Address Mask Request messages sent or received by the device.
addr mask reply	The number of Address Mask Replies messages sent or received by the device.
irdp advertisement	The number of ICMP Router Discovery Protocol (IRDP) Advertisement messages sent or received by the device.
irdp solicitation	The number of IRDP Solicitation messages sent or received by the device.
UDP statistics	
received	The number of UDP packets received by the device.
sent	The number of UDP packets sent by the device.
no port	The number of UDP packets dropped because the packet did not contain a valid UDP port number.

TABLE 110 CLI display of IP traffic statistics – Layer 2 Switch (Continued)

This field...	Displays...
input errors	This information is used by Dell customer support.
TCP statistics	
The TCP statistics are derived from RFC 793, "Transmission Control Protocol".	
current active tcbs	The number of TCP Control Blocks (TCBs) that are currently active.
tcbs allocated	The number of TCBs that have been allocated.
tcbs freed	The number of TCBs that have been freed.
tcbs protected	This information is used by Dell customer support.
active opens	The number of TCP connections opened by this device by sending a TCP SYN to another device.
passive opens	The number of TCP connections opened by this device in response to connection requests (TCP SYNs) received from other devices.
failed attempts	This information is used by Dell customer support.
active resets	The number of TCP connections this device reset by sending a TCP RESET message to the device at the other end of the connection.
passive resets	The number of TCP connections this device reset because the device at the other end of the connection sent a TCP RESET message.
input errors	This information is used by Dell customer support.
in segments	The number of TCP segments received by the device.
out segments	The number of TCP segments sent by the device.
retransmission	The number of segments that this device retransmitted because the retransmission timer for the segment had expired before the device at the other end of the connection had acknowledged receipt of the segment.

21 Displaying IP configuration information and statistics

Configuring RIP

RIP overview

Routing Information Protocol (RIP) is an IP route exchange protocol that uses a **distance vector** (a number representing distance) to measure the cost of a given route. The **cost** is a distance vector because the cost often is equivalent to the number of router hops between the Layer 3 Switch and the destination network.

A Layer 3 Switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the IP route table as the route to the destination. Typically, the best path is the path with the fewest hops. A hop is another router through which packets must travel to reach the destination. If the Layer 3 Switch receives a RIP update from another router that contains a path with fewer hops than the path stored in the Layer 3 Switch route table, the Layer 3 Switch replaces the older route with the newer one. The Layer 3 Switch then includes the new path in the updates it sends to other RIP routers, including Layer 3 Switches.

RIP routers, including the Layer 3 Switch, also can modify a route cost, generally by adding to it, to bias the selection of a route for a given destination. In this case, the actual number of router hops may be the same, but the route has an administratively higher cost and is thus less likely to be used than other, lower-cost routes.

A RIP route can have a maximum cost of 15. Any destination with a higher cost is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

Layer 3 Switches support the following RIP versions:

- Version 1
- V1 compatible with V2
- Version 2 (the default)

ICMP host unreachable message for undeliverable ARPs

If the router receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (router knows of no route to the destination address), the router sends an ICMP Host Unreachable message to the source.

RIP parameters and defaults

The following tables list the RIP parameters, their default values, and where to find configuration information.

RIP global parameters

Table 111 lists the global RIP parameters and their default values, and indicates where you can find configuration information.

TABLE 111 RIP global parameters

Parameter	Description	Default	See page...
RIP state	The global state of the protocol NOTE: You also must enable the protocol on individual interfaces. Globally enabling the protocol does not allow interfaces to send and receive RIP information. Refer to Table 112 on page 645.	Disabled	page 645
Administrative distance	The administrative distance is a numeric value assigned to each type of route on the router. When the router is selecting from among multiple routes (sometimes of different origins) to the same destination, the router compares the administrative distances of the routes and selects the route with the lowest administrative distance. This parameter applies to routes originated by RIP. The administrative distance stays with a route when it is redistributed into other routing protocols.	120	page 647
Redistribution	RIP can redistribute routes from other routing protocols such as OSPF and BGP4 into RIP. A redistributed route is one that a router learns through another protocol, then distributes into RIP.	Disabled	page 647
Redistribution metric	RIP assigns a RIP metric (cost) to each external route redistributed from another routing protocol into RIP. An external route is a route with at least one hop (packets must travel through at least one other router to reach the destination). This parameter applies to routes that are redistributed from other protocols into RIP.	1 (one)	page 649
Update interval	How often the router sends route updates to its RIP neighbors	30 seconds	page 650
Learning default routes	The router can learn default routes from its RIP neighbors. NOTE: You also can enable or disable this parameter on an individual interface basis. Refer to Table 112 on page 645.	Disabled	page 650
Advertising and learning with specific neighbors	The Layer 3 Switch learns and advertises RIP routes with all its neighbors by default. You can prevent the Layer 3 Switch from advertising routes to specific neighbors or learning routes from specific neighbors.	Learning and advertising permitted for all neighbors	page 651

RIP interface parameters

Table 112 lists the interface-level RIP parameters and their default values, and indicates where you can find configuration information.

TABLE 112 RIP interface parameters

Parameter	Description	Default	See page...
RIP state and version	The state of the protocol and the version that is supported on the interface. The version can be one of the following: <ul style="list-style-type: none"> • Version 1 only • Version 2 only • Version 1, but also compatible with version 2 NOTE: You also must enable RIP globally.	Disabled	page 645
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	page 646
Learning default routes	Locally overrides the global setting. Refer to Table 111 on page 644.	Disabled	page 650
Loop prevention	The method a router uses to prevent routing loops caused by advertising a route on the same interface as the one on which the router learned the route. <ul style="list-style-type: none"> • Split horizon – The router does not advertise a route on the same interface as the one on which the router learned the route. • Poison reverse – The router assigns a cost of 16 (“infinite” or “unreachable”) to a route before advertising it on the same interface as the one on which the router learned the route. 	Poison reverse NOTE: Enabling split horizon disables poison reverse on the interface.	page 651
Advertising and learning specific routes	You can control the routes that a Layer 3 Switch learns or advertises.	The Layer 3 Switch learns and advertises all RIP routes on all interfaces.	page 652

Configuring RIP parameters

Use the following procedures to configure RIP parameters on a system-wide and individual interface basis.

Enabling RIP

RIP is disabled by default. To enable it, use the following method.

NOTE

You must enable the protocol globally and also on individual interfaces on which you want to advertise RIP. Globally enabling the protocol does not enable it on individual interfaces.

To enable RIP globally, enter the following command.

```
PowerConnect(config)# router rip
```

Syntax: [no] router rip

After globally enabling the protocol, you must enable it on individual interfaces. You can enable the protocol on physical interfaces as well as virtual routing interfaces. To enable RIP on an interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# ip rip v1-only
```

Syntax: [no] ip rip v1-only | v1-compatible-v2 | v2-only

NOTE

You must specify the RIP version.

Configuring metric parameters

By default, a Layer 3 Switch port increases the cost of a RIP route that is learned on the port by one. You can configure individual ports to add more than one to a learned route cost. In addition, you can configure a RIP offset list to increase the metric for learned or advertised routes based on network address.

Changing the cost of routes learned on a port

By default, a Layer 3 Switch port increases the cost of a RIP route that is learned on the port. The Layer 3 Switch increases the cost by adding one to the route metric before storing the route.

You can change the amount that an individual port adds to the metric of RIP routes learned on the port. To do so, use the following method.

NOTE

RIP considers a route with a metric of 16 to be unreachable. Use this metric only if you do not want the route to be used. In fact, you can prevent the Layer 3 Switch from using a specific port for routes learned through that port by setting its metric to 16.

To increase the cost a port adds to RIP routes learned in that port, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# ip metric 5
```

These commands configure port 1 to add 5 to the cost of each route learned on the port.

Syntax: ip metric <1-16>

Configuring a RIP offset list

A RIP offset list allows you to add to the metric of specific inbound or outbound routes learned or advertised by RIP. RIP offset lists provide a simple method for adding to the cost of specific routes and therefore biasing the Layer 3 Switch route selection away from those routes.

An offset list consists of the following parameters:

- An ACL that specifies the routes to which to add the metric.
- The direction:
 - In applies to routes the Layer 3 Switch learns from RIP neighbors.
 - Out applies to routes the Layer 3 Switch is advertising to its RIP neighbors.
- The type and number of a specific port to which the offset list applies (optional).

The software adds the offset value to the routing metric (cost) of the routes that match the ACL. If a route matches both a global offset list and an interface-based offset list, the interface-based offset list takes precedence. The interface-based offset list metric is added to the route in this case.

You can configure up to 24 global RIP offset lists and up to 24 RIP offset lists on each interface.

To configure a global RIP offset list, enter commands such as the following.

```
PowerConnect(config)# access-list 21 deny 160.1.0.0 0.0.255.255
PowerConnect(config)# access-list 21 permit any
PowerConnect(config)# router rip
PowerConnect(config-rip-router)# offset-list 21 out 10
```

The commands in this example configure a standard ACL. The ACL matches on all IP networks except 160.1.x.x. When the Layer 3 Switch advertises a route that matches ACL 21, the offset list adds 10 to the route metric.

Syntax: [no] **offset-list** <ACL-number-or-name> **in** | **out offset** [**ethernet** <portnum>]

In the following example, the Layer 3 Switch uses ACL 21 to add 10 to the metric of routes received on Ethernet port 1.

```
PowerConnect(config-rip-router)# offset-list 21 in 10 ethernet 1
```

Changing the administrative distance

By default, the Layer 3 Switch assigns the default RIP administrative distance (120) to RIP routes. When comparing routes based on administrative distance, the Layer 3 Switch selects the route with the lower distance. You can change the administrative distance for RIP routes.

NOTE

Refer to [“Changing administrative distances”](#) on page 777 for the default distances for all route sources.

To change the administrative distance for RIP routes, enter a command such as the following.

```
PowerConnect(config-rip-router)# distance 140
```

This command changes the administrative distance to 140 for all RIP routes.

Syntax: [no] **distance** <num>

Configuring redistribution

You can configure the Layer 3 Switch to redistribute routes learned through Open Shortest Path First (OSPF) or Border Gateway Protocol version 4 (BGP4) into RIP. When you redistribute a route from one of these other protocols into RIP, the Layer 3 Switch can use RIP to advertise the route to its RIP neighbors.

To configure redistribution, perform the following tasks:

- Configure redistribution filters (optional). You can configure filters to permit or deny redistribution for a route based on its origin (OSPF, BGP4, and so on), the destination network address, and the route metric. You also can configure a filter to set the metric based on these criteria.

- Change the default redistribution metric (optional). The Layer 3 Switch assigns a RIP metric of one to each redistributed route by default. You can change the default metric to a value up to 16.
- Enable redistribution.

NOTE

Do not enable redistribution until you configure the other redistribution parameters.

Configuring redistribution filters

RIP redistribution filters apply to all interfaces. The software uses the filters in ascending numerical order and immediately takes the action specified by the filter. Thus, if filter 1 denies redistribution of a given route, the software does not redistribute the route, regardless of whether a filter with a higher ID would permit redistribution of that route.

NOTE

The default redistribution action is permit, even after you configure and apply redistribution filters to the virtual routing interface. If you want to tightly control redistribution, apply a filter to deny all routes as the last filter (the filter with the highest ID), then apply filters with lower filter IDs to allow specific routes.

To configure a redistribution filter, enter a command such as the following.

```
PowerConnect(config-rip-router)# deny redistribute 2 all address 207.92.0.0
255.255.0.0
```

This command denies redistribution for all types of routes to the 207.92.x.x network.

Syntax: [no] permit | deny redistribute <filter-num> all | bgp | ospf | static address <ip-addr>
<ip-mask>
[match-metric <value> | set-metric <value>]

The <filter-num> specifies the redistribution filter ID. The software uses the filters in ascending numerical order. Thus, if filter 1 denies a route from being redistributed, the software does not redistribute that route even if a filter with a higher ID permits redistribution of the route.

The **all** parameter applies redistribution to all route types.

The **bgp** parameter applies redistribution to BGP4 routes only.

The **ospf** parameter applies redistribution to OSPF routes only.

The **static** parameter applies redistribution to IP static routes only.

The **address <ip-addr> <ip-mask>** parameters apply redistribution to the specified network and subnet address. Use 0 to specify “any”. For example, “207.92.0.0 255.255.0.0” means “any 207.92.x.x subnet”. However, to specify any subnet (all subnets match the filter), enter “address 255.255.255.255 255.255.255.255”.

The **match-metric <value>** parameter applies the redistribution filter only to those routes with the specified metric value; possible values are from 1 – 15.

The **set-metric <value>** parameter sets the RIP metric value that will be applied to those routes imported into RIP.

The following command denies redistribution into RIP for all OSPF routes.

```
PowerConnect(config-rip-router)# deny redistribute 3 ospf address 207.92.0.0
255.255.0.0
```

The following command denies redistribution for all OSPF routes that have a metric of 10.

```
PowerConnect(config-rip-router)# deny redistribute 3 ospf address 207.92.0.0  
255.255.0.0 match-metric 10
```

The following commands deny redistribution of all routes except routes for 10.10.10.x and 20.20.20.x.

```
PowerConnect(config-rip-router)# deny redistribute 64 static address  
255.255.255.255 255.255.255.255  
PowerConnect(config-rip-router)# permit redistribute 1 static address 10.10.10.0  
255.255.255.0  
PowerConnect(config-rip-router)# permit redistribute 2 static address 20.20.20.0  
255.255.255.0
```

NOTE

This example assumes that the highest RIP redistribution filter ID configured on the device is 64.

Changing the redistribution metric

When the Layer 3 Switch redistributes a route into RIP, the software assigns a RIP metric (cost) to the route. By default, the software assigns a metric of one to each route that is redistributed into RIP. You can increase the metric that the Layer 3 Switch assigns, up to 15.

To change the RIP metric the Layer 3 Switch assigns to redistributed routes, enter a command such as the following.

```
PowerConnect(config-rip-router)# default-metric 10
```

This command assigns a RIP metric of 10 to each route that is redistributed into RIP.

Syntax: [no] default-metric <1-15>

Enabling redistribution

After you configure redistribution parameters, you need to enable redistribution.

To enable RIP redistribution, enter the following command.

```
PowerConnect(config-rip-router)# redistribution
```

Syntax: [no] redistribution

Removing a RIP redistribution deny filter

To remove a previously configured RIP redistribution deny filter, do the following.

1. Remove the RIP redistribution deny filter.
2. Disable the redistribution function.
3. Re-enable redistribution.

The following shows an example of how to remove a RIP redistribution deny filter.

```
PowerConnect(config-rip-router)# no deny redistribute 2 all address
207.92.0.0 255.255.0.0
PowerConnect(config-rip-router)# no redistribution
PowerConnect(config-rip-router)# redistribution
```

Configuring route learning and advertising parameters

By default, a Layer 3 Switch learns routes from all its RIP neighbors and advertises RIP routes to those neighbors.

You can configure the following learning and advertising parameters:

- Update interval – The update interval specifies how often the Layer 3 Switch sends RIP route advertisements to its neighbors. The default is 30 seconds. You can change the interval to a value from 1 – 1000 seconds.
- Learning and advertising of RIP default routes – The Layer 3 Switch learns and advertises RIP default routes by default. You can disable learning and advertising of default routes on a global or individual interface basis.
- Learning of standard RIP routes – By default, the Layer 3 Switch can learn RIP routes from all its RIP neighbors. You can configure RIP neighbor filters to explicitly permit or deny learning from specific neighbors.

Changing the update interval for route advertisements

The update interval specifies how often the Layer 3 Switch sends route advertisements to its RIP neighbors. You can specify an interval from 1 – 1000 seconds. The default is 30 seconds.

To change the RIP update interval, enter a command such as the following.

```
PowerConnect(config-rip-router)# update 120
```

This command configures the Layer 3 Switch to send RIP updates every 120 seconds.

Syntax: update-time <1-1000>

Enabling learning of RIP default routes

By default, the Layer 3 Switch does not learn RIP default routes. You can enable learning of RIP default routes on a global or interface basis.

To enable learning of default RIP routes on a global basis, enter the following command.

```
PowerConnect(config-rip-router)# learn-default
```

Syntax: [no] learn-default

To enable learning of default RIP routes on an interface basis, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# ip rip learn-default
```

Syntax: [no] ip rip learn-default

Configuring a RIP neighbor filter

By default, a Layer 3 Switch learns RIP routes from all its RIP neighbors. Neighbor filters allow you to specify the neighbor routers from which the device can receive RIP routes. Neighbor filters apply globally to all ports.

To configure a RIP neighbor filters, enter a command such as the following.

```
PowerConnect(config-rip-router)# neighbor 1 deny any
```

Syntax: [no] neighbor <filter-num> permit | deny <source-ip-address> | any

This command configures the Layer 3 Switch so that the device does not learn any RIP routes from any RIP neighbors.

The following commands configure the Layer 3 Switch to learn routes from all neighbors except 192.168.1.170. Once you define a RIP neighbor filter, the default action changes from learning all routes from all neighbors to denying all routes from all neighbors except the ones you explicitly permit. To deny learning from a specific neighbor but allow all other neighbors, you must add a filter that allows learning from all neighbors. Be sure to add the filter to permit all neighbors last (the one with the highest filter number). Otherwise, the software can match on the permit all filter instead of a filter that denies a specific neighbor, and learn routes from that neighbor.

```
PowerConnect(config-rip-router)# neighbor 2 deny 192.16.1.170
PowerConnect(config-rip-router)# neighbor 1024 permit any
```

Changing the route loop prevention method

RIP can use the following methods to prevent routing loops:

- Split horizon – The Layer 3 Switch does not advertise a route on the same interface as the one on which the router learned the route.
- Poison reverse – The Layer 3 Switch assigns a cost of 16 (“infinite” or “unreachable”) to a route before advertising it on the same interface as the one on which the router learned the route. This is the default.

These loop prevention methods are configurable on an individual interface basis. One of the methods is always in effect on an interface enabled for RIP. If you disable one method, the other method is enabled.

NOTE

These methods may be used in addition to RIP maximum valid route cost of 15.

To disable poison reverse and enable split horizon on an interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# no ip rip poison-reverse
```

Syntax: [no] ip rip poison-reverse

To disable split horizon and enable poison reverse on an interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-1)# ip rip poison-reverse
```

Suppressing RIP route advertisement on a VRRP or VRRPE backup interface

NOTE

This section applies only if you configure the Layer 3 Switch for Virtual Router Redundancy Protocol (VRRP) or VRRP Extended (VRRPE). Refer to [Chapter 24, “Configuring VRRP and VRRPE”](#).

Normally, a VRRP or VRRPE Backup includes route information for the virtual IP address (the backed up interface) in RIP advertisements. As a result, other routers receive multiple paths for the backed up interface and might sometimes unsuccessfully use the path to the Backup rather than the path to the Master.

You can prevent the Backups from advertising route information for the backed up interface by enabling suppression of the advertisements.

To suppress RIP advertisements for the backed up interface, enter the following commands.

```
PowerConnect(config)# router rip
PowerConnect(config-rip-router)# use-vrrp-path
```

Syntax: [no] use-vrrp-path

The syntax is the same for VRRP and VRRPE.

Configuring RIP route filters

You can configure RIP route filters to permit or deny learning or advertising of specific routes. Configure the filters globally, then apply them to individual interfaces. When you apply a RIP route filter to an interface, you specify whether the filter applies to learned routes (in) or advertised routes (out).

NOTE

A route is defined by the destination IP address and network mask.

NOTE

By default, routes that do not match a route filter are learned or advertised. To prevent a route from being learned or advertised, you must configure a filter to deny the route.

To configure RIP filters, enter commands such as the following.

```
PowerConnect(config-rip-router)# filter 1 permit 192.53.4.1 255.255.255.0
PowerConnect(config-rip-router)# filter 2 permit 192.53.5.1 255.255.255.0
PowerConnect(config-rip-router)# filter 3 permit 192.53.6.1 255.255.255.0
PowerConnect(config-rip-router)# filter 4 deny 192.53.7.1 255.255.255.0
```

These commands explicitly permit RIP routes to three networks, and deny the route to one network.

Since the default action is permit, all other routes (routes not explicitly permitted or denied by the filters) can be learned or advertised.

Syntax: filter <filter-num> permit | deny <source-ip-address> | any <source-mask> | any [log]

Applying a RIP route filter to an interface

Once you define RIP route filters, you must assign them to individual interfaces. The filters do not take effect until you apply them to interfaces. When you apply a RIP route filter, you also specify whether the filter applies to learned routes or advertised routes:

- Out filters apply to routes the Layer 3 Switch advertises to its neighbor on the interface.
- In filters apply to routes the Layer 3 Switch learns from its neighbor on the interface.

To apply RIP route filters to an interface, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 2
PowerConnect(config-if-2)# ip rip filter-group in 2 3 4
```

Syntax: [no] ip rip filter-group in | out <filter-list>

These commands apply RIP route filters 2, 3, and 4 to all routes learned from the RIP neighbor on port 2.

Displaying RIP filters

To display the RIP filters configured on the router, enter the following command at any CLI level.

```
PowerConnect# show ip rip

          RIP Route Filter Table
Index  Action  Route IP Address  Subnet Mask
  1     deny   any               any

          RIP Neighbor Filter Table
Index  Action  Neighbor IP Address
  1     permit any
```

Syntax: show ip rip

This display shows the following information.

TABLE 113 CLI display of RIP filter information

This field...	Displays...
Route filters	
The rows underneath "RIP Route Filter Table" list the RIP route filters. If no RIP route filters are configured on the device, the following message is displayed instead: "No Filters are configured in RIP Route Filter Table".	
Index	The filter number. You assign this number when you configure the filter.
Action	The action the router takes if a RIP route packet matches the IP address and subnet mask of the filter. The action can be one of the following: <ul style="list-style-type: none"> • deny – RIP route packets that match the address and network mask information in the filter are dropped. If applied to an interface outbound filter group, the filter prevents the router from advertising the route on that interface. If applied to an interface inbound filter group, the filter prevents the router from adding the route to its IP route table. • permit – RIP route packets that match the address and network mask information are accepted. If applied to an interface outbound filter group, the filter allows the router to advertise the route on that interface. If applied to an interface inbound filter group, the filter allows the router to add the route to its IP route table.
Route IP Address	The IP address of the route destination network or host.

TABLE 113 CLI display of RIP filter information (Continued)

This field...	Displays...
Subnet Mask	The network mask for the IP address.
Neighbor filters	
The rows underneath "RIP Neighbor Filter Table" list the RIP neighbor filters. If no RIP neighbor filters are configured on the device, the following message is displayed instead: "No Filters are configured in RIP Neighbor Filter Table".	
Index	The filter number. You assign this number when you configure the filter.
Action	The action the router takes for RIP route packets to or from the specified neighbor: <ul style="list-style-type: none"> deny – If the filter is applied to an interface outbound filter group, the filter prevents the router from advertising RIP routes to the specified neighbor on that interface. If the filter is applied to an interface inbound filter group, the filter prevents the router from receiving RIP updates from the specified neighbor. permit – If the filter is applied to an interface outbound filter group, the filter allows the router to advertise RIP routes to the specified neighbor on that interface. If the filter is applied to an interface inbound filter group, the filter allows the router to receive RIP updates from the specified neighbor.
Neighbor IP Address	The IP address of the RIP neighbor.

Displaying CPU utilization statistics

You can display CPU utilization statistics for RIP and other IP protocols.

To display CPU utilization statistics for RIP for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.

```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.03       0.09       0.22        9
BGP             0.04       0.06       0.08       0.14       13
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.00       0.00       0.00       0.00        0
IP             0.00       0.00       0.00       0.00        0
OSPF           0.00       0.00       0.00       0.00        0
RIP          0.04     0.07     0.08     0.09       7
STP            0.00       0.00       0.00       0.00        0
VRRP           0.00       0.00       0.00       0.00        0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example.

```
PowerConnect#show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP              0.01        0.00        0.00        0.00         0
BGP              0.00        0.00        0.00        0.00         0
GVRP            0.00        0.00        0.00        0.00         0
ICMP            0.01        0.00        0.00        0.00         1
IP              0.00        0.00        0.00        0.00         0
OSPF            0.00        0.00        0.00        0.00         0
RIP            0.00        0.00        0.00        0.00         0
STP            0.00        0.00        0.00        0.00         0
VRRP           0.00        0.00        0.00        0.00         0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.

```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name    Sec(%)    Time(ms)
ARP             0.00        0
BGP            0.00        0
GVRP          0.00        0
ICMP          0.01        1
IP            0.00        0
OSPF          0.00        0
RIP           0.00        0
STP           0.01        0
VRRP          0.00        0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The `<num>` parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

22 Displaying CPU utilization statistics

Configuring OSPF Version 2 (IPv4)

This chapter describes how to configure OSPF Version 2 on Layer 3 Switches using the CLI. OSPF Version 2 is supported on devices running IPv4.

NOTE

The terms *Layer 3 Switch* and *router* are used interchangeably in this chapter and mean the same thing.

Overview of OSPF

OSPF is a link-state routing protocol. The protocol uses link-state advertisements (LSA) to update neighboring routers regarding its interfaces and information on those interfaces. The router floods these LSAs to all neighboring routers to update them regarding the interfaces. Each router maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.

Layer 3 Switches support the following types of LSAs, which are described in RFC 1583:

- Router link
- Network link
- Summary link
- Autonomous system (AS) summary link
- AS external link
- Not-So-Stubby Area (NSSA) external link

OSPF is built upon a hierarchy of network components. The highest level of the hierarchy is the **Autonomous System (AS)**. An autonomous system is defined as a number of networks, all of which share the same routing and administration characteristics.

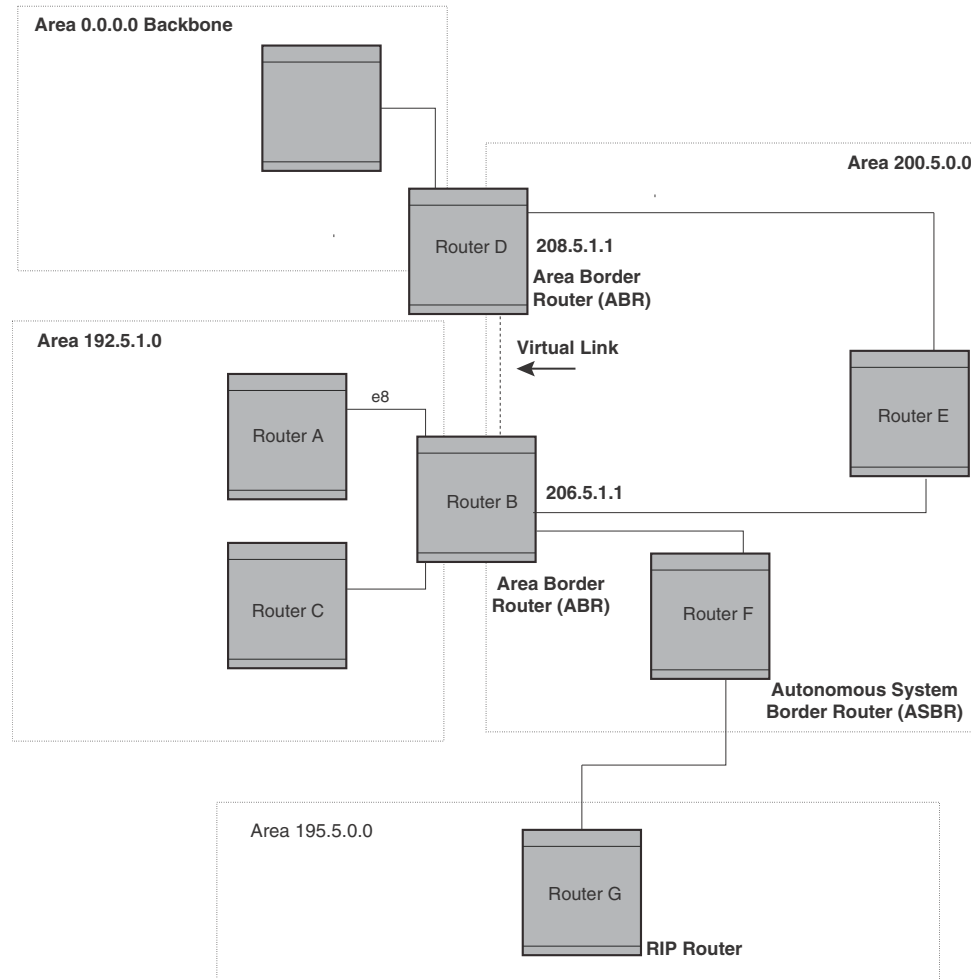
An AS can be divided into multiple **areas** as shown in [Figure 105](#) on page 658. Each area represents a collection of contiguous networks and hosts. Areas limit the area to which link-state advertisements are broadcast, thereby limiting the amount of flooding that occurs within the network. An area is represented in OSPF by either an IP address or a number.

You can further limit the broadcast area of flooding by defining an area range. The area range allows you to assign an aggregate value to a range of IP addresses. This aggregate value becomes the address that is advertised instead all of the individual addresses it represents being advertised. You can assign up to 32 ranges in an OSPF area.

An OSPF router can be a member of multiple areas. Routers with membership in multiple areas are known as **Area Border Routers (ABRs)**. Each ABR maintains a separate topological database for each area the router is in. Each topological database contains all of the LSA databases for each router within a given area. The routers within the same area have identical topological databases. The ABR is responsible for forwarding routing information or changes between its border areas.

An **Autonomous System Boundary Router (ASBR)** is a router that is running multiple protocols and serves as a gateway to routers outside an area and those operating with different protocols. The ASBR is able to import and translate different protocol routes into OSPF through a process known as **redistribution**. For more details on redistribution and configuration examples, refer to “[Enable route redistribution](#)” on page 686.

FIGURE 105 OSPF operating in a network



OSPF point-to-point Links

One important OSPF process is **Adjacency**. Adjacency occurs when a relationship is formed between neighboring routers for the purpose of exchanging routing information. Adjacent OSPF neighbor routers go beyond the simple Hello packet exchange; they exchange database information. In order to minimize the amount of information exchanged on a particular segment, one of the first steps in creating adjacency is to assign a Designated Router (DR) and a Backup Designated Router (BDR). The Designated Router ensures that there is a central point of contact, thereby improving convergence time within a multi-access segment.

In an OSPF point-to-point network, where a direct Layer 3 connection exists between a single pair of OSPF routers, there is no need for Designated and Backup Designated Routers, as is the case in OSPF multi-access networks. Without the need for Designated and Backup Designated routers, a point-to-point network establishes adjacency and converges faster. The neighboring routers become adjacent whenever they can communicate directly. In contrast, in broadcast and non-broadcast multi-access (NBMA) networks, the Designated Router and Backup Designated Router become adjacent to all other routers attached to the network.

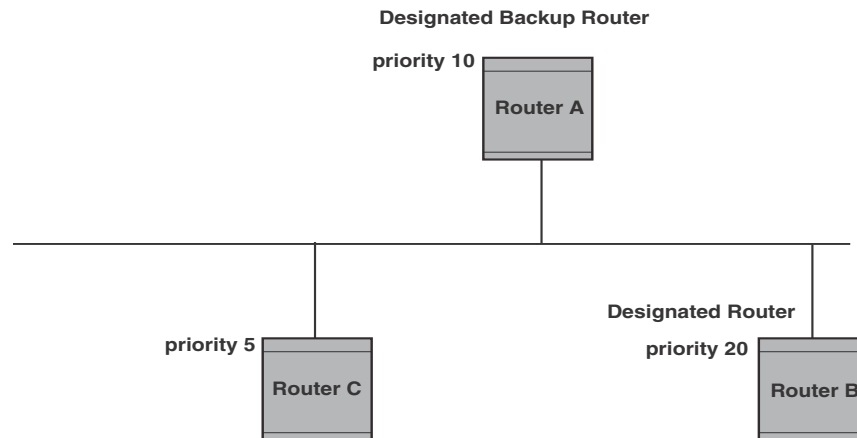
Designated routers in multi-access networks

In a network that has multiple routers attached, OSPF elects one router to serve as the designated router (DR) and another router on the segment to act as the backup designated router (BDR). This arrangement minimizes the amount of repetitive information that is forwarded on the network by forwarding all messages to the designated router and backup designated routers responsible for forwarding the updates throughout the network.

Designated router election in multi-access networks

In a network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next largest priority is elected as the BDR, as shown in [Figure 106](#)

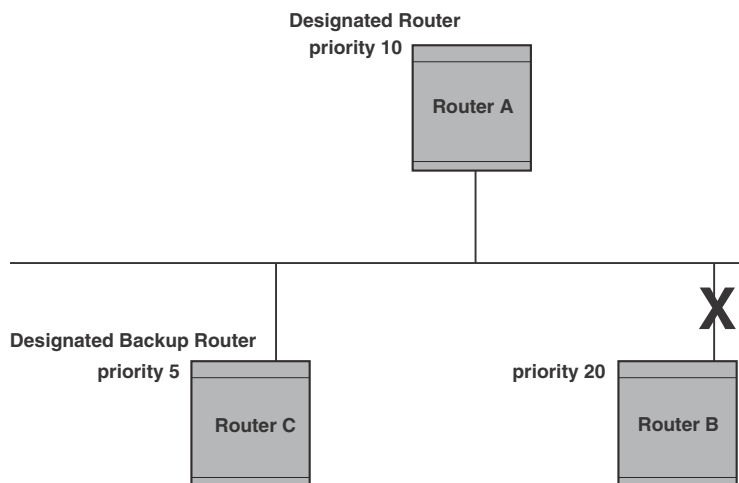
FIGURE 106 Designated and backup router election



If the DR goes off-line, the BDR automatically becomes the DR. The router with the next highest priority becomes the new BDR. This process is shown in [Figure 107](#).

NOTE

Priority is a configurable option at the interface level. You can use this parameter to help bias one router as the DR.

FIGURE 107 Backup designated router becomes designated router

If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR.

NOTE

By default, the router ID is the IP address configured on the lowest numbered loopback interface. If the Layer 3 Switch does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, refer to [“Changing the router ID”](#) on page 584.

When multiple routers on the same network are declaring themselves as DRs, then both priority and router ID are used to select the designated router and backup designated routers.

When only one router on the network claims the DR role despite neighboring routers with higher priorities or router IDs, this router remains the DR. This is also true for BDRs.

The DR and BDR election process is performed when one of the following events occurs:

- An interface is in a waiting state and the wait time expires
- An interface is in a waiting state and a hello packet is received that addresses the BDR
- A change in the neighbor state occurs, such as:
 - A neighbor state transitions from 2 or higher
 - Communication to a neighbor is lost
 - A neighbor declares itself to be the DR or BDR for the first time

OSPF RFC 1583 and 2178 compliance

Routers are configured, by default, to be compliant with the RFC 1583 OSPF V2 specification. Routers can also be configured to operate with the latest OSPF standard, RFC 2178.

NOTE

For details on how to configure the system to operate with the RFC 2178, refer to [“Modify OSPF standard compliance setting”](#) on page 694.

Reduction of equivalent AS External LSAs

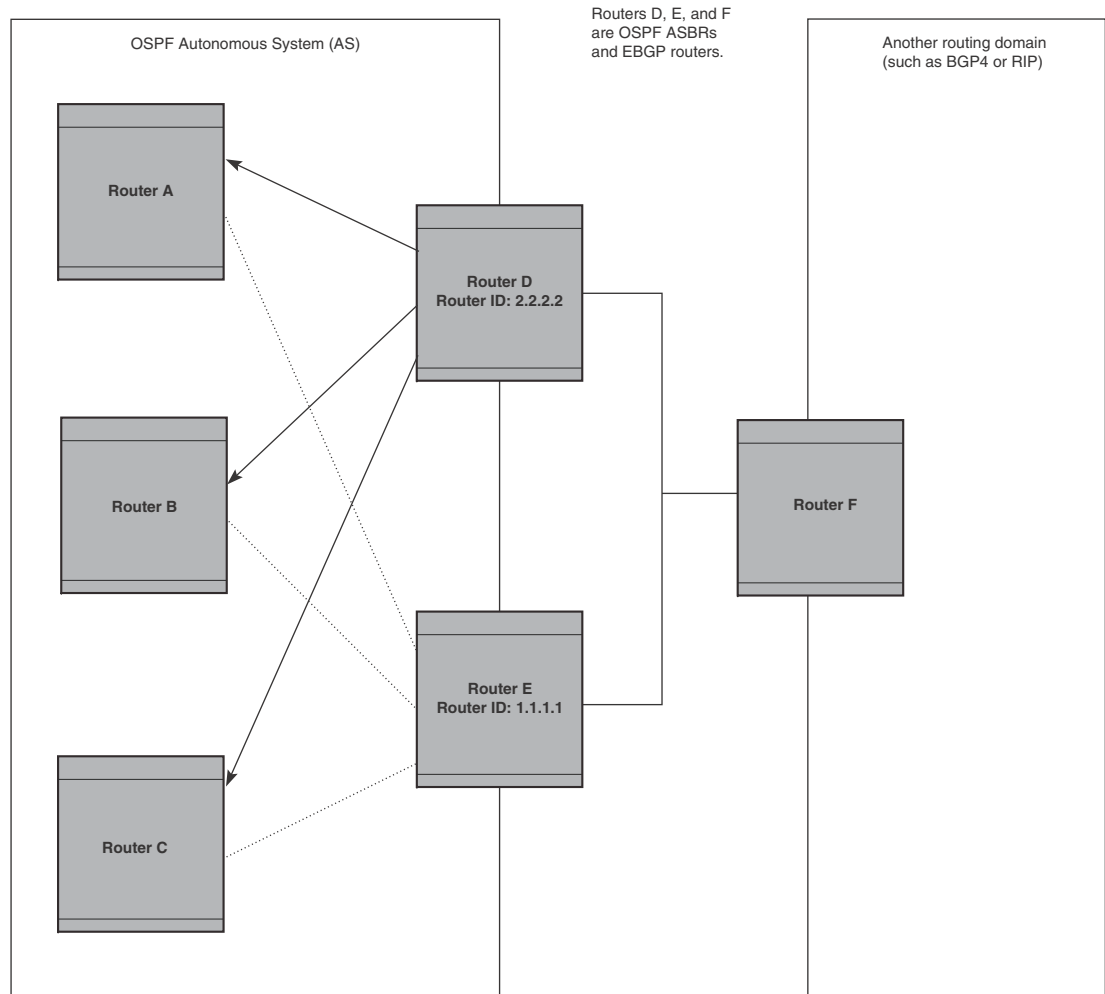
An OSPF ASBR uses AS External link advertisements (AS External LSAs) to originate advertisements of a route to another routing domain, such as a BGP4 or RIP domain. The ASBR advertises the route to the external domain by flooding AS External LSAs to all the other OSPF routers (except those inside stub networks) within the local OSPF Autonomous System (AS).

In some cases, multiple ASBRs in an AS can originate equivalent LSAs. The LSAs are equivalent when they have the same cost, the same next hop, and the same destination. devices optimize OSPF by eliminating duplicate AS External LSAs in this case. The Layer 3 Switch with the lower router ID flushes the duplicate External LSAs from its database and thus does not flood the duplicate External LSAs into the OSPF AS. AS External LSA reduction therefore reduces the size of the Layer 3 Switch link state database.

This enhancement implements the portion of RFC 2328 that describes AS External LSA reduction. This enhancement is enabled by default, requires no configuration, and cannot be disabled.

Figure 108 shows an example of the AS External LSA reduction feature. In this example, Layer 3 Switches D and E are OSPF ASBRs, and thus communicate route information between the OSPF AS, which contains Routers A, B, and C, and another routing domain, which contains Router F. The other routing domain is running another routing protocol, such as BGP4 or RIP. Routers D, E, and F, therefore, are each running both OSPF and either BGP4 or RIP.

FIGURE 108 AS External LSA reduction



Notice that both Router D and Router E have a route to the other routing domain through Router F. In earlier, if Routers D and E have equal-cost routes to Router F, then both Router D and Router E flood AS External LSAs to Routers A, B, and C advertising the route to Router F. Since both routers are flooding equivalent routes, Routers A, B, and C receive multiple routes with the same cost to the same destination (Router F). For Routers A, B, and C, either route to Router F (through Router D or through Router E) is equally good.

OSPF eliminates the duplicate AS External LSAs. When two or more Layer 3 Switches configured as ASBRs have equal-cost routes to the same next-hop router in an external routing domain, the ASBR with the highest router ID floods the AS External LSAs for the external domain into the OSPF AS, while the other ASBRs flush the equivalent AS External LSAs from their databases. As a result, the overall volume of route advertisement traffic within the AS is reduced and the Layer 3 Switches

that flush the duplicate AS External LSAs have more memory for other OSPF data. In [Figure 108](#), since Router D has a higher router ID than Router E, Router D floods the AS External LSAs for Router F to Routers A, B, and C. Router E flushes the equivalent AS External LSAs from its database.

Algorithm for AS External LSA reduction

[Figure 108](#) shows an example in which the normal AS External LSA reduction feature is in effect. The behavior changes under the following conditions:

- There is one ASBR advertising (originating) a route to the external destination, but one of the following happens:
 - A second ASBR comes on-line
 - A second ASBR that is already on-line begins advertising an equivalent route to the same destination.

In either case above, the router with the higher router ID floods the AS External LSAs and the other router flushes its equivalent AS External LSAs. For example, if Router D is offline, Router E is the only source for a route to the external routing domain. When Router D comes on-line, it takes over flooding of the AS External LSAs to Router F, while Router E flushes its equivalent AS External LSAs to Router F.

- One of the ASBRs starts advertising a route that is no longer equivalent to the route the other ASBR is advertising. In this case, the ASBRs each flood AS External LSAs. Since the LSAs either no longer have the same cost or no longer have the same next-hop router, the LSAs are no longer equivalent, and the LSA reduction feature no longer applies.
- The ASBR with the higher router ID becomes unavailable or is reconfigured so that it is no longer an ASBR. In this case, the other ASBR floods the AS External LSAs. For example, if Router D goes off-line, then Router E starts flooding the AS with AS External LSAs for the route to Router F.

Support for OSPF RFC 2328 Appendix E

Devices provide support for Appendix E in OSPF RFC 2328. Appendix E describes a method to ensure that an OSPF router (such as a Layer 3 Switch) generates unique link state IDs for type-5 (External) link state advertisements (LSAs) in cases where two networks have the same network address but different network masks.

NOTE

Support for Appendix E of RFC 2328 is enabled automatically and cannot be disabled. No user configuration is required.

Normally, an OSPF router uses the network address alone for the link state ID of the link state advertisement (LSA) for the network. For example, if the router needs to generate an LSA for network 10.1.2.3 255.0.0.0, the router generates ID 10.1.2.3 for the LSA.

However, suppose that an OSPF router needs to generate LSAs for all the following networks:

- 10.0.0.0 255.0.0.0
- 10.0.0.0 255.255.0.0
- 10.0.0.0 255.255.255.0

All three networks have the same network address, 10.0.0.0. Without support for RFC 2328 Appendix E, an OSPF router uses the same link state ID, 10.0.0.0, for the LSAs for all three networks. For example, if the router generates an LSA with ID 10.0.0.0 for network 10.0.0.0 255.0.0.0, this LSA conflicts with the LSA generated for network 10.0.0.0 255.255.0.0 or 10.0.0.0 255.255.255.0. The result is multiple LSAs that have the same ID but that contain different route information.

When appendix E is supported, the router generates the link state ID for a network as follows.

1. Does an LSA with the network address as its ID already exist?
 - No – Use the network address as the ID.
 - Yes – Go to [step 2](#).
2. Compare the networks that have the same network address, to determine which network is more specific. The more specific network is the one that has more contiguous one bits in its network mask. For example, network 10.0.0.0 255.255.0.0 is more specific than network 10.0.0.0 255.0.0.0, because the first network has 16 ones bits (255.255.0.0) whereas the second network has only 8 ones bits (255.0.0.0):
 - For the less specific network, use the networks address as the ID.
 - For the more specific network, use the network broadcast address as the ID. The broadcast address is the network address, with all ones bits in the host portion of the address. For example, the broadcast address for network 10.0.0.0 255.255.0.0 is 10.0.0.255.

If this comparison results in a change to the ID of an LSA that has already been generated, the router generates a new LSA to replace the previous one. For example, if the router has already generated an LSA for network with ID 10.0.0.0 for network 10.0.0.0 255.255.255.0, the router must generate a new LSA for the network, if the router needs to generate an LSA for network 10.0.0.0 255.255.0.0 or 10.0.0.0 255.0.0.0.

Dynamic OSPF activation and configuration

OSPF is automatically activated when you enable it. The protocol does not require a software reload.

You can configure and save the following OSPF changes without resetting the system:

- All OSPF interface-related parameters (for example: area, hello timer, router dead time cost, priority, re-transmission time, transit delay)
- All area parameters
- All area range parameters
- All virtual-link parameters
- All global parameters
- Creation and deletion of an area, interface or virtual link

In addition, you can make the following changes without a system reset by first disabling and then re-enabling OSPF operation:

- Changes to address ranges
- Changes to global values for redistribution
- Addition of new virtual links

You also can change the amount of memory allocated to various types of LSA entries. However, these changes require a system reset or reboot.

Configuring OSPF

Follow the steps given below to begin using OSPF on the router.

1. Enable OSPF on the router.
2. Assign the areas to which the router will be attached.
3. Assign individual interfaces to the OSPF areas.
4. Define redistribution filters, if desired.
5. Enable redistribution, if you defined redistribution filters.
6. Modify default global and port parameters as required.
7. Modify OSPF standard compliance, if desired.

NOTE

OSPF is automatically enabled without a system reset.

Configuration rules

- If a router is to operate as an ASBR, you must enable the ASBR capability at the system level.
- Redistribution must be enabled on routers configured to operate as ASBRs.
- All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding subnets on that port are automatically included in the assignment.

OSPF parameters

You can modify or set the following global and interface OSPF parameters.

Global parameters:

- Modify OSPF standard compliance setting.
- Assign an area.
- Define an area range.
- Define the area virtual link.
- Set global default metric for OSPF.
- Change the reference bandwidth for the default cost of OSPF interfaces.
- Disable or re-enable load sharing.
- Enable or disable default-information-originate.
- Modify Shortest Path First (SPF) timers
- Define external route summarization
- Define redistribution metric type.

- Define deny redistribution.
- Define permit redistribution.
- Enable redistribution.
- Change the LSA pacing interval.
- Modify OSPF Traps generated.
- Modify database overflow interval.

Interface parameters:

- Assign interfaces to an area.
- Define the authentication key for the interface.
- Change the authentication-change interval
- Modify the cost for a link.
- Modify the dead interval.
- Modify MD5 authentication key parameters.
- Modify the priority of the interface.
- Modify the retransmit interval for the interface.
- Modify the transit delay of the interface.

NOTE

When using the CLI, you set global level parameters at the OSPF CONFIG Level of the CLI. To reach that level, enter **router ospf...** at the global CONFIG Level. Interface parameters for OSPF are set at the interface CONFIG Level using the CLI command, **ip ospf...**

Enable OSPF on the router

When you enable OSPF on the router, the protocol is automatically activated. To enable OSPF on the router, enter the following CLI command.

```
PowerConnect(config)# router ospf
```

This command launches you into the OSPF router level where you can assign areas and modify OSPF global parameters.

Note regarding disabling OSPF

If you disable OSPF, the Layer 3 Switch removes all the configuration information for the disabled protocol from the running-config. Moreover, when you save the configuration to the startup-config file after disabling one of these protocols, all the configuration information for the disabled protocol is removed from the startup-config file.

NOTE

If you do not want to delete the OSPF configuration information, use the CLI command **clear ip ospf process** instead of **no router ospf**. Refer to “Resetting OSPF” on page 667.

When you enter the **no router ospf** command, the CLI displays a warning message such as the following.


```
PowerConnect(config-ospf-router)# no router ospf
router ospf mode now disabled. All ospf config data will be lost when writing to
flash!
```

If you have disabled the protocol but have not yet saved the configuration to the startup-config file and reloaded the software, you can restore the configuration information by re-entering the command to enable the protocol (ex: **router ospf**). If you have already saved the configuration to the startup-config file and reloaded the software, the information is gone.

If you are testing an OSPF configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup-config file containing the protocol configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

Resetting OSPF

The **clear ip ospf process all** command globally resets (disables then re-enables) OSPF without deleting the OSPF configuration information. This command is equivalent to entering the commands **no router ospf** followed by **router ospf**. Whereas the **no router ospf** command disables OSPF and removes all the configuration information for the disabled protocol from the running-config, the **router ospf** command re-enables OSPF and restores the OSPF configuration information.

The **clear ip ospf process all** command is useful if you are testing an OSPF configuration and are likely to disable and re-enable the protocol. This way, you do not have to save the configuration after disabling the protocol, and you do not have to restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

To reset OSPF without deleting the OSPF configuration, enter the following command at the Global CONFIG level or at the Router OSPF level of the CLI.

```
PowerConnect# clear ip ospf process all
```

Syntax: clear ip ospf process all

Assign OSPF areas

Once OSPF is enabled on the system, you can assign areas. Assign an IP address or number as the **area ID** for each area. The area ID is representative of all IP addresses (subnets) on a router port. Each port on a router can support one area.

An area can be **normal**, a **stub**, or a **Not-So-Stubby Area (NSSA)**:

- Normal – OSPF routers within a normal area can send and receive External Link State Advertisements (LSAs).
- Stub – OSPF routers within a stub area cannot send or receive External LSAs. In addition, OSPF routers in a stub area must use a default route to the area Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) to send traffic out of the area.
- NSSA – The ASBR of an NSSA can import external route information into the area:
 - ASBRs redistribute (import) external routes into the NSSA as type 7 LSAs. Type-7 External LSAs are a special type of LSA generated only by ASBRs within an NSSA, and are flooded to all the routers within only that NSSA.

- ABRs translate type 7 LSAs into type 5 External LSAs, which can then be flooded throughout the AS. You can configure address ranges on the ABR of an NSSA so that the ABR converts multiple type-7 External LSAs received from the NSSA into a single type-5 External LSA.

When an NSSA contains more than one ABR, OSPF elects one of the ABRs to perform the LSA translation for NSSA. OSPF elects the ABR with the highest router ID. If the elected ABR becomes unavailable, OSPF automatically elects the ABR with the next highest router ID to take over translation of LSAs for the NSSA. The election process for NSSA ABRs is automatic.

Example

To set up the OSPF areas shown in [Figure 105](#) on page 658, enter the following commands.

```
PowerConnect(config-ospf-router)# area 192.5.1.0
PowerConnect(config-ospf-router)# area 200.5.0.0
PowerConnect(config-ospf-router)# area 195.5.0.0
PowerConnect(config-ospf-router)# area 0.0.0.0
PowerConnect(config-ospf-router)# write memory
```

Syntax: `area <num> | <ip-addr>`

The `<num> | <ip-addr>` parameter specifies the area number, which can be a number or in IP address format. If you specify an number, the number can be from 0 – 2,147,483,647.

NOTE

You can assign one area on a router interface. For example, if the system or chassis module has 16 ports, 16 areas are supported on the chassis or module.

Assign a totally stubby area

By default, the Layer 3 Switch sends summary LSAs (LSA type 3) into stub areas. You can further reduce the number of link state advertisements (LSAs) sent into a stub area by configuring the Layer 3 Switch to stop sending summary LSAs (type 3 LSAs) into the area. You can disable the summary LSAs when you are configuring the stub area or later after you have configured the area.

This feature disables origination of summary LSAs, but the Layer 3 Switch still accepts summary LSAs from OSPF neighbors and floods them to other neighbors. The Layer 3 Switch can form adjacencies with other routers regardless of whether summarization is enabled or disabled for areas on each router.

When you enter a command to disable the summary LSAs, the change takes effect immediately. If you apply the option to a previously configured area, the Layer 3 Switch flushes all of the summary LSAs it has generated (as an ABR) from the area.

NOTE

This feature applies only when the Layer 3 Switch is configured as an Area Border Router (ABR) for the area. To completely prevent summary LSAs from being sent to the area, disable the summary LSAs on each OSPF router that is an ABR for the area.

This feature does not apply to Not So Stubby Areas (NSSAs).

To disable summary LSAs for a stub area, enter commands such as the following.

```
PowerConnect(config-ospf-router)# area 40 stub 99 no-summary
```

Syntax: `area <num> | <ip-addr> stub <cost> [no-summary]`

The `<num> | <ip-addr>` parameter specifies the area number, which can be a number or in IP address format. If you specify a number, the number can be from 0 – 2,147,483,647.

The `stub <cost>` parameter specifies an additional cost for using a route to or from this area and can be from 1 – 16777215. There is no default. Normal areas do not use the cost parameter.

The `no-summary` parameter applies only to stub areas and disables summary LSAs from being sent into the area.

NOTE

You can assign one area on a router interface. For example, if the system or chassis module has 16 ports, 16 areas are supported on the chassis or module.

Assign a Not-So-Stubby Area (NSSA)

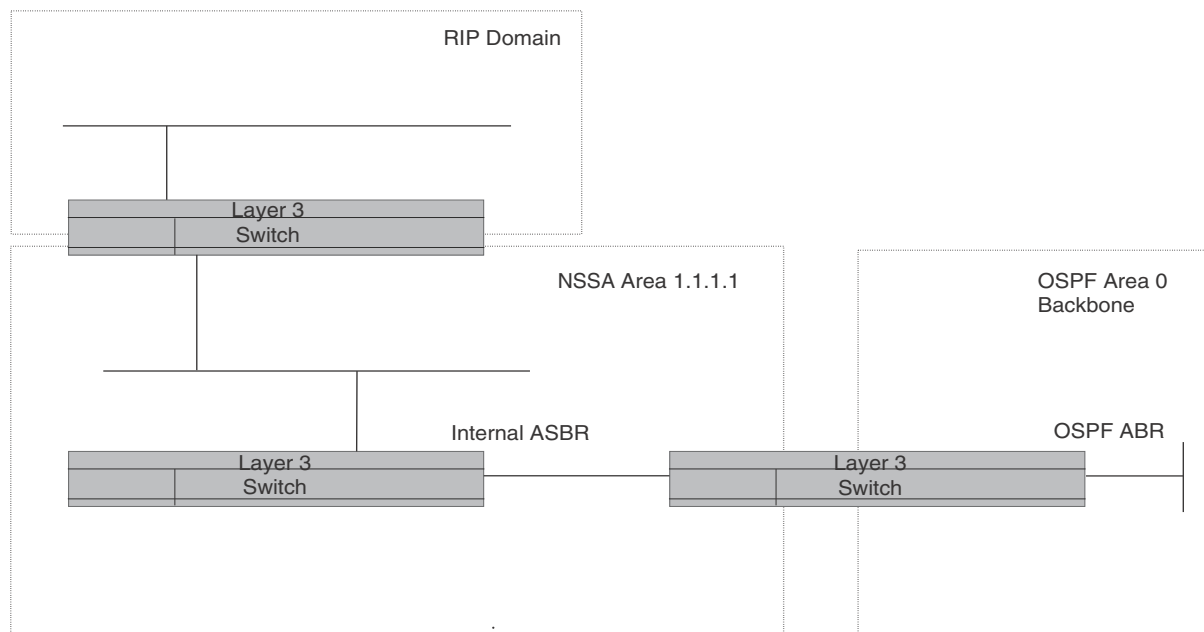
The OSPF Not So Stubby Area (NSSA) feature enables you to configure OSPF areas that provide the benefits of stub areas, but that also are capable of importing external route information. OSPF does not flood external routes from other areas into an NSSA, but does translate and flood route information from the NSSA into other areas such as the backbone.

NSSAs are especially useful when you want to summarize Type-5 External LSAs (external routes) before forwarding them into an OSPF area. The OSPF specification (RFC 2328) prohibits summarization of Type-5 LSAs and requires OSPF to flood Type-5 LSAs throughout a routing domain. When you configure an NSSA, you can specify an address range for aggregating the external routes that the NSSA's ABR exports into other areas.

The Dell implementation of NSSA is based on RFC 1587.

Figure 109 shows an example of an OSPF network containing an NSSA.

FIGURE 109 OSPF network containing an NSSA



This example shows two routing domains, a RIP domain and an OSPF domain. The ASBR inside the NSSA imports external routes from RIP into the NSSA as Type-7 LSAs, which the ASBR floods throughout the NSSA.

The ABR translates the Type-7 LSAs into Type-5 LSAs. If an area range is configured for the NSSA, the ABR also summarizes the LSAs into an aggregate LSA before flooding the Type-5 LSAs into the backbone.

Since the NSSA is partially “stubby” the ABR does not flood external LSAs from the backbone into the NSSA. To provide access to the rest of the Autonomous System (AS), the ABR generates a default Type-7 LSA into the NSSA.

Configuring an NSSA

To configure OSPF area 1.1.1.1 as an NSSA, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# area 1.1.1.1 nssa 1
PowerConnect(config-ospf-router)# write memory
```

Syntax: `area <num> | <ip-addr> nssa <cost> | default-information-originate`

The `<num> | <ip-addr>` parameter specifies the area number, which can be a number or in IP address format. If you specify an number, the number can be from 0 – 2,147,483,647.

The `nssa <cost> | default-information-originate` parameter specifies that this is a Not-So-Stubby-Area (NSSA). The `<cost>` specifies an additional cost for using a route to or from this NSSA and can be from 1 – 16777215. There is no default. Normal areas do not use the cost parameter. Alternatively, you can use the `default-information-originate` parameter causes the Layer 3 Switch to inject the default route into the NSSA.

NOTE

The Layer 3 Switch does not inject the default route into an NSSA by default.

NOTE

You can assign one area on a router interface. For example, if the system or chassis module has 16 ports, 16 areas are supported on the chassis or module.

To configure additional parameters for OSPF interfaces in the NSSA, use the `ip ospf area...` command at the interface level of the CLI.

Configuring a summary address for the NSSA

If you want the ABR that connects the NSSA to other areas to summarize the routes in the NSSA before translating them into Type-5 LSAs and flooding them into the other areas, configure a summary address. The ABR creates an aggregate value based on the summary address. The aggregate value becomes the address that the ABR advertises instead of advertising the individual addresses represented by the aggregate.

To configure a summary address in NSSA 1.1.1.1, enter the following commands. This example assumes that you have already configured NSSA 1.1.1.1.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# summary-address 209.157.22.1 255.255.0.0
PowerConnect(config-ospf-router)# write memory
```

Syntax: `[no] summary address <ip-addr> <ip-mask>`

The `<ip-addr>` parameter specifies the IP address portion of the range. The software compares the address with the significant bits in the mask. All network addresses that match this comparison are summarized in a single route advertised by the router.

The `<ip-mask>` parameter specifies the portions of the IP address that a route must contain to be summarized in the summary route. In the example above, all networks that begin with 209.157 are summarized into a single route.

Assigning an area range (optional)

You can assign a **range** for an area, but it is not required. Ranges allow a specific IP address and mask to represent a range of IP addresses within an area, so that only that reference range address is advertised to the network, instead of all the addresses within that range. Each area can have up to 32 range addresses.

Example

To define an area range for subnets on 193.45.5.1 and 193.45.6.2, enter the following command.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# area 192.45.5.1 range 193.45.0.0 255.255.0.0
PowerConnect(config-ospf-router)# area 193.45.6.2 range 193.45.0.0 255.255.0.0
```

Syntax: `area <num> | <ip-addr> range <ip-addr> <ip-mask>`

The `<num> | <ip-addr>` parameter specifies the area number, which can be in IP address format.

The **range** `<ip-addr>` parameter specifies the IP address portion of the range. The software compares the address with the significant bits in the mask. All network addresses that match this comparison are summarized in a single route advertised by the router.

The `<ip-mask>` parameter specifies the portions of the IP address that a route must contain to be summarized in the summary route. In the example above, all networks that begin with 193.45 are summarized into a single route.

Assigning interfaces to an area

Once you define OSPF areas, you can assign interfaces to the areas. All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding subnets on that port are automatically included in the assignment.

To assign interface 8 to area 195.5.0.0 and then save the changes, enter the following commands.

```
PowerConnect(config-ospf-router)# interface e 8
PowerConnect(config-if-8)# ip ospf area 195.5.0.0
PowerConnect(config-if-8)# write memory
```

Modify interface defaults

OSPF has interface parameters that you can configure. For simplicity, each of these parameters has a default value. No change to these default values is required except as needed for specific network configurations.

Port default values can be modified using the following commands at the interface configuration level of the CLI:

- `ip ospf area <ip-addr>`

- ip ospf auth-change-wait-time <secs>
- ip ospf authentication-key [0 | 1] <string>
- ip ospf cost <num>
- ip ospf dead-interval <value>
- ip ospf hello-interval <value>
- ip ospf md5-authentication key-activation-wait-time <num> | key-id <num> [0 | 1] key <string>
- ip ospf passive
- ip ospf priority <value>
- ip ospf retransmit-interval <value>
- ip ospf transmit-delay <value>

For a complete description of these parameters, see the summary of OSPF port parameters in the next section.

OSPF interface parameters

The following parameters apply to OSPF interfaces.

Area: Assigns an interface to a specific area. You can assign either an IP address or number to represent an OSPF Area ID. If you assign a number, it can be any value from 0 – 2,147,483,647.

Auth-change-wait-time: OSPF gracefully implements authentication changes to allow all routers to implement the change and thus prevent disruption to neighbor adjacencies. During the authentication-change interval, both the old and new authentication information is supported. The default authentication-change interval is 300 seconds (5 minutes). You change the interval to a value from 0 – 14400 seconds.

Authentication-key: OSPF supports three methods of authentication for each interface—none, simple password, and MD5. Only one method of authentication can be active on an interface at a time. The default authentication value is none, meaning no authentication is performed.

The simple password method of authentication requires you to configure an alphanumeric password on an interface. The simple password setting takes effect immediately. All OSPF packets transmitted on the interface contain this password. Any OSPF packet received on the interface is checked for this password. If the password is not present, then the packet is dropped. The password can be up to eight characters long.

The MD5 method of authentication requires you to configure a key ID and an MD5 Key. The key ID is a number from 1 – 255 and identifies the MD5 key that is being used. The MD5 key can be up to sixteen alphanumeric characters long.

Cost: Indicates the overhead required to send a packet across an interface. You can modify the cost to differentiate between 100 Mbps and 1000 Mbps (1 Gbps) links. The default cost is calculated by dividing 100 million by the bandwidth. For 10 Mbps links, the cost is 10. The cost for both 100 Mbps and 1000 Mbps links is 1, because the speed of 1000 Mbps was not in use at the time the OSPF cost formula was devised.

Dead-interval: Indicates the number of seconds that a neighbor router waits for a hello packet from the current router before declaring the router down. The value can be from 1 – 65535 seconds. The default is 40 seconds.

Hello-interval: Represents the length of time between the transmission of hello packets. The value can be from 1 – 65535 seconds. The default is 10 seconds.

MD5-authentication activation wait time: The number of seconds the Layer 3 Switch waits until placing a new MD5 key into effect. The wait time provides a way to gracefully transition from one MD5 key to another without disturbing the network. The wait time can be from 0 – 14400 seconds. The default is 300 seconds (5 minutes).

MD5-authentication key ID and key: A method of authentication that requires you to configure a key ID and an MD5 key. The key ID is a number from 1 – 255 and identifies the MD5 key that is being used. The MD5 key consists of up to 16 alphanumeric characters. The MD5 is encrypted and included in each OSPF packet transmitted.

Passive: When you configure an OSPF interface to be passive, that interface does not send or receive OSPF route updates. By default, all OSPF interfaces are active and thus can send and receive OSPF route information. Since a passive interface does not send or receive route information, the interface is in effect a stub network. OSPF interfaces are active by default.

NOTE

This option affects all IP subnets configured on the interface. If you want to disable OSPF updates only on some of the IP subnets on the interface, use the **ospf-ignore** or **ospf-passive** parameter with the **ip address** command. Refer to [“Assigning an IP address to an Ethernet port”](#) on page 579.

Priority: Allows you to modify the priority of an OSPF router. The priority is used when selecting the designated router (DR) and backup designated routers (BDRs). The value can be from 0 – 255. The default is 1. If you set the priority to 0, the Layer 3 Switch does not participate in DR and BDR election.

Retransmit-interval: The time between retransmissions of link-state advertisements (LSAs) to adjacent routers for this interface. The value can be from 0 – 3600 seconds. The default is 5 seconds.

Transit-delay: The time it takes to transmit Link State Update packets on this interface. The value can be from 0 – 3600 seconds. The default is 1 second.

Encrypted Display of the Authentication String or MD5 Authentication Key

The optional **0 | 1** parameter with the **authentication-key** and **md5-authentication key-id** parameters affects encryption.

For added security, devices encrypt display of the password or authentication string. Encryption is enabled by default. The software also provides an optional parameter to disable encryption of a password or authentication string, on an individual OSPF area or OSPF interface basis.

When encryption of the passwords or authentication strings is enabled, they are encrypted in the CLI regardless of the access level you are using.

The encryption option can be omitted (the default) or can be one of the following:

- **0** – Disables encryption for the password or authentication string you specify with the command. The password or string is shown as clear text in the running-config and the startup-config file. Use this option if you do not want display of the password or string to be encrypted.
- **1** – Assumes that the password or authentication string you enter is the encrypted form, and decrypts the value before using it.

NOTE

If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter **0** or **1**. Instead, omit the encryption option and allow the software to use the default behavior.

If you specify encryption option **1**, the software assumes that you are entering the encrypted form of the password or authentication string. In this case, the software decrypts the password or string you enter before using the value for authentication. If you accidentally enter option **1** followed by the clear-text version of the password or string, authentication will fail because the value used by the software will not match the value you intended to use.

Change the timer for OSPF authentication changes

When you make an OSPF authentication change, the software uses the authentication-change timer to gracefully implement the change. The software implements the change in the following ways:

- **Outgoing OSPF packets** – After you make the change, the software continues to use the old authentication to send packets, during the remainder of the current authentication-change interval. After this, the software uses the new authentication for sending packets.
- **Inbound OSPF packets** – The software accepts packets containing the new authentication and continues to accept packets containing the older authentication for two authentication-change intervals. After the second interval ends, the software accepts packets only if they contain the new authentication key.

The default authentication-change interval is 300 seconds (5 minutes). You change the interval to a value from 0 – 14400 seconds.

OSPF provides graceful authentication change for all the following types of authentication changes in OSPF:

- Changing authentication methods from one of the following to another of the following:
 - Simple text password
 - MD5 authentication
 - No authentication
- Configuring a new simple text password or MD5 authentication key
- Changing an existing simple text password or MD5 authentication key

To change the authentication-change interval, enter a command such as the following at the interface configuration level of the CLI.

```
PowerConnect(config-if-5)# ip ospf auth-change-wait-time 400
```

Syntax: [no] ip ospf auth-change-wait-time <secs>

The <secs> parameter specifies the interval and can be from 0 – 14400 seconds. The default is 300 seconds (5 minutes).

NOTE

For backward compatibility, the **ip ospf md5-authentication key-activation-wait-time <seconds>** command is still supported.

Block flooding of outbound LSAs on specific OSPF interfaces

By default, the Layer 3 Switch floods all outbound LSAs on all the OSPF interfaces within an area. You can configure a filter to block outbound LSAs on an OSPF interface. This feature is particularly useful when you want to block LSAs from some, but not all, of the interfaces attached to the area.

After you apply filters to block the outbound LSAs, the filtering occurs during the database synchronization and flooding.

If you remove the filters, the blocked LSAs are automatically re-flooded. You do not need to reset OSPF to re-flood the LSAs.

NOTE

You cannot block LSAs on virtual links.

To apply a filter to an OSPF interface to block flooding of outbound LSAs on the interface, enter the following command at the Interface configuration level for that interface.

```
PowerConnect(config-if-1)# ip ospf database-filter all out
```

The command in this example blocks all outbound LSAs on the OSPF interface configured on port 1.

Syntax: [no] ip ospf database-filter all out

To remove the filter, enter a command such as the following.

```
PowerConnect(config-if-1)# no ip ospf database-filter all out
```

Assign virtual links

All ABRs (area border routers) must have either a direct or indirect link to the OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to the area backbone, the ABR can configure a **virtual link** to another router within the same area, which has a physical connection to the area backbone.

The path for a virtual link is through an area shared by the neighbor ABR (router with a physical backbone connection), and the ABR requiring a logical connection to the backbone.

Two parameters fields must be defined for all virtual links—transit area ID and neighbor router:

- The **transit area ID** represents the shared area of the two ABRs and serves as the connection point between the two routers. This number should match the area ID value.
- The **neighbor router** field is the router ID (IP address) of the router that is physically connected to the backbone, when assigned from the router interface requiring a logical connection. When assigning the parameters from the router with the physical connection, the router ID is the IP address of the router requiring a logical connection to the backbone.

NOTE

By default, the router ID is the IP address configured on the lowest numbered loopback interface. If the Layer 3 Switch does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, refer to [“Changing the router ID”](#) on page 584.

NOTE

When you establish an area virtual link, you must configure it on both of the routers (both ends of the virtual link).

FIGURE 110 Defining OSPF virtual links within a network

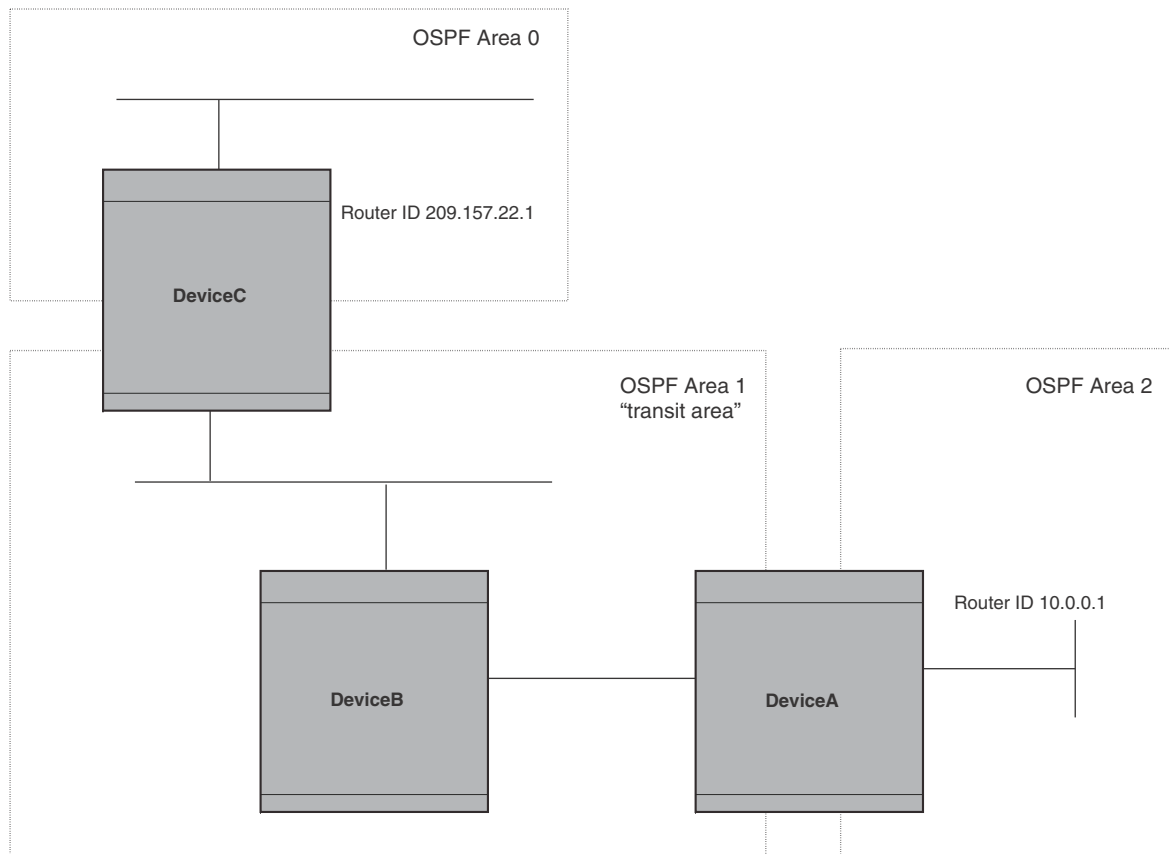
**Example**

Figure 110 shows an OSPF area border router, DeviceA, that is cut off from the backbone area (area 0). To provide backbone access to DeviceA, you can add a virtual link between DeviceA and DeviceC using area 1 as a transit area. To configure the virtual link, you define the link on the router that is at each end of the link. No configuration for the virtual link is required on the routers in the transit area.

To define the virtual link on DeviceA, enter the following commands.

```
PowerConnectA(config-ospf-router)# area 1 virtual-link 209.157.22.1
PowerConnectA(config-ospf-router)# write memory
```

Enter the following commands to configure the virtual link on DeviceC.

```
PowerConnectC(config-ospf-router)# area 1 virtual-link 10.0.0.1
PowerConnectC(config-ospf-router)# write memory
```

Syntax: `area <ip-addr> | <num> virtual-link <router-id>`
`[authentication-key | dead-interval | hello-interval | retransmit-interval | transmit-delay <value>]`

The **area** *<ip-addr> | <num>* parameter specifies the transit area.

The *<router-id>* parameter specifies the router ID of the OSPF router at the remote end of the virtual link. To display the router ID on a Layer 3 Switch, enter the **show ip** command.

Refer to “[Modify virtual link parameters](#)” on page 677 for descriptions of the optional parameters.

Modify virtual link parameters

OSPF has some parameters that you can modify for virtual links. Notice that these are the same parameters as the ones you can modify for physical interfaces.

You can modify default values for virtual links using the following CLI command at the **OSPF router level** of the CLI, as shown in the following syntax.

Syntax: **area** *<num> | <ip-addr>* **virtual-link** *<ip-addr>* [**authentication-key** [0 | 1] *<string>*]
 [**dead-interval** *<num>*]
 [**hello-interval** *<num>*] [**md5-authentication key-activation-wait-time** *<num>* | **key-id**
<num>] [0 | 1] **key** *<string>*]
 [**retransmit-interval** *<num>*] [**transmit-delay** *<num>*]

The parameters are described below.

Virtual link parameter descriptions

You can modify the following virtual link interface parameters.

Authentication Key: This parameter allows you to assign different authentication methods on a port-by-port basis. OSPF supports three methods of authentication for each interface—none, simple password, and MD5. Only one method of authentication can be active on an interface at a time.

The simple password method of authentication requires you to configure an alphanumeric password on an interface. The password can be up to eight characters long. The simple password setting takes effect immediately. All OSPF packets transmitted on the interface contain this password. All OSPF packets received on the interface are checked for this password. If the password is not present, the packet is dropped.

The MD5 method of authentication encrypts the authentication key you define. The authentication is included in each OSPF packet transmitted.

MD5 Authentication Key: When simple authentication is enabled, the key is an alphanumeric password of up to eight characters. When MD5 is enabled, the key is an alphanumeric password of up to 16 characters that is later encrypted and included in each OSPF packet transmitted. You must enter a password in this field when the system is configured to operate with either simple or MD5 authentication.

MD5 Authentication Key ID: The Key ID is a number from 1 – 255 and identifies the MD5 key that is being used. This parameter is required to differentiate among multiple keys defined on a router.

MD5 Authentication Wait Time: This parameter determines when a newly configured MD5 authentication key is valid. This parameter provides a graceful transition from one MD5 key to another without disturbing the network. All new packets transmitted after the key activation wait time interval use the newly configured MD5 Key. OSPF packets that contain the old MD5 key are accepted for up to five minutes after the new MD5 key is in operation.

The range for the key activation wait time is from 0 – 14400 seconds. The default value is 300 seconds.

Hello Interval: The length of time between the transmission of hello packets. The range is 1 – 65535 seconds. The default is 10 seconds.

Retransmit Interval: The interval between the re-transmission of link state advertisements to router adjacencies for this interface. The range is 0 – 3600 seconds. The default is 5 seconds.

Transmit Delay: The period of time it takes to transmit Link State Update packets on the interface. The range is 0 – 3600 seconds. The default is 1 second.

Dead Interval: The number of seconds that a neighbor router waits for a hello packet from the current router before declaring the router down. The range is 1 – 65535 seconds. The default is 40 seconds.

Changing the reference bandwidth for the cost on OSPF interfaces

Each interface on which OSPF is enabled has a cost associated with it. The Layer 3 Switch advertises its interfaces and their costs to OSPF neighbors. For example, if an interface has an OSPF cost of ten, the Layer 3 Switch advertises the interface with a cost of ten to other OSPF routers.

By default, an interface OSPF cost is based on the port speed of the interface. The cost is calculated by dividing the reference bandwidth by the port speed. The default reference bandwidth is 100 Mbps, which results in the following default costs:

- 10 Mbps port – 10
- All other port speeds – 1

You can change the reference bandwidth, to change the costs calculated by the software.

The software uses the following formula to calculate the cost.

$$\text{Cost} = \text{reference-bandwidth}/\text{interface-speed}$$

If the resulting cost is less than 1, the software rounds the cost up to 1. The default reference bandwidth results in the following costs:

- 10 Mbps port cost = $100/10 = 10$
- 100 Mbps port cost = $100/100 = 1$
- 1000 Mbps port cost = $100/1000 = 0.10$, which is rounded up to 1
- 155 Mbps port cost = $100/155 = 0.65$, which is rounded up to 1
- 622 Mbps port cost = $100/622 = 0.16$, which is rounded up to 1
- 2488 Mbps port cost = $100/2488 = 0.04$, which is rounded up to 1

For 10 Gbps OSPF interfaces, in order to differentiate the costs between 100 Mbps, 1000 Mbps, and 10,000 Mbps interfaces, you can set the auto-cost reference bandwidth to 10000, whereby each slower link is given a higher cost, as follows:

- 10 Mbps port cost = $10000/10 = 1000$
- 100 Mbps port cost = $10000/100 = 100$
- 1000 Mbps port cost = $10000/1000 = 10$
- 10000 Mbps port cost = $10000/10000 = 1$

The bandwidth for interfaces that consist of more than one physical port is calculated as follows:

- Trunk group – The combined bandwidth of all the ports.
- Virtual interface – The combined bandwidth of all the ports in the port-based VLAN that contains the virtual interface.

The default reference bandwidth is 100 Mbps. You can change the reference bandwidth to a value from 1 – 4294967.

If a change to the reference bandwidth results in a cost change to an interface, the Layer 3 Switch sends a link-state update to update the costs of interfaces advertised by the Layer 3 Switch.

NOTE

If you specify the cost for an individual interface, the cost you specify overrides the cost calculated by the software.

Interface types to which the reference bandwidth does not apply

Some interface types are not affected by the reference bandwidth and always have the same cost regardless of the reference bandwidth in use:

- The cost of a loopback interface is always 0.
- The cost of a virtual link is calculated using the Shortest Path First (SPF) algorithm and is not affected by the auto-cost feature.
- The bandwidth for tunnel interfaces is 9 Kbps and is not affected by the auto-cost feature.

Changing the reference bandwidth

To change the reference bandwidth, enter a command such as the following at the OSPF configuration level of the CLI.

```
PowerConnect(config-ospf-router)# auto-cost reference-bandwidth 500
```

The reference bandwidth specified in this example results in the following costs:

- 10 Mbps port cost = $500/10 = 50$
- 100 Mbps port cost = $500/100 = 5$
- 1000 Mbps port cost = $500/1000 = 0.5$, which is rounded up to 1
- 155 Mbps port cost = $500/155 = 3.23$, which is rounded up to 4
- 622 Mbps port cost = $500/622 = 0.80$, which is rounded up to 1
- 2488 Mbps port cost = $500/2488 = 0.20$, which is rounded up to 1

The costs for 10 Mbps, 100 Mbps, and 155 Mbps ports change as a result of the changed reference bandwidth. Costs for higher-speed interfaces remain the same.

Syntax: [no] auto-cost reference-bandwidth <num>

The <num> parameter specifies the reference bandwidth and can be a value from 1 – 4294967. The default is 100. For 10 Gbps OSPF interfaces, in order to differentiate the costs between 100 Mbps, 1000 Mbps, and 10,000 Mbps interfaces, set the auto-cost reference bandwidth to 10000, whereby each slower link is given a higher cost

To restore the reference bandwidth to its default value and thus restore the default costs of interfaces to their default values, enter the following command.

```
PowerConnect(config-ospf-router)# no auto-cost reference-bandwidth
```

Define redistribution filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On routers, redistribution is supported for static routes, OSPF, RIP, and BGP4. When you configure redistribution for RIP, you can specify that static, OSPF, or BGP4 routes are imported into RIP routes. Likewise, OSPF redistribution supports the import of static, RIP, and BGP4 routes into OSPF routes. BGP4 supports redistribution of static, RIP, and OSPF routes into BGP4.

NOTE

The Layer 3 Switch advertises the default route into OSPF even if redistribution is not enabled, and even if the default route is learned through an IBGP neighbor. IBGP routes (including the default route) are not redistributed into OSPF by OSPF redistribution (for example, by the OSPF **redistribute** command).

In [Figure 111](#) on page 681, an administrator wants to configure the Layer 3 Switch acting as the ASBR (Autonomous System Boundary Router) between the RIP domain and the OSPF domain to redistribute routes between the two domains.

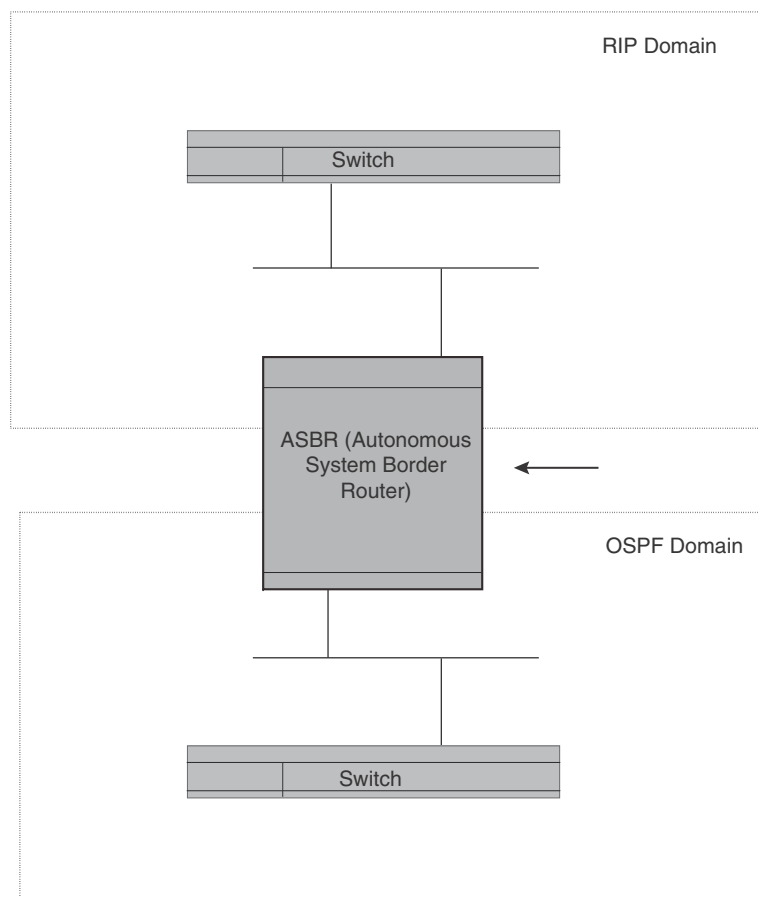
NOTE

The ASBR must be running both RIP and OSPF protocols to support this activity.

To configure for redistribution, define the redistribution tables with deny and permit redistribution filters. Use the **deny** and **permit** redistribute commands for OSPF at the OSPF router level.

NOTE

Do not enable redistribution until you have configured the redistribution filters. If you enable redistribution before you configure the redistribution filters, the filters will not take effect and all routes will be distributed.

FIGURE 111 Redistributing OSPF and static routes to RIP routes**Example**

To configure the Layer 3 Switch acting as an ASBR in [Figure 111](#) to redistribute OSPF, BGP4, and static routes into RIP, enter the following commands.

```
PowerConnectASBR(config)# router rip
PowerConnectASBR(config-rip-router)# permit redistribute 1 all
PowerConnectASBR(config-rip-router)# write memory
```

NOTE

Redistribution is permitted for all routes by default, so the **permit redistribute 1 all** command in the example above is shown for clarity but is not required.

You also have the option of specifying import of just OSPF, BGP4, or static routes, as well as specifying that only routes for a specific network or with a specific cost (metric) be imported, as shown in the command syntax below.

Syntax: `deny | permit redistribute <filter-num> all | bgp | connected | rip | static [address <ip-addr> <ip-mask> [match-metric <value> [set-metric <value>]]]`

Example

To redistribute RIP, static, and BGP4 routes into OSPF, enter the following commands on the Layer 3 Switch acting as an ASBR.

```
PowerConnectASBR(config)# router ospf
PowerConnectASBR(config-ospf-router)# permit redistribute 1 all
PowerConnectASBR(config-ospf-router)# write memory
```

Syntax: **deny** | **permit redistribute** <filter-num> **all** | **bgp** | **connected** | **rip** | **static**
address <ip-addr> <ip-mask>
[**match-metric** <value> | **set-metric** <value>]

NOTE

Redistribution is permitted for all routes by default, so the **permit redistribute 1 all** command in the example above is shown for clarity but is not required.

You also have the option of specifying import of just OSPF, BGP4, or static routes, as well as specifying that only routes for a specific network or with a specific cost (metric) be imported, as shown in the command syntax below.

Syntax: [**no**] **redistribution bgp** | **connected** | **rip** | **static** [**route-map** <map-name>]

For example, to enable redistribution of RIP and static IP routes into OSPF, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# redistribution rip
PowerConnect(config-ospf-router)# redistribution static
PowerConnect(config-ospf-router)# write memory
```

NOTE

The **redistribution** command does not perform the same function as the **permit redistribute** and **deny redistribute** commands. The **redistribute** commands allow you to control redistribution of routes by filtering on the IP address and network mask of a route. The **redistribution** commands enable redistribution for routes of specific types (static, directly connected, and so on). Configure all your redistribution filters before enabling redistribution.

NOTE

Do not enable redistribution until you have configured the redistribution filters. If you enable redistribution before you configure the redistribution filters, the filters will not take effect and all routes will be distributed.

Prevent specific OSPF routes from being installed in the IP route table

By default, all OSPF routes in the OSPF route table are eligible for installation in the IP route table. You can configure a distribution list to explicitly deny specific routes from being eligible for installation in the IP route table.

NOTE

This feature does not block receipt of LSAs for the denied routes. The Layer 3 Switch still receives the routes and installs them in the OSPF database. The feature only prevents the software from installing the denied OSPF routes into the IP route table.

To configure an OSPF distribution list:

- Configure a standard or extended ACL that identifies the routes you want to deny. Using a standard ACL lets you deny routes based on the destination network, but does not filter based on the network mask. To also filter based on the destination network network mask, use an extended ACL.
- Configure an OSPF distribution list that uses the ACL as input.

NOTE

If you change the ACL after you configure the OSPF distribution list, you must clear the IP route table to place the changed ACL into effect. To clear the IP route table, enter the **clear ip route** command at the Privileged EXEC level of the CLI.

The following sections show how to use the CLI to configure an OSPF distribution list. Separate examples are provided for standard and extended ACLs.

NOTE

The examples show named ACLs. However, you also can use a numbered ACL as input to the OSPF distribution list.

Using a standard ACL as input to the distribution list

To use a standard ACL to configure an OSPF distribution list for denying specific routes, enter commands such as the following.

```
PowerConnect(config)# ip access-list standard no_ip
PowerConnect(config-std-nACL)# deny 4.0.0.0 0.255.255.255
PowerConnect(config-std-nACL)# permit any any
PowerConnect(config-std-nACL)# exit
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# distribute-list no_ip in
```

The first three commands configure a standard ACL that denies routes to any 4.x.x.x destination network and allows all other routes for eligibility to be installed in the IP route table. The last three commands change the CLI to the OSPF configuration level and configure an OSPF distribution list that uses the ACL as input. The distribution list prevents routes to any 4.x.x.x destination network from entering the IP route table. The distribution list does not prevent the routes from entering the OSPF database.

Syntax: [no] **distribute-list** <ACL-name> | <ACL-id> **in** [<interface type>] [<interface number>]

Syntax: [no] **ip access-list standard** <ACL-name> | <ACL-id>

Syntax: **deny** | **permit** <source-ip> <wildcard>

The <ACL-name> | <ACL-id> parameter specifies the ACL name or ID.

The **in** command applies the ACL to incoming route updates.

The <interface number> parameter specifies the interface number on which to apply the ACL. Enter only one valid interface number. If necessary, use the **show interface brief** command to display a list of valid interfaces. If you do not specify an interface, the device applies the ACL to all incoming route updates.

If you do not specify an interface type and interface number, the device applies the OSPF distribution list to all incoming route updates.

The **deny** | **permit** parameter indicates whether packets that match the policy are dropped or forwarded.

The `<source-ip>` parameter specifies the source address for the policy. Since this ACL is input to an OSPF distribution list, the `<source-ip>` parameter actually is specifying the destination network of the route.

The `<wildcard>` parameter specifies the portion of the source address to match against. The `<wildcard>` is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the `<source-ip>`. Ones mean any value matches. For example, the `<source-ip>` and `<wildcard>` values 4.0.0.0 0.255.255.255 mean that all 4.x.x.x networks match the ACL.

If you want the policy to match on all destination networks, enter **any any**.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “4.0.0.0 0.255.255.255” as “4.0.0.0/8”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros.

NOTE

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/`<mask-bits>`” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** command.

Using an extended ACL as input to the distribution list

You can use an extended ACL with an OSPF distribution list to filter OSPF routes based on the network mask of the destination network.

To use an extended ACL to configure an OSPF distribution list for denying specific routes, enter commands such as the following.

```
PowerConnect(config)# ip access-list extended no_ip
PowerConnect(config-ext-nACL)# deny ip 4.0.0.0 0.255.255.255 255.255.0.0
0.0.255.255
PowerConnect(config-ext-nACL)# permit ip any any
PowerConnect(config-ext-nACL)# exit
PowerConnect(config)# router ospf
```

The first three commands configure an extended ACL that denies routes to any 4.x.x.x destination network with a 255.255.0.0 network mask and allows all other routes for eligibility to be installed in the IP route table. The last three commands change the CLI to the OSPF configuration level and configure an OSPF distribution list that uses the ACL as input. The distribution list prevents routes to any 4.x.x.x destination network with network mask 255.255.0.0 from entering the IP route table. The distribution list does not prevent the routes from entering the OSPF database.

Syntax: [no] **ip access-list extended** `<ACL-name>` | `<ACL-id>`

Syntax: **deny** | **permit** `<ip-protocol>` `<source-ip>` `<wildcard>` `<destination-ip>` `<wildcard>`

The `<ACL-name>` | `<ACL-id>` parameter specifies the ACL name or ID.

The **deny** | **permit** parameter indicates whether packets that match the policy are dropped or forwarded.

The `<ip-protocol>` parameter indicates the type of IP packet you are filtering. When using an extended ACL as input for an OSPF distribution list, specify **ip**.

Since this ACL is input to an OSPF distribution list, the `<source-ip>` parameter actually specifies the destination network of the route.

The `<wildcard>` parameter specifies the portion of the source address to match against. The `<wildcard>` is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the `<source-ip>`. Ones mean any value matches. For example, the `<source-ip>` and `<wildcard>` values 4.0.0.0 0.255.255.255 mean that all 4.x.x.x networks match the ACL.

If you want the policy to match on all network addresses, enter **any any**.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “4.0.0.0 0.255.255.255” as “4.0.0.0/8”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros.

NOTE

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/`<mask-bits>`” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show ip access-list** commands.

Since this ACL is input to an OSPF distribution list, the `<destination-ip>` parameter actually specifies the subnet mask of the route.

The `<wildcard>` parameter specifies the portion of the subnet mask to match against. For example, the `<destination-ip>` and `<wildcard>` values 255.255.255.255 0.0.0.255 mean that subnet mask /24 and longer match the ACL.

If you want the policy to match on all network masks, enter **any any**.

Modify default metric for redistribution

The default metric is a global parameter that specifies the cost applied to all OSPF routes by default. The default value is 10. You can assign a cost from 1 – 15.

NOTE

You also can define the cost on individual interfaces. The interface cost overrides the default cost.

To assign a default metric of 4 to all routes imported into OSPF, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# default-metric 4
```

Syntax: `default-metric <value>`

The `<value>` can be from 1 – 16,777,215. The default is 10.

Enable route redistribution

To enable route redistribution, use one of the following methods.

NOTE

Do not enable redistribution until you have configured the redistribution filters. Otherwise, you might accidentally overload the network with routes you did not intend to redistribute.

To enable redistribution of RIP and static IP routes into OSPF, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# redistribution rip
PowerConnect(config-ospf-router)# redistribution static
PowerConnect(config-ospf-router)# write memory
```

Example using a route map

To configure a route map and use it for redistribution of routes into OSPF, enter commands such as the following.

```
PowerConnect(config)# ip route 1.1.0.0 255.255.0.0 207.95.7.30
PowerConnect(config)# ip route 1.2.0.0 255.255.0.0 207.95.7.30
PowerConnect(config)# ip route 1.3.0.0 255.255.0.0 207.95.7.30
PowerConnect(config)# ip route 4.1.0.0 255.255.0.0 207.95.6.30
PowerConnect(config)# ip route 4.2.0.0 255.255.0.0 207.95.6.30
PowerConnect(config)# ip route 4.3.0.0 255.255.0.0 207.95.6.30
PowerConnect(config)# ip route 4.4.0.0 255.255.0.0 207.95.6.30 5
PowerConnect(config)# route-map abc permit 1
PowerConnect(config-routemap abc)# match metric 5
PowerConnect(config-routemap abc)# set metric 8
PowerConnect(config-routemap abc)# router ospf
PowerConnect(config-ospf-router)# redistribute static route-map abc
```

The commands in this example configure some static IP routes, then configure a route map and use the route map for redistributing static IP routes into OSPF.

The **ip route** commands configure the static IP routes. The **route-map** command begins configuration of a route map called “abc”. The number indicates the route map entry (called the “instance”) you are configuring. A route map can contain multiple entries. The software compares packets to the route map entries in ascending numerical order and stops the comparison once a match is found.

The **match** command in the route map matches on routes that have 5 for their metric value (cost). The **set** command changes the metric in routes that match the route map to 8.

The **redistribute static** command enables redistribution of static IP routes into OSPF, and uses route map “abc” to control the routes that are redistributed. In this example, the route map allows a static IP route to be redistributed into OSPF only if the route has a metric of 5, and changes the metric to 8 before placing the route into the OSPF route table.

The following command shows the result of the redistribution filter. Since only one of the static IP routes configured above matches the route map, only one route is redistributed. Notice that the route metric is 5 before redistribution but is 8 after redistribution.

```
PowerConnect# show ip ospf database external extensive
```

Index	Aging	LS ID	Router	Netmask	Metric	Flag
1	2	4.4.0.0	10.10.10.60	ffff0000	80000008	0000

Syntax: [no] redistribution bgp | connected | rip | static [route-map <map-name>]

The **bgp | connected | rip | static** parameter specifies the route source.

The **route-map <map-name>** parameter specifies the route map name. The following match parameters are valid for OSPF redistribution:

- **match ip address | next-hop <ACL-num>**
- **match metric <num>**
- **match tag <tag-value>**

The following set parameters are valid for OSPF redistribution:

- **set ip next hop <ip-addr>**
- **set metric [+ | -]<num> | none**
- **set metric-type type-1 | type-2**
- **set tag <tag-value>**

NOTE

You must configure the route map before you configure a redistribution filter that uses the route map.

NOTE

When you use a route map for route redistribution, the software disregards the permit or deny action of the route map.

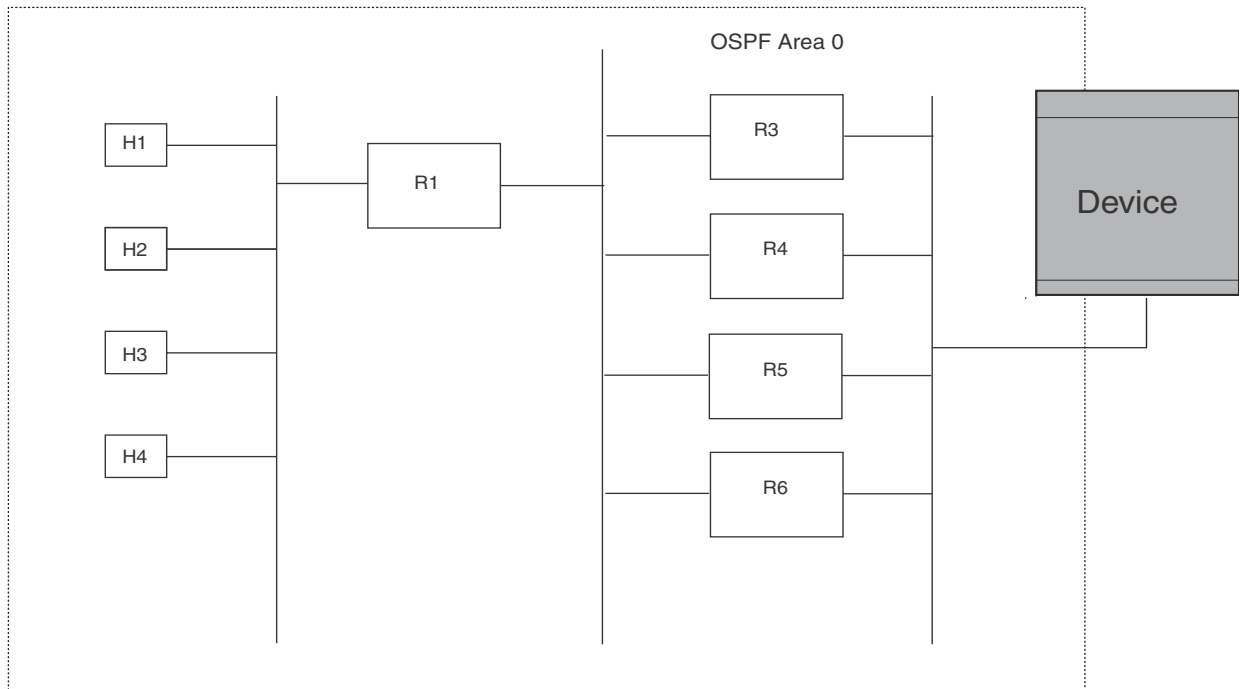
NOTE

For an external route that is redistributed into OSPF through a route map, the metric value of the route remains the same unless the metric is set by a **set metric** command inside the route map. The **default-metric <num>** command has no effect on the route. This behavior is different from a route that is redistributed without using a route map. For a route redistributed without using a route map, the metric is set by the **default-metric <num>** command.

Disable or re-enable load sharing

Routers can load share among up to eight equal-cost IP routes to a destination. By default, IP load sharing is enabled. The default is 4 equal-cost paths but you can specify from 2 – 6 paths.

The router software can use the route information it learns through OSPF to determine the paths and costs. Example OSPF network with four equal-cost paths



In the example in [Figure](#) , the switch has four paths to R1:

- Device->R3
- Device->R4
- Device->R5
- Device->R6

Normally, the switch will choose the path to the R1 with the lower metric. For example, if R3 metric is 1400 and R4 metric is 600, the switch will always choose R4.

However, suppose the metric is the same for all four routers in this example. If the costs are the same, the switch now has four equal-cost paths to R1. To allow the switch to load share among the equal cost routes, enable IP load sharing. The software supports four equal-cost OSPF paths by default when you enable load sharing. You can specify from 2 – 6 paths.

NOTE

The switch is not source routing in these examples. The switch is concerned only with the paths to the next-hop routers, not the entire paths to the destination hosts.

OSPF load sharing is enabled by default when IP load sharing is enabled. To configure IP load sharing parameters, refer to [“Configuring IP load sharing”](#) on page 605.

Configure external route summarization

When the Layer 3 Switch is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to advertise one external route as an aggregate for all redistributed routes that are covered by a specified address range.

When you configure an address range, the range takes effect immediately. All the imported routes are summarized according to the configured address range. Imported routes that have already been advertised and that fall within the range are flushed out of the AS and a single route corresponding to the range is advertised.

If a route that falls within a configured address range is imported by the Layer 3 Switch, no action is taken if the Layer 3 Switch has already advertised the aggregate route; otherwise the Layer 3 Switch advertises the aggregate route. If an imported route that falls within a configured address range is removed by the Layer 3 Switch, no action is taken if there are other imported routes that fall within the same address range; otherwise the aggregate route is flushed.

You can configure up to 32 address ranges. The Layer 3 Switch sets the forwarding address of the aggregate route to zero and sets the tag to zero.

If you delete an address range, the advertised aggregate route is flushed and all imported routes that fall within the range are advertised individually.

If an external LSDB overflow condition occurs, all aggregate routes are flushed out of the AS, along with other external routes. When the Layer 3 Switch exits the external LSDB overflow condition, all the imported routes are summarized according to the configured address ranges.

NOTE

If you use redistribution filters in addition to address ranges, the Layer 3 Switch applies the redistribution filters to routes first, then applies them to the address ranges.

NOTE

If you disable redistribution, all the aggregate routes are flushed, along with other imported routes.

To configure a summary address for OSPF routes, enter commands such as the following.

```
PowerConnect(config-ospf-router)# summary-address 10.1.0.0 255.255.0.0
```

The command in this example configures summary address 10.1.0.0, which includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. For all of these networks, only the address 10.1.0.0 is advertised in external LSAs.

Syntax: `summary-address <ip-addr> <ip-mask>`

The `<ip-addr>` parameter specifies the network address.

The `<ip-mask>` parameter specifies the network mask.

To display the configured summary addresses, enter the following command at any level of the CLI.

```
PowerConnect# show ip ospf config
OSPF Redistribution Address Ranges currently defined:
Range-Address      Subnetmask
1.0.0.0            255.0.0.0
1.0.1.0            255.255.255.0
1.0.2.0            255.255.255.0
```

Syntax: `show ip ospf config`

Configure default route origination

When the Layer 3 Switch is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to automatically generate a default external route into an OSPF routing domain. This feature is called “default route origination” or “default information origination”.

By default, Layer 3 Switches do not advertise the default route into the OSPF domain. If you want the Layer 3 Switch to advertise the OSPF default route, you must explicitly enable default route origination.

When you enable OSPF default route origination, the Layer 3 Switch advertises a type 5 default route that is flooded throughout the AS (except stub areas and NSSAs). In addition, internal NSSA ASBRs advertise their default routes as translatable type 7 default routes.

The Layer 3 Switch advertises the default route into OSPF even if OSPF route redistribution is not enabled, and even if the default route is learned through an IBGP neighbor.

NOTE

Layer 3 Switches never advertise the OSPF default route, regardless of other configuration parameters, unless you explicitly enable default route origination using the following method.

If the Layer 3 Switch is an ASBR, you can use the “always” option when you enable the default route origination. The always option causes the ASBR to create and advertise a default route if it does not already have one configured.

If default route origination is enabled and you disable it, the default route originated by the Layer 3 Switch is flushed. Default routes generated by other OSPF routers are not affected. If you re-enable the feature, the feature takes effect immediately and thus does not require you to reload the software.

NOTE

The ABR (Layer 3 Switch) will not inject the default route into an NSSA by default and the command described in this section will not cause the Layer 3 Switch to inject the default route into the NSSA. To inject the default route into an NSSA, use the **area <num> | <ip-addr> nssa default-information-originate** command. Refer to [“Assign a Not-So-Stubby Area \(NSSA\)”](#) on page 669.

To enable default route origination, enter the following command.

```
PowerConnect(config-ospf-router)# default-information-originate
```

To disable the feature, enter the following command.

```
PowerConnect(config-ospf-router)# no default-information-originate
```

Syntax: [no] default-information-originate [always] [metric <value>] [metric-type <type>]

The **always** parameter advertises the default route regardless of whether the router has a default route. This option is disabled by default.

The **metric <value>** parameter specifies a metric for the default route. If this option is not used, the default metric is used for the route.

The **metric-type <type>** parameter specifies the external link type associated with the default route advertised into the OSPF routing domain. The <type> can be one of the following:

- 1 – Type 1 external route
- 2 – Type 2 external route

If you do not use this option, the default redistribution metric type is used for the route type.

NOTE

If you specify a metric and metric type, the values you specify are used even if you do not use the **always** option.

Modify SPF timers

The Layer 3 Switch uses the following timers when calculating the shortest path for OSPF routes:

- **SPF delay** – When the Layer 3 Switch receives a topology change, the software waits before it starts a Shortest Path First (SPF) calculation. By default, the software waits five seconds. You can configure the SPF delay to a value from 0 – 65535 seconds. If you set the SPF delay to 0 seconds, the software immediately begins the SPF calculation after receiving a topology change.
- **SPF hold time** – The Layer 3 Switch waits for a specific amount of time between consecutive SPF calculations. By default, the Layer 3 Switch waits ten seconds. You can configure the SPF hold time to a value from 0 – 65535 seconds. If you set the SPF hold time to 0 seconds, the software does not wait between consecutive SPF calculations.

You can set the delay and hold time to lower values to cause the Layer 3 Switch to change to alternate paths more quickly in the event of a route failure. Note that lower values require more CPU processing time.

You can change one or both of the timers. To do so, enter commands such as the following.

```
PowerConnect(config-ospf-router)# timers spf 10 20
```

The command in this example changes the SPF delay to 10 seconds and changes the SPF hold time to 20 seconds.

Syntax: `timers spf <delay> <hold-time>`

The `<delay>` parameter specifies the SPF delay.

The `<hold-time>` parameter specifies the SPF hold time.

To set the timers back to their default values, enter a command such as the following.

```
PowerConnect(config-ospf-router)# no timers spf 10 20
```

Modify redistribution metric type

The redistribution metric type is used by default for all routes imported into OSPF unless you specify different metrics for individual routes using redistribution filters. Type 2 specifies a big metric (three bytes). Type 1 specifies a small metric (two bytes). The default value is type 2.

To modify the default value to type 1, enter the following command.

```
PowerConnect(config-ospf-router)# metric-type type1
```

Syntax: `metric-type type1 | type2`

The default is **type2**.

Modify administrative distance

Layer 3 Switches can learn about networks from various protocols, including Border Gateway Protocol version 4 (BGP4), RIP, and OSPF. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. The default administrative distance for OSPF routes is 110. Refer to [“Changing administrative distances”](#) on page 777 for a list of the default distances for all route sources.

The router selects one route over another based on the source of the route information. To do so, the router can use the administrative distances assigned to the sources. You can bias the Layer 3 Switch decision by changing the default administrative distance for RIP routes.

Configuring administrative distance based on route type

You can configure a unique administrative distance for each type of OSPF route. For example, you can use this feature to prefer a static route over an OSPF inter-area route but you also want to prefer OSPF intra-area routes to static routes.

The distance you specify influences the choice of routes when the Layer 3 Switch has multiple routes for the same network from different protocols. The Layer 3 Switch prefers the route with the lower administrative distance.

You can specify unique default administrative distances for the following route types:

- Intra-area routes
- Inter-area routes
- External routes

The default for all these OSPF route types is 110.

NOTE

This feature does not influence the choice of routes within OSPF. For example, an OSPF intra-area route is always preferred over an OSPF inter-area route, even if the intra-area route distance is greater than the inter-area route distance.

To change the default administrative distances for inter-area routes, intra-area routes, and external routes, enter the following command.

```
PowerConnect(config-ospf-router)# distance external 100
PowerConnect(config-ospf-router)# distance inter-area 90
PowerConnect(config-ospf-router)# distance intra-area 80
```

Syntax: `distance external | inter-area | intra-area <distance>`

The **external | inter-area | intra-area** parameter specifies the route type for which you are changing the default administrative distance.

The **<distance>** parameter specifies the new distance for the specified route type. Unless you change the distance for one of the route types using commands such as those shown above, the default is 110.

To reset the administrative distance to its system default (110), enter a command such as the following.

```
PowerConnect(config-ospf-router)# no distance external 100
```

Configure OSPF group Link State Advertisement (LSA) pacing

The Layer 3 Switch paces LSA refreshes by delaying the refreshes for a specified time interval instead of performing a refresh each time an individual LSA refresh timer expires. The accumulated LSAs constitute a group, which the Layer 3 Switch refreshes and sends out together in one or more packets.

The pacing interval, which is the interval at which the Layer 3 Switch refreshes an accumulated group of LSAs, is configurable to a range from 10 – 1800 seconds (30 minutes). The default is 240 seconds (four minutes). Thus, every four minutes, the Layer 3 Switch refreshes the group of accumulated LSAs and sends the group together in the same packets.

Usage guidelines

The pacing interval is inversely proportional to the number of LSAs the Layer 3 Switch is refreshing and aging. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval enhances performance. If you have a very small database (40 – 100 LSAs), increasing the pacing interval to 10 – 20 minutes might enhance performance slightly.

Changing the LSA pacing interval

To change the LSA pacing interval to two minutes (120 seconds), enter the following command.

```
PowerConnect(config-ospf-router)# timers lsa-group-pacing 120
```

Syntax: [no] **timers lsa-group-pacing** <secs>

The <secs> parameter specifies the number of seconds and can be from 10 – 1800 (30 minutes). The default is 240 seconds (four minutes).

To restore the pacing interval to its default value, enter the following command.

```
PowerConnect(config-ospf-router)# no timers lsa-group-pacing
```

Modify OSPF traps generated

OSPF traps as defined by RFC 1850 are supported on routers. OSPF trap generation is enabled on the router, by default.

When using the CLI, you can disable all or specific OSPF trap generation by entering the following CLI command.

```
PowerConnect(config-ospf-router)# no snmp-server trap ospf
```

To later re-enable the trap feature, enter **snmp-server trap ospf**.

To disable a specific OSPF trap, enter the command as **no snmp-server trap ospf <ospf-trap>**.

These commands are at the OSPF router Level of the CLI.

Here is a summary of OSPF traps supported on routers, their corresponding CLI commands, and their associated MIB objects from RFC 1850:

- **interface-state-change-trap** – [MIB object: OspfIfStateChange]
- **virtual-interface-state-change-trap** – [MIB object: OspfVirtIfStateChange]
- **neighbor-state-change-trap** – [MIB object: ospfNbrStateChange]

- **virtual-neighbor-state-change-trap** – [MIB object: ospfVirtNbrStateChange]
- **interface-config-error-trap** – [MIB object: ospflfConfigError]
- **virtual-interface-config-error-trap** – [MIB object: ospfVirtIfConfigError]
- **interface-authentication-failure-trap** – [MIB object: ospflfAuthFailure]
- **virtual-interface-authentication-failure-trap** – [MIB object: ospfVirtIfAuthFailure]
- **interface-receive-bad-packet-trap** – [MIB object: ospflfrxBadPacket]
- **virtual-interface-receive-bad-packet-trap** – [MIB object: ospfVirtIfRxBadPacket]
- **interface-retransmit-packet-trap** – [MIB object: ospfTxRetransmit]
- **virtual-interface-retransmit-packet-trap** – [MIB object: ospfVirtIfTxRetransmit]
- **originate-lsa-trap** – [MIB object: ospfOriginateLsa]
- **originate-maxage-lsa-trap** – [MIB object: ospfMaxAgeLsa]
- **link-state-database-overflow-trap** – [MIB object: ospfLsdbOverflow]
- **link-state-database-approaching-overflow-trap** – [MIB object: ospfLsdbApproachingOverflow]

Example

To stop an OSPF trap from being collected, use the CLI command: **no trap <ospf-trap>**, at the Router OSPF level of the CLI. To disable reporting of the neighbor-state-change-trap, enter the following command.

```
PowerConnect(config-ospf-router)# no trap neighbor-state-change-trap
```

Example

To reinstate the trap, enter the following command.

```
PowerConnect(config-ospf-router)# trap neighbor-state-change-trap
```

Syntax: [no] snmp-server trap ospf <ospf-trap>

Modify OSPF standard compliance setting

Routers are configured, by default, to be compliant with the RFC 1583 OSPF V2 specification.

To configure a router to operate with the latest OSPF standard, RFC 2178, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# no rfc1583-compatibility
```

Syntax: [no] rfc1583-compatibility

Modify exit overflow interval

If a database overflow condition occurs on a router, the router eliminates the condition by removing entries that originated on the router. The exit overflow interval allows you to set how often a Layer 3 Switch checks to see if the overflow condition has been eliminated. The default value is 0. The range is 0 – 86400 seconds (24 hours). If the configured value of the database overflow interval is zero, then the router never leaves the database overflow condition.

To modify the exit overflow interval to 60 seconds, enter the following command.

```
PowerConnect(config-ospf-router)# data-base-overflow-interval 60
```

Syntax: `database-overflow-interval <value>`

The `<value>` can be from 0 – 86400 seconds. The default is 0 seconds.

Specifying the types of OSPF Syslog messages to log

You can specify which kinds of OSPF-related Syslog messages are logged. By default, the only OSPF messages that are logged are those indicating possible system errors. If you want other kinds of OSPF messages to be logged, you can configure the device to log them.

For example, to specify that all OSPF-related Syslog messages be logged, enter the following commands.

```
PowerConnect(config)# router ospf
PowerConnect(config-ospf-router)# log all
```

Syntax: `[no] log all | adjacency | bad_packet [checksum] | database | memory | retransmit`

The **log** command has the following options:

The **all** option causes all OSPF-related Syslog messages to be logged. If you later disable this option with the **no log all** command, the OSPF logging options return to their default settings.

The **adjacency** option logs essential OSPF neighbor state changes, especially on error cases. This option is disabled by default.

The **bad_packet checksum** option logs all OSPF packets that have checksum errors. This option is enabled by default.

The **bad_packet** option logs all other bad OSPF packets. This option is disabled by default.

The **database** option logs OSPF LSA-related information. This option is disabled by default.

The **memory** option logs abnormal OSPF memory usage. This option is enabled by default.

The **retransmit** option logs OSPF retransmission activities. This option is disabled by default.

Clearing OSPF information

The following kinds of OSPF information can be cleared from a OSPF link state database and OSPF routing table:

- Routes received from OSPF neighbors. You can clear routes from all OSPF neighbors, or an individual OSPF neighbor, specified either by the neighbor IP address or its router ID
- OSPF topology information, including all routes in the OSPF routing table
- All routes in the OSPF routing table that were redistributed from other protocols
- OSPF area information, including routes received from OSPF neighbors within an area, as well as routes imported into the area. You can clear area information for all OSPF areas, or for a specified OSPF area

The OSPF information is cleared dynamically when you enter the command; you do not need to remove statements from the configuration or reload the software for the change to take effect.

Clearing OSPF neighbor information

To clear information on the device about all OSPF neighbors, enter the following command.

```
PowerConnect# clear ip ospf neighbor
```

Syntax: `clear ip ospf neighbor [ip <ip-addr> | id <ip-addr>]`

This command clears all OSPF neighbors and the OSPF routes exchanged with the neighbors in the OSPF link state database. After this information is cleared, adjacencies with all neighbors are re-established, and routes with these neighbors exchanged again.

To clear information on the device about OSPF neighbor 10.10.10.1, enter the following command.

```
PowerConnect# clear ip ospf neighbor ip 10.10.10.1
```

This command clears the OSPF neighbor and the OSPF routes exchanged with neighbor 10.10.10.1 in the OSPF link state database in the device. After this information is cleared, the adjacency with the neighbor is re-established, and routes are exchanged again.

The neighbor router can be specified either by its IP address or its router ID. To specify the neighbor router using its IP address, use the `ip <ip-addr>` parameter. To specify the neighbor router using its router ID, use the `id <ip-addr>` parameter.

Clearing OSPF topology information

To clear OSPF topology information on the device, enter the following command.

```
PowerConnect# clear ip ospf topology
```

Syntax: `clear ip ospf topology`

This command clears all OSPF routes from the OSPF routing table, including intra-area, (which includes ABR and ASBR intra-area routes), inter-area, external type 1, external type 2, OSPF default, and OSPF summary routes.

After you enter this command, the OSPF routing table is rebuilt, and valid routes are recomputed from the OSPF link state database. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated. If redistribution is enabled, the routes are imported again.

Clearing redistributed routes from the OSPF routing table

To clear all routes in the OSPF routing table that were redistributed from other protocols, enter the following command.

```
PowerConnect# clear ospf redistribution
```

Syntax: `clear ospf redistribution`

This command clears all routes in the OSPF routing table that are redistributed from other protocols, including direct connected, static, RIP, and BGP. To import redistributed routes from other protocols, use the `redistribution` command at the OSPF configuration level.

Clearing information for OSPF areas

To clear information on the device about all OSPF areas, enter the following command.

```
PowerConnect# clear ip ospf
```

This command clears all OSPF areas, all OSPF neighbors, and the entire OSPF routing table. After this information has been cleared, adjacencies with all neighbors are re-established, and all OSPF routes are re-learned.

To clear information on the device about OSPF area 1, enter the following command.

```
PowerConnect# clear ip ospf area 1
```

This command clears information about the specified area ID. Information about other OSPF areas is not affected. The command clears information about all OSPF neighbors belonging to the specified area, as well as all routes imported into the specified area. Adjacencies with neighbors belonging to the area are re-established, and routes imported into the area are re-learned.

Syntax: `clear ip ospf [area <area-id>]`

The <area-id> can be specified in decimal format or in IP address format.

Displaying OSPF information

You can use CLI commands to display the following OSPF information:

- Trap, area, and interface information – refer to [“Displaying general OSPF configuration information”](#) on page 697.
- CPU utilization statistics – refer to [“Displaying CPU utilization statistics”](#) on page 698.
- Area information – refer to [“Displaying OSPF area information”](#) on page 700.
- Neighbor information – refer to [“Displaying OSPF neighbor information”](#) on page 700.
- Interface information – refer to [“Displaying OSPF interface information”](#) on page 702.
- Route information – refer to [“Displaying OSPF route information”](#) on page 704.
- External link state information – refer to [“Displaying OSPF external link state information”](#) on page 706.
- Link state information – refer to [“Displaying OSPF link state information”](#) on page 707.
- Virtual Neighbor information – refer to [“Displaying OSPF virtual neighbor information”](#) on page 708.
- Virtual Link information – refer to [“Displaying OSPF virtual link information”](#) on page 708.
- ABR and ASBR information – refer to [“Displaying OSPF ABR and ASBR information”](#) on page 708.
- Trap state information – refer to [“Displaying OSPF trap status”](#) on page 709.

Displaying general OSPF configuration information

To display general OSPF configuration information, enter the following command at any CLI level.

23 Displaying OSPF information

```
PowerConnect# show ip ospf config
Router OSPF: Enabled
Redistribution: Disabled
Default OSPF Metric: 10
OSPF Redistribution Metric: Type2

OSPF External LSA Limit: 25000

OSPF Database Overflow Interval: 0

RFC 1583 Compatibility: Enabled

Router id: 207.95.11.128

Interface State Change Trap: Enabled
Virtual Interface State Change Trap: Enabled
Neighbor State Change Trap: Enabled
Virtual Neighbor State Change Trap: Enabled
Interface Configuration Error Trap: Enabled
Virtual Interface Configuration Error Trap: Enabled
Interface Authentication Failure Trap: Enabled
Virtual Interface Authentication Failure Trap: Enabled
Interface Receive Bad Packet Trap: Enabled
Virtual Interface Receive Bad Packet Trap: Enabled
Interface Retransmit Packet Trap: Enabled
Virtual Interface Retransmit Packet Trap: Enabled
Originate LSA Trap: Enabled
Originate MaxAge LSA Trap: Enabled
Link State Database Overflow Trap: Enabled
Link State Database Approaching Overflow Trap: Enabled

OSPF Area currently defined:
Area-ID          Area-Type Cost
0                 normal   0

OSPF Interfaces currently defined:
Ethernet Interface: 1-2
ip ospf md5-authentication-key-activation-wait-time 300
ip ospf cost 0
ip ospf area 0

Ethernet Interface: v1
ip ospf md5-authentication-key-activation-wait-time 300
ip ospf cost 0
ip ospf area 0
```

Syntax: show ip ospf config

Displaying CPU utilization statistics

You can display CPU utilization statistics for OSPF and other IP protocols.

To display CPU utilization statistics for OSPF for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.


```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.03       0.09       0.22        9
BGP             0.04       0.06       0.08       0.14       13
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.00       0.00       0.00       0.00        0
IP             0.00       0.00       0.00       0.00        0
OSPF         0.03     0.06     0.09     0.12     11
RIP            0.00       0.00       0.00       0.00        0
STP            0.00       0.00       0.00       0.00        0
VRRP           0.00       0.00       0.00       0.00        0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example.

```
PowerConnect# show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.00       0.00       0.00        0
BGP             0.00       0.00       0.00       0.00        0
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.01       0.00       0.00       0.00        1
IP             0.00       0.00       0.00       0.00        0
OSPF           0.00       0.00       0.00       0.00        0
RIP            0.00       0.00       0.00       0.00        0
STP            0.00       0.00       0.00       0.00        0
VRRP           0.00       0.00       0.00       0.00        0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.

```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name    Sec(%)    Time(ms)
ARP             0.00      0
BGP             0.00      0
GVRP           0.00      0
ICMP           0.01      1
IP             0.00      0
OSPF           0.00      0
RIP            0.00      0
STP            0.01      0
VRRP           0.00      0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The `<num>` parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Displaying OSPF area information

To display OSPF area information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf area
Indx Area      Type Cost  SPFR ABR ASBR LSA Chksum(Hex)
1   0.0.0.0    normal 0     1    0   0   1   0000781f
2  192.147.60.0 normal 0     1    0   0   1   0000fee6
3  192.147.80.0 stub   1     1    0   0   2   000181cd
```

Syntax: `show ip ospf area [<area-id>] | [<num>]`

The `<area-id>` parameter shows information for the specified area.

The `<num>` parameter displays the entry that corresponds to the entry number you enter. The entry number identifies the entry position in the area table.

This display shows the following information.

TABLE 114 CLI display of OSPF area information

This field...	Displays...
Indx	The row number of the entry in the router OSPF area table.
Area	The area number.
Type	The area type, which can be one of the following: <ul style="list-style-type: none"> • nssa • normal • stub
Cost	The area cost.
SPFR	The SPFR value.
ABR	The ABR number.
ASBR	The ABSR number.
LSA	The LSA number.
Chksum(Hex)	The checksum for the LSA packet. The checksum is based on all the fields in the packet except the age field. The Layer 3 Switch uses the checksum to verify that the packet is not corrupted.

Displaying OSPF neighbor information

To display OSPF neighbor information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf neighbor
Port Address      Pri State      Neigh Address  Neigh ID
8   212.76.7.251    1   full        212.76.7.200  173.35.1.220
```

To display detailed OSPF neighbor information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf neighbor detail
```

```
Port      Address      Pri State      Neigh Address  Neigh ID      Ev Op Cnt
1         20.2.0.2     1  FULL/DR    20.2.0.1      2.2.2.2      6  2  0
  Second-to-dead:39
1         20.3.0.2     1  FULL/BDR    20.3.0.1      3.3.3.3      5  2  0
  Second-to-dead:36
1-8      23.5.0.1     1  FULL/DR    23.5.0.2      16.16.16.16  6  2  0
  Second-to-dead:33
1-2      23.2.0.1     1  FULL/DR    23.2.0.2      15.15.15.15  6  2  0
  Second-to-dead:33
```

Syntax: `show ip ospf neighbor [router-id <ip-addr>] | [<num>] | [detail]`

The **router-id <ip-addr>** parameter displays only the neighbor entries for the specified router.

The **<num>** parameter displays only the entry in the specified index position in the neighbor table. For example, if you enter “1”, only the first entry in the table is displayed.

The **detail** parameter displays detailed information about the neighbor routers.

These displays show the following information.

TABLE 115 CLI display of OSPF neighbor information

Field	Description
Port	The port through which the Layer 3 Switch is connected to the neighbor. The port on which an OSPF point-to-point link is configured.
Address	The IP address of this Layer 3 Switch interface with the neighbor.
Pri	The OSPF priority of the neighbor: <ul style="list-style-type: none"> • For multi-access networks, the priority is used during election of the Designated Router (DR) and Backup designated Router (BDR). • For point-to-point links, this field shows one of the following values: • 1 = point-to-point link • 3 = point-to-point link with assigned subnet

TABLE 115 CLI display of OSPF neighbor information (Continued)

Field	Description
State	<p>The state of the conversation between the Layer 3 Switch and the neighbor. This field can have one of the following values:</p> <ul style="list-style-type: none"> • Down – The initial state of a neighbor conversation. This value indicates that there has been no recent information received from the neighbor. • Attempt – This state is only valid for neighbors attached to non-broadcast networks. It indicates that no recent information has been received from the neighbor. • Init – A Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor. (The router itself did not appear in the neighbor's Hello packet.) All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface. • 2-Way – Communication between the two routers is bidirectional. This is the most advanced state before beginning adjacency establishment. The Designated Router and Backup Designated Router are selected from the set of neighbors in the 2-Way state or greater. • ExStart – The first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description (DD) sequence number. Neighbor conversations in this state or greater are called adjacencies. • Exchange – The router is describing its entire link state database by sending Database Description packets to the neighbor. Each Database Description packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description packet can be outstanding at any time. In this state, Link State Request packets can also be sent asking for the neighbor's more recent advertisements. All adjacencies in Exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets. • Loading – Link State Request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the Exchange state. • Full – The neighboring routers are fully adjacent. These adjacencies will now appear in router links and network link advertisements.
Neigh Address	<p>The IP address of the neighbor:</p> <p>For point-to-point links, the value is as follows:</p> <ul style="list-style-type: none"> • If the Pri field is "1", this value is the IP address of the neighbor router interface. • If the Pri field is "3", this is the subnet IP address of the neighbor router interface.
Neigh ID	The neighbor router ID.
Ev	The number of times the neighbor state changed.
Opt	The sum of the option bits in the Options field of the Hello packet. This information is used by Dell technical support. Refer to Section A.2 in RFC 2178 for information about the Options field in Hello packets.
Cnt	The number of LSAs that were retransmitted.
Second-to-dead	The amount of time the device will wait for a HELLO message from each OSPF neighbor before assuming the neighbor is dead.

Displaying OSPF interface information

To display OSPF interface information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf interface 192.168.1.1
```

```
Ethernet 1,OSPF enabled
  IP Address 192.168.1.1, Area 0
  OSPF state ptr2ptr, Pri 1, Cost 1, Options 2, Type pt-2-pt Events 1
  Timers(sec): Transit 1, Retrans 5, Hello 10, Dead 40
  DR: Router ID 0.0.0.0           Interface Address 0.0.0.0
  BDR: Router ID 0.0.0.0         Interface Address 0.0.0.0
  Neighbor Count = 0, Adjacent Neighbor Count= 1
  Neighbor: 2.2.2.2
  Authentication-Key:None
  MD5 Authentication: Key None, Key-Id None, Auth-change-wait-time 300
```

Syntax: `show ip ospf interface [<ip-addr>]`

The `<ip-addr>` parameter displays the OSPF interface information for the specified IP address.

The following table defines the highlighted fields shown in the above example output of the `show ip ospf interface` command.

TABLE 116 Output of the `show ip ospf interface` command

This field	Displays
IP Address	The IP address of the interface.
OSPF state	ptr2ptr (point to point)
Pri	The link ID as defined in the router-LSA. This value can be one of the following: <ul style="list-style-type: none"> • 1 = point-to-point link • 3 = point-to-point link with an assigned subnet
Cost	The configured output cost for the interface.
Options	OSPF Options (Bit7 - Bit0): <ul style="list-style-type: none"> • unused:1 • opaque:1 • summary:1 • dont_propagate:1 • nssa:1 • multicast:1 • externals:1 • tos:1
Type	The area type, which can be one of the following: <ul style="list-style-type: none"> • Broadcast = 0x01 • NBMA = 0x02 • Point to Point = 0x03 • Virtual Link = 0x04 • Point to Multipoint = 0x05

TABLE 116 Output of the `show ip ospf interface` command (Continued)

This field	Displays
Events	OSPF Interface Event: <ul style="list-style-type: none"> • Interface_Up = 0x00 • Wait_Timer = 0x01 • Backup_Seen = 0x02 • Neighbor_Change = 0x03 • Loop_Indication = 0x04 • Unloop_Indication = 0x05 • Interface_Down = 0x06 • Interface_Passive = 0x07
Adjacent Neighbor Count	The number of adjacent neighbor routers.
Neighbor:	The neighbor router ID.

Displaying OSPF route information

To display OSPF route information for the router, enter the following command at any CLI level.

```
PowerConnect# show ip ospf routes
Index Destination      Mask                Path_Cost Type2_Cost Path_Type
1      212.95.7.0           255.255.255.0      1          0          Intra
      Adv_Router      Link_State          Dest_Type State      Tag      Flags
      173.35.1.220     212.95.7.251      Network Valid    00000000 7000
      Paths Out_Port Next_Hop            Type      Arp_Index State
      1          6          209.95.7.250     OSPF      8          84 00

Index Destination      Mask                Path_Cost Type2_Cost Path_Type
2      11.3.63.0            255.255.255.0      11         0          Inter
      Adv_Router      Link_State          Dest_Type State      Tag      Flags
      209.95.7.250    11.3.63.0          Network Valid    00000000 0000
      Paths Out_Port Next_Hop            Type      Arp_Index State
      1          6          209.95.7.250     OSPF      8          84 00
```

Syntax: `show ip ospf routes [<ip-addr>]`

The `<ip-addr>` parameter specifies a destination IP address. If you use this parameter, only the route entries for that destination are shown.

This display shows the following information.

TABLE 117 CLI Display of OSPF route information

This field...	Displays...
Index	The row number of the entry in the router OSPF route table.
Destination	The IP address of the route's destination.
Mask	The network mask for the route.
Path_Cost	The cost of this route path. (A route can have multiple paths. Each path represents a different exit port for the Layer 3 Switch.)
Type2_Cost	The type 2 cost of this path.

TABLE 117 CLI Display of OSPF route information (Continued)

This field...	Displays...
Path_Type	The type of path, which can be one of the following: <ul style="list-style-type: none"> • Inter – The path to the destination passes into another area. • Intra – The path to the destination is entirely within the local area. • External1 – The path to the destination is a type 1 external route. • External2 – The path to the destination is a type 2 external route.
Adv_Router	The OSPF router that advertised the route to this Layer 3 Switch.
Link-State	The link state from which the route was calculated.
Dest_Type	The destination type, which can be one of the following: <ul style="list-style-type: none"> • ABR – Area Border Router • ASBR – Autonomous System Boundary Router • Network – the network
State	The route state, which can be one of the following: <ul style="list-style-type: none"> • Changed • Invalid • Valid This information is used by Dell technical support.
Tag	The external route tag.
Flags	State information for the route entry. This information is used by Dell technical support.
Paths	The number of paths to the destination.
Out_Port	The router port through which the Layer 3 Switch reaches the next hop for this route path.
Next_Hop	The IP address of the next-hop router for this path.
Type	The route type, which can be one of the following: <ul style="list-style-type: none"> • OSPF • Static Replaced by OSPF
Arp_Index	The index position in the ARP table of the ARP entry for this path's IP address.
State	State information for the path. This information is used by Dell technical support.

Displaying the routes that have been redistributed into OSPF

You can display the routes that have been redistributed into OSPF. To display the redistributed routes, enter the following command at any level of the CLI.

```
PowerConnect# show ip ospf redistribute route
 4.3.0.0 255.255.0.0 static
 3.1.0.0 255.255.0.0 static
10.11.61.0 255.255.255.0 connected
 4.1.0.0 255.255.0.0 static
```

In this example, four routes have been redistributed. Three of the routes were redistributed from static IP routes and one route was redistributed from a directly connected IP route.

Syntax: `show ip ospf redistribute route [<ip-addr> <ip-mask>]`

The <ip-addr> <ip-mask> parameter specifies a network prefix and network mask. Here is an example.

```
PowerConnect# show ip ospf redistribute route 3.1.0.0 255.255.0.0
3.1.0.0 255.255.0.0 static
```

Displaying OSPF external link state information

To display external link state information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf database external-link-state
Index Aging  LS ID           Router           Netmask  Metric  Flag
1      1794  1.168.64.0     192.85.0.3     fffffe00 000003e8 b500 0.0.0.0
2      1794  3.215.0.0     192.85.0.3     ffff0000 000003e8 b500 0.0.0.0
3      1794  1.27.250.0    192.85.0.3     fffffe00 000003e8 b500 0.0.0.0
4      1794  1.24.23.0     192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
5      1794  1.21.52.0     192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
6      1794  1.18.81.0     192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
7      1794  1.15.110.0    192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
8      1794  1.12.139.0    192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
9      1794  1.9.168.0     192.85.0.3     ffffff00 000003e8 b500 0.0.0.0
```

Syntax: `show ip ospf database external-link-state [advertise <num>] | [extensive] | [link-state-id <ip-addr>] | [router-id <ip-addr>] | [sequence-number <num(Hex)>] | [status <num>]`

The **advertise <num>** parameter displays the hexadecimal data in the specified LSA packet. The **<num>** parameter identifies the LSA packet by its position in the router External LSA table. To determine an LSA packet position in the table, enter the **show ip ospf external-link-state** command to display the table. Refer to “[Displaying the data in an LSA](#)” on page 707 for an example.

The **extensive** option displays the LSAs in decrypted format.

NOTE

You cannot use the **extensive** option in combination with other display options. The entire database is displayed.

The **link-state-id <ip-addr>** parameter displays the External LSAs for the LSA source specified by **<IP-addr>**.

The **router-id <ip-addr>** parameter shows the External LSAs for the specified OSPF router.

The **sequence-number <num(Hex)>** parameter displays the External LSA entries for the specified hexadecimal LSA sequence number.

The **status <num>** option shows status information.

This display shows the following information.

TABLE 118 CLI display of OSPF external link state information

This field...	Displays...
Area ID	The OSPF area the router is in.
Aging	The age of the LSA, in seconds.
LS ID	The ID of the link-state advertisement from which the Layer 3 Switch learned this route.
Router	The router IP address.

TABLE 118 CLI display of OSPF external link state information (Continued)

This field...	Displays...
Seq(hex)	The sequence number of the LSA. The OSPF neighbor that sent the LSA stamps it with a sequence number to enable the Layer 3 Switch and other OSPF routers to determine which LSA for a given route is the most recent.
Chksum	A checksum for the LSA packet, which is based on all the fields in the packet except the age field. The Layer 3 Switch uses the checksum to verify that the packet is not corrupted.
Type	The route type, which is always EXTR (external).

Displaying OSPF link state information

To display link state information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf database link-state
```

Syntax: `show ip ospf database link-state [advertise <num>] | [asbr] | [extensive] | [link-state-id <ip-addr>] | [network] | [nssa] | [opaque-area] | [router] | [router-id <ip-addr>] | [sequence-number <num(Hex)>] | [status <num>] | [summary]`

The **advertise <num>** parameter displays the hexadecimal data in the specified LSA packet. The **<num>** parameter identifies the LSA packet by its position in the router External LSA table. To determine an LSA packet position in the table, enter the **show ip ospf external-link-state** command to display the table. Refer to “[Displaying the data in an LSA](#)” on page 707 for an example.

The **asbr** option shows ASBR information.

The **extensive** option displays the LSAs in decrypted format.

NOTE

You cannot use the **extensive** option in combination with other display options. The entire database is displayed.

The **link-state-id <ip-addr>** parameter displays the External LSAs for the LSA source specified by **<IP-addr>**.

The **network** option shows network information.

The **nssa** option shows network information.

The **opaque-area** option shows information for opaque areas.

The **router-id <ip-addr>** parameter shows the External LSAs for the specified OSPF router.

The **sequence-number <num(Hex)>** parameter displays the External LSA entries for the specified hexadecimal LSA sequence number.

The **status <num>** option shows status information.

The **summary** option shows summary information.

Displaying the data in an LSA

You can use the CLI to display the data the Layer 3 Switch received in a specific External LSA packet or other type of LSA packet. For example, to display the LSA data in entry 3 in the External LSA table, enter the following command.

```

Index Aging  LS ID           Router           Netmask  Metric   Flag
3      619    1.27.250.0      192.85.0.3      fffffe00 000003e8 b500 0.0.0.0
  LSA Header: age: 619, options: 0x02, seq-nbr: 0x80000003, length: 36
  NetworkMask: 255.255.254.0
  TOS 0: metric_type: 1, metric: 1000
         forwarding_address: 0.0.0.0
         external_route_tag: 0

```

Syntax: `show ip ospf database external-link-state [advertise <num>] | [[link-state-id <ip-addr>] | [router-id <ip-addr>] | [sequence-number <num(Hex)>] | [status <num>]`

To determine an external LSA or other type of LSA index number, enter one of the following commands to display the appropriate LSA table:

- **show ip ospf database link-state advertise <num>** – This command displays the data in the packet for the specified LSA.
- **show ip ospf database external-link-state advertise <num>** – This command displays the data in the packet for the specified external LSA.

For example, to determine an external LSA index number, enter the following command.

```

PowerConnect# show ip ospf external-link-state
Index Aging  LS ID           Router           Netmask  Metric   Flag
1      1809   1.18.81.0      103.103.103.6   ffffffff00 000003e8 b500 0.0.0.0
2       8    1.27.250.0     103.103.103.6   fffffe00 000003e8 b500 0.0.0.0
3       8    3.215.0.0      103.103.103.6   ffff0000 000003e8 b500 0.0.0.0
4       18    1.33.192.0     102.102.102.6   fffffc00 000003e8 b500 0.0.0.0
5      959    1.9.168.0      102.102.102.6   ffffffff00 00002710 b500 0.0.0.0
6     1807   1.3.226.0      192.85.0.3      ffffffff00 000003e8 b500 0.0.0.0
7     1809   1.6.197.0      192.85.3.3      ffffffff00 000003e8 b500 0.0.0.0

```

Displaying OSPF virtual neighbor information

To display OSPF virtual neighbor information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf virtual-neighbor
```

Syntax: `show ip ospf virtual-neighbor [<num>]`

The <num> parameter displays the table beginning at the specified entry number.

Displaying OSPF virtual link information

To display OSPF virtual link information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf virtual-link
```

Syntax: `show ip ospf virtual-link [<num>]`

The <num> parameter displays the table beginning at the specified entry number.

Displaying OSPF ABR and ASBR information

To display OSPF ABR and ASBR information, enter the following command at any CLI level.

```
PowerConnect# show ip ospf border-routers
```

Syntax: `show ip ospf border-routers [<ip-addr>]`

The `<ip-addr>` parameter displays the ABR and ASBR entries for the specified IP address.

Displaying OSPF trap status

All traps are enabled by default when you enable OSPF. To disable or re-enable an OSPF trap, refer to [“Modify OSPF traps generated”](#) on page 693.

To display the state of each OSPF trap, enter the following command at any CLI level.

```
PowerConnect# show ip ospf trap
Interface State Change Trap:           Enabled
Virtual Interface State Change Trap:    Enabled
Neighbor State Change Trap:           Enabled
Virtual Neighbor State Change Trap:     Enabled
Interface Configuration Error Trap:     Enabled
Virtual Interface Configuration Error Trap: Enabled
Interface Authentication Failure Trap:  Enabled
Virtual Interface Authentication Failure Trap: Enabled
Interface Receive Bad Packet Trap:     Enabled
Virtual Interface Receive Bad Packet Trap: Enabled
Interface Retransmit Packet Trap:       Enabled
Virtual Interface Retransmit Packet Trap: Enabled
Originate LSA Trap:                   Enabled
Originate MaxAge LSA Trap:            Enabled
Link State Database Overflow Trap:      Enabled
Link State Database Approaching Overflow Trap: Enabled
```

Syntax: `show ip ospf trap`

23 Displaying OSPF information

Configuring VRRP and VRRPE

This chapter describes how to configure Layer 3 Switches with the following router redundancy protocols:

- **Virtual Router Redundancy Protocol (VRRP)** – The standard router redundancy protocol described in RFC 2338.
- **VRRP Extended (VRRPE)** – An enhanced version of VRRP that overcomes limitations in the standard protocol.

NOTE

VRRP and VRRPE are separate protocols. You cannot use them together.

NOTE

You can use a Layer 3 Switch configured for VRRP with another Layer 3 Switch or a third-party router that is also configured for VRRP. However, you can use a Layer 3 Switch configured for VRRPE only with another Layer 3 Switch that also is configured for VRRPE.

For a summary of how these two router redundancy protocols differ, refer to [“Comparison of VRRP and VRRPE”](#) on page 719.

Overview

The following sections describe VRRP and VRRPE. The protocols both provide redundant paths for IP addresses. However, the protocols differ in a few important ways. For clarity, each protocol is described separately.


Overview of VRRP

NOTE

VRRP support in the base Layer 3 code is the same as in the full Layer 3 code.

VRRP is a protocol that provides redundancy to routers within a LAN. VRRP allows you to provide alternate router paths for a host without changing the IP address or MAC address by which the host knows its gateway. Consider the situation shown in [Figure 112](#).

FIGURE 112 Switch 1 is Host1 default gateway but is a single point of failure

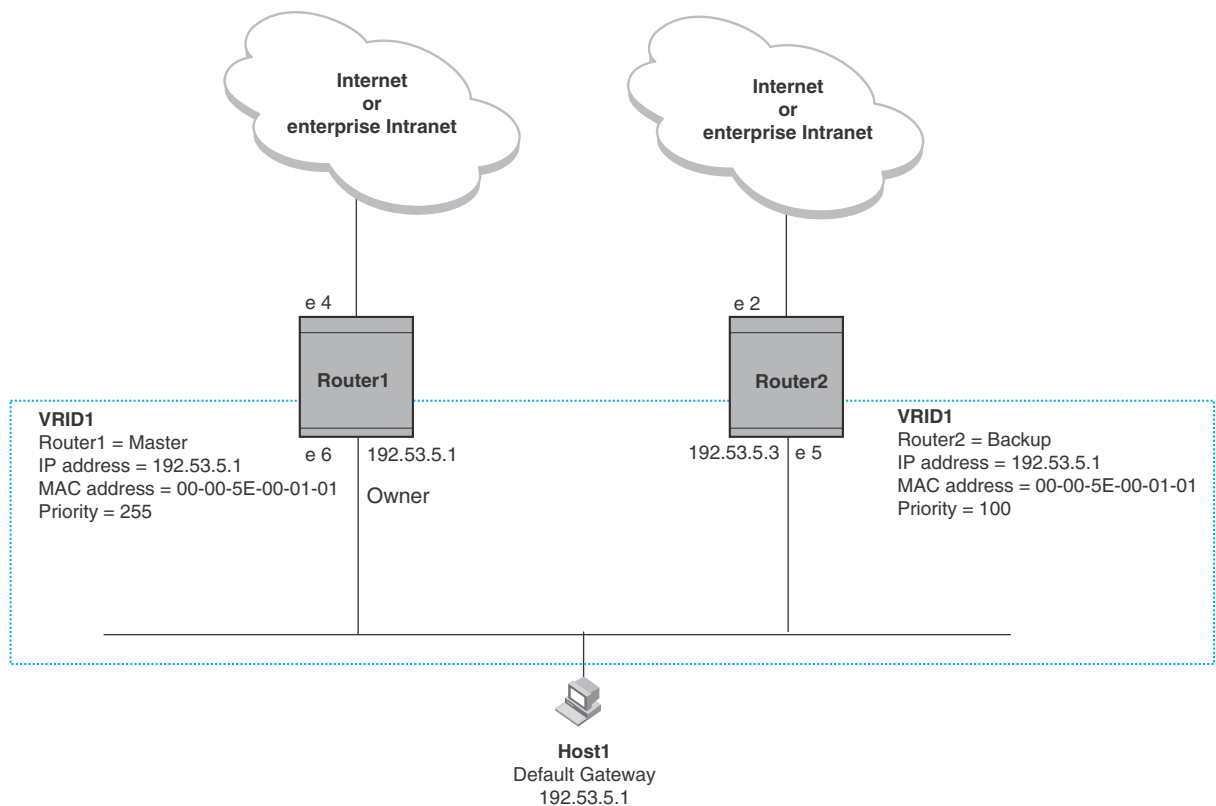


Switch 1 as the host default gateway out of the subnet. If this interface goes down, Host1 is cut off from the rest of the network. Switch 1 is thus a single point of failure for Host1 access to other networks.

If Switch 1 fails, you could configure Host1 to use Switch 2. Configuring one host with a different default gateway might not require too much extra administration. However, consider a more realistic network with dozens or even hundreds of hosts per subnet; reconfiguring the default gateways for all the hosts is impractical. It is much simpler to configure a VRRP virtual router on Switch 1 and Switch 2 to provide a redundant path for the hosts.

Figure 113 shows the same example network shown in Figure 112, but with a VRRP virtual router configured on Switch 1 and Switch 2.

FIGURE 113 Switch 1 and Switch 2 are configured as a VRRP virtual router for redundant network access for Host1



The dashed box in Figure 113 represents a VRRP virtual router. When you configure a virtual router, one of the configuration parameters is the virtual router ID (VRID), which can be a number from 1 – 255. In this example, the VRID is 1.

NOTE

You can provide more redundancy by also configuring a second VRID with Switch 2 as the Owner and Switch 1 as the Backup. This type of configuration is sometimes called Multigroup VRRP.

Virtual Router ID (VRID)

A **VRID** consists of one Master router and one or more Backup routers. The Master router is the router that owns the IP address(es) you associate with the VRID. For this reason, the Master router is sometimes called the “Owner”. Configure the VRID on the router that owns the default gateway interface. The other router in the VRID does not own the IP address(es) associated with VRID but provides the backup path if the Master router becomes unavailable.

Virtual router MAC address

Notice the MAC address associated with VRID1. The first five octets of the address are the standard MAC prefix for VRRP packets, as described in RFC 2338. The last octet is the VRID. THE VRID number becomes the final octet in the virtual MAC address associated with the virtual router.

When you configure a VRID, the software automatically assigns its MAC address. When a VRID becomes active, the Master router broadcasts a gratuitous ARP request containing the virtual router MAC address for each IP address associated with the virtual router. In [Figure 113](#), Switch 1 sends a gratuitous ARP with MAC address 00-00-5e-00-01-01 and IP address 192.53.5.1. Hosts use the virtual router MAC address in routed traffic they send to their default IP gateway (in this example, 192.53.5.1).

Virtual router IP address

VRRP does not use virtual IP addresses. Thus, there is no virtual IP address associated with a virtual router. Instead, you associate the virtual router with one or more real interface IP addresses configured on the router that owns the real IP address(es). In [Figure 113](#), the virtual router with VRID1 is associated with real IP address 192.53.5.1, which is configured on interface e 6 on Switch 1. VRIDs are interface-level parameters, not system-level parameters, so the IP address you associate with the VRID must already be a real IP address configured on the Owner interface.

NOTE

You also can associate a virtual router with a virtual interface. A virtual interface is a named set of physical interfaces.

When you configure the Backup router for the VRID, specify the same IP address as the one you specify on the Owner. This is the IP address used by the host as its default gateway. The IP address cannot also exist on the Backup router. The interface on which you configure the VRID on the Backup router must have an IP address in the same subnet.

NOTE

If you delete a real IP address used by a VRRP entry, the VRRP entry also is deleted automatically.

NOTE

When a Backup router takes over forwarding responsibilities from a failed Master router, the Backup forwards traffic addressed to the VRID MAC address, which the host believes is the MAC address of the router interface for its default gateway. However, the Backup cannot reply to IP pings sent to the IP address(es) associated with the VRID. Because the IP address(es) are owned by the Owner, if the Owner is unavailable, the IP addresses are unavailable as packet destinations.

Master negotiation

The routers within a VRID use the VRRP priority values associated with each router to determine which router becomes the Master. When you configure the VRID on a router interface, you specify whether the router is the Owner of the IP addresses you plan to associate with the VRID or a Backup. If you indicate that the router is the Owner of the IP addresses, the software automatically sets the router VRRP priority for the VRID to 255, the highest VRRP priority. The router with the highest priority becomes the Master.

Backup routers can have a priority from 3 – 254, which you assign when you configure the VRID on the Backup router interfaces. The default VRRP priority for Backup routers is 100.

Because the router that owns the IP addresses associated with the VRID always has the highest priority, when all the routers in the virtual router are operating normally, the negotiation process results in the Owner of the VRID IP addresses becoming the Master router. Thus, the VRRP negotiation results in the normal case, in which the hosts' path to the default route is to the router that owns the interface for that route.

Hello messages

VRRP routers use Hello messages for negotiation to determine the Master router. VRRP routers send Hello messages to IP Multicast address 224.0.0.18. The frequency with which the Master sends Hello messages is the Hello Interval. Only the Master sends Hello messages. However, a Backup uses the Hello interval you configure for the Backup if it becomes the Master.

The Backup routers wait for a period of time called the Dead Interval for a Hello message from the Master. If a Backup router does not receive a Hello message by the time the dead interval expires, the Backup router assumes that the Master router is dead and negotiates with the other Backups to select a new Master router. The Backup router with the highest priority becomes the new Master.

If the Owner becomes unavailable, but then comes back online, the Owner again becomes the Master router. The Owner becomes the Master router again because it has the highest priority. The Owner always becomes the Master again when the Owner comes back online.

NOTE

If you configure a track port on the Owner and the track port is down, the Owner priority is changed to the track priority. In this case, the Owner does not have a higher priority than the Backup that is acting as Master and the Owner therefore does not resume its position as Master. For more information about track ports, refer to [“Track ports and track priority”](#) on page 715.

By default, if a Backup is acting as the Master, and the Master is still unavailable, another Backup can “preempt” the Backup that is acting as the Master. This can occur if the new Backup has a higher priority than the Backup who is acting as Master. You can disable this behavior if you want. When you disable preemption, a Backup router that has a higher priority than the router who is currently acting as Master does not preempt the new Master by initiating a new Master negotiation. Refer to [“Backup preempt”](#) on page 729.

NOTE

Regardless of the setting for the preempt parameter, the Owner always becomes the Master again when it comes back online.

Track ports and track priority

The VRRP enhances the protocol by giving a VRRP router the capability to monitor the state of the interfaces on the other end of the route path through the router. For example, in [Figure 113](#) on page 713, interface e 6 on Switch 1 owns the IP address to which Host1 directs route traffic on its default gateway. The exit path for this traffic is through Router1 e 4 interface.

Suppose interface e 4 goes down. Even if interface e 6 is still up, Host1 is nonetheless cut off from other networks. In conventional VRRP, Switch 1 would continue to be the Master router despite the unavailability of the exit interface for the path the router is supporting. However, if you configure interface e 6 to track the state of interface e 4, if e 4 goes down, interface e 6 responds by changing Switch 1 VRRP priority to the value of the track priority. In the configuration shown in

Figure 113 on page 713, Switch 1 priority changes from 255 to 20. One of the parameters contained in the Hello messages the Master router sends to its Backups is the Master router priority. If the track port feature results in a change in the Master router priority, the Backup routers quickly become aware of the change and initiate a negotiation for Master router.

In Figure 113 on page 713, the track priority results in Switch 1 VRRP priority becoming lower than Switch 2 VRRP priority. As a result, when Switch 2 learns that it now has a higher priority than Switch 1, Switch 2 initiates negotiation for Master router and becomes the new Master router, thus providing an open path for Host1 traffic. To take advantage of the track port feature, make sure the track priorities are always lower than the VRRP priorities. The default track priority for the router that owns the VRID IP addresses is 2. The default track priority for Backup routers is 1. If you change the track port priorities, make sure you assign a higher track priority to the Owner of the IP addresses than the track priority you assign on the Backup routers.

Suppression of RIP advertisements for backed up interfaces

The VRRP allows you to configure the protocol to suppress RIP advertisements for the backed up paths from Backup routers. Normally, a VRRP Backup router includes route information for the interface it is backing up in RIP advertisements. As a result, other routers receive multiple paths for the interface and might sometimes unsuccessfully use the path to the Backup rather than the path to the Master. If you enable the Dell implementation of VRRP to suppress the VRRP Backup routers from advertising the backed up interface in RIP, other routers learn only the path to the Master router for the backed up interface.

Authentication

The VRRP can use simple passwords to authenticate VRRP packets. The VRRP authentication type is not a parameter specific to the VRID. Instead, VRRP uses the authentication type associated with the interfaces on which you define the VRID. For example, if you configure your router interfaces to use a simple password to authenticate traffic, VRRP uses the same simple password and VRRP packets that do not contain the password are dropped. If your interfaces do not use authentication, neither does VRRP.

NOTE

The MD5 authentication type is not supported for VRRP.

Independent operation of VRRP alongside RIP, OSPF, and BGP4

VRRP operation is independent of the RIP, OSPF, and BGP4 protocols. Their operation is unaffected when VRRP is enabled on a RIP, OSPF, or BGP4 interface.

Dynamic VRRP configuration

All VRRP global and interface parameters take effect immediately. You do not need to reset the system to place VRRP configuration parameters into effect.

Overview of VRRPE

NOTE

VRRPE is not supported in the base Layer 3 code.

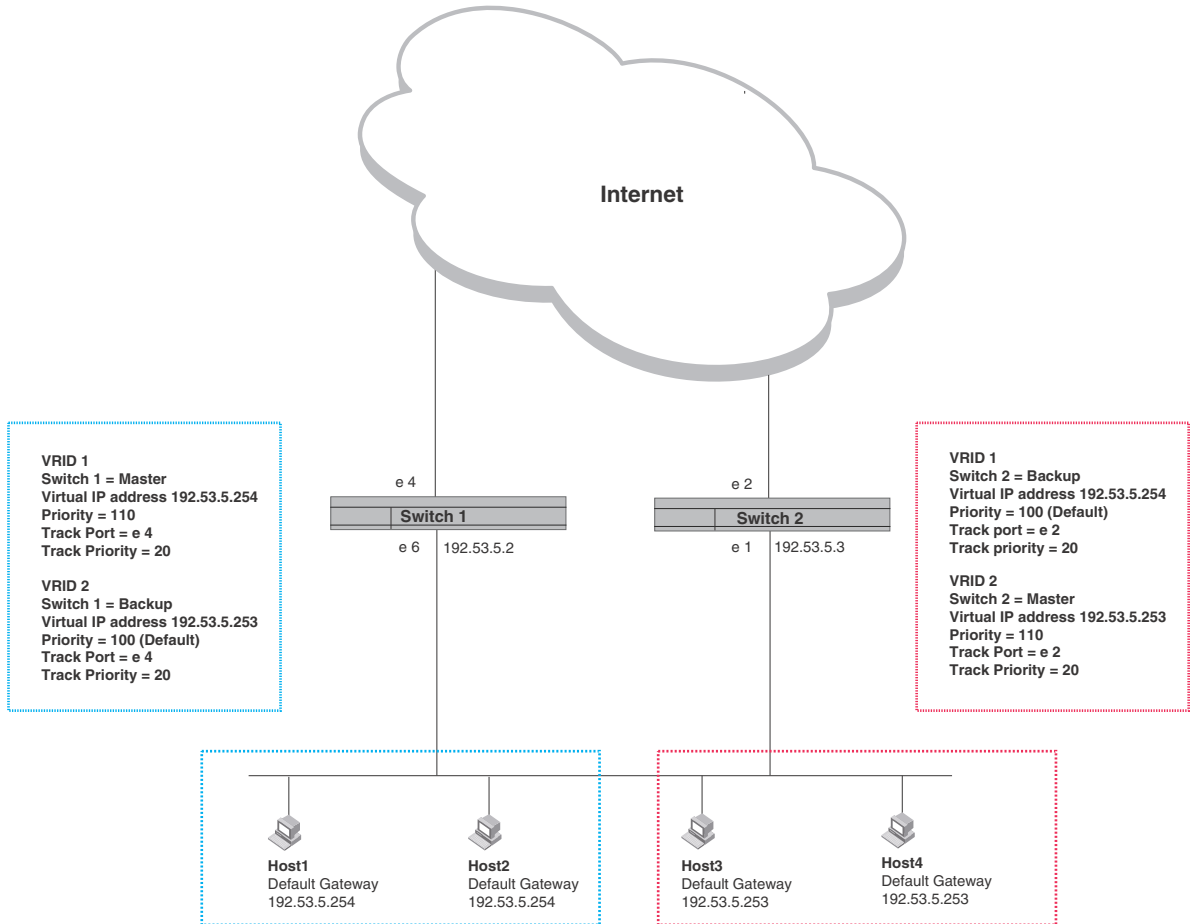
VRRPE is similar to VRRP, but differs in the following respects:

- Owners and Backup:
 - VRRP has an Owner and one or more Backups for each VRID. The Owner is the router on which the VRID's IP address is also configured as a real address. All the other routers supporting the VRID are Backups.
 - VRRPE does not use Owners. All routers are Backups for a given VRID. The router with the highest priority becomes Master. If there is a tie for highest priority, the router with the highest IP address becomes Master. The elected Master owns the virtual IP address and answers ping and ARP requests and so on.
- VRID's IP address:
 - VRRP requires that the VRID also be a real IP address configured on the VRID's interface on the Owner.
 - VRRPE requires only that the VRID be in the same subnet as an interface configured on the VRID's interface. In fact, VRRPE does not allow you to specify a real IP address configured on the interface as the VRID IP address.
- VRID's MAC Address:
 - VRRP source MAC is a virtual MAC address defined as 00-00-5E-00-01-*<vrid>*, where *<vrid>* is the VRID. The Master owns the Virtual MAC address.
 - VRRPE uses the interface actual MAC address as the source MAC address. The MAC address is 02-E0-52-*<hash-value>*-*<vrid>*, where *<hash-value>* is a two-octet hashed value for the IP address and *<vrid>* is the VRID.
- Hello packets:
 - VRRP sends Hello messages to IP Multicast address 224.0.0.18.
 - VRRPE uses UDP to send Hello messages in IP multicast messages. The Hello packets use the interface actual MAC address and IP address as the source addresses. The destination MAC address is 01-00-5E-00-00-02, and the destination IP address is 224.0.0.2 (the well-known IP multicast address for "all routers"). Both the source and destination UDP port number is 8888. VRRP messages are encapsulated in the data portion of the packet.
- Track ports and track priority:
 - VRRP changes the priority of the VRID to the track priority, which typically is lower than the VRID priority and lower than the VRID priorities configured on the Backups. For example, if the VRRP interface priority is 100 and a tracked interface with track priority 20 goes down, the software changes the VRRP interface priority to 20.
 - VRRPE reduces the priority of a VRRPE interface by the amount of a tracked interface priority if the tracked interface link goes down. For example, if the VRRPE interface priority is 200 and a tracked interface with track priority 20 goes down, the software changes the VRRPE interface priority to 180. If another tracked interface goes down, the software reduces the VRID priority again, by the amount of the tracked interface track priority.

The most important difference is that all VRRPE routers are Backups. There is no Owner router. VRRPE overcomes the limitations in standard VRRP by removing the Owner.

Figure 114 shows an example of a VRRPE configuration.

FIGURE 114 Router1 and Router2 are configured to provide dual redundant network access for the host



In this example, Switch 1 and Switch 2 use VRRPE to load share as well as provide redundancy to the hosts. The load sharing is accomplished by creating two VRRPE groups. Each group has its own virtual IP addresses. Half of the clients point to VRID 1's virtual IP address as their default gateway and the other half point to VRID 2's virtual IP address as their default gateway. This will enable some of the outbound Internet traffic to go through Switch 1 and the rest to go through Switch 2.

Switch 1 is the master for VRID 1 (backup priority = 110) and Switch 2 is the backup for VRID 1 (backup priority = 100). Switch 1 and Switch 2 both track the uplinks to the Internet. If an uplink failure occurs on Switch 1, its backup priority is decremented by 20 (track priority = 20), so that all traffic destined to the Internet is sent through Switch 2 instead.

Similarly, Switch 2 is the master for VRID 2 (backup priority = 110) and Switch 1 is the backup for VRID 2 (backup priority = 100). Switch 1 and Switch 2 are both tracking the uplinks to the Internet. If an uplink failure occurs on Switch 1, its backup priority is decremented by 20 (track priority = 20), so that all traffic destined to the internet is sent through Switch 2 instead.

Configuration note

VRRP-E is supported in the full Layer 3 code only. It is not supported in the Base Layer 3 code.

Comparison of VRRP and VRRPE

This section compares router redundancy protocols.

VRRP

VRRP is a standards-based protocol, described in RFC 2338. The VRRP contains the features in RFC 2338. It also provides the following additional features:

- **Track ports** – A feature that enables you to diagnose the health of all the Layer 3 Switch ports used by the backed-up VRID, instead of only the port connected to the client subnet. Refer to [“Track ports and track priority”](#) on page 715.
- **Suppression of RIP advertisements on Backup routes for the backed up interface** – You can enable the Layer 3 Switches to advertise only the path to the Master router for the backed up interface. Normally, a VRRP Backup router includes route information for the interface it is backing up in RIP advertisements.

Layer 3 Switches configured for VRRP can interoperate with third-party routers using VRRP.

VRRPE

VRRPE is a protocol that provides the benefits of VRRP without the limitations. VRRPE is unlike VRRP in the following ways:

- There is no “Owner” router. You do not need to use an IP address configured on one of the Layer 3 Switches as the virtual router ID (VRID), which is the address you are backing up for redundancy. The VRID is independent of the IP interfaces configured in the Layer 3 Switches. As a result, the protocol does not have an “Owner” as VRRP does.
- There is no restriction on which router can be the default master router. In VRRP, the “Owner” (the Layer 3 Switch on which the IP interface that is used for the VRID is configured) must be the default Master.

Layer 3 Switches configured for VRRPE can interoperate only with other Layer 3 Switches.

Architectural differences

The protocols have the following architectural differences:

Management protocol

- VRRP – VRRP routers send VRRP Hello and Hello messages to IP Multicast address 224.0.0.18.
- VRRPE – VRRPE sends messages to destination MAC address 01-00-5E-00-00-02 and destination IP address 224.0.0.2 (the standard IP multicast address for “all routers”).

Virtual router IP address (the address you are backing up)

- VRRP – The virtual router IP address is the same as an IP address or virtual interface configured on one of the Layer 3 Switches, which is the “Owner” and becomes the default Master.
- VRRPE – The virtual router IP address is the gateway address you want to backup, but does not need to be an IP interface configured on one of the Layer 3 Switch ports or a virtual interface.

Master and Backups

- VRRP – The “Owner” of the IP address of the VRID is the default Master and has the highest priority (255). The precedence of the Backups is determined by their priorities. The default Master is always the Owner of the IP address of the VRID.
- VRRPE – The Master and Backups are selected based on their priority. You can configure any of the Layer 3 Switches to be the Master by giving it the highest priority. There is no Owner.

VRRP and VRRPE parameters

Table 119 lists the VRRP and VRRPE parameters. Most of the parameters and default values are the same for both protocols. The exceptions are noted in the table.

TABLE 119 VRRP and VRRPE parameters

Parameter	Description	Default	See page...
Protocol	The Virtual Router Redundancy Protocol (VRRP) based on RFC 2338 or VRRP-Extended, the enhanced implementation of VRRP	Disabled NOTE: Only one of the protocols can be enabled at a time.	page 722 page 723
VRRP or VRRPE router	The Layer 3 Switch active participation as a VRRP or VRRPE router. Enabling the protocol does not activate the Layer 3 Switch for VRRP or VRRPE. You must activate the device as a VRRP or VRRPE router after you configure the VRRP or VRRPE parameters.	Inactive	page 722 page 723
Virtual Router ID (VRID)	The ID of the virtual router you are creating by configuring multiple routers to back up an IP interface. You must configure the same VRID on each router that you want to use to back up the address. No default.	None	page 713 page 722 page 723
Virtual Router IP address	This is the address you are backing up. No default: <ul style="list-style-type: none"> • VRRP – The virtual router IP address must be a real IP address configured on the VRID interface on one of the VRRP routers. This router is the IP address Owner and is the default Master. • VRRPE – The virtual router IP address must be in the same subnet as a real IP address configured on the VRRPE interface, but cannot be the same as a real IP address configured on the interface. 	None	page 714 page 722 page 723

TABLE 119 VRRP and VRRPE parameters (Continued)

Parameter	Description	Default	See page...
VRID MAC address	<p>The source MAC address in VRRP or VRRPE packets sent from the VRID interface, and the destination for packets sent to the VRID:</p> <ul style="list-style-type: none"> • VRRP – A virtual MAC address defined as 00-00-5e-00-01-<i><vrid></i>. The Master owns the Virtual MAC address. • VRRPE – A virtual MAC address defined as 02-E0-52-<i><hash-value></i>-<i><vrid></i>, where <i><hash-value></i> is a two-octet hashed value for the IP address and <i><vrid></i> is the VRID. 	Not configurable	page 714
Authentication type	<p>The type of authentication the VRRP or VRRPE routers use to validate VRRP or VRRPE packets. The authentication type must match the authentication type the VRID port uses with other routing protocols such as OSPF:</p> <ul style="list-style-type: none"> • No authentication – The interfaces do not use authentication. This is the VRRP default. • Simple – The interface uses a simple text-string as a password in packets sent on the interface. If the interface uses simple password authentication, the VRID configured on the interface must use the same authentication type and the same password. <p>NOTE: MD5 is not supported by VRRP or VRRPE.</p>	No authentication	page 716 page 725
Router type	<p>Whether the router is an Owner or a Backup.</p> <ul style="list-style-type: none"> • Owner (VRRP only) – The router on which the real IP address used by the VRID is configured. • Backup – Routers that can provide routing services for the VRID but do not have a real IP address matching the VRID. 	<p>VRRP – The Owner is always the router that has the real IP address used by the VRID. All other routers for the VRID are Backups.</p> <p>VRRPE – All routers for the VRID are Backups.</p>	page 726
Backup priority	<p>A numeric value that determines a Backup preferability for becoming the Master for the VRID. During negotiation, the router with the highest priority becomes the Master.</p> <ul style="list-style-type: none"> • VRRP – The Owner has the highest priority (255); other routers can have a priority from 3 – 254. • VRRPE – All routers are Backups and have the same priority by default. <p>If two or more Backups are tied with the highest priority, the Backup interface with the highest IP address becomes the Master for the VRID.</p>	<p>VRRP – 255 for the Owner; 100 for each Backup</p> <p>VRRPE – 100 for all Backups</p>	page 726
Suppression of RIP advertisements	<p>A router that is running RIP normally advertises routes to a backed up VRID even when the router is not currently the active router for the VRID. Suppression of these advertisements helps ensure that other routers do not receive invalid route paths for the VRID.</p>	Disabled	page 727
Hello interval	<p>The number of seconds between Hello messages from the Master to the Backups for a given VRID. The interval can from 1 – 84 seconds.</p>	One second	page 715 page 727

TABLE 119 VRRP and VRRPE parameters (Continued)

Parameter	Description	Default	See page...
Dead interval	The number of seconds a Backup waits for a Hello message from the Master for the VRID before determining that the Master is no longer active. If the Master does not send a Hello message before the dead interval expires, the Backups negotiate (compare priorities) to select a new Master for the VRID.	Three times the Hello Interval plus one-half second	page 715 page 728
Backup Hello interval	The number of seconds between Hello messages from a Backup to the Master. The message interval can be from 60 – 3600 seconds. You must enable the Backup to send the messages. The messages are disabled by default on Backups. The current Master (whether the VRRP Owner or a Backup) sends Hello messages by default.	Disabled 60 seconds when enabled	page 715 page 728
Track port	Another Layer 3 Switch port or virtual interface whose link status is tracked by the VRID interface. If the link for a tracked interface goes down, the VRRP or VRRPE priority of the VRID interface is changed, causing the devices to renegotiate for Master.	None	page 715 page 728
Track priority	A VRRP or VRRPE priority value assigned to the tracked ports. If a tracked port link goes down, the VRID port VRRP or VRRPE priority changes: <ul style="list-style-type: none"> • VRRP – The priority changes to the value of the tracked port priority. • VRRPE – The VRID port priority is reduced by the amount of the tracked port priority. 	VRRP – 2 VRRPE – 5	page 715 page 729
Backup preempt mode	Prevents a Backup with a higher VRRP priority from taking control of the VRID from another Backup that has a lower priority but has already assumed control of the VRID.	Enabled	page 729
Timer scale	Adjusts the timers for the Hello interval, Dead interval, Backup Hello interval, and Hold-down interval.	1	page 730
VRRP-E slow start timer	This feature causes a specified amount of time to elapse between the time the Master is restored and when it takes over from the Backup. This interval allows time for OSPF convergence when the Master is restored.	Disabled	page 731

Configuring basic VRRP parameters

To implement a simple VRRP configuration using all the default values, enter commands such as the following.

Configuring the Owner

```
Router1(config)# router vrrp
Router1(config)# inter e 6
Router1(config-if-6)# ip address 192.53.5.1
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# owner
Router1(config-if-6-vrid-1)# ip-address 192.53.5.1
Router1(config-if-6-vrid-1)# activate
```

Configuring a Backup

```
Router2(config)# router vrrp
Router2(config)# inter e 5
Router2(config-if-5)# ip address 192.53.5.3
Router2(config-if-5)# ip vrrp vrid 1
Router2(config-if-5-vrid-1)# backup
Router2(config-if-5-vrid-1)# advertise backup
Router2(config-if-5-vrid-1)# ip-address 192.53.5.1
Router2(config-if-5-vrid-1)# activate
```

Configuration rules for VRRP

- The interfaces of all routers in a VRID must be in the same IP subnet.
- The IP addresses associated with the VRID must already be configured on the router that will be the Owner router.
- An IP address associated with the VRID must be on only one router.
- The Hello interval must be set to the same value on both the Owner and Backups for the VRID.
- The Dead interval must be set to the same value on both the Owner and Backups for the VRID.
- The track priority on a router must be lower than the router VRRP priority. Also, the track priority on the Owner must be higher than the track priority on the Backups.

Configuring basic VRRPE parameters

To implement a simple VRRPE configuration using all the default values, enter commands such as the following on each Layer 3 Switch.

```
Router2(config)# router vrrp-extended
Router2(config)# inter e 5
Router2(config-if-5)# ip address 192.53.5.3
Router2(config-if-5)# ip vrrp-extended vrid 1
Router2(config-if-5-vrid-1)# backup
Router2(config-if-5-vrid-1)# advertise backup
Router2(config-if-5-vrid-1)# ip-address 192.53.5.254
Router2(config-if-5-vrid-1)# activate
```

NOTE

You also can use the **enable** command to activate the configuration. This command does the same thing as the **activate** command.

Configuration rules for VRRPE

- The interfaces of all routers in a VRID must be in the same IP subnet.
- The IP address associated with the VRID cannot be configured on any of the Layer 3 Switches.
- The Hello interval must be set to the same value on all the Layer 3 Switches.
- The Dead interval must be set to the same value on all the Layer 3 Switches.
- The track priority for a VRID must be lower than the VRRPE priority.

Note regarding disabling VRRP or VRRPE

If you disable VRRP or VRRPE, the Layer 3 Switch removes all the configuration information for the disabled protocol from the running-config. Moreover, when you save the configuration to the startup-config file after disabling one of these protocols, all the configuration information for the disabled protocol is removed from the startup-config file.

The CLI displays a warning message such as the following.

```
Router1(config-vrrp-router)# no router vrrp
router vrrp mode now disabled. All vrrp config data will be lost when writing to
flash!
```

If you have disabled the protocol but have not yet saved the configuration to the startup-config file and reloaded the software, you can restore the configuration information by re-entering the command to enable the protocol (ex: **router vrrp**). If you have already saved the configuration to the startup-config file and reloaded the software, the information is gone.

If you are testing a VRRP or VRRPE configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup-config file containing the protocol configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

Configuring additional VRRP and VRRPE parameters

You can modify the following VRRP and VRRPE parameters on an individual VRID basis. These parameters apply to both protocols:

- Authentication type (if the interfaces on which you configure the VRID use authentication)
- Router type (Owner or Backup)

NOTE

For VRRP, change the router type only if you have moved the real IP address from one router to another or you accidentally configured the IP address Owner as a Backup.

For VRRPE, the router type is always Backup. You cannot change the type to Owner.

- Backup priority
- Suppression of RIP advertisements on Backup routes for the backed up interface
- Hello interval
- Dead interval

- Backup Hello messages and message timer (Backup advertisement)
- Track port
- Track priority
- Backup preempt mode
- Timer scale
- VRRP-E slow start timer

For information about the fields, see the parameter descriptions in the following sections.

Refer to “[VRRP and VRRPE parameters](#)” on page 720 for a summary of the parameters and their defaults.

Authentication type

If the interfaces on which you configure the VRID use authentication, the VRRP or VRRPE packets on those interfaces also must use the same authentication. The VRRP and VRRPE supports the following authentication types:

- **No authentication** – The interfaces do not use authentication. This is the default for VRRP and VRRPE.
- **Simple** – The interfaces use a simple text-string as a password in packets sent on the interface. If the interfaces use simple password authentication, the VRID configured on the interfaces must use the same authentication type and the same password.

To configure the VRID interface on Router1 for simple-password authentication using the password “ourpword”, enter the following commands.

Configuring Router 1

```
Router1(config)# inter e 6
Router1(config-if-6)# ip vrrp auth-type simple-text-auth ourpword
```

Configuring Router 2

```
Router2(config)# inter e 5
Router2(config-if-5)# ip vrrp auth-type simple-text-auth ourpword
```

VRRP syntax

Syntax: `ip vrrp auth-type no-auth | simple-text-auth <auth-data>`

The **auth-type no-auth** parameter indicates that the VRID and the interface it is configured on do not use authentication.

The **auth-type simple-text-auth <auth-data>** parameter indicates that the VRID and the interface it is configured on use a simple text password for authentication. The **<auth-data>** parameter is the password. If you use this parameter, make sure all interfaces on all the routers supporting this VRID are configured for simple password authentication and use the same password.

VRRPE syntax

Syntax: `ip vrrp-extended auth-type no-auth | simple-text-auth <auth-data>`

The parameter values are the same as for VRRP.

Router type

A VRRP interface is either an Owner or a Backup for a given VRID. By default, the Owner becomes the Master following the negotiation. A Backup becomes the Master only if the Master becomes unavailable.

A VRRPE interface is always a Backup for its VRID. The Backup with the highest VRRP priority becomes the Master.

This section describes how to specify the interface type, how to change the type for VRRP, and how to set or change the interface VRRP or VRRPE priority and track priority for the VRID.

NOTE

You can force a VRRP master router to abdicate (give away control) of the VRID to a Backup by temporarily changing the Master VRRP priority to a value less than the Backup. Refer to [“Forcing a Master router to abdicate to a standby router”](#) on page 731.

NOTE

The type Owner is not applicable to VRRPE.

NOTE

The IP address(es) you associate with the Owner must be a real IP address (or addresses) on the interface on which you configure the VRID.

When you configure a Backup router, the router interface on which you are configuring the VRID must have a real IP address that is in the same subnet as the address associated with the VRID by the Owner. However, the address cannot be the same.

To configure Router1 as a VRRP VRID Owner, enter the following commands.

```
Router1(config)# inter e 6
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# owner
```

To configure Router2 as a VRRP Backup for the same VRID, enter the following commands.

```
Router2(config)# inter e 5
Router2(config-if-5)# ip vrrp vrid 1
Router2(config-if-5-vrid-1)# backup
Router2(config-if-5-vrid-1)# advertise backup
```

To configure a VRRPE interface as a Backup for a VRID and set its VRRPE priority and track priority, enter commands such as the following.

```
PowerConnect(config)# inter e 1
PowerConnect(config-if-1)# ip vrrp-extended vrid 1
PowerConnect(config-if-1-vrid-1)# backup priority 50 track-priority 10
Router2(config-if-1-vrid-1)# advertise backup
```

VRRP syntax

Syntax: `owner [track-priority <value>]`

The **track-priority <value>** parameter changes the track-port priority for this interface and VRID from the default (2) to a value from 1 - 254.

Syntax: `backup [priority <value>] [track-priority <value>]`

The **priority** *<value>* parameter specifies the VRRP priority for this interface and VRID. You can specify a value from 3 – 254. The default is 100.

The **track-priority** *<value>* parameter is the same as above.

NOTE

You cannot set the priority of a VRRP Owner. The Owner priority is always 255.

VRRPE syntax

Syntax: `backup [priority <value>] [track-priority <value>]`

The software requires you to identify a VRRPE interface as a Backup for its VRID before you can activate the interface for the VRID. However, after you configure the VRID, you can use this command to change its priority or track priority. The parameter values are the same as for VRRP.

Suppression of RIP advertisements on Backup routers for the Backup interface

Normally, a VRRP or VRRPE Backup includes route information for the virtual IP address (the backed up interface) in RIP advertisements. As a result, other routers receive multiple paths for the backed up interface and might sometimes unsuccessfully use the path to the Backup rather than the path to the Master.

You can prevent the Backups from advertising route information for the backed up interface by enabling suppression of the advertisements.

To suppress RIP advertisements for the backed up interface in Router2, enter the following commands.

```
Router2(config)# router rip
Router2(config-rip-router)# use-vrrp-path
```

Syntax: `use-vrrp-path`

The syntax is the same for VRRP and VRRPE.

Hello interval

The Master periodically sends Hello messages to the Backups. The Backups use the Hello messages as verification that the Master is still on-line. If the Backup routers stop receiving the Hello messages for the period of time specified by the Dead interval, the Backup routers determine that the Master router is dead. At this point, the Backup router with the highest priority becomes the new Master router. The Hello interval can be from 1 – 84 seconds. The default is 1 second.

NOTE

The default Dead interval is three times the Hello Interval plus one-half second. Generally, if you change the Hello interval, you also should change the Dead interval on the Backup routers.

To change the Hello interval on the Master to 10 seconds, enter the following commands.

```
Router1(config)# inter e 6
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# hello-interval 10
```

Syntax: `hello-interval <value>`

The syntax is the same for VRRP and VRRPE.

Dead interval

The Dead interval is the number of seconds a Backup waits for a Hello message from the Master before determining that the Master is dead. When Backups determine that the Master is dead, the Backup with the highest priority becomes the new Master. The Dead interval can be from 1 – 84 seconds. The default is 3.5 seconds. This is three times the default Hello interval (1 second) plus one-half second added by the router software. The software automatically adds one-half second to the Dead interval value you enter.

To change the Dead interval on a Backup to 30 seconds, enter the following commands.

```
Router2(config)# inter e 5
Router2(config-if-5)# ip vrrp vrid 1
Router2(config-if-5-vrid-1)# dead-interval 30
```

Syntax: `dead-interval <value>`

The syntax is the same for VRRP and VRRPE.

Backup Hello message state and interval

By default, Backup do not send Hello messages to advertise themselves to the Master. You can enable these messages if desired and also change the message interval.

To enable a Backup to send Hello messages to the Master, enter commands such as the following.

```
PowerConnect(config)# router vrrp
PowerConnect(config)# inter e 6
PowerConnect(config-if-6)# ip vrrp vrid 1
PowerConnect(config-if-6-vrid-1)# advertise backup
```

Syntax: `[no] advertise backup`

When you enable a Backup to send Hello messages, the Backup sends a Hello messages to the Master every 60 seconds by default. You can change the interval to be up to 3600 seconds. To do so, enter commands such as the following.

```
PowerConnect(config)# router vrrp
PowerConnect(config)# inter e 6
PowerConnect(config-if-6)# ip vrrp vrid 1
PowerConnect(config-if-6-vrid-1)# backup-hello-interval 180
```

Syntax: `[no] backup-hello-interval <num>`

The `<num>` parameter specifies the message interval and can be from 60 – 3600 seconds. The default is 60 seconds.

The syntax is the same for VRRP and VRRPE.

Track port

You can configure the VRID on one interface to track the link state of another interface on the Layer 3 Switch. This capability is quite useful for tracking the state of the exit interface for the path for which the VRID is providing redundancy. Refer to [“Track ports and track priority”](#) on page 715.

To configure 6 on Router1 to track interface 4, enter the following commands.

```
Router1(config)# inter e 6
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# track-port e 4
```

Syntax: `track-port ethernet <portnum> | ve <num>`

The syntax is the same for VRRP and VRRPE.

Track priority

When you configure a VRID to track the link state of other interfaces, if one of the tracked interface goes down, the software changes the VRRP or VRRPE priority of the VRID interface:

- For VRRP, the software changes the priority of the VRID to the track priority, which typically is lower than the VRID priority and lower than the VRID priorities configured on the Backups. For example, if the VRRPE interface priority is 100 and a tracked interface with track priority 60 goes down, the software changes the VRRPE interface priority to 60.
- For VRRPE, the software reduces the VRID priority by the amount of the priority of the tracked interface that went down. For example, if the VRRPE interface priority is 100 and a tracked interface with track priority 60 goes down, the software changes the VRRPE interface priority to 40. If another tracked interface goes down, the software reduces the VRID priority again, by the amount of the tracked interface track priority.

The default track priority for a VRRP Owner is 2. The default track priority for Backups is 1.

You enter the track priority as a parameter with the **owner** or **backup** command. Refer to “[Track port](#)” on page 728.

Syntax: `owner [track-priority <value>]`

Syntax: `backup [priority <value>] [track-priority <value>]`

The syntax is the same for VRRP and VRRPE.

Backup preempt

By default, a Backup that has a higher priority than another Backup that has become the Master can preempt the Master, and take over the role of Master. If you want to prevent this behavior, disable preemption.

Preemption applies only to Backups and takes effect only when the Master has failed and a Backup has assumed ownership of the VRID. The feature prevents a Backup with a higher priority from taking over as Master from another Backup that has a lower priority but has already become the Master of the VRID.

Preemption is especially useful for preventing flapping in situations where there are multiple Backups and a Backup with a lower priority than another Backup has assumed ownership, because the Backup with the higher priority was unavailable when ownership changed.

If you enable the non-preempt mode (thus disabling the preemption feature) on all the Backups, the Backup that becomes the Master following the disappearance of the Master continues to be the Master. The new Master is not preempted.

In the VRRP non-preempt mode, the VRRP router controlling the IP addresses associated with the virtual router becomes the Master. When the Master becomes unavailable, the Backup router with highest priority becomes the Master. If the Master router changes its priority to lower priority (due to track port going down), the Backup router in the non-preempt mode does not take over.

NOTE

In VRRP, regardless of the setting for the preempt parameter, the Owner always becomes the Master again when it comes back online.

24 Configuring additional VRRP and VRRPE parameters

To disable preemption on a Backup, enter commands such as the following.

```
Router1(config)# inter e 6
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# non-preempt-mode
```

Syntax: non-preempt-mode

The syntax is the same for VRRP and VRRPE.

Changing the timer scale

To achieve sub-second failover times, you can shorten the duration of all scale timers for VSRP, VRRP, and VRRP-E by adjusting the timer scale. The **timer scale** is a value used by the software to calculate the timers. By default, the scale value is 1. If you increase the timer scale, each timer's value is divided by the scale value. Using the timer scale to adjust timer values enables you to easily change all the timers while preserving the ratios among their values. Here is an example.

Timer	Timer scale	Timervalue
Hello interval	1	1 second
	2	0.5 seconds
Dead interval	1	3 seconds
	2	1.5 seconds
Backup Hello interval	1	60 seconds
	2	30 seconds
Hold-down interval	1	2 seconds
	2	1 second

If you configure the device to receive its timer values from the Master, the Backup also receives the timer scale value from the Master.

NOTE

The Backups always use the value of the timer scale received from the Master, regardless of whether the timer values that are saved in the configuration are the values configured on the Backup or the values received from the Master.

To change the timer scale, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# scale-timer 2
```

This command changes the scale to 2. All VSRP, VRRP, and VRRP-E timer values will be divided by 2.

Syntax: [no] scale-timer <num>

The <num> parameter specifies the multiplier. You can specify a timer scale from 1 – 10.

VRRP-E slow start timer

In a VRRP-E configuration, if a Master router goes down, the Backup router with the highest priority takes over. When the Master comes back up again, it takes over from the Backup. By default, this transition from Backup back to Master takes place immediately. However, you can configure the VRRP-E slow start timer feature, which causes a specified amount of time to elapse between the time the Master is restored and when it takes over from the Backup. This interval allows time for OSPF convergence when the Master is restored.

To set the VRRP-E slow start timer to 30 seconds, enter the following commands.

```
PowerConnect(config)# router vrrp-e
PowerConnect(config-vrrpe-router)# slow-start 30
```

Syntax: [no] **slow-start** <seconds>

For <seconds>, enter a value from 1 – 255.

When the VRRP-E slow start timer is enabled, if the Master goes down, the Backup takes over immediately. If the Master subsequently comes back up again, the amount of time specified by the VRRP-E slow start timer elapses (in this example, 30 seconds) before the Master takes over from the Backup.

The VRRP-E slow start timer is effective only if another VRRP-E Master (Standby) is detected. It is not effective during the initial boot up.

NOTE

The VRRP-E slow start timer applies only to VRRP-E configurations. It does not apply to VRRP configurations.

Forcing a Master router to abdicate to a standby router

You can force a VRRP Master to abdicate (give away control) of a VRID to a Backup by temporarily changing the Master priority to a value less than the Backup.

The VRRP Owner always has priority 255. You can even use this feature to temporarily change the Owner priority to a value from 1 – 254.

NOTE

When you change a VRRP Owner priority, the change takes effect only for the current power cycle. The change is not saved to the startup-config file when you save the configuration and is not retained across a reload or reboot. Following a reload or reboot, the VRRP Owner again has priority 255.

To temporarily change the Master priority, use the following CLI method.

To change the Master priority, enter commands such as the following.

```
PowerConnect(config)# ip int eth 6
PowerConnect(config-if-6)# ip vrrp vrid 1
PowerConnect(config-if-6-vrid-1)# owner priority 99
```

Syntax: [no] **owner priority** | **track-priority** <num>

The <num> parameter specifies the new priority and can be a number from 1 – 254.

24 Displaying VRRP and VRRPE information

When you press Enter, the software changes the priority of the Master to the specified priority. If the new priority is lower than at least one Backup priority for the same VRID, the Backup takes over and becomes the new Master until the next software reload or system reset.

To verify the change, enter the following command from any level of the CLI.

```
PowerConnect# show ip vrrp
Total number of VRRP routers defined: 1
Interface ethernet 6
auth-type no authentication
VRID 1
state backup
administrative-status enabled
mode owner
priority 99
current priority 99
hello-interval 1 sec
ip-address 192.53.5.1
backup routers 192.53.5.2
```

This example shows that even though this Layer 3 Switch is the Owner of the VRID (“mode owner”), the Layer 3 Switch priority for the VRID is only 99 and the state is now “backup” instead of “active”. In addition, the administrative status is “enabled”.

To change the Master priority back to the default Owner priority 255, enter “no” followed by the command you entered to change the priority. For example, to change the priority of a VRRP Owner back to 255 from 99, enter the following command.

```
PowerConnect(config-if-6-vrid-1)# no owner priority 99
```

You cannot set the priority to 255 using the **owner priority** command.

Displaying VRRP and VRRPE information

You can display the following information for VRRP or VRRPE:

- Summary configuration and status information
- Detailed configuration and status information
- VRRP and VRRPE Statistics
- CPU utilization statistics

Displaying summary information

To display summary information for a Layer 3 Switch, enter the following command at any level of the CLI.

```
PowerConnect# show ip vrrp brief

Total number of VRRP routers defined: 1
Interface VRID CurPri P State Master addr Backup addr VIP
6          1    255 P Init  192.53.5.1  192.53.5.3 192.53.5.1
```

The above example is for VRRP. Here is an example for VRRPE.

```
PowerConnect# show ip vrrp-extended brief
```

```
Total number of VRRP-Extended routers defined: 1
Interface VRID CurPri P State Master addr Backup addr VIP
6          1    255 P Init  192.53.5.2   192.53.5.3 192.53.5.254
```

Syntax: show ip vrrp brief | ethernet <portnum> | ve <num> | stat

Syntax: show ip vrrp-extended brief | ethernet <portnum> | ve <num> | stat

The **brief** parameter displays the summary information. If you do not use this parameter, detailed information is displayed instead. Refer to “[Displaying detailed information](#)” on page 734.

The <portnum> parameter specifies an Ethernet port. If you use this parameter, the command displays VRRP or VRRPE information only for the specified port.

The **ve <num>** parameter specifies a virtual interface. If you use this parameter, the command displays VRRP or VRRPE information only for the specified virtual interface.

The **stat** parameter displays statistics. Refer to “[Displaying statistics](#)” on page 739.

This display shows the following information.

TABLE 120 CLI display of VRRP or VRRPE summary information

This field...	Displays...
Total number of VRRP (or VRRP-Extended) routers defined	The total number of VRIDs configured on this Layer 3 Switch. NOTE: The total applies only to the protocol the Layer 3 Switch is running. For example, if the Layer 3 Switch is running VRRPE, the total applies only to VRRPE routers.
Interface	The interface on which VRRP or VRRPE is configured. If VRRP or VRRPE is configured on multiple interfaces, information for each interface is listed separately.
VRID	The VRID configured on this interface. If multiple VRIDs are configured on the interface, information for each VRID is listed in a separate row.
CurPri	The current VRRP or VRRPE priority of this Layer 3 Switch for the VRID.
P	Whether the backup preempt mode is enabled. If the backup preempt mode is enabled, this field contains a “P”. If the mode is disabled, this field is blank.
State	This Layer 3 Switch VRRP or VRRPE state for the VRID. The state can be one of the following: <ul style="list-style-type: none"> Init – The VRID is not enabled (activated). If the state remains Init after you activate the VRID, make sure that the VRID is also configured on the other routers and that the routers can communicate with each other. NOTE: If the state is Init and the mode is incomplete, make sure you have specified the IP address for the VRID. <ul style="list-style-type: none"> Backup – This Layer 3 Switch is a Backup for the VRID. Master – This Layer 3 Switch is the Master for the VRID.
Master addr	IP address of the router interface that is currently Master for the VRID.
Backup addr	IP addresses of router interfaces that are currently Backups for the VRID.
VIP	The virtual IP address that is being backed up by the VRID.

Displaying detailed information

To display detailed VRRP or VRRPE information, enter the following command at any level of the CLI.

```
PowerConnect# show ip vrrp

Total number of VRRP routers defined: 1
Interface ethernet 6
  auth-type no authentication
  VRID 1
    state master
    administrative-status enabled
    mode owner
    priority 255
    current priority 255
    hello-interval 10000 msec
    advertise backup: disabled
    track-port 4
```

This example is for a VRRP Owner. Here is an example for a VRRP Backup.

```
PowerConnect# show ip vrrp

Total number of VRRP routers defined: 1
Interface ethernet 5
  auth-type no authentication
  VRID 1
    state backup
    administrative-status enabled
    mode non-owner(backup)
    priority 100
    current priority 100
    hello-interval 10000 msec
    dead-interval 30000 msec
    current dead-interval 10000 msec
    preempt-mode true
    advertise backup: enabled
    backup router 192.53.5.3 expires in 00:00:03.0
    next hello sent in 00:00:02.0
    track-port 2
```

Here is an example for a VRRPE Backup.

```
PowerConnect# show ip vrrp-extended

Total number of VRRP-Extended routers defined: 1
Interface ethernet 6
  auth-type no authentication
  VRID 1
    state master
    administrative-status enabled
    priority 200
    current priority 200
    hello-interval 10000 msec
    dead-interval 30000 msec
    current dead-interval 30000 msec
    preempt-mode true
    virtual ip address 192.53.5.254
    advertise backup: enabled
    master router 192.53.5.2 expires in 00:00:03.0
    track-port 4
```

Syntax: show ip vrrp brief | ethernet <portnum> | ve <num> | stat

Syntax: show ip vrrp-extended brief | ethernet <portnum> | ve <num> | stat

The **brief** parameter displays summary information. Refer to “[Displaying summary information](#)” on page 732.

The <portnum> parameter specifies an Ethernet port. If you use this parameter, the command displays VRRP or VRRPE information only for the specified port.

The **ve <num>** parameter specifies a virtual interface. If you use this parameter, the command displays VRRP or VRRPE information only for the specified virtual interface.

The **stat** parameter displays statistics. Refer to “[Displaying statistics](#)” on page 739.

This display shows the following information.

TABLE 121 CLI display of VRRP or VRRPE detailed information

This field...	Displays...
Total number of VRRP (or VRRP-Extended) routers defined	The total number of VRIDs configured on this Layer 3 Switch. NOTE: The total applies only to the protocol the Layer 3 Switch is running. For example, if the Layer 3 Switch is running VRRPE, the total applies only to VRRPE routers.
Interface parameters	
Interface	The interface on which VRRP or VRRPE is configured. If VRRP or VRRPE is configured on multiple interfaces, information for each interface is listed separately.
auth-type	The authentication type enabled on the interface.
VRID parameters	
VRID	The VRID configured on this interface. If multiple VRIDs are configured on the interface, information for each VRID is listed separately.

TABLE 121 CLI display of VRRP or VRRPE detailed information (Continued)

This field...	Displays...
state	<p>This Layer 3 Switch VRRP or VRRPE state for the VRID. The state can be one of the following:</p> <ul style="list-style-type: none"> initialize – The VRID is not enabled (activated). If the state remains “initialize” after you activate the VRID, make sure that the VRID is also configured on the other routers and that the routers can communicate with each other. <p>NOTE: If the state is “initialize” and the mode is incomplete, make sure you have specified the IP address for the VRID.</p> <ul style="list-style-type: none"> backup – This Layer 3 Switch is a Backup for the VRID. master – This Layer 3 Switch is the Master for the VRID.
administrative-status	<p>The administrative status of the VRID. The administrative status can be one of the following:</p> <ul style="list-style-type: none"> disabled – The VRID is configured on the interface but VRRP or VRRPE has not been activated on the interface. enabled – VRRP or VRRPE has been activated on the interface.
mode	<p>Indicates whether the Layer 3 Switch is the Owner or a Backup for the VRID.</p> <p>NOTE: If “incomplete” appears after the mode, configuration for this VRID is incomplete. For example, you might not have configured the virtual IP address that is being backed up by the VRID.</p> <p>NOTE: This field applies only to VRRP. All Layer 3 Switches configured for VRRPE are Backups.</p>
priority	<p>The device preferability for becoming the Master for the VRID. During negotiation, the router with the highest priority becomes the Master. If two or more devices are tied with the highest priority, the Backup interface with the highest IP address becomes the active router for the VRID.</p>
current priority	<p>The current VRRP or VRRPE priority of this Layer 3 Switch for the VRID. The current priority can differ from the configured priority (see the row above) for the following reasons:</p> <ul style="list-style-type: none"> The VRID is still in the initialization stage and has not become a Master or Backup yet. In this case, the current priority is 0. The VRID is configured with track ports and the link on a tracked interface has gone down. Refer to “Track ports and track priority” on page 715.
hello-interval	<p>The configured value for the hello interval. This is the amount of time between Hello messages from the Master to the Backups for a given VRID:</p> <ul style="list-style-type: none"> It show the dead interval in number of milliseconds.
dead-interval	<p>The configured value for the dead interval. This is the amount of time a Backup waits for a Hello message from the Master for the VRID before determining that the Master is no longer active:</p> <ul style="list-style-type: none"> It show the dead interval in number of milliseconds. <p>If the Master does not send a Hello message before the dead interval expires, the Backups negotiate (compare priorities) to select a new Master for the VRID.</p> <p>NOTE: If the value is 0, then you have not configured this parameter.</p> <p>NOTE: This field does not apply to VRRP Owners.</p>
current dead-interval	<p>The current value of the dead interval. This is the value actually in use by this interface for the VRID:</p> <ul style="list-style-type: none"> It show the dead interval in number of milliseconds. <p>NOTE: This field does not apply to VRRP Owners.</p>

TABLE 121 CLI display of VRRP or VRRPE detailed information (Continued)

This field...	Displays...
preempt-mode	Whether the backup preempt mode is enabled. NOTE: This field does not apply to VRRP Owners.
virtual ip address	The virtual IP addresses that this VRID is backing up.
advertise backup	The IP addresses of Backups that have advertised themselves to this Layer 3 Switch by sending Hello messages. NOTE: Hello messages from Backups are disabled by default. You must enable the Hello messages on the Backup for the Backup to advertise itself to the current Master. Refer to “Hello messages” on page 715.
backup router <ip-addr> expires in <time>	The IP addresses of Backups that have advertised themselves to this Master by sending Hello messages. The <time> value indicates how long before the Backup expires. A Backup expires if you disable the advertise backup option on the Backup or the Backup becomes unavailable. Otherwise, the Backup next Hello message arrives before the Backup expires. The Hello message resets the expiration timer. An expired Backup does not necessarily affect the Master. However, if you have not disabled the advertise backup option on the Backup, then the expiration may indicate a problem with the Backup. NOTE: This field applies only when Hello messages are enabled on the Backups (using the advertise backup option).
next hello sent in <time>	How long until the Backup sends its next Hello message. NOTE: This field applies only when this Layer 3 Switch is the Master and the Backup is configured to send Hello messages (the advertise backup option is enabled).
master router <ip-addr> expires in <time>	The IP address of the Master and the amount of time until the Master dead interval expires. If the Backup does not receive a Hello message from the Master by the time the interval expires, either the IP address listed for the Master will change to the IP address of the new Master, or this Layer 3 Switch itself will become the Master. NOTE: This field applies only when this Layer 3 Switch is a Backup.
track port	The interfaces that the VRID interface is tracking. If the link for a tracked interface goes down, the VRRP or VRRPE priority of the VRID interface is changed, causing the devices to renegotiate for Master. NOTE: This field is displayed only if track interfaces are configured for this VRID.

Displaying detailed information for an individual VRID

You can display information about the settings configured for a specified VRRP Virtual Router ID (VRID). For example, to display information about VRID 1.

```
PowerConnect# show ip vrrp vrid 1
VRID 1
  Interface ethernet 11
  state initialize
  administrative-status disabled
  mode non-owner(backup)incomplete
  priority 12
  current priority 12
  track-priority 22
  hello-interval 1 sec
  dead-interval 0 sec
  current dead-interval 3.900 sec
  preempt-mode true
  advertise backup: disabled
```

Syntax: `show ip vrrp vrid <num> [ethernet <num> | ve <num>]`

The `<num>` parameter specifies the VRID.

The `ethernet <num> | ve <num>` specifies an interface on which the VRID is configured. If you specify an interface, VRID information is displayed for that interface only. Otherwise, information is displayed for all the interfaces on which the specified VRID is configured.

This display shows the following information.

TABLE 122 Output from the show ip vrrp vrid command

This field...	Displays...
VRID	The specified VRID.
Interface	The interface on which VRRP is configured.
State	This Layer 3 Switch VRRP state for the VRID. The state can be one of the following: <ul style="list-style-type: none"> Init – The VRID is not enabled (activated). If the state remains Init after you activate the VRID, make sure that the VRID is also configured on the other routers and that the routers can communicate with each other. <p>NOTE: If the state is Init and the mode is incomplete, make sure you have specified the IP address for the VRID:</p> <ul style="list-style-type: none"> Backup – This Layer 3 Switch is a Backup for the VRID. Master – This Layer 3 Switch is the Master for the VRID.
priority	The configured VRRP priority of this Layer 3 Switch for the VRID.
current priority	The current VRRP priority of this Layer 3 Switch for the VRID.
track-priority	The new VRRP priority that the router receives for this VRID if the interface goes down
hello-interval	How often the Master router sends Hello messages to the Backups.
dead-interval	The amount of time a Backup waits for a Hello message from the Master before determining that the Master is dead.
current dead-interval	The current Dead interval. The software automatically adds one-half second to the Dead interval value you enter.
preempt-mode	Whether the backup preempt mode is enabled. If the backup preempt mode is enabled, this field contains “true”. If the mode is disabled, this field contains “false”.
advertise backup	Whether Backup routers send Hello messages to the Master.

Displaying statistics

To display statistics on most devices, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip vrrp statistic

Interface ethernet 5
  rxd vrrp header error count = 0
  rxd vrrp auth error count = 0
  rxd vrrp auth passwd mismatch error count = 0
  rxd vrrp vrid not found error count = 0
  VRID 1
  rxd arp packet drop count = 0
  rxd ip packet drop count = 0
  rxd vrrp port mismatch count = 0
  rxd vrrp ip address mismatch count = 0
  rxd vrrp hello interval mismatch count = 0
  rxd vrrp priority zero from master count = 0
  rxd vrrp higher priority count = 0
  transitioned to master state count = 1
  transitioned to backup state count = 1
```

The same statistics are listed for VRRP and VRRPE.

Syntax: `show ip vrrp brief | ethernet <portnum> | ve <num> | statistic`

Syntax: `show ip vrrp-extended brief | ethernet <portnum> | ve <num> | stat`

The **brief** parameter displays summary information. Refer to [“Displaying summary information”](#) on page 732.

The **<portnum>** parameter specifies an Ethernet port. If you use this parameter, the command displays detailed VRRP or VRRPE information only for the specified port. Refer to [“Displaying detailed information”](#) on page 734.

The **ve <num>** parameter specifies a virtual interface. If you use this parameter, the command displays detailed VRRP or VRRPE information only for the specified virtual interface. Refer to [“Displaying detailed information”](#) on page 734.

The **statistic** parameter displays statistics. This parameter is required for displaying the statistics.

This display shows the following information.

TABLE 123 CLI display of VRRP or VRRPE statistics

This field...	Displays...
Interface statistics	
Interface	The interface on which VRRP or VRRPE is configured. If VRRP or VRRPE is configured on more than one interface, the display lists the statistics separately for each interface.
rxed vrrp header error count	The number of VRRP or VRRPE packets received by the interface that had a header error.
rxed vrrp auth error count	The number of VRRP or VRRPE packets received by the interface that had an authentication error.
rxed vrrp auth passwd mismatch error count	The number of VRRP or VRRPE packets received by the interface that had a password value that does not match the password used by the interface for authentication.

TABLE 123 CLI display of VRRP or VRRPE statistics (Continued)

This field...	Displays...
rxed vrrp vrid not found error count	The number of VRRP or VRRPE packets received by the interface that contained a VRID that is not configured on this interface.
VRID statistics	
rxed arp packet drop count	The number of ARP packets addressed to the VRID that were dropped.
rxed ip packet drop count	The number of IP packets addressed to the VRID that were dropped.
rxed vrrp port mismatch count	The number of packets received that did not match the configuration for the receiving interface.
rxed vrrp ip address mismatch count	The number of packets received that did not match the configured IP addresses.
rxed vrrp hello interval mismatch count	The number of packets received that did not match the configured Hello interval.
rxed vrrp priority zero from master count	The current Master has resigned.
rxed vrrp higher priority count	The number of VRRP or VRRPE packets received by the interface that had a higher backup priority for the VRID than this Layer 3 Switch backup priority for the VRID.
transitioned to master state count	The number of times this Layer 3 Switch has changed from the backup state to the master state for the VRID.
transitioned to backup state count	The number of times this Layer 3 Switch has changed from the master state to the backup state for the VRID.

Clearing VRRP or VRRPE statistics

Use the following methods to clear VRRP or VRRPE statistics.

To clear VRRP or VRRPE statistics, enter the following command at the Privileged EXEC level or any configuration level of the CLI.

```
Router1# clear ip vrrp-stat
```

Syntax: clear ip vrrp-stat

Displaying CPU utilization statistics

You can display CPU utilization statistics for VRRP and other IP protocols.

To display CPU utilization statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.

```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP              0.01        0.03        0.09        0.22         9
BGP              0.04        0.06        0.08        0.14        13
GVRP             0.00        0.00        0.00        0.00         0
ICMP             0.00        0.00        0.00        0.00         0
IP               0.00        0.00        0.00        0.00         0
OSPF             0.00        0.00        0.00        0.00         0
RIP              0.00        0.00        0.00        0.00         0
STP              0.00        0.00        0.00        0.00         0
VRRP           0.03        0.07        0.09        0.10         8
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example.

```
PowerConnect# show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP              0.01        0.00        0.00        0.00         0
BGP              0.00        0.00        0.00        0.00         0
GVRP             0.00        0.00        0.00        0.00         0
ICMP             0.01        0.00        0.00        0.00         1
IP               0.00        0.00        0.00        0.00         0
OSPF             0.00        0.00        0.00        0.00         0
RIP              0.00        0.00        0.00        0.00         0
STP              0.00        0.00        0.00        0.00         0
VRRP             0.00        0.00        0.00        0.00         0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.

```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name    Sec(%)    Time(ms)
ARP              0.00        0
BGP              0.00        0
GVRP             0.00        0
ICMP             0.01        1
IP               0.00        0
OSPF             0.00        0
RIP              0.00        0
STP              0.01        0
VRRP             0.00        0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The <num> parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Configuration examples

The following sections contain the CLI commands for implementing the VRRP and VRRPE configurations shown in [Figure 113](#) on page 713 and [Figure 114](#) on page 718.

VRRP example

To implement the VRRP configuration shown in [Figure 113](#) on page 713, use the following method.

Configuring Router1

To configure VRRP Router1, enter the following commands.

```
Router1(config)# router vrrp
Router1(config)# inter e 6
Router1(config-if-6)# ip address 192.53.5.1
Router1(config-if-6)# ip vrrp vrid 1
Router1(config-if-6-vrid-1)# owner track-priority 20
Router1(config-if-6-vrid-1)# track-port ethernet 4
Router1(config-if-6-vrid-1)# ip-address 192.53.5.1
Router1(config-if-6-vrid-1)# activate
```

NOTE

When you configure the Master (Owner), the address you enter with the **ip-address** command must already be configured on the interface.

The **ip vrrp owner** command specifies that this router owns the IP address you are associating with the VRID. Because this router owns the IP address, this router is the default Master router and its VRRP priority is thus 255.

Configuring Router2

To configure Router2 in [Figure 113](#) on page 713 after enabling VRRP, enter the following commands.

```
Router2(config)# router vrrp
Router2(config)# inter e 5
Router2(config-if-5)# ip address 192.53.5.3
Router2(config-if-5)# ip vrrp vrid 1
Router2(config-if-5-vrid-1)# backup priority 100 track-priority 19
Router2(config-if-5-vrid-1)# track-port ethernet 2
Router2(config-if-5-vrid-1)# ip-address 192.53.5.1
Router2(config-if-5-vrid-1)# activate
```

The **backup** command specifies that this router is a VRRP Backup for virtual router VRID1. The IP address entered with the **ip-address** command is the same IP address as the one entered when configuring Router1. In this case, the IP address cannot also exist on Router2, but the interface on which you are configuring the VRID Backup must have an IP address in the same subnet. By entering the same IP address as the one associated with this VRID on the Owner, you are configuring the Backup to back up the address, but you are not duplicating the address.

NOTE

When you configure a Backup router, the router interface on which you are configuring the VRID must have a real IP address that is in the same subnet as the address associated with the VRID by the Owner. However, the address cannot be the same.

The **priority** parameter establishes the router VRRP priority in relation to the other VRRP routers in this virtual router. The **track-priority** parameter specifies the new VRRP priority that the router receives for this VRID if the interface goes down. Refer to “[Track ports and track priority](#)” on page 715.

The **activate** command activates the VRID configuration on this interface. The interface does not provide backup service for the virtual IP address until you activate the VRRP configuration.

Syntax: router vrrp

Syntax: ip vrrp vrid <vrid>

Syntax: owner [track-priority <value>]

Syntax: backup [priority <value>] [track-priority <value>]

Syntax: track-port ethernet <portnum> | ve <num>

Syntax: ip-address <ip-addr>

Syntax: activate

VRRPE example

To implement the VRRPE configuration shown in [Figure 114](#) on page 718, use the following CLI method.

Configuring Router1

To configure VRRP Router1 in [Figure 114](#) on page 718, enter the following commands.

```
Router1(config)# router vrrp-extended
Router1(config)# interface ethernet 6
Router1(config-if-6)# ip address 192.53.5.2/24
Router1(config-if-6)# ip vrrp-extended vrid 1
Router1(config-if-6-vrid-1)# backup priority 110 track-priority 20
Router1(config-if-6-vrid-1)# track-port ethernet 4
Router1(config-if-6-vrid-1)# ip-address 192.53.5.254
Router1(config-if-6-vrid-1)# activate
VRRP router 1 for this interface is activating
Router1(config-if-6-vrid-1)# exit
Router1(config)# interface ethernet 6
Router1(config-if-6)# ip vrrp-extended vrid 2
Router1(config-if-6-vrid-1)# backup priority 100 track-priority 20
Router1(config-if-6-vrid-1)# track-port ethernet 4
Router1(config-if-6-vrid-1)# ip-address 192.53.5.253
Router1(config-if-6-vrid-1)# activate
VRRP router 2 for this interface is activating
```

NOTE

The address you enter with the **ip-address** command cannot be the same as a real IP address configured on the interface.

Configuring Router2

To configure Router2, enter the following commands.

```
Router1(config)# router vrrp-extended
Router1(config)# interface ethernet 5
Router1(config-if-5)# ip address 192.53.5.3/24
Router1(config-if-5)# ip vrrp-extended vrid 1
Router1(config-if-5-vrid-1)# backup priority 100 track-priority 20
Router1(config-if-5-vrid-1)# track-port ethernet 2
Router1(config-if-5-vrid-1)# ip-address 192.53.5.254
Router1(config-if-5-vrid-1)# activate
VRRP router 1 for this interface is activating
Router1(config-if-5-vrid-1)# exit
Router1(config)# interface ethernet 5
Router1(config-if-5)# ip vrrp-extended vrid 2
Router1(config-if-5-vrid-1)# backup priority 110 track-priority 20
Router1(config-if-5-vrid-1)# track-port ethernet 4
Router1(config-if-5-vrid-1)# ip-address 192.53.5.253
Router1(config-if-5-vrid-1)# activate
VRRP router 2 for this interface is activating
```

The **backup** command specifies that this router is a VRRPE Backup for virtual router VRID1. The IP address entered with the **ip-address** command is the same IP address as the one entered when configuring Router1. In this case, the IP address cannot also exist on Router2, but the interface on which you are configuring the VRID Backup must have an IP address in the same subnet. By entering the same IP address as the one associated with this VRID on the Owner, you are configuring the Backup to back up the address, but you are not duplicating the address.

NOTE

When you configure a Backup router, the router interface on which you are configuring the VRID must have a real IP address that is in the same subnet as the address associated with the VRID by the Owner. However, the address cannot be the same.

The **priority** parameter establishes the router VRRPE priority in relation to the other VRRPE routers in this virtual router. The **track-priority** parameter specifies the new VRRPE priority that the router receives for this VRID if the interface goes down. Refer to [“Track ports and track priority”](#) on page 715.

The **activate** command activates the VRID configuration on this interface. The interface does not provide backup service for the virtual IP address until you activate the VRRPE configuration. Alternatively, you can use the **enable** command. The **activate** and **enable** commands do the same thing.

Syntax: router vrrp-extended

Syntax: ip vrrp-extended vrid <vrid>

Syntax: backup [priority <value>] [track-priority <value>]

Syntax: track-port ethernet<portnum> | ve <num>

Syntax: ip-address <ip-addr>

Syntax: activate

Configuring BGP4

This chapter provides details on how to configure **Border Gateway Protocol version 4 (BGP4)** using the CLI.

BGP4 is described in RFC 1771. The Dell implementation fully complies with RFC 1771. The Dell BGP4 implementation also supports the following RFCs:

- RFC 1745 (OSPF Interactions)
- RFC 1997 (BGP Communities Attributes)
- RFC 2385 (TCP MD5 Signature Option)
- RFC 2439 (Route Flap Dampening)
- RFC 2796 (Route Reflection)
- RFC 2842 (Capability Advertisement)

To display BGP4 configuration information and statistics, refer to [“Displaying BGP4 information”](#) on page 816.

This chapter shows the commands you need in order to configure the Layer 3 Switch for BGP4.

NOTE

Your Layer 3 Switch management module must have 32MB or higher to run BGP4.

NOTE

PowerConnect B-Series TI24X devices support up to 12,000 BGP routes.

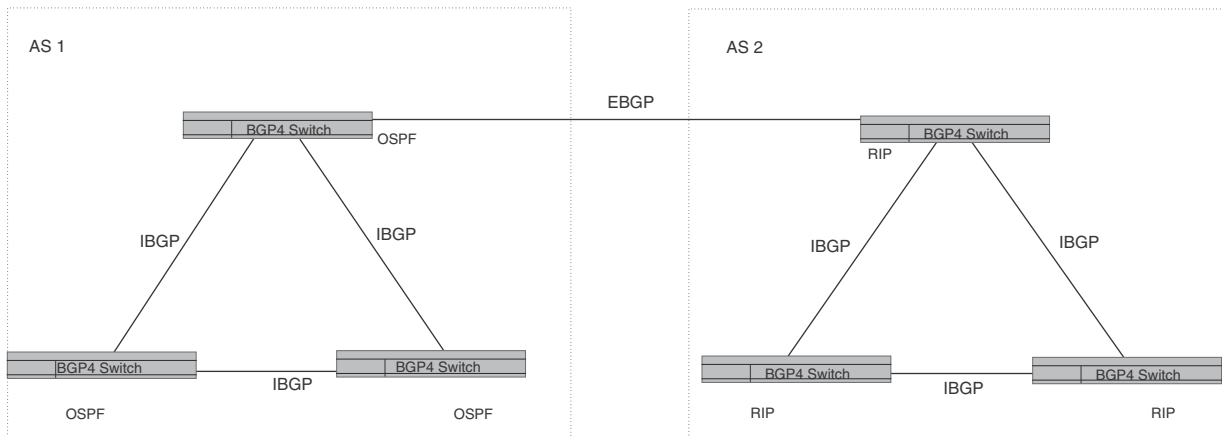
Overview of BGP4

BGP4 is the standard Exterior Gateway Protocol (EGP) used on the Internet to route traffic between **Autonomous Systems (AS)** and to maintain loop-free routing. An autonomous system is a collection of networks that share the same routing and administration characteristics. For example, a corporate intranet consisting of several networks under common administrative control might be considered an AS. The networks in an AS can but do not need to run the same routing protocol to be in the same AS, nor do they need to be geographically close.

Routers within an AS can use different Interior Gateway Protocols (IGPs) such as RIP and OSPF to communicate with one another. However, for routers in different ASs to communicate, they need to use an EGP. BGP4 is the standard EGP used by Internet routers and therefore is the EGP implemented on Layer 3 Switches.

Figure 115 on page 746 shows a simple example of two BGP4 ASs. Each AS contains three BGP4 switches. All of the BGP4 switches within an AS communicate using IBGP. BGP4 switches communicate with other ASs using EBGP. Notice that each of the switches also is running an Interior Gateway Protocol (IGP). The switches in AS1 are running OSPF and the switches in AS2 are running RIP. Layer 3 Switches can be configured to redistribute routes among BGP4, RIP, and OSPF. They also can redistribute static routes.

FIGURE 115 Example BGP4 ASs



Relationship between the BGP4 route table and the IP route table

The Layer 3 Switch BGP4 route table can have multiple routes to the same destination, which are learned from different BGP4 neighbors. A BGP4 neighbor is another switch that also is running BGP4. BGP4 neighbors communicate using Transmission Control Protocol (TCP) port 179 for BGP communication. When you configure the Layer 3 Switch for BGP4, one of the configuration tasks you perform is to identify the Layer 3 Switch BGP4 neighbors.

Although a Layer 3 Switch BGP4 route table can have multiple routes to the same destination, the BGP4 protocol evaluates the routes and chooses only one of the routes to send to the IP route table. The route that BGP4 chooses and sends to the IP route table is the **preferred route** and will be used by the Layer 3 Switch. If the preferred route goes down, BGP4 updates the route information in the IP route table with a new BGP4 preferred route.

NOTE

If IP load sharing is enabled and you enable multiple equal-cost paths for BGP4, BGP4 can select more than one equal-cost path to a destination.

A BGP4 route consists of the following information:

- **Network number (prefix)** – A value comprised of the network mask bits and an IP address (*<IP address>/ <mask bits>*); for example, 192.215.129.0/18 indicates a network mask of 18 bits applied to the IP address 192.215.129.0. When a BGP4 Layer 3 Switch advertises a route to one of its neighbors, the route is expressed in this format.

- **AS-path** – A list of the other ASs through which a route passes. BGP4 routers can use the AS-path to detect and eliminate routing loops. For example, if a route received by a BGP4 router contains the AS that the router is in, the router does not add the route to its own BGP4 table. (The BGP4 RFCs refer to the AS-path as “AS_PATH”.)
- **Additional path attributes** – A list of additional parameters that describe the route. The route origin and next hop are examples of these additional path attributes.

NOTE

The Layer 3 Switch re-advertises a learned best BGP4 route to the Layer 3 Switch neighbors even when the software does not select that route for installation in the IP route table. The best BGP4 route is the route that the software selects based on comparison of the BGP4 route path attributes.

After a Layer 3 Switch successfully negotiates a BGP4 session with a neighbor (a BGP4 peer), the Layer 3 Switch exchanges complete BGP4 route tables with the neighbor. After this initial exchange, the Layer 3 Switch and all other RFC 1771-compliant BGP4 routers send UPDATE messages to inform neighbors of new, changed, or no longer feasible routes. BGP4 routers do not send regular updates. However, if configured to do so, a BGP4 router does regularly send KEEPALIVE messages to its peers to maintain BGP4 sessions with them if the router does not have any route information to send in an UPDATE message. Refer to “[BGP4 message types](#)” on page 748 for information about BGP4 messages.

How BGP4 selects a path for a route

When multiple paths for the same route are known to a BGP4 router, the router uses the following algorithm to weigh the paths and determine the optimal path for the route. The optimal path depends on various parameters, which can be modified. (Refer to “[Optional configuration tasks](#)” on page 767.)

1. Is the next hop accessible through an Interior Gateway Protocol (IGP) route? If not, ignore the route.

NOTE

The device does not use the default route to resolve BGP4 next hop. Also refer to “[Enabling next-hop recursion](#)” on page 774.

2. Use the path with the largest weight.
3. If the weights are the same, prefer the route with the largest local preference.
4. If the routes have the same local preference, prefer the route that was originated locally (by this BGP4 Layer 3 Switch).
5. If the local preferences are the same, prefer the route with the shortest AS-path. An AS-SET counts as 1. A confederation path length, if present, is not counted as part of the path length.
6. If the AS-path lengths are the same, prefer the route with the lowest origin type. From low to high, route origin types are valued as follows:
 - IGP is lowest
 - EGP is higher than IGP but lower than INCOMPLETE
 - INCOMPLETE is highest

7. If the routes have the same origin type, prefer the route with the lowest MED. For a definition of MED, refer to [“Configuring the Layer 3 Switch to always compare Multi-Exit Discriminators \(MEDs\)”](#) on page 779.

BGP4 compares the MEDs of two otherwise equivalent paths if and only if the routes were learned from the same neighboring AS. This behavior is called deterministic MED.

Deterministic MED is always enabled and cannot be disabled. In addition, you can enable the Layer 3 Switch to always compare the MEDs, regardless of the AS information in the paths. To enable this comparison, enter the **always-compare-med** command at the BGP4 configuration level of the CLI. This option is disabled by default.

NOTE

By default, value 0 (most favorable) is used in MED comparison when the MED attribute is not present. The default MED comparison results in the Layer 3 Switch favoring the route paths that are missing their MEDs. You can use the **med-missing-as-worst** command to make the Layer 3 Switch regard a BGP route with a missing MED attribute as the least favorable route, when comparing the MEDs of the routes.

NOTE

MED comparison is not performed for internal routes originated within the local AS or confederation.

8. Prefer routes in the following order:
 - Routes received through EBGP from a BGP4 neighbor outside of the confederation
 - Routes received through EBGP from a BGP4 router within the confederation
 - Routes received through IBGP
9. If all the comparisons above are equal, prefer the route with the lowest IGP metric to the BGP4 next hop. This is the closest internal path inside the AS to reach the destination.
10. If the internal paths also are the same and BGP4 load sharing is enabled, load share among the paths. Otherwise, prefer the route that comes from the BGP4 router with the lowest router ID.

NOTE

Layer 3 Switches support BGP4 load sharing among multiple equal-cost paths. BGP4 load sharing enables the Layer 3 Switch to balance the traffic across the multiple paths instead of choosing just one path based on router ID. For EBGP routes, load sharing applies only when the paths are from neighbors within the same remote AS. EBGP paths from neighbors in different ASs are not compared.

BGP4 message types

BGP4 routers communicate with their neighbors (other BGP4 routers) using the following types of messages:

- OPEN
- UPDATE
- KEEPALIVE
- NOTIFICATION

OPEN message

After a BGP4 router establishes a TCP connection with a neighboring BGP4 router, the routers exchange OPEN messages. An OPEN message indicates the following:

- **BGP version** – Indicates the version of the protocol that is in use on the router. BGP version 4 supports Classless Interdomain Routing (CIDR) and is the version most widely used in the Internet. Version 4 also is the only version supported on Layer 3 Switches.
- **AS number** – A two-byte number that identifies the AS to which the BGP4 router belongs.
- **Hold Time** – The number of seconds a BGP4 router will wait for an UPDATE or KEEPALIVE message (described below) from a BGP4 neighbor before assuming that the neighbor is dead. BGP4 routers exchange UPDATE and KEEPALIVE messages to update route information and maintain communication. If BGP4 neighbors are using different Hold Times, the lowest Hold Time is used by the neighbors. If the Hold Time expires, the BGP4 router closes its TCP connection to the neighbor and clears any information it has learned from the neighbor and cached.

You can configure the Hold Time to be 0, in which case a BGP4 router will consider its neighbors to always be up. For directly-attached neighbors, you can configure the Layer 3 Switch to immediately close the TCP connection to the neighbor and clear entries learned from an EBGP neighbor if the interface to that neighbor goes down. This capability is provided by the fast external fallover feature, which is disabled by default.

- **BGP Identifier** – The router ID. The BGP Identifier (router ID) identifies the BGP4 router to other BGP4 routers. Layer 3 Switches use the same router ID for OSPF and BGP4. If you do not set a router ID, the software uses the IP address on the lowest numbered loopback interface configured on the router. If the Layer 3 Switch does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, refer to “[Changing the router ID](#)” on page 584.
- **Parameter list** – An optional list of additional parameters used in peer negotiation with BGP4 neighbors.

UPDATE message

After BGP4 neighbors establish a BGP4 connection over TCP and exchange their BGP4 routing tables, they do not send periodic routing updates. Instead, a BGP4 neighbor sends an update to its neighbor when it has a new route to advertise or routes have changed or become unfeasible. An UPDATE message can contain the following information:

- **Network Layer Reachability Information (NLRI)** – The mechanism by which BGP4 supports Classless Interdomain Routing (CIDR). An NLRI entry consists of an IP prefix that indicates a network being advertised by the UPDATE message. The prefix consists of an IP network number and the length of the network portion of the number. For example, an UPDATE message with the NLRI entry 192.215.129.0/18 indicates a route to IP network 192.215.129.0 with network mask 255.255.192.0. The binary equivalent of this mask is 18 consecutive one bits, thus “18” in the NLRI entry.
- **Path attributes** – Parameters that indicate route-specific information such as path information, route preference, next hop values, and aggregation information. BGP4 uses the path attributes to make filtering and routing decisions.

- **Unreachable routes** – A list of routes that have been in the sending router BGP4 table but are no longer feasible. The UPDATE message lists unreachable routes in the same format as new routes.
<IP address>/<CIDR prefix>.

KEEPALIVE message

BGP4 routers do not regularly exchange UPDATE messages to maintain the BGP4 sessions. For example, if a Layer 3 Switch configured to perform BGP4 routing has already sent the latest route information to its peers in UPDATE messages, the router does not send more UPDATE messages. Instead, BGP4 routers send KEEPALIVE messages to maintain the BGP4 sessions. KEEPALIVE messages are 19 bytes long and consist only of a message header; they contain no routing data.

BGP4 routers send KEEPALIVE messages at a regular interval, the Keep Alive Time. The default Keep Alive Time on Layer 3 Switches is 60 seconds.

A parameter related to the Keep Alive Time is the Hold Time. A BGP4 router Hold Time determines how many seconds the router will wait for a KEEPALIVE or UPDATE message from a BGP4 neighbor before deciding that the neighbor is dead. The Hold Time is negotiated when BGP4 routers exchange OPEN messages; the lower Hold Time is then used by both neighbors. For example, if BGP4 Router A sends a Hold Time of 5 seconds and BGP4 Router B sends a Hold Time of 4 seconds, both routers use 4 seconds as the Hold Time for their BGP4 session. The default Hold Time is 180 seconds. Generally, the Hold Time is configured to three times the value of the Keep Alive Time.

If the Hold Time is 0, a BGP4 router assumes that its neighbor is alive regardless of how many seconds pass between receipt of UPDATE or KEEPALIVE messages.

NOTIFICATION message

When you close the router BGP4 session with a neighbor, or the router detects an error in a message received from the neighbor, or an error occurs on the router, the router sends a NOTIFICATION message to the neighbor. No further communication takes place between the BGP4 router that sent the NOTIFICATION and the neighbors that received the NOTIFICATION.

Basic configuration and activation for BGP4

BGP4 is disabled by default. Follow the steps given below to enable BGP4 and place your Layer 3 Switch into service as a BGP4 router.

1. Enable the BGP4 protocol.
2. Set the local AS number.

NOTE

You must specify the local AS number for BGP4 to become functional.

3. Add each BGP4 neighbor (peer BGP4 router) and identify the AS the neighbor is in.
4. Save the BGP4 configuration information to the system configuration file.

NOTE

By default, the router ID is the IP address configured on the lowest numbered loopback interface. If the Layer 3 Switch does not have a loopback interface, the default router ID is the lowest numbered IP interface address configured on the device. For more information or to change the router ID, refer to “[Changing the router ID](#)” on page 584. If you change the router ID, all current BGP4 sessions are cleared.

```
PowerConnect># enable
PowerConnect# configure terminal
PowerConnect(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
PowerConnect(config-bgp-router)# local-as 10
PowerConnect(config-bgp-router)# neighbor 209.157.23.99 remote-as 100
PowerConnect(config-bgp-router)# write memory
```

NOTE

When BGP4 is enabled on a Layer 3 Switch, you do not need to reset the system. The protocol is activated as soon as you enable it. Moreover, the router begins a BGP4 session with a BGP4 neighbor as soon as you add the neighbor.

Note regarding disabling BGP4

If you disable BGP4, the Layer 3 Switch removes all the running configuration information for the disabled protocol from the running-config. To restore the BGP4 configuration, you must reload the software to load the configuration from the startup-config. Moreover, when you save the configuration to the startup-config file after disabling the protocol, all the configuration information for the disabled protocol is removed from the startup-config file.

The CLI displays a warning message such as the following.

```
PowerConnect(config-bgp-router)# no router bgp
router bgp mode now disabled. All bgp config data will be lost when writing to
flash!
```

If you are testing a BGP4 configuration and are likely to disable and re-enable the protocol, you might want to make a backup copy of the startup-config file containing the protocol configuration information. This way, if you remove the configuration information by saving the configuration after disabling the protocol, you can restore the configuration by copying the backup copy of the startup-config file onto the flash memory.

NOTE

To disable BGP4 without losing the BGP4 configuration information, remove the local AS (for example, by entering the **no local-as <num>** command). In this case, BGP4 retains the other configuration information but is not operational until you set the local AS again.

BGP4 parameters

You can modify or set the following BGP4 parameters:

- Optional – Define the router ID. (The same router ID also is used by OSPF.)
- Required – Specify the local AS number.
- Optional – Add a loopback interface for use with neighbors.

- Required – Identify BGP4 neighbors.
- Optional – Change the Keep Alive Time and Hold Time.
- Optional – Change the update timer for route changes.
- Optional – Enable fast external fallover.
- Optional – Specify a list of individual networks in the local AS to be advertised to remote ASs using BGP4.
- Optional – Change the default local preference for routes.
- Optional – Enable the default route (default-information-originate).
- Optional – Enable use of a default route to resolve a BGP4 next-hop route.
- Optional – Change the default MED (metric).
- Optional – Enable next-hop recursion.
- Optional – Change the default administrative distances for EBGp, IBGP, and locally originated routes.
- Optional – Require the first AS in an Update from an EBGp neighbor to be the neighbor AS.
- Optional – Change MED comparison parameters.
- Optional – Disable comparison of the AS-Path length.
- Optional – Enable comparison of the router ID.
- Optional – Enable auto summary to summarize routes at an IP class boundary (A, B, or C).
- Optional – Aggregate routes in the BGP4 route table into CIDR blocks.
- Optional – Configure the router as a BGP4 router reflector.
- Optional – Configure the Layer 3 Switch as a member of a BGP4 confederation.
- Optional – Change the default metric for routes that BGP4 redistributes into RIP or OSPF.
- Optional – Change the parameters for RIP, OSPF, or static routes redistributed into BGP4.
- Optional – Change the number of paths for BGP4 load sharing.
- Optional – Change other load-sharing parameters
- Optional – Define BGP4 address filters.
- Optional – Define BGP4 AS-path filters.
- Optional – Define BGP4 community filters.
- Optional – Define IP prefix lists.
- Optional – Define neighbor distribute lists.
- Optional – Define BGP4 route maps for filtering routes redistributed into RIP and OSPF.
- Optional – Define route flap dampening parameters.

NOTE

When using CLI, you set global level parameters at the BGP CONFIG Level of the CLI. You can reach the BGP CONFIG level by entering **router bgp...** at the global CONFIG level.

When parameter changes take effect

Some parameter changes take effect immediately while others do not take full effect until the router sessions with its neighbors are reset. Some parameters do not take effect until the router is rebooted.

Immediately

The following parameter changes take effect immediately:

- Enable or disable BGP.
- Set or change the local AS.
- Add neighbors.
- Change the update timer for route changes.
- Disable or enable fast external fallover.
- Specify individual networks that can be advertised.
- Change the default local preference, default information originate setting, or administrative distance.
- Enable or disable use of a default route to resolve a BGP4 next-hop route.
- Enable or disable MED (metric) comparison.
- Require the first AS in an Update from an EBGp neighbor to be the neighbor AS.
- Change MED comparison parameters.
- Disable comparison of the AS-Path length.
- Enable comparison of the router ID.
- Enable next-hop recursion.
- Enable or disable auto summary.
- Change the default metric.
- Disable or re-enable route reflection.
- Configure confederation parameters.
- Disable or re-enable load sharing.
- Change the maximum number of load-sharing paths.
- Change other load-sharing parameters.
- Define route flap dampening parameters.
- Add, change, or negate redistribution parameters (except changing the default MED; see below).
- Add, change, or negate route maps (when used by the **network** command or a redistribution command).

After resetting neighbor sessions

The following parameter changes take effect only after the router BGP4 sessions are cleared, or reset using the “soft” clear option. (Refer to [“Closing or resetting a neighbor session”](#) on page 851.)

The parameter are as follows:

- Change the Hold Time or Keep Alive Time.
- Aggregate routes.
- Add, change, or negate filter tables.

After disabling and re-enabling redistribution

The following parameter change takes effect only after you disable and then re-enable redistribution:

- Change the default MED (metric).

Memory considerations

BGP4 handles a very large number of routes and therefore requires a lot of memory. For example, in a typical configuration with just a single BGP4 neighbor, a BGP4 router may need to be able to hold up to 80,000 routes. Many configurations, especially those involving more than one neighbor, can require the router to hold even more routes. Layer 3 Switches provide dynamic memory allocation for BGP4 data. These devices automatically allocate memory when needed to support BGP4 neighbors, routes, and route attribute entries. Dynamic memory allocation is performed automatically by the software and does not require a reload.

The memory amounts, including routes received from neighbors, BGP route advertisements (routes sent to neighbors), and BGP route attribute entries. The routes sent to and received from neighbors use the most BGP4 memory. Generally, the actual limit to the number of neighbors, routes, or route attribute entries the device can accommodate depends on how many routes the Layer 3 Switch sends to and receives from the neighbors.

In some cases, where most of the neighbors do not send or receive a full BGP route table (about 80,000 routes), the memory can support a larger number of BGP4 neighbors. However, if most of the BGP4 neighbors send or receive full BGP route tables, the number of BGP neighbors the memory can support is less than in configurations where the neighbors send smaller route tables.

As a guideline, Layer 3 Switches with a 512 MB Management 4 module can accommodate 150 – 200 neighbors, with the assumption that the Layer 3 Switch receives about one million routes total from all neighbors and sends about eight million routes total to neighbors. For each additional one million incoming routes, the capacity for outgoing routes decreases by around two million.

Memory configuration options obsoleted by dynamic memory

Devices that support dynamic BGP4 memory allocation do not require or even support static configuration of memory for BGP4 neighbors, routes, or route attributes. Consequently, the following CLI commands are not supported on these devices:

- **max-neighbors** <num>
- **max-routes** <num>
- **max-attribute-entries** <num>

If you boot a device that has a startup-config file that contains these commands, the software ignores the commands and uses dynamic memory allocation for BGP4. The first time you save the device running configuration (running-config) to the startup-config file, the commands are removed from the file.

Basic configuration tasks

The following sections describe how to perform the configuration tasks that are required to use BGP4 on the Layer 3 Switch. You can modify many parameters in addition to the ones described in this section. Refer to “[Optional configuration tasks](#)” on page 767.

Enabling BGP4 on the router

When you enable BGP4 on the router, BGP4 is automatically activated. To enable BGP4 on the router, enter the following commands.

```
PowerConnect> enable
PowerConnect# configure terminal
PowerConnect(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
PowerConnect(config-bgp-router)# local-as 10
PowerConnect(config-bgp-router)# neighbor 209.157.23.99 remote-as 100
PowerConnect(config-bgp-router)# write memory
```

Changing the router ID

The OSPF and BGP4 protocols use router IDs to identify the routers that are running the protocols. A router ID is a valid, unique IP address and sometimes is an IP address configured on the router. The router ID cannot be an IP address in use by another device.

By default, the router ID on a Layer 3 Switch is one of the following:

- If the router has loopback interfaces, the default router ID is the IP address configured on the lowest numbered loopback interface configured on the Layer 3 Switch. For example, if you configure loopback interfaces 1, 2, and 3 as follows, the default router ID is 9.9.9.9/24:
 - Loopback interface 1, 9.9.9.9/24
 - Loopback interface 2, 4.4.4.4/24
 - Loopback interface 3, 1.1.1.1/24
- If the device does not have any loopback interfaces, the default router ID is the lowest numbered IP interface address configured on the device.

NOTE

Layer 3 Switches use the same router ID for both OSPF and BGP4. If the router is already configured for OSPF, you may want to use the router ID that is already in use on the router rather than set a new one. To display the router ID, enter the **show ip** CLI command at any CLI level.

To change the router ID, enter a command such as the following.

```
PowerConnect(config)# ip router-id 209.157.22.26
```

Syntax: **ip router-id** <ip-addr>

The <ip-addr> can be any valid, unique IP address.

NOTE

You can specify an IP address used for an interface on the Layer 3 Switch, but do not specify an IP address in use by another device.

Setting the local AS number

The local AS number identifies the AS the BGP4 router is in. The AS number can be from 1 – 65535. There is no default. AS numbers 64512 – 65535 are the well-known private BGP4 AS numbers and are not advertised to the Internet community.

To set the local AS number, enter commands such as the following.

```
PowerConnect(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
PowerConnect(config-bgp-router)# local-as 10
PowerConnect(config-bgp-router)# write memory
```

Syntax: [no] local-as <num>

The <num> parameter specifies the local AS number.

Adding a loopback interface

You can configure the router to use a loopback interface instead of a specific port or virtual routing interface to communicate with a BGP4 neighbor. A loopback interface adds stability to the network by working around route flap problems that can occur due to unstable links between the router and its neighbors.

Loopback interfaces are always up, regardless of the states of physical interfaces. Loopback interfaces are especially useful for IBGP neighbors (neighbors in the same AS) that are multiple hops away from the router. When you configure a BGP4 neighbor on the router, you can specify whether the router uses the loopback interface to communicate with the neighbor. As long as a path exists between the router and its neighbor, BGP4 information can be exchanged. The BGP4 session is not associated with a specific link but instead is associated with the virtual interfaces.

You can add up to 24 IP addresses to each loopback interface.

NOTE

If you configure the Layer 3 Switch to use a loopback interface to communicate with a BGP4 neighbor, the peer IP address on the remote router pointing to your loopback address must be configured.

To add a loopback interface, enter commands such as those shown in the following example.

```
PowerConnect(config-bgp-router)# exit
PowerConnect(config)# int loopback 1
PowerConnect(config-lbif-1)# ip address 10.0.0.1/24
```

Syntax: interface loopback <num>

The <num> value can be from 1 – 8 on Chassis Layer 3 Switches. The value can be from 1 – 4 on the Compact Layer 3 Switch.

Adding BGP4 neighbors

The BGP4 protocol does not contain a peer discovery process. Therefore, for each of the router BGP4 neighbors (peers), you must indicate the neighbor IP address and the AS each neighbor is in. Neighbors that are in different ASs communicate using EBGP. Neighbors within the same AS communicate using IBGP.

NOTE

If the Layer 3 Switch has multiple neighbors with similar attributes, you can simplify configuration by configuring a peer group, then adding individual neighbors to it. The configuration steps are similar, except you specify a peer group name instead of a neighbor IP address when configuring the neighbor parameters, then add individual neighbors to the peer group. Refer to [“Adding a BGP4 peer group”](#) on page 763.

NOTE

The Layer 3 Switch attempts to establish a BGP4 session with a neighbor as soon as you enter a command specifying the neighbor IP address. If you want to completely configure the neighbor parameters before the Layer 3 Switch establishes a session with the neighbor, you can administratively shut down the neighbor. Refer to [“Administratively shutting down a session with a BGP4 neighbor”](#) on page 766.

To add a BGP4 neighbor with IP address 209.157.22.26, enter the following command.

```
PowerConnect(config-bgp-router)# neighbor 209.157.22.26
```

The neighbor *<ip-addr>* must be a valid IP address.

The **neighbor** command has some additional parameters, as shown in the following syntax:

```
Syntax: [no] neighbor <ip-addr> | <peer-group-name>
[advertisement-interval <num>]
[capability orf prefixlist [send | receive]]
[default-originate [route-map <map-name>]]
[description <string>]
[distribute-list in | out <num,num,...> | <ACL-num> in | out]
[ebgp-multihop [<num>]]
[filter-list in | out <num,num,...> | <ACL-num> in | out | weight]
[maximum-prefix <num> [<threshold>] [teardown]]
[next-hop-self]
[nlri multicast | unicast | multicast unicast]
[password [0 | 1] <string>]
[prefix-list <string> in | out]
[remote-as <as-number>]
[remove-private-as]
[route-map in | out <map-name>]
[route-reflector-client]
[send-community]
[soft-reconfiguration inbound]
[shutdown]
[timers keep-alive <num> hold-time <num>]
[unsuppress-map <map-name>]
[update-source <ip-addr> | ethernet<portnum> | loopback <num> | ve <num>]
[weight <num>]
```

The *<ip-addr>* | *<peer-group-name>* parameter indicates whether you are configuring an individual neighbor or a peer group. If you specify a neighbor IP address, you are configuring that individual neighbor. If you specify a peer group name, you are configuring a peer group. Refer to [“Adding a BGP4 peer group”](#) on page 763.

advertisement-interval *<num>* specifies the minimum delay (in seconds) between messages to the specified neighbor. The default is 30 for EBGp neighbors (neighbors in other ASs). The default is 5 for IBGP neighbors (neighbors in the same AS). The range is 0 – 600.

NOTE

The Layer 3 Switch applies the advertisement interval only under certain conditions. The Layer 3 Switch does not apply the advertisement interval when sending initial updates to a BGP4 neighbor. As a result, the Layer 3 Switch sends the updates one immediately after another, without waiting for the advertisement interval.

capability orf prefixlist [send | receive] configures cooperative router filtering. The **send | receive** parameter specifies the support you are enabling:

- **send** – The Layer 3 Switch sends the IP prefix lists as Outbound Route Filters (ORFs) to the neighbor.
- **receive** – The Layer 3 Switch accepts filters as Outbound Route Filters (ORFs) from the neighbor.

If you do not specify the capability, both capabilities are enabled. The **prefixlist** parameter specifies the type of filter you want to send to the neighbor.

For more information, refer to [“Configuring cooperative BGP4 route filtering”](#) on page 806.

NOTE

The current release supports cooperative filtering only for filters configured using IP prefix lists.

default-originate [route-map <map-name>] configures the Layer 3 Switch to send the default route 0.0.0.0 to the neighbor. If you use the route-map <map-name> parameter, the route map injects the default route conditionally, based on the match conditions in the route map.

description <string> specifies a name for the neighbor. You can enter an alphanumeric text string up to 80 characters long.

distribute-list in | out <num,num,...> specifies a distribute list to be applied to updates to or from the specified neighbor. The **in | out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. The <num,num,...> parameter specifies the list of address-list filters. The router applies the filters in the order in which you list them and stops applying the filters in the distribute list when a match is found.

Alternatively, you can specify **distribute-list <ACL-num> in | out** to use an IP ACL instead of a distribute list. In this case, <ACL-num> is an IP ACL.

NOTE

By default, if a route does not match any of the filters, the Layer 3 Switch denies the route. To change the default behavior, configure the last filter as “permit any any”.

NOTE

The address filter must already be configured. Refer to [“Filtering specific IP addresses”](#) on page 788.

ebgp-multihop [<num>] specifies that the neighbor is more than one hop away and that the session type with the neighbor is thus EBGp-multihop. This option is disabled by default. The <num> parameter specifies the TTL you are adding for the neighbor. You can specify a number from 0 – 255. The default is 0. If you leave the EBGp TTL value set to 0, the software uses the IP TTL value.

filter-list in | out <num,num,...> specifies an AS-path filter list or a list of AS-path ACLs. The **in | out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. If you specify in or out, The <num,num,...> parameter specifies the list of AS-path filters. The router applies the filters in the order in which you list them and stops applying the filters in the AS-path filter list when a match is found. The **weight** <num> parameter specifies a weight that the Layer 3 Switch applies to routes received from the neighbor that match the AS-path filter or ACL. You can specify a number from 0 – 65535.

Alternatively, you can specify filter-list <ACL-num> **in | out | weight** to use an AS-path ACL instead of an AS-path filter list. In this case, <ACL-num> is an AS-path ACL.

NOTE

By default, if an AS-path does not match any of the filters or ACLs, the Layer 3 Switch denies the route. To change the default behavior, configure the last filter or ACL as “permit any any”.

NOTE

The AS-path filter or ACL must already be configured. Refer to “[Filtering AS-paths](#)” on page 790.

maximum-prefix <num> specifies the maximum number of IP network prefixes (routes) that can be learned from the specified neighbor or peer group. You can specify a value from 0 – 4294967295. The default is 0 (unlimited):

- The <num> parameter specifies the maximum number. You can specify a value from 0 – 4294967295. The default is 0 (unlimited).
- The <threshold> parameter specifies the percentage of the value you specified for the **maximum-prefix** <num>, at which you want the software to generate a Syslog message. You can specify a value from 1 (one percent) to 100 (100 percent). The default is 100.
- The **teardown** parameter tears down the neighbor session if the maximum-prefix limit is exceeded. The session remains shutdown until you clear the prefixes using the **clear ip bgp neighbor all** or **clear ip bgp neighbor** <ip-addr> command, or change the neighbor maximum-prefix configuration. The software also generates a Syslog message.

next-hop-self specifies that the router should list itself as the next hop in updates sent to the specified neighbor. This option is disabled by default.

The **nlri multicast | unicast | multicast** unicast parameter specifies whether the neighbor is a multicast neighbor or a unicast neighbor. Optionally, you also can specify unicast if you want the Layer 3 Switch to exchange unicast (BGP4) routes as well as multicast routes with the neighbor. The default is unicast only.

password [0 | 1] <string> specifies an MD5 password for securing sessions between the Layer 3 Switch and the neighbor. You can enter a string up to 80 characters long. The string can contain any alphanumeric characters, but the first character cannot be a number. If the password contains a number, do not enter a space following the number.

The **0 | 1** parameter is the encryption option, which you can omit (the default) or which can be one of the following:

- **0** – Disables encryption for the authentication string you specify with the command. The password or string is shown as clear text in the output of commands that display neighbor or peer group configuration information.
- **1** – Assumes that the authentication string you enter is the encrypted form, and decrypts the value before using it.

For more information, refer to “[Encryption of BGP4 MD5 authentication keys](#)” on page 761.

NOTE

If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter 0 or 1. Instead, omit the encryption option and allow the software to use the default behavior.

If you specify encryption option 1, the software assumes that you are entering the encrypted form of the password or authentication string. In this case, the software decrypts the password or string you enter before using the value for authentication. If you accidentally enter option 1 followed by the clear-text version of the password or string, authentication will fail because the value used by the software will not match the value you intended to use.

prefix-list <string> **in** | **out** specifies an IP prefix list. You can use IP prefix lists to control routes to and from the neighbor. IP prefix lists are an alternative method to AS-path filters. The **in** | **out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. You can configure up to 1000 prefix list filters. The filters can use the same prefix list or different prefix lists. To configure an IP prefix list, refer to [“Defining IP prefix lists”](#) on page 795.

remote-as <as-number> specifies the AS the remote neighbor is in. The <as-number> can be a number from 1 – 65535. There is no default.

remove-private-as configures the router to remove private AS numbers from UPDATE messages the router sends to this neighbor. The router will remove AS numbers 64512 – 65535 (the well-known BGP4 private AS numbers) from the AS-path attribute in UPDATE messages the Layer 3 Switch sends to the neighbor. This option is disabled by default.

route-map in | **out** <map-name> specifies a route map the Layer 3 Switch will apply to updates sent to or received from the specified neighbor. The **in** | **out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor.

NOTE

The route map must already be configured. Refer to [“Defining route maps”](#) on page 797.

route-reflector-client specifies that this neighbor is a route-reflector client of the router. Use the parameter only if this router is going to be a route reflector. For information, refer to [“Configuring route reflection parameters”](#) on page 780. This option is disabled by default.

send-community enables sending the community attribute in updates to the specified neighbor. By default, the router does not send the community attribute.

shutdown administratively shuts down the session with this neighbor. Shutting down the session allows you to completely configure the neighbor and save the configuration without actually establishing a session with the neighbor. This option is disabled by default.

soft-reconfiguration inbound enables the soft reconfiguration feature, which stores all the route updates received from the neighbor. If you request a soft reset of inbound routes, the software performs the reset by comparing the policies against the stored route updates, instead of requesting the neighbor BGP4 route table or resetting the session with the neighbor. Refer to [“Using soft reconfiguration”](#) on page 845.

timers keep-alive <num> **hold-time** <num> overrides the global settings for the Keep Alive Time and Hold Time. For the Keep Alive Time, you can specify from 0 – 65535 seconds. For the Hold Time, you can specify 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the router waits indefinitely for

messages from a neighbor without concluding that the neighbor is dead. The defaults for these parameters are the currently configured global Keep Alive Time and Hold Time. For more information about these parameters, refer to [“Changing the Keep Alive Time and Hold Time”](#) on page 767.

unsuppress-map <map-name> removes route dampening from a neighbor routes when those routes have been dampened due to aggregation. Refer to [“Removing route dampening from a neighbor routes suppressed due to aggregation”](#) on page 812.

update-source <ip-addr> | **ethernet** <portnum> | **loopback** <num> | **ve** <num> configures the router to communicate with the neighbor through the specified interface. There is no default.

weight <num> specifies a weight the Layer 3 Switch will add to routes received from the specified neighbor. BGP4 prefers larger weights over smaller weights. The default weight is 0.

Encryption of BGP4 MD5 authentication keys

When you configure a BGP4 neighbor or neighbor peer group, you can specify an MD5 authentication string for authenticating packets exchanged with the neighbor or peer group of neighbors.

For added security, the software encrypts display of the authentication string by default. The software also provides an optional parameter to disable encryption of the authentication string, on an individual neighbor or peer group basis. By default, the MD5 authentication strings are displayed in encrypted format in the output of the following commands:

- **show running-config** (or **write terminal**)
- **show configuration**
- **show ip bgp config**

When encryption of the authentication string is enabled, the string is encrypted in the CLI regardless of the access level you are using.

If you display the running-config after reloading, the BGP4 commands that specify an authentication string show the string in encrypted form.

In addition, when you save the configuration to the startup-config file, the file contains the new BGP4 command syntax and encrypted passwords or strings.

NOTE

Dell recommends that you save a copy of the startup-config file for each switch you plan to upgrade.

Encryption example

The following commands configure a BGP4 neighbor and a peer group, and specify MD5 authentication strings (passwords) for authenticating packets exchanged with the neighbor or peer group.

```
PowerConnect(config-bgp-router)# local-as 2
PowerConnect(config-bgp-router)# neighbor xyz peer-group
PowerConnect(config-bgp-router)# neighbor xyz password abc
PowerConnect(config-bgp-router)# neighbor 10.10.200.102 peer-group xyz
PowerConnect(config-bgp-router)# neighbor 10.10.200.102 password test
```

Here is how the commands appear when you display the BGP4 configuration commands.

```
PowerConnect# show ip bgp config
Current BGP configuration:
router bgp
  local-as 2
  neighbor xyz peer-group
  neighbor xyz password 1 $!2d
  neighbor 10.10.200.102 peer-group xyz
  neighbor 10.10.200.102 remote-as 1
  neighbor 10.10.200.102 password 1 $on-o
```

Notice that the software has converted the commands that specify an authentication string into the new syntax (described below), and has encrypted display of the authentication strings.

Command syntax

Since the default behavior does not affect the BGP4 configuration itself but does encrypt display of the authentication string, the CLI does not list the encryption options.

Syntax: [no] **neighbor** <ip-addr> | <peer-group-name> **password** [0 | 1] <string>

The <ip-addr> | <peer-group-name> parameter indicates whether you are configuring an individual neighbor or a peer group. If you specify a neighbor IP address, you are configuring that individual neighbor. If you specify a peer group name, you are configuring a peer group.

The **password** <string> parameter specifies an MD5 authentication string for securing sessions between the Layer 3 Switch and the neighbor. You can enter a string up to 80 characters long. The string can contain any alphanumeric characters, but the first character cannot be a number. If the password contains a number, do not enter a space following the number.

The **0** | **1** parameter is the encryption option, which you can omit (the default) or which can be one of the following:

- **0** – Disables encryption for the authentication string you specify with the command. The password or string is shown as clear text in the output of commands that display neighbor or peer group configuration information.
- **1** – Assumes that the authentication string you enter is the encrypted form, and decrypts the value before using it.

NOTE

If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter 0 or 1. Instead, omit the encryption option and allow the software to use the default behavior.

If you specify encryption option 1, the software assumes that you are entering the encrypted form of the password or authentication string. In this case, the software decrypts the password or string you enter before using the value for authentication. If you accidentally enter option 1 followed by the clear-text version of the password or string, authentication will fail because the value used by the software will not match the value you intended to use.

Displaying the Authentication String

If you want to display the authentication string, enter the following commands.

```
PowerConnect(config)# enable password-display
PowerConnect# show ip bgp neighbors
```


The **enable password-display** command enables display of the authentication string, but only in the output of the **show ip bgp neighbors** command. Display of the string is still encrypted in the startup-config file and running-config. Enter the command at the global CONFIG level of the CLI.

NOTE

The command also displays SNMP community strings in clear text, in the output of the **show snmp server** command.

Adding a BGP4 peer group

A **peer group** is a set of BGP4 neighbors that share common parameters. Peer groups provide the following benefits:

- **Simplified neighbor configuration** – You can configure a set of neighbor parameters and then apply them to multiple neighbors. You do not need to individually configure the common parameters individually on each neighbor.
- **Flash memory conservation** – Using peer groups instead of individually configuring all the parameters for each neighbor requires fewer configuration commands in the startup-config file.

You can perform the following tasks on a peer-group basis:

- Reset neighbor sessions
- Perform soft-outbound resets (the Layer 3 Switch updates outgoing route information to neighbors but does not entirely reset the sessions with those neighbors)
- Clear BGP message statistics
- Clear error buffers

Peer group parameters

You can set all neighbor parameters in a peer group. When you add a neighbor to the peer group, the neighbor receives all the parameter settings you set in the group, except parameter values you have explicitly configured for the neighbor. If you do not set a neighbor parameter in the peer group and the parameter also is not set for the individual neighbor, the neighbor uses the default value.

Configuration rules

The following rules apply to peer group configuration:

- You must configure a peer group before you can add neighbors to the peer group.
- If you remove a parameter from a peer group, the value for that parameter is reset to the default for all the neighbors within the peer group, unless you have explicitly set that parameter on individual neighbors. In this case, the value you set on the individual neighbors applies to those neighbors, while the default value applies to neighbors for which you have not explicitly set the value.

NOTE

If you enter a command to remove the remote AS parameter from a peer group, the software checks to ensure that the peer group does not contain any neighbors. If the peer group does contain neighbors, the software does not allow you to remove the remote AS. The software prevents removing the remote AS in this case so that the neighbors in the peer group that are using the remote AS do not lose connectivity to the Layer 3 Switch.

- Once you add a neighbor to a peer group, you cannot configure the following outbound parameters (the parameters governing outbound traffic) for the neighbor:
 - Default-information-originate
 - Next-hop-self
 - Outbound route map
 - Outbound filter list
 - Outbound distribute list
 - Outbound prefix list
 - Remote AS, if configured for the peer group
 - Remove private AS
 - Route reflector client
 - Send community
 - Timers
 - Update source

If you want to change an outbound parameter for an individual neighbor, you must first remove the neighbor from the peer group. In this case, you cannot re-add the neighbor to the same peer group, but you can add the neighbor to a different peer group. All the neighbors within a peer group must have the same values for the outbound parameters. To change an outbound parameter to the same value for all neighbors within a peer group, you can change the parameter on a peer-group basis. In this case, you do not need to remove the neighbors and change the parameter individually for each neighbor.

- If you add an outbound parameter to a peer group, that parameter is automatically applied to all neighbors within the peer group.
- When you add a neighbor to a peer group, the software removes any outbound parameters for that neighbor from the running configuration (running-config). As a result, when you save the configuration to the startup-config file, the file does not contain any outbound parameters for the individual neighbors you have placed in a peer group. The only outbound parameters the startup-config file contains for neighbors within a peer group are the parameters associated with the peer group itself. However, the running-config and the startup-config file can contain individual parameters listed in the previous section as well as the settings for those parameters within a peer group.

You can override neighbor parameters that do not affect outbound policy on an individual neighbor basis.

- If you do not specify a parameter for an individual neighbor, the neighbor uses the value in the peer group.
- If you set the parameter for the individual neighbor, that value overrides the value you set in the peer group.

- If you add a parameter to a peer group that already contains neighbors, the parameter value is applied to neighbors that do not already have the parameter explicitly set. If a neighbor has the parameter explicitly set, the explicitly set value overrides the value you set for the peer group.
- If you remove the setting for a parameter from a peer group, the value for that parameter changes to the default value for all the neighbors in the peer group that do not have that parameter individually set.

Configuring a peer group

To configure a BGP4 peer group, enter commands such as the following at the BGP configuration level.

```
PowerConnect(config-bgp-router)# neighbor PeerGroup1 peer-group
PowerConnect(config-bgp-router)# neighbor PeerGroup1 description "EastCoast
Neighbors"
PowerConnect(config-bgp-router)# neighbor PeerGroup1 remote-as 100
PowerConnect(config-bgp-router)# neighbor PeerGroup1 distribute-list out 1
```

The commands in this example configure a peer group called "PeerGroup1" and set the following parameters for the peer group:

- A description, "EastCoast Neighbors"
- A remote AS number, 100
- A distribute list for outbound traffic

The software applies these parameters to each neighbor you add to the peer group. You can override the description parameter for individual neighbors. If you set the description parameter for an individual neighbor, the description overrides the description configured for the peer group. However, you cannot override the remote AS and distribute list parameters for individual neighbors. Since these parameters control outbound traffic, the parameters must have the same values for all neighbors within the peer group.

Syntax: `neighbor <peer-group-name> peer-group`

The `<peer-group-name>` parameter specifies the name of the group and can be up to 80 characters long. The name can contain special characters and internal blanks. If you use internal blanks, you must use quotation marks around the name. For example, the command `neighbor "My Three Peers" peer-group` is valid, but the command `neighbor My Three Peers peer-group` is not valid.

Syntax: `[no] neighbor <ip-addr> | <peer-group-name>`
`[advertisement-interval <num>]`
`[default-originate [route-map <map-name>]]`
`[description <string>]`
`[distribute-list in | out <num,num,...> | <ACL-num> in | out]`
`[ebgp-multihop [<num>]]`
`[filter-list in | out <num,num,...> | <ACL-num> in | out | weight]`
`[maximum-prefix <num> [<threshold>] [teardown]]`
`[next-hop-self]`
`[password [0 | 1] <string>]`
`[prefix-list <string> in | out]`
`[remote-as <as-number>]`
`[remove-private-as]`
`[route-map in | out <map-name>]`
`[route-reflector-client]`

```

[send-community]
[soft-reconfiguration inbound]
[shutdown]
[timers keep-alive <num> hold-time <num>]
[update-source loopback <num>]
[weight <num>]

```

The `<ip-addr>` | `<peer-group-name>` parameter indicates whether you are configuring a peer group or an individual neighbor. You can specify a peer group name or IP address with the **neighbor** command. If you specify a peer group name, you are configuring a peer group. If you specify a neighbor IP address, you are configuring that individual neighbor. Use the `<ip-addr>` parameter if you are configuring an individual neighbor instead of a peer group. Refer to [“Adding BGP4 neighbors”](#) on page 756.

The remaining parameters are the same ones supported for individual neighbors. Refer to [“Adding BGP4 neighbors”](#) on page 756.

Applying a peer group to a neighbor

After you configure a peer group, you can add neighbors to the group. When you add a neighbor to a peer group, you are applying all the neighbor attributes specified in the peer group to the neighbor.

To add neighbors to a peer group, enter commands such as the following.

```

PowerConnect(config-bgp-router)# neighbor 192.168.1.12 peer-group PeerGroup1
PowerConnect(config-bgp-router)# neighbor 192.168.2.45 peer-group PeerGroup1
PowerConnect(config-bgp-router)# neighbor 192.168.3.69 peer-group PeerGroup1

```

The commands in this example add three neighbors to the peer group “PeerGroup1”. As members of the peer group, the neighbors automatically receive the neighbor parameter values configured for the peer group. You also can override the parameters (except parameters that govern outbound traffic) on an individual neighbor basis. For neighbor parameters not specified for the peer group, the neighbors use the default values.

Syntax: `neighbor <ip-addr> peer-group <peer-group-name>`

The `<ip-addr>` parameter specifies the IP address of the neighbor.

The `<peer-group-name>` parameter specifies the peer group name.

NOTE

You must add the peer group before you can add neighbors to it.

Administratively shutting down a session with a BGP4 neighbor

You can prevent the Layer 3 Switch from starting a BGP4 session with a neighbor by administratively shutting down the neighbor. This option is very useful for situations in which you want to configure parameters for a neighbor but are not ready to use the neighbor. You can shut the neighbor down as soon as you have added it the Layer 3 Switch, configure the neighbor parameters, then allow the Layer 3 Switch to re-establish a session with the neighbor by removing the shutdown option from the neighbor.

When you apply the new option to shut down a neighbor, the option takes place immediately and remains in effect until you remove the option. If you save the configuration to the startup-config file, the shutdown option remains in effect even after a software reload.

NOTE

The software also contains an option to end the session with a BGP4 neighbor and thus clear the routes learned from the neighbor. Unlike this clear option, the option for shutting down the neighbor can be saved in the startup-config file and thus can prevent the Layer 3 Switch from establishing a BGP4 session with the neighbor even after reloading the software.

NOTE

If you notice that a particular BGP4 neighbor never establishes a session with the Layer 3 Switch, check the Layer 3 Switch running-config and startup-config files to see whether the configuration contains a command that is shutting down the neighbor. The neighbor may have been shut down previously by an administrator.

To shut down a BGP4 neighbor, enter commands such as the following.

```
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# neighbor 209.157.22.26 shutdown
PowerConnect(config-bgp-router)# write memory
```

Syntax: [no] neighbor <ip-addr> shutdown

The <ip-addr> parameter specifies the IP address of the neighbor.

Optional configuration tasks

The following sections describe how to perform optional BGP4 configuration tasks.

Changing the Keep Alive Time and Hold Time

The Keep Alive Time specifies how frequently the router will send KEEPALIVE messages to its BGP4 neighbors. The Hold Time specifies how long the router will wait for a KEEPALIVE or UPDATE message from a neighbor before concluding that the neighbor is dead. When the router concludes that a BGP4 neighbor is dead, the router ends the BGP4 session and closes the TCP connection to the neighbor.

The default Keep Alive time is 60 seconds. The default Hold Time is 180 seconds. To change the timers, use either of the following methods.

NOTE

Generally, you should set the Hold Time to three times the value of the Keep Alive Time.

NOTE

You can override the global Keep Alive Time and Hold Time on individual neighbors. Refer to [“Adding BGP4 neighbors”](#) on page 756.

To change the Keep Alive Time to 30 and Hold Time to 90, enter the following command.

```
PowerConnect(config-bgp-router)# timers keep-alive 30 hold-time 90
```

Syntax: timers keep-alive <num> hold-time <num>

For each keyword, *<num>* indicates the number of seconds. The Keep Alive Time can be 0 – 65535. The Hold Time can be 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the router waits indefinitely for messages from a neighbor without concluding that the neighbor is dead.

Changing the BGP4 next-hop update timer

By default, the Layer 3 Switch updates its BGP4 next-hop tables and affected BGP4 routes five seconds after IGP route changes. You can change the update timer to a value from 1 – 30 seconds.

To change the BGP4 update timer value, enter a command such as the following at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# update-time 15
```

This command changes the update timer to 15 seconds.

Syntax: [no] **update-time** <secs>

The <secs> parameter specifies the number of seconds and can be from 1 – 30. The default is 5.

Enabling fast external fallover

BGP4 routers rely on KEEPALIVE and UPDATE messages from neighbors to signify that the neighbors are alive. For BGP4 neighbors that are two or more hops away, such messages are the only indication that the BGP4 protocol has concerning the alive state of the neighbors. As a result, if a neighbor dies, the router will wait until the Hold Time expires before concluding that the neighbor is dead and closing its BGP4 session and TCP connection with the neighbor.

The router waits for the Hold Time to expire before ending the connection to a directly-attached BGP4 neighbor that dies.

For directly attached neighbors, the router immediately senses loss of a connection to the neighbor from a change of state of the port or interface that connects the router to its neighbor. For directly attached EBGP neighbors, the router can use this information to immediately close the BGP4 session and TCP connection to locally attached neighbors that die.

NOTE

The fast external fallover feature applies only to directly attached EBGP neighbors. The feature does not apply to IBGP neighbors.

If you want to enable the router to immediately close the BGP4 session and TCP connection to locally attached neighbors that die, use either of the following methods.

To enable fast external fallover, enter the following command.

```
PowerConnect(config-bgp-router)# fast-external-fallover
```

To disable fast external fallover again, enter the following command.

```
PowerConnect(config-bgp-router)# no fast-external-fallover
```

Syntax: [no] **fast-external-fallover**

Changing the maximum number of paths for BGP4 load sharing

Load sharing enables the Layer 3 Switch to balance traffic to a route across multiple equal-cost paths of the same type (EBGP or IBGP) for the route.

To configure the Layer 3 Switch to perform BGP4 load sharing:

- Enable IP load sharing if it is disabled.
- Set the maximum number of paths. The default maximum number of BGP4 load sharing paths is 1, which means no BGP4 load sharing takes place by default.

NOTE

The maximum number of BGP4 load sharing paths cannot be greater than the maximum number of IP load sharing paths.

How load sharing affects route selection

During evaluation of multiple paths to select the best path to a given destination for installment in the IP route table, the last comparison the Layer 3 Switch performs is a comparison of the internal paths:

- When IP load sharing is disabled, the Layer 3 Switch prefers the path to the router with the lower router ID.
- When IP load sharing and BGP4 load sharing are enabled, the Layer 3 Switch balances the traffic across the multiple paths instead of choosing just one path based on router ID.

Refer to [“How BGP4 selects a path for a route”](#) on page 747 for a description of the BGP4 algorithm.

When you enable IP load sharing, the Layer 3 Switch can load balance BGP4 or OSPF routes across up to four equal paths by default. You can change the number of IP load sharing paths to a value from 2 – 6.

How load sharing works

Load sharing is performed in round-robin fashion and is based on the destination IP address only. The first time the router receives a packet destined for a specific IP address, the router uses a round-robin algorithm to select the path that was not used for the last newly learned destination IP address. Once the router associates a path with a particular destination IP address, the router will always use that path as long as the router contains the destination IP address in its cache.

NOTE

The Layer 3 Switch does not perform source routing. The router is concerned only with the paths to the next-hop routers, not the entire paths to the destination hosts.

A BGP4 destination can be learned from multiple BGP4 neighbors, leading to multiple BGP4 paths to reach the same destination. Each of the paths may be reachable through multiple IGP paths (multiple OSPF or RIP paths). In this case, the software installs all the multiple equal-cost paths in the BGP4 route table, up to the maximum number of BGP4 equal-cost paths allowed. The IP load sharing feature then distributes traffic across the equal-cost paths to the destination.

If an IGP path used by a BGP4 next-hop route path installed in the IP route table changes, then the BGP4 paths and IP paths are adjusted accordingly. For example, if one of the OSPF paths to reach the BGP4 next hop goes down, the software removes this path from the BGP4 route table and the IP route table. Similarly, if an additional OSPF path becomes available to reach the BGP4 next-hop router for a particular destination, the software adds the additional path to the BGP4 route table and the IP route table.

Changing the maximum number of shared BGP4 paths

When IP load sharing is enabled, BGP4 can balance traffic to a specific destination across up to four equal paths. You can set the maximum number of paths to a value from 1 – 4. The default is 1.

NOTE

The maximum number of BGP4 load sharing paths cannot be greater than the maximum number of IP load sharing paths. To increase the maximum number of IP load sharing paths, use the **ip load sharing <num>** command at the global CONFIG level of the CLI.

To change the maximum number of shared paths, enter commands such as the following.

```
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# maximum-paths 4
PowerConnect(config-bgp-router)# write memory
```

Syntax: [no] maximum-paths <num>

The <num> parameter specifies the maximum number of paths across which the Layer 3 Switch can balance traffic to a given BGP4 destination. You can change the maximum number of paths to a value from 2 – 4. The default is 1.

Customizing BGP4 load sharing

By default, when BGP4 load sharing is enabled, both IBGP and EBGP paths are eligible for load sharing, while paths from different neighboring ASs are not eligible. You can change load sharing to apply only to IBGP or EBGP paths, or to support load sharing among paths from different neighboring ASs.

To enable load sharing of IBGP paths only, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# multipath ibgp
```

To enable load sharing of EBGP paths only, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# multipath ebgp
```

To enable load sharing of paths from different neighboring ASs, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# multipath multi-as
```

Syntax: [no] multipath ebgp | ibgp | multi-as

The **ebgp** | **ibgp** | **multi-as** parameter specifies the change you are making to load sharing:

- **ebgp** – Load sharing applies only to EBGP paths. Load sharing is disabled for IBGP paths.
- **ibgp** – Load sharing applies only to IBGP paths. Load sharing is disabled for EBGP paths.

- **multi-as** – Load sharing is enabled for paths from different ASs.

By default, load sharing applies to EBGP and IBGP paths, and does not apply to paths from different neighboring ASs.

Specifying a list of networks to advertise

By default, the router sends BGP4 routes only for the networks you identify using the **network** command or that are redistributed into BGP4 from RIP or OSPF. You can specify up to 600 networks.

To specify a network to be advertised, use either of the following methods.

NOTE

The exact route must exist in the IP route table before the Layer 3 Switch can create a local BGP route.

To configure the Layer 3 Switch to advertise network 209.157.22.0/24, enter the following command.

```
PowerConnect(config-bgp-router)# network 209.157.22.0 255.255.255.0
```

Syntax: **network** <ip-addr> <ip-mask> [**nlri multicast** | **unicast** | **multicast unicast**] [**route-map** <map-name>] | [**weight** <num>] | [**backdoor**]

The <ip-addr> is the network number and the <ip-mask> specifies the network mask.

The **nlri multicast** | **unicast** | **multicast unicast** parameter specifies whether the neighbor is a multicast neighbor or a unicast neighbor. Optionally, you also can specify **unicast** if you want the Layer 3 Switch to exchange unicast (BGP4) routes as well as multicast routes with the neighbor. The default is **unicast** only.

The **route-map** <map-name> parameter specifies the name of the route map you want to use to set or change BGP4 attributes for the network you are advertising. The route map must already be configured.

The **weight** <num> parameter specifies a weight to be added to routes to this network.

The **backdoor** parameter changes the administrative distance of the route to this network from the EBGP administrative distance (20 by default) to the Local BGP weight (200 by default), thus tagging the route as a backdoor route. Use this parameter when you want the router to prefer IGP routes such as RIP or OSPF routes over the EBGP route for the network.

Specifying a route map name when configuring BGP4 network information

You can specify a route map as one of the parameters when you configure a BGP4 network to be advertised. The Layer 3 Switch can use the route map to set or change BGP4 attributes when creating a local BGP4 route.

To configure network information and use a route map to set or change BGP4 attributes, use the following CLI method.

NOTE

You must configure the route map before you can specify the route map name in a BGP4 network configuration.

To configure a route map, and use it to set or change route attributes for a network you define for BGP4 to advertise, enter commands such as the following.

```
PowerConnect(config)# route-map set_net permit 1
PowerConnect(config-route-map set_net)# set community no-export
PowerConnect(config-route-map set_net)# exit
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# network 100.100.1.0/24 route-map set_net
```

The first two commands in this example create a route map named “set_net” that sets the community attribute for routes that use the route map to “NO_EXPORT”. The next two commands change the CLI to the BGP4 configuration level. The last command configures a network for advertising from BGP4, and associates the “set_net” route map with the network. When BGP4 originates the 100.100.1.0/24 network, BGP4 also sets the community attribute for the network to “NO_EXPORT”.

Syntax: `network <ip-addr> <ip-mask> [route-map <map-name>] | [weight <num>] | [backdoor]`

The `route-map <map-name>` parameter specifies the name of the route map you want to use to set or change BGP4 attributes for the network you are advertising. The route map must already be configured.

For information about the other parameters, refer to [“Defining route maps”](#) on page 797.

Changing the default local preference

When the router uses the BGP4 algorithm to select a route to send to the IP route table, one of the parameters the algorithm uses is the local preference. Local preference is an attribute that indicates a degree of preference for a route relative to other routes. BGP4 neighbors can send the local preference value as an attribute of a route in an UPDATE message.

Local preference applies only to routes within the local AS. BGP4 routers can exchange local preference information with neighbors who also are in the local AS, but BGP4 routers do not exchange local preference information with neighbors in remote ASs.

The default local preference is 100. For routes learned from EBGp neighbors, the default local preference is assigned to learned routes. For routes learned from IBGP neighbors, the local preference value is not changed for the route.

When the BGP4 algorithm compares routes on the basis of local preferences, the route with the higher local preference is chosen.

NOTE

To set the local preference for individual routes, use route maps. Refer to [“Defining route maps”](#) on page 797. Refer to [“How BGP4 selects a path for a route”](#) on page 747 for information about the BGP4 algorithm.

To change the default local preference to 200, enter the following command.

```
PowerConnect(config-bgp-router)# default-local-preference 200
```

Syntax: `default-local-preference <num>`

The `<num>` parameter indicates the preference and can be a value from 0 – 4294967295.

Using the IP default route as a valid next hop for a BGP4 route

By default, the Layer 3 Switch does not use a default route to resolve a BGP4 next-hop route. If the IP route lookup for the BGP4 next hop does not result in a valid IGP route (including static or direct routes), the BGP4 next hop is considered to be unreachable and the BGP4 route is not used.

In some cases, such as when the Layer 3 Switch is acting as an edge router, you might want to allow the device to use the default route as a valid next hop. To do so, enter the following command at the BGP4 configuration level of the CLI.

```
PowerConnect(config-bgp-router)# next-hop-enable-default
```

Syntax: [no] next-hop-enable-default

Advertising the default route

By default, the Layer 3 Switch does not originate and advertise a default route using BGP4. A BGP4 default route is the IP address 0.0.0.0 and the route prefix 0 or network mask 0.0.0.0. For example, 0.0.0.0/0 is a default route. You can enable the router to advertise a default BGP4 route using either of the following methods.

NOTE

The Layer 3 Switch checks for the existence of an IGP route for 0.0.0.0/0 in the IP route table before creating a local BGP route for 0.0.0.0/0.

To enable the router to originate and advertise a default BGP4 route, enter the following command.

```
PowerConnect(config-bgp-router)# default-information-originate
```

Syntax: [no] default-information-originate

Changing the default MED (Metric) used for route redistribution

The Layer 3 Switch can redistribute directly connected routes, static IP routes, RIP routes, and OSPF routes into BGP4. The MED (metric) is a global parameter that specifies the cost that will be applied to all routes by default when they are redistributed into BGP4. When routes are selected, lower metric values are preferred over higher metric values. The default BGP4 MED value is 0 and can be assigned a value from 0 – 4294967295.

NOTE

RIP and OSPF also have default metric parameters. The parameters are set independently for each protocol and have different ranges.

To change the default metric to 40, enter the following command.

```
PowerConnect(config-bgp-router)# default-metric 40
```

Syntax: default-metric <num>

The <num> indicates the metric and can be a value from 0 – 4294967295.

Enabling next-hop recursion

For each BGP4 route a Layer 3 Switch learns, the Layer 3 Switch performs a route lookup to obtain the IP address of the route next hop. A BGP4 route becomes eligible for installation into the IP route table only if the following conditions are true:

- The lookup succeeds in obtaining a valid next-hop IP address for the route.
- The path to the next-hop IP address is an Interior Gateway Protocol (IGP) path or a static route path.

By default, the software performs only one lookup for a BGP route next-hop IP address. If the next-hop lookup does not result in a valid next-hop IP address or the path to the next-hop IP address is a BGP path, the software considers the BGP route destination to be unreachable. The route is not eligible to be installed in the IP route table.

It is possible for the BGP route table to contain a route whose next-hop IP address is not reachable through an IGP route, even though a hop farther away can be reached by the Layer 3 Switch through an IGP route. This can occur when the IGP does not learn a complete set of IGP routes, resulting in the Layer 3 Switch learning about an internal route through IBGP instead of through an IGP. In this case, the IP route table does not contain a route that can be used to reach the BGP route destination.

To enable the Layer 3 Switch to find the IGP route to a BGP route next-hop gateway, enable recursive next-hop lookups. When you enable recursive next-hop lookup, if the first lookup for a BGP route results in an IBGP path originated within the same Autonomous System (AS), rather than an IGP path or static route path, the Layer 3 Switch performs a lookup on the next-hop gateway next-hop IP address. If this second lookup results in an IGP path, the software considers the BGP route to be valid and thus eligible for installation in the IP route table. Otherwise, the Layer 3 Switch performs a lookup on the next-hop IP address of the next-hop gateway next hop, and so on, until one of the lookups results in an IGP route.

NOTE

The software does not support using the default route to resolve a BGP4 route's next hop. Instead, you must configure a static route or use an IGP to learn the route to the EBGp multihop peer.

Previous software releases support use of the default route to resolve routes learned from EBGp multihop neighbors. However, even in this case Dell recommends that you use a static route for the EBGp multihop neighbor instead. In general, we recommend that you do not use the default route as the next hop for BGP4 routes, especially when there are two or more BGP4 neighbors. Using the default route can cause loops.

Example when recursive route lookups are disabled

Here is an example of the results of an unsuccessful next-hop lookup for a BGP route. In this case, next-hop recursive lookups are disabled. The example is for the BGP route to network 240.0.0.0/24.

```
PowerConnect# show ip bgp route
Total number of BGP Routes: 5
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
Prefix          Next Hop          Metric      LocPrf      Weight  Status
1      0.0.0.0/0        10.1.0.2        0           100         0      BI
   AS_PATH: 65001 4355 701 80
2      102.0.0.0/24    10.0.0.1         1           100         0      BI
   AS_PATH: 65001 4355 1
3      104.0.0.0/24    10.1.0.2        0           100         0      BI
   AS_PATH: 65001 4355 701 1 189
4      240.0.0.0/24    102.0.0.1      1          100        0      I
   AS_PATH: 65001 4355 3356 7170 1455
5      250.0.0.0/24    209.157.24.1    1           100         0      I
   AS_PATH: 65001 4355 701
```

In this example, the Layer 3 Switch cannot reach 240.0.0.0/24, because the next-hop IP address for the route is an IBGP route instead of an IGP route, and thus is considered unreachable by the Layer 3 Switch. Here is the IP route table entry for the BGP route next-hop gateway (102.0.0.1/24).

```
PowerConnect# show ip route 102.0.0.1
Total number of IP routes: 37
Network Address  NetMask          Gateway          Port    Cost  Type
102.0.0.0        255.255.255.0    10.0.0.1        1       1     B
```

The route to the next-hop gateway is a BGP route, not an IGP route, and thus cannot be used to reach 240.0.0.0/24. In this case, the Layer 3 Switch tries to use the default route, if present, to reach the subnet that contains the BGP route next-hop gateway.

```
PowerConnect# show ip route 240.0.0.0/24
Total number of IP routes: 37
Network Address  NetMask          Gateway          Port    Cost  Type
0.0.0.0          0.0.0.0          10.0.0.202      1       1     S
```

Example when recursive route lookups are enabled

When recursive next-hop lookups are enabled, the Layer 3 Switch recursively looks up the next-hop gateways along the route until the Layer 3 Switch finds an IGP route to the BGP route destination. Here is an example.

```
PowerConnect# show ip bgp route
Total number of BGP Routes: 5
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
Prefix          Next Hop          Metric      LocPrf      Weight Status
1  0.0.0.0/0      10.1.0.2         0           100         0      BI
   AS_PATH: 65001 4355 701 80
2  102.0.0.0/24   10.0.0.1         1           100         0      BI
   AS_PATH: 65001 4355 1
3  104.0.0.0/24   10.1.0.2         0           100         0      BI
   AS_PATH: 65001 4355 701 1 189
4  240.0.0.0/24 102.0.0.1       1          100        0      BI
   AS_PATH: 65001 4355 3356 7170 1455
5  250.0.0.0/24   209.157.24.1    1           100         0      I
   AS_PATH: 65001 4355 701
```

The first lookup results in an IBGP route, to network 102.0.0.0/24.

```
PowerConnect# show ip route 102.0.0.1
Total number of IP routes: 38
Network Address  NetMask          Gateway          Port    Cost  Type
102.0.0.0      255.255.255.0  10.0.0.1      1     1    B
   AS_PATH: 65001 4355 1
```

Since the route to 102.0.0.1/24 is not an IGP route, the Layer 3 Switch cannot reach the next hop through IP, and thus cannot use the BGP route. In this case, since recursive next-hop lookups are enabled, the Layer 3 Switch next performs a lookup for 102.0.0.1 next-hop gateway, 10.0.0.1.

```
PowerConnect# show ip bgp route 102.0.0.0
Number of BGP Routes matching display condition : 1
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
Prefix          Next Hop          Metric      LocPrf      Weight Status
1  102.0.0.0/24 10.0.0.1       1          100        0      BI
   AS_PATH: 65001 4355 1
```

The next-hop IP address for 102.0.0.1 is not an IGP route, which means the BGP route destination still cannot be reached through IP. The recursive next-hop lookup feature performs a lookup on 10.0.0.1 next-hop gateway.

```
PowerConnect# show ip route 10.0.0.1
Total number of IP routes: 38
Network Address  NetMask          Gateway          Port    Cost  Type
10.0.0.0      255.255.255.0  0.0.0.0      1     1    D
   AS_PATH: 65001 4355 1
```

This lookup results in an IGP route. In fact, this route is a directly-connected route. As a result, the BGP route destination is now reachable through IGP, which means the BGP route is eligible for installation in the IP route table. Here is the BGP route in the IP route table.

```
PowerConnect# show ip route 240.0.0.0/24
Total number of IP routes: 38
Network Address  NetMask          Gateway          Port    Cost  Type
240.0.0.0      255.255.255.0  10.0.0.1      1     1    B
   AS_PATH: 65001 4355 1
```

This Layer 3 Switch can use this route because the Layer 3 Switch has an IP route to the next-hop gateway. Without recursive next-hop lookups, this route would not be in the IP route table.

Enabling recursive next-hop lookups

The recursive next-hop lookups feature is disabled by default. To enable recursive next-hop lookups, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# next-hop-recursion
```

Syntax: [no] next-hop-recursion

Changing administrative distances

BGP4 routers can learn about networks from various protocols, including the EBGp portion of BGP4 and IGP's such as OSPF and RIP. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned.

To select one route over another based on the source of the route information, the Layer 3 Switch can use the administrative distances assigned to the sources. The administrative distance is a protocol-independent metric that IP routers use to compare routes from different sources.

The Layer 3 Switch re-advertises a learned best BGP4 route to the Layer 3 Switch neighbors even when the software does not also select that route for installation in the IP route table. The best BGP4 routes is the BGP4 path that the software selects based on comparison of the paths' BGP4 route parameters. Refer to "[How BGP4 selects a path for a route](#)" on page 747.

When selecting a route from among different sources (BGP4, OSPF, RIP, static routes, and so on), the software compares the routes on the basis of each route administrative distance. If the administrative distance of the paths is lower than the administrative distance of paths from other sources (such as static IP routes, RIP, or OSPF), the BGP4 paths are installed in the IP route table.

NOTE

The software will replace a statically configured default route with a learned default route if the learned route administrative distance is lower than the statically configured default route distance. However, the default administrative distance for static routes is changed to 1, so only directly-connected routes are preferred over static routes when the default administrative distances for the routes are used.

The following default administrative distances are found on the Layer 3 Switch:

- Directly connected – 0 (this value is not configurable)
- Static – 1 (applies to all static routes, including default routes)
- EBGp – 20
- OSPF – 110
- RIP – 120
- IBGP – 200
- Local BGP – 200
- Unknown – 255 (the router will not use this route)

Lower administrative distances are preferred over higher distances. For example, if the router receives routes for the same network from OSPF and from RIP, the router will prefer the OSPF route by default. The administrative distances are configured in different places in the software. The Layer 3 Switch re-advertises a learned best BGP4 route to neighbors by default, regardless of whether the route administrative distance is lower than other routes from different route sources to the same destination.

- To change the EBGp, IBGP, and Local BGP default administrative distances, see the instructions in this section.
- To change the default administrative distance for OSPF, refer to “[Modify administrative distance](#)” on page 692.
- To change the default administrative distance for RIP, refer to “[Changing the administrative distance](#)” on page 647.
- To change the default administrative distance for static routes, refer to “[Configuring static routes](#)” on page 596.

You can change the default EBGp, IBGP, and Local BGP administrative distances using either of the following methods.

To change the default administrative distances for EBGp, IBGP, and Local BGP, enter a command such as the following.

```
PowerConnect(config-bgp-router)# distance 180 160 40
```

Syntax: `distance <external-distance> <internal-distance> <local-distance>`

The `<external-distance>` sets the EBGp distance and can be a value from 1 – 255.

The `<internal-distance>` sets the IBGP distance and can be a value from 1 – 255.

The `<local-distance>` sets the Local BGP distance and can be a value from 1 – 255.

Requiring the first AS to be the neighbor AS

By default, the device does not require the first AS listed in the AS_SEQUENCE field of an AS path Update from an EBGp neighbor to be the AS that the neighbor who sent the Update is in. You can enable the device for this requirement.

When you enable the device to require the AS that an EBGp neighbor is in to be the same as the first AS in the AS_SEQUENCE field of an Update from the neighbor, the device accepts the Update only if the ASs match. If the ASs do not match, the device sends a Notification message to the neighbor and closes the session. The requirement applies to all Updates received from EBGp neighbors.

To enable this feature, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# enforce-first-as
```

Syntax: `[no] enforce-first-as`

Disabling or re-enabling comparison of the AS-Path length

AS-Path comparison is Step 5 in the algorithm BGP4 uses to select the next path for a route. Comparison of the AS-Path length is enabled by default. To disable it, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# as-path-ignore
```


This command disables comparison of the AS-Path lengths of otherwise equal paths. When you disable AS-Path length comparison, the BGP4 algorithm shown in [“How BGP4 selects a path for a route”](#) on page 747 skips from Step 4 to Step 6.

Syntax: [no] as-path-ignore

Enabling or disabling comparison of the router IDs

Router ID comparison is Step 10 in the algorithm BGP4 uses to select the next path for a route.

NOTE

Comparison of router IDs is applicable only when BGP4 load sharing is disabled.

When router ID comparison is enabled, the path comparison algorithm compares the router IDs of the neighbors that sent the otherwise equal paths:

- If BGP4 load sharing is disabled (maximum-paths 1), the Layer 3 Switch selects the path that came from the neighbor with the lower router ID.
- If BGP4 load sharing is enabled, the Layer 3 Switch load shares among the remaining paths. In this case, the router ID is not used to select a path.

NOTE

Router ID comparison is disabled by default. In previous releases, router ID comparison is enabled by default and cannot be disabled.

To enable router ID comparison, enter the following command at the BGP configuration level of the CLI.

```
PowerConnect(config-bgp-router)# compare-routerid
```

Syntax: [no] compare-routerid

For more information, refer to [“How BGP4 selects a path for a route”](#) on page 747.

Configuring the Layer 3 Switch to always compare Multi-Exit Discriminators (MEDs)

A Multi-Exit Discriminator (MED) is a value that the BGP4 algorithm uses when comparing multiple paths received from different BGP4 neighbors in the same AS for the same route. In BGP4, a route MED is equivalent to its “metric”:

- BGP4 compares the MEDs of two otherwise equivalent paths **if and only if** the routes were learned from the same neighboring AS. This behavior is called **deterministic MED**. Deterministic MED is always enabled and cannot be disabled.

In addition, you can enable the Layer 3 Switch to always compare the MEDs, regardless of the AS information in the paths. To enable this comparison, enter the **always-compare-med** command at the BGP4 configuration level of the CLI. This option is disabled by default.

- The Layer 3 Switch compares the MEDs based on one or more of the following conditions. By default, the Layer 3 Switch compares the MEDs of paths **only if** the first AS in the paths is the same. (The Layer 3 Switch skips over the AS-CONFED-SEQUENCE if present.)

You can enable the Layer 3 Switch to always compare the MEDs, regardless of the AS information in the paths. For example, if the router receives UPDATES for the same route from neighbors in three ASs, the router would compare the MEDs of all the paths together, rather than comparing the MEDs for the paths in each AS individually.

NOTE

By default, value 0 (most favorable) is used in MED comparison when the MED attribute is not present. The default MED comparison results in the Layer 3 Switch favoring the route paths that are missing their MEDs. You can use the **med-missing-as-worst** command to make the Layer 3 Switch regard a BGP route with a missing MED attribute as the least favorable route, when comparing the MEDs of the routes.

NOTE

MED comparison is not performed for internal routes originated within the local AS or confederation.

To configure the router to always compare MEDs, enter the following command.

```
PowerConnect(config-bgp-router)# always-compare-med
```

Syntax: [no] always-compare-med

Treating missing MEDs as the worst MEDs

By default, the Layer 3 Switch favors a lower MED over a higher MED during MED comparison. Since the Layer 3 Switch assigns the value 0 to a route path MED if the MED value is missing, the default MED comparison results in the Layer 3 Switch favoring the route paths that are missing their MEDs.

To change this behavior so that the Layer 3 Switch favors a route that has a MED over a route that is missing its MED, enter the following command at the BGP4 configuration level of the CLI.

```
PowerConnect(config-bgp-router)# med-missing-as-worst
```

Syntax: [no] med-missing-as-worst

NOTE

This command affects route selection only when route paths are selected based on MED comparison. It is still possible for a route path that is missing its MED to be selected based on other criteria. For example, a route path with no MED can be selected if its weight is larger than the weights of the other route paths.

Configuring route reflection parameters

Normally, all the BGP routers within an AS are fully meshed. Each of the routers has an IBGP session with each of the other BGP routers in the AS. Each IBGP router thus has a route for each of its IBGP neighbors. For large ASs containing many IBGP routers, the IBGP route information in each of the fully-meshed IBGP routers can introduce too much administrative overhead.

To avoid this problem, you can hierarchically organize your IGP routers into clusters:

- A **cluster** is a group of IGP routers organized into route reflectors and route reflector clients. You configure the cluster by assigning a cluster ID on the route reflector and identifying the IGP neighbors that are members of that cluster. All the configuration for route reflection takes place on the route reflectors. The clients are unaware that they are members of a route reflection cluster. All members of the cluster must be in the same AS. The cluster ID can be any number from 0 – 4294967295. The default is the router ID, expressed as a 32-bit number.

NOTE

If the cluster contains more than one route reflector, you need to configure the same cluster ID on all the route reflectors in the cluster. The cluster ID helps route reflectors avoid loops within the cluster.

- A **route reflector** is an IGP router configured to send BGP route information to all the clients (other BGP4 routers) within the cluster. Route reflection is enabled on all BGP4 routers by default but does not take effect unless you add route reflector clients to the router.
- A **route reflector client** is an IGP router identified as a member of a cluster. You identify a router as a route reflector client on the router that is the route reflector, not on the client. The client itself requires no additional configuration. In fact, the client does not know that it is a route reflector client. The client just knows that it receives updates from its neighbors and does not know whether one or more of those neighbors are route reflectors.


NOTE

Route reflection applies only among IBGP routers within the same AS. You cannot configure a cluster that spans multiple ASs.

[Figure 116](#) shows an example of a route reflector configuration. In this example, two Layer 3 Switches are configured as route reflectors for the same cluster. The route reflectors provide redundancy in case one of the reflectors becomes unavailable. Without redundancy, if a route reflector becomes unavailable, its clients are cut off from BGP4 updates.

AS1 contains a cluster with two route reflectors and two clients. The route reflectors are fully meshed with other BGP4 routers, but the clients are not fully meshed. They rely on the route reflectors to propagate BGP4 route updates.

FIGURE 116 Example of a route reflector configuration



Support for RFC 2796

Route reflection on devices is based on RFC 2796. This updated RFC helps eliminate routing loops that are possible in some implementations of the older specification, RFC 1966.

NOTE

The configuration procedure for route reflection is the same regardless of whether your software release is using RFC 1966 or RFC 2796. However, the operation of the feature is different as explained below.

RFC 2796 provides more details than RFC 1966 regarding the use of the route reflection attributes, `ORIGINATOR_ID` and `CLUSTER_LIST`, to help prevent loops:

- **ORIGINATOR_ID** – Specifies the router ID of the BGP4 switch that originated the route. The route reflector inserts this attribute when reflecting a route to an IBGP neighbor. If a BGP4 switch receives an advertisement that contains its own router ID as the `ORIGINATOR_ID`, the switch discards the advertisement and does not forward it.
- **CLUSTER_LIST** – A list of the route reflection clusters through which the advertisement has passed. A cluster contains a route reflector and its clients. When a route reflector reflects a route, the route reflector adds its cluster ID to the front of the `CLUSTER_LIST`. If a route reflector receives a route that has its own cluster ID, the switch discards the advertisement and does not forward it.

The device handles the attributes as follows:

- The Layer 3 Switch adds the attributes only if it is a route reflector, and only when advertising IBGP route information to other IBGP neighbors. The attributes are not used when communicating with EBGP neighbors.
- A Layer 3 Switch configured as a route reflector sets the ORIGINATOR_ID attribute to the router ID of the router that originated the route. Moreover, the route reflector sets the attribute only if this is the first time the route is being reflected (sent by a route reflector). In previous software releases, the route reflector set the attribute to the router ID of the route reflector itself. When a Layer 3 Switch receives a route that already has the ORIGINATOR_ID attribute set, the Layer 3 Switch does not change the value of the attribute.
- If a Layer 3 Switch receives a route whose ORIGINATOR_ID attribute has the value of the Layer 3 Switch own router ID, the Layer 3 Switch discards the route and does not advertise it. By discarding the route, the Layer 3 Switch prevents a routing loop. The Layer 3 Switch did not discard the route in previous software releases.
- The first time a route is reflected by a Layer 3 Switch configured as a route reflector, the route reflector adds the CLUSTER_LIST attribute to the route. Other route reflectors who receive the route from an IBGP neighbor add their cluster IDs to the front of the route CLUSTER_LIST. If the route reflector does not have a cluster ID configured, the Layer 3 Switch adds its router ID to the front of the CLUSTER_LIST.
- If Layer 3 Switch configured as a route reflector receives a route whose CLUSTER_LIST contains the route reflector own cluster ID, the route reflector discards the route and does not forward it.

Configuration procedures

To configure a Layer 3 Switch to be a BGP4 route reflector, use either of the following methods.

NOTE

All configuration for route reflection takes place on the route reflectors, not on the clients.

Enter the following commands to configure a Layer 3 Switch as route reflector 1 in [Figure 116](#) on page 782. To configure route reflector 2, enter the same commands on the Layer 3 Switch that will be route reflector 2. The clients require no configuration for route reflection.

```
PowerConnect(config-bgp-router)# cluster-id 1
PowerConnect(config-bgp-router)# neighbor 10.0.1.0 route-reflector-client
PowerConnect(config-bgp-router)# neighbor 10.0.2.0 route-reflector-client
```

Syntax: [no] cluster-id <num> | <ip-addr>

The <num> | <ip-addr> parameter specifies the cluster ID and can be a number from 0 – 4294967295 or an IP address. The default is the router ID. You can configure one cluster ID on the router. All route-reflector clients for the router are members of the cluster.

NOTE

If the cluster contains more than one route reflector, you need to configure the same cluster ID on all the route reflectors in the cluster. The cluster ID helps route reflectors avoid loops within the cluster.

To add an IBGP neighbor to the cluster, enter the following command.

Syntax: neighbor <ip-addr> route-reflector-client

For more information about the **neighbor** command, refer to [“Adding BGP4 neighbors”](#) on page 756.

By default, the clients of a route reflector are not required to be fully meshed; the routes from a client are reflected to other clients. However, if the clients are fully meshed, route reflection is not required between clients.

If you need to disable route reflection between clients, enter the following command. When the feature is disabled, route reflection does not occur between clients but reflection does still occur between clients and non-clients.

```
PowerConnect(config-bgp-router)# no client-to-client-reflection
```

Enter the following command to re-enable the feature.

```
PowerConnect(config-bgp-router)# client-to-client-reflection
```

Syntax: [no] client-to-client-reflection

Aggregating routes advertised to BGP4 neighbors

By default, the Layer 3 Switch advertises individual routes for all the networks. The aggregation feature allows you to configure the Layer 3 Switch to aggregate routes in a range of networks into a single CIDR number. For example, without aggregation, the Layer 3 Switch will individually advertise routes for networks 207.95.1.0, 207.95.2.0, and 207.95.3.0. You can configure the Layer 3 Switch to instead send a single, aggregate route for the networks. The aggregate route would be advertised as 207.95.0.0.

NOTE

To summarize CIDR networks, you must use the aggregation feature. The auto summary feature does not summarize networks that use CIDR numbers instead of class A, B, or C numbers.

To aggregate routes for 209.157.22.0, 209.157.23.0, and 209.157.24.0, enter the following command.

```
PowerConnect(config-bgp-router)# aggregate-address 209.157.0.0 255.255.0.0
```

Syntax: **aggregate-address** <ip-addr> <ip-mask> [**as-set**] [**nlri multicast | unicast | multicast unicast**] [**summary-only**] [**suppress-map** <map-name>] [**advertise-map** <map-name>] [**attribute-map** <map-name>]

The <ip-addr> and <ip-mask> parameters specify the aggregate value for the networks. Specify 0 for the host portion and for the network portion that differs among the networks in the aggregate. For example, to aggregate 10.0.1.0, 10.0.2.0, and 10.0.3.0, enter the IP address 10.0.0.0 and the network mask 255.255.0.0.

The **as-set** parameter causes the router to aggregate AS-path information for all the routes in the aggregate address into a single AS-path.

The **nlri multicast | unicast | multicast unicast** parameter specifies whether the neighbor is a multicast neighbor or a unicast neighbor. Optionally, you also can specify **unicast** if you want the Layer 3 Switch to exchange unicast (BGP4) routes as well as multicast routes with the neighbor. The default is **unicast** only.

The **summary-only** parameter prevents the router from advertising more specific routes contained within the aggregate route.

The **suppress-map** <map-name> parameter prevents the more specific routes contained in the specified route map from being advertised.

The **advertise-map** *<map-name>* parameter configures the router to advertise the more specific routes in the specified route map.

The **attribute-map** *<map-name>* parameter configures the router to set attributes for the aggregate routes based on the specified route map.

NOTE

For the **suppress-map**, **advertise-map**, and **attribute-map** parameters, the route map must already be defined. Refer to [“Defining route maps”](#) on page 797 for information on defining a route map.

Modifying redistribution parameters

By default, the Layer 3 Switch does not redistribute route information between BGP4 and the IP IGP (RIP and OSPF). You can configure the switch to redistribute OSPF routes, RIP routes, directly connected routes, or static routes into BGP4 by using the following methods.

To enable redistribution of all OSPF routes and directly attached routes into BGP4, enter the following commands.

```
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# redistribute ospf
PowerConnect(config-bgp-router)# redistribute connected
PowerConnect(config-bgp-router)# write memory
```

Syntax: [no] redistribute connected | ospf | rip | static

The **connected** parameter indicates that you are redistributing routes to directly attached devices into BGP.

The **ospf** parameter indicates that you are redistributing OSPF routes into BGP4.

NOTE

Entering **redistribute ospf** simply redistributes internal OSPF routes. If you want to redistribute external OSPF routes also, you must use the **redistribute ospf match external...** command. Refer to [“Redistributing OSPF external routes”](#) on page 786.

The **rip** parameter indicates that you are redistributing RIP routes into BGP4.

The **static** parameter indicates that you are redistributing static routes into BGP.

Refer to the following sections for details on redistributing specific routes using the CLI:

- [“Redistributing connected routes”](#) on page 785
- [“Redistributing RIP routes”](#) on page 786
- [“Redistributing OSPF external routes”](#) on page 786
- [“Redistributing static routes”](#) on page 787

Redistributing connected routes

To configure BGP4 to redistribute directly connected routes, enter the following command.

```
PowerConnect(config-bgp-router)# redistribute connected
```

Syntax: redistribute connected [metric *<num>*] [route-map *<map-name>*]

The **connected** parameter indicates that you are redistributing routes to directly attached devices into BGP4.

The **metric** *<num>* parameter changes the metric. You can specify a value from 0 – 4294967295. The default is 0.

The **route-map** *<map-name>* parameter specifies a route map to be consulted before adding the RIP route to the BGP4 route table.

NOTE

The route map you specify must already be configured on the switch. Refer to [“Defining route maps”](#) on page 797 for information about defining route maps.

Redistributing RIP routes

To configure BGP4 to redistribute RIP routes and add a metric of 10 to the redistributed routes, enter the following command.

```
PowerConnect(config-bgp-router)# redistribute rip metric 10
```

Syntax: **redistribute rip** [**metric** *<num>*] [**route-map** *<map-name>*]

The **rip** parameter indicates that you are redistributing RIP routes into BGP4.

The **metric** *<num>* parameter changes the metric. Specify a value from 0 – 4294967295. The default is 0.

The **route-map** *<map-name>* parameter specifies a route map to be consulted before adding the RIP route to the BGP4 route table.

NOTE

The route map you specify must already be configured on the switch. Refer to [“Defining route maps”](#) on page 797 for information about defining route maps.

Redistributing OSPF external routes

To configure the Layer 3 Switch to redistribute OSPF external type 1 routes, enter the following command.

```
PowerConnect(config-bgp-router)# redistribute ospf match external1
```

Syntax: **redistribute ospf** [**match internal** | **external1** | **external2**] [**metric** *<num>*] [**route-map** *<map-name>*]

The **ospf** parameter indicates that you are redistributing OSPF routes into BGP4.

The **match internal** | **external1** | **external2** parameter applies only to OSPF. This parameter specifies the types of OSPF routes to be redistributed into BGP4. The default is internal.

NOTE

If you do not enter a value for the **match** parameter, (for example, you enter **redistribute ospf** only) then only internal OSPF routes will be redistributed.

The **metric** *<num>* parameter changes the metric. Specify a value from 0 – 4294967295. The default is 0.

The **route-map** *<map-name>* parameter specifies a route map to be consulted before adding the OSPF route to the BGP4 route table.

NOTE

The route map you specify must already be configured on the switch. Refer to [“Defining route maps”](#) on page 797 for information about defining route maps.

NOTE

If you use both the **redistribute ospf route-map** *<map-name>* command and the **redistribute ospf match internal | external1 | external2** command, the software uses only the route map for filtering.

Redistributing static routes

To configure the Layer 3 Switch to redistribute static routes, enter the following command.

```
PowerConnect(config-bgp-router)# redistribute static
```

Syntax: **redistribute static** [**metric** *<num>*] [**route-map** *<map-name>*]

The **static** parameter indicates that you are redistributing static routes into BGP4.

The **metric** *<num>* parameter changes the metric. Specify a value from 0 – 4294967295. The default is 0.

The **route-map** *<map-name>* parameter specifies a route map to be consulted before adding the static route to the BGP4 route table.

NOTE

The route map you specify must already be configured on the switch. Refer to [“Defining route maps”](#) on page 797 for information about defining route maps.

Disabling or re-enabling re-advertisement of all learned BGP4 routes to all BGP4 neighbors

By default, the Layer 3 Switch re-advertises all learned best BGP4 routes to BGP4 neighbors, unless the routes are discarded or blocked by route maps or other filters.

If you want to prevent the Layer 3 Switch from re-advertising a learned best BGP4 route unless that route also is installed in the IP route table, use the following CLI method.

To disable re-advertisement of BGP4 routes to BGP4 neighbors except for routes that the software also installs in the route table, enter the following command.

```
PowerConnect(config-bgp-router)# no readvertise
```

Syntax: [no] readvertise

To re-enable re-advertisement, enter the following command.

```
PowerConnect(config-bgp-router)# readvertise
```

Redistributing IBGP routes into RIP and OSPF

By default, the Layer 3 Switch does not redistribute IBGP routes from BGP4 into RIP or OSPF. This behavior helps eliminate routing loops. However, if your network can benefit from redistributing the IBGP routes from BGP4 into OSPF or RIP, you can enable the Layer 3 Switch to redistribute the routes. To do so, use the following CLI method.

To enable the Layer 3 Switch to redistribute BGP4 routes into OSPF and RIP, enter the following command.

```
PowerConnect(config-bgp-router)# bgp-redistribute-internal
```

Syntax: [no] bgp-redistribute-internal

To disable redistribution of IBGP routes into RIP and OSPF, enter the following command.

```
PowerConnect(config-bgp-router)# no bgp-redistribute-internal
```

Filtering

This section describes the following:

- [“Filtering specific IP addresses”](#) on page 788
- [“Filtering AS-paths”](#) on page 790
- [“Filtering communities”](#) on page 793
- [“Defining IP prefix lists”](#) on page 795
- [“Defining neighbor distribute lists”](#) on page 796
- [“Defining route maps”](#) on page 797
- [“Using a table map to set the rag value”](#) on page 805
- [“Configuring cooperative BGP4 route filtering”](#) on page 806

Filtering specific IP addresses

You can configure the router to explicitly permit or deny specific IP addresses received in updates from BGP4 neighbors by defining IP address filters. The router permits all IP addresses by default. You can define up to 100 IP address filters for BGP4.

- If you want **permit** to remain the default behavior, define individual filters to deny specific IP addresses.
- If you want to change the default behavior to **deny**, define individual filters to permit specific IP addresses.

NOTE

Once you define a filter, the default action for addresses that do not match a filter is “deny”. To change the default action to “permit”, configure the last filter as “permit any any”.

Address filters can be referred to by a BGP neighbor's distribute list number as well as by match statements in a route map.

NOTE

If the filter is referred to by a route map match statement, the filter is applied in the order in which the filter is listed in the match statement.

NOTE

You also can filter on IP addresses by using IP ACLs.

To define an IP address filter to deny routes to 209.157.0.0, enter the following command.

```
PowerConnect(config-bgp-router)# address-filter 1 deny 209.157.0.0 255.255.0.0
```

Syntax: `address-filter <num> permit | deny <ip-addr> <wildcard> <mask> <wildcard>`

The `<num>` parameter is the filter number.

The `permit | deny` parameter indicates the action the Layer 3 Switch takes if the filter match is true.

- If you specify **permit**, the Layer 3 Switch permits the route into the BGP4 table if the filter match is true.
- If you specify **deny**, the Layer 3 Switch denies the route from entering the BGP4 table if the filter match is true.

NOTE

Once you define a filter, the default action for addresses that do not match a filter is “deny”. To change the default action to “permit”, configure the last filter as “permit any any”.

The `<ip-addr>` parameter specifies the IP address. If you want the filter to match on all addresses, enter **any**.

The `<wildcard>` parameter specifies the portion of the IP address to match against. The `<wildcard>` is in dotted-decimal notation (IP address format). It is a four-part value, where each part is 8 bits (one byte) separated by dots, and each bit is a one or a zero. Each part is a number ranging from 0 to 255, for example 0.0.0.255. Zeros in the mask mean the packet source address must match the `<source-ip>`. Ones mean any value matches. For example, the `<ip-addr>` and `<wildcard>` values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “209.157.22.26 0.0.0.255” as “209.157.22.26/24”. The CLI automatically converts the CIDR number into the appropriate mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of subnet lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in “/`<mask-bits>`” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the filter regardless of whether the software is configured to display the masks in CIDR format.

The `<mask>` parameter specifies the network mask. If you want the filter to match on all destination addresses, enter **any**. The wildcard works the same as described above.

Filtering AS-paths

You can filter updates received from BGP4 neighbors based on the contents of the AS-path list accompanying the updates. For example, if you want to deny routes that have the AS 4.3.2.1 in the AS-path from entering the BGP4 route table, you can define a filter to deny such routes.

The Layer 3 Switch provides the following methods for filtering on AS-path information:

- AS-path filters
- AS-path ACLs

NOTE

The Layer 3 Switch cannot actively support AS-path filters and AS-path ACLs at the same time. Use one method or the other but do not mix methods.

NOTE

Once you define a filter or ACL, the default action for updates that do not match a filter is “deny”. To change the default action to “permit”, configure the last filter or ACL as “permit any any”.

AS-path filters or AS-path ACLs can be referred to by a BGP neighbor's filter list number as well as by match statements in a route map.

Defining an AS-path filter

To define AS-path filter 4 to permit AS 2500, enter the following command.

```
PowerConnect(config-bgp-router)# as-path-filter 4 permit 2500
```

Syntax: `as-path-filter <num> permit | deny <as-path>`

The `<num>` parameter identifies the filter position in the AS-path filter list and can be from 1 – 100. Thus, the AS-path filter list can contain up to 100 filters. The Layer 3 Switch applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the Layer 3 Switch stops and does not continue applying filters from the list.

NOTE

If the filter is referred to by a route map match statement, the filter is applied in the order in which the filter is listed in the match statement.

The `permit | deny` parameter indicates the action the router takes if the filter match is true.

- If you specify `permit`, the router permits the route into the BGP4 table if the filter match is true.
- If you specify `deny`, the router denies the route from entering the BGP4 table if the filter match is true.

The `<as-path>` parameter indicates the AS-path information. You can enter an exact AS-path string if you want to filter for a specific value. You also can use regular expressions in the filter string.

Defining an AS-path ACL

To configure an AS-path list that uses ACL 1, enter a command such as the following.

```
PowerConnect(config)# ip as-path access-list 1 permit 100
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# neighbor 10.10.10.1 filter-list 1 in
```

The **ip as-path** command configures an AS-path ACL that permits routes containing AS number 100 in their AS paths. The **neighbor** command then applies the AS-path ACL to advertisements and updates received from neighbor 10.10.10.1. In this example, the only routes the Layer 3 Switch permits from neighbor 10.10.10.1 are those whose AS-paths contain AS-path number 100.

Syntax: `ip as-path access-list <string> [seq <seq-value>] deny | permit <regular-expression>`

The `<string>` parameter specifies the ACL name. (If you enter a number, the CLI interprets the number as a text string.)

The `seq <seq-value>` parameter is optional and specifies the AS-path list sequence number. You can configure up to 199 entries in an AS-path list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in an AS-path list in numerical order, beginning with the lowest sequence number.

The **deny | permit** parameter specifies the action the software takes if a route AS-path list matches a match statement in this ACL. To configure the AS-path match statements in a route map, use the **match as-path** command. Refer to [“Matching based on AS-path ACL”](#) on page 800.

The `<regular-expression>` parameter specifies the AS path information you want to permit or deny to routes that match any of the match statements within the ACL. You can enter a specific AS number or use a regular expression. For the regular expression syntax, refer to [“Using regular expressions”](#) on page 791.

The **neighbor** command uses the **filter-list** parameter to apply the AS-path ACL to the neighbor. Refer to [“Adding BGP4 neighbors”](#) on page 756.

Using regular expressions

You use a regular expression for the `<as-path>` parameter to specify a single character or multiple characters as a filter pattern. If the AS-path matches the pattern specified in the regular expression, the filter evaluation is true; otherwise, the evaluation is false.

In addition, you can include special characters that influence the way the software matches the AS-path against the filter value.

To filter on a specific single-character value, enter the character for the `<as-path>` parameter. For example, to filter on AS-paths that contain the letter “z”, enter the following command.

```
PowerConnect(config-bgp-router)# as-path-filter 1 permit z
```

To filter on a string of multiple characters, enter the characters in brackets. For example, to filter on AS-paths that contain “x”, “y”, or “z”, enter the following command.

```
PowerConnect(config-bgp-router)# as-path-filter 1 permit [xyz]
```

Special characters

When you enter as single-character expression or a list of characters, you also can use the following special characters. [Table 124](#) on page 792 lists the special characters. The description for each special character includes an example. Notice that you place some special characters in front of the characters they control but you place other special characters after the characters they control. In each case, the examples show where to place the special character.

TABLE 124 BGP4 special characters for regular expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches for "aa", "ab", "ac", and so on, but not just "a". a.
*	The asterisk matches on zero or more sequences of a pattern. For example, the following regular expression matches on an AS-path that contains the string "1111" followed by any value: 1111*
+	The plus sign matches on one or more sequences of a pattern. For example, the following regular expression matches on an AS-path that contains a sequence of "g"s, such as "deg", "degg", "deggg", and so on: deg+
?	The question mark matches on zero occurrences or one occurrence of a pattern. For example, the following regular expression matches on an AS-path that contains "dg" or "deg": de?g
^	A caret (when not used within brackets) matches on the beginning of an input string. For example, the following regular expression matches on an AS-path that begins with "3": ^3
\$	A dollar sign matches on the end of an input string. For example, the following regular expression matches on an AS-path that ends with "deg": deg\$
_	An underscore matches on one or more of the following: <ul style="list-style-type: none"> • , (comma) • { (left curly brace) • } (right curly brace) • ((left parenthesis) •) (right parenthesis) • The beginning of the input string • The end of the input string • A blank space For example, the following regular expression matches on "100" but not on "1002", "2100", and so on. _100_
[]	Square brackets enclose a range of single-character patterns. For example, the following regular expression matches on an AS-path that contains "1", "2", "3", "4", or "5": [1-5] You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets: <ul style="list-style-type: none"> • ^ - The caret matches on any characters except the ones in the brackets. For example, the following regular expression matches on an AS-path that does not contain "1", "2", "3", "4", or "5": [^1-5] • - The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. See the example above.

TABLE 124 BGP4 special characters for regular expressions (Continued)

Character	Operation
	A vertical bar (sometimes called a pipe or a “logical or”) separates two alternative values or sets of values. The AS-path can match one or the other value. For example, the following regular expression matches on an AS-path that contains either “abc” or “defg”: (abc) (defg) NOTE: The parentheses group multiple characters to be treated as one value. See the following row for more information about parentheses.
()	Parentheses allow you to create complex expressions. For example, the following complex expression matches on “abc”, “abcabc”, or “abcabcabcdefg”, but not on “abcdefgdefg”: ((abc)+) ((defg)?)

If you want to filter for a special character instead of using the special character, enter “\” (backslash) in front of the character. For example, to filter on AS-path strings containing an asterisk, enter the asterisk portion of the regular expression as “*”.

```
PowerConnect(config-bgp-router)# as-path-filter 2 deny \*
```

To use the backslash as a string character, enter two slashes. For example, to filter on AS-path strings containing a backslash, enter the backslash portion of the regular expression as “\\”.

```
PowerConnect(config-bgp-router)# as-path-filter 2 deny \\
```

Filtering communities

You can filter routes received from BGP4 neighbors based on community names. Use either of the following methods to do so.

A community is an optional attribute that identifies the route as a member of a user-defined class of routes. Community names are arbitrary values made of two five-digit integers joined by a colon. You determine what the name means when you create the community name as one of a route attributes. Each string in the community name can be a number from 0 – 65535.

This format allows you to easily classify community names. For example, a common convention used in community naming is to configure the first string as the local AS and the second string as the unique community within that AS. Using this convention, communities 1:10, 1:20, and 1:30 can be easily identified as member communities of AS 1.

The Layer 3 Switch provides the following methods for filtering on community information:

- Community filters
- Community list ACLs

NOTE

The Layer 3 Switch cannot actively support community filters and community list ACLs at the same time. Use one method or the other but do not mix methods.

NOTE

Once you define a filter or ACL, the default action for communities that do not match a filter or ACL is “deny”. To change the default action to “permit”, configure the last filter or ACL entry as “permit any any”.

Community filters or ACLs can be referred to by match statements in a route map.

Defining a community filter

To define filter 3 to permit routes that have the NO_ADVERTISE community, enter the following command.

```
PowerConnect(config-bgp-router)# community-filter 3 permit no-advertise
```

Syntax: `community-filter <num> permit | deny <num>:<num> | internet | local-as | no-advertise | no-export`

The `<num>` parameter identifies the filter position in the community filter list and can be from 1 – 100. Thus, the community filter list can contain up to 100 filters. The router applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the router stops and does not continue applying filters from the list.

NOTE

If the filter is referred to by a route map match statement, the filter is applied in the order in which the filter is listed in the match statement.

The `permit | deny` parameter indicates the action the router takes if the filter match is true.

- If you specify **permit**, the router permits the route into the BGP4 table if the filter match is true.
- If you specify **deny**, the router denies the route from entering the BGP4 table if the filter match is true.

The `<num>:<num>` parameter indicates a specific community number to filter. Use this parameter to filter for a private (administrator-defined) community. You can enter up to 20 community numbers with the same command.

If you want to filter for the well-known communities “LOCAL_AS”, “NO_EXPORT” or “NO_ADVERTISE”, use the corresponding keyword (described below).

The **internet** keyword checks for routes that do not have the community attribute. Routes without a specific community are considered by default to be members of the largest community, the Internet.

The **local-as** keyword checks for routes with the well-known community “LOCAL_AS”. The Layer 3 Switch advertises the route only within the sub-AS.

The **no-advertise** keyword filters for routes with the well-known community “NO_ADVERTISE”. A route in this community should not be advertised to any BGP4 neighbors.

The **no-export** keyword filters for routes with the well-known community “NO_EXPORT”. A route in this community should not be advertised to any BGP4 neighbors outside the local AS. If the router is a member of a confederation, the Layer 3 Switch advertises the route only within the confederation.

Defining a community ACL

To configure community ACL 1, enter a command such as the following.

```
PowerConnect(config)# ip community-list 1 permit 123:2
```

This command configures a community ACL that permits routes that contain community 123:2.

NOTE

Refer to [“Matching based on community ACL”](#) on page 800 for information about how to use a community list as a match condition in a route map.

Syntax: `ip community-list standard <string> [seq <seq-value>] deny | permit <community-num>`

Syntax: `ip community-list extended <string> [seq <seq-value>] deny | permit <community-num> | <regular-expression>`

The `<string>` parameter specifies the ACL name. (If you enter a number, the CLI interprets the number as a text string.)

The **standard** or **extended** parameter specifies whether you are configuring a standard community ACL or an extended one. A standard community ACL does not support regular expressions whereas an extended one does. This is the only difference between standard and extended IP community lists.

The **seq** `<seq-value>` parameter is optional and specifies the community list sequence number. You can configure up to 199 entries in a community list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in a community list in numerical order, beginning with the lowest sequence number.

The **deny | permit** parameter specifies the action the software takes if a route community list matches a match statement in this ACL. To configure the community-list match statements in a route map, use the **match community** command. Refer to [“Matching based on community ACL”](#) on page 800.

The `<community-num>` parameter specifies the community type or community number. This parameter can have the following values:

- `<num>:<num>` – A specific community number
- **internet** – The Internet community
- **no-export** – The community of sub-ASs within a confederation. Routes with this community can be exported to other sub-ASs within the same confederation but cannot be exported outside the confederation to other ASs or otherwise sent to EBGp neighbors.
- **local-as** – The local sub-AS within the confederation. Routes with this community can be advertised only within the local subAS.
- **no-advertise** – Routes with this community cannot be advertised to any other BGP4 routers at all.

The `<regular-expression>` parameter specifies a regular expression for matching on community names. For information about regular expression syntax, refer to [“Using regular expressions”](#) on page 791. You can specify a regular expression only in an extended community ACL.

Defining IP prefix lists

An IP prefix list specifies a list of networks. When you apply an IP prefix list to a neighbor, the Layer 3 Switch sends or receives only a route whose destination is in the IP prefix list. You can configure up to 100 prefix lists. The software interprets the prefix lists in order, beginning with the lowest sequence number.

To configure an IP prefix list and apply it to a neighbor, enter commands such as the following.

```
PowerConnect(config)# ip prefix-list Routesfor20 permit 20.20.0.0/24
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# neighbor 10.10.10.1 prefix-list Routesfor20 out
```

These commands configure an IP prefix list named `Routesfor20`, which permits routes to network `20.20.0.0/24`. The `neighbor` command configures the Layer 3 Switch to use IP prefix list `Routesfor20` to determine which routes to send to neighbor `10.10.10.1`. The Layer 3 Switch sends routes that go to `20.20.x.x` to neighbor `10.10.10.1` because the IP prefix list explicitly permits these routes to be sent to the neighbor.

Syntax: `ip prefix-list <name> [seq <seq-value>] [description <string>] deny | permit <network-addr>/<mask-bits> [ge <ge-value>] [le <le-value>]`

The `<name>` parameter specifies the prefix list name. You use this name when applying the prefix list to a neighbor.

The `description <string>` parameter is a text string describing the prefix list.

The `seq <seq-value>` parameter is optional and specifies the IP prefix list sequence number. You can configure up to 100 prefix list entries. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with prefix list entry 5. The software interprets the prefix list entries in numerical order, beginning with the lowest sequence number.

The `deny | permit` parameter specifies the action the software takes if a neighbor route is in this prefix list.

The prefix-list matches only on this network unless you use the `ge <ge-value>` or `le <le-value>` parameters. (See below.)

The `<network-addr>/<mask-bits>` parameter specifies the network number and the number of bits in the network mask.

You can specify a range of prefix length for prefixes that are more specific than `<network-addr>/<mask-bits>`.

- If you specify only `ge <ge-value>`, then the mask-length range is from `<ge-value>` to 32.
- If you specify only `le <le-value>`, then the mask-length range is from length to `<le-value>`.

The `<ge-value>` or `<le-value>` you specify must meet the following condition.

length < ge-value <= le-value <= 32

If you do not specify `ge <ge-value>` or `le <le-value>`, the prefix list matches only on the exact network prefix you specify with the `<network-addr>/<mask-bits>` parameter.

For the syntax of the `neighbor` command shown in the example above, refer to [“Adding BGP4 neighbors”](#) on page 756.

Defining neighbor distribute lists

A neighbor distribute list is a list of BGP4 address filters or ACLs that filter the traffic to or from a neighbor. To configure a neighbor distribute list, use either of the following methods.

To configure a distribute list that uses ACL 1, enter a command such as the following.

```
PowerConnect(config-bgp-router)# neighbor 10.10.10.1 distribute-list 1 in
```

This command configures the Layer 3 Switch to use ACL 1 to select the routes that the Layer 3 Switch will accept from neighbor `10.10.10.1`.

Syntax: `neighbor <ip-addr> distribute-list <name-or-num> in | out`

The `<ip-addr>` parameter specifies the neighbor.

The *<name-or-num>* parameter specifies the name or number of a standard, extended, or named ACL.

The **in** | **out** parameter specifies whether the distribute list applies to inbound or outbound routes:

- **in** – controls the routes the Layer 3 Switch will accept from the neighbor.
- **out** – controls the routes sent to the neighbor.

NOTE

The command syntax shown above is new. However, the **neighbor <ip-addr> distribute-list in | out <num>** command (where the direction is specified before the filter number) is the same as in earlier software releases. Use the new syntax when you are using an IP ACL with the distribute list. Use the old syntax when you are using a BGP4 address filter with the distribute list.

Defining route maps

A **route map** is a named set of match conditions and parameter settings that the router can use to modify route attributes and to control redistribution of the routes into other protocols. A route map consists of a sequence of up to 50 **instances**. If you think of a route map as a table, an instance is a row in that table. The router evaluates a route according to a route map instances in ascending numerical order. The route is first compared against instance 1, then against instance 2, and so on. As soon as a match is found, the router stops evaluating the route against the route map instances.

Route maps can contain **match** statements and **set** statements. Each route map contains a “permit” or “deny” action for routes that match the match statements:

- If the route map contains a permit action, a route that matches a match statement is permitted; otherwise, the route is denied.
- If the route map contains a deny action, a route that matches a match statement is denied.
- If a route does not match any match statements in the route map, the route is denied. This is the default action. To change the default action, configure the last match statement in the last instance of the route map to “permit any any”.
- If there is no match statement, the software considers the route to be a match.
- For route maps that contain address filters, AS-path filters, or community filters, if the action specified by a filter conflicts with the action specified by the route map, the route map action takes precedence over the individual filter action.

If the route map contains set statements, routes that are permitted by the route map match statements are modified according to the set statements.

Match statements compare the route against one or more of the following:

- The route BGP4 MED (metric)
- A sequence of AS-path filters
- A sequence of community filters
- A sequence of address filters
- The IP address of the next hop router
- The route tag
- For OSPF routes only, the route type (internal, external type-1, or external type-2)
- An AS-path ACL

- A community ACL
- An IP prefix list
- An IP ACL

For routes that match all of the match statements, the route map set statements can perform one or more of the following modifications to the route attributes:

- Prepend AS numbers to the front of the route AS-path. By adding AS numbers to the AS-path, you can cause the route to be less preferred when compared to other routes on the basis of the length of the AS-path.
- Add a user-defined tag to the route or add an automatically calculated tag to the route.
- Set the community value.
- Set the local preference.
- Set the MED (metric).
- Set the IP address of the next hop router.
- Set the origin to IGP or INCOMPLETE.
- Set the weight.

For example, when you configure parameters for redistributing routes into RIP, one of the optional parameters is a route map. If you specify a route map as one of the redistribution parameters, the router will match the route against the match statements in the route map. If a match is found and if the route map contains set statements, the router will set attributes in the route according to the set statements.

To create a route map, you define instances of the map. Each instance is identified by a sequence number. A route map can contain up to 50 instances.

To define a route map, use the procedures in the following sections.

Entering the route map into the software

To add instance 1 of a route map named “GET_ONE” with a permit action, enter the following command.

```
PowerConnect(config)# route-map GET_ONE permit 1
PowerConnect(config-routemap GET_ONE)#
```

Syntax: [no] route-map <map-name> permit | deny <num>

As shown in this example, the command prompt changes to the Route Map level. You can enter the match and set statements at this level. Refer to “[Specifying the match conditions](#)” on page 799 and “[Setting parameters in the routes](#)” on page 802.

The <map-name> is a string of characters that names the map. Map names can be up to 32 characters in length.

The **permit** | **deny** parameter specifies the action the router will take if a route matches a match statement.

- If you specify **deny**, the Layer 3 Switch does not advertise or learn the route.
- If you specify **permit**, the Layer 3 Switch applies the match and set statements associated with this route map instance.

The <num> parameter specifies the instance of the route map you are defining. Each route map can have up to 50 instances.

To delete a route map, enter a command such as the following. When you delete a route map, all the permit and deny entries in the route map are deleted.

```
PowerConnect(config)# no route-map Map1
```

This command deletes a route map named “Map1”. All entries in the route map are deleted.

To delete a specific instance of a route map without deleting the rest of the route map, enter a command such as the following.

```
PowerConnect(config)# no route-map Map1 permit 10
```

This command deletes the specified instance from the route map but leaves the other instances of the route map intact.

Specifying the match conditions

Use the following command to define the match conditions for instance 1 of the route map GET_ONE. This instance compares the route updates against BGP4 address filter 11.

```
PowerConnect(config-routemap GET_ONE)# match address-filters 11
```

Syntax: `match [as-path <num>] | [address-filters | as-path-filters | community-filters <num,num,...>] | [community <num>] | [community <ACL> exact-match] | [ip address <ACL> | prefix-list <string>] | [ip route-source <ACL> | prefix <name>] [metric <num>] | [next-hop <address-filter-list>] | [nlri multicast | unicast | multicast unicast] | [route-type internal | external-type1 | external-type2] | [tag <tag-value>]`

The **as-path <num>** parameter specifies an AS-path ACL. You can specify up to five AS-path ACLs. To configure an AS-path ACL, use the **ip as-path access-list** command. Refer to “[Defining an AS-path ACL](#)” on page 790.

The **address-filters | as-path-filters | community-filters <num,num,...>** parameter specifies a filter or list of filters to be matched for each route. The router treats the first match as the best match. If a route does not match any filter in the list, then the router considers the match condition to have failed. To configure these types of filters, use commands at the BGP configuration level:

- To configure an address filter, refer to “[Filtering specific IP addresses](#)” on page 788.
- To configure an AS-path filter or AS-path ACL, refer to “[Filtering AS-paths](#)” on page 790.
- To configure a community filter or community ACL, refer to “[Filtering communities](#)” on page 793.

You can enter up to six community names on the same command line.

NOTE

The filters must already be configured.

The **community <num>** parameter specifies a community ACL.

NOTE

The ACL must already be configured.

The **community <ACL> exact-match** parameter matches a route if (and only if) the route's community attributes field contains the same community numbers specified in the match statement.

The **ip address | next-hop <ACL-num> | prefix-list <string>** parameter specifies an ACL or IP prefix list. Use this parameter to match based on the destination network or next-hop gateway. To configure an IP ACL for use with this command, use the **ip access-list** command. Refer to [“ACL overview”](#) on page 361. To configure an IP prefix list, use the **ip prefix-list** command. Refer to [“Defining IP prefix lists”](#) on page 795.

The **ip route-source <ACL> | prefix <name>** parameter matches based on the source of a route (the IP address of the neighbor from which the device learned the route).

The **metric <num>** parameter compares the route MED (metric) to the specified value.

The **next-hop <address-filter-list>** parameter compares the IP address of the route next hop to the specified IP address filters. The filters must already be configured.

The **nlri multicast | unicast | multicast unicast** parameter specifies whether you want the route map to match on multicast routes, unicast routes, or both route types.

NOTE

By default, route maps apply to both unicast and multicast traffic.

The **route-type internal | external-type1 | external-type2** parameter applies only to OSPF routes. This parameter compares the route type to the specified value.

The **tag <tag-value>** parameter compares the route tag to the specified value.

Match examples using ACLs

The following sections show some detailed examples of how to configure route maps that include match statements that match on ACLs.

Matching based on AS-path ACL

To construct a route map that matches based on AS-path ACL 1, enter the following commands.

```
PowerConnect(config)# route-map PathMap permit 1
PowerConnect(config-routemap PathMap)# match as-path 1
```

Syntax: **match as-path <num>**

The **<num>** parameter specifies an AS-path ACL and can be a number from 1 – 199. You can specify up to five AS-path ACLs. To configure an AS-path ACL, use the **ip as-path access-list** command. Refer to [“Defining an AS-path ACL”](#) on page 790.

Matching based on community ACL

To construct a route map that matches based on community ACL 1, enter the following commands.

```
PowerConnect(config)# ip community-list 1 permit 123:2
PowerConnect(config)# route-map CommMap permit 1
PowerConnect(config-routemap CommMap)# match community 1
```

Syntax: **match community <string>**

The **<string>** parameter specifies a community list ACL. To configure a community list ACL, use the **ip community-list** command. Refer to [“Defining a community ACL”](#) on page 794.

Matching based on destination network

To construct match statements for a route map that match based on destination network, use the following method. You can use the results of an IP ACL or an IP prefix list as the match condition.

```
PowerConnect(config)# route-map NetMap permit 1
PowerConnect(config-routemap NetMap)# match ip address 1
```

Syntax: `match ip address <name-or-num>`

Syntax: `match ip address prefix-list <name>`

The *<name-or-num>* parameter with the first command specifies an IP ACL and can be a number from 1 – 199 or the ACL name if it is a named ACL. To configure an IP ACL, use the **ip access-list** or **access-list** command. Refer to [Chapter 13, “Configuring Rule-Based IP Access Control Lists”](#).

The *<name>* parameter with the second command specifies an IP prefix list name. To configure an IP prefix list, refer to [“Defining IP prefix lists”](#) on page 795.

Matching based on next-hop router

To construct match statements for a route map that match based on the IP address of the next-hop router, use either of the following methods. You can use the results of an IP ACL or an IP prefix list as the match condition.

To construct a route map that matches based on the next-hop router, enter commands such as the following.

```
PowerConnect(config)# route-map HopMap permit 1
PowerConnect(config-routemap HopMap)# match ip next-hop 2
```

Syntax: `match ip next-hop <num>`

Syntax: `match ip next-hop prefix-list <name>`

The *<num>* parameter with the first command specifies an IP ACL and can be a number from 1 – 199 or the ACL name if it is a named ACL. To configure an IP ACL, use the **ip access-list** or **access-list** command. Refer to [Chapter 13, “Configuring Rule-Based IP Access Control Lists”](#).

The *<name>* parameter with the second command specifies an IP prefix list name. To configure an IP prefix list, refer to [“Defining IP prefix lists”](#) on page 795.

Matching based on the route source

To match a BGP4 route based on its source, use the **match ip route-source** statement. Here is an example.

```
PowerConnect(config)# access-list 10 permit 192.168.6.0 0.0.0.255
PowerConnect(config)# route-map bgp1 permit 1
PowerConnect(config-routemap bgp1)# match ip route-source 10
```

The first command configures an IP ACL that matches on routes received from 192.168.6.0/24. The remaining commands configure a route map that matches on all BGP4 routes advertised by the BGP4 neighbors whose addresses match addresses in the IP prefix list. You can add a set statement to change a route attribute in the routes that match. You also can use the route map as input for other commands, such as the **neighbor** and **network** commands and some **show** commands.

Syntax: `match ip route-source <ACL> | prefix <name>`

The *<ACL> | prefix <name>* parameter specifies the name or ID of an IP ACL, or an IP prefix list.

Matching on routes containing a specific set of communities

Device software enables you to match routes based on the presence of a community name or number in a route, and to match when a route contains exactly the set of communities you specify. To match based on a set of communities, configure a community ACL that lists the communities, then compare routes against the ACL.

Here is an example.

```
PowerConnect(config)# ip community-list standard std_1 permit 12:34 no-export
PowerConnect(config)# route-map bgp2 permit 1
PowerConnect(config-routemap bgp2)# match community std_1 exact-match
```

The first command configures a community ACL that contains community number 12:34 and community name no-export. The remaining commands configure a route map that matches the community attributes field in BGP4 routes against the set of communities in the ACL. A route matches the route map only if the route contains all the communities in the ACL and no other communities.

Syntax: match community <ACL> exact-match

The <ACL> parameter specifies the name of a community list ACL. You can specify up to five ACLs. Separate the ACL names or IDs with spaces.

Here is another example.

```
PowerConnect(config)# ip community-list standard std_2 permit 23:45 56:78
PowerConnect(config)# route-map bgp3 permit 1
PowerConnect(config-routemap bgp3)# match community std_1 std_2 exact-match
```

These commands configure an additional community ACL, std_2, that contains community numbers 23:45 and 57:68. Route map bgp3 compares each BGP4 route against the sets of communities in ACLs std_1 and std_2. A BGP4 route that contains *either but not both* sets of communities matches the route map. For example, a route containing communities 23:45 and 57:68 matches. However, a route containing communities 23:45, 57:68 and 12:34, or communities 23:45, 57:68, 12:34, and no-export does not match. To match, the route communities must be the same as those in exactly one of the community ACLs used by the match community statement.

Setting parameters in the routes

Use the following command to define a set statement that prepends an AS number to the AS path on each route that matches the corresponding match statement.

```
PowerConnect(config-routemap GET_ONE)# set as-path prepend 65535
```

Syntax: set [as-path [prepend <as-num,as-num,...>]] | [automatic-tag] | [comm-list <ACL> delete] | [community <num>:<num> | <num> | internet | local-as | no-advertise | no-export] | [dampening [<half-life> <reuse> <suppress> <max-suppress-time>]] [[default] interface null0 | [ip [default] next hop <ip-addr>] [ip next-hop peer-address] | [local-preference <num>] | [metric [+ | -]<num> | none] | [metric-type type-1 | type-2] | [metric-type internal] | [next-hop <ip-addr>] | [nlri multicast | unicast | multicast unicast] | [origin igp | incomplete] | [tag <tag-value>] | [weight <num>]

The **as-path prepend <num,num,...>** parameter adds the specified AS numbers to the front of the AS-path list for the route.

The **automatic-tag** parameter calculates and sets an automatic tag value for the route.

NOTE

This parameter applies only to routes redistributed into OSPF.

The **comm-list** parameter deletes a community from a BGP4 route community attributes field.

The **community** parameter sets the community attribute for the route to the number or well-known type you specify.

The **dampening** [*<half-life>* *<reuse>* *<suppress>* *<max-suppress-time>*] parameter sets route dampening parameters for the route. The *<half-life>* parameter specifies the number of minutes after which the route penalty becomes half its value. The *<reuse>* parameter specifies how low a route penalty must become before the route becomes eligible for use again after being suppressed. The *<suppress>* parameter specifies how high a route penalty can become before the Layer 3 Switch suppresses the route. The *<max-suppress-time>* parameter specifies the maximum number of minutes that a route can be suppressed regardless of how unstable it is. For information and examples, refer to “[Configuring route flap dampening](#)” on page 809.

The [**default**] **interface null0** parameter redirects the traffic to the specified interface. You can send the traffic to the null0 interface, which is the same as dropping the traffic. You can specify more than one interface, in which case the Layer 3 Switch uses the first available port. If the first port is unavailable, the Layer 3 Switch sends the traffic to the next port in the list. If you specify **default**, the route map redirects the traffic to the specified interface only if the Layer 3 Switch does not already have explicit routing information for the traffic. This option is used in Policy-Based Routing (PBR).

The **ip [default] next hop** *<ip-addr>* parameter sets the next-hop IP address for traffic that matches a match statement in the route map. If you specify **default**, the route map sets the next-hop gateway only if the Layer 3 Switch does not already have explicit routing information for the traffic. This option is used in Policy-Based Routing (PBR).

The **ip next-hop peer-address** parameter sets the BGP4 next hop for a route to the specified neighbor address.

The **local-preference** *<num>* parameter sets the local preference for the route. You can set the preference to a value from 0 – 4294967295.

The **metric** [+ | -]*<num>* | none parameter sets the MED (metric) value for the route. The default MED value is 0. You can set the preference to a value from 0 – 4294967295.

- **set metric** *<num>* – Sets the route metric to the number you specify.
- **set metric** +*<num>* – Increases route metric by the number you specify.
- **set metric** -*<num>* – Decreases route metric by the number you specify.
- **set metric none** – Removes the metric from the route (removes the MED attribute from the BGP4 route).

The **metric-type type-1 | type-2** parameter changes the metric type of a route redistributed into OSPF.

The **metric-type internal** parameter sets the route's MED to the same value as the IGP metric of the BGP4 next-hop route. The parameter does this when advertising a BGP4 route to an EBGp neighbor.

The **next-hop** *<ip-addr>* parameter sets the IP address of the route next hop router.

The **nlri multicast | unicast | multicast unicast** parameter redistributes routes into the multicast Routing Information Base (RIB) instead of the unicast RIB.

NOTE

Setting the NLRI type to multicast applies only when you are using the route map to redistribute directly-connected routes. Otherwise, the set option is ignored.

The **origin igp | incomplete** parameter sets the route origin to IGP or INCOMPLETE.

The **tag <tag-value>** parameter sets the route tag. You can specify a tag value from 0 - 4294967295.

NOTE

This parameter applies only to routes redistributed into OSPF.

NOTE

You also can set the tag value using a table map. The table map changes the value only when the Layer 3 Switch places the route in the IP route table instead of changing the value in the BGP route table. Refer to [“Using a table map to set the tag value”](#) on page 805.

The **weight <num>** parameter sets the weight for the route. You can specify a weight value from 0 - 4294967295.

Setting a BGP4 route MED to the same value as the IGP metric of the next-hop route

To set a route's MED to the same value as the IGP metric of the BGP4 next-hop route, when advertising the route to a neighbor, enter commands such as the following.

```
PowerConnect(config)# access-list 1 permit 192.168.9.0 0.0.0.255
PowerConnect(config)# route-map bgp4 permit 1
PowerConnect(config-routemap bgp4)# match ip address 1
PowerConnect(config-routemap bgp4)# set metric-type internal
```

The first command configures an ACL that matches on routes with destination network 192.168.9.0. The remaining commands configure a route map that matches on the destination network in ACL 1, then sets the metric type for those routes to the same value as the IGP metric of the BGP4 next-hop route.

Syntax: set metric-type internal

Setting the next hop of a BGP4 route

To set the next hop address of a BGP4 route to a neighbor address, enter commands such as the following.

```
PowerConnect(config)# route-map bgp5 permit 1
PowerConnect(config-routemap bgp5)# match ip address 1
PowerConnect(config-routemap bgp5)# set ip next-hop peer-address
```

These commands configure a route map that matches on routes whose destination network is specified in ACL 1, and sets the next hop in the routes to the neighbor address (inbound filtering) or the local IP address of the BGP4 session (outbound filtering).

Syntax: set ip next-hop peer-address

The value that the software substitutes for **peer-address** depends on whether the route map is used for inbound filtering or outbound filtering:

- When you use the **set ip next-hop peer-address** command in an inbound route map filter, **peer-address** substitutes for the neighbor IP address.

- When you use the **set ip next-hop peer-address** command in an outbound route map filter, **peer-address** substitutes for the local IP address of the BGP4 session.

NOTE

You can use this command for a peer group configuration.

Deleting a community from a BGP4 route

To delete a community from a BGP4 route community attributes field, enter commands such as the following.

```
PowerConnect(config)# ip community-list standard std_3 permit 12:99 12:86
PowerConnect(config)# route-map bgp6 permit 1
PowerConnect(config-routemap bgp6)# match ip address 1
PowerConnect(config-routemap bgp6)# set comm-list std_3 delete
```

The first command configures a community ACL containing community numbers 12:99 and 12:86. The remaining commands configure a route map that matches on routes whose destination network is specified in ACL 1, and deletes communities 12:99 and 12:86 from those routes. The route does not need to contain all the specified communities in order for them to be deleted. For example, if a route contains communities 12:86, 33:44, and 66:77, community 12:86 is deleted.

Syntax: set comm-list <ACL> delete

The <ACL> parameter specifies the name of a community list ACL.

Using a table map to set the rag value

Route maps that contain set statements change values in routes when the routes are accepted by the route map. For inbound route maps (route maps that filter routes received from neighbors), this means that the routes are changed before they enter the BGP4 route table.

For tag values, if you do not want the value to change until a route enters the IP route table, you can use a table map to change the value. A table map is a route map that you have associated with the IP routing table. The Layer 3 Switch applies the set statements for tag values in the table map to routes before adding them to the route table.

To configure a table map, you configure the route map, then identify it as a table map. The table map does not require separate configuration. You create it simply by calling an existing route map a table map. You can have one table map.

NOTE

Use table maps only for setting the tag value. Do not use table maps to set other attributes. To set other route attributes, use route maps or filters.

To create a route map and identify it as a table map, enter commands such as following. These commands create a route map that uses an address filter. For routes that match the address filter, the route map changes the tag value to 100. This route map is then identified as a table map. As a result, the route map is applied only to routes that the Layer 3 Switch places in the IP route table. The route map is not applied to all routes. This example assumes that address filter 11 has already been configured.

```
PowerConnect(config)# route-map TAG_IP permit 1
PowerConnect(config-routemap TAG_IP)# match address-filters 11
PowerConnect(config-routemap TAG_IP)# set tag 100
PowerConnect(config-routemap TAG_IP)# router bgp
PowerConnect(config-bgp-router)# table-map TAG_IP
```

Configuring cooperative BGP4 route filtering

By default, the Layer 3 Switch performs all filtering of incoming routes locally, on the Layer 3 Switch itself. You can use cooperative BGP4 route filtering to cause the filtering to be performed by a neighbor before it sends the routes to the Layer 3 Switch. Cooperative filtering conserves resources by eliminating unnecessary route updates and filter processing. For example, the Layer 3 Switch can send a deny filter to its neighbor, which the neighbor uses to filter out updates before sending them to the Layer 3 Switch. The neighbor saves the resources it would otherwise use to generate the route updates, and the Layer 3 Switch saves the resources it would use to filter out the routes.

When you enable cooperative filtering, the Layer 3 Switch advertises this capability in its Open message to the neighbor when initiating the neighbor session. The Open message also indicates whether the Layer 3 Switch is configured to send filters, receive filters or both, and the types of filters it can send or receive. The Layer 3 Switch sends the filters as Outbound Route Filters (ORFs) in Route Refresh messages.

To configure cooperative filtering, perform the following tasks on the Layer 3 Switch and on its BGP4 neighbor:

- Configure the filter.

NOTE

The current release supports cooperative filtering only for filters configured using IP prefix lists.

- Apply the filter as an *inbound* filter to the neighbor.
- Enable the cooperative route filtering feature on the Layer 3 Switch. You can enable the Layer 3 Switch to send ORFs to the neighbor, to receive ORFs from the neighbor, or both. The neighbor uses the ORFs you send as outbound filters when it sends routes to the Layer 3 Switch. Likewise, the Layer 3 Switch uses the ORFs it receives from the neighbor as outbound filters when sending routes to the neighbor.
- Reset the BGP4 neighbor session to send and receive ORFs.
- Perform these steps on the other device.

NOTE

If the Layer 3 Switch has inbound filters, the filters are still processed even if equivalent filters have been sent as ORFs to the neighbor.

Enabling cooperative filtering

To configure cooperative filtering, enter commands such as the following.

```
PowerConnect(config)# ip prefix-list Routesfrom1234 deny 20.20.0.0/24
PowerConnect(config)# ip prefix-list Routesfrom1234 permit 0.0.0.0/0 le 32
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# neighbor 1.2.3.4 prefix-list Routesfrom1234 in
PowerConnect(config-bgp-router)# neighbor 1.2.3.4 capability orf prefixlist send
```

The first two commands configure statements for the IP prefix list `Routesfrom1234`. The first command configures a statement that denies routes to `20.20.0.0/24`. The second command configures a statement that permits all other routes. (Once you configure an IP prefix list statement, all routes not explicitly permitted by statements in the prefix list are denied.)

The next two commands change the CLI to the BGP4 configuration level, then apply the IP prefix list to neighbor 1.2.3.4. The last command enables the Layer 3 Switch to send the IP prefix list as an ORF to neighbor 1.2.3.4. When the Layer 3 Switch sends the IP prefix list to the neighbor, the neighbor filters out the 20.20.0.x routes from its updates to the Layer 3 Switch. (This assumes that the neighbor also is configured for cooperative filtering.)

The `<ip-addr> | <peer-group-name>` parameter specifies the IP address of a neighbor or the name of a peer group of neighbors.

The **send | receive** parameter specifies the support you are enabling:

- **send** – The Layer 3 Switch sends the IP prefix lists to the neighbor.
- **receive** – The Layer 3 Switch accepts filters from the neighbor.

If you do not specify the capability, both capabilities are enabled.

The **prefixlist** parameter specifies the type of filter you want to send to the neighbor.

NOTE

The current release supports cooperative filtering only for filters configured using IP prefix lists.

Sending and receiving ORFs

Cooperative filtering affects neighbor sessions that start after the filtering is enabled, but do not affect sessions that are already established.

To activate cooperative filtering, reset the session with the neighbor. This is required because the cooperative filtering information is exchanged in Open messages during the start of a session.

To place a prefix-list change into effect after activating cooperative filtering, perform a soft reset of the neighbor session. A soft reset does not end the current session, but sends the prefix list to the neighbor in the next route refresh message.

NOTE

Make sure cooperative filtering is enabled on the Layer 3 Switch and on the neighbor before you send the filters.

To reset a neighbor session and send ORFs to the neighbor, enter a command such as the following.

```
PowerConnect# clear ip bgp neighbor 1.2.3.4
```

This command resets the BGP4 session with neighbor 1.2.3.4 and sends the ORFs to the neighbor. If the neighbor sends ORFs to the Layer 3 Switch, the Layer 3 Switch accepts them if the send capability is enabled.

To perform a soft reset of a neighbor session and send ORFs to the neighbor, enter a command such as the following.

```
PowerConnect# clear ip bgp neighbor 1.2.3.4 soft in prefix-list
```

Syntax: `clear ip bgp neighbor <ip-addr> [soft in prefix-filter]`

If you use the **soft in prefix-filter** parameter, the Layer 3 Switch sends the updated IP prefix list to the neighbor as part of its route refresh message to the neighbor.

NOTE

If the Layer 3 Switch or the neighbor is not configured for cooperative filtering, the command sends a normal route refresh message.

Displaying cooperative filtering information

You can display the following cooperative filtering information:

- The cooperative filtering configuration on the Layer 3 Switch.
- The ORFs received from neighbors.

To display the cooperative filtering configuration on the Layer 3 Switch, enter a command such as the following. The line shown in bold type shows the cooperative filtering status.

```
PowerConnect# show ip bgp neighbor 10.10.10.1
1  IP Address: 10.10.10.1, AS: 65200 (IBGP), RouterID: 10.10.10.1
   State: ESTABLISHED, Time: 0h0m7s, KeepAliveTime: 60, HoldTime: 180
   RefreshCapability: Received
   CooperativeFilteringCapability: Received
   Messages:      Open      Update  KeepAlive  Notification  Refresh-Req
   Sent          : 1        0       1          0              1
   Received: 1    0       1          0              1
   Last Update Time: NLRI      Withdraw      NLRI      Withdraw
                   Tx: ---      ---          Rx: ---      ---
   Last Connection Reset Reason:Unknown
   Notification Sent:      Unspecified
   Notification Received: Unspecified
   TCP Connection state: ESTABLISHED
   Byte Sent:      110, Received: 110
   Local host:    10.10.10.2, Local Port: 8138
   Remote host:  10.10.10.1, Remote Port: 179
   ISentSeq:      460  SendNext:      571  TotUnAck:      0
   TotSent:       111  ReTrans:       0   UnAckSeq:      571
   IRcvSeq:       7349 RcvNext:      7460  SendWnd:       16384
   TotalRcv:      111  DupliRcv:     0   RcvWnd:        16384
   SendQue:        0   RcvQue:       0   CngstWnd:      5325
```

Syntax: `show ip bgp neighbor <ip-addr>`

To display the ORFs received from a neighbor, enter a command such as the following.

```
PowerConnect# show ip bgp neighbor 10.10.10.1 received prefix-filter
ip prefix-list 10.10.10.1: 4 entries
  seq 5 permit 10.10.0.0/16 ge 18 le 28
  seq 10 permit 20.20.10.0/24
  seq 15 permit 30.0.0.0/8 le 32
  seq 20 permit 40.10.0.0/16 ge 18
```

Syntax: `show ip bgp neighbor <ip-addr> received prefix-filter`

Configuring route flap dampening

A “route flap” is the change in a route state, from up to down or down to up. When a route state changes, the state change causes changes in the route tables of the routers that support the route. Frequent changes in a route state can cause Internet instability and add processing overhead to the routers that support the route.

Route flap dampening is a mechanism that reduces the impact of route flap by changing a BGP4 router response to route state changes. When route flap dampening is configured, the Layer 3 Switch suppresses unstable routes until the route state changes reduce enough to meet an acceptable degree of stability. The Dell implementation of route flap dampening is based on RFC 2439.

Route flap dampening is disabled by default. You can enable the feature globally or on an individual route basis using route maps.

NOTE

The Layer 3 Switch applies route flap dampening only to routes learned from EBGp neighbors.

The route flap dampening mechanism is based on penalties. When a route exceeds a configured penalty value, the Layer 3 Switch stops using that route and also stops advertising it to other routers. The mechanism also allows a route penalties to reduce over time if the route stability improves. The route flap dampening mechanism uses the following parameters:

- **Suppression threshold** – Specifies the penalty value at which the Layer 3 Switch stops using the route. Each time a route becomes unreachable or is withdrawn by a BGP4 UPDATE from a neighbor, the route receives a penalty of 1000. By default, when a route has a penalty value greater than 2000, the Layer 3 Switch stops using the route. Thus, by default, if a route goes down more than twice, the Layer 3 Switch stops using the route. You can set the suppression threshold to a value from 1 – 20000. The default is 2000.
- **Half-life** – Once a route has been assigned a penalty, the penalty decreases exponentially and decreases by half after the half-life period. The default half-life period is 15 minutes. The software reduces route penalties every five seconds. For example, if a route has a penalty of 2000 and does not receive any more penalties (it does not go down again) during the half-life, the penalty is reduced to 1000 after the half-life expires. You can configure the half-life to be from 1 – 45 minutes. The default is 15 minutes.
- **Reuse threshold** – Specifies the minimum penalty a route can have and still be suppressed by the Layer 3 Switch. If the route's penalty falls below this value, the Layer 3 Switch un-suppresses the route and can use it again. The software evaluates the dampened routes every ten seconds and un-suppresses the routes that have penalties below the reuse threshold. You can set the reuse threshold to a value from 1 – 20000. The default is 750.
- **Maximum suppression time** – Specifies the maximum number of minutes a route can be suppressed regardless of how unstable the route has been before this time. You can set the parameter to a value from 1 – 20000 minutes. The default is four times the half-life. When the half-life value is set to its default (15 minutes), the maximum suppression time defaults to 60 minutes.

You can configure route flap dampening globally or for individual routes using route maps. If you configure route flap dampening parameters globally and also use route maps, the settings in the route maps override the global values.

Globally configuring route flap dampening

To enable route flap dampening using the default values, enter the following command.

```
PowerConnect(config-bgp-router)# dampening
```

Syntax: `dampening` [`<half-life>` `<reuse>` `<suppress>` `<max-suppress-time>`]

The `<half-life>` parameter specifies the number of minutes after which the route penalty becomes half its value. The route penalty allows routes that have remained stable for a while despite earlier instability to eventually become eligible for use again. The decay rate of the penalty is proportional to the value of the penalty. After the half-life expires, the penalty decays to half its value. Thus, a dampened route that is no longer unstable can eventually become eligible for use again. You can configure the half-life to be from 1 - 45 minutes. The default is 15 minutes.

The `<reuse>` parameter specifies how low a route penalty must become before the route becomes eligible for use again after being suppressed. You can set the reuse threshold to a value from 1 - 20000. The default is 750 (0.75, or three-fourths, of the penalty assessed for a one “flap”).

The `<suppress>` parameter specifies how high a route penalty can become before the Layer 3 Switch suppresses the route. You can set the suppression threshold to a value from 1 - 20000. The default is 2000 (two “flaps”).

The `<max-suppress-time>` parameter specifies the maximum number of minutes that a route can be suppressed regardless of how unstable it is. You can set the maximum suppression time to a value from 1 - 20000 minutes. The default is four times the half-life setting. Thus, if you use the default half-life of 15 minutes, the maximum suppression time is 60 minutes.

The following example shows how to change the dampening parameters.

```
PowerConnect(config-bgp-router)# dampening 20 200 2500 40
```

This command changes the half-life to 20 minutes, the reuse threshold to 200, the suppression threshold to 2500, and the maximum number of minutes a route can be dampened to 40.

NOTE

To change any of the parameters, you must specify all the parameters with the command. If you want to leave some parameters unchanged, enter their default values.

Using a route map to configure route flap dampening for specific routes

Route maps enable you to fine tune route flap dampening parameters for individual routes. To configure route flap dampening parameters using route maps, configure BGP4 address filters for each route you want to set the dampening parameters for, then configure route map entries that set the dampening parameters for those routes. The following sections show examples.

To configure address filters and a route map for dampening specific routes, enter commands such as the following.


```

PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# address-filter 9 permit 209.157.22.0
255.255.255.0 255.255.255.0 255.255.255.0
PowerConnect(config-bgp-router)# address-filter 10 permit 209.157.23.0
255.255.255.0 255.255.255.0 255.255.255.0
PowerConnect(config-bgp-router)# exit
PowerConnect(config)# route-map DAMPENING_MAP permit 9
PowerConnect(config-routemap DAMPENING_MAP)# match address-filters 9
PowerConnect(config-routemap DAMPENING_MAP)# set dampening 10 200 2500 40
PowerConnect(config-routemap DAMPENING_MAP)# exit
PowerConnect(config)# route-map DAMPENING_MAP permit 10
PowerConnect(config-routemap DAMPENING_MAP)# match address-filters 10
PowerConnect(config-routemap DAMPENING_MAP)# set dampening 20 200 2500 60
PowerConnect(config-routemap DAMPENING_MAP)# router bgp
PowerConnect(config-bgp-router)# dampening route-map DAMPENING_MAP

```

The **address-filter** commands in this example configure two BGP4 address filters, for networks 209.157.22.0 and 209.157.23.0. The first route-map command creates an entry in a route map called "DAMPENING_MAP". Within this entry of the route map, the **match** command matches based on address filter 9, and the **set** command sets the dampening parameters for the route that matches. Thus, for BGP4 routes to 209.157.22.0, the Layer 3 Switch uses the route map to set the dampening parameters. These parameters override the globally configured dampening parameters.

The commands for the second entry in the route map (instance 10 in this example) perform the same functions for route 209.157.23.0. Notice that the dampening parameters are different for each route.

Using a route map to configure route flap dampening for a specific neighbor

You can use a route map to configure route flap dampening for a specific neighbor by performing the following tasks:

- Configure an empty route map with no match or set statements. This route map does not specify particular routes for dampening but does allow you to enable dampening globally when you refer to this route map from within the BGP configuration level.
- Configure another route map that explicitly enables dampening. Use a set statement within the route map to enable dampening. When you associate this route map with a specific neighbor, the route map enables dampening for all routes associated with the neighbor. You also can use match statements within the route map to selectively perform dampening on some routes from the neighbor.

NOTE

You still need to configure the first route map to enable dampening globally. The second route map does not enable dampening by itself; it just applies dampening to a neighbor.

- Apply the route map to the neighbor.

To enable route flap dampening for a specific BGP4 neighbor, enter commands such as the following.

```
PowerConnect(config)# route-map DAMPENING_MAP_ENABLE permit 1
PowerConnect(config-routemap DAMPENING_MAP_ENABLE)# exit
PowerConnect(config)# route-map DAMPENING_MAP_NEIGHBOR_A permit 1
PowerConnect(config-routemap DAMPENING_MAP_NEIGHBOR_A)# set dampening
PowerConnect(config-routemap DAMPENING_MAP_NEIGHBOR_A)# exit
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# dampening route-map DAMPENING_MAP_ENABLE
PowerConnect(config-bgp-router)# neighbor 10.10.10.1 route-map in
DAMPENING_MAP_NEIGHBOR_A
```

In this example, the first command globally enables route flap dampening. This route map does not contain any match or set statements. At the BGP configuration level, the **dampening route-map** command refers to the DAMPENING_MAP_ENABLE route map created by the first command, thus enabling dampening globally.

The third and fourth commands configure a second route map that explicitly enables dampening. Notice that the route map does not contain a match statement. The route map implicitly applies to all routes. Since the route map will be applied to a neighbor at the BGP configuration level, the route map will apply to all routes associated with the neighbor.

Although the second route map enables dampening, the first route map is still required. The second route map enables dampening for the neighbors to which the route map is applied. However, unless dampening is already enabled globally by the first route map, the second route map has no effect.

The last two commands apply the route maps. The **dampening route-map** command applies the first route map, which enables dampening globally. The **neighbor** command applies the second route map to neighbor 10.10.10.1. Since the second route map does not contain match statements for specific routes, the route map enables dampening for all routes received from the neighbor.

Removing route dampening from a route

You can un-suppress routes by removing route flap dampening from the routes. The Layer 3 Switch allows you to un-suppress all routes at once or un-suppress individual routes.

To un-suppress all the suppressed routes, enter the following command at the Privileged EXEC level of the CLI.

```
PowerConnect# clear ip bgp damping
```

Syntax: `clear ip bgp damping [<ip-addr> <ip-mask>]`

The *<ip-addr>* parameter specifies a particular network.

The *<ip-mask>* parameter specifies the network mask.

To un-suppress a specific route, enter a command such as the following.

```
PowerConnect# clear ip bgp damping 209.157.22.0 255.255.255.0
```

This command un-suppresses only the routes for network 209.157.22.0/24.

Removing route dampening from a neighbor routes suppressed due to aggregation

You can selectively unsuppress more-specific routes that have been suppressed due to aggregation, and allow the routes to be advertised to a specific neighbor or peer group.

Here is an example.

```
PowerConnect(config-bgp-router)# aggregate-address 209.1.0.0 255.255.0.0
summary-only
PowerConnect(config-bgp-router)# show ip bgp route 209.1.0.0/16 longer
Number of BGP Routes matching display condition : 2
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
Prefix      Next Hop      Metric      LocPrf      Weight Status
1   209.1.0.0/16  0.0.0.0      101         32768  BAL
    AS_PATH:
2   209.1.44.0/24 10.2.0.1      1           101         32768  BLS
    AS_PATH:
```

The **aggregate-address** command configures an aggregate address. The **summary-only** parameter prevents the Layer 3 Switch from advertising more specific routes contained within the aggregate route. The **show ip bgp route** command shows that the more specific routes aggregated into 209.1.0.0/16 have been suppressed. In this case, the route to 209.1.44.0/24 has been suppressed. The following command indicates that the route is not being advertised to the Layer 3 Switch BGP4 neighbors.

```
PowerConnect# show ip bgp route 209.1.44.0/24
Number of BGP Routes matching display condition : 1
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
Prefix      Next Hop      Metric      LocPrf      Weight Status
1   209.1.44.0/24 10.2.0.1      1           101         32768  BLS
    AS_PATH:
Route is not advertised to any peers
```

If you want to override the **summary-only** parameter and allow a specific route to be advertised to a neighbor, enter commands such as the following.

```
PowerConnect(config)# ip prefix-list Unsuppress1 permit 209.1.44.0/24
PowerConnect(config)# route-map RouteMap1 permit 1
PowerConnect(config-routemap RouteMap1)# match prefix-list Unsuppress1
PowerConnect(config-routemap RouteMap1)# exit
PowerConnect(config)# router bgp
PowerConnect(config-bgp-router)# neighbor 10.1.0.2 unsuppress-map RouteMap1
PowerConnect(config-bgp-router)# clear ip bgp neighbor 10.1.0.2 soft-out
```

The **ip prefix-list** command configures an IP prefix list for network 209.1.44.0/24, which is the route you want to unsuppress. The next two commands configure a route map that uses the prefix list as input. The **neighbor** command enables the Layer 3 Switch to advertise the routes specified in the route map to neighbor 10.1.0.2. The **clear** command performs a soft reset of the session with the neighbor so that the Layer 3 Switch can advertise the unsuppressed route.

Syntax: [no] **neighbor** <ip-addr> | <peer-group-name> **unsuppress-map** <map-name>

The following command verifies that the route has been unsuppressed.

```
PowerConnect# show ip bgp route 209.1.44.0/24
Number of BGP Routes matching display condition : 1
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
      E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
Prefix          Next Hop          Metric      LocPrf      Weight Status
1      209.1.44.0/24      10.2.0.1      1           101         32768 BLS
      AS_PATH:
Route is advertised to 1 peers:
      10.1.0.2(4)
```

Displaying and clearing route flap dampening statistics

The software provides many options for displaying and clearing route flap statistics. To display the statistics, use either of the following methods.

Displaying route flap dampening statistics

To display route dampening statistics or all the dampened routes, enter the following command at any level of the CLI.

```
PowerConnect# show ip bgp flap-statistics
Total number of flapping routes: 414
      Status Code >:best d:damped h:history *:valid
      Network          From          Flaps Since      Reuse      Path
h> 192.50.206.0/23     166.90.213.77  1      0 :0 :13 0 :0 :0 65001 4355 1 701
h> 203.255.192.0/20    166.90.213.77  1      0 :0 :13 0 :0 :0 65001 4355 1 7018
h> 203.252.165.0/24    166.90.213.77  1      0 :0 :13 0 :0 :0 65001 4355 1 7018
h> 192.50.208.0/23     166.90.213.77  1      0 :0 :13 0 :0 :0 65001 4355 1 701
h> 133.33.0.0/16       166.90.213.77  1      0 :0 :13 0 :0 :0 65001 4355 1 701
*> 204.17.220.0/24     166.90.213.77  1      0 :1 :4  0 :0 :0 65001 4355 701 62
```

Syntax: `show ip bgp flap-statistics [regular-expression <regular-expression> | <address> <mask> [longer-prefixes] | neighbor <ip-addr>]`

The **regular-expression** *<regular-expression>* parameter is a regular expression. The regular expressions are the same ones supported for BGP4 AS-path filters. Refer to [“Using regular expressions”](#) on page 791.

The **<address> <mask>** parameter specifies a particular route. If you also use the optional **longer-prefixes** parameter, then all statistics for routes that match the specified route or have a longer prefix than the specified route are displayed. For example, if you specify **209.157.0.0 longer**, then all routes with the prefix 209.157. or that have a longer prefix (such as 209.157.22.) are displayed.

The **neighbor <ip-addr>** parameter displays route flap dampening statistics only for routes learned from the specified neighbor. You also can display route flap statistics for routes learned from a neighbor by entering the following command: **show ip bgp neighbor <ip-addr> flap-statistics**.

This display shows the following information.

TABLE 125 Route flap dampening statistics

This field...	Displays...
Total number of flapping routes	Total number of routes in the Layer 3 Switch BGP4 route table that have changed state and thus have been marked as flapping routes.
Status code	Indicates the dampening status of the route, which can be one of the following: <ul style="list-style-type: none"> • > – This is the best route among those in the BGP4 route table to the route destination. • d – This route is currently dampened, and thus unusable. • h – The route has a history of flapping and is unreachable now. • * – The route has a history of flapping but is currently usable.
Network	The destination network of the route.
From	The neighbor that sent the route to the Layer 3 Switch.
Flaps	The number of flaps (state changes) the route has experienced.
Since	The amount of time since the first flap of this route.
Reuse	The amount of time remaining until this route will be un-suppressed and thus be usable again.
Path	Shows the AS-path information for the route.

You also can display all the dampened routes by entering the following command.

```
show ip bgp dampened-paths
```

Clearing route flap dampening statistics

To clear route flap dampening statistics, use the following CLI method.

NOTE

Clearing the dampening statistics for a route does not change the dampening status of the route.

To clear all the route dampening statistics, enter the following command at any level of the CLI.

```
PowerConnect# clear ip bgp flap-statistics
```

Syntax: `clear ip bgp flap-statistics [regular-expression <regular-expression> | <address> <mask> | neighbor <ip-addr>]`

The parameters are the same as those for the **show ip bgp flap-statistics** command (except the **longer-prefixes** option is not supported). Refer to “[Displaying route flap dampening statistics](#)” on page 814.

NOTE

The **clear ip bgp damping** command not only clears statistics but also un-suppresses the routes. Refer to “[Displaying route flap dampening statistics](#)” on page 814.

Generating traps for BGP

You can enable and disable SNMP traps for BGP. BGP traps are enabled by default.

To enable BGP traps after they have been disabled, enter the following command.

```
PowerConnect(config)# snmp-server enable traps bgp
```

Syntax: [no] snmp-server enable traps bgp

Use the **no** form of the command to disable BGP traps.

Displaying BGP4 information

You can display the following configuration information and statistics for the BGP4 protocol on the router:

- Summary BGP4 configuration information for the router
- Active BGP4 configuration information (the BGP4 information in the running-config)
- CPU utilization statistics
- Neighbor information
- Peer-group information
- Information about the paths from which BGP4 selects routes
- Summary BGP4 route information
- The router BGP4 route table
- Route flap dampening statistics
- Active route maps (the route map configuration information in the running-config)

Displaying summary BGP4 information

You can display the local AS number, the maximum number of routes and neighbors supported, and some BGP4 statistics.

To view summary BGP4 information for the router, enter the following command at any CLI prompt.

```
PowerConnect# show ip bgp summary
BGP4 Summary
Router ID: 101.0.0.1   Local AS Number : 4
Confederation Identifier : not configured
Confederation Peers: 4 5
Maximum Number of Paths Supported for Load Sharing : 1
Number of Neighbors Configured : 11
Number of Routes Installed : 2
Number of Routes Advertising to All Neighbors : 8
Number of Attribute Entries Installed : 6
Neighbor Address  AS#  State  Time      Rt:Accepted  Filtered  Sent  ToSend
1.2.3.4           200  ADMDN  0h44m56s  0            0          0     2
10.0.0.2          5    ADMDN  0h44m56s  0            0          0     0
10.1.0.2          5    ESTAB  0h44m56s  1            11         0     0
10.2.0.2          5    ESTAB  0h44m55s  1            0          0     0
10.3.0.2          5    ADMDN  0h25m28s  0            0          0     0
10.4.0.2          5    ADMDN  0h25m31s  0            0          0     0
10.5.0.2          5    CONN   0h 0m 8s  0            0          0     0
10.7.0.2          5    ADMDN  0h44m56s  0            0          0     0
100.0.0.1         4    ADMDN  0h44m56s  0            0          0     2
102.0.0.1         4    ADMDN  0h44m56s  0            0          0     2
150.150.150.150  0    ADMDN  0h44m56s  0            0          0     2
```

This display shows the following information.

TABLE 126 BGP4 summary information

This field...	Displays...
Router ID	The Layer 3 Switch router ID.
Local AS Number	The BGP4 AS number the router is in.
Confederation Identifier	The AS number of the confederation the Layer 3 Switch is in.
Confederation Peers	The numbers of the local ASs contained in the confederation. This list matches the confederation peer list you configure on the Layer 3 Switch.
Maximum Number of Paths Supported for Load Sharing	The maximum number of route paths across which the device can balance traffic to the same destination. The feature is enabled by default but the default number of paths is 1. You can increase the number from 2 – 4 paths. Refer to “Changing the maximum number of paths for BGP4 load sharing” on page 769.
Number of Neighbors Configured	The number of BGP4 neighbors configured on this Layer 3 Switch.
Number of Routes Installed	The number of BGP4 routes in the router BGP4 route table. To display the BGP4 route table, refer to “Displaying the BGP4 route table” on page 835.
Number of Routes Advertising to All Neighbors	The total of the RtSent and RtToSend columns for all neighbors.
Number of Attribute Entries Installed	The number of BGP4 route-attribute entries in the router route-attributes table. To display the route-attribute table, refer to “Displaying BGP4 route-attribute entries” on page 841.
Neighbor Address	The IP addresses of this router BGP4 neighbors.
AS#	The AS number.

TABLE 126 BGP4 summary information (Continued)

This field...	Displays...
State	<p>The state of this router neighbor session with each neighbor. The states are from this router perspective of the session, not the neighbor perspective. The state values are based on the BGP4 state machine values described in RFC 1771 and can be one of the following for each router:</p> <ul style="list-style-type: none"> • IDLE – The BGP4 process is waiting to be started. Usually, enabling BGP4 or establishing a neighbor session starts the BGP4 process. A minus sign (-) indicates that the session has gone down and the software is clearing or removing routes. • ADMND – The neighbor has been administratively shut down. Refer to “Administratively shutting down a session with a BGP4 neighbor” on page 766. A minus sign (-) indicates that the session has gone down and the software is clearing or removing routes. • CONNECT – BGP4 is waiting for the connection process for the TCP neighbor session to be completed. • ACTIVE – BGP4 is waiting for a TCP connection from the neighbor. <p>NOTE: If the state frequently changes between CONNECT and ACTIVE, there may be a problem with the TCP connection.</p> <ul style="list-style-type: none"> • OPEN SENT – BGP4 is waiting for an Open message from the neighbor. • OPEN CONFIRM – BGP4 has received an OPEN message from the neighbor and is now waiting for either a KEEPALIVE or NOTIFICATION message. If the router receives a KEEPALIVE message from the neighbor, the state changes to Established. If the message is a NOTIFICATION, the state changes to Idle. • ESTABLISHED – BGP4 is ready to exchange UPDATE packets with the neighbor. If there is more BGP data in the TCP receiver queue, a plus sign (+) is also displayed. <p>NOTE: If you display information for the neighbor using the show ip bgp neighbor <ip-addr> command, the TCP receiver queue value will be greater than 0.</p>
Time	The time that has passed since the state last changed.
Accepted	The number of routes received from the neighbor that this router installed in the BGP4 route table. Usually, this number is lower than the RoutesRcvd number. The difference indicates that this router filtered out some of the routes received in the UPDATE messages.
Filtered	<p>The routes or prefixes that have been filtered out:</p> <ul style="list-style-type: none"> • If soft reconfiguration is enabled, this field shows how many routes were filtered out (not placed in the BGP4 route table) but retained in memory. • If soft reconfiguration is not enabled, this field shows the number of BGP4 routes that have been filtered out.
Sent	The number of BGP4 routes that the Layer 3 Switch has sent to the neighbor.
ToSend	The number of routes the Layer 3 Switch has queued to send to this neighbor.

Displaying the active BGP4 configuration

To view the active BGP4 configuration information contained in the running-config without displaying the entire running-config, use the following CLI method.

To display the device active BGP4 configuration, enter the following command at any level of the CLI.


```
PowerConnect# show ip bgp config
Current BGP configuration:
router bgp
  address-filter 1 deny any any
  as-path-filter 1 permit ^65001$
  local-as 65002
  maximum-paths 4
  neighbor pgl peer-group
  neighbor pgl remote-as 65001
  neighbor pgl description "PowerConnect group 1"
  neighbor pgl distribute-list out 1
  neighbor 192.169.100.1 peer-group pgl
  neighbor 192.169.101.1 peer-group pgl
  neighbor 192.169.102.1 peer-group pgl
  neighbor 192.169.201.1 remote-as 65101
  neighbor 192.169.201.1 shutdown
  neighbor 192.169.220.3 remote-as 65432
  network 1.1.1.0 255.255.255.0
  network 2.2.2.0 255.255.255.0
  redistribute connected
```

Syntax: show ip bgp config

Displaying CPU utilization statistics

You can display CPU utilization statistics for BGP4 and other IP protocols.

To display CPU utilization statistics for BGP4 for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI.

```
PowerConnect# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP              0.01        0.03        0.09        0.22         9
BGP            0.04       0.06       0.08       0.14        13
GVRP            0.00        0.00        0.00        0.00         0
ICMP            0.00        0.00        0.00        0.00         0
IP              0.00        0.00        0.00        0.00         0
OSPF            0.00        0.00        0.00        0.00         0
RIP             0.00        0.00        0.00        0.00         0
STP             0.00        0.00        0.00        0.00         0
VRRP           0.00        0.00        0.00        0.00         0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example.

```
PowerConnect# show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.00       0.00       0.00        0
BGP             0.00       0.00       0.00       0.00        0
GVRP           0.00       0.00       0.00       0.00        0
ICMP           0.01       0.00       0.00       0.00        1
IP             0.00       0.00       0.00       0.00        0
OSPF           0.00       0.00       0.00       0.00        0
RIP            0.00       0.00       0.00       0.00        0
STP            0.00       0.00       0.00       0.00        0
VRRP           0.00       0.00       0.00       0.00        0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following.

```
PowerConnect# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name    Sec(%)    Time(ms)
ARP             0.00      0
BGP             0.00      0
GVRP           0.00      0
ICMP           0.01      1
IP             0.00      0
OSPF           0.00      0
RIP            0.00      0
STP            0.01      0
VRRP           0.00      0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: `show process cpu [<num>]`

The `<num>` parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Displaying summary neighbor information

To display summary neighbor information, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.211 routes-summary
1 IP Address: 192.168.4.211
Routes Accepted/Installed:1, Filtered/Kept:11, Filtered:11
  Routes Selected as BEST Routes:1
    BEST Routes not Installed in IP Forwarding Table:0
  Unreachable Routes (no IGP Route for NEXTHOP):0
  History Routes:0

NLRI Received in Update Message:24, Withdraws:0 (0), Replacements:1
  NLRI Discarded due to
    Maximum Prefix Limit:0, AS Loop:0
    Invalid Nexthop:0, Invalid Nexthop Address:0.0.0.0
    Duplicated Originator_ID:0, Cluster_ID:0

Routes Advertised:0, To be Sent:0, To be Withdrawn:0
NLRI Sent in Update Message:0, Withdraws:0, Replacements:0

Peer Out of Memory Count for:
  Receiving Update Messages:0, Accepting Routes(NLRI):0
  Attributes:0, Outbound Routes(RIB-out):0
```

Syntax: show ip bgp neighbors [<ip-addr>] | [route-summary]

This display shows the following information.

TABLE 127 BGP4 route summary information for a neighbor

This field...	Displays...
IP Address	The IP address of the neighbor
Routes Received	How many routes the Layer 3 Switch has received from the neighbor during the current BGP4 session: <ul style="list-style-type: none"> Accepted/Installed – Indicates how many of the received routes the Layer 3 Switch accepted and installed in the BGP4 route table. Filtered/Kept – Indicates how many routes were filtered out, but were nonetheless retained in memory for use by the soft reconfiguration feature. Filtered – Indicates how many of the received routes were filtered out.
Routes Selected as BEST Routes	The number of routes that the Layer 3 Switch selected as the best routes to their destinations.
BEST Routes not Installed in IP Forwarding Table	The number of routes received from the neighbor that are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).
Unreachable Routes	The number of routes received from the neighbor that are unreachable because the Layer 3 Switch does not have a valid RIP, OSPF, or static route to the next hop.
History Routes	The number of routes that are down but are being retained for route flap dampening purposes.
NLRIs Received in Update Message	The number of routes received in Network Layer Reachability (NLRI) format in UPDATE messages: <ul style="list-style-type: none"> Withdraws – The number of withdrawn routes the Layer 3 Switch has received. Replacements – The number of replacement routes the Layer 3 Switch has received.

TABLE 127 BGP4 route summary information for a neighbor (Continued)

This field...	Displays...
NLRIs Discarded due to	Indicates the number of times the Layer 3 Switch discarded an NLRI for the neighbor due to the following reasons: <ul style="list-style-type: none"> • Maximum Prefix Limit – The Layer 3 Switch configured maximum prefix amount had been reached. • AS Loop – An AS loop occurred. An AS loop occurs when the BGP4 AS-path attribute contains the local AS number. • Invalid Nexthop – The next hop value was not acceptable. • Duplicated Originator_ID – The originator ID was the same as the local router ID. • Cluster_ID – The cluster list contained the local cluster ID, or contained the local router ID (see above) if the cluster ID is not configured.
Routes Advertised	The number of routes the Layer 3 Switch has advertised to this neighbor: <ul style="list-style-type: none"> • To be Sent – The number of routes the Layer 3 Switch has queued to send to this neighbor. • To be Withdrawn – The number of NLRIs for withdrawing routes the Layer 3 Switch has queued up to send to this neighbor in UPDATE messages.
NLRIs Sent in Update Message	The number of NLRIs for new routes the Layer 3 Switch has sent to this neighbor in UPDATE messages: <ul style="list-style-type: none"> • Withdraws – The number of routes the Layer 3 Switch has sent to the neighbor to withdraw. • Replacements – The number of routes the Layer 3 Switch has sent to the neighbor to replace routes the neighbor already has.
Peer Out of Memory Count for	Statistics for the times the Layer 3 Switch has run out of BGP4 memory for the neighbor during the current BGP4 session: <ul style="list-style-type: none"> • Receiving Update Messages – The number of times UPDATE messages were discarded because there was no memory for attribute entries. • Accepting Routes(NLRI) – The number of NLRIs discarded because there was no memory for NLRI entries. This count is not included in the Receiving Update Messages count. • Attributes – The number of times there was no memory for BGP4 attribute entries. • Outbound Routes(RIB-out) – The number of times there was no memory to place a “best” route into the neighbor's route information base (Adj-RIB-Out) for routes to be advertised.

Displaying BGP4 neighbor information

To view BGP4 neighbor information including the values for all the configured parameters, enter the following command.

NOTE

The display shows all the configured parameters for the neighbor. Only the parameters that have values different from their defaults are shown.

```

PowerConnect# show ip bgp neighbor 10.4.0.2
1  IP Address: 10.4.0.2, AS: 5 (EBGP), RouterID: 100.0.0.1
   Description: neighbor 10.4.0.2
   State: ESTABLISHED, Time: 0h1m0s, KeepAliveTime: 0, HoldTime: 0
   PeerGroup: pgl
   Multihop-EBGP: yes, ttl: 1
   RouteReflectorClient: yes
   SendCommunity: yes
   NextHopSelf: yes
   DefaultOriginate: yes (default sent)
   MaximumPrefixLimit: 90000
   RemovePrivateAs: : yes
   RefreshCapability: Received
Route Filter Policies:
  Distribute-list: (out) 20
  Filter-list: (in) 30
  Prefix-list: (in) pfl
  Route-map: (in) setnp1 (out) setnp2
Messages:      Open      Update  KeepAlive  Notification  Refresh-Req
Sent          : 1        1        1           0              0
Received: 1    8        1           0              0
Last Update Time: NLRI          Withdraw      NLRI          Withdraw
                  Tx: 0h0m59s    ---          Rx: 0h0m59s    ---
Last Connection Reset Reason:Unknown
Notification Sent:      Unspecified
Notification Received: Unspecified
TCP Connection state: ESTABLISHED
Local host: 10.4.0.1, Local Port: 179
Remote host: 10.4.0.2, Remote Port: 8053
ISentSeq: 52837276 SendNext: 52837392 TotUnAck: 0
TotSent: 116 ReTrans: 0 UnAckSeq: 52837392
IRcvSeq: 2155052043 RcvNext: 2155052536 SendWnd: 16384
TotalRcv: 493 DupliRcv: 0 RcvWnd: 16384
SendQue: 0 RcvQue: 0 CngstWnd: 1460

```

This example shows how to display information for a specific neighbor, by specifying the neighbor IP address with the command. None of the other display options are used; thus, all of the information is displayed for the neighbor. The number in the far left column indicates the neighbor for which information is displayed. When you list information for multiple neighbors, this number makes the display easier to read.

The TCP statistics at the end of the display show status for the TCP session with the neighbor. Most of the fields show information stored in the Layer 3 Switch Transmission Control Block (TCB) for the TCP session between the Layer 3 Switch and its neighbor. These fields are described in detail in section 3.2 of RFC 793, "Transmission Control Protocol Functional Specification".

Syntax: `show ip bgp neighbors [<ip-addr> [advertised-routes [detail [<ip-addr>/<mask-bits>]]] | [attribute-entries [detail]] | [flap-statistics] | [last-packet-with-error] | [received prefix-filter] | [received-routes] | [routes [best] | [detail [best] | [not-installed-best] | [unreachable]]] | [rib-out-routes [<ip-addr>/<mask-bits> | <ip-addr> <net-mask> | detail]] | [routes-summary]]`

The `<ip-addr>` option lets you narrow the scope of the command to a specific neighbor.

The `advertised-routes` option displays only the routes that the Layer 3 Switch has advertised to the neighbor during the current BGP4 neighbor session.

The **attribute-entries** option shows the attribute-entries associated with routes received from the neighbor.

The **flap-statistics** option shows the route flap statistics for routes received from or sent to the neighbor.

The **last-packet-with-error** option displays the last packet from the neighbor that contained an error. The packet's contents are displayed in decoded (human-readable) format.

The **received prefix-filter** option shows the Outbound Route Filters (ORFs) received from the neighbor. This option applies to cooperative route filtering.

The **received-routes option** lists all the route information received in route updates from the neighbor since the soft reconfiguration feature was enabled. Refer to [“Using soft reconfiguration”](#) on page 845.

The **routes** option lists the routes received in UPDATE messages from the neighbor. You can specify the following additional options:

- **best** – Displays the routes received from the neighbor that the Layer 3 Switch selected as the best routes to their destinations.
- **not-installed-best** – Displays the routes received from the neighbor that are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).
- **unreachable** – Displays the routes that are unreachable because the Layer 3 Switch does not have a valid RIP, OSPF, or static route to the next hop.
- **detail** – Displays detailed information for the specified routes. You can refine your information request by also specifying one of the options above (**best**, **not-installed-best**, or **unreachable**).

The **rib-out-routes** option lists the route information base (RIB) for outbound routes. You can display all the routes or specify a network address.

The **routes-summary** option displays a summary of the following information:

- Number of routes received from the neighbor
- Number of routes accepted by this Layer 3 Switch from the neighbor
- Number of routes this Layer 3 Switch filtered out of the UPDATES received from the neighbor and did not accept
- Number of routes advertised to the neighbor
- Number of attribute entries associated with routes received from or advertised to the neighbor.

This display shows the following information.

TABLE 128 BGP4 neighbor information

This field...	Displays...
IP Address	The IP address of the neighbor.
AS	The AS the neighbor is in.
EBGP/IBGP	Whether the neighbor session is an IBGP session, an EBGP session, or a confederation EBGP session: <ul style="list-style-type: none"> • EBGP – The neighbor is in another AS. • EBGP_Confed – The neighbor is a member of another sub-AS in the same confederation. • IBGP – The neighbor is in the same AS.

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
RouterID	The neighbor router ID.
Description	The description you gave the neighbor when you configured it on the Layer 3 Switch.
State	<p>The state of the router session with the neighbor. The states are from this router perspective of the session, not the neighbor perspective. The state values are based on the BGP4 state machine values described in RFC 1771 and can be one of the following for each router:</p> <ul style="list-style-type: none"> • IDLE – The BGP4 process is waiting to be started. Usually, enabling BGP4 or establishing a neighbor session starts the BGP4 process. A minus sign (-) indicates that the session has gone down and the software is clearing or removing routes. • ADMND – The neighbor has been administratively shut down. Refer to “Administratively shutting down a session with a BGP4 neighbor” on page 766. A minus sign (-) indicates that the session has gone down and the software is clearing or removing routes. • CONNECT – BGP4 is waiting for the connection process for the TCP neighbor session to be completed. • ACTIVE – BGP4 is waiting for a TCP connection from the neighbor. <p>NOTE: If the state frequently changes between CONNECT and ACTIVE, there may be a problem with the TCP connection.</p> <ul style="list-style-type: none"> • OPEN SENT – BGP4 is waiting for an Open message from the neighbor. • OPEN CONFIRM – BGP4 has received an OPEN message from the neighbor and is now waiting for either a KEEPALIVE or NOTIFICATION message. If the router receives a KEEPALIVE message from the neighbor, the state changes to Established. If the message is a NOTIFICATION, the state changes to Idle. • ESTABLISHED – BGP4 is ready to exchange UPDATE messages with the neighbor. If there is more BGP data in the TCP receiver queue, a plus sign (+) is also displayed. <p>NOTE: If you display information for the neighbor using the show ip bgp neighbor <ip-addr> command, the TCP receiver queue value will be greater than 0.</p>
Time	The amount of time this session has been in its current state.
KeepAliveTime	The keep alive time, which specifies how often this router sends keep alive messages to the neighbor. Refer to “Changing the Keep Alive Time and Hold Time” on page 767.
HoldTime	The hold time, which specifies how many seconds the router will wait for a KEEPALIVE or UPDATE message from a BGP4 neighbor before deciding that the neighbor is dead. Refer to “Changing the Keep Alive Time and Hold Time” on page 767.
PeerGroup	The name of the peer group the neighbor is in, if applicable.
Multihop-EBGP	Whether this option is enabled for the neighbor.
RouteReflectorClient	Whether this option is enabled for the neighbor.
SendCommunity	Whether this option is enabled for the neighbor.
NextHopSelf	Whether this option is enabled for the neighbor.
DefaultOriginate	Whether this option is enabled for the neighbor.
MaximumPrefixLimit	Lists the maximum number of prefixes the Layer 3 Switch will accept from this neighbor.

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
RemovePrivateAs	Whether this option is enabled for the neighbor.
RefreshCapability	Whether this Layer 3 Switch has received confirmation from the neighbor that the neighbor supports the dynamic refresh capability.
CooperativeFilteringCapability	Whether the neighbor is enabled for cooperative route filtering.
Distribute-list	Lists the distribute list parameters, if configured.
Filter-list	Lists the filter list parameters, if configured.
Prefix-list	Lists the prefix list parameters, if configured.
Route-map	Lists the route map parameters, if configured.
Messages Sent	The number of messages this router has sent to the neighbor. The display shows statistics for the following message types: <ul style="list-style-type: none"> • Open • Update • KeepAlive • Notification • Refresh-Req
Messages Received	The number of messages this router has received from the neighbor. The message types are the same as for the Message Sent field.
Last Update Time	Lists the last time updates were sent and received for the following: <ul style="list-style-type: none"> • NLRIs • Withdraws

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
Last Connection Reset Reason	<p>The reason the previous session with this neighbor ended. The reason can be one of the following.</p> <p>Reasons described in the BGP specifications:</p> <ul style="list-style-type: none"> • Message Header Error • Connection Not Synchronized • Bad Message Length • Bad Message Type • OPEN Message Error • Unsupported Version Number • Bad Peer AS Number • Bad BGP Identifier • Unsupported Optional Parameter • Authentication Failure • Unacceptable Hold Time • Unsupported Capability • UPDATE Message Error • Malformed Attribute List • Unrecognized Well-known Attribute • Missing Well-known Attribute • Attribute Flags Error • Attribute Length Error • Invalid ORIGIN Attribute • Invalid NEXT_HOP Attribute • Optional Attribute Error • Invalid Network Field • Malformed AS_PATH • Hold Timer Expired • Finite State Machine Error • Rcv Notification
Last Connection Reset Reason (cont.)	<p>Reasons specific to the Dell implementation:</p> <ul style="list-style-type: none"> • Reset All Peer Sessions • User Reset Peer Session • Port State Down • Peer Removed • Peer Shutdown • Peer AS Number Change • Peer AS Confederation Change • TCP Connection KeepAlive Timeout • TCP Connection Closed by Remote • TCP Data Stream Error Detected

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
Notification Sent	<p>If the router receives a NOTIFICATION message from the neighbor, the message contains an error code corresponding to one of the following errors. Some errors have subcodes that clarify the reason for the error. Where applicable, the subcode messages are listed underneath the error code messages.</p> <p>Message Header Error:</p> <ul style="list-style-type: none"> • Connection Not Synchronized • Bad Message Length • Bad Message Type • Unspecified <p>Open Message Error:</p> <ul style="list-style-type: none"> • Unsupported Version • Bad Peer As • Bad BGP Identifier • Unsupported Optional Parameter • Authentication Failure • Unacceptable Hold Time • Unspecified <p>Update Message Error:</p> <ul style="list-style-type: none"> • Malformed Attribute List • Unrecognized Attribute • Missing Attribute • Attribute Flag Error • Attribute Length Error • Invalid Origin Attribute • Invalid NextHop Attribute • Optional Attribute Error • Invalid Network Field • Malformed AS Path • Unspecified <p>Hold Timer Expired</p> <p>Finite State Machine Error</p> <p>Cease</p> <p>Unspecified</p>
Notification Received	See above.

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
TCP Connection state	The state of the connection with the neighbor. The connection can have one of the following states: <ul style="list-style-type: none"> • LISTEN – Waiting for a connection request. • SYN-SENT – Waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED – Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED – Data can be sent and received over the connection. This is the normal operational state of the connection. • FIN-WAIT-1 – Waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 – Waiting for a connection termination request from the remote TCP. • CLOSE-WAIT – Waiting for a connection termination request from the local user. • CLOSING – Waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK – Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request). • TIME-WAIT – Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request. • CLOSED – There is no connection state.
Byte Sent	The number of bytes sent.
Byte Received	The number of bytes received.
Local host	The IP address of the Layer 3 Switch.
Local port	The TCP port the Layer 3 Switch is using for the BGP4 TCP session with the neighbor.
Remote host	The IP address of the neighbor.
Remote port	The TCP port the neighbor is using for the BGP4 TCP session with the Layer 3 Switch.
ISentSeq	The initial send sequence number for the session.
SendNext	The next sequence number to be sent.
TotUnAck	The number of sequence numbers sent by the Layer 3 Switch that have not been acknowledged by the neighbor.
TotSent	The number of sequence numbers sent to the neighbor.
ReTrans	The number of sequence numbers that the Layer 3 Switch retransmitted because they were not acknowledged.
UnAckSeq	The current acknowledged sequence number.
IRcvSeq	The initial receive sequence number for the session.
RcvNext	The next sequence number expected from the neighbor.
SendWnd	The size of the send window.
TotalRcv	The number of sequence numbers received from the neighbor.
DupliRcv	The number of duplicate sequence numbers received from the neighbor.

TABLE 128 BGP4 neighbor information (Continued)

This field...	Displays...
RcvWnd	The size of the receive window.
SendQue	The number of sequence numbers in the send queue.
RcvQue	The number of sequence numbers in the receive queue.
CngstWnd	The number of times the window has changed.

Displaying route information for a neighbor

You can display routes based on the following criteria:

- A summary of the routes for a specific neighbor.
- The routes received from the neighbor that the Layer 3 Switch selected as the best routes to their destinations.
- The routes received from the neighbor that are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).
- The routes that are unreachable because the Layer 3 Switch does not have a valid RIP, OSPF, or static route to the next hop.
- Routes for a specific network advertised by the Layer 3 Switch to the neighbor.
- The Routing Information Base (RIB) for a specific network advertised to the neighbor. You can display the RIB regardless of whether the Layer 3 Switch has already sent it to the neighbor.

To display route information for a neighbor, use the following CLI methods.

Displaying summary route information

To display summary route information, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 10.1.0.2 routes-summary
1  IP Address: 10.1.0.2
Routes Accepted/Installed:1,  Filtered/Kept:11,  Filtered:11
  Routes Selected as BEST Routes:1
    BEST Routes not Installed in IP Forwarding Table:0
  Unreachable Routes (no IGP Route for NEXTHOP):0
  History Routes:0

NLRI's Received in Update Message:24,  Withdraws:0 (0),  Replacements:1
  NLRI's Discarded due to
    Maximum Prefix Limit:0,  AS Loop:0
    Invalid Nexthop:0,  Invalid Nexthop Address:0.0.0.0
    Duplicated Originator_ID:0,  Cluster_ID:0

Routes Advertised:0,  To be Sent:0,  To be Withdrawn:0
NLRI's Sent in Update Message:0,  Withdraws:0,  Replacements:0

Peer Out of Memory Count for:
  Receiving Update Messages:0,  Accepting Routes(NLRI):0
  Attributes:0,  Outbound Routes(RIB-out):0
```

This display shows the following information.

TABLE 129 BGP4 route summary information for a neighbor

This field...	Displays...
Routes Received	How many routes the Layer 3 Switch has received from the neighbor during the current BGP4 session: <ul style="list-style-type: none"> • Accepted/Installed – Indicates how many of the received routes the Layer 3 Switch accepted and installed in the BGP4 route table. • Filtered – Indicates how many of the received routes the Layer 3 Switch did not accept or install because they were denied by filters on the Layer 3 Switch.
Routes Selected as BEST Routes	The number of routes that the Layer 3 Switch selected as the best routes to their destinations.
BEST Routes not Installed in IP Forwarding Table	The number of routes received from the neighbor that are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).
Unreachable Routes	The number of routes received from the neighbor that are unreachable because the Layer 3 Switch does not have a valid RIP, OSPF, or static route to the next hop.
History Routes	The number of routes that are down but are being retained for route flap dampening purposes.
NLRIs Received in Update Message	The number of routes received in Network Layer Reachability (NLRI) format in UPDATE messages: <ul style="list-style-type: none"> • Withdraws – The number of withdrawn routes the Layer 3 Switch has received. • Replacements – The number of replacement routes the Layer 3 Switch has received.
NLRIs Discarded due to	Indicates the number of times the Layer 3 Switch discarded an NLRI for the neighbor due to the following reasons: <ul style="list-style-type: none"> • Maximum Prefix Limit – The Layer 3 Switch configured maximum prefix amount had been reached. • AS Loop – An AS loop occurred. An AS loop occurs when the BGP4 AS-path attribute contains the local AS number. • Invalid Nexthop – The next hop value was not acceptable. • Duplicated Originator_ID – The originator ID was the same as the local router ID. • Cluster_ID – The cluster list contained the local cluster ID, or contained the local router ID (see above) if the cluster ID is not configured.
Routes Advertised	The number of routes the Layer 3 Switch has advertised to this neighbor: <ul style="list-style-type: none"> • To be Sent – The number of routes the Layer 3 Switch has queued to send to this neighbor. • To be Withdrawn – The number of NLRIs for withdrawing routes the Layer 3 Switch has queued up to send to this neighbor in UPDATE messages.

TABLE 129 BGP4 route summary information for a neighbor (Continued)

This field...	Displays...
NLRIs Sent in Update Message	The number of NLRIs for new routes the Layer 3 Switch has sent to this neighbor in UPDATE messages: <ul style="list-style-type: none"> • Withdraws – The number of routes the Layer 3 Switch has sent to the neighbor to withdraw. • Replacements – The number of routes the Layer 3 Switch has sent to the neighbor to replace routes the neighbor already has.
Peer Out of Memory Count for	Statistics for the times the Layer 3 Switch has run out of BGP4 memory for the neighbor during the current BGP4 session: <ul style="list-style-type: none"> • Receiving Update Messages – The number of times UPDATE messages were discarded because there was no memory for attribute entries. • Accepting Routes(NLRI) – The number of NLRIs discarded because there was no memory for NLRI entries. This count is not included in the Receiving Update Messages count. • Attributes – The number of times there was no memory for BGP4 attribute entries. • Outbound Routes(RIB-out) – The number of times there was no memory to place a “best” route into the neighbor’s route information base (Adj-RIB-Out) for routes to be advertised.

Displaying advertised routes

To display the routes the Layer 3 Switch has advertised to a specific neighbor for a specific network, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbors 192.168.4.211 advertised-routes
      There are 2 routes advertised to neighbor 192.168.4.211
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST I:IBGP L:LOCAL
      Network      Next Hop      Metric      LocPrf      Weight      Status
1      102.0.0.0/24   192.168.2.102  12          32768       BL
2      200.1.1.0/24   192.168.2.102  0          32768       BL
```

You also can enter a specific route, as in the following example.

```
PowerConnect# show ip bgp neighbors 192.168.4.211 advertised 200.1.1.0/24
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST I:IBGP L:LOCAL
      Network      Next Hop      Metric      LocPrf      Weight      Status
1      200.1.1.0/24   192.168.2.102  0          32768       BL
```

Syntax: `show ip bgp neighbor <ip-addr> advertised-routes [<ip-addr>/<prefix>]`

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the `show ip bgp` display.

Displaying the best routes

To display the routes received from a specific neighbor that are the “best” routes to their destinations, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.211 routes best
```

Syntax: `show ip bgp neighbor <ip-addr> routes best`

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the `show ip bgp` display.

Displaying the best routes that were nonetheless not installed in the IP route table

To display the BGP4 routes received from a specific neighbor that are the “best” routes to their destinations but are not installed in the Layer 3 Switch IP route table, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.211 routes not-installed-best
```

Each of the displayed routes is a valid path to its destination, but the Layer 3 Switch received another path from a different source (such as OSPF, RIP, or a static route) that has a lower administrative distance. The Layer 3 Switch always selects the path with the lowest administrative distance to install in the IP route table.

Syntax: `show ip bgp neighbor <ip-addr> routes not-installed-best`

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the `show ip bgp` display.

Displaying the routes whose destinations are unreachable

To display BGP4 routes whose destinations are unreachable using any of the BGP4 paths in the BGP4 route table, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.211 routes unreachable
```

Syntax: `show ip bgp neighbor <ip-addr> routes unreachable`

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the `show ip bgp` display.

Displaying the Adj-RIB-Out for a neighbor

To display the Layer 3 Switch current BGP4 Routing Information Base (Adj-RIB-Out) for a specific neighbor and a specific destination network, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.211 rib-out-routes 192.168.1.0/24
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST I:IBGP L:LOCAL
      Prefix          Next Hop      Metric      LocPrf      Weight Status
1      200.1.1.0/24    0.0.0.0      0           101         32768 BL
```

The Adj-RIB-Out contains the routes that the Layer 3 Switch either has most recently sent to the neighbor or is about to send to the neighbor.

Syntax: `show ip bgp neighbor <ip-addr> rib-out-routes [<ip-addr>/<prefix>]`

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the `show ip bgp` display.

Displaying peer group information

You can display configuration information for peer groups.

To display peer-group information, enter a command such as the following at the Privileged EXEC level of the CLI.

```
PowerConnect# show ip bgp peer-group pg1
1  BGP peer-group is pg
   Description: peer group abc
   SendCommunity: yes
   NextHopSelf: yes
   DefaultOriginate: yes
   Members:
     IP Address: 192.168.10.10, AS: 65111
```

Syntax: `show ip bgp peer-group [<peer-group-name>]`

Only the parameters that have values different from their defaults are listed.

Displaying summary route information

To display summary statistics for all the routes in the Layer 3 Switch BGP4 route table, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp routes summary
Total number of BGP routes (NLRIs) Installed      : 20
Distinct BGP destination networks                 : 20
Filtered BGP routes for soft reconfig             : 100178
Routes originated by this router                  : 2
Routes selected as BEST routes                   : 19
BEST routes not installed in IP forwarding table  : 1
Unreachable routes (no IGP route for NEXTHOP)    : 1
IBGP routes selected as best routes              : 0
EBGP routes selected as best routes              : 17
```

Syntax: `show ip bgp routes summary`

This display shows the following information.

TABLE 130 BGP4 summary route information

This field...	Displays...
Total number of BGP routes (NLRIs) Installed	The number of BGP4 routes the Layer 3 Switch has installed in the BGP4 route table.
Distinct BGP destination networks	The number of destination networks the installed routes represent. The BGP4 route table can have multiple routes to the same network.
Filtered BGP routes for soft reconfig	The number of route updates received from soft-reconfigured neighbors or peer groups that have been filtered out but retained. For information about soft reconfiguration, refer to “Using soft reconfiguration” on page 845.
Routes originated by this router	The number of routes in the BGP4 route table that this Layer 3 Switch originated.
Routes selected as BEST routes	The number of routes in the BGP4 route table that this Layer 3 Switch has selected as the best routes to the destinations.
BEST routes not installed in IP forwarding table	The number of BGP4 routes that are the best BGP4 routes to their destinations but were not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).
Unreachable routes (no IGP route for NEXTHOP)	The number of routes in the BGP4 route table whose destinations are unreachable because the next hop is unreachable.

TABLE 130 BGP4 summary route information (Continued)

This field...	Displays...
IBGP routes selected as best routes	The number of “best” routes in the BGP4 route table that are IBGP routes.
EBGP routes selected as best routes	The number of “best” routes in the BGP4 route table that are EBGP routes.

Displaying the BGP4 route table

BGP4 uses filters you define as well as the algorithm described in “How BGP4 selects a path for a route” on page 747 to determine the preferred route to a destination. BGP4 sends only the preferred route to the router IP table. However, if you want to view all the routes BGP4 knows about, you can display the BGP4 table using either of the following methods.

To view the BGP4 route table, enter the following command.

```
PowerConnect# show ip bgp routes
Total number of BGP Routes: 97371
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
      E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
Prefix          Next Hop          Metric      LocPrf      Weight  Status
1      3.0.0.0/8          192.168.4.106          100          0      BE
      AS_PATH: 65001 4355 701 80
2      4.0.0.0/8          192.168.4.106          100          0      BE
      AS_PATH: 65001 4355 1
3      4.60.212.0/22      192.168.4.106          100          0      BE
      AS_PATH: 65001 4355 701 1 189
4      6.0.0.0/8          192.168.4.106          100          0      BE
      AS_PATH: 65001 4355 3356 7170 1455
5      8.8.1.0/24          192.168.4.106          0            100          0      BE
      AS_PATH: 65001
```

Syntax: `show ip bgp routes` [[network] <ip-addr>] | <num> | [age <secs>] | [as-path-access-list <num>] | [best] | [cidr-only] | [community <num>] | no-export | no-advertise | internet | local-as | [community-access-list <num>] | [community-list <num>] | [detail <option>] | [filter-list <num, num,...>] | [next-hop <ip-addr>] | [no-best] | [not-installed-best] | [prefix-list <string>] | [regular-expression <regular-expression>] | [route-map <map-name>] | [summary] | [unreachable]

The <ip-addr> option displays routes for a specific network. The **network** keyword is optional. You can enter the network address without entering “network” in front of it.

The <num> option specifies the table entry with which you want the display to start. For example, if you want to list entries beginning with table entry 100, specify 100.

The **age** <secs> parameter displays only the routes that have been received or updated more recently than the number of seconds you specify.

The **as-path-access-list** <num> parameter filters the display using the specified AS-path ACL.

The **best** parameter displays the routes received from the neighbor that the Layer 3 Switch selected as the best routes to their destinations.

The **cidr-only** option lists only the routes whose network masks do not match their class network length.

The **community** option lets you display routes for a specific community. You can specify **local-as**, **no-export**, **no-advertise**, **internet**, or a private community number. You can specify the community number as either two five-digit integer values of up to 1– 65535, separated by a colon (for example, 12345:6789) or a single long integer value.

The **community-access-list** *<num>* parameter filters the display using the specified community ACL.

The **community-list** option lets you display routes that match a specific community filter.

The **detail** option lets you display more details about the routes. You can refine your request by also specifying one of the other display options after the detail keyword.

The **filter-list** option displays routes that match a specific address filter list.

The **next-hop** *<ip-addr>* option displays the routes for a given next-hop IP address.

The **no-best** option displays the routes for which none of the routes to a given prefix were selected as the best route. The IP route table does not contain a BGP4 route for any of the routes listed by the command.

The **not-installed-best** option displays the routes received from the neighbor that are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes).

The **prefix-list** *<string>* parameter filters the display using the specified IP prefix list.

The **regular-expression** *<regular-expression>* option filters the display based on a regular expression. Refer to “Using regular expressions” on page 791.

The **route-map** *<map-name>* parameter filters the display using the specified route map. The software displays only the routes that match the match statements in the route map. The software disregards the route map set statements.

The **summary** option displays summary information for the routes.

The **unreachable** option displays the routes that are unreachable because the Layer 3 Switch does not have a valid RIP, OSPF, or static route to the next hop.

Displaying the best BGP4 routes

To display all the BGP4 routes in the Layer 3 Switch BGP4 route table that are the best routes to their destinations, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp routes best
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
      E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
      Prefix           Next Hop           Metric           LocPrf           Weight Status
1      3.0.0.0/8         192.168.4.106     0                100              0          BE
      AS_PATH: 65001 4355 701 80
2      4.0.0.0/8         192.168.4.106     0                100              0          BE
      AS_PATH: 65001 4355 1
3      4.60.212.0/22     192.168.4.106     0                100              0          BE
      AS_PATH: 65001 4355 701 1 189
4      6.0.0.0/8         192.168.4.106     0                100              0          BE
      AS_PATH: 65001 4355 3356 7170 1455
5      9.2.0.0/16        192.168.4.106     0                100              0          BE
      AS_PATH: 65001 4355 701
```

Syntax: show ip bgp routes best

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the **show ip bgp** display.

Displaying those best BGP4 routes that are nonetheless not in the IP route table

When the Layer 3 Switch has multiple routes to a destination from different sources (such as BGP4, OSPF, RIP, or static routes), the Layer 3 Switch selects the route with the lowest administrative distance as the best route, and installs that route in the IP route table.

To display the BGP4 routes are the “best” routes to their destinations but are not installed in the Layer 3 Switch IP route table, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp routes not-installed-best
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
        E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
        Prefix          Next Hop          Metric          LocPrf          Weight Status
1         192.168.4.0/24    192.168.4.106    0                100             0        bE
        AS_PATH: 65001
```

Each of the displayed routes is a valid path to its destination, but the Layer 3 Switch received another path from a different source (such as OSPF, RIP, or a static route) that has a lower administrative distance. The Layer 3 Switch always selects the path with the lowest administrative distance to install in the IP route table.

Notice that the route status in this example is the new status, “b”. Refer to [Table 131](#) on page 838 for a description.

Syntax: show ip bgp routes not-installed-best

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the **show ip bgp** display.

NOTE

To display the routes that the Layer 3 Switch has selected as the best routes and installed in the IP route table, display the IP route table using the **show ip route** command.

Displaying BGP4 routes whose destinations are unreachable

To display BGP4 routes whose destinations are unreachable using any of the BGP4 paths in the BGP4 route table, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp routes unreachable
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
        H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
        Prefix          Next Hop          Metric          LocPrf          Weight Status
1         8.8.8.0/24    192.168.5.1      0                101             0
        AS_PATH: 65001 4355 1
```

Syntax: show ip bgp routes unreachable

For information about the fields in this display, refer to [Table 131](#) on page 838. The fields in this display also appear in the **show ip bgp** display.

Displaying information for a specific route

To display BGP4 network information by specifying an IP address within the network, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp 9.3.4.0
Number of BGP Routes matching display condition : 1
Status codes: s suppressed, d damped, h history, * valid, > best, i internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*> 9.3.4.0/24       192.168.4.106      100    0      65001 4355 1 1221 ?
   Last update to IP routing table: 0h11m38s, 1 path(s) installed:
     Gateway        Port
     192.168.2.1    1
   Route is advertised to 1 peers:
     20.20.20.2(65300)
```

Syntax: `show ip bgp [route] <ip-addr>/<prefix> [longer-prefixes] | <ip-addr>`

If you use the **route** option, the display for the information is different, as shown in the following example.

```
PowerConnect# show ip bgp route 9.3.4.0
Number of BGP Routes matching display condition : 1
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
   Prefix          Next Hop          Metric LocPrf Weight Status
1    9.3.4.0/24     192.168.4.106      100    0      BE
   AS_PATH: 65001 4355 1 1221
   Last update to IP routing table: 0h12m1s, 1 path(s) installed:
     Gateway        Port
     192.168.2.1    1
   Route is advertised to 1 peers:
     20.20.20.2(65300)
```

These displays show the following information.

TABLE 131 BGP4 network information

This field...	Displays...
Number of BGP Routes matching display condition	The number of routes that matched the display parameters you entered. This is the number of routes displayed by the command.
Status codes	A list of the characters the display uses to indicate the route status. The status code appears in the left column of the display, to the left of each route. The status codes are described in the command output. NOTE: This field appears only if you <i>do not</i> enter the route option.
Prefix	The network address and prefix.
Next Hop	The next-hop router for reaching the network from the Layer 3 Switch.
Metric	The value of the route MED attribute. If the route does not have a metric, this field is blank.
LocPrf	The degree of preference for this route relative to other routes in the local AS. When the BGP4 algorithm compares routes on the basis of local preferences, the route with the higher local preference is chosen. The preference can have a value from 0 - 4294967295.

TABLE 131 BGP4 network information (Continued)

This field...	Displays...
Weight	The value that this router associates with routes from a specific neighbor. For example, if the router receives routes to the same destination from two BGP4 neighbors, the router prefers the route from the neighbor with the larger weight.
Path	The route AS path. NOTE: This field appears only if you <i>do not</i> enter the route option.
Origin code	A character the display uses to indicate the route origin. The origin code appears to the right of the AS path (Path field). The origin codes are described in the command output. NOTE: This field appears only if you <i>do not</i> enter the route option.
Status	The route status, which can be one or more of the following: <ul style="list-style-type: none"> • A – AGGREGATE. The route is an aggregate route for multiple networks. • B – BEST. BGP4 has determined that this is the optimal route to the destination. NOTE: If the “b” is shown in lowercase, the software was not able to install the route in the IP route table. <ul style="list-style-type: none"> • b – NOT-INSTALLED-BEST. The routes received from the neighbor are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes). • C – CONFED_EBGP. The route was learned from a neighbor in the same confederation and AS, but in a different sub-AS within the confederation. • D – DAMPED. This route has been dampened (by the route dampening feature), and is currently unusable. • H – HISTORY. Route dampening is configured for this route, and the route has a history of flapping and is unreachable now. • I – INTERNAL. The route was learned through BGP4. • L – LOCAL. The route originated on this Layer 3 Switch. • M – MULTIPATH. BGP4 load sharing is enabled and this route was selected as one of the best ones to the destination. The best route among the multiple paths also is marked with “B”. NOTE: If the “m” is shown in lowercase, the software was not able to install the route in the IP route table. <ul style="list-style-type: none"> • S – SUPPRESSED. This route was suppressed during aggregation and thus is not advertised to neighbors. NOTE: This field appears only if you enter the route option.

Displaying route details

Here is an example of the information displayed when you use the **detail** option. In this example, the information for one route is shown.

```
PowerConnect# show ip bgp routes detail
Total number of BGP Routes: 2
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
1 Prefix: 10.5.0.0/24, Status: BME, Age: 0h28m28s
NEXT_HOP: 201.1.1.2, Learned from Peer: 10.1.0.2 (5)
LOCAL_PREF: 101, MED: 0, ORIGIN: igp, Weight: 10
AS_PATH: 5
Adj_RIB_out count: 4, Admin distance 20
```

These displays show the following information.

TABLE 132 BGP4 route information

This field...	Displays...
Total number of BGP Routes	The number of BGP4 routes.
Status codes	A list of the characters the display uses to indicate the route status. The status code is appears in the left column of the display, to the left of each route. The status codes are described in the command output.
Prefix	The network prefix and mask length.
Status	<p>The route status, which can be one or more of the following:</p> <ul style="list-style-type: none"> • A – AGGREGATE. The route is an aggregate route for multiple networks. • B – BEST. BGP4 has determined that this is the optimal route to the destination. <p>NOTE: If the “b” is shown in lowercase, the software was not able to install the route in the IP route table.</p> <ul style="list-style-type: none"> • b – NOT-INSTALLED-BEST. The routes received from the neighbor are the best BGP4 routes to their destinations, but were nonetheless not installed in the IP route table because the Layer 3 Switch received better routes from other sources (such as OSPF, RIP, or static IP routes). • C – CONFED_EBGP. The route was learned from a neighbor in the same confederation and AS, but in a different sub-AS within the confederation. • D – DAMPED. This route has been dampened (by the route dampening feature), and is currently unusable. • H – HISTORY. Route dampening is configured for this route, and the route has a history of flapping and is unreachable now. • I – INTERNAL. The route was learned through BGP4. • L – LOCAL. The route originated on this Layer 3 Switch. • M – MULTIPATH. BGP4 load sharing is enabled and this route was selected as one of the best ones to the destination. The best route among the multiple paths also is marked with “B”. <p>NOTE: If the “m” is shown in lowercase, the software was not able to install the route in the IP route table.</p> <ul style="list-style-type: none"> • S – SUPPRESSED. This route was suppressed during aggregation and thus is not advertised to neighbors.
Age	The last time an update occurred.
Next_Hop	The next-hop router for reaching the network from the Layer 3 Switch.
Learned from Peer	The IP address of the neighbor that sent this route.
Local_Pref	The degree of preference for this route relative to other routes in the local AS. When the BGP4 algorithm compares routes on the basis of local preferences, the route with the higher local preference is chosen. The preference can have a value from 0 – 4294967295.
MED	The route metric. If the route does not have a metric, this field is blank.
Origin	<p>The source of the route information. The origin can be one of the following:</p> <ul style="list-style-type: none"> • EGP – The routes with this set of attributes came to BGP through EGP. • IGP – The routes with this set of attributes came to BGP through IGP. • INCOMPLETE – The routes came from an origin other than one of the above. For example, they may have been redistributed from OSPF or RIP. <p>When BGP4 compares multiple routes to a destination to select the best route, IGP is preferred over EGP and both are preferred over INCOMPLETE.</p>

TABLE 132 BGP4 route information (Continued)

This field...	Displays...
Weight	The value that this router associates with routes from a specific neighbor. For example, if the router receives routes to the same destination from two BGP4 neighbors, the router prefers the route from the neighbor with the larger weight.
Atomic	Whether network information in this route has been aggregated <i>and</i> this aggregation has resulted in information loss. NOTE: Information loss under these circumstances is a normal part of BGP4 and does not indicate an error.
Aggregation ID	The router that originated this aggregator.
Aggregation AS	The AS in which the network information was aggregated. This value applies only to aggregated routes and is otherwise 0.
Originator	The originator of the route in a route reflector environment.
Cluster List	The route-reflector clusters through which this route has passed.
Learned From	The IP address of the neighbor from which the Layer 3 Switch learned the route.
Admin Distance	The administrative distance of the route.
Adj_RIB_out	The number of neighbors to which the route has been or will be advertised. This is the number of times the route has been selected as the best route and placed in the Adj-RIB-Out (outbound queue) for a BGP4 neighbor.
Communities	The communities the route is in.

Displaying BGP4 route-attribute entries

The route-attribute entries table lists the sets of BGP4 attributes stored in the router memory. Each set of attributes is unique and can be associated with one or more routes. In fact, the router typically has fewer route attribute entries than routes. To display the route-attribute entries table, use one of the following methods.

To display the IP route table, enter the following command.

```
PowerConnect# show ip bgp attribute-entries
```

Syntax: show ip bgp attribute-entries

Here is an example of the information displayed by this command. A zero value indicates that the attribute is not set.

```
PowerConnect# show ip bgp attribute-entries
      Total number of BGP Attribute Entries: 7753
1      Next Hop :192.168.11.1      Metric :0      Origin:IGP
      Originator:0.0.0.0      Cluster List:None
      Aggregator:AS Number :0      Router-ID:0.0.0.0      Atomic:FALSE
      Local Pref:100      Communities:Internet
      AS Path : (65002) 65001 4355 2548 3561 5400 6669 5548
2      Next Hop :192.168.11.1      Metric :0      Origin:IGP
      Originator:0.0.0.0      Cluster List:None
      Aggregator:AS Number :0      Router-ID:0.0.0.0      Atomic:FALSE
      Local Pref:100      Communities:Internet
      AS Path : (65002) 65001 4355 2548
```

This display shows the following information.

TABLE 133 BGP4 route-attribute entries information

This field...	Displays...
Total number of BGP Attribute Entries	The number of routes contained in this router BGP4 route table.
Next Hop	The IP address of the next hop router for routes that have this set of attributes.
Metric	The cost of the routes that have this set of attributes.
Origin	The source of the route information. The origin can be one of the following: <ul style="list-style-type: none"> • EGP – The routes with this set of attributes came to BGP through EGP. • IGP – The routes with this set of attributes came to BGP through IGP. • INCOMPLETE – The routes came from an origin other than one of the above. For example, they may have been redistributed from OSPF or RIP. When BGP4 compares multiple routes to a destination to select the best route, IGP is preferred over EGP and both are preferred over INCOMPLETE.
Originator	The originator of the route in a route reflector environment.
Cluster List	The route-reflector clusters through which this set of attributes has passed.
Aggregator	Aggregator information: <ul style="list-style-type: none"> • AS Number shows the AS in which the network information in the attribute set was aggregated. This value applies only to aggregated routes and is otherwise 0. • Router-ID shows the router that originated this aggregator.
Atomic	Whether the network information in this set of attributes has been aggregated <i>and</i> this aggregation has resulted in information loss: <ul style="list-style-type: none"> • TRUE – Indicates information loss has occurred • FALSE – Indicates no information loss has occurred NOTE: Information loss under these circumstances is a normal part of BGP4 and does not indicate an error.
Local Pref	The degree of preference for routes that use this set of attributes relative to other routes in the local AS.
Communities	The communities that routes with this set of attributes are in.
AS Path	The ASs through which routes with this set of attributes have passed. The local AS is shown in parentheses.

Displaying the routes BGP4 has placed in the IP route table

The IP route table indicates the routes it has received from BGP4 by listing “BGP” as the route type.

To display the IP route table, enter the following command.

```
PowerConnect# show ip route
```

Syntax: `show ip route [<ip-addr> | <num> | bgp | ospf | rip]`

Here is an example of the information displayed by this command. Notice that most of the routes in this example have type “B”, indicating that their source is BGP4.


```
PowerConnect# show ip route
Total number of IP routes: 50834
B:BGP D:Directly-Connected O:OSPF R:RIP S:Static
  Network Address  NetMask          Gateway          Port    Cost    Type
  3.0.0.0          255.0.0.0        192.168.13.2    1       0       B
  4.0.0.0          255.0.0.0        192.168.13.2    1       0       B
  9.20.0.0         255.255.128.0    192.168.13.2    1       0       B
  10.1.0.0         255.255.0.0      0.0.0.0 0       1       D
  10.10.11.0       255.255.255.0    0.0.0.0 0       1       D
  12.2.97.0        255.255.255.0    192.168.13.2    1       0       B
  12.3.63.0        255.255.255.0    192.168.13.2    1       0       B
  12.3.123.0       255.255.255.0    192.168.13.2    1       0       B
  12.5.252.0       255.255.254.0    192.168.13.2    1       0       B
  12.6.42.0        255.255.254.0    192.168.13.2    1       0       B
remaining 50824 entries not shown...
```

Displaying route flap dampening statistics

To display route dampening statistics or all the dampened routes, enter the following command at any level of the CLI.

```
PowerConnect# show ip bgp flap-statistics
Total number of flapping routes: 414
  Status Code  >:best d:damped h:history *:valid
  Network      From          Flaps Since    Reuse    Path
h> 192.50.206.0/23 166.90.213.77 1 0 :0 :13 0 :0 :0 65001 4355 1 701
h> 203.255.192.0/20 166.90.213.77 1 0 :0 :13 0 :0 :0 65001 4355 1 7018
h> 203.252.165.0/24 166.90.213.77 1 0 :0 :13 0 :0 :0 65001 4355 1 7018
h> 192.50.208.0/23 166.90.213.77 1 0 :0 :13 0 :0 :0 65001 4355 1 701
h> 133.33.0.0/16 166.90.213.77 1 0 :0 :13 0 :0 :0 65001 4355 1 701
*> 204.17.220.0/24 166.90.213.77 1 0 :1 :4 0 :0 :0 65001 4355 701 62
```

Syntax: `show ip bgp flap-statistics [regular-expression <regular-expression> | <address> <mask> [longer-prefixes] | neighbor <ip-addr> | filter-list <num>...]`

The **regular-expression** <regular-expression> parameter is a regular expression. The regular expressions are the same ones supported for BGP4 AS-path filters. Refer to [“Using regular expressions”](#) on page 791.

The <address> <mask> parameter specifies a particular route. If you also use the optional **longer-prefixes** parameter, then all statistics for routes that match the specified route or have a longer prefix than the specified route are displayed. For example, if you specify **209.157.0.0 longer**, then all routes with the prefix 209.157 or that have a longer prefix (such as 209.157.22) are displayed.

The **neighbor** <ip-addr> parameter displays route flap dampening statistics only for routes learned from the specified neighbor. You also can display route flap statistics for routes learned from a neighbor by entering the following command: **show ip bgp neighbor <ip-addr> flap-statistics**.

The **filter-list** <num> parameter specifies one or more filters. Only the routes that have been dampened and that match the specified filters are displayed.

This display shows the following information.

TABLE 134 Route flap dampening statistics

This field...	Displays...
Total number of flapping routes	The total number of routes in the Layer 3 Switch BGP4 route table that have changed state and thus have been marked as flapping routes.
Status code	Indicates the dampening status of the route, which can be one of the following: <ul style="list-style-type: none"> • > - This is the best route among those in the BGP4 route table to the route destination. • d - This route is currently dampened, and thus unusable. • h - The route has a history of flapping and is unreachable now. • * - The route has a history of flapping but is currently usable.
Network	The destination network of the route.
From	The neighbor that sent the route to the Layer 3 Switch.
Flaps	The number of flaps (state changes) the route has experienced.
Since	The amount of time since the first flap of this route.
Reuse	The amount of time remaining until this route will be un-suppressed and thus be usable again.
Path	Shows the AS-path information for the route.

You also can display all the dampened routes by entering the following command.

```
show ip bgp dampened-paths.
```

Displaying the active route map configuration

To view the device active route map configuration (contained in the running-config) without displaying the entire running-config, enter the following command at any level of the CLI.

```
PowerConnect# show route-map
route-map permitnet4 permit 10
  match ip address prefix-list plist1
route-map permitnet1 permit 1
  match ip address prefix-list plist2
route-map setcomm permit 1
  set community 1234:2345 no-export
route-map test111 permit 111
  match address-filters 11
  set community 11:12 no-export
route-map permit1122 permit 12
  match ip address 11
route-map permit1122 permit 13
  match ip address std_22
```

This example shows that the running-config contains six route maps. Notice that the match and set statements within each route map are listed beneath the command for the route map itself. In this simplified example, each route map contains only one match or set statement.

To display the active configuration for a specific route map, enter a command such as the following, which specifies a route map name.

```
PowerConnect# show route-map setcomm
route-map setcomm permit 1
  set community 1234:2345 no-export
```

This example shows the active configuration for a route map called “setcomm”.

Syntax: `show route-map [<map-name>]`

Updating route information and resetting a neighbor session

The following sections describe ways to update route information with a neighbor, reset the session with a neighbor, and close a session with a neighbor.

Whenever you change a policy (ACL, route map, and so on) that affects the routes that the Layer 3 Switch learns from a BGP4 neighbor or peer group of neighbors, you must enter a command to place the changes into effect. The changes take place automatically, but only affect new route updates. To make changes retroactive for routes received or sent before the changes were made, you need to enter a clear command.

You can update the learned routes using either of the following methods:

- Request the complete BGP4 route table from the neighbor or peer group. You can use this method if the neighbor supports the refresh capability (RFCs 2842 and 2858).
- Clear (reset) the session with the neighbor or peer group. This is the only method you can use if the neighbor does not support the refresh capability.

Each of these methods is effective, but can be disruptive to the network. The first method adds overhead while the Layer 3 Switch learns and filters the neighbor or group entire route table, while the second method adds more overhead while the devices re-establish their BGP4 sessions.

You also can clear and reset the BGP4 routes that have been installed in the IP route table. Refer to [“Clearing and resetting BGP4 routes in the IP route table”](#) on page 851.

Using soft reconfiguration

The **soft reconfiguration** feature places policy changes into effect without resetting the BGP4 session. Soft reconfiguration does not request the neighbor or group to send its entire BGP4 table, nor does the feature reset the session with the neighbor or group. Instead, the soft reconfiguration feature stores all the route updates received from the neighbor or group. When you request a soft reset of inbound routes, the software performs route selection by comparing the policies against the stored route updates, instead of requesting the neighbor BGP4 route table or resetting the session with the neighbor.

When you enable the soft reconfiguration feature, it sends a refresh message to the neighbor or group if the neighbor or group supports dynamic refresh. Otherwise, the feature resets the neighbor session. This step is required to ensure that the soft reconfiguration feature has a complete set of updates to use, and occurs only once, when you enable the feature. The feature accumulates all the route updates from the neighbor, eliminating the need for additional refreshes or resets when you change policies in the future.

To use soft reconfiguration:

- Enable the feature.
- Make the policy changes.
- Apply the changes by requesting a soft reset of the inbound updates from the neighbor or group.

Use the following CLI methods to configure soft configuration, apply policy changes, and display information for the updates that are filtered out by the policies.

Enabling soft reconfiguration

To configure a neighbor for soft reconfiguration, enter a command such as the following.

```
PowerConnect(config-bgp-router)# neighbor 10.10.200.102 soft-reconfiguration inbound
```

This command enables soft reconfiguration for updates received from 10.10.200.102. The software dynamically refreshes or resets the session with the neighbor, then retains all route updates from the neighbor following the reset.

Syntax: [no] neighbor <ip-addr> | <peer-group-name> soft-reconfiguration inbound

NOTE

The syntax related to soft reconfiguration is shown. For complete command syntax, refer to [“Adding BGP4 neighbors”](#) on page 756.

Placing a policy change into effect

To place policy changes into effect, enter a command such as the following.

```
PowerConnect(config-bgp-router)# clear ip bgp neighbor 10.10.200.102 soft in
```

This command updates the routes by comparing the route policies against the route updates that the Layer 3 Switch has stored. The command does not request additional updates from the neighbor or otherwise affect the session with the neighbor.

Syntax: clear ip bgp neighbor <ip-addr> | <peer-group-name> soft in

NOTE

If you do not specify “in”, the command applies to both inbound and outbound updates.

NOTE

The syntax related to soft reconfiguration is shown. For complete command syntax, refer to [“Dynamically refreshing routes”](#) on page 848.

Displaying the filtered routes received from the neighbor or peer group

When you enable soft reconfiguration, the Layer 3 Switch saves all updates received from the specified neighbor or peer group. This includes updates that contain routes that are filtered out by the BGP4 route policies in effect on the Layer 3 Switch. To display the routes that have been filtered out, enter the following command at any level of the CLI.

```
PowerConnect# show ip bgp filtered-routes
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
Prefix          Next Hop          Metric      LocPrf      Weight Status
1      3.0.0.0/8        192.168.4.106
   AS_PATH: 65001 4355 701 80
2      4.0.0.0/8        192.168.4.106
   AS_PATH: 65001 4355 1
3      4.60.212.0/22   192.168.4.106
   AS_PATH: 65001 4355 701 1 189
           100          0          EF
```

The routes displayed by the command are the routes that the Layer 3 Switch BGP4 policies filtered out. The Layer 3 Switch did not place the routes in the BGP4 route table, but did keep the updates. If a policy change causes these routes to be permitted, the Layer 3 Switch does not need to request the route information from the neighbor, but instead uses the information in the updates.

Syntax: `show ip bgp filtered-routes [<ip-addr>] | [as-path-access-list <num>] | [detail] | [prefix-list <string>]`

The *<ip-addr>* parameter specifies the IP address of the destination network.

The *as-path-access-list <num>* parameter specifies an AS-path ACL. Only the routes permitted by the AS-path ACL are displayed.

The *detail* parameter displays detailed information for the routes. (The example above shows summary information.) You can specify any of the other options after *detail* to further refine the display request.

The *prefix-list <string>* parameter specifies an IP prefix list. Only the routes permitted by the prefix list are displayed.

NOTE

The syntax for displaying filtered routes is shown. For complete command syntax, refer to [“Displaying the BGP4 route table”](#) on page 835.

Displaying all the routes received from the neighbor

To display all the route information received in route updates from a neighbor since you enabled soft reconfiguration, enter a command such as the following at any level of the CLI.

```
PowerConnect# show ip bgp neighbor 192.168.4.106 received-routes
There are 97345 received routes from neighbor 192.168.4.106
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED F:FILTERED
Prefix          Next Hop          Metric      LocPrf      Weight Status
1      3.0.0.0/8        192.168.4.106
   AS_PATH: 65001 4355 701 80
2      4.0.0.0/8        192.168.4.106
   AS_PATH: 65001 4355 1
3      4.60.212.0/22   192.168.4.106
   AS_PATH: 65001 4355 701 1 189
4      6.0.0.0/8        192.168.4.106
           100          0          BE
```

Syntax: `show ip bgp neighbors <ip-addr> received-routes [detail]`

The **detail** parameter displays detailed information for the routes. The example above shows summary information.

NOTE

The syntax for displaying received routes is shown. For complete command syntax, refer to [“Displaying BGP4 neighbor information”](#) on page 822.

NOTE

The **show ip bgp neighbor <ip-addr> received-routes** syntax supported in previous software releases is changed to the following syntax: **show ip bgp neighbor <ip-addr> routes**.

Dynamically requesting a route refresh from a BGP4 neighbor

You can easily apply changes to filters that control BGP4 routes received from or advertised to a neighbor, without resetting the BGP4 session between the Layer 3 Switch and the neighbor. For example, if you add, change, or remove a BGP4 address filter that denies specific routes received from a neighbor, you can apply the filter change by requesting a route refresh from the neighbor. If the neighbor also supports dynamic route refreshes, the neighbor resends its Adj-RIB-Out, its table of BGP4 routes. Using the route refresh feature, you do not need to reset the session with the neighbor.

The route refresh feature is based on the following specifications:

- RFC 2842. This RFC specifies the Capability Advertisement, which a BGP4 router uses to dynamically negotiate a capability with a neighbor.
- RFC 2858 for Multi-protocol Extension.

NOTE

The Dell implementation of dynamic route refresh supports negotiation of IP version 4 unicasts only.

- RFC 2918, which describes the dynamic route refresh capability

The dynamic route refresh capability is enabled by default and cannot be disabled. When the Layer 3 Switch sends a BGP4 OPEN message to a neighbor, the Layer 3 Switch includes a Capability Advertisement to inform the neighbor that the Layer 3 Switch supports dynamic route refresh.

NOTE

The option for dynamically refreshing routes received from a neighbor requires the neighbor to support dynamic route refresh. If the neighbor does not support this feature, the option does not take effect and the software displays an error message. The option for dynamically re-advertising routes to a neighbor does not require the neighbor to support dynamic route refresh.

To use the dynamic refresh feature, use either of the following methods.

Dynamically refreshing routes

The following sections describe how to dynamically refresh BGP4 routes to place new or changed filters into effect.

To request a dynamic refresh of all routes from a neighbor, enter a command such as the following.

```
PowerConnect(config-bgp-router)# clear ip bgp neighbor 192.168.1.170 soft in
```

This command asks the neighbor to send its BGP4 table (Adj-RIB-Out) again. The Layer 3 Switch applies its filters to the incoming routes and adds, modifies, or removes BGP4 routes as necessary.

Syntax: `clear ip bgp neighbor all | <ip-addr> | <peer-group-name> | <as-num> [soft-outbound | soft [in | out]]`

The `all | <ip-addr> | <peer-group-name> | <as-num>` specifies the neighbor. The `<ip-addr>` parameter specifies a neighbor by its IP interface with the Layer 3 Switch. The `<peer-group-name>` specifies all neighbors in a specific peer group. The `<as-num>` parameter specifies all neighbors within the specified AS. The `all` parameter specifies all neighbors.

The `soft-outbound` parameter updates all outbound routes by applying the new or changed filters, but sends only the existing routes affected by the new or changed filters to the neighbor.

The `soft [in | out]` parameter specifies whether you want to refresh the routes received from the neighbor or sent to the neighbor:

- **soft in** does one of the following:
 - If you enabled soft reconfiguration for the neighbor or peer group, **soft in** updates the routes by comparing the route policies against the route updates that the Layer 3 Switch has stored. Soft reconfiguration does not request additional updates from the neighbor or otherwise affect the session with the neighbor. Refer to [“Using soft reconfiguration”](#) on page 845.
 - If you did not enable soft reconfiguration, **soft in** requests the neighbor entire BGP4 route table (Adj-RIB-Out), then applies the filters to add, change, or exclude routes.
 - If a neighbor does not support dynamic refresh, **soft in** resets the neighbor session.
- **soft out** updates all outbound routes, then sends the Layer 3 Switch entire BGP4 route table (Adj-RIB-Out) to the neighbor, after changing or excluding the routes affected by the filters.

If you do not specify `in` or `out`, the Layer 3 Switch performs both options.

NOTE

The `soft-outbound` parameter updates all outbound routes by applying the new or changed filters, but sends only the existing routes affected by the new or changed filters to the neighbor. The `soft out` parameter updates all outbound routes, then sends the Layer 3 Switch entire BGP4 route table (Adj-RIB-Out) to the neighbor, after changing or excluding the routes affected by the filters. Use `soft-outbound` if only the outbound policy is changed.

To dynamically resend all the Layer 3 Switch BGP4 routes to a neighbor, enter a command such as the following.

```
PowerConnect(config-bgp-router)# clear ip bgp neighbor 192.168.1.170 soft out
```

This command applies its filters for outgoing routes to the Layer 3 Switch BGP4 route table (Adj-RIB-Out), changes or excludes routes accordingly, then sends the resulting Adj-RIB-Out to the neighbor.

NOTE

The Layer 3 Switch does not automatically update outbound routes using a new or changed outbound policy or filter when a session with the neighbor goes up or down. Instead, the Layer 3 Switch applies a new or changed policy or filter when a route is placed in the outbound queue (Adj-RIB-Out).

To place a new or changed outbound policy or filter into effect, you must enter a **clear ip bgp neighbor** command regardless of whether the neighbor session is up or down. You can enter the command without optional parameters or with the **soft out** or **soft-outbound** option. Either way, you must specify a parameter for the neighbor (*<ip-addr>*, *<as-num>*, *<peer-group-name>*, or **all**).

Displaying dynamic refresh information

You can use the **show ip bgp neighbors** command to display information for dynamic refresh requests. For each neighbor, the display lists the number of dynamic refresh requests the Layer 3 Switch has sent to or received from the neighbor and indicates whether the Layer 3 Switch received confirmation from the neighbor that the neighbor supports dynamic route refresh.

The RefreshCapability field indicates whether this Layer 3 Switch has received confirmation from the neighbor that the neighbor supports the dynamic refresh capability. The statistics in the Message Sent and Message Received rows under Refresh-Req indicate how many dynamic refreshes have been sent to and received from the neighbor. The statistic is cumulative across sessions.

```
PowerConnect# show ip bgp neighbor 10.4.0.2
1  IP Address: 10.4.0.2, AS: 5 (EBGP), RouterID: 100.0.0.1
   Description: neighbor 10.4.0.2
   State: ESTABLISHED, Time: 0h1m0s, KeepAliveTime: 0, HoldTime: 0
   PeerGroup: pgl
   Mutihop-EBGP: yes, ttl: 1
   RouteReflectorClient: yes
   SendCommunity: yes
   NextHopSelf: yes
   DefaultOriginate: yes (default sent)
   MaximumPrefixLimit: 90000
   RemovePrivateAs: : yes
   RefreshCapability: Received
Route Filter Policies:
  Distribute-list: (out) 20
  Filter-list: (in) 30
  Prefix-list: (in) pfl
  Route-map: (in) setnp1 (out) setnp2
Messages:   Open      Update  KeepAlive  Notification  Refresh-Req
  Sent      : 1         1         1           0              0
  Received: 1         8         1           0              0
Last Update Time: NLRI          Withdraw          NLRI          Withdraw
                  Tx: 0h0m59s    ---              Rx: 0h0m59s    ---
Last Connection Reset Reason:Unknown
Notification Sent:      Unspecified
Notification Received: Unspecified
TCP Connection state: ESTABLISHED
Byte Sent: 115, Received: 492
Local host: 10.4.0.1, Local Port: 179
Remote host: 10.4.0.2, Remote Port: 8053
ISentSeq: 52837276  SendNext: 52837392  TotUnAck: 0
TotSent: 116  ReTrans: 0  UnAckSeq: 52837392
IRcvSeq: 2155052043  RcvNext: 2155052536  SendWnd: 16384
TotalRcv: 493  DupliRcv: 0  RcvWnd: 16384
SendQue: 0  RcvQue: 0  CngstWnd: 1460
```


Closing or resetting a neighbor session

You can close a neighbor session or resend route updates to a neighbor.

If you make changes to filters or route maps and the neighbor does not support dynamic route refresh, use the following methods to ensure that neighbors contain only the routes you want them to contain:

- If you close a neighbor session, the Layer 3 Switch and the neighbor clear all the routes they learned from each other. When the Layer 3 Switch and neighbor establish a new BGP4 session, they exchange route tables again. Use this method if you want the Layer 3 Switch to relearn routes from the neighbor and resend its own route table to the neighbor.
- If you use the soft-outbound option, the Layer 3 Switch compiles a list of all the routes it would normally send to the neighbor at the beginning of a session. However, before sending the updates, the Layer 3 Switch also applies the filters and route maps you have configured to the list of routes. If the filters or route maps result in changes to the list of routes, the Layer 3 Switch sends updates to advertise, change, or even withdraw routes on the neighbor as needed. This ensures that the neighbor receives only the routes you want it to contain. Even if the neighbor already contains a route learned from the Layer 3 Switch that you later decided to filter out, using the soft-outbound option removes that route from the neighbor.

You can specify a single neighbor or a peer group.

To close a neighbor session and thus flush all the routes exchanged by the Layer 3 Switch and the neighbor, enter the following command.

```
PowerConnect# clear ip bgp neighbor all
```

Syntax: `clear ip bgp neighbor all | <ip-addr> | <peer-group-name> | <as-num> [soft-outbound | soft [in | out]]`

The `all | <ip-addr> | <peer-group-name> | <as-num>` specifies the neighbor. The `<ip-addr>` parameter specifies a neighbor by its IP interface with the Layer 3 Switch. The `<peer-group-name>` specifies all neighbors in a specific peer group. The `<as-num>` parameter specifies all neighbors within the specified AS. The `all` parameter specifies all neighbors.

To resend routes to a neighbor without closing the neighbor session, enter a command such as the following.

```
PowerConnect# clear ip bgp neighbor 10.0.0.1 soft out
```

Clearing and resetting BGP4 routes in the IP route table

To clear BGP4 routes from the IP route table and reset the routes, enter a command such as the following.

```
PowerConnect# clear ip bgp routes
```

Syntax: `clear ip bgp routes [<ip-addr>/<prefix-length>]`

NOTE

The `clear ip bgp routes` command has the same effect as the `clear ip route` command, but applies only to routes that come from BGP4.

Clearing traffic counters

You can clear the counters (reset them to 0) for BGP4 messages. To do so, use one of the following methods.

To clear the BGP4 message counter for all neighbors, enter the following command.

```
PowerConnect# clear ip bgp traffic
```

Syntax: clear ip bgp traffic

To clear the BGP4 message counter for a specific neighbor, enter a command such as the following.

```
PowerConnect# clear ip bgp neighbor 10.0.0.1 traffic
```

To clear the BGP4 message counter for all neighbors within a peer group, enter a command such as the following.

```
PowerConnect# clear ip bgp neighbor PeerGroup1 traffic
```

Syntax: clear ip bgp neighbor all | <ip-addr> | <peer-group-name> | <as-num> traffic

The **all** | <ip-addr> | <peer-group-name> | <as-num> specifies the neighbor. The <ip-addr> parameter specifies a neighbor by its IP interface with the Layer 3 Switch. The <peer-group-name> specifies all neighbors in a specific peer group. The <as-num> parameter specifies all neighbors within the specified AS. The **all** parameter specifies all neighbors.

Clearing route flap dampening statistics

To clear route flap dampening statistics, use the following CLI method.

NOTE

Clearing the dampening statistics for a route does not change the dampening status of the route.

To clear all the route dampening statistics, enter the following command at any level of the CLI.

```
PowerConnect# clear ip bgp flap-statistics
```

Syntax: clear ip bgp flap-statistics [regular-expression <regular-expression> | <address> <mask> | neighbor <ip-addr>]

The parameters are the same as those for the **show ip bgp flap-statistics** command (except the **longer-prefixes** option is not supported). Refer to “[Displaying route flap dampening statistics](#)” on page 814.

NOTE

The **clear ip bgp damping** command not only clears statistics but also un-suppresses the routes. Refer to “[Displaying route flap dampening statistics](#)” on page 814.

Removing route flap dampening

You can un-suppress routes by removing route flap dampening from the routes. The Layer 3 Switch allows you to un-suppress all routes at once or un-suppress individual routes.

To un-suppress all the suppressed routes, enter the following command at the Privileged EXEC level of the CLI.

```
PowerConnect# clear ip bgp damping
```

Syntax: `clear ip bgp damping [<ip-addr> <ip-mask>]`

The `<ip-addr>` parameter specifies a particular network.

The `<ip-mask>` parameter specifies the network mask.

To un-suppress a specific route, enter a command such as the following.

```
PowerConnect# clear ip bgp damping 209.157.22.0 255.255.255.0
```

This command un-suppresses only the routes for network 209.157.22.0/24.

Clearing diagnostic buffers

The Layer 3 Switch stores the following BGP4 diagnostic information in buffers:

- The first 400 bytes of the last packet that contained an error
- The last NOTIFICATION message either sent or received by the Layer 3 Switch

To display these buffers, use options with the `show ip bgp neighbors` command. Refer to [“Displaying BGP4 neighbor information”](#) on page 822.

This information can be useful if you are working with Dell Technical Support to resolve a problem. The buffers do not identify the system time when the data was written to the buffer. If you want to ensure that diagnostic data in a buffer is recent, you can clear the buffers. You can clear the buffers for a specific neighbor or for all neighbors.

If you clear the buffer containing the first 400 bytes of the last packet that contained errors, all the bytes are changed to zeros. The Last Connection Reset Reason field of the BGP neighbor table also is cleared.

If you clear the buffer containing the last NOTIFICATION message sent or received, the buffer contains no data.

You can clear the buffers for all neighbors, for an individual neighbor, or for all the neighbors within a specific peer group.

To clear these buffers for neighbor 10.0.0.1, enter the following commands.

```
PowerConnect# clear ip bgp neighbor 10.0.0.1 last-packet-with-error  
PowerConnect# clear ip bgp neighbor 10.0.0.1 notification-errors
```

Syntax: `clear ip bgp neighbor all | <ip-addr> | <peer-group-name> | <as-num>
last-packet-with-error | notification-errors`

The `all | <ip-addr> | <peer-group-name> | <as-num>` specifies the neighbor. The `<ip-addr>` parameter specifies a neighbor by its IP interface with the Layer 3 Switch. The `<peer-group-name>` specifies all neighbors in a specific peer group. The `<as-num>` parameter specifies all neighbors within the specified AS. The `all` parameter specifies all neighbors.

25 Clearing diagnostic buffers

Securing Access to Management Functions

This chapter explains how to secure access to management functions on a device.

NOTE

For all devices, RADIUS Challenge is supported for 802.1x authentication but not for login authentication. Also, multiple challenges are supported for TACACS+ login authentication.

Securing access methods

The following table lists the management access methods available on a device, how they are secured by default, and the ways in which they can be secured.

TABLE 135 Ways to secure management access to devices

Access method	How the access method is secured by default	Ways to secure the access method	See page
Serial access to the CLI	Not secured	Establish passwords for management privilege levels	page 866
Access to the Privileged EXEC and CONFIG levels of the CLI	Not secured	Establish a password for Telnet access to the CLI	page 866
		Establish passwords for management privilege levels	page 866
		Set up local user accounts	page 870
		Configure TACACS/TACACS+ security	page 877
		Configure RADIUS security	page 892

TABLE 135 Ways to secure management access to devices (Continued)

Access method	How the access method is secured by default	Ways to secure the access method	See page
Telnet access	Not secured	Regulate Telnet access using ACLs	page 857
		Allow Telnet access only from specific IP addresses	page 860
		Restrict Telnet access based on a client MAC address	page 861
		Allow Telnet access only from specific MAC addresses	page 862
		Specify the maximum number of login attempts for Telnet access	page 861
		Disable Telnet access	page 864
		Establish a password for Telnet access	page 866
		Establish passwords for privilege levels of the CLI	page 866
		Set up local user accounts	page 870
		Configure TACACS/TACACS+ security	page 877
Configure RADIUS security	page 892		
Secure Shell (SSH) access	Not configured	Configure SSH	page 945
		Regulate SSH access using ACLs	page 858
		Allow SSH access only from specific IP addresses	page 860
		Allow SSH access only from specific MAC addresses	page 861
		Establish passwords for privilege levels of the CLI	page 866
		Set up local user accounts	page 870
		Configure TACACS/TACACS+ security	page 877
		Configure RADIUS security	page 892
SNMP (Brocade Network Advisor) access	SNMP read or read-write community strings and the password to the Super User privilege level NOTE: SNMP read or read-write community strings are always required for SNMP access to the device.	Regulate SNMP access using ACLs	page 858
		Allow SNMP access only from specific IP addresses	page 860
		Disable SNMP access	page 865
		Allow SNMP access only to clients connected to a specific VLAN	page 862
		Establish passwords to management levels of the CLI	page 866
		Set up local user accounts	page 870
		Establish SNMP read or read-write community strings	page 877

TABLE 135 Ways to secure management access to devices (Continued)

Access method	How the access method is secured by default	Ways to secure the access method	See page
TFTP access	Not secured	Allow TFTP access only to clients connected to a specific VLAN	page 862
		Disable TFTP access	page 865

Restricting remote access to management functions

You can restrict access to management functions from remote sources, including Telnet, , and SNMP. The following methods for restricting remote access are supported:

- Using ACLs to restrict Telnet, or SNMP access
- Allowing remote access only from specific IP addresses
- Allowing Telnet and SSH access only from specific MAC addresses
- Allowing remote access only to clients connected to a specific VLAN
- Specifically disabling Telnet, or SNMP access to the device

The following sections describe how to restrict remote access to a device using these methods.

Using ACLs to restrict remote access

You can use standard ACLs to control the following access methods to management functions on a device:

- Telnet
- SSH
- SNMP

Consider the following to configure access control for these management access methods.

1. Configure an ACL with the IP addresses you want to allow to access the device.
2. Configure a Telnet access group, SSH access group, and SNMP community strings. Each of these configuration items accepts an ACL as a parameter. The ACL contains entries that identify the IP addresses that can use the access method.

The following sections present examples of how to secure management access using ACLs. Refer to [Chapter 13, “Configuring Rule-Based IP Access Control Lists”](#) for more information on configuring ACLs.

Using an ACL to restrict Telnet access

To configure an ACL that restricts Telnet access to the device, enter commands such as the following.

```
PowerConnect(config)# access-list 10 deny host 209.157.22.32 log
PowerConnect(config)# access-list 10 deny 209.157.23.0 0.0.0.255 log
PowerConnect(config)# access-list 10 deny 209.157.24.0 0.0.0.255 log
PowerConnect(config)# access-list 10 deny 209.157.25.0/24 log
PowerConnect(config)# access-list 10 permit any
PowerConnect(config)# telnet access-group 10
PowerConnect(config)# write memory
```

Syntax: `telnet access-group <num>`

The `<num>` parameter specifies the number of a standard ACL and must be from 1 – 99.

The commands above configure ACL 10, then apply the ACL as the access list for Telnet access. The device allows Telnet access to all IP addresses except those listed in ACL 10.

To configure a more restrictive ACL, create permit entries and omit the **permit any** entry at the end of the ACL.

Example

```
PowerConnect(config)# access-list 10 permit host 209.157.22.32
PowerConnect(config)# access-list 10 permit 209.157.23.0 0.0.0.255
PowerConnect(config)# access-list 10 permit 209.157.24.0 0.0.0.255
PowerConnect(config)# access-list 10 permit 209.157.25.0/24
PowerConnect(config)# telnet access-group 10
PowerConnect(config)# write memory
```

The ACL in this example permits Telnet access only to the IP addresses in the **permit** entries and denies Telnet access from all other IP addresses.

Using an ACL to restrict SSH access

To configure an ACL that restricts SSH access to the device, enter commands such as the following.

```
PowerConnect(config)# access-list 12 deny host 209.157.22.98 log
PowerConnect(config)# access-list 12 deny 209.157.23.0 0.0.0.255 log
PowerConnect(config)# access-list 12 deny 209.157.24.0/24 log
PowerConnect(config)# access-list 12 permit any
PowerConnect(config)# ssh access-group 12
PowerConnect(config)# write memory
```

Syntax: `ssh access-group <num>`

The `<num>` parameter specifies the number of a standard ACL and must be from 1 – 99.

These commands configure ACL 12, then apply the ACL as the access list for SSH access. The device denies SSH access from the IP addresses listed in ACL 12 and permits SSH access from all other IP addresses. Without the last ACL entry for permitting all packets, this ACL would deny SSH access from all IP addresses.

NOTE

In this example, the command **ssh access-group 10** could have been used to apply the ACL configured in the example for Telnet access. You can use the same ACL multiple times.

Using ACLs to restrict SNMP access

To restrict SNMP access to the device using ACLs, enter commands such as the following.

NOTE

The syntax for using ACLs for SNMP access is different from the syntax for controlling Telnet and SSH access using ACLs.

```
PowerConnect(config)# access-list 25 deny host 209.157.22.98 log
PowerConnect(config)# access-list 25 deny 209.157.23.0 0.0.0.255 log
PowerConnect(config)# access-list 25 deny 209.157.24.0 0.0.0.255 log
PowerConnect(config)# access-list 25 permit any
PowerConnect(config)# access-list 30 deny 209.157.25.0 0.0.0.255 log
PowerConnect(config)# access-list 30 deny 209.157.26.0/24 log
PowerConnect(config)# access-list 30 permit any
PowerConnect(config)# snmp-server community public ro 25
PowerConnect(config)# snmp-server community private rw 30
PowerConnect(config)# write memory
```

Syntax: `snmp-server community <string> ro | rw <num>`

The *<string>* parameter specifies the SNMP community string the user must enter to gain SNMP access.

The **ro** parameter indicates that the community string is for read-only (“get”) access. The **rw** parameter indicates the community string is for read-write (“set”) access.

The *<num>* parameter specifies the number of a standard ACL and must be from 1 – 99.

These commands configure ACLs 25 and 30, then apply the ACLs to community strings.

ACL 25 is used to control read-only access using the “public” community string. ACL 30 is used to control read-write access using the “private” community string.

NOTE

When **snmp-server community** is configured, all incoming SNMP packets are validated first by their community strings and then by their bound ACLs.

Defining the console idle time

By default, a device does not time out serial console sessions. A serial session remains open indefinitely until you close it. You can however define how many minutes a serial management session can remain idle before it is timed out.

NOTE

You must enable AAA support for console commands, AAA authentication, and Exec authorization in order to set the console idle time.

To configure the idle time for a serial console session, use the following command.

```
PowerConnect(config)# console timeout 120
```

Syntax: `[no] console timeout <0 – 240>`

Possible values: 0 – 240 minutes

Default value: 0 minutes (no timeout)

NOTE

In RADIUS, the standard attribute Idle-Timeout is used to define the console session timeout value. The attribute Idle-Timeout value is specified in seconds. Within the switch, it is truncated to the nearest minute, because the switch configuration is defined in minutes.

Restricting remote access to the device to specific IP addresses

By default, a device does not control remote management access based on the IP address of the managing device. You can restrict remote management access to a single IP address for the following access methods:

- Telnet access
- SSH access
- SNMP access

In addition, you can restrict all access methods to the same IP address using a single command.

The following examples show the CLI commands for restricting remote access. You can specify only one IP address with each command. However, you can enter each command ten times to specify up to ten IP addresses.

Restricting Telnet access to a specific IP address

To allow Telnet access to the device only to the host with IP address 209.157.22.39, enter the following command.

```
PowerConnect(config)# telnet-client 209.157.22.39
```

Syntax: [no] telnet-client <ip-addr> | <ipv6-addr>

Restricting SSH access to a specific IP address

To allow SSH access to the device only to the host with IP address 209.157.22.39, enter the following command.

```
PowerConnect(config)# ip ssh client 209.157.22.39
```

Syntax: [no] ip ssh client <ip-addr> | <ipv6-addr>

Restricting SNMP access to a specific IP address

To allow SNMP access (which includes Brocade Network Advisor) to the device only to the host with IP address 209.157.22.14, enter the following command.

```
PowerConnect(config)# snmp-client 209.157.22.14
```

Syntax: [no] snmp-client <ip-addr> | <ipv6-addr>

Restricting all remote management access to a specific IP address

To allow Telnet, and SNMP management access to the device only to the host with IP address 209.157.22.69, enter three separate commands (one for each access type) or enter the following command.

```
PowerConnect(config)# all-client 209.157.22.69
```

Syntax: [no] all-client <ip-addr> | <ipv6-addr>

Restricting access to the device based on IP or MAC address

You can restrict remote management access to the device, using Telnet, SSH, HTTP, and HTTPS, based on the connecting client IP or MAC address.

Restricting Telnet connection

You can restrict Telnet connection to a device based on the client IP address or MAC address.

To allow Telnet access to the device only to the host with IP address 209.157.22.39 **and** MAC address 0007.e90f.e9a0, enter the following command.

```
PowerConnect(config)# telnet client 209.157.22.39 0007.e90f.e9a0
```

Syntax: [no] telnet client <ip-addr> | <ipv6-addr> <mac-addr>

The following command allows Telnet access to the device to a host with any IP address and MAC address 0007.e90f.e9a0.

```
PowerConnect(config)# telnet client any 0007.e90f.e9a0
```

Syntax: [no] telnet client any <mac-addr>

Restricting SSH connection

You can restrict SSH connection to a device based on the client IP address or MAC address.

To allow SSH access to the device only to the host with IP address 209.157.22.39 **and** MAC address 0007.e90f.e9a0, enter the following command.

```
PowerConnect(config)# ip ssh client 209.157.22.39 0007.e90f.e9a0
```

Syntax: [no] ip ssh client <ip-addr> | <ipv6-addr> <mac-addr>

To allow SSH access to the device to a host with any IP address and MAC address 0007.e90f.e9a0, enter the following command.

```
PowerConnect(config)# ip ssh client any 0007.e90f.e9a0
```

Syntax: [no] ip ssh client any <mac-addr>

Specifying the maximum number of login attempts for Telnet access

If you are connecting to the device using Telnet, the device prompts you for a username and password. By default, you have up to 4 chances to enter a correct username and password. If you do not enter a correct username or password after 4 attempts, the device disconnects the Telnet session.

You can specify the number of attempts a Telnet user has to enter a correct username and password before the device disconnects the Telnet session. For example, to allow a Telnet user up to 5 chances to enter a correct username and password, enter the following command.

```
PowerConnect(config)# telnet login-retries 5
```

Syntax: [no] telnet login-retries <number>

You can specify from 0 – 5 attempts. The default is 4 attempts.

Restricting remote access to the device to specific VLAN IDs

You can restrict management access to a device to ports within a specific port-based VLAN. VLAN-based access control applies to the following access methods:

- Telnet access
- SNMP access
- TFTP access

By default, access is allowed for all the methods listed above on all ports. Once you configure security for a given access method based on VLAN ID, access to the device using that method is restricted to only the ports within the specified VLAN.

VLAN-based access control works in conjunction with other access control methods. For example, suppose you configure an ACL to permit Telnet access only to specific client IP addresses, and you also configure VLAN-based access control for Telnet access. In this case, the only Telnet clients that can access the device are clients that have one of the IP addresses permitted by the ACL **and** are connected to a port that is in a permitted VLAN. Clients who have a permitted IP address but are connected to a port in a VLAN that is not permitted still cannot access the device through Telnet.

Restricting Telnet access to a specific VLAN

To allow Telnet access only to clients in a specific VLAN, enter a command such as the following.

```
PowerConnect(config)# telnet server enable vlan 10
```

The command in this example configures the device to allow Telnet management access only to clients connected to ports within port-based VLAN 10. Clients connected to ports that are not in VLAN 10 are denied management access.

Syntax: [no] telnet server enable vlan <vlan-id>

Restricting SNMP access to a specific VLAN

To allow SNMP access only to clients in a specific VLAN, enter a command such as the following.

```
PowerConnect(config)# snmp-server enable vlan 40
```

The command in this example configures the device to allow SNMP access only to clients connected to ports within port-based VLAN 40. Clients connected to ports that are not in VLAN 40 are denied access.

Syntax: [no] snmp-server enable vlan <vlan-id>

Restricting TFTP access to a specific VLAN

To allow TFTP access only to clients in a specific VLAN, enter a command such as the following.

```
PowerConnect(config)# tftp client enable vlan 40
```

The command in this example configures the device to allow TFTP access only to clients connected to ports within port-based VLAN 40. Clients connected to ports that are not in VLAN 40 are denied access.

Syntax: [no] tftp client enable vlan <vlan-id>

Designated VLAN for Telnet management sessions to a Layer 2 Switch

By default, the management IP address you configure on a Layer 2 Switch applies globally to all the ports on the device. This is true even if you divide the device ports into multiple port-based VLANs.

If you want to restrict the IP management address to a specific port-based VLAN, you can make that VLAN the designated management VLAN for the device. When you configure a VLAN to be the designated management VLAN, the management IP address you configure on the device is associated only with the ports in the designated VLAN. To establish a Telnet management session with the device, a user must access the device through one of the ports in the designated VLAN.

You also can configure up to five default gateways for the designated VLAN, and associate a metric with each one. The software uses the gateway with the lowest metric. The other gateways reside in the configuration but are not used. To use one of the other gateways, modify the configuration so that the gateway you want to use has the lowest metric.

If more than one gateway has the lowest metric, the gateway that appears first in the running-config is used.

NOTE

If you have already configured a default gateway globally and you do not configure a gateway in the VLAN, the software uses the globally configured gateway and gives the gateway a metric value of 1.

To configure a designated management VLAN, enter commands such as the following.

```
PowerConnect(config)# vlan 10 by port
PowerConnect(config-vlan-10)# untag ethernet 1 to 4
PowerConnect(config-vlan-10)# management-vlan
PowerConnect(config-vlan-10)# default-gateway 10.10.10.1 1
PowerConnect(config-vlan-10)# default-gateway 20.20.20.1 2
```

These commands configure port-based VLAN 10 to consist of ports 1 – 4 and to be the designated management VLAN. The last two commands configure default gateways for the VLAN. Since the 10.10.10.1 gateway has a lower metric, the software uses this gateway. The other gateway remains in the configuration but is not used. You can use the other one by changing the metrics so that the 20.20.20.1 gateway has the lower metric.

Syntax: [no] default-gateway <ip-addr> <metric>

The <ip-addr> parameters specify the IP address of the gateway router.

The <metric> parameter specifies the metric (cost) of the gateway. You can specify a value from 1 – 5. There is no default. The software uses the gateway with the lowest metric.

Device management security

By default, all management access is disabled. Each of the following management access methods must be specifically enabled as required in your installation:

- SSHv2

- SNMP

The commands for granting access to each of these management interfaces is described in the following.

SSHv2

To allow SSHv2 access to the device, you must generate a Crypto Key as shown in the following command.

```
PowerConnect(config)# crypto key generate
```

Syntax: crypto key [generate | zeroize]

The **generate** parameter generates a dsa key pair.

The **zeroize** parameter deletes the currently operative dsa key pair.

In addition, you must use AAA authentication to create a password to allow SSHv2 access. For example the following command configures AAA authentication to use TACACS+ for authentication as the default or local if TACACS+ is not available.

```
PowerConnect(config)# aaa authentication login default tacacs+ local
```

SNMP

To allow SNMP access to the device, enter the following command.

```
PowerConnect(config)# snmp-server
```

Syntax: [no] snmp-server

Disabling specific access methods

You can specifically disable the following access methods:

- Telnet access
- SNMP access
- TFTP

NOTE

If you disable Telnet access, you will not be able to access the CLI except through a serial connection to the management module. If you disable SNMP access, you will not be able to use Brocade Network Advisor or third-party SNMP management applications.

Disabling Telnet access

You can use a Telnet client to access the CLI on the device over the network. If you do not plan to use the CLI over the network and want to disable Telnet access to prevent others from establishing CLI sessions with the device, enter the following command.

```
PowerConnect(config)# no telnet server
```

To re-enable Telnet operation, enter the following command.

```
PowerConnect(config)# telnet server
```

Syntax: [no] telnet server

NOTE

The Telnet server is enabled by default.

Disabling SNMP access

SNMP is required if you want to manage a device using Brocade Network Advisor.

To disable SNMP management of the device.

```
PowerConnect(config)# no snmp-server
```

To later re-enable SNMP management of the device.

```
PowerConnect(config)# snmp-server
```

Syntax: no snmp-server

Disabling TFTP access

On PowerConnect B-Series TI24X devices, You can globally disable TFTP to block TFTP client access. By default, TFTP client access is enabled.

To disable TFTP client access, enter the following command at the Global CONFIG level of the CLI.

```
PowerConnect(config)# tftp disable
```

When TFTP is disabled, users are prohibited from using the **copy tftp** command to copy files to the system flash. If users enter this command while TFTP is disabled, the system will reject the command and display an error message.

To re-enable TFTP client access once it is disabled, enter the following command.

```
PowerConnect(config)# no tftp disable
```

Syntax: [no] tftp disable

Setting passwords

Passwords can be used to secure the following access methods:

- Telnet access can be secured by setting a Telnet password. Refer to [“Setting a Telnet password”](#) on page 866.
- Access to the Privileged EXEC and CONFIG levels of the CLI can be secured by setting passwords for management privilege levels. Refer to [“Setting passwords for management privilege levels”](#) on page 866.

This section also provides procedures for enhancing management privilege levels, recovering from a lost password, and disabling password encryption.

NOTE

You also can configure up to 16 user accounts consisting of a user name and password, and assign each user account a management privilege level. Refer to [“Setting up local user accounts”](#) on page 870.

Setting a Telnet password

By default, the device does not require a user name or password when you log in to the CLI using Telnet. You can assign a password for Telnet access using one of the following methods.

Set the password “letmein” for Telnet access to the CLI using the following command at the global CONFIG level.

```
PowerConnect(config)# enable telnet password letmein
```

Syntax: [no] enable telnet password <string>

Suppressing Telnet connection rejection messages

By default, if a device denies Telnet management access to the device, the software sends a message to the denied Telnet client. You can optionally suppress the rejection message. When you enable the option, a denied Telnet client does not receive a message from the device. Instead, the denied client simply does not gain access.

To suppress the connection rejection message, use the following CLI method.

To suppress the connection rejection message sent by the device to a denied Telnet client, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# telnet server suppress-reject-message
```

Syntax: [no] telnet server suppress-reject-message

Setting passwords for management privilege levels

You can set one password for each of the following management privilege levels:

- **Super User level** – Allows complete read-and-write access to the system. This is generally for system administrators and is the only management privilege level that allows you to configure passwords.
- **Port Configuration level** – Allows read-and-write access for specific ports but not for global (system-wide) parameters.
- **Read Only level** – Allows access to the Privileged EXEC mode and CONFIG mode of the CLI but only with read access.

You can assign a password to each management privilege level. You also can configure up to 16 user accounts consisting of a user name and password, and assign each user account to one of the three privilege levels. Refer to [“Setting up local user accounts”](#) on page 870.

NOTE

You must use the CLI to assign a password for management privilege levels.

If you configure user accounts in addition to privilege level passwords, the device will validate a user access attempt using one or both methods (local user account or privilege level password), depending on the order you specify in the authentication-method lists. Refer to [“Configuring authentication-method lists”](#) on page 907.

Follow the steps given below to set passwords for management privilege levels.

1. At the opening CLI prompt, enter the following command to change to the Privileged level of the EXEC mode.

```
PowerConnect> enable
PowerConnect#
```

2. Access the CONFIG level of the CLI by entering the following command.

```
PowerConnect# configure terminal
PowerConnect(config)#
```

3. Enter the following command to set the Super User level password.

```
PowerConnect(config)# enable super-user-password <text>
```

NOTE

You must set the Super User level password before you can set other types of passwords. The Super User level password can be an alphanumeric string, but cannot begin with a number.

4. Enter the following commands to set the Port Configuration level and Read Only level passwords.

```
PowerConnect(config)# enable port-config-password <text>
PowerConnect(config)# enable read-only-password <text>
```

Syntax: `enable super-user-password <text>`

Syntax: `enable port-config-password <text>`

Syntax: `enable read-only-password <text>`

NOTE

If you forget your Super User level password, refer to [“Recovering from a lost password”](#) on page 868.

Augmenting management privilege levels

Each management privilege level provides access to specific areas of the CLI by default:

- Super User level provides access to all commands and displays.
- Port Configuration level gives access to:
 - The User EXEC and Privileged EXEC levels
 - The port-specific parts of the CONFIG level
 - All interface configuration levels
- Read Only level gives access to:
 - The User EXEC and Privileged EXEC levels

You can grant additional access to a privilege level on an individual command basis. To grant the additional access, you specify the privilege level you are enhancing, the CLI level that contains the command, and the individual command.

NOTE

This feature applies only to management privilege levels on the CLI.

Enhance the Port Configuration privilege level so users also can enter IP commands at the global CONFIG level.

```
PowerConnect(config)# privilege configure level 4 ip
```

In this command, **configure** specifies that the enhanced access is for a command at the global CONFIG level of the CLI. The **level 4** parameter indicates that the enhanced access is for management privilege level 4 (Port Configuration). All users with Port Configuration privileges will have the enhanced access. The **ip** parameter indicates that the enhanced access is for the IP commands. Users who log in with valid Port Configuration level user names and passwords can enter commands that begin with “ip” at the global CONFIG level.

Syntax: [no] **privilege** <cli-level> **level** <privilege-level> <command-string>

The <cli-level> parameter specifies the CLI level and can be one of the following values:

- **exec** – EXEC level; for example, PowerConnect> or PowerConnect#
- **configure** – CONFIG level; for example, PowerConnect(config)#
- **interface** – Interface level; for example, PowerConnect(config-if-6)#
- **loopback-interface** – loopback interface level
- **virtual-interface** – Virtual-interface level; for example, PowerConnect(config-vif-6)#
- **dot1x** – 802.1X configuration level
- **ipv6-access-list** – IPv6 access list configuration level
- **rip-router** – RIP router level; for example, PowerConnect(config-rip-router)#
- **ospf-router** – OSPF router level; for example, PowerConnect(config-ospf-router)#
- **pim-router** – PIM router level; for example, PowerConnect(config-pim-router)#
- **bgp-router** – BGP4 router level; for example, PowerConnect(config-bgp-router)#
- **vrrp-router** – VRRP configuration level
- **gvrp** – GVRP configuration level
- **trunk** – trunk configuration level
- **port-vlan** – Port-based VLAN level; for example, PowerConnect(config-vlan)#
- **protocol-vlan** – Protocol-based VLAN level

The <privilege-level> indicates the number of the management privilege level you are augmenting. You can specify one of the following:

- **0** – Super User level (full read-write access)
- **4** – Port Configuration level
- **5** – Read Only level

The <command-string> parameter specifies the command you are allowing users with the specified privilege level to enter. To display a list of the commands at a CLI level, enter “?” at that level's command prompt.

Recovering from a lost password

Recovery from a lost password requires direct access to the serial port and a system reset.

NOTE

You can perform this procedure only from the CLI.

Follow the steps given below to recover from a lost password.

1. Start a CLI session over the serial interface to the device.
2. Reboot the device.
3. At the initial boot prompt at system startup, enter **b** to enter the boot monitor mode.
4. Enter **no password** at the prompt. (You cannot abbreviate this command.) This command will cause the device to bypass the system password check.
5. Enter **boot system flash primary** at the prompt.
6. After the console prompt reappears, assign a new password.

Displaying the SNMP community string

If you want to display the SNMP community string, enter the following commands.

```
PowerConnect(config)# enable password-display
PowerConnect# show snmp server
```

The **enable password-display** command enables display of the community string, but only in the output of the **show snmp server** command. Display of the string is still encrypted in the startup-config file and running-config. Enter the command at the global CONFIG level of the CLI.

Disabling password encryption

When you configure a password, then save the configuration to the flash memory on the device, the password is also saved to flash as part of the configuration file. By default, the passwords are encrypted so that the passwords cannot be observed by another user who displays the configuration file. Even if someone observes the file while it is being transmitted over TFTP, the password is encrypted.

If you want to remove the password encryption, you can disable encryption by entering the following command.

```
PowerConnect(config)# no service password-encryption
```

Syntax: [no] service password-encryption

Specifying a minimum password length

By default, the device imposes no minimum length on the Line (Telnet), Enable, or Local passwords. You can configure the device to require that Line, Enable, and Local passwords be at least a specified length.

For example, to specify that the Line, Enable, and Local passwords be at least 8 characters, enter the following command.

```
PowerConnect(config)# enable password-min-length 8
```

Syntax: enable password-min-length <number-of-characters>

The <number-of-characters> can be from 1 – 48.

NOTE

On PowerConnect B-Series TI24X devices, you can set local password.

Setting up local user accounts

You can define up to 16 local user accounts on a device. User accounts regulate who can access the management functions in the CLI using the following methods:

- Telnet access
- SNMP access

Local user accounts provide greater flexibility for controlling management access to devices than do management privilege level passwords and SNMP community strings of SNMP versions 1 and 2. You can continue to use the privilege level passwords and the SNMP community strings as additional means of access authentication. Alternatively, you can choose not to use local user accounts and instead continue to use only the privilege level passwords and SNMP community strings. Local user accounts are backward-compatible with configuration files that contain privilege level passwords. Refer to [“Setting passwords for management privilege levels”](#) on page 866.

If you configure local user accounts, you also need to configure an authentication-method list for Telnet access, and SNMP access. Refer to [“Configuring authentication-method lists”](#) on page 907.

For each local user account, you specify a user name. You also can specify the following parameters:

- A password
- A management privilege level, which can be one of the following:
 - **Super User level (default)** – Allows complete read-and-write access to the system. This is generally for system administrators and is the only privilege level that allows you to configure passwords.
 - **Port Configuration level** – Allows read-and-write access for specific ports but not for global parameters.
 - **Read Only level** – Allows access to the Privileged EXEC mode and CONFIG mode with read access only.

Enhancements to username and password

This section describes the enhancements to the username and password features introduced in the releases listed above.

The following rules are enabled by default:

- Users are required to accept the message of the day.
- Users are locked out (disabled) if they fail to login after three attempts. Use the **disable-on-login-failure** command to change the number of login attempts (up to 10) before users are locked out.

The following rules are disabled by default:

- Enhanced user password combination requirements
- User password masking
- Quarterly updates of user passwords
- You can configure the system to store up to 15 previously configured passwords for each user.
- You can use the **disable-on-login-failure** command to change the number of login attempts (up to 10) before users are locked out.

- A password can now be set to expire.

Enabling enhanced user password combination requirements

When strict password enforcement is enabled on the device, you must enter a minimum of eight characters containing the following combinations when you create an enable and a user password:

- At least two upper case characters
- At least two lower case characters
- At least two numeric characters
- At least two special characters

NOTE

Password minimum and combination requirements are strictly enforced.

Use the **enable strict-password-enforcement** command to enable the password security feature.

```
PowerConnect(config)# enable strict-password-enforcement
```

Syntax: [no] enable strict-password-enforcement

This feature is disabled by default.

The following security upgrades apply to the **enable strict-password-enforcement** command:

- Passwords must not share four or more concurrent characters with any other password configured on the router. If the user tries to create a password with four or more concurrent characters, the following error message will be returned.

```
Error - The substring <str> within the password has been used earlier, please choose a different password.
```

For example, the previous password was Mali4aYa&, the user cannot use any of the following as his or her new password:

- Malimai\$D because “Mail” were used consecutively in the previous password
- &3B9aYa& because “aYa&” were used consecutively in the previous password
- i4aYEv#8 because “i4aY” were used consecutively in the previous password
- If the user tries to configure a password that was previously used, the Local User Account configuration will not be allowed and the following message will be displayed.

```
This password was used earlier for same or different user, please choose a different password.
```

Enabling user password masking

By default, when you use the CLI to create a user password, the password displays on the console as you type it. For enhanced security, you can configure the device to mask the password characters entered at the CLI. When password masking is enabled, the CLI displays asterisks (*) on the console instead of the actual password characters entered.

The following shows the default CLI behavior when configuring a username and password.

```
PowerConnect(config)# username kelly password summertime
```

The following shows the CLI behavior when configuring a username and password when **password-masking** is enabled.

```
PowerConnect(config)# username kelly password
Enter Password: *****
```

NOTE

When password masking is enabled, press the [Enter] key before entering the password.

Syntax: `username <name> password [Enter]`

For [Enter], press the Enter key. Enter the password when prompted.

If **strict-password-enforcement** is enabled, enter a password which contains the required character combination. Refer to [“Enabling enhanced user password combination requirements”](#) on page 871.

To enable password masking, enter the following command.

```
PowerConnect(config)# enable user password-masking
```

Syntax: `[no] enable user password-masking`

Enabling user password aging

For enhanced security, password aging enforces quarterly updates of all user passwords. After 180 days, the CLI will automatically prompt users to change their passwords when they attempt to sign on.

When password aging is enabled, the software records the system time that each user password was configured or last changed. The time displays in the output of the **show running configuration** command, indicated by `set-time <time>`.

Example

```
PowerConnect# show run
Current configuration:
....
username waldo password .....
username raveen set-time 2086038248
....
```

The password aging feature uses the SNTP server clock to record the set-time. If the network does not have an SNTP server, then set-time will appear as **set-time 0** in the output of the **show running configuration** command.

A username set-time configuration is removed when:

- The username and password is deleted from the configuration
- The username password expires

When a username set-time configuration is removed, it no longer appears in the **show running configuration** output.

Note that if a username does not have an assigned password, the username will not have a set-time configuration.

Password aging is disabled by default. To enable it, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# enable user password-aging
```

Syntax: `[no] enable user password-aging`

Enhanced login lockout

The CLI provides up to three login attempts. If a user fails to login after three attempts, that user is locked out (disabled). If desired, you can increase or decrease the number of login attempts before the user is disabled. To do so, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# enable user disable-on-login-failure 7
```

Syntax: `enable user disable-on-login-failure <1 - 10>`

To re-enable a user that has been locked out, do one of the following:

- Reboot the device to re-enable all disabled users.
- Enable the user by entering the following command.

```
PowerConnect(config)# username sandy enable
```

Example

```
PowerConnect(config)# user sandy enable
PowerConnect# show user
Username Password Encrypt Priv Status Expire Time
=====
==
sandy $1$Gz...uX/$wQ44fVGtsqbKWkQknzAZ6. enabled 0 enabled 90 days
```

Syntax: `username <name> enable`

Setting passwords to expire

You can set a user password to expire. Once a password expires, the administrator must assign a new password to the user. To configure a user password to expire, enter the following.

```
PowerConnect(config)# username sandy expires 20
```

Syntax: `username <name> expires <days>`

Enter 1 – 365 for number of days. The default is 90 days.

Example

```
PowerConnect(config)# username sandy expires 20
PowerConnect# show user
Username Password Encrypt Priv Status Expire
Time
=====
==
sandy $1$Gz...uX/$wQ44fVGtsqbKWkQknzAZ6. enabled 0 enabled 20 days
```

Requirement to accept the message of the day

If a message of the day (MOTD) is configured, a user will be required to press the Enter key before he or she can login. MOTD is configured using the **banner motd** command.

There are no new CLI commands for this feature.

NOTE

This requirement is disabled by default, unless configured. Users are not required to press Enter after the MOTD banner is displayed.

Configuring a local user account

You can create accounts for local users with or without passwords. Accounts with passwords can have encrypted or unencrypted passwords.

You can assign privilege levels to local user accounts, but on a new device, you must create a local user account that has a Super User privilege before you can create accounts with other privilege levels.

NOTE

You must grant Super User level privilege to at least one account before you add accounts with other privilege levels. You need the Super User account to make further administrative changes.

Local user accounts with no passwords

To create a user account without a password, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# username wonka nopassword
```

Syntax: [no] **username** <user-string> **privilege** <privilege-level> **nopassword**

Local user accounts with unencrypted passwords

If you want to use unencrypted passwords for local user accounts, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# username wonka password willy
```

If password masking is enabled, press the [Enter] key before entering the password.

```
PowerConnect(config)# username wonka
Enter Password: willy
```

The above commands add a local user account with the user name “wonka” and the password “willy”. This account has the Super User privilege level; this user has full access to all configuration and display features.

```
PowerConnect(config)# username waldo privilege 5 password whereis
```

This command adds a user account for user name “waldo”, password “whereis”, with the Read Only privilege level. Waldo can look for information but cannot make configuration changes.

Syntax: [no] **username** <user-string> **privilege** <privilege-level> **password** | **nopassword** <password-string>

You can enter up to 48 characters for <user-string>.

The **privilege** <privilege-level> parameter specifies the privilege level for the account. You can specify one of the following:

- **0** – Super User level (full read-write access)
- **4** – Port Configuration level

- **5** – Read Only level

The default privilege level is **0**. If you want to assign Super User level access to the account, you can enter the command without **privilege 0**, as shown in the command example above.

The **password** | **nopassword** parameter indicates whether the user must enter a password. If you specify **password**, enter the string for the user's password. You can enter up to 48 characters for *<password-string>*. If **strict password enforcement** is enabled on the device, you must enter a minimum of eight characters containing the following combinations:

- At least two upper case characters
- At least two lower case characters
- At least two numeric characters
- At least two special characters

NOTE

You must be logged on with Super User access (privilege level 0) to add user accounts or configure other access parameters.

To display user account information, enter the following command.

```
PowerConnect# show users
```

Syntax: show users

Local accounts with encrypted passwords

You can create local user accounts with MD5 encrypted passwords using one of the following methods:

- Issuing the **service password-encryption** command after creating the local user account with a **username <user-string> [privilege <privilege-level>] password 0** command

NOTE

To create an encrypted all-numeric password, use the **username <user-string> create-password** command.

If you create a local user account using the commands discussed in [“Local user accounts with unencrypted passwords”](#) on page 874, you can issue the **service password-encryption** command to encrypt all passwords that have been previously entered.

Example

```
PowerConnect(config)# username wonka privilege 5 password willy
PowerConnect(config)# service password-encryption
```

If password masking is enabled, enter the commands this way.

```
PowerConnect(config)# username wonka privilege 5 password
Enter Password: willy
PowerConnect(config)# service password-encryption
```

Syntax: [no] service password-encryption

Create password option

As an alternative to the commands above, the **create-password** option allows you to create an encrypted password in one line of command. Also, this new option allows you to create an all-numeric, encrypted password.

You can enter.

```
PowerConnect(config)# username wonka privilege 5 create-password willy
```

Syntax: **[no] username** <user-string> **[privilege** <privilege-level>] **create-password** <password-string>

You can enter up to 48 characters for <user-string>. This string can be alphanumeric or all-numeric.

The **privilege** parameter specifies the privilege level for the account. You can specify one of the following:

- **0** – Super User level (full read-write access)
- **4** – Port Configuration level
- **5** – Read Only level

Enter up to 48 alphanumeric characters for <password-string>.

Changing a local user password

To change a local user password for an existing local user account, enter a command such as the following at the global CONFIG level of the CLI.

NOTE

You must be logged on with Super User access (privilege level 0) to change user passwords.

```
PowerConnect(config)# username wonka password willy
```

If password masking is enabled, enter the username, press the [Enter] key, then enter the password.

```
PowerConnect(config)# username wonka password
Enter Password: willy
```

The above commands change wonka's user name password to "willy".

Syntax: **[no] username** <user-string> **password** <password-string>

Enter up to 48 characters for <user-string>.

The <password-string> parameter is the user password. The password can be up to 48 characters and must differ from the current password and two previously configured passwords.

When a password is changed, a message such as the following is sent to the Syslog.

```
SYSLOG: <14>Jan 1 00:00:00 10.44.9.11 Security: Password has been changed for user
tester from console session.
```

The message includes the name of the user whose password was changed and during which session type, such as Console, Telnet, SSH, SNMP, or others, the password was changed.

Configuring TACACS/TACACS+ security

You can use the security protocol Terminal Access Controller Access Control System (TACACS) or TACACS+ to authenticate the following kinds of access to the device:

- Telnet access
- SSH access
- Console access
- Access to the Privileged EXEC level and CONFIG levels of the CLI

The TACACS and TACACS+ protocols define how authentication, authorization, and accounting information is sent between a device and an authentication database on a TACACS/TACACS+ server. TACACS/TACACS+ services are maintained in a database, typically on a UNIX workstation or PC with a TACACS/TACACS+ server running.

NOTE

On PowerConnect B-Series T124X devices, the TACACS+ security feature is supported.

How TACACS+ differs from TACACS

TACACS is a simple UDP-based access control protocol originally developed by BBN for MILNET. TACACS+ is an enhancement to TACACS and uses TCP to ensure reliable delivery.

TACACS+ is an enhancement to the TACACS security protocol. TACACS+ improves on TACACS by separating the functions of authentication, authorization, and accounting (AAA) and by encrypting all traffic between the device and the TACACS+ server. TACACS+ allows for arbitrary length and content authentication exchanges, which allow any authentication mechanism to be utilized with the device. TACACS+ is extensible to provide for site customization and future development features. The protocol allows the device to request very precise access control and allows the TACACS+ server to respond to each component of that request.

NOTE

TACACS+ provides for authentication, authorization, and accounting, but an implementation or configuration is not required to employ all three.

TACACS/TACACS+ authentication, authorization, and accounting

When you configure a device to use a TACACS/TACACS+ server for **authentication**, the device prompts users who are trying to access the CLI for a user name and password, then verifies the password with the TACACS/TACACS+ server.

If you are using TACACS+, Dell recommends that you also configure **authorization**, in which the device consults a TACACS+ server to determine which management privilege level (and which associated set of commands) an authenticated user is allowed to use. You can also optionally configure **accounting**, which causes the device to log information on the TACACS+ server when specified events occur on the device.

NOTE

By default, a user logging into the device from Telnet or SSH would first enter the User EXEC level. The user can enter the **enable** command to get to the Privileged EXEC level.

A user that is successfully authenticated can be automatically placed at the Privileged EXEC level after login. Refer to [“Entering privileged EXEC mode after a Telnet or SSH login”](#) on page 886.

TACACS authentication

NOTE

Also, multiple challenges are supported for TACACS+ login authentication.

When TACACS authentication takes place, the following events occur.

1. A user attempts to gain access to the device by doing one of the following:
 - Logging into the device using Telnet, or SSH.
 - Entering the Privileged EXEC level or CONFIG level of the CLI
2. The user is prompted for a username and password.
3. The user enters a username and password.
4. The device sends a request containing the username and password to the TACACS server.
5. The username and password are validated in the TACACS server database.
6. If the password is valid, the user is authenticated.

TACACS+ authentication

When TACACS+ authentication takes place, the following events occur.

1. A user attempts to gain access to the device by doing one of the following:
 - Logging into the device using Telnet, or SSH.
 - Entering the Privileged EXEC level or CONFIG level of the CLI
2. The user is prompted for a username.
3. The user enters a username.
4. The device obtains a password prompt from a TACACS+ server.
5. The user is prompted for a password.
6. The user enters a password.
7. The device sends the password to the TACACS+ server.
8. The password is validated in the TACACS+ server database.
9. If the password is valid, the user is authenticated.

TACACS+ authorization

Devices support two kinds of TACACS+ authorization:

- Exec authorization determines a user privilege level when they are authenticated

- Command authorization consults a TACACS+ server to get authorization for commands entered by the user

When TACACS+ exec authorization takes place, the following events occur.

1. A user logs into the device using Telnet, or SSH.
2. The user is authenticated.
3. The device consults the TACACS+ server to determine the privilege level of the user.
4. The TACACS+ server sends back a response containing an A-V (Attribute-Value) pair with the privilege level of the user.
5. The user is granted the specified privilege level.

When TACACS+ command authorization takes place, the following events occur.

1. A Telnet, SSH user previously authenticated by a TACACS+ server enters a command on the device.
2. The device looks at its configuration to see if the command is at a privilege level that requires TACACS+ command authorization.
3. If the command belongs to a privilege level that requires authorization, the device consults the TACACS+ server to see if the user is authorized to use the command.
4. If the user is authorized to use the command, the command is executed.

TACACS+ accounting

TACACS+ accounting works as follows.

1. One of the following events occur on the device:
 - A user logs into the management interface using Telnet or SSH
 - A user enters a command for which accounting has been configured
 - A system event occurs, such as a reboot or reloading of the configuration file
2. The device checks the configuration to see if the event is one for which TACACS+ accounting is required.
3. If the event requires TACACS+ accounting, the device sends a TACACS+ Accounting Start packet to the TACACS+ accounting server, containing information about the event.
4. The TACACS+ accounting server acknowledges the Accounting Start packet.
5. The TACACS+ accounting server records information about the event.
6. When the event is concluded, the device sends an Accounting Stop packet to the TACACS+ accounting server.
7. The TACACS+ accounting server acknowledges the Accounting Stop packet.

AAA operations for TACACS/TACACS+

The following table lists the sequence of authentication, authorization, and accounting operations that take place when a user gains access to a device that has TACACS/TACACS+ security configured.

TABLE 136

User action	Applicable AAA operations
User attempts to gain access to the Privileged EXEC and CONFIG levels of the CLI	Enable authentication: aaa authentication enable default <method-list>
	Exec authorization (TACACS+): aaa authorization exec default tacacs+
	System accounting start (TACACS+): aaa accounting system default start-stop <method-list>
User logs in using Telnet/SSH	Login authentication: aaa authentication login default <method-list>
	Exec authorization (TACACS+): aaa authorization exec default tacacs+
	Exec accounting start (TACACS+): aaa accounting exec default <method-list>
	System accounting start (TACACS+): aaa accounting system default start-stop <method-list>
User logs out of Telnet/SSH session	Command accounting (TACACS+): aaa accounting commands <privilege-level> default start-stop <method-list>
	EXEC accounting stop (TACACS+): aaa accounting exec default start-stop <method-list>
User enters system commands (for example, reload , boot system)	Command authorization (TACACS+): aaa authorization commands <privilege-level> default <method-list>
	Command accounting (TACACS+): aaa accounting commands <privilege-level> default start-stop <method-list>
	System accounting stop (TACACS+): aaa accounting system default start-stop <method-list>
User enters the command: [no] aaa accounting system default start-stop <method-list>	Command authorization (TACACS+): aaa authorization commands <privilege-level> default <method-list>
	Command accounting (TACACS+): aaa accounting commands <privilege-level> default start-stop <method-list>
	System accounting start (TACACS+): aaa accounting system default start-stop <method-list>
User enters other commands	Command authorization (TACACS+): aaa authorization commands <privilege-level> default <method-list>
	Command accounting (TACACS+): aaa accounting commands <privilege-level> default start-stop <method-list>

AAA security for commands pasted into the running-config

If AAA security is enabled on the device, commands pasted into the running-config are subject to the same AAA operations as if they were entered manually.

When you paste commands into the running-config, and AAA command authorization or accounting, or both, are configured on the device, AAA operations are performed on the pasted commands. The AAA operations are performed before the commands are actually added to the running-config. The server performing the AAA operations should be reachable when you paste the commands into the running-config file. If the device determines that a pasted command is invalid, AAA operations are halted on the remaining commands. The remaining commands may not be executed if command authorization is configured.

TACACS/TACACS+ configuration considerations

- You must deploy at least one TACACS/TACACS+ server in your network.
- Devices support authentication using up to eight TACACS/TACACS+ servers. The device tries to use the servers in the order you add them to the device configuration.
- You can select only one primary authentication method for each type of access to a device (CLI through Telnet, CLI Privileged EXEC and CONFIG levels). For example, you can select TACACS+ as the primary authentication method for Telnet CLI access, but you cannot also select RADIUS authentication as a primary method for the same type of access. However, you can configure backup authentication methods for each access type.
- You can configure the device to authenticate using a TACACS or TACACS+ server, not both.

TACACS configuration procedure

Follow the procedure given below for TACACS configurations.

1. Identify TACACS servers. Refer to [“Identifying the TACACS/TACACS+ servers”](#) on page 882.
2. Set optional parameters. Refer to [“Setting optional TACACS/TACACS+ parameters”](#) on page 883.
3. Configure authentication-method lists. Refer to [“Configuring authentication-method lists for TACACS/TACACS+”](#) on page 884.

TACACS+ configuration procedure

Follow the procedure given below for TACACS+ configurations.

1. Identify TACACS+ servers. Refer to [“Identifying the TACACS/TACACS+ servers”](#) on page 882.
2. Set optional parameters. Refer to [“Setting optional TACACS/TACACS+ parameters”](#) on page 883.
3. Configure authentication-method lists. Refer to [“Configuring authentication-method lists for TACACS/TACACS+”](#) on page 884.
4. Optionally configure TACACS+ authorization. Refer to [“Configuring TACACS+ authorization”](#) on page 886.
5. Optionally configure TACACS+ accounting. Refer to [“Configuring TACACS+ accounting”](#) on page 889.

Enabling TACACS

TACACS is disabled by default. To configure TACACS/TACACS+ authentication parameters, you must enable TACACS by entering the following command.

```
PowerConnect(config)# enable snmp config-tacacs
```

Syntax: [no] enable snmp <config-radius | config-tacacs>

The <config-radius> parameter specifies the RADIUS configuration mode. RADIUS is disabled by default.

The <config-tacacs> parameter specifies the TACACS configuration mode. TACACS is disabled by default.

Identifying the TACACS/TACACS+ servers

To use TACACS/TACACS+ servers to authenticate access to a device, you must identify the servers to the device.

For example, to identify three TACACS/TACACS+ servers, enter commands such as the following.

```
PowerConnect(config)# tacacs-server host 207.94.6.161
PowerConnect(config)# tacacs-server host 207.94.6.191
PowerConnect(config)# tacacs-server host 207.94.6.122
```

Syntax: tacacs-server host <ip-addr> | <ipv6-addr> | <hostname> [auth-port <number>]

The <ip-addr> | <ipv6-addr> | <hostname> parameter specifies the IP address or host name of the server. You can enter up to eight **tacacs-server host** commands to specify up to eight different servers.

NOTE

To specify the server's host name instead of its IP address, you must first identify a DNS server using the **ip dns server-address** <ip-addr> command at the global CONFIG level.

If you add multiple TACACS/TACACS+ authentication servers to the device, the device tries to reach them in the order you add them. For example, if you add three servers in the following order, the software tries the servers in the same order.

1. 207.94.6.161
2. 207.94.6.191
3. 207.94.6.122

You can remove a TACACS/TACACS+ server by entering **no** followed by the **tacacs-server** command. For example, to remove 207.94.6.161, enter the following command.

```
PowerConnect(config)# no tacacs-server host 207.94.6.161
```

NOTE

If you erase a **tacacs-server** command (by entering “no” followed by the command), make sure you also erase the **aaa** commands that specify TACACS/TACACS+ as an authentication method. (Refer to “[Configuring authentication-method lists for TACACS/TACACS+](#)” on page 884.) Otherwise, when you exit from the CONFIG mode or from a Telnet session, the system continues to believe it is TACACS/TACACS+ enabled and you will not be able to access the system.

The **auth-port** parameter specifies the UDP (for TACACS) or TCP (for TACACS+) port number of the authentication port on the server. The default port number is 49.

Specifying different servers for individual AAA functions

In a TACACS+ configuration, you can designate a server to handle a specific AAA task. For example, you can designate one TACACS+ server to handle authorization and another TACACS+ server to handle accounting. You can set the TACACS+ key for each server.

To specify different TACACS+ servers for authentication, authorization, and accounting, enter the command such as following.

```
PowerConnect(config)# tacacs-server host 1.2.3.4 auth-port 49
authentication-only key abc
PowerConnect(config)# tacacs-server host 1.2.3.5 auth-port 49 authorization-only
key def
PowerConnect(config)# tacacs-server host 1.2.3.6 auth-port 49 accounting-only
key ghi
```

Syntax: `tacacs-server host <ip-addr> | <ipv6-addr> | <server-name> [auth-port <num>] [authentication-only | authorization-only | accounting-only | default] [key 0 | 1 <string>]`

The default parameter causes the server to be used for all AAA functions.

After authentication takes place, the server that performed the authentication is used for authorization and accounting. If the authenticating server cannot perform the requested function, then the next server in the configured list of servers is tried; this process repeats until a server that can perform the requested function is found, or every server in the configured list has been tried.

Setting optional TACACS/TACACS+ parameters

You can set the following optional parameters in a TACACS/TACACS+ configuration:

- **TACACS+ key** – This parameter specifies the value that the device sends to the TACACS+ server when trying to authenticate user access.
- **Retransmit interval** – This parameter specifies how many times the device will resend an authentication request when the TACACS/TACACS+ server does not respond. The retransmit value can be from 1 – 5 times. The default is 3 times.
- **Dead time** – This parameter specifies how long the device waits for the primary authentication server to reply before deciding the server is dead and trying to authenticate using the next server. The dead-time value can be from 1 – 5 seconds. The default is 3 seconds.
- **Timeout** – This parameter specifies how many seconds the device waits for a response from a TACACS/TACACS+ server before either retrying the authentication request, or determining that the TACACS/TACACS+ servers are unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 – 15 seconds. The default is 3 seconds.
- **TACACS/TACACS+ over IPv6** – This parameter enables the device to send TACACS/TACACS+ packets over IPv6.

Setting the TACACS+ key

The **key** parameter in the `tacacs-server` command is used to encrypt TACACS+ packets before they are sent over the network. The value for the **key** parameter on the device should match the one configured on the TACACS+ server. The key can be from 1 – 32 characters in length and cannot include any space characters.

NOTE

The **tacacs-server key** command applies only to TACACS+ servers, not to TACACS servers. If you are configuring TACACS, do not configure a key on the TACACS server and do not enter a key on the device.

To specify a TACACS+ server key, enter a command such as following.

```
PowerConnect(config)# tacacs-server key rkwong
```

Syntax: **tacacs-server key** [0 | 1] <string>

When you display the configuration of the device, the TACACS+ keys are encrypted. For example.

```
PowerConnect(config)# tacacs-server key 1 abc
PowerConnect(config)# write terminal
...
tacacs-server host 1.2.3.5 auth-port 49
tacacs key 1 $!2d
```

NOTE

Encryption of the TACACS+ keys is done by default. The **0** parameter disables encryption. The **1** parameter is not required; it is provided for backwards compatibility.

Setting the retransmission limit

The **retransmit** parameter specifies how many times the device will resend an authentication request when the TACACS/TACACS+ server does not respond. The retransmit limit can be from 1 – 5 times. The default is 3 times.

To set the TACACS/TACACS+ retransmit limit, enter a command such as the following.

```
PowerConnect(config)# tacacs-server retransmit 5
```

Syntax: **tacacs-server retransmit** <number>

Setting the timeout parameter

The **timeout** parameter specifies how many seconds the device waits for a response from the TACACS/TACACS+ server before either retrying the authentication request, or determining that the TACACS/TACACS+ server is unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 – 15 seconds. The default is 3 seconds.

```
PowerConnect(config)# tacacs-server timeout 5
```

Syntax: **tacacs-server timeout** <number>

Configuring authentication-method lists for TACACS/TACACS+

You can use TACACS/TACACS+ to authenticate Telnet/SSH access and access to Privileged EXEC level and CONFIG levels of the CLI. When configuring TACACS/TACACS+ authentication, you create authentication-method lists specifically for these access methods, specifying TACACS/TACACS+ as the primary authentication method.

Within the authentication-method list, TACACS/TACACS+ is specified as the primary authentication method and up to six backup authentication methods are specified as alternates. If TACACS/TACACS+ authentication fails due to an error, the device tries the backup authentication methods in the order they appear in the list.

When you configure authentication-method lists for TACACS/TACACS+ authentication, you must create a separate authentication-method list for Telnet/SSH CLI access, and for access to the Privileged EXEC level and CONFIG levels of the CLI.

To create an authentication method list that specifies TACACS/TACACS+ as the primary authentication method for securing Telnet/SSH access to the CLI.

```
PowerConnect(config)# enable telnet authentication
PowerConnect(config)# aaa authentication login default tacacs local
```

The commands above cause TACACS/TACACS+ to be the primary authentication method for securing Telnet/SSH access to the CLI. If TACACS/TACACS+ authentication fails due to an error with the server, authentication is performed using local user accounts instead.

To create an authentication-method list that specifies TACACS/TACACS+ as the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI.

```
PowerConnect(config)# aaa authentication enable default tacacs local none
```

The command above causes TACACS/TACACS+ to be the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI. If TACACS/TACACS+ authentication fails due to an error with the server, local authentication is used instead. If local authentication fails, no authentication is used; the device automatically permits access.

Syntax: [no] **aaa authentication enable** | **login default** <method1> [<method2>] [<method3>] [<method4>] [<method5>] [<method6>] [<method7>]

The **web-server** | **enable** | **login** parameter specifies the type of access this authentication-method list controls. You can configure one authentication-method list for each type of access.

The <method1> parameter specifies the primary authentication method. The remaining optional <method> parameters specify additional methods to try if an error occurs with the primary method. A method can be one of the values listed in the Method Parameter column in the following table.

TABLE 137 Authentication method values

Method parameter	Description
line	Authenticate using the password you configured for Telnet access. The Telnet password is configured using the enable telnet password... command. Refer to “Setting a Telnet password” on page 866.
enable	Authenticate using the password you configured for the Super User privilege level. This password is configured using the enable super-user-password... command. Refer to “Setting passwords for management privilege levels” on page 866.
local	Authenticate using a local user name and password you configured on the device. Local user names and passwords are configured using the username... command. Refer to “Configuring a local user account” on page 874.
tacacs	Authenticate using the database on a TACACS server. You also must identify the server to the device using the tacacs-server command.
tacacs+	Authenticate using the database on a TACACS+ server. You also must identify the server to the device using the tacacs-server command.
radius	Authenticate using the database on a RADIUS server. You also must identify the server to the device using the radius-server command.
none	Do not use any authentication method. The device automatically permits access.

NOTE

For examples of how to define authentication-method lists for types of authentication other than TACACS/TACACS+, refer to “[Configuring authentication-method lists](#)” on page 907.

Entering privileged EXEC mode after a Telnet or SSH login

By default, a user enters User EXEC mode after a successful login through Telnet or SSH. Optionally, you can configure the device so that a user enters Privileged EXEC mode after a Telnet or SSH login. To do this, use the following command.

```
PowerConnect(config)# aaa authentication login privilege-mode
```

Syntax: aaa authentication login privilege-mode

The user privilege level is based on the privilege level granted during login.

Configuring enable authentication to prompt for password only

If Enable authentication is configured on the device, when a user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI, by default he or she is prompted for a username and password. You can configure the device to prompt only for a password. The device uses the username entered at login, if one is available. If no username was entered at login, the device prompts for both username and password.

To configure the device to prompt only for a password when a user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI.

```
PowerConnect(config)# aaa authentication enable implicit-user
```

Syntax: [no] aaa authentication enable implicit-user

Telnet/SSH prompts when the TACACS+ Server is unavailable

When TACACS+ is the first method in the authentication method list, the device displays the login prompt received from the TACACS+ server. If a user attempts to login through Telnet or SSH, but none of the configured TACACS+ servers are available, the following takes place:

- If the next method in the authentication method list is "enable", the login prompt is skipped, and the user is prompted for the Enable password (that is, the password configured with the **enable super-user-password** command).
- If the next method in the authentication method list is "line", the login prompt is skipped, and the user is prompted for the Line password (that is, the password configured with the **enable telnet password** command).

Configuring TACACS+ authorization

Devices support TACACS+ authorization for controlling access to management functions in the CLI. Two kinds of TACACS+ authorization are supported:

- Exec authorization determines a user privilege level when they are authenticated
- Command authorization consults a TACACS+ server to get authorization for commands entered by the user

Configuring exec authorization

When TACACS+ exec authorization is performed, the device consults a TACACS+ server to determine the privilege level of the authenticated user. To configure TACACS+ exec authorization on the device, enter the following command.

```
PowerConnect(config)# aaa authorization exec default tacacs+
```

Syntax: aaa authorization exec default tacacs+ | none

If you specify **none**, or omit the **aaa authorization exec** command from the device configuration, no exec authorization is performed.

A user privilege level is obtained from the TACACS+ server in the “foundry-privlvl” A-V pair. If the **aaa authorization exec default tacacs** command exists in the configuration, the device assigns the user the privilege level specified by this A-V pair. If the command does not exist in the configuration, then the value in the “foundry-privlvl” A-V pair is ignored, and the user is granted Super User access.

NOTE

If the **aaa authorization exec default tacacs+** command exists in the configuration, following successful authentication the device assigns the user the privilege level specified by the “foundry-privlvl” A-V pair received from the TACACS+ server. If the **aaa authorization exec default tacacs+** command does not exist in the configuration, then the value in the “foundry-privlvl” A-V pair is ignored, and the user is granted Super User access.

Also note that in order for the **aaa authorization exec default tacacs+** command to work, either the **aaa authentication enable default tacacs+** command, or the **aaa authentication login privilege-mode** command must also exist in the configuration.

Configuring an Attribute-Value pair on the TACACS+ server

During TACACS+ exec authorization, the device expects the TACACS+ server to send a response containing an A-V (Attribute-Value) pair that specifies the privilege level of the user. When the device receives the response, it extracts an A-V pair configured for the Exec service and uses it to determine the user privilege level.

To set a user privilege level, you can configure the “foundry-privlvl” A-V pair for the Exec service on the TACACS+ server.

Example

```
user=bob {
  default service = permit
  member admin
  #Global password
  global = cleartext "cat"
  service = exec {
    foundry-privlvl = 0
  }
}
```

In this example, the A-V pair `foundry-privlvl = 0` grants the user full read-write access. The value in the `foundry-privlvl` A-V pair is an integer that indicates the privilege level of the user. Possible values are 0 for super-user level, 4 for port-config level, or 5 for read-only level. If a value other than 0, 4, or 5 is specified in the `foundry-privlvl` A-V pair, the default privilege level of 5 (read-only) is used. The `foundry-privlvl` A-V pair can also be embedded in the group configuration for the user. See your TACACS+ documentation for the configuration syntax relevant to your server.

If the foundry-privlvl A-V pair is not present, the device extracts the last A-V pair configured for the Exec service that has a numeric value. The device uses this A-V pair to determine the user privilege level.

Example

```
user=bob {
  default service = permit
  member admin
  #Global password
  global = cleartext "cat"
  service = exec {
    priv-lvl = 15
  }
}
```

The attribute name in the A-V pair is not significant; the device uses the last one that has a numeric value. However, the device interprets the value for a non-“foundry-privlvl” A-V pair differently than it does for a “foundry-privlvl” A-V pair. The following table lists how the device associates a value from a non-“foundry-privlvl” A-V pair with a privilege level.

TABLE 138 Dell equivalents for non-“foundry-privlvl” A-V pair values

Value for non-“foundry-privlvl” A-V pair	Dell privilege level
15	0 (super-user)
From 14 - 1	4 (port-config)
Any other number or 0	5 (read-only)

In the example above, the A-V pair configured for the Exec service is `priv-lvl = 15`. The device uses the value in this A-V pair to set the user privilege level to 0 (super-user), granting the user full read-write access.

In a configuration that has both a “foundry-privlvl” A-V pair and a non-“foundry-privlvl” A-V pair for the Exec service, the non-“foundry-privlvl” A-V pair is ignored.

Example

```
user=bob {
  default service = permit
  member admin
  #Global password
  global = cleartext "cat"
  service = exec {
    foundry-privlvl = 4
    priv-lvl = 15
  }
}
```

In this example, the user would be granted a privilege level of 4 (port-config level). The `priv-lvl = 15` A-V pair is ignored by the device.

If the TACACS+ server has no A-V pair configured for the Exec service, the default privilege level of 5 (read-only) is used.

Configuring command authorization

When TACACS+ command authorization is enabled, the device consults a TACACS+ server to get authorization for commands entered by the user.

You enable TACACS+ command authorization by specifying a privilege level whose commands require authorization. For example, to configure the device to perform authorization for the commands available at the Super User privilege level (that is, all commands on the device), enter the following command.

```
PowerConnect(config)# aaa authorization commands 0 default tacacs+
```

Syntax: `aaa authorization commands <privilege-level> default tacacs+ | radius | none`

The `<privilege-level>` parameter can be one of the following:

- **0** – Authorization is performed for commands available at the Super User level (all commands)
- **4** – Authorization is performed for commands available at the Port Configuration level (port-config and read-only commands)
- **5** – Authorization is performed for commands available at the Read Only level (read-only commands)

NOTE

TACACS+ command authorization can be performed only for commands entered from Telnet or SSH sessions, or from the console. No authorization is performed for commands entered at the Brocade Network Advisor.

TACACS+ command authorization is not performed for the following commands:

- At all levels: **exit**, **logout**, **end**, and **quit**.
- At the Privileged EXEC level: **enable** or **enable <text>**, where `<text>` is the password configured for the Super User privilege level.

If configured, command accounting is performed for these commands.

AAA support for console commands

AAA support for commands entered at the console includes the following:

- Login prompt that uses AAA authentication, using authentication-method Lists
- Exec Authorization
- Exec Accounting
- Command authorization
- Command accounting
- System Accounting

To enable AAA support for commands entered at the console, enter the following command.

```
PowerConnect(config)# enable aaa console
```

Syntax: `[no] enable aaa console`

Configuring TACACS+ accounting

Devices support TACACS+ accounting for recording information about user activity and system events. When you configure TACACS+ accounting on a device, information is sent to a TACACS+ accounting server when specified events occur, such as when a user logs into the device or the system is rebooted.

Configuring TACACS+ accounting for Telnet/SSH (Shell) access

To send an Accounting Start packet to the TACACS+ accounting server when an authenticated user establishes a Telnet or SSH session on the device, and an Accounting Stop packet when the user logs out.

```
PowerConnect(config)# aaa accounting exec default start-stop tacacs+
```

Syntax: aaa accounting exec default start-stop radius | tacacs+ | none

Configuring TACACS+ accounting for CLI commands

You can configure TACACS+ accounting for CLI commands by specifying a privilege level whose commands require accounting. For example, to configure the device to perform TACACS+ accounting for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
PowerConnect(config)# aaa accounting commands 0 default start-stop tacacs+
```

An Accounting Start packet is sent to the TACACS+ accounting server when a user enters a command, and an Accounting Stop packet is sent when the service provided by the command is completed.

NOTE

If authorization is enabled, and the command requires authorization, then authorization is performed before accounting takes place. If authorization fails for the command, no accounting takes place.

Syntax: aaa accounting commands <privilege-level> default start-stop radius | tacacs+ | none

The <privilege-level> parameter can be one of the following:

- **0** – Records commands available at the Super User level (all commands)
- **4** – Records commands available at the Port Configuration level (port-config and read-only commands)
- **5** – Records commands available at the Read Only level (read-only commands)

Configuring TACACS+ accounting for system events

You can configure TACACS+ accounting to record when system events occur on the device. System events include rebooting and when changes to the active configuration are made.

The following command causes an Accounting Start packet to be sent to the TACACS+ accounting server when a system event occurs, and a Accounting Stop packet to be sent when the system event is completed.

```
PowerConnect(config)# aaa accounting system default start-stop tacacs+
```

Syntax: aaa accounting system default start-stop radius | tacacs+ | none

Configuring an interface as the source for all TACACS/TACACS+ packets

You can designate the lowest-numbered IP address configured on an Ethernet port, loopback interface, or virtual interface as the source IP address for all TACACS/TACACS+ packets from the Layer 3 Switch. Identifying a single source IP address for TACACS/TACACS+ packets provides the following benefits:

- If your TACACS/TACACS+ server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the TACACS/TACACS+ server by configuring the device to always send the TACACS/TACACS+ packets from the same link or source address.
- If you specify a loopback interface as the single source for TACACS/TACACS+ packets, TACACS/TACACS+ servers can receive the packets regardless of the states of individual links. Thus, if a link to the TACACS/TACACS+ server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS/TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an Ethernet, loopback, or virtual interface as the source for all TACACS/TACACS+ packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the port or interface as the source IP address for TACACS/TACACS+ packets originated by the device.

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all TACACS/TACACS+ packets, enter commands such as the following.

```
PowerConnect(config)# int ve 1
PowerConnect(config-vif-1)# ip address 10.0.0.3/24
PowerConnect(config-vif-1)# exit
PowerConnect(config)# ip tacacs source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all TACACS/TACACS+ packets from the Layer 3 Switch.

Syntax: `ip tacacs source-interface ethernet <portnum> | loopback <num> | ve <num>`

The `<portnum>` parameter is a valid port number.

The `<num>` parameter is a loopback interface or virtual interface number.

Displaying TACACS/TACACS+ statistics and configuration information

The **show aaa** command displays information about all TACACS+ and RADIUS servers identified on the device.

```

PowerConnect# show aaa
Tacacs+ key: Brocade
Tacacs+ retries: 1
Tacacs+ timeout: 15 seconds
Tacacs+ dead-time: 3 minutes
Tacacs+ Server: 207.95.6.90 Port:49:
                opens=6 closes=3 timeouts=3 errors=0
                packets in=4 packets out=4
no connection
Radius key: networks
Radius retries: 3
Radius timeout: 3 seconds
Radius dead-time: 3 minutes
Radius Server: 207.95.6.90 Auth Port=1645 Acct Port=1646:
                opens=2 closes=1 timeouts=1 errors=0
                packets in=1 packets out=4
no connection

```

The following table describes the TACACS/TACACS+ information displayed by the **show aaa** command.

TABLE 139 Output of the show aaa command for TACACS/TACACS+

Field	Description
Tacacs+ key	The setting configured with the tacacs-server key command. At the Super User privilege level, the actual text of the key is displayed. At the other privilege levels, a string of periods (....) is displayed instead of the text.
Tacacs+ retries	The setting configured with the tacacs-server retransmit command.
Tacacs+ timeout	The setting configured with the tacacs-server timeout command.
Tacacs+ dead-time	The setting configured with the tacacs-server dead-time command.
Tacacs+ Server	For each TACACS/TACACS+ server, the IP address, port, and the following statistics are displayed: <ul style="list-style-type: none"> • opens - Number of times the port was opened for communication with the server • closes - Number of times the port was closed normally • timeouts - Number of times port was closed due to a timeout • errors - Number of times an error occurred while opening the port • packets in - Number of packets received from the server • packets out - Number of packets sent to the server
connection	The current connection status. This can be “no connection” or “connection active”.

Configuring RADIUS security

You can use a Remote Authentication Dial In User Service (RADIUS) server to secure the following types of access to the Layer 2 Switch or Layer 3 Switch:

- Telnet access
- SSH access
- Access to the Privileged EXEC level and CONFIG levels of the CLI

NOTE

Devices do not support RADIUS security for SNMP (Brocade Network Advisor) access.

RADIUS authentication, authorization, and accounting

When RADIUS *authentication* is implemented, the device consults a RADIUS server to verify user names and passwords. You can optionally configure RADIUS *authorization*, in which the device consults a list of commands supplied by the RADIUS server to determine whether a user can execute a command he or she has entered, as well as *accounting*, which causes the device to log information on a RADIUS accounting server when specified events occur on the device.

RADIUS authentication

When RADIUS authentication takes place, the following events occur.

1. A user attempts to gain access to the device by doing one of the following:
 - Logging into the device using Telnet, or SSH.
 - Entering the Privileged EXEC level or CONFIG level of the CLI
2. The user is prompted for a username and password.
3. The user enters a username and password.
4. The device sends a RADIUS Access-Request packet containing the username and password to the RADIUS server.
5. The RADIUS server validates the device using a shared secret (the RADIUS key).
6. The RADIUS server looks up the username in its database.
7. If the username is found in the database, the RADIUS server validates the password.
8. If the password is valid, the RADIUS server sends an Access-Accept packet to the device, authenticating the user. Within the Access-Accept packet are three Dell vendor-specific attributes that indicate:
 - The privilege level of the user
 - A list of commands
 - Whether the user is allowed or denied usage of the commands in the listThe last two attributes are used with RADIUS authorization, if configured.
9. The user is authenticated, and the information supplied in the Access-Accept packet for the user is stored on the device. The user is granted the specified privilege level. If you configure RADIUS authorization, the user is allowed or denied usage of the commands in the list.

RADIUS authorization

When RADIUS authorization takes place, the following events occur.

1. A user previously authenticated by a RADIUS server enters a command on the device.
2. The device looks at its configuration to see if the command is at a privilege level that requires RADIUS command authorization.

3. If the command belongs to a privilege level that requires authorization, the device looks at the list of commands delivered to it in the RADIUS Access-Accept packet when the user was authenticated. (Along with the command list, an attribute was sent that specifies whether the user is permitted or denied usage of the commands in the list.)

NOTE

After RADIUS authentication takes place, the command list resides on the device. The RADIUS server is not consulted again once the user has been authenticated. This means that any changes made to the user command list on the RADIUS server are not reflected until the next time the user is authenticated by the RADIUS server, and the new command list is sent to the device.

4. If the command list indicates that the user is authorized to use the command, the command is executed.

RADIUS accounting

RADIUS accounting works as follows.

1. One of the following events occur on the device:
 - A user logs into the management interface using Telnet or SSH
 - A user enters a command for which accounting has been configured
 - A system event occurs, such as a reboot or reloading of the configuration file
2. The device checks its configuration to see if the event is one for which RADIUS accounting is required.
3. If the event requires RADIUS accounting, the device sends a RADIUS Accounting Start packet to the RADIUS accounting server, containing information about the event.
4. The RADIUS accounting server acknowledges the Accounting Start packet.
5. The RADIUS accounting server records information about the event.
6. When the event is concluded, the device sends an Accounting Stop packet to the RADIUS accounting server.
7. The RADIUS accounting server acknowledges the Accounting Stop packet.

AAA operations for RADIUS

The following table lists the sequence of authentication, authorization, and accounting operations that take place when a user gains access to a device that has RADIUS security configured.

TABLE 140

User action	Applicable AAA operations
User attempts to gain access to the Privileged EXEC and CONFIG levels of the CLI	Enable authentication: aaa authentication enable default <method-list>
	System accounting start: aaa accounting system default start-stop <method-list>

TABLE 140

User action	Applicable AAA operations
User logs in using Telnet/SSH	Login authentication: aaa authentication login default <method-list> EXEC accounting Start: aaa accounting exec default start-stop <method-list> System accounting Start: aaa accounting system default start-stop <method-list>
User logs out of Telnet/SSH session	Command authorization for logout command: aaa authorization commands <privilege-level> default <method-list> Command accounting: aaa accounting commands <privilege-level> default start-stop <method-list> EXEC accounting stop: aaa accounting exec default start-stop <method-list>
User enters system commands (for example, reload , boot system)	Command authorization: aaa authorization commands <privilege-level> default <method-list> Command accounting: aaa accounting commands <privilege-level> default start-stop <method-list> System accounting stop: aaa accounting system default start-stop <method-list>
User enters the command: [no] aaa accounting system default start-stop <method-list>	Command authorization: aaa authorization commands <privilege-level> default <method-list> Command accounting: aaa accounting commands <privilege-level> default start-stop <method-list> System accounting start: aaa accounting system default start-stop <method-list>
User enters other commands	Command authorization: aaa authorization commands <privilege-level> default <method-list> Command accounting: aaa accounting commands <privilege-level> default start-stop <method-list>

AAA security for commands pasted into the running-config

If AAA security is enabled on the device, commands pasted into the running-config are subject to the same AAA operations as if they were entered manually.

When you paste commands into the running-config, and AAA command authorization or accounting, or both, are configured on the device, AAA operations are performed on the pasted commands. The AAA operations are performed before the commands are actually added to the running-config. The server performing the AAA operations should be reachable when you paste the commands into the running-config file. If the device determines that a pasted command is invalid, AAA operations are halted on the remaining commands. The remaining commands may not be executed if command authorization is configured.

NOTE

Since RADIUS command authorization relies on a list of commands received from the RADIUS server when authentication is performed, it is important that you use RADIUS authentication when you also use RADIUS command authorization.

RADIUS configuration considerations

- You must deploy at least one RADIUS server in your network.
- Devices support authentication using up to eight RADIUS servers, including those used for 802.1X authentication and for management. The device tries to use the servers in the order you add them to the device configuration. If one RADIUS server times out (does not respond), the device tries the next one in the list. Servers are tried in the same sequence each time there is a request.
- You can select only one primary authentication method for each type of access to a device (CLI through Telnet, CLI Privileged EXEC and CONFIG levels). For example, you can select RADIUS as the primary authentication method for Telnet CLI access, but you cannot also select TACACS+ authentication as the primary method for the same type of access. However, you can configure backup authentication methods for each access type.

RADIUS configuration procedure

Follow the procedure given below to configure a device for RADIUS.

1. Configure Dell vendor-specific attributes on the RADIUS server. Refer to [“Configuring Dell-specific attributes on the RADIUS server”](#) on page 896.
2. Identify the RADIUS server to the device. Refer to [“Identifying the RADIUS server to the device”](#) on page 898.
3. Optionally specify different servers for individual AAA functions. Refer to [“Specifying different servers for individual AAA functions”](#) on page 898.
4. Optionally configure the RADIUS server as a “port only” server. Refer to [“Configuring a RADIUS server per port”](#) on page 898.
5. Optionally bind the RADIUS servers to ports on the device. Refer to [“Mapping a RADIUS server to individual ports”](#) on page 899.
6. Set RADIUS parameters. Refer to [“Setting RADIUS parameters”](#) on page 900.
7. Configure authentication-method lists. Refer to [“Configuring authentication-method lists for RADIUS”](#) on page 901.
8. Optionally configure RADIUS authorization. Refer to [“Configuring RADIUS authorization”](#) on page 903.
9. Optionally configure RADIUS accounting. [“Configuring RADIUS accounting”](#) on page 905.

Configuring Dell-specific attributes on the RADIUS server

NOTE

For all devices, RADIUS Challenge is supported for 802.1x authentication but not for login authentication.

During the RADIUS authentication process, if a user supplies a valid username and password, the RADIUS server sends an Access-Accept packet to the device, authenticating the user. Within the Access-Accept packet are three Dell vendor-specific attributes that indicate:

- The privilege level of the user
- A list of commands

- Whether the user is allowed or denied usage of the commands in the list

You must add these three Dell vendor-specific attributes to your RADIUS server configuration, and configure the attributes in the individual or group profiles of the users that will access the device.

Dell Vendor-ID is 1991, with Vendor-Type 1. The following table describes the Dell vendor-specific attributes.

TABLE 141 Dell vendor-specific attributes for RADIUS

Attribute name	Attribute ID	Data type	Description
foundry-privilege-level	1	integer	Specifies the privilege level for the user. This attribute can be set to one of the following: <ul style="list-style-type: none"> • 0 - Super User level – Allows complete read-and-write access to the system. This is generally for system administrators and is the only management privilege level that allows you to configure passwords. • 4 - Port Configuration level – Allows read-and-write access for specific ports but not for global (system-wide) parameters. • 5 - Read Only level – Allows access to the Privileged EXEC mode and CONFIG mode of the CLI but only with read access.
foundry-command-string	2	string	Specifies a list of CLI commands that are permitted or denied to the user when RADIUS authorization is configured. The commands are delimited by semi-colons (;). You can specify an asterisk (*) as a wildcard at the end of a command string. For example, the following command list specifies all show and debug ip commands, as well as the write terminal command: show *; debug ip *; write term*
foundry-command-exception-flag	3	integer	Specifies whether the commands indicated by the foundry-command-string attribute are permitted or denied to the user. This attribute can be set to one of the following: <ul style="list-style-type: none"> • 0 - Permit execution of the commands indicated by foundry-command-string, deny all other commands. • 1 - Deny execution of the commands indicated by foundry-command-string, permit all other commands.

Enabling SNMP to configure RADIUS

To enable SNMP access to RADIUS MIB objects on the device, enter a command such as the following.

```
PowerConnect(config)# enable snmp config-radius
```

Syntax: [no] enable snmp <config-radius | config-tacacs>

The <config-radius> parameter specifies the RADIUS configuration mode. RADIUS is disabled by default.

The `<config-tacacs>` parameter specifies the TACACS configuration mode. TACACS is disabled by default.

Identifying the RADIUS server to the device

To use a RADIUS server to authenticate access to a device, you must identify the server to the device.

Example

```
PowerConnect(config)# radius-server host 209.157.22.99
```

Syntax: `radius-server host <ip-addr> | <ipv6-addr> | <server-name> [auth-port <number>] [acct-port <number>]`

The `host <ip-addr> | <ipv6-addr> | <server-name>` parameter is either an IP address or an ASCII text string.

The `<auth-port>` parameter is the Authentication port number. The default is 1645.

The `<acct-port>` parameter is the Accounting port number. The default is 1646.

Specifying different servers for individual AAA functions

In a RADIUS configuration, you can designate a server to handle a specific AAA task. For example, you can designate one RADIUS server to handle authorization and another RADIUS server to handle accounting. You can specify individual servers for authentication and accounting, but not for authorization. You can set the RADIUS key for each server.

To specify different RADIUS servers for authentication, authorization, and accounting, enter commands such as the following.

```
PowerConnect(config)# radius-server host 1.2.3.4 authentication-only key abc
PowerConnect(config)# radius-server host 1.2.3.5 authorization-only key def
PowerConnect(config)# radius-server host 1.2.3.6 accounting-only key ghi
```

Syntax: `radius-server host <ip-addr> | <ipv6-addr> | <server-name> [auth-port <number>] [acct-port <number>] [authentication-only | accounting-only | default] [key 0 | 1 <string>]`

The **default** parameter causes the server to be used for all AAA functions.

After authentication takes place, the server that performed the authentication is used for authorization and accounting. If the authenticating server cannot perform the requested function, then the next server in the configured list of servers is tried; this process repeats until a server that can perform the requested function is found, or every server in the configured list has been tried.

Configuring a RADIUS server per port

You can optionally configure a RADIUS server per port, indicating that it will be used only to authenticate users on ports to which it is mapped. A RADIUS server that is not explicitly configured as a RADIUS server per port is a **global server**, and can be used to authenticate users on ports to which no RADIUS servers are mapped.

Configuration notes

- This feature works with 802.1X and multi-device port authentication only.
- As in previous releases, you can define up to eight RADIUS servers per device.

Configuration example and command syntax

The following shows an example configuration.

```
PowerConnect(config)# radius-server host 10.10.10.103 auth-port 1812 acct-port
1813 default key mykeyword dot1x port-only
PowerConnect(config)# radius-server host 10.10.10.104 auth-port 1812 acct-port
1813 default key mykeyword dot1x port-only
PowerConnect(config)# radius-server host 10.10.10.105 auth-port 1812 acct-port
1813 default key mykeyword dot1x
PowerConnect(config)# radius-server host 10.10.10.106 auth-port 1812 acct-port
1813 default key mykeyword dot1x
```

The above configuration has the following affect:

- RADIUS servers 10.10.10.103 and 10.10.10.104 will be used only to authenticate users on ports to which the servers are mapped. To map a RADIUS server to a port, refer to [“Mapping a RADIUS server to individual ports”](#) on page 899.
- RADIUS servers 10.10.10.105 and 10.10.10.106 will be used to authenticate users on ports to which no RADIUS servers are mapped. For example, port e 9, to which no RADIUS servers are mapped, will send a RADIUS request to the first configured RADIUS server, 10.10.10.105. If the request fails, it will go to the second configured RADIUS server, 10.10.10.106. It will not send requests to 10.10.10.103 or 10.10.10.104, since these servers are configured as port servers.

Syntax: `radius-server host <ip-addr> | <server-name> [auth-port <number>] [acct-port <number>] [default key <string> dot1x] [port-only]`

The **host** `<ip-addr>` is the IPv4 address.

The **auth-port** `<number>` parameter is the Authentication port number; it is an optional parameter. The default is 1645.

The **acct-port** `<number>` parameter is the Accounting port number; it is an optional parameter. The default is 1646.

The **default key** `<string> dot1x` parameter indicates that this RADIUS server supports the 802.1X standard. A RADIUS server that supports the 802.1X standard can also be used to authenticate non-802.1X authentication requests.

The **port-only** parameter is optional and specifies that the server will be used only to authenticate users on ports to which it is mapped.

Mapping a RADIUS server to individual ports

You can map up to eight RADIUS servers to each port on the device. The port will authenticate users using only the RADIUS servers to which the port is mapped. If there are no RADIUS servers mapped to a port, it will use the “global” servers for authentication.

As in previous releases, a port goes through the list of servers in the order in which it was mapped or configured, until a server that can perform the requested function is found, or until every server in the list has been tried.

Configuration notes

- This feature works with 802.1X and multic-device port authentication only.
- You can map a RADIUS server to a physical port only. You cannot map a RADIUS server to a VE.

Configuration example and command syntax

To map a RADIUS server to a port, enter commands such as the following.

```
PowerConnect(config)# int e 3
PowerConnect(config-if-e10000-3)# dot1x port-control auto
PowerConnect(config-if-e10000-3)# use-radius-server 10.10.10.103
PowerConnect(config-if-e10000-3)# use-radius-server 10.10.10.110
```

With the above configuration, port e 3 would send a RADIUS request to 10.10.10.103 first, since it is the first server mapped to the port. If it fails, it will go to 10.10.10.110.

Syntax: `use-radius-server <ip-addr>`

The **host** `<ip-addr>` is an IPv4 address.

Setting RADIUS parameters

You can set the following parameters in a RADIUS configuration:

- **RADIUS key** – This parameter specifies the value that the device sends to the RADIUS server when trying to authenticate user access.
- **Retransmit interval** – This parameter specifies how many times the device will resend an authentication request when the RADIUS server does not respond. The retransmit value can be from 1 – 5 times. The default is 3 times.
- **Timeout** – This parameter specifies how many seconds the device waits for a response from a RADIUS server before either retrying the authentication request, or determining that the RADIUS servers are unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 – 15 seconds. The default is 3 seconds.

Setting the RADIUS key

The **key** parameter in the **radius-server** command is used to encrypt RADIUS packets before they are sent over the network. The value for the **key** parameter on the device should match the one configured on the RADIUS server. The key can be from 1 – 32 characters in length and cannot include any space characters.

To specify a RADIUS server key, enter a command such as the following.

```
PowerConnect(config)# radius-server key mirabeau
```

Syntax: `radius-server key [0 | 1] <string>`

When you display the configuration of the device, the RADIUS key is encrypted.

Example

```
PowerConnect(config)# radius-server key 1 abc
PowerConnect(config)# write terminal
...
radius-server host 1.2.3.5
radius key 1 $!2d
```

NOTE

Encryption of the RADIUS keys is done by default. The **0** parameter disables encryption. The **1** parameter is not required; it is provided for backwards compatibility.

Setting the retransmission limit

The **retransmit** parameter specifies the maximum number of retransmission attempts. When an authentication request times out, the software will retransmit the request up to the maximum number of retransmissions configured. The default retransmit value is 3 retries. The range of retransmit values is from 1 – 5.

To set the RADIUS retransmit limit, enter a command such as the following.

```
PowerConnect(config)# radius-server retransmit 5
```

Syntax: **radius-server retransmit** <number>

Setting the timeout parameter

The **timeout** parameter specifies how many seconds the device waits for a response from the RADIUS server before either retrying the authentication request, or determining that the RADIUS server is unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 – 15 seconds. The default is 3 seconds.

```
PowerConnect(config)# radius-server timeout 5
```

Syntax: **radius-server timeout** <number>

Configuring authentication-method lists for RADIUS

You can use RADIUS to authenticate Telnet/SSH access and access to Privileged EXEC level and CONFIG levels of the CLI. When configuring RADIUS authentication, you create authentication-method lists specifically for these access methods, specifying RADIUS as the primary authentication method.

Within the authentication-method list, RADIUS is specified as the primary authentication method and up to six backup authentication methods are specified as alternates. If RADIUS authentication fails due to an error, the device tries the backup authentication methods in the order they appear in the list.

When you configure authentication-method lists for RADIUS, you must create a separate authentication-method list for Telnet or SSH CLI access and for CLI access to the Privileged EXEC level and CONFIG levels of the CLI.

To create an authentication-method list that specifies RADIUS as the primary authentication method for securing Telnet access to the CLI.

```
PowerConnect(config)# enable telnet authentication  
PowerConnect(config)# aaa authentication login default radius local
```

The commands above cause RADIUS to be the primary authentication method for securing Telnet access to the CLI. If RADIUS authentication fails due to an error with the server, local authentication is used instead.

To create an authentication-method list that specifies RADIUS as the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI.

```
PowerConnect(config)# aaa authentication enable default radius local none
```

The command above causes RADIUS to be the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI. If RADIUS authentication fails due to an error with the server, local authentication is used instead. If local authentication fails, no authentication is used; the device automatically permits access.

Syntax: [no] **aaa authentication enable** | **login default** <method1> [<method2>] [<method3>] [<method4>] [<method5>] [<method6>] [<method7>]

The **web-server** | **enable** | **login** parameter specifies the type of access this authentication-method list controls. You can configure one authentication-method list for each type of access.

The <method1> parameter specifies the primary authentication method. The remaining optional <method> parameters specify additional methods to try if an error occurs with the primary method. A method can be one of the values listed in the Method Parameter column in the following table.

TABLE 142 Authentication method values

Method parameter	Description
line	Authenticate using the password you configured for Telnet access. The Telnet password is configured using the enable telnet password... command. Refer to “Setting a Telnet password” on page 866.
enable	Authenticate using the password you configured for the Super User privilege level. This password is configured using the enable super-user-password... command. Refer to “Setting passwords for management privilege levels” on page 866.
local	Authenticate using a local user name and password you configured on the device. Local user names and passwords are configured using the username... command. Refer to “Configuring a local user account” on page 874.
tacacs	Authenticate using the database on a TACACS server. You also must identify the server to the device using the tacacs-server command.
tacacs+	Authenticate using the database on a TACACS+ server. You also must identify the server to the device using the tacacs-server command.
radius	Authenticate using the database on a RADIUS server. You also must identify the server to the device using the radius-server command.
none	Do not use any authentication method. The device automatically permits access.

NOTE

For examples of how to define authentication-method lists for types of authentication other than RADIUS, refer to [“Configuring authentication-method lists”](#) on page 907.

Entering privileged EXEC mode after a Telnet or SSH login

By default, a user enters User EXEC mode after a successful login through Telnet or SSH. Optionally, you can configure the device so that a user enters Privileged EXEC mode after a Telnet or SSH login. To do this, use the following command.

```
PowerConnect(config)# aaa authentication login privilege-mode
```

Syntax: **aaa authentication login privilege-mode**

The user privilege level is based on the privilege level granted during login.

Configuring enable authentication to prompt for password only

If Enable authentication is configured on the device, when a user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI, by default he or she is prompted for a username and password. In this release, you can configure the device to prompt only for a password. The device uses the username entered at login, if one is available. If no username was entered at login, the device prompts for both username and password.

To configure the device to prompt only for a password when a user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI.

```
PowerConnect(config)# aaa authentication enable implicit-user
```

Syntax: [no] aaa authentication enable implicit-user

Configuring RADIUS authorization

Devices support RADIUS authorization for controlling access to management functions in the CLI. Two kinds of RADIUS authorization are supported:

- Exec authorization determines a user privilege level when they are authenticated
- Command authorization consults a RADIUS server to get authorization for commands entered by the user

Configuring exec authorization

When RADIUS exec authorization is performed, the device consults a RADIUS server to determine the privilege level of the authenticated user. To configure RADIUS exec authorization on the device, enter the following command.

```
PowerConnect(config)# aaa authorization exec default radius
```

Syntax: aaa authorization exec default radius | none

If you specify **none**, or omit the **aaa authorization exec** command from the device configuration, no exec authorization is performed.

NOTE

If the **aaa authorization exec default radius** command exists in the configuration, following successful authentication the device assigns the user the privilege level specified by the foundry-privilege-level attribute received from the RADIUS server. If the **aaa authorization exec default radius** command does not exist in the configuration, then the value in the foundry-privilege-level attribute is ignored, and the user is granted Super User access.

Also note that in order for the **aaa authorization exec default radius** command to work, either the **aaa authentication enable default radius** command, or the **aaa authentication login privilege-mode** command must also exist in the configuration.

Configuring command authorization

When RADIUS command authorization is enabled, the device consults the list of commands supplied by the RADIUS server during authentication to determine whether a user can execute a command he or she has entered.

You enable RADIUS command authorization by specifying a privilege level whose commands require authorization. For example, to configure the device to perform authorization for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
PowerConnect(config)# aaa authorization commands 0 default radius
```

Syntax: `aaa authorization commands <privilege-level> default radius | tacacs+ | none`

The `<privilege-level>` parameter can be one of the following:

- **0** – Authorization is performed (that is, the device looks at the command list) for commands available at the Super User level (all commands)
- **4** – Authorization is performed for commands available at the Port Configuration level (port-config and read-only commands)
- **5** – Authorization is performed for commands available at the Read Only level (read-only commands)

NOTE

RADIUS command authorization can be performed only for commands entered from Telnet or SSH sessions, or from the console. No authorization is performed for commands entered at the Brocade Network Advisor.

NOTE

Since RADIUS command authorization relies on the command list supplied by the RADIUS server during authentication, you cannot perform RADIUS authorization without RADIUS authentication.

Command authorization and accounting for console commands

The device supports command authorization and command accounting for CLI commands entered at the console. To configure the device to perform command authorization and command accounting for console commands, enter the following.

```
PowerConnect(config)# enable aaa console
```

Syntax: `enable aaa console`



CAUTION

If you have previously configured the device to perform command authorization using a RADIUS server, entering the `enable aaa console` command may prevent the execution of any subsequent commands entered on the console.

This happens because RADIUS command authorization requires a list of allowable commands from the RADIUS server. This list is obtained during RADIUS authentication. For console sessions, RADIUS authentication is performed only if you have configured Enable authentication and specified RADIUS as the authentication method (for example, with the `aaa authentication enable default radius` command). If RADIUS authentication is never performed, the list of allowable commands is never obtained from the RADIUS server. Consequently, there would be no allowable commands on the console.

Configuring RADIUS accounting

Devices support RADIUS accounting for recording information about user activity and system events. When you configure RADIUS accounting on a device, information is sent to a RADIUS accounting server when specified events occur, such as when a user logs into the device or the system is rebooted.

Configuring RADIUS accounting for Telnet/SSH (Shell) access

To send an Accounting Start packet to the RADIUS accounting server when an authenticated user establishes a Telnet or SSH session on the device, and an Accounting Stop packet when the user logs out.

```
PowerConnect(config)# aaa accounting exec default start-stop radius
```

Syntax: aaa accounting exec default start-stop radius | tacacs+ | none

Configuring RADIUS accounting for CLI commands

You can configure RADIUS accounting for CLI commands by specifying a privilege level whose commands require accounting. For example, to configure the device to perform RADIUS accounting for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
PowerConnect(config)# aaa accounting commands 0 default start-stop radius
```

An Accounting Start packet is sent to the RADIUS accounting server when a user enters a command, and an Accounting Stop packet is sent when the service provided by the command is completed.

NOTE

If authorization is enabled, and the command requires authorization, then authorization is performed before accounting takes place. If authorization fails for the command, no accounting takes place.

Syntax: aaa accounting commands <privilege-level> default start-stop radius | tacacs | none

The <privilege-level> parameter can be one of the following:

- **0** – Records commands available at the Super User level (all commands)
- **4** – Records commands available at the Port Configuration level (port-config and read-only commands)
- **5** – Records commands available at the Read Only level (read-only commands)

Configuring RADIUS accounting for system events

You can configure RADIUS accounting to record when system events occur on the device. System events include rebooting and when changes to the active configuration are made.

The following command causes an Accounting Start packet to be sent to the RADIUS accounting server when a system event occurs, and a Accounting Stop packet to be sent when the system event is completed.

```
PowerConnect(config)# aaa accounting system default start-stop radius
```

Syntax: aaa accounting system default start-stop radius | tacacs+ | none

Configuring an interface as the source for all RADIUS packets

You can designate the lowest-numbered IP address configured on an Ethernet port, loopback interface, or virtual interface as the source IP address for all RADIUS packets from the Layer 3 Switch. Identifying a single source IP address for RADIUS packets provides the following benefits:

- If your RADIUS server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the RADIUS server by configuring the device to always send the RADIUS packets from the same link or source address.
- If you specify a loopback interface as the single source for RADIUS packets, RADIUS servers can receive the packets regardless of the states of individual links. Thus, if a link to the RADIUS server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS/TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an Ethernet, loopback or virtual interface as the source for all RADIUS packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the port or interface as the source IP address for RADIUS packets originated by the device.

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all RADIUS packets, enter commands such as the following.

```
PowerConnect(config)# int ve 1
PowerConnect(config-vif-1)# ip address 10.0.0.3/24
PowerConnect(config-vif-1)# exit
PowerConnect(config)# ip radius source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all RADIUS packets from the Layer 3 Switch.

Syntax: `ip radius source-interface ethernet <portnum> | loopback <num> | ve <num>`

The `<portnum>` parameter is a valid port number.

The `<num>` parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the `<portnum>` is the port number.

Displaying RADIUS configuration information

The `show aaa` command displays information about all TACACS/TACACS+ and RADIUS servers identified on the device.

Example

```
PowerConnect# show aaa
Tacacs+ key: Brocade
Tacacs+ retries: 1
Tacacs+ timeout: 15 seconds
Tacacs+ dead-time: 3 minutes
Tacacs+ Server: 207.95.6.90 Port:49:
                opens=6 closes=3 timeouts=3 errors=0
                packets in=4 packets out=4
no connection
Radius key: networks
Radius retries: 3
Radius timeout: 3 seconds
Radius dead-time: 3 minutes
Radius Server: 207.95.6.90 Auth Port=1645 Acct Port=1646:
                opens=2 closes=1 timeouts=1 errors=0
                packets in=1 packets out=4
no connection
```

The following table describes the RADIUS information displayed by the **show aaa** command.

TABLE 143 Output of the show aaa command for RADIUS

Field	Description
Radius key	The setting configured with the radius-server key command. At the Super User privilege level, the actual text of the key is displayed. At the other privilege levels, a string of periods (....) is displayed instead of the text.
Radius retries	The setting configured with the radius-server retransmit command.
Radius timeout	The setting configured with the radius-server timeout command.
Radius dead-time	The setting configured with the radius-server dead-time command.
Radius Server	For each RADIUS server, the IP address, and the following statistics are displayed: Auth PortRADIUS authentication port number (default 1645) Acct PortRADIUS accounting port number (default 1646) <ul style="list-style-type: none"> • opens - Number of times the port was opened for communication with the server • closes - Number of times the port was closed normally • timeouts - Number of times port was closed due to a timeout • errors - Number of times an error occurred while opening the port • packets in - Number of packets received from the server • packets out - Number of packets sent to the server
connection	The current connection status. This can be "no connection" or "connection active".

Configuring authentication-method lists

To implement one or more authentication methods for securing access to the device, you configure authentication-method lists that set the order in which the authentication methods are consulted.

In an authentication-method list, you specify the access method (Telnet, SNMP, and so on) and the order in which the device tries one or more of the following authentication methods:

- Local Telnet login password
- Local password for the Super User privilege level

- Local user accounts configured on the device
- Database on a TACACS or TACACS+ server
- Database on a RADIUS server
- No authentication

NOTE

The TACACS/TACACS+, RADIUS, and Telnet login password authentication methods are not supported for SNMP access.

NOTE

To authenticate Telnet access to the CLI, you also must enable the authentication by entering the **enable telnet authentication** command at the global CONFIG level of the CLI.

NOTE

You do not need an authentication-method list to secure access based on ACLs or a list of IP addresses. Refer to [“Using ACLs to restrict remote access”](#) on page 857 or [“Restricting remote access to the device to specific IP addresses”](#) on page 860.

In an authentication-method list for a particular access method, you can specify up to seven authentication methods. If the first authentication method is successful, the software grants access and stops the authentication process. If the access is rejected by the first authentication method, the software denies access and stops checking.

However, if an error occurs with an authentication method, the software tries the next method on the list, and so on. For example, if the first authentication method is the RADIUS server, but the link to the server is down, the software will try the next authentication method in the list.

NOTE

If an authentication method is working properly and the password (and user name, if applicable) is not known to that method, this is not an error. The authentication attempt stops, and the user is denied access.

The software will continue this process until either the authentication method is passed or the software reaches the end of the method list. If the Super User level password is not rejected after all the access methods in the list have been tried, access is granted.

Configuration considerations for authentication- method lists

- For CLI access, you must configure authentication-method lists if you want the device to authenticate access using local user accounts or a RADIUS server. Otherwise, the device will authenticate using only the locally based password for the Super User privilege level.
- For devices that can be managed using Brocade Network Advisor, the default authentication method (if no authentication-method list is configured for SNMP) is the CLI Super User level password. If no Super User level password is configured, then access through Brocade Network Advisor is not authenticated.

Examples of authentication-method lists

The following examples show how to configure authentication-method lists. In these examples, the primary authentication method for each is “local”. The device will authenticate access attempts using the locally configured usernames and passwords.

The command syntax for each of the following examples is provided in “[Command Syntax](#)” on page 909.

Example 1

To configure an authentication-method list for SNMP, enter a command such as the following.

```
PowerConnect(config)# aaa authentication snmp-server default local
```

This command allows certain incoming SNMP SET operations to be authenticated using the locally configured usernames and passwords. When this command is enabled, community string validation is not performed for incoming SNMP V1 and V2c packets. This command takes effect as long as the first varbind for SNMP packets is set to one of the following:

- snAgGblPassword="*<username> <password>*" (for AAA method local)
- snAgGblPassword="*<password>*" (for AAA method line, enable)

NOTE

Certain SNMP objects need additional validation. These objects include but are not limited to: **snAgReload**, **snAgWriteNVRAM**, **snAgConfigFromNVRAM**, **snAgImgLoad**, **snAgCfgLoad** and **snAgGblTelnetPassword**. For more information, see **snAgGblPassword** in the *IronWare MIB Reference*.

If AAA is set up to check both the username and password, the string contains the username, followed by a space then the password. If AAA is set up to authenticate with the current Enable or Line password, the string contains the password only.

Note that the above configuration can be overridden by the command **no snmp-server pw-check**, which disables password checking for SNMP SET requests.

Example 2

To configure an authentication-method list for the Privileged EXEC and CONFIG levels of the CLI, enter the following command.

```
PowerConnect(config)# aaa authentication enable default local
```

This command configures the device to use the local user accounts to authenticate attempts to access the Privileged EXEC and CONFIG levels of the CLI.

Example 3

To configure the device to consult a RADIUS server first to authenticate attempts to access the Privileged EXEC and CONFIG levels of the CLI, then consult the local user accounts if the RADIUS server is unavailable, enter the following command.

```
PowerConnect(config)# aaa authentication enable default radius local
```

Command Syntax

The following is the command syntax for the preceding examples.

Syntax: [no] **aaa authentication snmp-server | web-server | enable | login default** *<method1>* [*<method2>*] [*<method3>*] [*<method4>*] [*<method5>*] [*<method6>*] [*<method7>*]

The **snmp-server** | **web-server** | **enable** | **login** parameter specifies the type of access this authentication-method list controls. You can configure one authentication-method list for each type of access.

NOTE

TACACS/TACACS+ and RADIUS are supported only with the **enable** and **login** parameters.

The `<method1>` parameter specifies the primary authentication method. The remaining optional `<method>` parameters specify additional methods to try if an error occurs with the primary method. A method can be one of the values listed in the Method Parameter column in the following table.

TABLE 144 Authentication method values

Method parameter	Description
line	Authenticate using the password you configured for Telnet access. The Telnet password is configured using the enable telnet password... command. Refer to “Setting a Telnet password” on page 866.
enable	Authenticate using the password you configured for the Super User privilege level. This password is configured using the enable super-user-password... command. Refer to “Setting passwords for management privilege levels” on page 866.
local	Authenticate using a local user name and password you configured on the device. Local user names and passwords are configured using the username... command. Refer to “Configuring a local user account” on page 874.
tacacs	Authenticate using the database on a TACACS server. You also must identify the server to the device using the tacacs-server command.
tacacs+	Authenticate using the database on a TACACS+ server. You also must identify the server to the device using the tacacs-server command.
radius	Authenticate using the database on a RADIUS server. You also must identify the server to the device using the radius-server command. Refer to “Configuring RADIUS security” on page 892.
none	Do not use any authentication method. The device automatically permits access.

Configuring SSH2 and SCP

SSH version 2 support

Secure Shell (SSH) is a mechanism for allowing secure remote access to management functions on a device. SSH provides a function similar to Telnet. Users can log into and configure the device using a publicly or commercially available SSH client program, just as they can with Telnet. However, unlike Telnet, which provides no security, SSH provides a secure, encrypted connection to the device.

The SSH2 implementation is compatible with all versions of the SSH2 protocol (2.1, 2.2, and so on). At the beginning of an SSH session, the device negotiates the version of SSH2 to be used. The highest version of SSH2 supported by both the device and the client is the version that is used for the session. Once the SSH2 version is negotiated, the encryption algorithm with the highest security ranking is selected to be used for the session.

PowerConnect device also support Secure Copy (SCP) for securely transferring files between a device and SCP-enabled remote hosts.

NOTE

The SSH feature includes software that is copyright Allegro Software Development Corporation.

SSH2 is supported in the Layer 2 and Layer 3 codes, and SSH version 1 (SSH1) is no longer supported. Refer to [Chapter 30, “Configuring SSH1 and SCP”](#).

SSH2 is a substantial revision of Secure Shell, comprising the following hybrid protocols and definitions:

- SSH Transport Layer Protocol
- SSH Authentication Protocol
- SSH Connection Protocol
- SECSH Public Key File Format
- SSH Fingerprint Format
- SSH Protocol Assigned Numbers
- SSH Transport Layer Encryption Modes
- SCP/SFTP/SSH URI Format

Tested SSH2 clients

The following SSH clients have been tested with SSH2:

- SSH Secure Shell 3.2.3
- Van Dyke SecureCRT 4.0 and 4.1
- F-Secure SSH Client 5.3 and 6.0
- PuTTY 0.54 and 0.56

- OpenSSH 3.5_p1 and 3.6.1p2
- Solaris Sun-SSH-1.0

NOTE

The PowerConnect B-Series T124X devices support client public key sizes of 2048 bits or less.

Supported features

SSH2 (Secure Shell version 2 protocol) provides an SSH server. The SSH server allows secure remote access management functions on a device. SSH provides a function that is similar to Telnet, but unlike Telnet, SSH provides a secure, encrypted connection.

SSH2 support includes the following:

- Key exchange methods are **diffie-hellman-group1-sha1**
- The public key algorithm is **ssh-dss**.
- Encryption is provided with **3des-cbc**, **aes128-cbc**, **aes192-cbc** or **aes256-cbc**. AES encryption has been adopted by the U.S. Government as an encryption standard. Refer to [“AES encryption for SSH2.”](#) on page 912.
- Data integrity is ensured with **hmac-sha1**.
- Supported authentication methods are **Password** and **publickey**.

Unsupported features

The following are not supported with SSH2

- Compression
- TCP/IP port forwarding, X11 forwarding, and secure file transfer
- SSH version 1

AES encryption for SSH2.

Encryption is provided with **3des-cbc**, **aes128-cbc**, **aes192-cbc** or **aes256-cbc**. AES encryption has been adopted by the U.S. Government as an encryption standard.

A total of five SSH connections can be active on a device. To display information about SSH connections, enter the following command.

```
PowerConnect# show ip ssh
Connection  Version  Encryption  Username
1          SSH-2    3des-cbc    Raymond
2          SSH-2    3des-cbc    Ron
3          SSH-2    aes128-cbc  David
4          SSH-2    aes192-cbc  Francesca
5          SSH-2    aes256-cbc  Bob
```

You can also use the **show who** command to display information about SSH connections

```
PowerConnect# show who
  Console connections:
  Established
  you are connecting to this session
  2 minutes 56 seconds in idle

SSH connections:
1. established, client ip address 2.2.2.1, user is Raymond
   1 minutes 15 seconds in idle
2. established, client ip address 2.2.2.2, user is Ron
   2 minutes 25 seconds in idle
3. established, client ip address 2.2.2.1, user is David
   1 minutes 8 seconds in idle
4. established, client ip address 2.2.2.1, user is Franchesca
   2 minutes 32 seconds in idle
5. established, client ip address 2.2.2.3, user is Bob
   5 minutes 17 seconds in idle
```

To terminate an active connection, enter the following command

```
PowerConnect# kill ssh 1
```

Syntax: `kill ssh <connection-id>`

Configuring SSH2

The Dell implementation of SSH2 supports two kinds of user authentication:

- **DSA challenge-response authentication**, where a collection of public keys are stored on the device. Only clients with a private key that corresponds to one of the stored public keys can gain access to the device using SSH.

NOTE

SSH2 supports and validates DSA keys only. It does not support or validate SSH1 RSA keys.

- **Password authentication**, where users attempting to gain access to the device using an SSH client are authenticated with passwords stored on the device or on a TACACS/TACACS+ or RADIUS server

Both kinds of user authentication are enabled by default. You can configure the device to use one or both of them.

Follow the steps given below to configure Secure Shell on a device.

1. If necessary, recreate the SSH keys
2. Generate a host DSA public and private key pair for the device
3. Configure DSA challenge-response authentication
4. Set optional parameters

You can also view information about active SSH connections on the device as well as terminate them.

Recreating SSH keys

You must recreate SSH keys after any one of the following events:

- After upgrading from a software release that supports SSH1, to a software release that supports SSH2.
- After downgrading a software release that supports SSH2, to a software release that supports SSH1

To recreate SSH keys, enter the following command.

```
PowerConnect(config)# crypto key generate
```

Syntax: crypto key generate

Generating a host key pair

When SSH is configured, a public and private **host DSA key pair** is generated for the device. The SSH server on the device uses this host DSA key pair, along with a dynamically generated **server DSA key pair**, to negotiate a session key and encryption method with the client trying to connect to it.

The host DSA key pair is stored in the system-config file of the device. Only the public key is readable. The public key should be added to a “known hosts” file (for example, \$HOME/.ssh/known_hosts on UNIX systems) on the clients who want to access the device. Some SSH client programs add the public key to the known hosts file automatically; in other cases, you must manually create a known hosts file and place the public key of the device in it.

While the SSH listener exists at all times, sessions can not be started from clients until a key is generated. Once a key is generated, clients can start sessions. The keys are also not displayed in the configuration file by default. To display the keys, use the **ssh show-host-keys** command in Privileged EXEC mode.

To generate a public and private DSA host key pair on a device, enter the following command.

```
PowerConnect(config)# crypto key generate
```

When a host key pair is generated, it is saved to the flash memory of all management modules.

To disable SSH2 on a device, enter the following command.

```
PowerConnect(config)# crypto key zeroize
```

When SSH is disabled, it is deleted from the flash memory of all management modules.

Syntax: crypto key generate | zeroize

The **generate** keyword places a DSA host key pair in the flash memory and enables SSH on the device.

The **zeroize** keyword deletes the DSA host key pair from the flash memory and disables SSH on the device.

By default, public keys are hidden in the running configuration. You can optionally configure the device to display the DSA host key pair in the running configuration file, by entering the following command.

```
PowerConnect# ssh show-host-keys
```

Syntax: ssh show-host-keys

To hide the public keys in the running configuration file, enter the following command.

```
PowerConnect# ssh no-show-host-keys
```

Syntax: ssh no-show-host-keys

Providing the public key to clients

If you are using SSH to connect to a device from a UNIX system, you may need to add the public key on the device to a “known hosts” file; for example, \$HOME/.ssh/known_hosts. The following is an example of an entry in a known hosts file.

```
AAAAAB3NzaC1kc3MAAACBAPY8ZOHY2yF5SJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/
z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI140m
leg9e4NnCR1leaQoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfXOD2s9Rd7NBvQAAAIEALN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eolD+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uLlJn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRhtUAAACB
AN7CY+KKvlgHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXGlvo+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEbl1juqnF0GD1B3VvmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
```

Configuring DSA challenge-response authentication

With DSA challenge-response authentication, a collection of clients’ public keys are stored on the device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH.

When DSA challenge-response authentication is enabled, the following events occur when a client attempts to gain access to the device using SSH.

1. The client sends its public key to the device.
2. The device compares the client public key to those stored in memory.
3. If there is a match, the device uses the public key to encrypt a random sequence of bytes.
4. The device sends these encrypted bytes to the client.
5. The client uses its private key to decrypt the bytes.
6. The client sends the decrypted bytes back to the device.
7. The device compares the decrypted bytes to the original bytes it sent to the client. If the two sets of bytes match, it means that the client private key corresponds to an authorized public key, and the client is authenticated.

Setting up DSA challenge-response authentication consists of the following steps.

1. Importing authorized public keys into the device.
2. Enabling DSA challenge response authentication

Importing authorized public keys into the device

SSH clients that support DSA authentication normally provide a utility to generate an DSA key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected. You should collect one public key from each client to be granted access to the device and place all of these keys into one file. This public key file is imported into the device.

The following is an example of a public key file containing one public key.

```
---- BEGIN SSH2 PUBLIC KEY ----
Comment: DSA Public Key
AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yF5JA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/
z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
leg9e4NnCRleaQoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfXOD2s9Rd7NBvQAAAEALN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eolD+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uLlJn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKW0ocJmJ80qadxTRHtUAAACB
AN7CY+KKvlgHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEbl1juqnF0GD1B3VvmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
---- END SSH2 PUBLIC KEY ----
```

You can import the authorized public keys into the active configuration by loading them from a file on a TFTP server. If you import a public key file from a TFTP server, the file is automatically loaded into the active configuration the next time the device is booted.

To cause a public key file called pkeys.txt to be loaded from a TFTP server each time the device is booted, enter a command such as the following.

```
PowerConnect(config)# ip ssh pub-key-file tftp 192.168.1.234 pkeys.txt
```

Syntax: `ip ssh pub-key-file tftp | <tftp-server-ip-addr> <filename> [remove]`

The `<tftp-server-ip-addr>` variable is the IP address of the tftp server that contains the public key file that you want to import into the device.

The `<filename>` variable is the name of the dsa public key file that you want to import into the device.

The **remove** parameter deletes the key from the system.

To display the currently loaded public keys, enter the following command.

```
PowerConnect# show ip client-pub-key
---- BEGIN SSH2 PUBLIC KEY ----
Comment: DSA Public Key
AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yF5JA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/
z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI140m
leg9e4NnCRleaQoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eolD+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uLlJn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXGlV0+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb1ljuqnF0GD1B3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
---- END SSH2 PUBLIC KEY ----
```

Syntax: `show ip client-pub-key [begin <expression> | exclude <expression> | include <expression>]`

To clear the public keys from the buffers, enter the following command.

```
PowerConnect# clear public-key
```

Syntax: `clear public-key`

Use the `ip ssh pub-key remove` command to delete the public key from the system.

Enabling DSA challenge-response authentication

DSA challenge-response authentication is enabled by default. You can disable or re-enable it manually.

To enable DSA challenge-response authentication.

```
PowerConnect(config)# ip ssh key-authentication yes
```

To disable DSA challenge-response authentication.

```
PowerConnect(config)# ip ssh key-authentication no
```

Syntax: `ip ssh key-authentication yes | no`

Setting optional parameters

You can adjust the following SSH settings on the device:

- The number of SSH authentication retries
- The user authentication method the device uses for SSH connections
- Whether the device allows users to log in without supplying a password
- The port number for SSH connections
- The SSH login timeout value
- A specific interface to be used as the source for all SSH traffic from the device
- The maximum idle time for SSH sessions

Setting the number of SSH authentication retries

By default, the device attempts to negotiate a connection with the connecting host three times. The number of authentication retries can be changed to between 1 – 5.

For example, the following command changes the number of authentication retries to 5.

```
PowerConnect(config)# ip ssh authentication-retries 5
```

Syntax: `ip ssh authentication-retries <number>`

Deactivating user authentication

After the SSH server on the device negotiates a session key and encryption method with the connecting client, user authentication takes place. The Dell implementation of SSH supports DSA challenge-response authentication and password authentication.

With DSA challenge-response authentication, a collection of clients' public keys are stored on the device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH.

With password authentication, users are prompted for a password when they attempt to log into the device (provided empty password logins are not allowed). If there is no user account that matches the user name and password supplied by the user, the user is not granted access.

You can deactivate one or both user authentication methods for SSH. Note that deactivating both authentication methods essentially disables the SSH server entirely.

To disable DSA challenge-response authentication, enter the following command.

```
PowerConnect(config)# ip ssh key-authentication no
```

Syntax: `ip ssh key-authentication yes | no`

The default is **yes**.

To deactivate password authentication, enter the following command.

```
PowerConnect(config)# ip ssh password-authentication no
```

Syntax: `ip ssh password-authentication no | yes`

The default is **yes**.

Enabling empty password logins

By default, empty password logins are not allowed. This means that users with an SSH client are always prompted for a password when they log into the device. To gain access to the device, each user must have a user name and password. Without a user name and password, a user is not granted access.

If you enable empty password logins, users are **not** prompted for a password when they log in. Any user with an SSH client can log in without being prompted for a password.

To enable empty password logins, enter the following command.

```
PowerConnect(config)# ip ssh permit-empty-passwd yes
```

Syntax: `ip ssh permit-empty-passwd no | yes`

Setting the SSH port number

By default, SSH traffic occurs on TCP port 22. You can change this port number. For example, the following command changes the SSH port number to 2200.

```
PowerConnect(config)# ip ssh port 2200
```

Note that if you change the default SSH port number, you must configure SSH clients to connect to the new port. Also, you should be careful not to assign SSH to a port that is used by another service. If you change the SSH port number, Dell recommends that you change it to a port number greater than 1024.

Syntax: `ip ssh port <number>`

Setting the SSH login timeout value

When the SSH server attempts to negotiate a session key and encryption method with a connecting client, it waits a maximum of 120 seconds for a response from the client. If there is no response from the client after 120 seconds, the SSH server disconnects. You can change this timeout value to between 1 – 120 seconds. For example, to change the timeout value to 60 seconds, enter the following command.

```
PowerConnect(config)# ip ssh timeout 60
```

Syntax: `ip ssh timeout <seconds>`

Designating an interface as the source for all SSH packets (Layer 3 code only)

You can designate a loopback interface, virtual interface, or Ethernet port as the source for all SSH packets from the device. The software uses the IP address with the numerically lowest value configured on the port or interface as the source IP address for SSH packets originated by the device.

NOTE

When you specify a single SSH source, you can use only that source address to establish SSH management sessions with the device.

To specify the numerically lowest IP address configured on a loopback interface as the device source for all SSH packets, enter commands such as the following.

```
PowerConnect(config)# int loopback 2
PowerConnect(config-lbif-2)# ip address 10.0.0.2/24
PowerConnect(config-lbif-2)# exit
PowerConnect(config)# ip ssh source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all SSH packets from the device.

Syntax: `ip ssh source-interface ethernet <port> | loopback <num> | ve <num>`

The <num> parameter is a loopback interface or virtual interface number. The <port> parameter specifies an ethernet port number.

Example

```
PowerConnect(config)# interface ethernet 4
PowerConnect(config-if-e10000-4)# ip address 209.157.22.110/24
PowerConnect(config-if-e10000-4)# exit
PowerConnect(config)# ip ssh source-interface ethernet 4
```

Configuring the maximum idle time for SSH sessions

By default, SSH sessions do not time out. Optionally, you can set the amount of time an SSH session can be inactive before the device closes it. For example, to set the maximum idle time for SSH sessions to 30 minutes, enter the following command.

```
PowerConnect(config)# ip ssh idle-time 30
```

Syntax: `ip ssh idle-time <minutes>`

If an established SSH session has no activity for the specified number of minutes, the device closes it. An idle time of 0 minutes (the default value) means that SSH sessions never time out. The maximum idle time for SSH sessions is 240 minutes.

Filtering SSH access using ACLs

You can permit or deny SSH access to the device using ACLs. To use ACLs, first create the ACLs you want to use. You can specify a numbered standard IPv4 ACL, a named standard IPv4 ACL

Enter commands such as the following.

```
PowerConnect(config)# access-list 10 permit host 192.168.144.241
PowerConnect(config)# access-list 10 deny host 192.168.144.242 log
PowerConnect(config)# access-list 10 permit host 192.168.144.243
PowerConnect(config)# access-list 10 deny any
PowerConnect(config)# ssh access-group 10
```

Syntax: `ssh access-group <standard-named-acl> | <standard-numbered-acl>`

Terminating an active SSH connection

To terminate one of the active SSH connections, enter the following command

```
PowerConnect# kill ssh 1
```

Syntax: `kill ssh <connection-id>`

Displaying SSH connection information

Up to five SSH connections can be active on the device. To display information about SSH connections, enter the following command.

```
PowerConnect# show ip ssh
Connection Version Encryption Username
1          SSH-2    3des-cbc  Hanuma
2          SSH-2    3des-cbc  Mikaila
3          SSH-2    3des-cbc  Jenny
4          SSH-2    3des-cbc  Mariah
5          SSH-2    3des-cbc  Logan
```

Syntax: `show ip ssh [begin <expression> | exclude <expression> | include <expression>]`

This display shows the following information about the active SSH connections.

TABLE 145 SSH connection information

This field...	Displays...
Connection	The SSH connection ID. This can be from 1 - 5.
Version	The SSH version number. This should always be 1.5.
Encryption	The encryption method used for the connection.
Username	The user name for the connection.

The **show who** command also displays information about SSH connections.

Example

```
PowerConnect# show who
Console connections:
established, monitor enabled, in config mode
2 minutes 17 seconds in idle
Telnet connections (inbound):
1 closed
2 closed
3 closed
4 closed
5 closed
Telnet connection (outbound):
6 closed
SSH connections:
1 established, client ip address 192.168.144.241, user is hanuma
1 minutes 16 seconds in idle
2 established, client ip address 192.168.144.241, user is Mikaila
you are connecting to this session
18 seconds in idle
3 established, client ip address 192.168.144.241, user is Jenny
1 minutes 39 seconds in idle
4 established, client ip address 192.168.144.242, user is Mariah
41 seconds in idle
5 established, client ip address 192.168.144.241, user is Logan
23 seconds in idle
```

Syntax: `show who [begin <expression> | exclude <expression> | include <expression>]`

Using Secure copy with SSH2

Secure Copy (SCP) uses security built into SSH to transfer image and configuration files to and from the device. SCP automatically uses the authentication methods, encryption algorithm, and data compression level configured for SSH. For example, if password authentication is enabled for SSH, the user is prompted for a user name and password before SCP allows a file to be transferred. No additional configuration is required for SCP on top of SSH.

You can use SCP to copy files on the device, including the startup configuration and running configuration files, to or from an SCP-enabled remote host.

Enabling and disabling SCP

SCP is enabled by default and can be disabled. To disable SCP, enter the following command.

```
PowerConnect(config)# ip ssh scp disable
```

Syntax: ip ssh scp disable | enable

NOTE

If you disable SSH, SCP is also disabled.

NOTE

When using SCP, enter the **scp** commands on the SCP-enabled client, rather than the console on the device.

NOTE

Certain SCP client options, including **-p** and **-r**, are ignored by the SCP server on the device. If an option is ignored, the client is notified.

Example file transfers using SCP

The following are examples of using SCP to transfer files to and from a device.

Copying a file to the running config

To copy a configuration file (c:\cfg\brocade.cfg) to the running configuration file on a device at 192.168.1.50 and log in as user terry, enter the following command on the SCP-enabled client.

```
C:\> scp c:\cfg\brocade.cfg terry@192.168.1.50:runConfig
```

If password authentication is enabled for SSH, the user is prompted for user terry password before the file transfer takes place.

Copying a file to the startup config

To copy the configuration file to the startup configuration file, enter the following command.

```
C:\> scp c:\cfg\brocade.cfg terry@192.168.1.50:startConfig
```


Copying the running config file to an SCP-enabled client

To copy the running configuration file on the device to a file called `c:\cfg\fdryrun.cfg` on the SCP-enabled client, enter the following command.

```
C:\> scp terry@192.168.1.50:runConfig c:\cfg\fdryrun.cfg
```

Copying the startup config file to an SCP-enabled client

To copy the startup configuration file on the device to a file called `c:\cfg\fdrystart.cfg` on the SCP-enabled client, enter the following command.

```
C:\> scp terry@192.168.1.50:startConfig c:\cfg\fdry
```

To overwrite the running configuration file

```
C:\> scp c:\cfg\brocade.cfg terry@192.168.1.50:runConfig-overwrite
```

Copying a software image file to flash memory

To copy a software image file from an SCP-enabled client to the primary flash on a device use the following command.

```
C:\> scp TIR04200.bin terry@192.168.1.50:flash:primary
```

To copy a software image file from an SCP-enabled client to the secondary flash on a device , enter the following command.

```
C:\> scp TIR04200.bin terry@192.168.1.50:flash:secondary
```

NOTE

The device supports only one SCP copy session at a time.

Copying a Software Image file from flash memory

To copy a software image from the primary flash on the device to an SCP-enabled client, enter a command such as the following.

```
C:\> scp terry@192.168.1.50:flash:primary TIR04200.bin
```

To copy a software image from the secondary flash on the device to an SCP-enabled client, enter a command such as the following.

```
C:\> scp terry@192.168.1.50:flash:secondary TIR04200.bin
```

NOTE

The device supports only one SCP copy session at a time.

27 Using Secure copy with SSH2

Configuring 802.1X Port Security

IETF RFC support

When a user logs on to a network that uses 802.1X port security, the device grants (or does not grant) access to network services after the user is authenticated by an authentication server. The user-based authentication in 802.1X port security provides an alternative to granting network access based on a user IP address, MAC address, or subnetwork.

The Dell implementation of 802.1X port security supports the following RFCs:

- RFC 2284 PPP Extensible Authentication Protocol (EAP)
- RFC 2865 Remote Authentication Dial In User Service (RADIUS)
- RFC 2869 RADIUS Extensions

How 802.1X port security works

This section explains the basic concepts behind 802.1X port security, including device roles, how the devices communicate, and the procedure used for authenticating clients.

NOTE

802.1X Port Security cannot be configured on MAC Port Security-enabled ports.


Device roles in an 802.1X configuration

The 802.1X standard defines the roles of *Client/Supplicant*, *Authenticator*, and *Authentication Server* in a network.

The Client (known as a **Supplicant** in the 802.1X standard) provides username/password information to the Authenticator. The Authenticator sends this information to the Authentication Server. Based on the Client's information, the Authentication Server determines whether the Client can use services provided by the Authenticator. The Authentication Server passes this information to the Authenticator, which then provides services to the Client, based on the authentication result.

Figure 117 illustrates these roles.

FIGURE 117 Authenticator, client/supplicant, and authentication server in an 802.1X configuration



Authenticator – The device that controls access to the network. In an 802.1X configuration, the device serves as the Authenticator. The Authenticator passes messages between the Client and the Authentication Server. Based on the identity information supplied by the Client, and the authentication information supplied by the Authentication Server, the Authenticator either grants or does not grant network access to the Client.

Client/Supplicant – The device that seeks to gain access to the network. Clients must be running software that supports the 802.1X standard (for example, the Windows XP operating system). Clients can either be directly connected to a port on the Authenticator, or can be connected by way of a hub.


Authentication server – The device that validates the Client and specifies whether or not the Client may access services on the device. Dell supports Authentication Servers running RADIUS.

Communication between the devices

For communication between the devices, 802.1X port security uses the **Extensible Authentication Protocol** (EAP), defined in RFC 2284. The 802.1X standard specifies a method for encapsulating EAP messages so that they can be carried over a LAN. This encapsulated form of EAP is known as EAP over LAN (**EAPOL**). The standard also specifies a means of transferring the EAPOL information between the Client/Supplicant, Authenticator, and Authentication Server.

EAPOL messages are passed between the **Port Access Entity (PAE)** on the Supplicant and the Authenticator. [Figure 118](#) shows the relationship between the Authenticator PAE and the Supplicant PAE.

FIGURE 118 Authenticator PAE and supplicant PAE



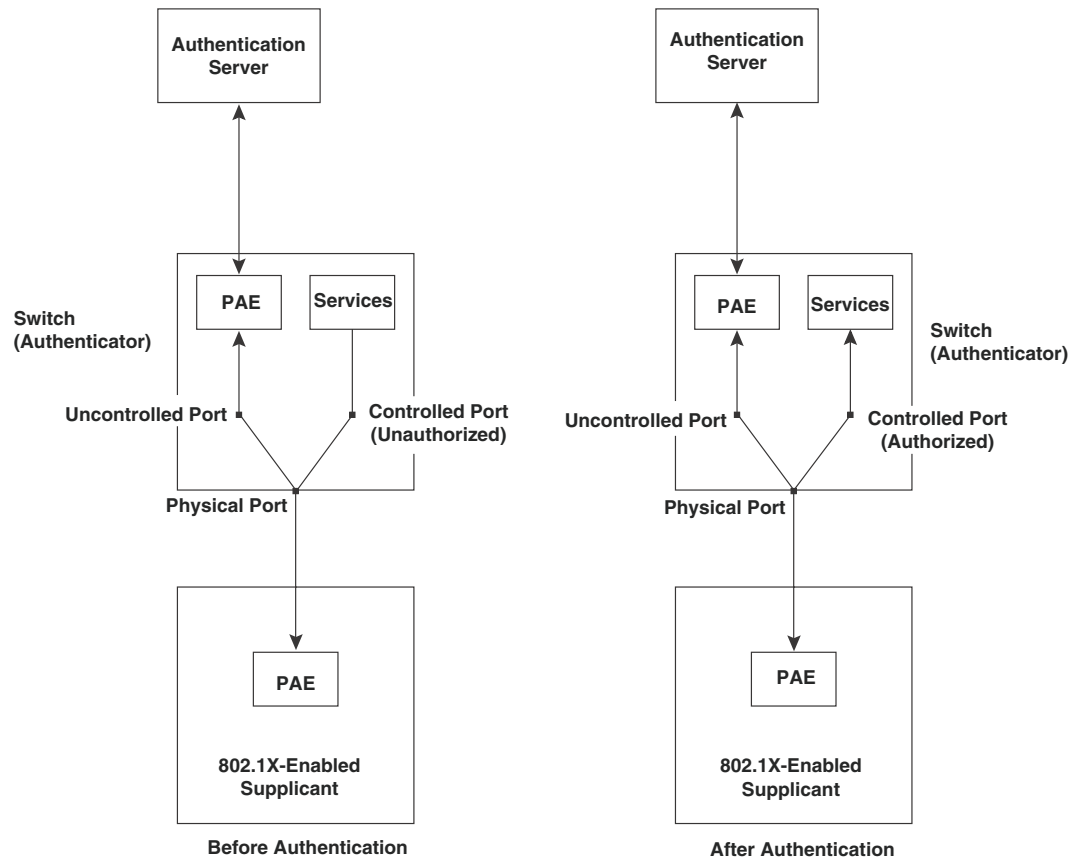
Authenticator PAE – The Authenticator PAE communicates with the Supplicant PAE, receiving identifying information from the Supplicant. Acting as a RADIUS client, the Authenticator PAE passes the Supplicant information to the Authentication Server, which decides whether the Supplicant can gain access to the port. If the Supplicant passes authentication, the Authenticator PAE grants it access to the port.

Supplicant PAE – The Supplicant PAE supplies information about the Client to the Authenticator PAE and responds to requests from the Authenticator PAE. The Supplicant PAE can also initiate the authentication procedure with the Authenticator PAE, as well as send log off messages.

Controlled and uncontrolled ports

A physical port on the device used with 802.1X port security has two virtual access points a **controlled** port and an **uncontrolled** port. The controlled port provides full access to the network. The uncontrolled port provides access only for EAPOL traffic between the Client and the Authentication Server. When a Client is successfully authenticated, the controlled port is opened to the Client. [Figure 119](#) illustrates this concept.

FIGURE 119 Controlled and uncontrolled ports before and after client authentication



Before a Client is authenticated, only the uncontrolled port on the Authenticator is open. The uncontrolled port allows only EAPOL frames to be exchanged between the Client and the Authentication Server. The controlled port is in the unauthorized state and allows no traffic to pass through.

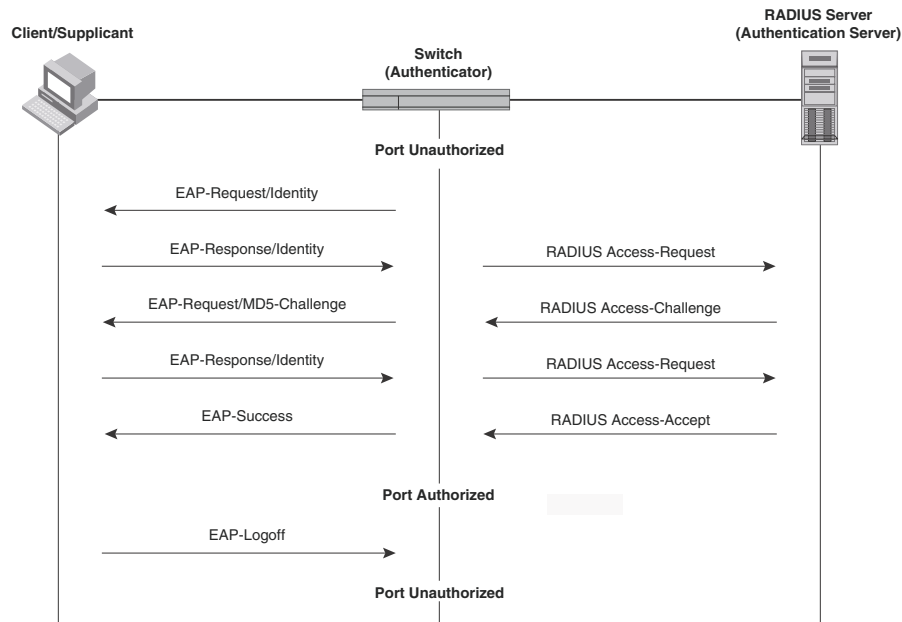
During authentication, EAPOL messages are exchanged between the Supplicant PAE and the Authenticator PAE, and RADIUS messages are exchanged between the Authenticator PAE and the Authentication Server. Refer to [“Message exchange during authentication”](#) on page 929 for an example of this process. If the Client is successfully authenticated, the controlled port becomes authorized, and traffic from the Client can flow through the port normally.

By default, all controlled ports on the device are placed in the authorized state, allowing all traffic. When authentication is activated on an 802.1X-enabled interface, the interface controlled port is placed initially in the unauthorized state. When a Client connected to the port is successfully authenticated, the controlled port is then placed in the authorized state until the Client logs off. Refer to [“Enabling 802.1X port security”](#) on page 945 for more information.

Message exchange during authentication

Figure 120 illustrates a sample exchange of messages between an 802.1X-enabled Client, a switch acting as Authenticator, and a RADIUS server acting as an Authentication Server.

FIGURE 120 Message exchange between client/supplicant, authenticator, and authentication server



In this example, the Authenticator initiates communication with an 802.1X-enabled Client. When the Client responds, it is prompted for a username (48 characters maximum) and password. The Authenticator passes this information to the Authentication Server, which determines whether the Client can access services provided by the Authenticator. When the Client is successfully authenticated by the RADIUS server, the port is authorized. When the Client logs off, the port becomes unauthorized again.

The Dell 802.1X implementation supports dynamic VLAN assignment. If one of the attributes in the Access-Accept message sent by the RADIUS server specifies a VLAN identifier, and this VLAN is available on the device, the client port is moved from its default VLAN to the specified VLAN. When the client disconnects from the network, the port is placed back in its default VLAN. Refer to [“Configuring dynamic VLAN assignment for 802.1X ports”](#) on page 938 for more information.

If a Client does not support 802.1X, authentication cannot take place. The device sends EAP-Request/Identity frames to the Client, but the Client does not respond to them.

When a Client that supports 802.1X attempts to gain access through a non-802.1X-enabled port, it sends an EAP start frame to the device. When the device does not respond, the Client considers the port to be authorized, and starts sending normal traffic.

Devices support Identity and MD5-challenge requests in EAP Request/Response messages.

NOTE

Refer to also [“EAP pass-through support”](#) on page 930.

- **EAP-TLS (RFC 2716)** – EAP Transport Level Security (TLS) provides strong security by requiring both client and authentication server to be identified and validated through the use of public key infrastructure (PKI) digital certificates. EAP-TLS establishes a tunnel between the client and the authentication server to protect messages from unauthorized users' eavesdropping activities. Since EAP-TLS requires PKI digital certificates on both the clients and the authentication servers, the roll out, maintenance, and scalability of this authentication method is much more complex than other methods. EAP-TLS is best for installations with existing PKI certificate infrastructures.
- **EAP-TTLS (Internet-Draft)** – The EAP Tunnelled Transport Level Security (TTLS) is an extension of EAP-TLS. Like TLS, EAP-TTLS provides strong authentication; however it requires only the authentication server to be validated by the client through a certificate exchange between the server and the client. Clients are authenticated by the authentication server using user names and passwords.

A TLS tunnel can be used to protect EAP messages and existing user credential services such as Active Directory, RADIUS, and LDAP. Backward compatibility for other authentication protocols such as PAP, CHAP, MS-CHAP, and MS-CHAP-V2 are also provided by EAP-TTLS. EAP-TTLS is not considered foolproof and can be fooled into sending identity credentials if TLS tunnels are not used. EAP-TTLS is suited for installations that require strong authentication without the use of mutual PKI digital certificates.

- **PEAP (Internet-Draft)** – Protected EAP Protocol (PEAP) is an Internet-Draft that is similar to EAP-TTLS. PEAP client authenticates directly with the backend authentication server. The authenticator acts as a pass-through device, which does not need to understand the specific EAP authentication protocols.

Unlike EAP-TTLS, PEAP does not natively support user name and password to authenticate clients against an existing user database such as LDAP. PEAP secures the transmission between the client and authentication server with a TLS encrypted tunnel. PEAP also allows other EAP authentication protocols to be used. It relies on the mature TLS keying method for its key creation and exchange. PEAP is best suited for installations that require strong authentication without the use of mutual certificates.

NOTE

If the 802.1X Client will be sending a packet that is larger than 1500 bytes, you must enable **jumbo** at the Global config level of the CLI.

Configuration for these challenge types is the same as for the EAP-MD5 challenge type.

EAP pass-through support

EAP pass-through support is fully compliant with RFC 3748, in which, by default, compliant pass-through authenticator implementations forward EAP challenge request packets of any type.


Configuration notes

- If the 802.1X supplicant or authentication server will be sending packets that are greater than 1500 MTU, you should configure the device to accommodate a bigger buffer size.
- EAP pass-through is supported on the PowerConnect B-Series TI24X devices.

Authenticating multiple hosts connected to the same port

Devices support 802.1X authentication for ports with more than one host connected to them. [Figure 121](#) illustrates a sample configuration where multiple hosts are connected to a single 802.1X port.

FIGURE 121 Multiple hosts connected to a single 802.1X-enabled port



By default, traffic from hosts that cannot be authenticated by the RADIUS server is dropped in hardware. You can optionally configure the device to assign the port to a “restricted” VLAN if authentication of the Client is unsuccessful.

How 802.1X Multiple-host authentication works

When multiple hosts are connected to a single 802.1X-enabled port on a device (as in [Figure 121](#)), 802.1X authentication is performed in the following way.

1. One of the 802.1X-enabled Clients attempts to log into a network in which a device serves as an Authenticator.
2. The device creates an internal session (called a **dot1x-mac-session**) for the Client. A dot1x-mac-session serves to associate a Client MAC address and username with its authentication status.
3. The device performs 802.1X authentication for the Client. Messages are exchanged between the device and the Client, and between the device and the Authentication Server (RADIUS server). The result of this process is that the Client is either successfully authenticated or not authenticated, based on the username and password supplied by the client.
4. If the Client is successfully authenticated, the Client dot1x-mac-session is set to “access-is-allowed”. This means that traffic from the Client can be forwarded normally.

5. If authentication for the Client is unsuccessful the first time, multiple attempts to authenticate the client will be made as determined by the **attempts** variable in the **auth-fail-max-attempts** command.
 - Refer to [“Specifying the number of authentication attempts the device makes before dropping packets”](#) on page 950 for information on how to do this.
6. If authentication for the Client is unsuccessful more than the number of times specified by the attempts variable in the **auth-fail-max-attempts** command, an **authentication-failure** action is taken. The authentication-failure action can be either to drop traffic from the Client, or to place the port in a “restricted” VLAN:
 - If the authentication-failure action is to drop traffic from the Client, then the Client dot1x-mac-session is set to “access-denied”, causing traffic from the Client to be dropped in hardware.
 - If the authentication-failure action is to place the port in a “restricted” VLAN, If the Client dot1x-mac-session is set to “access-restricted” then the port is moved to the specified restricted VLAN, and traffic from the Client is forwarded normally.
7. When the Client disconnects from the network, the device deletes the Client dot1x-mac-session. This does not affect the dot1x-mac-session or authentication status (if any) of the other hosts connected on the port.

Configuration notes

- The Client dot1x-mac-session establishes a relationship between the username and MAC address used for authentication. If a user attempts to gain access from different Clients (with different MAC addresses), he or she would need to be authenticated from each Client.
- If a Client has been denied access to the network (that is, the Client dot1x-mac-session is set to “access-denied”), then you can cause the Client to be re-authenticated by manually disconnecting the Client from the network, or by using the **clear dot1x mac-session** command. Refer to [“Clearing a dot1x-mac-session for a MAC address”](#) on page 951 for information on this command.
- When a Client has been denied access to the network, its dot1x-mac-session is aged out if no traffic is received from the Client MAC address over a fixed hardware aging period (70 seconds), plus a configurable software aging period. You can optionally change the software aging period for dot1x-mac-sessions or disable aging altogether. After the denied Client dot1x-mac-session is aged out, traffic from that Client is no longer blocked, and the Client can be re-authenticated.

In addition, you can configure disable aging for the dot1x-mac-session of Clients that have been granted either full access to the network, or have been placed in a restricted VLAN. After a Client dot1x-mac-session ages out, the Client must be re-authenticated. Refer to [“Disabling aging for dot1x-mac-sessions”](#) on page 950 for more information.

- Dynamic IP ACL and MAC address filter assignment is supported in an 802.1X multiple-host configuration. Refer to [“Dynamically applying IP ACLs and MAC filters to 802.1X ports”](#) on page 941.
- 802.1X multiple-host authentication has the following additions:
 - Configurable hardware aging period for denied client dot1x-mac-sessions. Refer to [“Configurable hardware aging period for denied client dot1x-mac-sessions”](#) on page 933.
 - Dynamic ACL and MAC address filter assignment in 802.1X multiple-host configurations. Refer to [“Dynamically applying IP ACLs and MAC filters to 802.1X ports”](#) on page 941.

- Dynamic multiple VLAN assignment for 802.1X ports. Refer [“Dynamic multiple VLAN assignment for 802.1X ports”](#) on page 939.
- Configure a restriction to forward authenticated and unauthenticated tagged and untagged clients to a restricted VLAN.
- Configure an override to send failed dot1x and non-dot1x clients to a restricted VLAN.
- Configure VLAN assignments for clients attempting to gain access through dual-mode ports.
- Enhancements to some **show** commands.
- Differences in command syntax for saving dynamic VLAN assignments to the startup-config file.

Configurable hardware aging period for denied client dot1x-mac-sessions

When one of the 802.1X-enabled Clients in a multiple-host configuration attempts to log into a network in which a device serves as an Authenticator, the device creates a dot1x-mac-session for the Client.

When a Client has been denied access to the network, its dot1x-mac-session is aged out if no traffic is received from the Client MAC address over a period of time. After a denied Client dot1x-mac-session ages out, the Client can be re-authenticated. Aging of a denied Client's dot1x-mac-session occurs in two phases, known as hardware aging and software aging.

The hardware aging period for a denied Client's dot1x-mac-session is not fixed at 70 seconds. The hardware aging period for a denied Client's dot1x-mac-session is equal to the length of time specified with the dot1x **timeout quiet-period** command. By default, the hardware aging time is 60 seconds. Once the hardware aging period ends, the software aging period begins. When the software aging period ends, the denied Client's dot1x-mac-session ages out, and the Client can be authenticated again.

802.1X port security and sFlow

sFlow is a standards-based protocol that allows network traffic to be sampled at a user-defined rate for the purpose of monitoring traffic flow patterns and identifying packet transfer rates on user-specified interfaces.

When you enable sFlow forwarding on an 802.1X-enabled interface, the samples taken from the interface include the user name string at the inbound or outbound port, or both, if that information is available.

For more information on sFlow, refer to [Appendix A, “Network Monitoring”](#).

Configuring 802.1X port security

Configuring 802.1X port security on a device consists of the following tasks.

1. Configure the device interaction with the Authentication Server:
 - [“Configuring an authentication method list for 802.1X”](#) on page 934
 - [“Setting RADIUS parameters”](#) on page 934
 - [“Configuring dynamic VLAN assignment for 802.1X ports”](#) on page 938 (optional)

- “Dynamically applying IP ACLs and MAC filters to 802.1X ports” on page 941
2. Configure the device role as the Authenticator:
 - “Enabling 802.1X port security” on page 945
 - “Initializing 802.1X on a port” on page 949 (optional)
 3. Configure the device interaction with Clients:
 - “Configuring periodic re-authentication” on page 946 (optional)
 - “Re-authenticating a port manually” on page 947 (optional)
 - “Setting the quiet period” on page 947 (optional)
 - “Setting the wait interval for EAP frame retransmissions” on page 947 (optional)
 - “Setting the maximum number of EAP frame retransmissions” on page 948 (optional)
 - “Specifying a timeout for retransmission of messages to the authentication server” on page 949 (optional)
 - “Allowing access to multiple hosts” on page 949 (optional)

Configuring an authentication method list for 802.1X

To use 802.1X port security, you must specify an authentication method to be used to authenticate Clients. Dell supports RADIUS authentication with 802.1X port security. To use RADIUS authentication with 802.1X port security, you create an authentication method list for 802.1X and specify RADIUS as an authentication method, then configure communication between the device and RADIUS server.

Example

```
PowerConnect(config)# aaa authentication dot1x default radius
```

Syntax: [no] **aaa authentication dot1x default** <method-list>

For the <method-list>, enter at least one of the following authentication methods

radius – Use the list of all RADIUS servers that support 802.1X for authentication.

none – Use no authentication. The Client is automatically authenticated without the device using information supplied by the Client.

NOTE

If you specify both **radius** and **none**, make sure **radius** comes before **none** in the method list.

Setting RADIUS parameters

To use a RADIUS server to authenticate access to a device, you must identify the server to the device.

Example

```
PowerConnect(config)# radius-server host 209.157.22.99 auth-port 1812 acct-port 1813 default key mirabeau dot1x
```

Syntax: **radius-server host** <ip-addr> | <ipv6-addr> | <server-name> [**auth-port** <num> | **acct-port** <num> | **default**] [**key** 0 | 1 <string>] [**dot1x**]

The **host** `<ip-addr> | <ipv6-addr> | <server-name>` parameter is either an IP address or an ASCII text string.

The **dot1x** parameter indicates that this RADIUS server supports the 802.1X standard. A RADIUS server that supports the 802.1X standard can also be used to authenticate non-802.1X authentication requests.

NOTE

To implement 802.1X port security, at least one of the RADIUS servers identified to the device must support the 802.1X standard.

Supported RADIUS attributes

Many IEEE 802.1X Authenticators will function as RADIUS clients. Some of the RADIUS attributes may be received as part of IEEE 802.1X authentication. devices support the following RADIUS attributes for IEEE 802.1X authentication:

- Username (1) – RFC 2865
- NAS-IP-Address (4) – RFC 2865
- NAS-Port (5) – RFC 2865
- Service-Type (6) – RFC 2865
- FilterId (11) – RFC 2865
- Framed-MTU (12) – RFC 2865
- State (24) – RFC 2865
- Vendor-Specific (26) – RFC 2865
- Session-Timeout (27) – RFC 2865
- Termination-Action (29) – RFC 2865
- Calling-Station-ID (31) – RFC 2865
- NAS-Port-Type (61) § RFC 2865
- Tunnel-Type (64) – RFC 2868
- Tunnel-Medium-Type (65) – RFC 2868
- EAP Message (79) – RFC 2579
- Message-Authenticator (80) RFC 3579
- Tunnel-Private-Group-Id (81) – RFC 2868
- NAS-Port-id (87) – RFC 2869

Specifying the RADIUS timeout action

A RADIUS timeout occurs when the device does not receive a response from a RADIUS server within a specified time limit and after a certain number of retries. The time limit and number of retries can be manually configured using the CLI commands **radius-server timeout** and **radius-server retransmit**, respectively. If the parameters are not manually configured, the device applies the default value of three seconds time limit with a maximum of three retries.

A **pass** essentially bypasses the authentication process and permits user access to the network. A **fail** bypasses the authentication process and blocks user access to the network, unless **restrict-vlan** is configured, in which case, the user is placed into a VLAN with restricted or limited access. By default, the device will reset the authentication process and retry to authenticate the user.

Specify the RADIUS timeout action at the Interface level of the CLI.

Permit user access to the network after a RADIUS timeout

To set the RADIUS timeout behavior to bypass 802.1X authentication and *permit* user access to the network, enter commands such as the following

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x auth-timeout-action success
```

Syntax: [no] **dot1x auth-timeout-action success**

Once the *success* timeout action is enabled, use the **no** form of the command to reset the RADIUS timeout behavior to *retry*.

Re-authenticate a user

To configure RADIUS timeout behavior to bypass multi-device port authentication and *permit* user access to the network, enter commands similar to the following

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x re-auth-timeout-success 60
```

Syntax: [no] **dot1x re-auth-timeout-success <seconds>**

The **<seconds>** parameter specifies the number of seconds the device will wait to re-authenticate a user after a timeout. The minimum value is 10 seconds. The maximum value is $2^{16}-1$ (maximum unsigned 16-bit value).

Deny user access to the network after a RADIUS timeout

To set the RADIUS timeout behavior to bypass 802.1X authentication and *block* user access to the network, enter commands such as the following

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x auth-timeout-action failure
```

Syntax: [no] **dot1x auth-timeout-action failure**

Once the failure timeout action is enabled, use the **no** form of the command to reset the RADIUS timeout behavior to *retry*.

NOTE

If **restrict-vlan** is configured along with **auth-timeout-action failure**, the user will be placed into a VLAN with restricted or limited access. Refer to [“Allow user access to a restricted VLAN after a RADIUS timeout”](#) on page 936.

Allow user access to a restricted VLAN after a RADIUS timeout

To set the RADIUS timeout behavior to bypass 802.1X authentication and place the user in a VLAN with restricted or limited access, enter commands such as the following

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x auth-fail-action restrict-vlan 100
PowerConnect(config-if-e100-1)# dot1x auth-timeout-action failure
```

Syntax: [no] dot1x auth-fail-action restrict-vlan [<vlan-id>]

Syntax: [no] dot1x auth-timeout-action failure

Send a failed Dot1X client to a restricted VLAN

In [Figure 122](#), a VoIP phone sends both tagged and untagged traffic to dual-mode port e 3. Assuming the VoIP phone is authenticated to a voice VLAN as tagged, a MAC session for the VoIP phone is learned on the voice VLAN. In addition, since the phone sends untagged traffic, a MAC session is also learned on the native untagged VLAN (based on the VLAN dual-mode configuration).

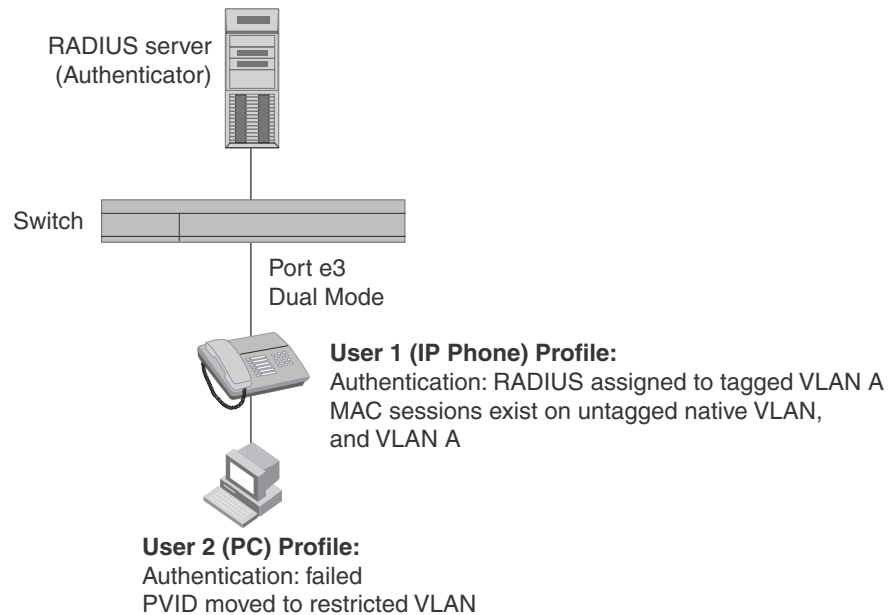
Use the **auth-fail-force-restrict** command to override the VoIP MAC session on the native VLAN, and move the PVID for the port to the restricted VLAN. Future untagged traffic from both phone and client establishes MAC sessions on the restricted VLAN, for restricted access.

This command is configured under the global **dot1x-enable** command as follows

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# auth-fail-force-restrict
```

Syntax: auth-fail-force-restrict

FIGURE 122 Redirecting clients to a restricted VLAN



After authentication fails for User 2, and the PVID moves to the restricted VLAN, there will be a total of 3 MAC sessions on port e 3:

- one tagged MAC session on VLAN A for the phone
- one untagged MAC session on the restricted VLAN for the phone
- one untagged MAC session on the restricted VLAN for the client

Configuring dynamic VLAN assignment for 802.1X ports

When a client successfully completes the EAP authentication process, the Authentication Server (the RADIUS server) sends the Authenticator (the Dell device) a RADIUS Access-Accept message that grants the client access to the network. The RADIUS Access-Accept message contains attributes set for the user in the user's access profile on the RADIUS server.

If one of the attributes in the Access-Accept message specifies a VLAN identifier, and if this VLAN is available on the Dell device, the client port is moved from its default VLAN to this specified VLAN.

NOTE

This feature is supported on port-based VLANs only. This feature cannot be used to place an 802.1X-enabled port into a Layer 3 protocol VLAN.

Automatic removal of dynamic VLAN assignments for 802.1X ports

For increased security, this feature removes any association between a port and a dynamically-assigned VLAN when all 802.1x sessions for that VLAN have expired on the port.

NOTE

When a **show run** command is issued during a session, the dynamically-assigned VLAN is not displayed.

Enable 802.1X VLAN ID support by adding the following attributes to a user profile on the RADIUS server.

Table 6:

Attribute name	Type	Value
Tunnel-Type	064	13 (decimal) – VLAN
Tunnel-Medium-Type	065	6 (decimal) – 802
Tunnel-Private-Group-ID	081	<vlan-name> (string) – either the name or the number of a VLAN configured on the Dell device.

The device reads the attributes as follows:

- If the Tunnel-Type or the Tunnel-Medium-Type attributes in the Access-Accept message do not have the values specified above, the Dell device ignores the three Attribute-Value pairs. The client becomes authorized, but the client port is not dynamically placed in a VLAN.
- If the Tunnel-Type or the Tunnel-Medium-Type attributes in the Access-Accept message do have the values specified above, but there is no value specified for the Tunnel-Private-Group-ID attribute, the client will not become authorized.
- When the Dell device receives the value specified for the Tunnel-Private-Group-ID attribute, it checks whether the <vlan-name> string matches the name of a VLAN configured on the device. If there is a VLAN on the device whose name matches the <vlan-name> string, then the client port is placed in the VLAN whose ID corresponds to the VLAN name.
- If the <vlan-name> string does not match the name of a VLAN, the Dell device checks whether the string, when converted to a number, matches the ID of a VLAN configured on the device. If it does, then the client port is placed in the VLAN with that ID.
- If the <vlan-name> string does not match either the name or the ID of a VLAN configured on the device, then the client will not become authorized.

The **show interface** command displays the VLAN to which an 802.1X-enabled port has been dynamically assigned, as well as the port from which it was moved (that is, the port default VLAN). Refer to [“Displaying dynamically assigned VLAN information”](#) on page 957 for sample output indicating the port dynamically assigned VLAN.

Dynamic multiple VLAN assignment for 802.1X ports

When you add attributes to a user profile on the RADIUS server, the `<vlan-name>` value for the Tunnel-Private-Group-ID attribute can specify the name or number of one or more VLANs configured on the Dell device.

For example, to specify one VLAN, configure the following for the `<vlan-name>` value in the Tunnel-Private-Group-ID attribute on the RADIUS server.

"10" or "marketing"

In this example, the port on which the Client is authenticated is assigned to VLAN 10 or the VLAN named "marketing". The VLAN to which the port is assigned must have previously been configured on the Dell device.

To specify an untagged VLAN, use the following.

"U:10" or "U:marketing"

When the RADIUS server specifies an untagged VLAN ID, the port default VLAN ID (or **PVID**) is changed from the system DEFAULT-VLAN (VLAN 1) to the specified VLAN ID. The port transmits only untagged traffic on its PVID. In this example, the port PVID is changed from VLAN 1 (the DEFAULT-VLAN) to VLAN 10 or the VLAN named "marketing".

The PVID for a port can be changed only once through RADIUS authentication. For example, if RADIUS authentication for a Client causes a port PVID to be changed from 1 to 10, and then RADIUS authentication for another Client on the same port specifies that the port PVID be moved to 20, then the second PVID assignment from the RADIUS server is ignored.

If the link goes down, or the dot1x-mac-session for the Client that caused the initial PVID assignment ages out, then the port reverts back to its original (non-RADIUS-specified) PVID, and subsequent RADIUS authentication can change the PVID assignment for the port.

If a port PVID is assigned through the multi-device port authentication feature, and 802.1X authentication subsequently specifies a different PVID, then the PVID specified through 802.1X authentication overrides the PVID specified through multi-device port authentication.

To specify tagged VLANs, use the following.

"T:12;T:20" or "T:12;T:marketing"

In this example, the port is added to VLANs 12 and 20 or VLANs 12 and the VLAN named "marketing". When a tagged packet is authenticated, and a list of VLANs is specified on the RADIUS server for the MAC address, then the packet tag must match one of the VLANs in the list in order for the Client to be successfully authenticated. If authentication is successful, then the port is added to all of the VLANs specified in the list.

Unlike with a RADIUS-specified untagged VLAN, if the dot1x-mac-session for the Client ages out, the port membership in RADIUS-specified tagged VLANs is not changed. In addition, if multi-device port authentication specifies a different list of tagged VLANs, then the port is added to the specified list of VLANs. Membership in the VLANs specified through 802.1X authentication is not changed.

To specify an untagged VLAN and multiple tagged VLANs, use the following.

"U:10;T:12;T:marketing"

When the RADIUS server returns a value specifying both untagged and tagged VLAN IDs, the port becomes a dual-mode port, accepting and transmitting both tagged traffic and untagged traffic at the same time. A dual-mode port transmits only untagged traffic on its default VLAN (PVID) and only tagged traffic on all other VLANs.

In this example, the port VLAN configuration is changed so that it transmits untagged traffic on VLAN 10, and transmits tagged traffic on VLAN 12 and the VLAN named "marketing".

For a configuration example, refer to ["802.1X Authentication with dynamic VLAN assignment"](#) on page 965.

Saving dynamic VLAN assignments to the running-config file

You can configure the Dell device to save the RADIUS-specified VLAN assignments to the device's running-config file. Enter commands such as the following.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# save-dynamicvlan-to-config
```

Syntax: save-dynamicvlan-to-config

By default, the dynamic VLAN assignments are not saved to the running-config file. Entering the **show running-config** command does not display dynamic VLAN assignments, although they can be displayed with the **show vlan** and **show authenticated-mac-address detail** commands.

NOTE

When this feature is enabled, issuing the command **write mem** will save any dynamic VLAN assignments to the startup configuration file.

Considerations for dynamic VLAN assignment in an 802.1X multiple-host configuration

The following considerations apply when a Client in a 802.1X multiple-host configuration is successfully authenticated, and the RADIUS Access-Accept message specifies a VLAN for the port:

- If the port is not already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of a valid VLAN on the Dell device, then the port is placed in that VLAN.
- If the port is already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of a different VLAN, then it is considered an authentication failure. The port VLAN membership is not changed.
- If the port is already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of that same VLAN, then traffic from the Client is forwarded normally.
- If the RADIUS Access-Accept message specifies the name or ID of a VLAN that does not exist on the Dell device, then it is considered an authentication failure.
- If the port is a tagged or dual-mode port, and the RADIUS Access-Accept message specifies the name or ID of a valid VLAN on the Dell device, then the port is placed in that VLAN. If the port is already a member of the RADIUS-specified VLAN, no further action is taken. Note that the Client dot1x-mac-session is set to "access-is-allowed" for the RADIUS-specified VLAN only. If traffic from the Client MAC address is received on any other VLAN, it is dropped.

- If the RADIUS Access-Accept message does not contain any VLAN information, the Client dot1x-mac-session is set to “access-is-allowed”. If the port is already in a RADIUS-specified VLAN, it remains in that VLAN.

Using dynamic VLAN assignment with the MAC port security feature

MAC port security allows the Dell device to learn a limited number of “secure” MAC addresses on an interface. The interface forwards only packets with source MAC addresses that match these secure addresses. If the interface receives a packet with a source MAC address that is different from any of the secure addresses, it is considered a security violation, and subsequent packets from the violating MAC address can be dropped, or the port can be disabled entirely.

If a port is disabled due to a MAC port security violation, 802.1X clients attempting to connect over the port cannot be authorized. In addition, 802.1X clients connecting from non-secure MAC addresses cannot be authorized.

To use 802.1X dynamic VLAN assignment with the MAC port security feature on an interface, you must set the number of secure MAC addresses to two or more.

Example

```
PowerConnect(config)# int e 2
PowerConnect(config-if-e10000-2)# port security
PowerConnect(config-port-security-e10000-2)# maximum 2
PowerConnect(config-port-security-e10000-2)# exit
```

Refer to [Chapter 29, “Using the MAC Port Security Feature”](#) for more information.

Dynamically applying IP ACLs and MAC filters to 802.1X ports

The Dell 802.1X implementation supports dynamically applying an IP ACL or MAC address filter to a port, based on information received from an Authentication Server.

When a client/supplicant successfully completes the EAP authentication process, the Authentication Server (the RADIUS server) sends the Authenticator (the Dell device) a RADIUS Access-Accept message that grants the client access to the network. The RADIUS Access-Accept message contains attributes set for the user in the user's access profile on the RADIUS server.

If the Access-Accept message contains Filter-ID (type 11) or Vendor-Specific (type 26), or both attributes, the Dell device can use information in these attributes to apply an IP ACL or MAC address filter to the authenticated port. This IP ACL or MAC address filter applies to the port for as long as the client is connected to the network. When the client disconnects from the network, the IP ACL or MAC address filter is no longer applied to the port. If an IP ACL or MAC address filter had been applied to the port prior to 802.1X authentication, it is then re-applied to the port.

The Dell device uses information in the Filter ID and Vendor-Specific attributes as follows:

- The Filter-ID attribute can specify the number of an existing IP ACL or MAC address filter configured on the Dell device. In this case, the IP ACL or MAC address filter with the specified number is applied to the port.
- The Vendor-Specific attribute can specify actual syntax for a Dell IP ACL or MAC address filter, which is then applied to the authenticated port. Configuring a Vendor-Specific attribute in this way allows you to create IP ACLs and MAC filters that apply to individual users; that is, **per-user** IP ACLs or MAC address filters.

Configuration considerations

The following restrictions apply to dynamic IP ACLs or MAC address filters:

- Inbound dynamic IP ACLs are supported. Outbound dynamic ACLs are not supported.
- Inbound Vendor-Specific attributes are supported. Outbound Vendor-Specific attributes are not supported.
- A maximum of one IP ACL can be configured in the inbound direction on an interface.
- MAC address filters cannot be configured in the outbound direction on an interface.
- Concurrent operation of MAC address filters and IP ACLs is not supported.

Disabling and enabling strict security mode for dynamic filter assignment

By default, 802.1X dynamic filter assignment operates in **strict security mode**. When strict security mode is enabled, 802.1X authentication for a port fails if the Filter-ID attribute contains invalid information, or if insufficient system resources are available to implement the per-user IP ACLs or MAC address filters specified in the Vendor-Specific attribute.

When strict security mode is enabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the port will not be authenticated, regardless of any other information in the message (for example, if the Tunnel-Private-Group-ID attribute specifies a VLAN on which to assign the port).
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port will not be authenticated.
- If the device does not have the system resources available to dynamically apply a filter to a port, then the port will not be authenticated.

NOTE

If the Access-Accept message contains values for both the Filter-ID and Vendor-Specific attributes, then the value in the Vendor-Specific attribute (the per-user filter) takes precedence.

Also, if authentication for a port fails because the Filter-ID attribute referred to a non-existent filter, or there were insufficient system resources to implement the filter, then a Syslog message is generated.

When strict security mode is disabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the port is still authenticated, but no filter is dynamically applied to it.
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port is still authenticated, but the filter specified in the Vendor-Specific attribute is not applied to the port.

By default, strict security mode is enabled for all 802.1X-enabled interfaces, but you can manually disable or enable it, either globally or for specific interfaces.

To disable strict security mode globally, enter the following commands.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# no global-filter-strict-security
```

After you globally disable strict security mode, you can re-enable it by entering the following command.

```
PowerConnect(config-dot1x)# global-filter-strict-security
```

Syntax: [no] global-filter-strict-security

To disable strict security mode for a specific interface, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# dot1x disable-filter-strict-security
```

To re-enable strict security mode for an interface, enter the following command.

```
PowerConnect(config-if-e10000-1)# no dot1x disable-filter-strict-security
```

Syntax: [no] dot1x disable-filter-strict-security

The output of the **show dot1x** and **show dot1x config** commands has been enhanced to indicate whether strict security mode is enabled or disabled globally and on an interface. Refer to [“Displaying 802.1X multiple-host authentication information”](#) on page 959.

Dynamically applying existing ACLs or MAC address filters

When a port is authenticated using 802.1X security, an IP ACL or MAC address filter that exists in the running-config on the Dell device can be dynamically applied to the port. To do this, you configure the Filter-ID (type 11) attribute on the RADIUS server. The Filter-ID attribute specifies the name or number of the Dell IP ACL or MAC address filter.

The following is the syntax for configuring the Filter-ID attribute to refer to a Dell IP ACL or MAC address filter.

Table 7:

Value	Description
ip.<number>.in	Applies the specified numbered ACL to the 802.1X authenticated port in the inbound direction.
ip.<name>.in	Applies the specified named ACL to the 802.1X authenticated port in the inbound direction.
mac.<number>.in	Applies the specified numbered MAC address filter to the 802.1X authenticated port in the inbound direction.

The following table lists examples of values you can assign to the Filter-ID attribute on the RADIUS server to refer to IP ACLs and MAC address filters configured on a Dell device.

Table 8:

Possible values for the filter ID attribute on the RADIUS server	ACL or MAC address filter configured on the Dell device
ip.2.in	access-list 2 permit host 36.48.0.3 access-list 2 permit 36.0.0.0 0.255.255.255
ip.102.in	access-list 102 permit ip 36.0.0.0 0.255.255.255 any
ip.fdry_filter.in	ip access-list standard fdry_filter permit host 36.48.0.3

Table 8:

Possible values for the filter ID attribute on the RADIUS server	ACL or MAC address filter configured on the Dell device
mac.2.in	mac filter 2 permit 3333.3333.3333 ffff.ffff.ffff any etype eq 0800
mac.2.in	mac filter 2 permit 3333.3333.3333 ffff.ffff.ffff any etype eq 0800
mac.3.in	mac filter 3 permit 2222.2222.2222 ffff.ffff.ffff any etype eq 0800

Notes

- The <name> in the Filter ID attribute is case-sensitive.
- You can specify only numbered MAC address filters in the Filter ID attribute. Named MAC address filters are not supported.
- Dynamic ACL filters are supported only for the inbound direction. Dynamic outbound ACL filters are not supported.
- MAC address filters are supported only for the inbound direction. Outbound MAC address filters are not supported.
- Dynamically assigned IP ACLs and MAC address filters are subject to the same configuration restrictions as non-dynamically assigned IP ACLs and MAC address filters.

Configuring per-user IP ACLs or MAC address filters

Per-user IP ACLs and MAC address filters make use of the Vendor-Specific (type 26) attribute to dynamically apply filters to ports. Defined in the Vendor-Specific attribute are Dell ACL or MAC address filter statements. When the RADIUS server returns the Access-Accept message granting a client access to the network, the Dell device reads the statements in the Vendor-Specific attribute and applies these IP ACLs or MAC address filters to the client port. When the client disconnects from the network, the dynamically applied filters are no longer applied to the port. If any filters had been applied to the port previous to the client connecting, then those filters are reapplied to the port.

NOTE

Dynamic IP ACL filters and MAC address filters are not supported on the same port at the same time.

The following table shows the syntax for configuring the Dell Vendor-Specific attributes with ACL or MAC address filter statements.

Table 9:

Value	Description
ipACL.e.in=<extended-ACL-entries>	Applies the specified extended ACL entries to the 802.1X authenticated port in the inbound direction.
macfilter.in=<mac-filter-entries>	Applies the specified MAC address filter entries to the 802.1X authenticated port in the inbound direction.

The following table shows examples of IP ACLs and MAC address filters configured in the Dell Vendor-Specific attribute on a RADIUS server. These IP ACLs and MAC address filters follow the same syntax as other Dell ACLs and MAC address filters. Refer to the related chapters in this book for information on syntax.

Table 10:

ACL or MAC address filter	Vendor-specific attribute on RADIUS server
MAC address filter with one entry	macfilter.in= deny any any
MAC address filter with two entries	macfilter.in= permit 0000.0000.3333 ffff.ffff.0000 any, macfilter.in= permit 0000.0000.4444 ffff.ffff.0000 any

The RADIUS server allows one instance of the Vendor-Specific attribute to be sent in an Access-Accept message.

Enabling 802.1X port security

By default, 802.1X port security is disabled on Dell devices. To enable the feature on the device and enter the dot1x configuration level, enter the following command.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)#
```

Syntax: [no] dot1x-enable

At the dot1x configuration level, you can enable 802.1X port security on all interfaces at once, on individual interfaces, or on a range of interfaces.

For example, to enable 802.1X port security on all interfaces on the device, enter the following command.

```
PowerConnect(config-dot1x)# enable all
```

Syntax: [no] enable all

To enable 802.1X port security on interface 11, enter the following command.

```
PowerConnect(config-dot1x)# enable ethernet 11
```

Syntax: [no] enable ethernet <portnum>

The <portnum> parameter is a valid port number.

To enable 802.1X port security on interfaces 11 through 16, enter the following command.

```
PowerConnect(config-dot1x)# enable ethernet 11 to 16
```

Syntax: [no] enable ethernet<portnum> to <portnum>

The <portnum> parameter is a valid port number.

Setting the port control

To activate authentication on an 802.1X-enabled interface, you specify the kind of **port control** to be used on the interface. An interface used with 802.1X port security has two virtual access points: a controlled port and an uncontrolled port:

- The controlled port can be either the authorized or unauthorized state. In the authorized state, it allows normal traffic to pass between the Client and the Authenticator. In the unauthorized state, no traffic is allowed to pass.
- The uncontrolled port allows only EAPOL traffic between the Client and the Authentication Server.

Refer to [Figure 119](#) for an illustration of this concept.

By default, all controlled ports on the device are in the authorized state, allowing all traffic. When you activate authentication on an 802.1X-enabled interface, its controlled port is placed in the unauthorized state. When a Client connected to the interface is successfully authenticated, the controlled port is then placed in the authorized state. The controlled port remains in the authorized state until the Client logs off.

To activate authentication on an 802.1X-enabled interface, you configure the interface to place its controlled port in the authorized state when a Client is authenticated by an Authentication Server. To do this, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-1)# dot1x port-control auto
```

Syntax: [no] dot1x port-control [force-authorized | force-unauthorized | auto]

When an interface control type is set to **auto**, the controlled port is initially set to unauthorized, but is changed to authorized when the connecting Client is successfully authenticated by an Authentication Server.

The port control type can be one of the following

force-authorized – The controlled port is placed unconditionally in the authorized state, allowing all traffic. This is the default state for ports on the Dell device.

force-unauthorized – The controlled port is placed unconditionally in the unauthorized state.

auto – The controlled port is unauthorized until authentication takes place between the Client and Authentication Server. Once the Client passes authentication, the port becomes authorized. This activates authentication on an 802.1X-enabled interface.

NOTE

You cannot enable 802.1X port security on ports that have any of the following features enabled:

- Link aggregation
- Metro Ring Protocol (MRP)
- Mirror port
- Trunk port

Configuring periodic re-authentication

You can configure the device to periodically re-authenticate Clients connected to 802.1X-enabled interfaces. When you enable periodic re-authentication, the device re-authenticates Clients every 3,600 seconds by default. You can optionally specify a different re-authentication interval of between 1 – 4294967295 seconds.

To configure periodic re-authentication using the default interval of 3,600 seconds, enter the following command.

```
PowerConnect(config-dot1x)# re-authentication
```

Syntax: [no] re-authentication

To configure periodic re-authentication with an interval of 2,000 seconds, enter the following commands.

```
PowerConnect(config-dot1x)# re-authentication
PowerConnect(config-dot1x)# timeout re-authperiod 2000
```


Syntax: [no] **timeout re-authperiod** <seconds>

The re-authentication interval is a global setting, applicable to all 802.1X-enabled interfaces. To re-authenticate Clients connected to a specific port manually, use the **dot1x re-authenticate** command. Refer to [“Re-authenticating a port manually”](#), below.

Re-authenticating a port manually

When periodic re-authentication is enabled, by default the Dell device re-authenticates Clients connected to an 802.1X-enabled interface every 3,600 seconds (or the time specified by the **dot1x timeout re-authperiod** command). You can also manually re-authenticate Clients connected to a specific port.

For example, to re-authenticate Clients connected to interface 1, enter the following command.

```
PowerConnect# dot1x re-authenticate e 1
```

Syntax: **dot1x re-authenticate ethernet** <portnum>

The <portnum> parameter is a valid port number.

Setting the quiet period

If the Dell device is unable to authenticate the Client, the Dell device waits a specified amount of time before trying again. The amount of time the Dell device waits is specified with the **quiet-period** parameter. The **quiet-period** parameter can be from 1 – 4294967295 seconds. The default is 60 seconds.

For example, to set the quiet period to 30 seconds, enter the following command.

```
PowerConnect(config-dot1x)# timeout quiet-period 30
```

Syntax: [no] **timeout quiet-period** <seconds>

Specifying the wait interval and number of EAP-request/ identity frame retransmissions from the PowerConnect device

When the Dell device sends an EAP-request/identity frame to a Client, it expects to receive an EAP-response/identity frame from the Client. By default, if the Dell device does not receive an EAP-response/identity frame from a Client, the device waits 30 seconds, then retransmits the EAP-request/identity frame. Also by default, the Dell device retransmits the EAP-request/identity frame a maximum of two times. You can optionally configure the amount of time the device will wait before retransmitting an EAP-request/identity frame, and the number of times the EAP-request/identity frame will be transmitted. This section provides the command syntax for these features.

Setting the wait interval for EAP frame retransmissions

By default, if the Dell device does not receive an EAP-response/identity frame from a Client, the device waits 30 seconds, then retransmits the EAP-request/identity frame. You can optionally change the amount of time the Dell device waits before retransmitting the EAP-request/identity frame to the Client.

For example, to cause the Dell device to wait 60 seconds before retransmitting an EAP-request/identity frame to a Client, enter the following command.

```
PowerConnect(config-dot1x)# timeout tx-period 60
```

If the Client does not send back an EAP-response/identity frame within 60 seconds, the device will transmit another EAP-request/identity frame.

Syntax: [no] **timeout tx-period** <seconds>

where <seconds> is a value from 1 – 4294967295. The default is 30 seconds.

Setting the maximum number of EAP frame retransmissions

The Dell device retransmits the EAP-request/identity frame a maximum of two times. If no EAP-response/identity frame is received from the Client after two EAP-request/identity frame retransmissions (or the amount of time specified with the **auth-max** command), the device restarts the authentication process with the Client.

You can optionally change the number of times the Dell device should retransmit the EAP-request/identity frame. You can specify between 1 – 10 frame retransmissions. For example, to configure the device to retransmit an EAP-request/identity frame to a Client a maximum of three times, enter the following command:

```
PowerConnect(config-dot1x)# auth-max 3
```

Syntax: **auth-max** <value>

<value> is a number from 1 – 10. The default is 2.

Specifying the wait interval and number of EAP-request/identity frame retransmissions from the RADIUS server

Acting as an intermediary between the RADIUS Authentication Server and the Client, the Dell device receives RADIUS messages from the RADIUS server, encapsulates them as EAPOL frames, and sends them to the Client. By default, when the Dell device relays an EAP-Request frame from the RADIUS server to the Client, it expects to receive a response from the Client within 30 seconds. If the Client does not respond within the allotted time, the device retransmits the EAP-Request frame to the Client. Also by default, the Dell device retransmits the EAP-request frame twice. If no EAP-response frame is received from the Client after two EAP-request frame retransmissions, the device restarts the authentication process with the Client.

You can optionally configure the amount of time the device will wait before retransmitting an EAP-request/identity frame, and the number of times the EAP-request/identity frame will be transmitted. This section provides the command syntax for these features.

Setting the wait interval for EAP frame retransmissions

By default, when the Dell device relays an EAP-Request frame from the RADIUS server to the Client, it expects to receive a response from the Client within 30 seconds. You can optionally specify the wait interval using the **supptimeout** command.

For example, to configure the device to retransmit an EAP-Request frame if the Client does not respond within 45 seconds, enter the following command.

```
PowerConnect(config-dot1x)# supptimeout 45
```

Syntax: `supertimeout <seconds>`

`<seconds>` is a number from 1 – 4294967295 seconds. The default is 30 seconds.

Setting the maximum number of EAP frame retransmissions

You can optionally specify the number of times the Dell device will retransmit the EAP-request frame. You can specify between 1 – 10 frame retransmissions. For example, to configure the device to retransmit an EAP-request frame to a Client a maximum of three times, enter the following command.

```
PowerConnect(config-dot1x)# max-req 3
```

Syntax: `max-req <value>`

`<value>` is a number from 1 – 10. The default is 2.

Specifying a timeout for retransmission of messages to the authentication server

When performing authentication, the Dell device receives EAPOL frames from the Client and passes the messages on to the RADIUS server. The device expects a response from the RADIUS server within 30 seconds. If the RADIUS server does not send a response within 30 seconds, the Dell device retransmits the message to the RADIUS server. The time constraint for retransmission of messages to the Authentication Server can be between 0 – 4294967295 seconds.

For example, to configure the device to retransmit a message if the Authentication Server does not respond within 45 seconds, enter the following command.

```
PowerConnect(config-dot1x)# servertimeout 45
```

Syntax: `servertimeout <seconds>`

Initializing 802.1X on a port

To initialize 802.1X port security on a port, enter a command such as the following.

```
PowerConnect# dot1x initialize e 1
```

Syntax: `dot1x initialize ethernet<portnum>`

The `<portnum>` parameter is a valid port number.

Allowing access to multiple hosts

Dell devices support 802.1X authentication for ports with more than one host connected to them. If there are multiple hosts connected to a single 802.1X-enabled port, the Dell device authenticates each of them individually. Refer to [“Configuring 802.1X multiple-host authentication”](#) on page 949.

Configuring 802.1X multiple-host authentication

When multiple hosts are connected to the same 802.1X-enabled port, the functionality described in [“How 802.1X Multiple-host authentication works”](#) on page 931 is enabled by default. You can optionally do the following:

- Specify the authentication-failure action
- Specify the number of authentication attempts the device makes before dropping packets
- Disabling aging for dot1x-mac-sessions
- Configure aging time for blocked Clients
- Clear the dot1x-mac-session for a MAC address

Specifying the authentication-failure action

In an 802.1X multiple-host configuration, if RADIUS authentication for a Client is unsuccessful, traffic from that Client is either dropped in hardware (the default), or the Client port is placed in a “restricted” VLAN. You can specify which of these two authentication-failure actions is to be used. If the authentication-failure action is to place the port in a restricted VLAN, you can specify the ID of the restricted VLAN.

To specify that the authentication-failure action is to place the Client port in a restricted VLAN, enter the following command.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# auth-fail-action restricted-vlan
```

Syntax: [no] auth-fail-action restricted-vlan

To specify the ID of the restricted VLAN as VLAN 300, enter the following command.

```
PowerConnect(config-dot1x)# auth-fail-vlanid 300
```

Syntax: [no] auth-fail-vlanid <vlan-id>

Specifying the number of authentication attempts the device makes before dropping packets

When the authentication-failure action is to drop traffic from the Client, and the initial authentication attempt made by the device to authenticate the Client is unsuccessful, the Dell device immediately retries to authenticate the Client. After three unsuccessful authentication attempts, the Client dot1x-mac-session is set to “access-denied”, causing traffic from the Client to be dropped in hardware.

You can optionally configure the number of authentication attempts the device makes before dropping traffic from the Client. To do so, enter a command such as the following.

```
PowerConnect(config-dot1x)# auth-fail-max-attempts 2
```

Syntax: [no] auth-fail-max-attempts <attempts>

By default, the device makes 3 attempts to authenticate a Client before dropping packets from the Client. You can specify between 1 – 10 authentication attempts.

Disabling aging for dot1x-mac-sessions

The dot1x-mac-sessions for Clients authenticated or denied by a RADIUS server are aged out if no traffic is received from the Client MAC address for a certain period of time. After a Client dot1x-mac-session is aged out, the Client must be re-authenticated:

- **Permitted** dot1x-mac-sessions, which are the dot1x-mac-sessions for authenticated Clients, as well as for non-authenticated Clients whose ports have been placed in the restricted VLAN, are aged out if no traffic is received from the Client MAC address over the normal MAC aging interval on the Dell device.
- **Denied** dot1x-mac-sessions, which are the dot1x-mac-sessions for non-authenticated Clients that are blocked by the Dell device are aged out over a configurable software aging period. (Refer to the next section for more information on configuring the software aging period).

You can optionally disable aging of the permitted or denied dot1x-mac-sessions, or both, on the Dell device.

To disable aging of the permitted dot1x-mac-sessions, enter the following command.

```
PowerConnect(config-dot1x)# mac-session-aging no-aging permitted-mac-only
```

Syntax: [no] mac-session-aging no-aging permitted-mac-only

To disable aging of the denied dot1x-mac-sessions, enter the following command.

```
PowerConnect(config-dot1x)# mac-session-aging no-aging denied-mac-only
```

Syntax: [no] mac-session-aging no-aging denied-mac-only

NOTE

This command enables aging of permitted sessions.

As a shortcut, use the command **[no] mac-session-aging** to enable or disable aging for permitted and denied sessions.

Specifying the aging time for blocked clients

When the Dell device is configured to drop traffic from non-authenticated Clients, traffic from the blocked Clients is dropped in hardware, without being sent to the CPU. A Layer 2 CAM entry is created that drops traffic from the blocked Client MAC address in hardware. If no traffic is received from the blocked Client MAC address for a certain amount of time, this Layer 2 CAM entry is aged out. If traffic is subsequently received from the Client MAC address, then an attempt can be made to authenticate the Client again.

Aging of the Layer 2 CAM entry for a blocked Client MAC address occurs in two phases, known as **hardware aging** and **software aging**. The hardware aging period is fixed at 70 seconds and is non-configurable. The software aging time is configurable through the CLI.

Once the Dell device stops receiving traffic from a blocked Client MAC address, the hardware aging begins and lasts for a fixed period of time. After the hardware aging period ends, the software aging period begins. The software aging period lasts for a configurable amount of time (by default 120 seconds). After the software aging period ends, the blocked Client MAC address ages out, and can be authenticated again if the Dell device receives traffic from the Client MAC address.

Change the length of the software aging period for a blocked Client MAC address by entering a command such as the following.

```
PowerConnect(config)# mac-session-aging max-age 180
```

Syntax: [no] mac-session-aging max-age <seconds>

You can specify from 1 – 65535 seconds. The default is 120 seconds.

Clearing a dot1x-mac-session for a MAC address

You can clear the dot1x-mac-session for a specified MAC address, so that the Client with that MAC address can be re-authenticated by the RADIUS server.

Example

```
PowerConnect# clear dot1x mac-session 00e0.1234.abd4
```

Syntax: clear dot1x mac-session <mac-address>

Configuring VLAN access for non-EAP-capable clients

You can configure the Dell device to grant "guest" or restricted VLAN access to clients that do not support Extensible EAP. The restricted VLAN limits access to the network or applications, instead of blocking access to these services altogether.

When the Dell device receives the first packet (non-EAP packet) from a client, the device waits for 10 seconds or the amount of time specified with the **timeout restrict-fwd-period** command. If the Dell device does not receive subsequent packets after the timeout period, the device places the client on the restricted VLAN.

This feature is disabled by default. To enable this feature and change the timeout period, enter commands such as the following.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# restrict-forward-non-dot1x
PowerConnect(config-dot1x)# timeout restrict-fwd-period 15
```

Once the *success* timeout action is enabled, use the **no** form of the command to reset the RADIUS timeout behavior to *retry*.

Syntax: `timeout restrict-fwd-period <num>`

The `<num>` parameter is a value from 0 to 32767. The default value is 10.

Configuring a timeout action to cancel 802.1X authentication for Non-802.1x clients

Normally, the Dell-specific attribute obtained from the RADIUS server identifies a client as not 802.1X-capable and tells the switch not to perform 802.1X authentication for this client.

However, if you configure an **auth-timeout-action** at the global level, the Dell-specific attribute from the RADIUS server is no longer required to cancel 802.1X authentication for a non-802.1X user. To configure the timeout action, enter commands similar to the following at the global level.

```
PowerConnect(config)# dot1x-enable
PowerConnect(config-dot1x)# restrict-forward-non-dot1x auth-timeout-action
```

Syntax: `restrict-forward-non-dot1x [auth-timeout-action]`

To set the RADIUS timeout behavior to bypass dot.1X authentication and **permit** client access to the network, enter commands similar to the following (at the interface level).

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x auth-timeout-action success
```

To set the RADIUS timeout behavior to bypass 802.1X authentication and return a **failure**, which limits access to the network and moves the client to the restricted VLAN, enter commands similar to the following (at the interface level).

```
PowerConnect(config)# interface ethernet 1
PowerConnect(config-if-e100-1)# dot1x auth-timeout-action failure
```

Syntax: `[no] dot1x auth-timeout-action success`

Syntax: `[no] dot1x auth-timeout-action failure`

NOTE

The success or failure of multi-device port authentication can change the effect of these commands.

Displaying 802.1X information

You can display the following 802.1X-related information:

- The 802.1X configuration on the device and on individual ports
- Statistics about the EAPOL frames passing through the device
- 802.1X-enabled ports dynamically assigned to a VLAN
- User-defined and dynamically applied MAC filters and IP ACLs currently active on the device
- The 802.1X multiple-host configuration

Displaying 802.1X configuration information

To display information about the 802.1X configuration on the Dell device, enter the following command.

```
PowerConnect# show dot1x
PAE Capability:    Authenticator Only
system-auth-control: Enable
re-authentication: Disable
global-filter-strict-security: Enable
quiet-period:    60 Seconds
tx-period:       30 Seconds
supptimeout:    30 Seconds
servertimeout:   30 Seconds
maxreq:         2
re-authperiod:   3600 Seconds
Protocol Version: 1
```

Syntax: show dot1x

The following table describes the information displayed by the **show dot1x** command.

TABLE 146 Output from the **show dot1x** command

This field...	Displays...
PAE Capability	The Port Access Entity (PAE) role for the Dell device. This is always "Authenticator Only".
system-auth-control	Whether system authentication control is enabled on the device. The dot1x-enable command enables system authentication control on the device.
re-authentication	Whether periodic re-authentication is enabled on the device. Refer to "Configuring periodic re-authentication" on page 946. When periodic re-authentication is enabled, the device automatically re-authenticates Clients every 3,600 seconds by default.
global-filter-strict-security	Whether strict security mode is enabled or disabled globally. Refer to "Disabling and enabling strict security mode for dynamic filter assignment" on page 942.
quiet-period	When the Dell device is unable to authenticate a Client, the amount of time the Dell device waits before trying again (default 60 seconds). Refer to "Setting the quiet period" on page 947 for information on how to change this setting.

TABLE 146 Output from the **show dot1x** command (Continued)

This field...	Displays...
tx-period	When a Client does not send back an EAP-response/identity frame, the amount of time the Dell device waits before retransmitting the EAP-request/identity frame to a Client (default 30 seconds). Refer to “Setting the wait interval for EAP frame retransmissions” on page 947 for information on how to change this setting.
supp-timeout	When a Client does not respond to an EAP-request frame, the amount of time before the Dell device retransmits the frame. Refer to “Setting the wait interval for EAP frame retransmissions” on page 948 for information on how to change this setting.
server-timeout	When the Authentication Server does not respond to a message sent from the Client, the amount of time before the Dell device retransmits the message. Refer to “Specifying a timeout for retransmission of messages to the authentication server” on page 949 for information on how to change this setting.
max-req	The number of times the Dell device retransmits an EAP-request/identity frame if it does not receive an EAP-response/identity frame from a Client (default 2 times). Refer to “Setting the maximum number of EAP frame retransmissions” on page 948 for information on how to change this setting.
re-authperiod	How often the device automatically re-authenticates Clients when periodic re-authentication is enabled (default 3,600 seconds). Refer to “Configuring periodic re-authentication” on page 946 for information on how to change this setting.
Protocol Version	The version of the 802.1X protocol in use on the device.

To display information about the 802.1X configuration on an individual port, enter a command such as the following.

```
PowerConnect# show dot1x configuration ethernet 3
Port-Control                : control-auto
filter strict security      : Enable
Action on RADIUS timeout   : Treat as a failed authentication
  re-authenticate          : 150 seconds
PVID State                  : Normal (101)
Original PVID               : 101
PVID mac total              : 1
PVID mac authorized         : 1
num mac sessions            : 1
num mac authorized          : 1
Number of Auth filter       : 0
```

Syntax: **show dot1x config ethernet** <portnum>

The <portnum> parameter is a valid port number.

The following additional information is displayed in the **show dot1x config** command for an interface.

TABLE 147 Output from the **show dot1x config** command for an interface

This field...	Displays...
Authenticator PAE state	The current status of the Authenticator PAE state machine. This can be INITIALIZE, DISCONNECTED, CONNECTING, AUTHENTICATING, AUTHENTICATED, ABORTING, HELD, FORCE_AUTH, or FORCE_UNAUTH. NOTE: When the Authenticator PAE state machine is in the AUTHENTICATING state, if the reAuthenticate, eapStart, eapLogoff, or authTimeout parameters are set to TRUE, it may place the Authenticator PAE state machine indefinitely in the ABORTING state. If this should happen, use the dot1x initialize command to initialize 802.1X port security on the port, or unplug the Client or hub connected to the port, then reconnect it.
Backend Authentication state	The current status of the Backend Authentication state machine. This can be REQUEST, RESPONSE, SUCCESS, FAIL, TIMEOUT, IDLE, or INITIALIZE.
AdminControlledDirections	Indicates whether an unauthorized controlled port exerts control over communication in both directions (disabling both reception of incoming frames and transmission of outgoing frames), or just in the incoming direction (disabling only reception of incoming frames). On Dell devices, this parameter is set to BOTH.
OperControlledDirections	The setting for the OperControlledDirections parameter, as defined in the 802.1X standard. According to the 802.1X standard, if the AdminControlledDirections parameter is set to BOTH, the OperControlledDirections parameter is unconditionally set to BOTH. Since the AdminControlledDirections parameter on Dell devices is always set to BOTH, the OperControlledDirections parameter is also set to BOTH.
AuthControlledPortControl	The port control type configured for the interface. If set to auto, authentication is activated on the 802.1X-enabled interface.
AuthControlledPortStatus	The current status of the interface controlled port either authorized or unauthorized.
multiple-hosts	Whether the port is configured to allow multiple Supplicants accessing the interface on the Dell device through a hub. Refer to “Allowing access to multiple hosts” on page 949 for information on how to change this setting.

Displaying 802.1X statistics

To display 802.1X statistics for an individual port, enter a command such as the following

```
PowerConnect# show dot1x statistics e 3
```

```
Port 3 Statistics:
RX EAPOL Start:      0
RX EAPOL Logoff:    0
RX EAPOL Invalid:   0
RX EAPOL Total:     0
RX EAP Resp/Id:     0
RX EAP Resp other than Resp/Id: 0
RX EAP Length Error: 0
Last EAPOL Version: 0
Last EAPOL Source:  0007.9550.0B83
TX EAPOL Total:     217
TX EAP Req/Id:      163
TX EAP Req other than Req/Id: 0
```

Syntax: `show dot1x statistics ethernet <portnum>`

The *<portnum>* parameter is a valid port number.

The following table describes the information displayed by the **show dot1x statistics** command for an interface.

TABLE 148 Output from the **show dot1x statistics** command

This field...	Displays...
RX EAPOL Start	The number of EAPOL-Start frames received on the port.
RX EAPOL Logoff	The number of EAPOL-Logoff frames received on the port.
RX EAPOL Invalid	The number of invalid EAPOL frames received on the port.
RX EAPOL Total	The total number of EAPOL frames received on the port.
RX EAP Resp/Id	The number of EAP-Response/Identity frames received on the port
RX EAP Resp other than Resp/Id	The total number of EAPOL-Response frames received on the port that were not EAP-Response/Identity frames.
RX EAP Length Error	The number of EAPOL frames received on the port that have an invalid packet body length.
Last EAPOL Version	The version number of the last EAPOL frame received on the port.
Last EAPOL Source	The source MAC address in the last EAPOL frame received on the port.
TX EAPOL Total	The total number of EAPOL frames transmitted on the port.
TX EAP Req/Id	The number of EAP-Request/Identity frames transmitted on the port.
TX EAP Req other than Req/Id	The number of EAP-Request frames transmitted on the port that were not EAP-Request/Identity frames.

Clearing 802.1X statistics

You can clear the 802.1X statistics counters on all interfaces at once, on individual interfaces, or on a range of interfaces.

For example, to clear the 802.1X statistics counters on all interfaces on the device, enter the following command.

```
PowerConnect# clear dot1x statistics all
```

Syntax: clear dot1x statistics all

To clear the 802.1X statistics counters on interface e 11, enter the following command.

```
PowerConnect# clear dot1x statistics e 11
```

Syntax: clear dot1x statistics ethernet <portnum>

The <portnum> parameter is a valid port number.

Displaying dynamically assigned VLAN information

The **show interface** command displays the VLAN to which an 802.1X-enabled port has been dynamically assigned, as well as the port from which it was moved (that is, the port default VLAN).

The following example of the **show interface** command indicates the port dynamically assigned VLAN. Information about the dynamically assigned VLAN is shown in bold type.

```
PowerConnect# show interface e 2
FastEthernet2 is up, line protocol is up
  Hardware is FastEthernet, address is 0204.80a0.4681 (bia 0204.80a0.4681)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Member of L2 VLAN ID 2 (dot1x-RADIUS assigned), original L2 VLAN ID is 1,
  port is untagged, port state is FORWARDING
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  MTU 1518 bytes
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 256 bits/sec, 0 packets/sec, 0.00% utilization
  3 packets input, 192 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 3 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants, DMA received 3 packets
  919 packets output, 58816 bytes, 0 underruns
  Transmitted 1 broadcasts, 916 multicasts, 2 unicasts
  0 output errors, 0 collisions, DMA transmitted 919 packets
```

In this example, the 802.1X-enabled port has been moved from VLAN 1 to VLAN 2. When the client disconnects, the port will be moved back to VLAN 1.

The **show run** command also indicates the VLAN to which the port has been dynamically assigned. The output can differ depending on whether GARP VLAN Registration Protocol (GVRP) is enabled on the device:

- **Without GVRP** – When you enter the **show run** command, the output indicates that the port is a member of the VLAN to which it was dynamically assigned through 802.1X. If you then enter the **write memory** command, the VLAN to which the port is currently assigned becomes the port default VLAN in the device configuration.
- **With GVRP** – When you enter the **show run** command, if the VLAN name supplied by the RADIUS server corresponds to a VLAN learned through GVRP, then the output indicates that the port is a member of the VLAN to which it was originally assigned (not the VLAN to which it was dynamically assigned).

If the VLAN name supplied by the RADIUS server corresponds to a statically configured VLAN, the output indicates that the port is a member of the VLAN to which it was dynamically assigned through 802.1X. If you then enter the **write memory** command, the VLAN to which the port is currently assigned becomes the port default VLAN in the device configuration.

Displaying information about dynamically applied MAC filters and IP ACLs

You can display information about currently active user-defined and dynamically applied MAC filters and IP ACLs.

Displaying user-defined MAC filters and IP ACLs

To display the user-defined MAC filters active on the device, enter the following command.

```
PowerConnect# show dot1x mac-address filter
```

```
Port 3 (User defined MAC Address Filter) :
    mac filter 1 permit any any
```

Syntax: show dot1x mac-address-filter

To display the user-defined IP ACLs active on the device, enter the following command.

```
PowerConnect# show dot1x ip-ACL
```

```
Port 3 (User defined IP ACLs):
Extended IP access list Port_3_E_IN
permit udp any any
```

```
Extended IP access list Port_3_E_OUT
```

Syntax: show dot1x ip-ACL

Displaying dynamically applied MAC filters and IP ACLs

To display the dynamically applied MAC address filters active on an interface, enter a command such as the following.

```
PowerConnect# show dot1x mac-address-filter e 3
```

```
Port 3 MAC Address Filter information:
 802.1X Dynamic MAC Address Filter :
    mac filter-group 2
Port default MAC Address Filter:
    No mac address filter is set
```

Syntax: show dot1x mac-address-filter all | ethernet <portnum>

The **all** keyword displays all dynamically applied MAC address filters active on the device.

The <portnum> parameter is a valid port number.

To display the dynamically applied IP ACLs active on an interface, enter a command such as the following.

```
PowerConnect# show dot1x ip-ACL e 3

Port 3 IP ACL information:
 802.1X dynamic IP ACL (user defined) in:
   ip access-list extended Port_3_E_IN in
Port default IP ACL in:
  No inbound ip access-list is set
 802.1X dynamic IP ACL (user defined) out:
   ip access-list extended Port_3_E_OUT out
Port default IP ACL out:
  No outbound ip access-list is set
```

Syntax: `show dot1x ip-ACL all | ethernet<portnum>`

The **all** keyword displays all dynamically applied IP ACLs active on the device.

The `<portnum>` parameter is a valid port number.

Displaying 802.1X multiple-host authentication information

You can display the following information about 802.1X multiple-host authentication:

- Information about the 802.1X multiple-host configuration
- The dot1x-mac-sessions on each port
- The number of users connected on each port in a 802.1X multiple-host configuration

Displaying 802.1X multiple-host configuration information

The output of the **show dot1x** and **show dot1x config** commands displays information related to 802.1X multiple-host authentication.

The following is an example of the output of the **show dot1x** command. The information related to multiple-host authentication is highlighted in bold.

```
PowerConnect# show dot1x

Number of Ports enabled           : 2
Re-Authentication                 : Enabled
Authentication-fail-action      : Restricted VLAN
Authentication Failure VLAN   : 111
Mac Session Aging             : Disabled for permitted MAC sessions
Mac Session max-age           : 60 seconds
Protocol Version                  : 1
quiet-period                      : 5 Seconds
tx-period                         : 30 Seconds
supptimeout                      : 30 Seconds
servertimeout                    : 30 Seconds
maxreq                            : 2
re-authperiod                    : 3600 Seconds
security-hold-time                : 60 Seconds
re-authentication                 : Enable
Flow based multi-user policy    : Disable
```

Syntax: `show dot1x`

Table 149 describes the bold fields in the display.

TABLE 149 Output from the **show dot1x** command for multiple host authentication

This field...	Displays...
Authentication-fail-action	The configured authentication-failure action. This can be Restricted VLAN or Block Traffic.
Authentication Failure VLAN	If the authentication-failure action is Restricted VLAN, the ID of the VLAN to which unsuccessfully authenticated Client ports are assigned.
Mac Session Aging	Whether aging for dot1x-mac-sessions has been enabled or disabled for permitted or denied dot1x-mac-sessions.
Mac Session max-age	The configured software aging time for dot1x-mac-sessions.
Flow based multi-user policy	The dynamically assigned IP ACLs and MAC address filters used in the 802.1X multiple-host configuration.

The output of the **show dot1x config** command for an interface displays the configured port control for the interface. This command also displays information related to 802.1X multiple host-authentication.

The following is an example of the output of the **show dot1x config** command for an interface.

```
PowerConnect# show dot1x config e 1

Port-Control                : control-auto
filter strict security      : Enable
PVID State                  : Restricted (10)
Original PVID               : 10
PVID mac total              : 1
PVID mac authorized         : 0
num mac sessions           : 1
num mac authorized         : 0
```

Syntax: **show dot1x config ethernet** <portnum>

The <portnum> parameter is a valid port number.

The following table lists the fields in the display.

TABLE 150 Output from the **show dot1x config** command

This field...	Displays...
Port-Control	The configured port control type for the interface. This can be one of the following: force-authorized – The controlled port is placed unconditionally in the authorized state, allowing all traffic. This is the default state for ports on the Dell device. force-unauthorized – The controlled port is placed unconditionally in the unauthorized state. No authentication takes place for any connected 802.1X Clients. auto – The authentication status for each 802.1X Client depends on the authentication status returned from the RADIUS server.
filter strict security	Whether strict security mode is enabled or disabled on the interface.
PVID State	The port default VLAN ID (PVID) and the state of the port PVID. The PVID state can be one of the following Normal – The port PVID is not set by a RADIUS server, nor is it the restricted VLAN. RADIUS – The port PVID was dynamically assigned by a RADIUS server. RESTRICTED – The port PVID is the restricted VLAN.

TABLE 150 Output from the **show dot1x config** command (Continued)

This field...	Displays...
Original PVID	The originally configured (not dynamically assigned) PVID for the port.
PVID mac total	The number of devices transmitting untagged traffic on the port PVID.
PVID mac authorized	The number of devices transmitting untagged traffic on the port PVID as a result of dynamic VLAN assignment.
num mac sessions	The number of dot1x-mac-sessions on the port.
num mac authorized	The number of authorized dot1x-mac-sessions on the port.

Displaying information about the dot1x MAC sessions on each port

The **show dot1x mac-session** command displays information about the dot1x-mac-sessions on each port on the device. The output also shows the authenticator PAE state.

Example

```
PowerConnect# show dot1x mac-session
```

Port	MAC/(username)	Vlan	Auth State	ACL	Age	PAE State
1	0010.a498.24f7 :User	10	permit	none	S20	AUTHENTICATED

Syntax: show dot1x mac-session

[Table 151](#) lists the new fields in the display.

TABLE 151 Output from the **show dot1x mac-session** command

This field...	Displays...
Port	The port on which the dot1x-mac-session exists.
MAC/ (username)	The MAC address of the Client and the username used for RADIUS authentication.
Vlan	The VLAN to which the port is currently assigned.
Auth-State	The authentication state of the dot1x-mac-session. This can be one of the following permit - The Client has been successfully authenticated, and traffic from the Client is being forwarded normally. blocked - Authentication failed for the Client, and traffic from the Client is being dropped in hardware. restricted - Authentication failed for the Client, but traffic from the Client is allowed in the restricted VLAN only. init - The Client is in is in the process of 802.1X authentication, or has not started the authentication process.

TABLE 151 Output from the **show dot1x mac-session** command (Continued)

This field...	Displays...
Age	The software age of the dot1x-mac-session.
PAE State	The current status of the Authenticator PAE state machine. This can be INITIALIZE, DISCONNECTED, CONNECTING, AUTHENTICATING, AUTHENTICATED, ABORTING, HELD, FORCE_AUTH, or FORCE_UNAUTH. NOTE: When the Authenticator PAE state machine is in the AUTHENTICATING state, if the reAuthenticate, eapStart, eapLogoff, or authTimeout parameters are set to TRUE, it may place the Authenticator PAE state machine indefinitely in the ABORTING state. If this should happen, use the dot1x initialize command to initialize 802.1X port security on the port, or unplug the Client or hub connected to the port, then reconnect it.

Displaying information about the ports in an 802.1X multiple-host configuration

To display information about the ports in an 802.1X multiple-host configuration, enter the following command.

```
PowerConnect(config-dot1x)# sh do mac-s br
Port          Number of  Number of  Dynamic  Dynamic  Dynamic
              users     Authorized users  VLAN    ACL      MAC-Filt
-----
1             0         0          no      no       no
2             0         0          no      no       no
3             0         0          no      no       no
4             0         0          no      no       no
5             0         0          no      no       no
6             0         0          no      no       no
7             0         0          no      no       no
8             0         0          no      no       no
9             0         0          no      no       no
10            0         0          no      no       no
11            0         0          no      no       no
12            0         0          no      no       no
13            0         0          no      no       no
14            0         0          no      no       no
15            0         0          no      no       no
16            0         0          no      no       no
```

Syntax: show dot1x mac-session brief

The following table describes the information displayed by the **show dot1x mac-session brief** command.

TABLE 152 Output from the **show dot1x mac-session brief** command

This field...	Displays...
Port	Information about the users connected to each port.
Number of users	The number of users connected to the port.
Number of Authorized users	The number of users connected to the port that have been successfully authenticated.

TABLE 152 Output from the `show dot1x mac-session brief` command (Continued)

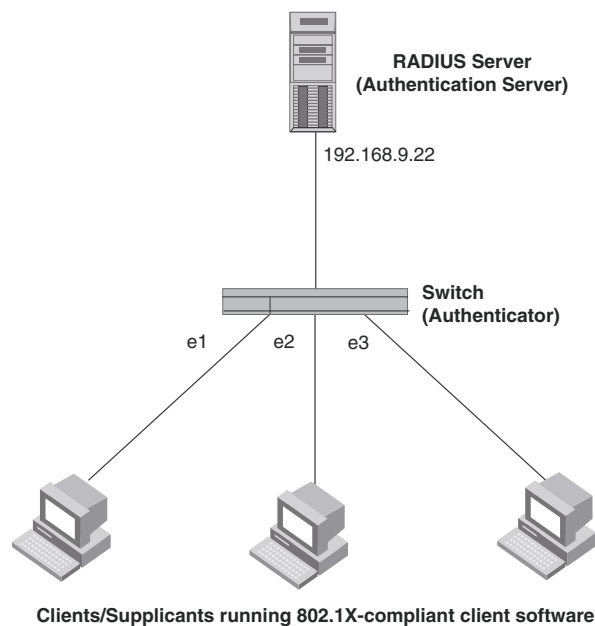
This field...	Displays...
Dynamic VLAN	Whether the port is a member of a RADIUS-specified VLAN.
Dynamic Filters	Whether RADIUS-specified IP ACLs or MAC address filters have been applied to the port.

Sample 802.1X configurations

This section illustrates a sample point-to-point configuration and a sample hub configuration that use 802.1X port security.

Point-to-point configuration

Figure 123 illustrates a sample 802.1X configuration with Clients connected to three ports on the Dell device. In a point-to-point configuration, only one 802.1X Client can be connected to each port.

FIGURE 123 Sample point-to-point 802.1X configuration

The following commands configure the Dell device in Figure 123

```
PowerConnect(config)# aaa authentication dot1x default radius
PowerConnect(config)# radius-server host 192.168.9.22 auth-port 1812 acct-port
1813 default key mirabeau dot1x
PowerConnect(config)# dot1x-enable e 1 to 3
PowerConnect(config-dot1x)# re-authentication
PowerConnect(config-dot1x)# timeout re-authperiod 2000
PowerConnect(config-dot1x)# timeout quiet-period 30
PowerConnect(config-dot1x)# timeout tx-period 60
PowerConnect(config-dot1x)# max-req 6
PowerConnect(config-dot1x)# exit
```

```


PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# dot1x port-control auto
PowerConnect(config-if-e10000-1)# exit
PowerConnect(config)# interface e 2
PowerConnect(config-if-e10000-2)# dot1x port-control auto
PowerConnect(config-if-e10000-2)# exit
PowerConnect(config)# interface e 3
PowerConnect(config-if-e10000-3)# dot1x port-control auto
PowerConnect(config-if-e10000-3)# exit

```

Hub configuration

Figure 124 illustrates a configuration where three 802.1X-enabled Clients are connected to a hub, which is connected to a port on the Dell device. The configuration is similar to that in Figure 123, except that 802.1X port security is enabled on only one port, and the **multiple-hosts** command is used to allow multiple Clients on the port.

FIGURE 124 Sample 802.1X configuration using a hub



The following commands configure the Dell device in Figure 124

```

PowerConnect(config)# aaa authentication dot1x default radius
PowerConnect(config)# radius-server host 192.168.9.22 auth-port 1812 acct-port
1813 default key mirabeau dot1x
PowerConnect(config)# dot1x-enable e 1
PowerConnect(config-dot1x)# re-authentication
PowerConnect(config-dot1x)# timeout re-authperiod 2000
PowerConnect(config-dot1x)# timeout quiet-period 30
PowerConnect(config-dot1x)# timeout tx-period 60
PowerConnect(config-dot1x)# max-req 6
PowerConnect(config-dot1x)# exit


```

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# dot1x port-control auto
PowerConnect(config-if-e10000-1)# dot1x multiple-hosts
PowerConnect(config-if-e10000-1)# exit
```

802.1X Authentication with dynamic VLAN assignment

Figure 125 illustrates 802.1X authentication with dynamic VLAN assignment. In this configuration, two user PCs are connected to a hub, which is connected to port e2. Port e2 is configured as a dual-mode port. Both PCs transmit untagged traffic. The profile for User 1 on the RADIUS server specifies that User 1 PC should be dynamically assigned to VLAN 3. The RADIUS profile for User 2 on the RADIUS server specifies that User 2 PC should be dynamically assigned to VLAN 20.

FIGURE 125 Sample configuration using 802.1X authentication with dynamic VLAN assignment



In this example, the PVID for port e2 would be changed based on the first host to be successfully authenticated. If User 1 is authenticated first, then the PVID for port e2 is changed to VLAN 3. If User 2 is authenticated first, then the PVID for port e2 is changed to VLAN 20. Since a PVID cannot be changed by RADIUS authentication after it has been dynamically assigned, if User 2 is authenticated after the port PVID was changed to VLAN 3, then User 2 would not be able to gain access to the network.

If there were only one device connected to the port, and authentication failed for that device, it could be placed into the restricted VLAN, where it could gain access to the network.

The part of the running-config related to 802.1X authentication would be as follows.

```
dot1x-enable
 re-authentication
 servertimeout 10
 timeout re-authperiod 10
 auth-fail-action restricted-vlan
 auth-fail-vlanid 1023
 mac-session-aging no-aging permitted-mac-only
 enable ethe 2 to 4
!
!
!
interface ethernet 2
 dot1x port-control auto
 dual-mode
```

If User 1 is successfully authenticated before User 2, the PVID for port e2 would be changed from the default VLAN to VLAN 3.

Had User 2 been the first to be successfully authenticated, the PVID would be changed to 20, and User 1 would not be able to gain access to the network. If there were only one device connected to the port that was sending untagged traffic, and 802.1X authentication failed for that device, it would be placed in the restricted VLAN 1023, and would be able to gain access to the network.

Using multi-device port authentication and 802.1X security on the same port

You can configure the Dell device to use multi-device port authentication and 802.1X security on the same port:

- The multi-device port authentication feature allows you to configure a Dell device to forward or block traffic from a MAC address based on information received from a RADIUS server. Incoming traffic originating from a given MAC address is switched or forwarded by the device only if the source MAC address is successfully authenticated by a RADIUS server. The MAC address itself is used as the username and password for RADIUS authentication. A connecting user does not need to provide a specific username and password to gain access to the network.
- The IEEE 802.1X standard is a means for authenticating devices attached to LAN ports. Using 802.1X port security, you can configure a Dell device to grant access to a port based on information supplied by a client to an authentication server.

For information on configuring the multi-device port authentication feature and 802.1X security on Dell devices, Refer to the related chapters in this book.

When both of these features are enabled on the same port, multi-device port authentication is performed prior to 802.1X authentication. If multi-device port authentication is successful, 802.1X authentication may be performed, based on the configuration of a vendor-specific attribute (VSA) in the profile for the MAC address on the RADIUS server.

When both features are configured on a port, a device connected to the port is authenticated as follows.

1. Multi-device port authentication is performed on the device to authenticate the device MAC address.
2. If multi-device port authentication is successful for the device, then the Dell device checks whether the RADIUS server included the Foundry-802_1x-enable VSA (described in [Table 153](#)) in the Access-Accept message that authenticated the device.
3. If the Foundry-802_1x-enable VSA is not present in the Access-Accept message, or is present and set to 1, then 802.1X authentication is performed for the device.
4. If the Foundry-802_1x-enable VSA is present in the Access-Accept message, and is set to 0, then 802.1X authentication is skipped. The device is authenticated, and any dynamic VLANs specified in the Access-Accept message returned during multi-device port authentication are applied to the port.
5. If 802.1X authentication is performed on the device, and is successful, then dynamic VLANs or ACLs specified in the Access-Accept message returned during 802.1X authentication are applied to the port.

If multi-device port authentication fails for a device, then by default traffic from the device is either blocked in hardware, or the device is placed in a restricted VLAN. You can optionally configure the Dell device to perform 802.1X authentication on a device when it fails multi-device port authentication. Refer to [“Example 2”](#) on page 970 for a sample configuration where this is used.

Configuring Dell-specific attributes on the RADIUS server

If the RADIUS authentication process is successful, the RADIUS server sends an Access-Accept message to the Dell device, authenticating the device. The Access-Accept message can include Vendor-Specific Attributes (VSAs) that specify additional information about the device. If you are configuring multi-device port authentication and 802.1X authentication on the same port, then you can configure the Dell VSAs listed in [Table 153](#) on the RADIUS server.

Add these Dell vendor-specific attributes to your RADIUS server configuration, and configure the attributes in the individual or group profiles of the devices that will be authenticated. The Dell Vendor-ID is 1991, with Vendor-Type 1.

TABLE 153 Dell vendor-specific attributes for RADIUS

Attribute Name	Attribute ID	Data Type	Description
Foundry-802_1x-enable	6	integer	Specifies whether 802.1X authentication is performed when multi-device port authentication is successful for a device. This attribute can be set to one of the following: 0 Do not perform 802.1X authentication on a device that passes multi-device port authentication. Set the attribute to zero for devices that do not support 802.1X authentication. 1 Perform 802.1X authentication when a device passes multi-device port authentication. Set the attribute to one for devices that support 802.1X authentication.
Foundry-802_1x-valid	7	integer	Specifies whether the RADIUS record is valid only for multi-device port authentication, or for both multi-device port authentication and 802.1X authentication. This attribute can be set to one of the following: 0 The RADIUS record is valid only for multi-device port authentication. Set this attribute to zero to prevent a user from using their MAC address as username and password for 802.1X authentication 1 The RADIUS record is valid for both multi-device port authentication and 802.1X authentication.

If neither of these VSAs exist in a device profile on the RADIUS server, then by default the device is subject to multi-device port authentication (if configured), then 802.1X authentication (if configured). The RADIUS record can be used for both multi-device port authentication and 802.1X authentication.

Example configurations

The following examples show configurations that use multi-device port authentication and 802.1X authentication on the same port.


Example 1

[Figure 126](#) illustrates an example configuration that uses multi-device port authentication and 802.1X authentication on the same port. In this configuration, a PC and an IP phone are connected to port e 3 on a Dell device. Port e 3 is configured as a dual-mode port.

The profile for the PC MAC address on the RADIUS server specifies that the PC should be dynamically assigned to VLAN "Login-VLAN", and the RADIUS profile for the IP phone specifies that it should be dynamically assigned to the VLAN named "IP-Phone-VLAN". When User 1 is successfully authenticated using 802.1X authentication, the PC is then placed in the VLAN named "User-VLAN".

NOTE

This example assumes that the IP phone initially transmits untagged packets (for example, CDP or DHCP packets), which trigger the authentication process on the Dell device and client lookup on the RADIUS server. If the phone sends only tagged packets and the port (e 3) is not a member of that VLAN, authentication would not occur. In this case, port e 3 must be added to that VLAN prior to authentication.

FIGURE 126 Multi-device port authentication and 802.1X authentication on the same port

When the devices attempt to connect to the network, they are first subject to multi-device port authentication.

When the MAC address of the IP phone is authenticated, the Access-Accept message from the RADIUS server specifies that the IP phone port be placed into the VLAN named "IP-Phone-VLAN", which is VLAN 7. The Foundry-802_1x-enable attribute is set to 0, meaning that 802.1X authentication is skipped for this MAC address. Port 3 is placed in VLAN 7 as a tagged port. No further authentication is performed.

When the PC MAC address is authenticated, the Access-Accept message from the RADIUS server specifies that the PVID for the PC port be changed to the VLAN named "Login-VLAN", which is VLAN 1024. The Foundry-802_1x-enable attribute is set to 1, meaning that 802.1X authentication is required for this MAC address. The PVID of the port 3 is temporarily changed to VLAN 1024, pending 802.1X authentication.

When User 1 attempts to connect to the network from the PC, he is subject to 802.1X authentication. If User 1 is successfully authenticated, the Access-Accept message from the RADIUS server specifies that the PVID for User 1 port be changed to the VLAN named "User-VLAN", which is VLAN 3. If 802.1X authentication for User 1 is unsuccessful, the PVID for port 3 is changed to that of the restricted VLAN, which is 1023, or untagged traffic from port e 3 can be blocked in hardware.

The part of the running-config related to port e 3 would be as follows.

```
interface ethernet 3
  dot1x port-control auto
  mac-authentication enable
  dual-mode
```

When the PC is authenticated using multi-device port authentication, the port PVID is changed to "Login-VLAN", which is VLAN 1024 in this example.

When User 1 is authenticated using 802.1X authentication, the port PVID is changed to "User-VLAN", which is VLAN 3 in this example.

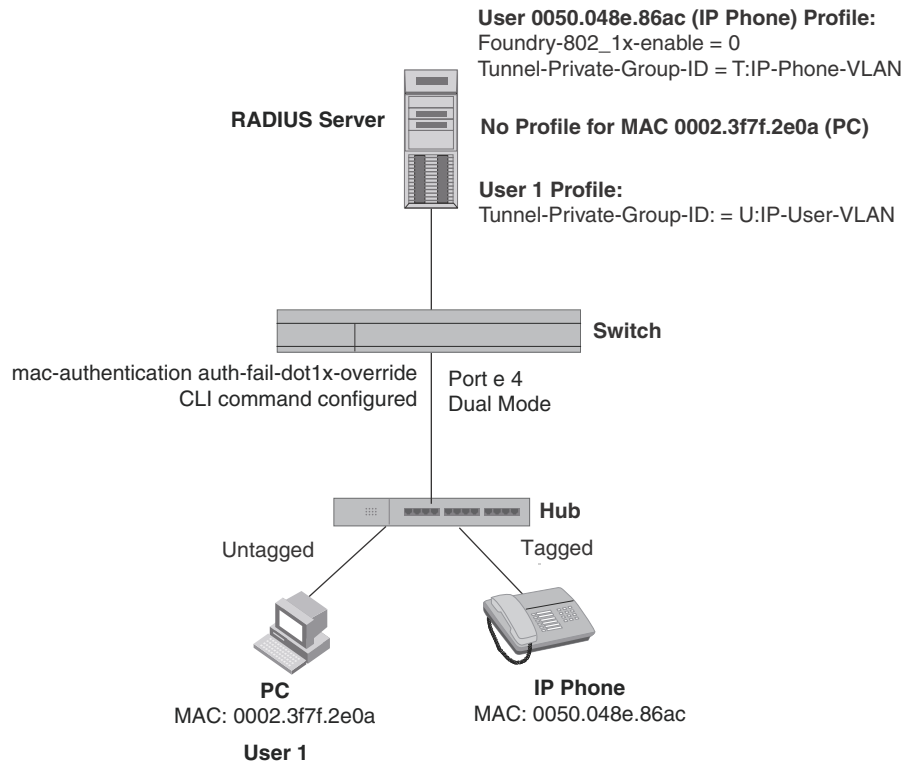
Example 2

The configuration in [Figure 126](#) requires that you create a profile on the RADIUS server for each MAC address to which a device or user can connect to the network. In a large network, this can be difficult to implement and maintain.

As an alternative, you can create MAC address profiles only for those devices that do not support 802.1X authentication, such as IP phones and printers, and configure the Dell device to perform 802.1X authentication for the other devices that do not have MAC address profiles, such as user PCs. To do this, you configure the Dell device to perform 802.1X authentication when a device fails multi-device port authentication.

Figure 127 shows a configuration where multi-device port authentication is performed for an IP phone, and 802.1X authentication is performed for a user PC. There is a profile on the RADIUS server for the IP phone MAC address, but not for the PC MAC address.

FIGURE 127 802.1X Authentication is performed when a device fails multi-device port authentication



Multi-device port authentication is initially performed for both devices. The IP phone MAC address has a profile on the RADIUS server. This profile indicates that 802.1X authentication should be skipped for this device, and that the device port be placed into the VLAN named “IP-Phone-VLAN”.

Since there is no profile for the PC MAC address on the RADIUS server, multi-device port authentication for this MAC address fails. Ordinarily, this would mean that the PVID for the port would be changed to that of the restricted VLAN, or traffic from this MAC would be blocked in hardware.

NOTE

This example assumes that the IP phone initially transmits untagged packets (for example, CDP or DHCP packets), which trigger the authentication process on the Dell device and client lookup on the RADIUS server. If the phone sends only tagged packets and the port (e 4) is not a member of that VLAN, authentication would not occur. In this case, port e 4 must be added to that VLAN prior to authentication.

To configure the Dell device to perform 802.1X authentication when a device fails multi-device port authentication, enter the following command.

```
PowerConnect(config)# mac-authentication auth-fail-dot1x-override
```

Syntax: [no] mac-authentication auth-fail-dot1x-override

28 Using multi-device port authentication and 802.1X security on the same port

Using the MAC Port Security Feature

This chapter describes how to configure devices to learn “secure” MAC addresses on an interface so that the interface will forward only packets that match the secure addresses.

Overview

You can configure the device to learn “secure” MAC addresses on an interface. The interface will forward only packets with source MAC addresses that match these learned secure addresses. The secure MAC addresses can be specified manually, or the device can learn them automatically. After the device reaches the limit for the number of secure MAC addresses it can learn on the interface, if the interface then receives a packet with a source MAC address that does not match the learned addresses, it is considered a security violation.

When a security violation occurs, a Syslog entry and an SNMP trap are generated. In addition, the device takes one of two actions; it either drops packets from the violating address (and allows packets from the secure addresses), or disables the port for a specified amount of time. You specify which of these actions takes place.

The secure MAC addresses are not flushed when an interface is disabled and re-enabled. The secure addresses can be kept secure permanently (the default), or can be configured to age out, at which time they are no longer secure. You can configure the device to automatically save the secure MAC address list to the startup-config file at specified intervals, allowing addresses to be kept secure across system restarts.

Local and global resources

The port security feature uses a concept of local and global “resources” to determine how many MAC addresses can be secured on each interface. In this context, a “resource” is the ability to store one secure MAC address entry. Each interface is allocated 64 local resources. Additional global resources are shared among all interfaces on the device.

When the port security feature is enabled on an interface, the interface can store one secure MAC address. You can increase the number of MAC addresses that can be secured using local resources to a maximum of 64.

Besides the maximum of 64 local resources available to an interface, there are additional global resources. Depending on flash memory size, a device can have 1024, 2048, or 4096 global resources available. When an interface has secured enough MAC addresses to reach its limit for local resources, it can secure additional MAC addresses by using global resources. Global resources are shared among all the interfaces on a first-come, first-served basis.

The maximum number of MAC addresses any single interface can secure is 64 (the maximum number of local resources available to the interface), plus the number of global resources not allocated to other interfaces.

Configuration notes and feature limitations

The following limitations apply to this feature:

- MAC port security applies only to Ethernet interfaces.
- MAC port security is not supported on static trunk group members or ports that are configured for link aggregation.
- MAC port security is not supported on 802.1X port security-enabled ports.
- Devices do not support the **reserved-vlan-id** <num> command, which changes the default VLAN ID for the MAC port security feature.
- The SNMP trap generated for restricted MAC addresses indicates the VLAN ID associated with the MAC address, as well as the port number and MAC address.
- MAC port security is not supported on ports that have multi-device port authentication enabled.

Configuring the MAC port security feature

To configure the MAC port security feature, perform the following tasks:

- Enable the MAC port security feature
- Set the maximum number of secure MAC addresses for an interface
- Set the port security age timer
- Specify secure MAC addresses
- Configure the device to automatically save secure MAC addresses to the startup-config file
- Specify the action taken when a security violation occurs
- Deny specific MAC addresses

Enabling the MAC port security feature

By default, the MAC port security feature is disabled on all interfaces. You can enable or disable the feature globally on all interfaces at once, or on individual interfaces.

To enable the feature on all interfaces at once, enter the following commands.

```
PowerConnect(config)# port security
PowerConnect(config-port-security)# enable
```

To disable the feature on all interfaces at once, enter the following commands.

```
PowerConnect(config)# port security
PowerConnect(config-port-security)# no enable
```

To enable the feature on a specific interface, enter the following commands.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# enable
```

Syntax: port security

Syntax: [no] enable

Setting the maximum number of secure MAC addresses for an interface

When port security is enabled, an interface can store one secure MAC address. You can increase the number of MAC addresses that can be stored to a maximum of 64, plus the total number of global resources available.

For example, to configure interface 11 to have a maximum of 10 secure MAC addresses, enter the following commands.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# maximum 10
```

Syntax: maximum <number-of-addresses>

The <number-of-addresses> parameter can be set to a number from 0 – (64 + the total number of global resources available). The total number of global resources is 2048 or 4096, depending on flash memory size. Setting the parameter to 0 prevents any addresses from being learned. The default is 1.

Setting the port security age timer

By default, learned MAC addresses stay secure indefinitely. You can optionally configure the device to age out secure MAC addresses after a specified amount of time.

To set the port security age timer to 10 minutes on all interfaces, enter the following commands.

```
PowerConnect(config)# port security
PowerConnect(config-port-security)# age 10
```

To set the port security age timer to 10 minutes on a specific interface, enter the following commands.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# age 10
```

Syntax: [no] age <minutes>

The default is 0 (never age out secure MAC addresses).

Specifying secure MAC addresses

This section describes how to configure secure MAC addresses on tagged and untagged interfaces.

On an untagged interface

To specify a secure MAC address on an untagged interface, enter commands such as the following.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# secure-mac-address 0050.DA18.747C
```

Syntax: [no] secure-mac-address <mac-address>

On a tagged interface

When specifying a secure MAC address on a tagged interface, you must also specify the VLAN ID. To do so, enter commands such as the following.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# secure-mac-address 0050.DA18.747C 2
```

Syntax: [no] **secure-mac-address** <mac-address> <vlan-ID>

NOTE

If MAC port security is enabled on a port and you change the VLAN membership of the port, make sure that you also change the VLAN ID specified in the **secure-mac-address** configuration statement for the port.

When a secure MAC address is applied to a tagged port, the **vlan-id** is generated for both tagged and untagged ports. When you display the configuration, you will see an entry for the secure MAC addresses **secure-mac-address** <address> <vlan>. For example, you may see the following line.

```
secure-mac-address 0000.1111.2222 10
```

This line means that MAC address 0000.1111.2222 on VLAN 10 is a secure MAC address.

Autosaving secure MAC addresses to the startup-config file

Learned MAC addresses can automatically be saved to the startup-config file at specified intervals. For example, to automatically save learned secure MAC addresses every twenty minutes, enter the following commands.

```
PowerConnect(config)# port security
PowerConnect(config-port-security)# autosave 20
```

Syntax: [no] **autosave** <minutes>

You can specify from 15 – 1440 minutes. By default, secure MAC addresses are not autosaved to the startup-config file.

Specifying the action taken when a security violation occurs

A security violation can occur when a user tries to connect to a port where a MAC address is already locked, or the maximum number of secure MAC addresses has been exceeded. When a security violation occurs, an SNMP trap and Syslog message are generated.

You can configure the device to take one of two actions when a security violation occurs; either drop packets from the violating address (and allow packets from secure addresses), or disable the port for a specified time.

Dropping packets from a violating address

To configure the device to drop packets from a violating address and allow packets from secure addresses, enter the following commands.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# violation restrict
```

Syntax: violation restrict

NOTE

When the **restrict** option is used, the maximum number of MAC addresses that can be restricted is 128. If the number of violating MAC addresses exceeds this number, the port is shut down. An SNMP trap and the following Syslog message are generated "Port Security violation restrict limit 128 exceeded on interface ethernet <port_id>". This is followed by a port shutdown Syslog message and trap.

Specifying the period of time to drop packets from a violating address

To specify the number of minutes that the device drops packets from a violating address, use commands similar to the following.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# violation restrict 5
```

Syntax: violation restrict <age>

<age> can be from 0 - 1440 minutes. The default is 5 minutes. Specifying 0 drops packets from the violating address permanently.

Aging for restricted MAC addresses is done in software. There can be a worst case inaccuracy of one minute from the specified time.

The restricted MAC addresses are denied in hardware.

Disabling the port for a specified amount of time

You can configure the device to disable the port for a specified amount of time when a security violation occurs.

To shut down the port for 5 minutes when a security violation occurs, enter the following commands.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# port security
PowerConnect(config-port-security-e10000-11)# violation shutdown 5
```

Syntax: violation shutdown <minutes>

You can specify from 0 - 1440 minutes. Specifying 0 shuts down the port permanently when a security violation occurs.

Clearing port security statistics

You can clear restricted MAC addresses and violation statistics from ports globally (on all ports) or on individual ports.

Clearing restricted MAC addresses

To clear all restricted MAC addresses globally, enter the following command.

```
PowerConnect#clear port security restricted-macs all
```

To clear restricted MAC addresses on a specific port, enter a command such as the following.

```
PowerConnect#clear port security restricted-macs e 5
```

Syntax: `clear port security restricted-macs all | ethernet <port-num>`

The `<portnum>` parameter is a valid port number.

Clearing violation statistics

To clear violation statistics globally, enter the following command.

```
PowerConnect#clear port security statistics all
```

To clear violation statistics on a specific port, enter a command such as the following.

```
PowerConnect#clear port security statistics e 5
```

Syntax: `clear port security statistics all | ethernet <port-num>`

The `<portnum>` parameter is a valid port number.

Displaying port security information

You can display the following information about the port security feature:

- The port security settings for an individual port or for all the ports on a specified module
- The secure MAC addresses configured on the device
- Port security statistics for an interface or for a module

Displaying port security settings

You can display the port security settings for an individual port or for all the ports on a specified module. For example, to display the port security settings for port 11, enter the following command.

```
PowerConnect# show port security e 11
Port Security Violation Shutdown-Time Age-Time Max-MAC
-----
11 disabled shutdown 10 10 1
```

Syntax: `show port security ethernet <portnum>`.

The <portnum> parameter is a valid port number.

TABLE 154 Output from the **show port security** command

This field...	Displays...
Port	The port number of the interface.
Security	Whether the port security feature has been enabled on the interface.
Violation	The action to be undertaken when a security violation occurs, either “shutdown” or “restrict”.
Shutdown-Time	The number of seconds a port is shut down following a security violation, if the port is set to “shutdown” when a violation occurs.
Age-Time	The amount of time, in minutes, MAC addresses learned on the port will remain secure.
Max-MAC	The maximum number of secure MAC addresses that can be learned on the interface.

Displaying the secure MAC addresses

To list the secure MAC addresses configured on the device, enter the following command.

```
PowerConnect# show port security mac
Port  Num-Addr  Secure-Src-Addr  Resource  Age-Left  Shutdown/Time-Left
-----
11      1    0050.da18.747c   Local      10        no
```

Syntax: show port security mac

This command displays the following information.

TABLE 155 Output from the **show port security mac** command

This field...	Displays...
Port	The port number of the interface.
Num-Addr	The number of MAC addresses secured on this interface.
Secure-Src-Addr	The secure MAC address.
Resource	Whether the address was secured using a local or global resource. Refer to “Local and global resources” on page 973 for more information.
Age-Left	The number of minutes the MAC address will remain secure.
Shutdown/Time-Left	Whether the interface has been shut down due to a security violation and the number of seconds before it is enabled again.

Displaying port security statistics

You can display port security statistics for an interface or for a module.

For example, to display port security statistics for interface 11, enter the following command.

```
PowerConnect# show port security statistics e 11
Port  Total-Addrs  Maximum-Addrs  Violation  Shutdown/Time-Left
-----
11      1              1              0         no
```

Syntax: show port security statistics <portnum>

The `<portnum>` parameter is a valid port number.

TABLE 156 Output from the `show port security statistics <portnum>` command

This field...	Displays...
Port	The port number of the interface.
Total-Addrs	The total number of secure MAC addresses on the interface.
Maximum-Addrs	The maximum number of secure MAC addresses on the interface.
Violation	The number of security violations on the port.
Shutdown/Time-Left	Whether the port has been shut down due to a security violation and the number of seconds before it is enabled again.

To display port security statistics for an interface module, enter the following command.

```
PowerConnect# show port security statistics 7
Module 7:
  Total ports: 0
  Total MAC address(es): 0
  Total violations: 0
  Total shutdown ports 0
```

Syntax: `show port security statistics <module>`

TABLE 157 Output from the `show port security statistics <module>` command

This field...	Displays...
Total ports	The number of ports on the module.
Total MAC address(es)	The total number of secure MAC addresses on the module.
Total violations	The number of security violations encountered on the module.
Total shutdown ports	The number of ports on the module shut down as a result of security violations.

Displaying restricted MAC addresses on a port

To display a list of restricted MAC addresses on a port, enter a command such as the following.

```
PowerConnect# show port security e 5 restricted-macs
```

Syntax: `show port security ethernet <portnum> restricted-macs`

The `<portnum>` parameter is a valid port number.

Configuring Multi-Device Port Authentication

Multi-device port authentication is a way to configure a device to forward or block traffic from a MAC address based on information received from a RADIUS server.

How multi-device port authentication works

Multi-device port authentication is a way to configure a device to forward or block traffic from a MAC address based on information received from a RADIUS server.

The multi-device port authentication feature is a mechanism by which incoming traffic originating from a specific MAC address is switched or forwarded by the device only if the source MAC address is successfully authenticated by a RADIUS server. The MAC address itself is used as the username and password for RADIUS authentication; the user does not need to provide a specific username and password to gain access to the network. If RADIUS authentication for the MAC address is successful, traffic from the MAC address is forwarded in hardware.

If the RADIUS server cannot validate the user's MAC address, then it is considered an authentication failure, and a specified authentication-failure action can be taken. The default authentication-failure action is to drop traffic from the non-authenticated MAC address in hardware. You can also configure the device to move the port on which the non-authenticated MAC address was learned into a restricted or "guest" VLAN, which may have limited access to the network.

NOTE

PowerConnect devices do not support:

- multi-device authentication and port security configured on the same port
 - multi-device authentication and lock-address configured on the same port
-

RADIUS authentication

The multi-device port authentication feature communicates with the RADIUS server to authenticate a newly found MAC address. The device supports multiple RADIUS servers; if communication with one of the RADIUS servers times out, the others are tried in sequential order. If a response from a RADIUS server is not received within a specified time (by default, 3 seconds) the RADIUS session times out, and the device retries the request up to three times. If no response is received, the next RADIUS server is chosen, and the request is sent for authentication.

The RADIUS server is configured with the usernames and passwords of authenticated users. For multi-device port authentication, the username and password is the MAC address itself; that is, the device uses the MAC address for both the username and the password in the request sent to the RADIUS server. For example, given a MAC address of 0007e90feaa1, the users file on the RADIUS server would be configured with a username and password both set to 0007e90feaa1. When

traffic from this MAC address is encountered on a MAC-authentication-enabled interface, the device sends the RADIUS server an Access-Request message with 0007e90feaa1 as both the username and password. The format of the MAC address sent to the RADIUS server is configurable through the CLI.

The request for authentication from the RADIUS server is successful only if the username and password provided in the request matches an entry in the users database on the RADIUS server. When this happens, the RADIUS server returns an Access-Accept message back to the device. When the RADIUS server returns an Access-Accept message for a MAC address, that MAC address is considered authenticated, and traffic from the MAC address is forwarded normally by the device.

Authentication-failure actions

If the MAC address does not match the username and password of an entry in the users database on the RADIUS server, then the RADIUS server returns an Access-Reject message. When this happens, it is considered an authentication failure for the MAC address. When an authentication failure occurs, the device can either drop traffic from the MAC address in hardware (the default), or move the port on which the traffic was received to a restricted VLAN.

Supported RADIUS attributes

Devices support the following RADIUS attributes for multi-device port authentication:

- Username (1) – RFC 2865
- NAS-IP-Address (4) – RFC 2865
- NAS-Port (5) – RFC 2865
- Service-Type (6) – RFC 2865
- FilterId (11) – RFC 2865
- Framed-MTU (12) – RFC 2865
- State (24) – RFC 2865
- Vendor-Specific (26) – RFC 2865
- Session-Timeout (27) – RFC 2865
- Termination-Action (29) – RFC 2865
- Calling-Station-ID (31) – RFC 2865
- Tunnel-Type (64) – RFC 2868
- Tunnel-Medium-Type (65) – RFC 2868
- EAP Message (79) – RFC 2579
- Message-Authenticator (80) RFC 3579
- Tunnel-Private-Group-Id (81) – RFC 2868
- NAS-Port-id (87) – RFC2869

Support for dynamic VLAN assignment

The multi-device port authentication feature supports **dynamic VLAN assignment**, where a port can be placed in one or more VLANs based on the MAC address learned on that interface. For details about this feature, refer to [“Configuring the RADIUS server to support dynamic VLAN assignment”](#) on page 989.

Support for dynamic ACLs

The multi-device port authentication implementation supports the assignment of a MAC address to a specific ACL, based on the MAC address learned on the interface. For details about this feature, refer to [“Dynamically applying IP ACLs to authenticated MAC addresses”](#) on page 990.

Support for authenticating multiple MAC addresses on an interface

The multi-device port authentication feature allows multiple MAC addresses to be authenticated or denied authentication on each interface. The maximum number of MAC addresses that can be authenticated on each interface is limited only by the amount of system resources available on the device.

Using multi-device port authentication and 802.1X security on the same port

Multi-device port authentication and 802.1X security can be configured on the same port. When both of these features are enabled on the same port, multi-device port authentication is performed prior to 802.1X authentication. If multi-device port authentication is successful, 802.1X authentication may be performed, based on the configuration of a vendor-specific attribute (VSA) in the profile for the MAC address on the RADIUS server.

NOTE

When multi-device port authentication and 802.1X security are configured together on the same port, Dell recommends that dynamic VLANs and dynamic ACLs are done at the multi-device port authentication level, and not at the 802.1X level.

When both features are configured on a port, a device connected to the port is authenticated as follows.

1. Multi-device port authentication is performed on the device to authenticate the device MAC address.
2. If multi-device port authentication is successful for the device, then the device checks whether the RADIUS server included the Foundry-802_1x-enable VSA (described in [Table 158](#)) in the Access-Accept message that authenticated the device.
3. If the Foundry-802_1x-enable VSA is not present in the Access-Accept message, or is present and set to 1, then 802.1X authentication is performed for the device.

4. If the Foundry-802_1x-enable VSA is present in the Access-Accept message, and is set to 0, then 802.1X authentication is skipped. The device is authenticated, and any dynamic VLANs specified in the Access-Accept message returned during multi-device port authentication are applied to the port.
5. If 802.1X authentication is performed on the device, and is successful, then dynamic VLANs or ACLs specified in the Access-Accept message returned during 802.1X authentication are applied to the port.

If multi-device port authentication fails for a device, then by default traffic from the device is either blocked in hardware, or the device is placed in a restricted VLAN. You can optionally configure the device to perform 802.1X authentication on a device when it fails multi-device port authentication.

Configuring Dell-specific attributes on the RADIUS server

If the RADIUS authentication process is successful, the RADIUS server sends an Access-Accept message to the device, authenticating the device. The Access-Accept message can include Vendor-Specific Attributes (VSAs) that specify additional information about the device. If you are configuring multi-device port authentication and 802.1X authentication on the same port, then you can configure the Dell VSAs listed in [Table 158](#) on the RADIUS server.

You add these Dell vendor-specific attributes to your RADIUS server configuration, and configure the attributes in the individual or group profiles of the devices that will be authenticated. The Dell Vendor-ID is 1991, with Vendor-Type 1.

TABLE 158 Dell vendor-specific attributes for RADIUS

Attribute name	Attribute ID	Data type	Description
Foundry-802_1x-enable	6	integer	Specifies whether 802.1X authentication is performed when multi-device port authentication is successful for a device. This attribute can be set to one of the following: 0 - Do not perform 802.1X authentication on a device that passes multi-device port authentication. Set the attribute to zero for devices that do not support 802.1X authentication. 1 - Perform 802.1X authentication when a device passes multi-device port authentication. Set the attribute to one for devices that support 802.1X authentication.
Foundry-802_1x-valid	7	integer	Specifies whether the RADIUS record is valid only for multi-device port authentication, or for both multi-device port authentication and 802.1X authentication. This attribute can be set to one of the following: 0 - The RADIUS record is valid only for multi-device port authentication. Set this attribute to zero to prevent a user from using their MAC address as username and password for 802.1X authentication 1 - The RADIUS record is valid for both multi-device port authentication and 802.1X authentication.

If neither of these VSAs exist in a device profile on the RADIUS server, then by default the device is subject to multi-device port authentication (if configured), then 802.1X authentication (if configured). The RADIUS record can be used for both multi-device port authentication and 802.1X authentication.

Configuring multi-device port authentication

Configuring multi-device port authentication on the device consists of the following tasks:

- Enabling multi-device port authentication globally and on individual interfaces
- Specifying the format of the MAC addresses sent to the RADIUS server (optional)
- Specifying the authentication-failure action (optional)
- Enabling and disabling SNMP traps for multi-device port authentication
- Defining MAC address filters (optional)
- Configuring dynamic VLAN assignment (optional)
- Dynamically Applying IP ACLs to authenticated MAC addresses
- Enabling denial of service attack protection (optional)
- Clearing authenticated MAC addresses (optional)
- Disabling aging for authenticated MAC addresses (optional)
- Configuring the hardware aging period for blocked MAC addresses
- Specifying the aging time for blocked MAC addresses (optional)

NOTE

On PowerConnect B-Series T124X devices, the 802.1X port security feature is supported.

Enabling multi-device port authentication

To enable multi-device port authentication, you first enable the feature globally on the device. On some devices, you can then enable the feature on individual interfaces.

Globally enabling multi-device port authentication

To globally enable multi-device port authentication on the device, enter the following command.

```
PowerConnect(config)# mac-authentication enable
```

Syntax: [no] mac-authentication enable

Enabling multi-device port authentication on an interface

To enable multi-device port authentication on an individual interface, enter a command such as the following.

```
PowerConnect(config)# mac-authentication enable ethernet 1
```

Syntax: [no] mac-authentication enable <portnum> | all

The <portnum> parameter is a valid port number.

The **all** option enables the feature on all interfaces at once.

You can enable the feature on an interface at the interface CONFIG level.

Example

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication enable
```

Syntax: [no] mac-authentication enable

You can also configure multi-device port authentication commands on a range of interfaces.

Example

```
PowerConnect(config)# int e 1 to 12
PowerConnect(config-mif-1-12)# mac-authentication enable
```

Specifying the format of the MAC addresses sent to the RADIUS server

When multi-device port authentication is configured, the Dell device authenticates MAC addresses by sending username and password information to a RADIUS server. The username and password is the MAC address itself; that is, the device uses the MAC address for both the username and the password in the request sent to the RADIUS server.

By default, the MAC address is sent to the RADIUS server in the format xxxxxxxxxxxx. You can optionally configure the device to send the MAC address to the RADIUS server in the format xx-xx-xx-xx-xx-xx, or the format xxx.xxxx.xxxx. To do this, enter a command such as the following.

```
PowerConnect(config)# mac-authentication auth-passwd-format xxx.xxxx.xxxx
```

Syntax: [no] mac-authentication auth-passwd-format xxx.xxxx.xxxx | xx-xx-xx-xx-xx-xx | xxxxxxxxxxxx

Specifying the authentication-failure action

When RADIUS authentication for a MAC address fails, you can configure the device to perform one of two actions:

- Drop traffic from the MAC address in hardware (the default)
- Move the port on which the traffic was received to a restricted VLAN

To configure the device to move the port to a restricted VLAN when multi-device port authentication fails, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication auth-fail-action
restrict-vlan 100
```

Syntax: [no] mac-authentication auth-fail-action restrict-vlan [*<vlan-id>*]

If the ID for the restricted VLAN is not specified at the interface level, the global restricted VLAN ID applies for the interface.

To specify the VLAN ID of the restricted VLAN globally, enter the following command.

```
PowerConnect(config)# mac-authentication auth-fail-vlan-id 200
```

Syntax: [no] mac-authentication auth-fail-vlan-id *<vlan-id>*

The command above applies globally to all MAC-authentication-enabled interfaces.

Note that the restricted VLAN must already exist on the device. You cannot configure the restricted VLAN to be a non-existent VLAN. If the port is a tagged or dual-mode port, you cannot use a restricted VLAN as the authentication-failure action.

To configure the device to drop traffic from non-authenticated MAC addresses in hardware, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication auth-fail-action
block-traffic
```

Syntax: [no] mac-authentication auth-fail-action block-traffic

Dropping traffic from non-authenticated MAC addresses is the default behavior when multi-device port authentication is enabled.

Generating traps for multi-device port authentication

You can enable and disable SNMP traps for multi-device port authentication. SNMP traps are enabled by default.

To enable SNMP traps for multi-device port authentication after they have been disabled, enter the following command.

```
PowerConnect(config)# snmp-server enable traps mac-authentication
```

Syntax: [no] snmp-server enable traps mac-authentication

Use the **no** form of the command to disable SNMP traps for multi-device port authentication.

Defining MAC address filters

You can specify MAC addresses that do not have to go through multi-device port authentication. These MAC addresses are considered pre-authenticated, and are not subject to RADIUS authentication. To do this, you can define MAC address filters that specify the MAC addresses to exclude from multi-device port authentication.

You should use a MAC address filter when the RADIUS server itself is connected to an interface where multi-device port authentication is enabled. If a MAC address filter is not defined for the MAC address of the RADIUS server and applied on the interface, the RADIUS authentication process would fail since the device would drop all packets from the RADIUS server itself.

For example, the following command defines a MAC address filter for address 0010.dc58.aca4.

```
PowerConnect(config)# mac-authentication mac-filter 1 0010.dc58.aca4
```

Syntax: [no] mac-authentication mac-filter <filter>

The following commands apply the MAC address filter on an interface so that address 0010.dc58.aca4 is excluded from multi-device port authentication.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication apply-mac-auth-filter 1
```

Syntax: [no] mac-authentication apply-mac-auth-filter <filter-id>

Configuring dynamic VLAN assignment

An interface can be dynamically assigned to one or more VLANs based on the MAC address learned on that interface. When a MAC address is successfully authenticated, the RADIUS server sends the Dell device a RADIUS Access-Accept message that allows the Dell device to forward traffic from that MAC address. The RADIUS Access-Accept message can also contain attributes set for the MAC address in its access profile on the RADIUS server.

If one of the attributes in the Access-Accept message specifies one or more VLAN identifiers, and the VLAN is available on the Dell device, the port is moved from its default VLAN to the specified VLAN.

To enable dynamic VLAN assignment for authenticated MAC addresses, you must add attributes to the profile for the MAC address on the RADIUS server, then enable dynamic VLAN assignment on multi-device port authentication-enabled interfaces. Refer to [“Configuring the RADIUS server to support dynamic VLAN assignment”](#) on page 989 for a list of the attributes that must be set on the RADIUS server.

To enable dynamic VLAN assignment on a multi-device port authentication-enabled interface, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication enable-dynamic-vlan
```

Syntax: [no] mac-authentication enable-dynamic-vlan

Configuring a port to remain in the restricted VLAN after a successful authentication attempt

If a previous authentication attempt for a MAC address failed, and as a result the port was placed in the restricted VLAN, but a subsequent authentication attempt was successful, the RADIUS Access-Accept message may specify a VLAN for the port. By default, the Dell device moves the port out of the restricted VLAN and into the RADIUS-specified VLAN. You can optionally configure the device to leave the port in the restricted VLAN. To do this, enter the following command.

```
PowerConnect(config-if-e10000-1)# mac-authentication no-override-restrict-vlan
```

When the above command is applied, if the RADIUS-specified VLAN configuration is tagged (e.g., T:1024) and the VLAN is valid, then the port is placed in the RADIUS-specified VLAN as a tagged port and left in the restricted VLAN. If the RADIUS-specified VLAN configuration is untagged (e.g., U:1024), the configuration from the RADIUS server is ignored, and the port is left in the restricted VLAN.

Syntax: [no] mac-authentication no-override-restrict-vlan

Configuration notes

- If you configure dynamic VLAN assignment on a multi-device port authentication enabled interface, and the Access-Accept message returned by the RADIUS server contains a Tunnel-Type and Tunnel-Medium-Type, but does not contain a Tunnel-Private-Group-ID attribute, then it is considered an authentication failure, and the configured authentication failure action is performed for the MAC address.
- If the <vlan-name> string does not match either the name or the ID of a VLAN configured on the device, then it is considered an authentication failure, and the configured authentication failure action is performed for the MAC address.

- For tagged or dual-mode ports, if the VLAN ID provided by the RADIUS server does not match the VLAN ID in the tagged packet that contains the authenticated MAC address as its source address, then it is considered an authentication failure, and the configured authentication failure action is performed for the MAC address.
- If an untagged port had previously been assigned to a VLAN through dynamic VLAN assignment, and then another MAC address is authenticated on the same port, but the RADIUS Access-Accept message for the second MAC address specifies a different VLAN, then it is considered an authentication failure for the second MAC address, and the configured authentication failure action is performed. Note that this applies only if the first MAC address has not yet aged out. If the first MAC address has aged out, then dynamic VLAN assignment would work as expected for the second MAC address.
- For dual mode ports, if the RADIUS server returns T:<vlan-name>, the traffic will still be forwarded in the statically assigned PVID. If the RADIUS server returns U:<vlan-name>, the traffic will not be forwarded in the statically assigned PVID.

Configuring the RADIUS server to support dynamic VLAN assignment

To specify VLAN identifiers on the RADIUS server, add the following attributes to the profile for the MAC address on the RADIUS server, then enable dynamic VLAN assignment on multi-device port authentication-enabled interfaces.

Table 11:

Attribute name	Type	Value
Tunnel-Type	064	13 (decimal) – VLAN
Tunnel-Medium-Type	065	6 (decimal) – 802
Tunnel-Private-Group-ID	081	<vlan-name> (string) The <vlan-name> value can specify either the name or the number of one or more VLANs configured on the Dell device.

For information about the attributes, refer to [“Dynamic multiple VLAN assignment for 802.1X ports”](#) on page 939.

Specifying to which VLAN a port is moved after its RADIUS-specified VLAN assignment expires

When a port is dynamically assigned to a VLAN through the authentication of a MAC address, and the MAC session for that address is deleted on the Dell device, then by default the port is removed from its RADIUS-assigned VLAN and placed back in the VLAN where it was originally assigned.

A port can be removed from its RADIUS-assigned VLAN when any of the following occur:

- The link goes down for the port
- The MAC session is manually deleted with the **mac-authentication clear-mac-session** command
- The MAC address that caused the port to be dynamically assigned to a VLAN ages out

For example, say port 1 is currently in VLAN 100, to which it was assigned when MAC address 0007.eaa1.e90f was authenticated by a RADIUS server. The port was originally configured to be in VLAN 111. If the MAC session for address 0007.eaa1.e90f is deleted, then port 1 is moved from VLAN 100 back into VLAN 111.

You can optionally specify an alternate VLAN to which to move the port when the MAC session for the address is deleted. For example, to place the port in the restricted VLAN, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-auth move-back-to-old-vlan
port-restrict-vlan
```

Syntax: [no] mac-authentication move-back-to-old-vlan disable | port-configured-vlan | system-default-vlan

The **disable** keyword disables moving the port back to its original VLAN. The port would stay in its RADIUS-assigned VLAN.

The **port-configured-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it back in the VLAN where it was originally assigned. This is the default.

The **port-restrict-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it in the restricted VLAN.

The **system-default-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it in the DEFAULT-VLAN.

NOTE

When a MAC session is deleted, if the port is moved back to a VLAN that is different than the running-config file, the system will update the running-config file to reflect the changes. This will occur even if **mac-authentication save-dynamicvlan-to-config** is not configured.

Saving dynamic VLAN assignments to the running-config file

By default, dynamic VLAN assignments are not saved to the running-config file of the device. However, you can configure the device to do so by entering the following command.

```
PowerConnect(config)# mac-authentication save-dynamicvlan-to-config
```

When the above command is applied, dynamic VLAN assignments are saved to the running-config file and are displayed when the **show run** command is issued. Dynamic VLAN assignments can also be displayed with the **show vlan**, **show auth-mac-addresses detail**, and **show auth-mac-addresses authorized-mac** commands.

Syntax: [no] mac-authentication save-dynamicvlan-to-config

Dynamically applying IP ACLs to authenticated MAC addresses

The multi-device port authentication implementation supports the assignment of a MAC address to a specific ACL, based on the MAC address learned on the interface.

When a MAC address is successfully authenticated, the RADIUS server sends the device a RADIUS Access-Accept message that allows the device to forward traffic from that MAC address. The RADIUS Access-Accept message can also contain, among other attributes, the Filter-ID (type 11) attribute for the MAC address. When the Access-Accept message containing the Filter-ID (type 11) attribute is received by the device, it will use the information in these attributes to apply an IP ACL on a per-MAC (per user) basis.

The dynamic IP ACL is active as long as the client is connected to the network. When the client disconnects from the network, the IP ACL is no longer applied to the port. If an IP ACL had been applied to the port prior to multi-device port authentication; it will be re-applied to the port.

The device uses information in the Filter ID to apply an IP ACL on a per-user basis. The Filter-ID attribute can specify the number of an existing IP ACL configured on the device. If the Filter-ID is an ACL number, the specified IP ACL is applied on a per-user basis.

Configuration considerations and guidelines

- Dynamic IP ACLs with multi-device port authentication are supported. Dynamic MAC filters with multi-device port authentication are not supported.
- In the Layer 2 switch code, dynamic IP ACLs are not supported when **ACL-per-port-per-vlan** is enabled on a global-basis.
- The RADIUS Filter ID (type 11) attribute is supported. The Vendor-Specific (type 26) attribute is not supported.
- The dynamic ACL must be an extended ACL. Standard ACLs are not supported.
- Multi-device port authentication and 802.1x can be used together on the same port. However, Dell does not recommend the use of multi-device port authentication and 802.1x with dynamic ACLs together on the same port. If a single supplicant requires both 802.1x and multi-device port authentication, and if both 802.1x and multi-device port authentication try to install different dynamic ACLs for the same supplicant, the supplicant will fail authentication.
- Dynamically assigned IP ACLs are subject to the same configuration restrictions as non-dynamically assigned IP ACLs. One caveat is that ports with VE interfaces cannot have assigned user-defined ACLs. For example, a user-defined ACL bound to a VE or a port on a VE is not allowed. There are no restrictions on ports that do not have VE interfaces.
- Dynamic ACL filters are supported only for the inbound direction. Dynamic outbound ACL filters are not supported.
- Dynamic ACL assignment with multi-device port authentication is not supported in conjunction with any of the following features:
 - IP source guard
 - Rate limiting
 - Protection against ICMP or TCP Denial-of-Service (DoS) attacks
 - Policy-based routing
 - 802.1X dynamic filter
- On PowerConnect B-Series TI24X devices, the dynamic ACLs are not supported on ports that have 802.1x and MAC authentication enabled along with the **auth-fail-dot1x-override** option enabled.
- On PowerConnect devices, MAC authentication and 802.1X authentication can be configured on the same port. When both of these features are enabled on the same port, MAC authentication is performed prior to 802.1X authentication. If MAC authentication is successful, 802.1X authentication may be performed. If 802.1X authentication is successful, dynamic VLAN and ACL filters for 802.1X are applied.
- PowerConnect devices authenticate devices attached to VLAN ports. The MAC address is used as the username and password for port authentication. The 802.1X authentication uses the username and password for port authentication.

Configuring the RADIUS server to support dynamic IP ACLs

When a port is authenticated using multi-device port authentication, an IP ACL filter that exists in the running-config file on the device can be dynamically applied to the port. To do this, you configure the Filter-ID (type 11) attribute on the RADIUS server. The Filter-ID attribute specifies the name or number of the IP ACL.

The following is the syntax for configuring the Filter-ID attribute on the RADIUS server to refer to a IP ACL.

Table 12:

Value	Description
ip.<number>.in ¹	Applies the specified numbered ACL to the authenticated port in the inbound direction.
ip.<name>.in ^{1,2}	Applies the specified named ACL to the authenticated port in the inbound direction.

1. The ACL must be an extended ACL. Standard ACLs are not supported.
2. The <name> in the Filter ID attribute is case-sensitive

The following table lists examples of values you can assign to the Filter-ID attribute on the RADIUS server to refer to IP ACLs configured on a device.

Table 13:

Possible values for the filter ID attribute on the RADIUS server	ACLs configured on the Dell device
ip.102.in	access-list 102 permit ip 36.0.0.0 0.255.255.255 any
ip.fdry_filter.in	ip access-list standard fdry_filter permit host 36.48.0.3

Enabling denial of service attack protection

The Dell device does not start forwarding traffic from an authenticated MAC address in hardware until the RADIUS server authenticates the MAC address; traffic from the non-authenticated MAC addresses is sent to the CPU. A denial of service (DoS) attack could be launched against the device where a high volume of new source MAC addresses is sent to the device, causing the CPU to be overwhelmed with performing RADIUS authentication for these MAC addresses. In addition, the high CPU usage in such an attack could prevent the RADIUS response from reaching the CPU in time, causing the device to make additional authentication attempts.

To limit the susceptibility of the device to such attacks, you can configure the device to use multiple RADIUS servers, which can share the load when there are a large number of MAC addresses that need to be authenticated. The device can run a maximum of 10 RADIUS clients per server and will attempt to authenticate with a new RADIUS server if current one times out.

In addition, you can configure the device to limit the rate of authentication attempts sent to the RADIUS server. When the multi-device port authentication feature is enabled, it keeps track of the number of RADIUS authentication attempts made per second. When you also enable the DoS protection feature, if the number of RADIUS authentication attempts for MAC addresses learned on an interface per second exceeds a configurable rate (by default 512 authentication attempts per second), the device considers this a possible DoS attack and disables the port. You must then manually re-enable the port.

The DoS protection feature is disabled by default. To enable it on an interface, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication dos-protection enable
```

To specify a maximum rate for RADIUS authentication attempts, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication dos-protection mac-limit 256
```

Syntax: [no] **mac-authentication dos-protection mac-limit** <number>

You can specify a rate from 1 – 65535 authentication attempts per second. The default is a rate of 512 authentication attempts per second.

Clearing authenticated MAC addresses

The Dell device maintains an internal table of the authenticated MAC addresses (viewable with the **show authenticated-mac-address** command). You can clear the contents of the authenticated MAC address table either entirely, or just for the entries learned on a specified interface. In addition, you can clear the MAC session for an address learned on a specific interface.

To clear the entire contents of the authenticated MAC address table, enter the following command.

```
PowerConnect# clear auth-mac-table
```

Syntax: clear auth-mac-table

To clear the authenticated MAC address table of entries learned on a specified interface, enter a command such as the following.

```
PowerConnect# clear auth-mac-table e1
```

Syntax: clear auth-mac-table <portnum>

The <portnum> parameter is a valid port number.

To clear the MAC session for an address learned on a specific interface, enter commands such as the following.

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication clear-mac-session
00e0.1234.abd4
```

Syntax: mac-authentication clear-mac-session <mac-address>

This command removes the Layer 2 CAM entry created for the specified MAC address. If the device receives traffic from the MAC address again, the MAC address is authenticated again.

Disabling aging for authenticated MAC addresses

MAC addresses that have been authenticated or denied by a RADIUS server are aged out if no traffic is received from the MAC address for a certain period of time:

- Authenticated MAC addresses or non-authenticated MAC addresses that have been placed in the restricted VLAN are aged out if no traffic is received from the MAC address over the device normal MAC aging interval.

- Non-authenticated MAC addresses that are blocked by the device are aged out if no traffic is received from the address over a fixed hardware aging period (70 seconds), plus a configurable software aging period. (Refer to the next section for more information on configuring the software aging period).

You can optionally disable aging for MAC addresses subject to authentication, either for all MAC addresses or for those learned on a specified interface.

Globally disabling aging of MAC addresses

On most devices, you can disable aging for all MAC addresses on all interfaces where multi-device port authentication has been enabled by entering the following command.

```
PowerConnect(config)# mac-authentication disable-aging
```

Syntax: mac-authentication disable-aging

Enter the command at the global or interface configuration level.

The **denied-only** parameter prevents denied sessions from being aged out, but ages out permitted sessions.

The **permitted-only** parameter prevents permitted (authenticated and restricted) sessions from being aged out and ages denied sessions.

Disabling the aging of MAC addresses on interfaces

To disable aging for all MAC addresses subject to authentication on a specific interface where multi-device port authentication has been enabled, enter the command at the interface level.

Example

```
PowerConnect(config)# interface e 1
PowerConnect(config-if-e10000-1)# mac-authentication disable-aging
```

Syntax: [no] mac-authentication disable-aging

Changing the hardware aging period for blocked MAC addresses

When the Dell device is configured to drop traffic from non-authenticated MAC addresses, traffic from the blocked MAC addresses is dropped in hardware, without being sent to the CPU. A Layer 2 hardware entry is created that drops traffic from the MAC address in hardware. If no traffic is received from the MAC address for a certain amount of time, this Layer 2 hardware entry is aged out. If traffic is subsequently received from the MAC address, then an attempt can be made to authenticate the MAC address again.

Aging of the Layer 2 hardware entry for a blocked MAC address occurs in two phases, known as hardware aging and software aging.

The software aging period for blocked MAC addresses is configurable through the CLI, with the **mac-authentication max-age** command. Once the hardware aging period ends, the software aging period begins. When the software aging period ends, the blocked MAC address ages out, and can be authenticated again if the device receives traffic from the MAC address.

To change the hardware aging period for blocked MAC addresses, enter a command such as the following.


```
PowerConnect(config)# mac-authentication hw-deny-age 10
```

Syntax: [no] **mac-authentication hw-deny-age** <num>

The <num> parameter is a value from 1 to 65535 seconds. The default is 70 seconds.

Specifying the aging time for blocked MAC addresses

When the Dell device is configured to drop traffic from non-authenticated MAC addresses, traffic from the blocked MAC addresses is dropped in hardware, without being sent to the CPU. A Layer 2 CAM entry is created that drops traffic from the blocked MAC address in hardware. If no traffic is received from the blocked MAC address for a certain amount of time, this Layer 2 CAM entry is aged out. If traffic is subsequently received from the MAC address, then an attempt can be made to authenticate the MAC address again.

Aging of the Layer 2 CAM entry for a blocked MAC address occurs in two phases, known as **hardware aging** and **software aging**. The hardware aging period is fixed at 70 seconds and is non-configurable. The software aging time is configurable through the CLI.

Once the Dell device stops receiving traffic from a blocked MAC address, the hardware aging begins and lasts for a fixed period of time. After the hardware aging period ends, the software aging period begins. The software aging period lasts for a configurable amount of time (by default 120 seconds). After the software aging period ends, the blocked MAC address ages out, and can be authenticated again if the device receives traffic from the MAC address.

To change the length of the software aging period for blocked MAC addresses, enter a command such as the following.

```
PowerConnect(config)# mac-authentication max-age 180
```

Syntax: [no] **mac-authentication max-age** <seconds>

You can specify from 1 – 65535 seconds. The default is 120 seconds.

Specifying the RADIUS timeout action

A RADIUS timeout occurs when the Dell device does not receive a response from a RADIUS server within a specified time limit and after a certain number of retries. The time limit and number of retries can be manually configured using the CLI commands **radius-server timeout** and **radius-server retransmit**, respectively. If the parameters are not manually configured, the device applies the default value of three seconds with a maximum of three retries.

You can better control port behavior when a RADIUS timeout occurs by configuring a port on the device to automatically pass or fail user authentication. A **pass** essentially bypasses the authentication process and permits user access to the network. A **fail** bypasses the authentication process and blocks user access to the network, unless restrict-vlan is configured, in which case, the user is placed into a VLAN with restricted or limited access. By default, the device will reset the authentication process and retry to authenticate the user.

Specify the RADIUS timeout action at the Interface level of the CLI.

Permit User access to the network after a RADIUS timeout

To set the RADIUS timeout behavior to bypass multi-device port authentication and *permit* user access to the network, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 3
PowerConnect(config-if-e100-3)# mac-authentication auth-timeout-action success
```

Syntax: [no] mac-authentication auth-timeout-action success

Once the *success* timeout action is enabled, use the *no* form of the command to reset the RADIUS timeout behavior to *retry*.

Deny User access to the network after a RADIUS timeout

To set the RADIUS timeout behavior to bypass multi-device port authentication and block user access to the network, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 3
PowerConnect(config-if-e100-3)# mac-authentication auth-timeout-action failure
```

Syntax: [no] mac-authentication auth-timeout-action failure

Once the failure timeout action is enabled, use the **no** form of the command to reset the RADIUS timeout behavior to *retry*.

NOTE

If **restrict-vlan** is configured along with **auth-timeout-action failure**, the user will be placed into a VLAN with restricted or limited access. Refer to [“Allow user access to a restricted VLAN after a RADIUS timeout”](#) on page 996.

Allow user access to a restricted VLAN after a RADIUS timeout

To set the RADIUS timeout behavior to bypass multi-device port authentication and place the user in a VLAN with restricted or limited access, enter commands such as the following.

```
PowerConnect(config)# interface ethernet 3
PowerConnect(config-if-e100-3)# mac-authentication auth-fail-action restrict-vlan
100
PowerConnect(config-if-e100-3)# mac-authentication auth-timeout-action failure
```

Syntax: [no] mac-authentication auth-fail-action restrict-vlan [*<vlan-id>*]

Syntax: [no] mac-authentication auth-timeout-action failure

Multi-device port authentication password override

The multi-device port authentication feature communicates with the RADIUS server to authenticate a newly found MAC address. The RADIUS server is configured with the usernames and passwords of authenticated users. For multi-device port authentication, the username and password is the MAC address itself; that is, the device uses the MAC address for both the username and the password in the request sent to the RADIUS server. For example, given a MAC address of 0007e90feaa1, the users file on the RADIUS server would be configured with a username and password both set to 0007e90feaa1. When traffic from this MAC address is encountered on a MAC-authentication-enabled interface, the device sends the RADIUS server an Access-Request message with 0007e90feaa1 as both the username and password.

To change the password for multi-device port authentication, enter a command such as the following at the GLOBAL Config Level of the CLI.

```
PowerConnect(config)# mac-authentication password-override
```

Syntax: [no] **mac-authentication password-override** <password>

where <password> can have up to 32 alphanumeric characters, but cannot include blank spaces.

Limiting the number of authenticated MAC addresses

You cannot enable MAC port security on the same port that has multi-device port authentication enabled. To simulate the function of MAC port security, you can enter a command such as the following.

```
PowerConnect(config-if-e10000-2)# mac-authentication max-accepted-session 5
```

Syntax: [no] **mac-authentication max-accepted-session** <session-number>

This command limits the number of successfully authenticated MAC addresses. Enter a value from 1 - 250 for session-number

Displaying multi-device port authentication information

You can display the following information about the multi-device port authentication configuration:

- Information about authenticated MAC addresses
- Information about the multi-device port authentication configuration
- Authentication Information for a specific MAC address or port
- Multi-device port authentication settings and authenticated MAC addresses for each port where the multi-device port authentication feature is enabled
- The MAC addresses that have been successfully authenticated
- The MAC addresses for which authentication was not successful

Displaying authenticated MAC address information

To display information about authenticated MAC addresses on the ports where the multi-device port authentication feature is enabled, enter the following command.

```
PowerConnect# show auth-mac-address
```

Port	Vlan	Accepted MACs	Rejected MACs	Attempted-MACs
18	100	1	100	0
20	40	0	0	0
22	100	0	0	0
5	30	0	0	0

Syntax: **show auth-mac-address**

The following table describes the information displayed by the **show auth-mac-address** command.

TABLE 159 Output from the **show authenticated-mac-address** command

This field...	Displays...
Port	The port number where the multi-device port authentication feature is enabled.
Vlan	The VLAN to which the port has been assigned.

TABLE 159 Output from the **show authenticated-mac-address** command (Continued)

This field...	Displays...
Accepted MACs	The number of MAC addresses that have been successfully authenticated
Rejected MACs	The number of MAC addresses for which authentication has failed.
Attempted-MACs	The rate at which authentication attempts are made for MAC addresses.

Displaying multi-device port authentication configuration information

To display information about the multi-device port authentication configuration, enter the following command.

```
PowerConnect# show auth-mac-address configuration
```

```
Feature enabled           : Yes
Number of Ports enabled  : 4
```

```
-----
Port  Fail-Action      Fail-vlan  Dyn-vlan  MAC-filter
-----
18  Block Traffic    1          No        No
20  Block Traffic    1          No        No
22  Block Traffic    1          No        Yes
5   Block Traffic    1          No        No
```

Syntax: show auth-mac-address configuration

The following table describes the output from the **show auth-mac-address configuration** command.

TABLE 160 Output from the **show authenticated-mac-address** configuration command

This field...	Displays...
Feature enabled	Whether multi-device port authentication is enabled on the device.
Number of Ports enabled	The number of ports on which the multi-device port authentication feature is enabled.
Port	Information for each multi-device port authentication-enabled port.
Fail-Action	What happens to traffic from a MAC address for which RADIUS authentication has failed either block the traffic or assign the MAC address to a restricted VLAN.
Fail-vlan	The restricted VLAN to which non-authenticated MAC addresses are assigned, if the Fail-Action is to assign the MAC address to a restricted VLAN.
Dyn-vlan	Whether RADIUS dynamic VLAN assignment is enabled for the port.
MAC-filter	Whether a MAC filter has been applied to specify pre-authenticated MAC addresses.

Displaying multi-device port authentication information for a specific MAC address or port

To display authentication information for a specific MAC address or port, enter a command such as the following.

```
PowerConnect# show auth-mac-address 0007.e90f.eaa1
```

```
-----
MAC/IP Address                Port          Vlan  Authenticated  Time      Age  CAM
-----
0007.e90f.eaa1 : 25.25.25.25    18          100   Yes           00d01h10m06s 0   N/A
-----
```

Syntax: `show auth-mac-address <mac-address> | <ip-addr> | <portnum>`

The `<ip-addr>` parameter lists the MAC address associated with the specified IP address.

The `<portnum>` parameter is a valid port number.

The following table describes the information displayed by the `show authenticated-mac-address` command for a specified MAC address or port.

TABLE 161 Output from the `show authenticated-mac-address <address>` command

This field...	Displays...
MAC/IP Address	The MAC address for which information is displayed. If the packet for which multi-device port authentication was performed also contained an IP address, then the IP address is displayed as well.
Port	The port on which the MAC address was learned.
Vlan	The VLAN to which the MAC address was assigned.
Authenticated	Whether the MAC address was authenticated.
Time	The time at which the MAC address was authenticated. If the clock is set on the device, then the actual date and time are displayed. If the clock has not been set, then the time is displayed relative to when the device was last restarted.
Age	The age of the MAC address entry in the authenticated MAC address list.
CAM Index	If the MAC address is blocked, this is the index entry for the Layer 2 CAM entry created for this MAC address. If the MAC address is not blocked, either through successful authentication or through being placed in the restricted VLAN, then "N/A" is displayed. If the hardware aging period has expired, then "ffff" is displayed for the MAC address during the software aging period.

Displaying the authenticated MAC addresses

To display the MAC addresses that have been successfully authenticated, enter the `show auth-mac-addresses authorized-mac` command.

Syntax: `show auth-mac-addresses authorized-mac`

Displaying the non-authenticated MAC addresses

To display the MAC addresses for which authentication was not successful, enter the following command

```
PowerConnect# show auth-mac-addresses unauthorized-mac
```

MAC Address	Port	Vlan	Authenticated	Time	Age	dot1x
000f.ed00.0321	1	87	No	00d01h03m17s	H44	Ena
000f.ed00.0259	1	87	No	00d01h03m17s	H44	Ena
000f.ed00.0385	1	87	No	00d01h03m17s	H44	Ena
000f.ed00.02bd	1	87	No	00d01h03m17s	H44	Ena
000f.ed00.00c9	1	87	No	00d01h03m17s	H44	Ena

Syntax: show auth-mac-addresses unauthorized-mac

[Table 162](#) explains the information in the output.

Displaying multi-device port authentication information for a port

To display a summary of Multi-Device Port Authentication for ports on a device, enter the following command

```
PowerConnect# show auth-mac-addresses ethernet 1
```

MAC Address	Port	Vlan	Authenticated	Time	Age	Dot1x
000f.ed00.0001	1	87	Yes	00d01h03m17s	Ena	Ena
000f.ed00.012d	1	87	Yes	00d01h03m17s	Ena	Ena
000f.ed00.0321	1	87	No	00d01h03m17s	H52	Ena
000f.ed00.0259	1	87	No	00d01h03m17s	H52	Ena
000f.ed00.0065	1	87	Yes	00d01h03m17s	Ena	Ena
000f.ed00.0385	1	87	No	00d01h03m17s	H52	Ena
000f.ed00.0191	1	87	Yes	00d01h03m17s	Ena	Ena
000f.ed00.02bd	1	87	No	00d01h03m17s	H52	Ena
000f.ed00.00c9	1	87	No	00d01h03m17s	H52	Ena
000f.ed00.01f5	1	87	Yes	00d01h03m17s	Ena	Ena

Syntax: show auth-mac-address ethernet <portnum>

[Table 162](#) explains the information in the output.

TABLE 162 Output of show auth-mac-address

This field...	Displays...
MAC Address	The MAC addresses learned on the port. If the packet for which multi-device port authentication was performed also contained an IP address, the IP address is also displayed.
Port	ID of the port on which the MAC address was learned.
VLAN	VLAN of which the port is a member.
Authenticated	Whether the MAC address has been authenticated by the RADIUS server.
Time	The time the MAC address was authenticated. If the clock is set on the device, then the actual date and time are displayed. If the clock has not been set, the time is displayed relative to when the device was last restarted.
Age	The age of the MAC address entry in the authenticated MAC address list.
Dot1x	Indicates if 802.1X authentication is enabled or disabled for the MAC address

Displaying multi-device port authentication settings and authenticated MAC addresses

To display the multi-device port authentication settings and authenticated MAC addresses for a port where the feature is enabled, enter the following command.

Syntax: `show auth-mac-address [detail] [ethernet <portnum>]`

The `<portnum>` parameter is a valid port number.

Omitting the `<portnum>` parameter displays information for all interfaces where the multi-device port authentication feature is enabled.

```
PowerConnect# show auth-mac-addresses detailed ethernet 23
Port : 23
Dynamic-Vlan Assignment : Enabled
RADIUS failure action : Block Traffic
  Failure restrict use dot1x : No
Override-restrict-vlan : Yes
Port Default VLAN : 101 ( RADIUS assigned: No) (101)
Port Vlan State : DEFAULT
802.1x override Dynamic PVID : YES
  override return to PVID : 101
Original PVID : 101
DOS attack protection : Disabled
Accepted Mac Addresses : 1
Rejected Mac Addresses : 0
Authentication in progress : 0
Authentication attempts : 0
RADIUS timeouts : 0
RADIUS timeouts action : Success
MAC Address on PVID : 1
MAC Address authorized on PVID : 1
Aging of MAC-sessions : Enabled
Port move-back vlan : Port-configured-vlan
Max-Age of sw mac session : 120 seconds
hw age for denied mac : 70 seconds
MAC Filter applied : No
Dynamic ACL applied : No
num Dynamic Tagged Vlan : 2
Dynamic Tagged Vlan list : 1025 (1/1) 4060 (1/0)
```

```
-----
MAC Address      RADIUS Server  Authenticated  Time           Age  Dot1x
-----
0030.4874.3181  64.12.12.5    Yes           00d01h03m17s  Ena  Ena
-----
```

The following table describes the information displayed by the `show auth-mac-addresses detailed` command.

TABLE 163 Output from the `show auth-mac-addresses detailed` command

This field...	Displays...
Port	The port to which this information applies.
Dynamic-Vlan Assignment	Whether RADIUS dynamic VLAN assignment has been enabled for the port.

TABLE 163 Output from the **show auth-mac-addresses** detailed **command** (Continued)

This field...	Displays...
RADIUS failure action	What happens to traffic from a MAC address for which RADIUS authentication has failed either block the traffic or assign the MAC address to a restricted VLAN.
Failure restrict use dot1x	Indicates if 802.1x traffic that failed multi-device port authentication, but succeeded 802.1x authentication to gain access to the network.
Override-restrict-vlan	Whether a port can be dynamically assigned to a VLAN specified by a RADIUS server, if the port had been previously placed in the restricted VLAN because a previous attempt at authenticating a MAC address on that port failed.
Port Default Vlan	The VLAN to which the port is assigned, and whether the port had been dynamically assigned to the VLAN by a RADIUS server.
Port VLAN state	Indicates the state of the port VLAN. The State can be one of the following "Default", "RADIUS Assigned" or "Restricted".
802.1X override Dynamic PVID	Indicates if 802.1X can dynamically assign a Port VLAN ID (PVID).
override return to PVID	If a port PVID is assigned through the multi-device port authentication feature, and 802.1X authentication subsequently specifies a different PVID, then the PVID specified through 802.1X authentication overrides the PVID specified through multi-device port authentication. This line indicates the PVID the port will use if 802.1X dynamically assigns PVID.
Original PVID	The originally configured (not dynamically assigned) PVID for the port.
DOS attack protection	Whether denial of service attack protection has been enabled for multi-device port authentication, limiting the rate of authentication attempts sent to the RADIUS server.
Accepted Mac Addresses	The number of MAC addresses that have been successfully authenticated.
Rejected Mac Addresses	The number of MAC addresses for which authentication has failed.
Authentication in progress	The number of MAC addresses for which authentication is pending. This is the number of MAC addresses for which an Access-Request message has been sent to the RADIUS server, and for which the RADIUS server has not yet sent an Access-Accept message.
Authentication attempts	The total number of authentication attempts made for MAC addresses on an interface, including pending authentication attempts.
RADIUS timeouts	The number of times the session between the device and the RADIUS server timed out.
RADIUS timeout action	Action to be taken by the RADIUS server if it times out.
MAC address on the PVID	Number of MAC addresses on the PVID.
MAC address authorized on PVID	Number of authorized MAC addresses on the PVID.
Aging of MAC-sessions	Whether software aging of MAC addresses is enabled.
Port move-back VLAN	Indicates the destination VLAN when a RADIUS assigned VLAN is removed. By default, it would return the configured VLAN.
Max-Age of sw MAC-sessions	The configured software aging period for MAC addresses.
hw age for denied MAC	The hardware aging period for blocked MAC addresses. The MAC addresses are dropped in hardware ones the aging period expires.
MAC Filter applied	Indicates whether a MAC filter has been applied to this port to specify pre-authenticated MAC addresses.

TABLE 163 Output from the **show auth-mac-addresses** detailed **command** (Continued)

This field...	Displays...
Dynamic ACL applied	Indicates whether a dynamic ACL was applied to this port.
num Dynamic Tagged Vlan	The number of dynamically tagged VLANs on this port.
Dynamic Tagged Vlan list	The list of dynamically tagged VLANs on this port. In this example, 1025 (1/1) indicates that there was one MAC session and one learned MAC address for VLAN 1025. Likewise, 4060 (1/0) indicates that there was one MAC session and no learned MAC addresses for VLAN 4060.
MAC Address	The MAC addresses learned on the port. If the packet for which multi-device port authentication was performed also contained an IP address, then the IP address is displayed as well.
RADIUS Server	The IP address of the RADIUS server used for authenticating the MAC addresses.
Authenticated	Whether the MAC address has been authenticated by the RADIUS server.
Time	The time at which the MAC address was authenticated. If the clock is set on the device, then the actual date and time are displayed. If the clock has not been set, then the time is displayed relative to when the device was last restarted.
Age	The age of the MAC address entry in the authenticated MAC address list.
Dot1x	Indicated if 802.1X authentication is enabled or disabled for the MAC address

30 Displaying multi-device port authentication information

Protecting Against Denial of Service Attacks

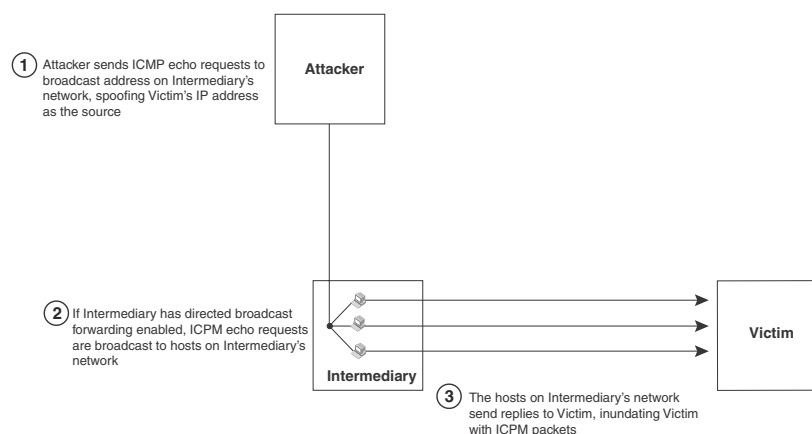
Protecting against Smurf attacks

This chapter explains how to protect your devices from Denial of Service (DoS) attacks.

In a Denial of Service (DoS) attack, a router is flooded with useless packets, hindering normal operation. Devices include measures for defending against two types of DoS attacks Smurf attacks and TCP SYN attacks.

A **Smurf attack** is a kind of DoS attack where an attacker causes a victim to be flooded with ICMP echo (Ping) replies sent from another network. [Figure 128](#) illustrates how a Smurf attack works.

FIGURE 128 How a Smurf attack floods a victim with ICMP replies



The attacker sends an ICMP echo request packet to the broadcast address of an intermediary network. The ICMP echo request packet contains the spoofed address of a victim network as its source. When the ICMP echo request reaches the intermediary network, it is converted to a Layer 2 broadcast and sent to the hosts on the intermediary network. The hosts on the intermediary network then send ICMP replies to the victim network.

For each ICMP echo request packet sent by the attacker, a number of ICMP replies equal to the number of hosts on the intermediary network are sent to the victim. If the attacker generates a large volume of ICMP echo request packets, and the intermediary network contains a large number of hosts, the victim can be overwhelmed with ICMP replies.

Avoiding being an intermediary in a Smurf attack

A Smurf attack relies on the intermediary to broadcast ICMP echo request packets to hosts on a target subnet. When the ICMP echo request packet arrives at the target subnet, it is converted to a Layer 2 broadcast and sent to the connected hosts. This conversion takes place only when directed broadcast forwarding is enabled on the device.

To avoid being an intermediary in a Smurf attack, make sure forwarding of directed broadcasts is disabled on the device. Directed broadcast forwarding is disabled by default. To disable directed broadcast forwarding, do one of the following.

```
PowerConnect(config)# no ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

Avoiding being a victim in a Smurf attack

You can configure the device to drop ICMP packets when excessive numbers are encountered, as is the case when the device is the victim of a Smurf attack. You can set threshold values for ICMP packets that are targeted at the router itself or passing through an interface, and drop them when the thresholds are exceeded.

Protection against ICMP attacks in PowerConnect devices

The ICMP flood attack protection is implemented in hardware on PowerConnect B-Series TI24X devices. This feature can coexist with port-based rate-limiting, MAC filters, Layer 4 ACLs, and other features.

You can set threshold values for ICMP packets that are targeted at the router itself or passing through an interface, and drop them when the thresholds are exceeded.

The syntax to set threshold values for ICMP packets targeted on a PowerConnect device is as follows.

Syntax: ip icmp attack-rate burst-normal <value> burst-max <value> lockup <seconds>

The **attack-rate** keyword indicates that the normal burst value and maximum burst values to be specified in kilobits per second (kbps).

The **burst-normal** value ranges from 20 through 10000000.

The **burst-max** value ranges from 20 through 10000000.

The **lockup** value ranges from 1 through 10000.

The number of incoming ICMP packets per second are measured and compared to the threshold values as follows:

- If the number of ICMP packets exceeds the **burst-normal** value, the excess ICMP packets are dropped.
- If the number of ICMP packets exceeds the **burst-max** value, all ICMP packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

Configuration notes

Consider the following statements when DoS attack protection is implemented at port level or VLAN level.

- The ACL based ingress rate-limiting for ICMP flow on a port is not accurate if ICMP Dos attack protection is enabled on the same port. Non-ICMP flows are not affected.

- ICMP DoS attack protection considers packet marked as drop by port-based ingress rate limiting. In this case, even if the port-based ingress rate-limiting reduces the packet per byte rate, DoS attack is still detected by using actual ingress packet per byte rate on a port.

NOTE

If you configure both DoS attack protection and ACL or MAC filter, the DoS attack statistics for dropped ICMP or TCP SYN packet increments even if the ACL or MAC filter denies the traffic.

Protecting against TCP SYN attacks

TCP SYN attacks exploit the process of how TCP connections are established in order to disrupt normal traffic flow. When a TCP connection starts, the connecting host first sends a TCP SYN packet to the destination host. The destination host responds with a SYN ACK packet, and the connecting host sends back an ACK packet. This process, known as a “TCP three-way handshake”, establishes the TCP connection.

While waiting for the connecting host to send an ACK packet, the destination host keeps track of the as-yet incomplete TCP connection in a connection queue. When the ACK packet is received, information about the connection is removed from the connection queue. Usually there is not much time between the destination host sending a SYN ACK packet and the source host sending an ACK packet, so the connection queue clears quickly.

In a TCP SYN attack, an attacker floods a host with TCP SYN packets that have random source IP addresses. For each of these TCP SYN packets, the destination host responds with a SYN ACK packet and adds information to the connection queue. However, since the source host does not exist, no ACK packet is sent back to the destination host, and an entry remains in the connection queue until it ages out (after around a minute). If the attacker sends enough TCP SYN packets, the connection queue can fill up, and service can be denied to legitimate TCP connections.

Protection against TCP-SYN attacks in PowerConnect devices

The TCP-SYN flood attack protection is implemented in hardware on PowerConnect B-Series TI24X devices. The protection against TCP SYN flood assume that the TCP SYN packet size is 74 bytes, which includes L2, IPv4, and TCP header. If packet size of the attack exceeds the limit, the TCP attack protection takes effect faster than the configured burst values.

To protect against TCP SYN attacks, you can configure the device to drop TCP SYN packets when excessive numbers are encountered. You can set threshold values for TCP SYN packets that are targeted at the router itself or passing through an interface, and drop them when the thresholds are exceeded.

For example, to set threshold values for TCP SYN packets targeted at the router, enter the following command in CONFIG mode.

```
PowerConnect(config)# ip tcp burst-normal 30 burst-max 100 lockup 300
```

To set threshold values for TCP SYN packets received on interface 11, enter the following command.

```
PowerConnect(config)# int e 11
PowerConnect(config-if-e10000-11)# ip tcp burst-normal 30 burst-max 100 lockup
300
```

Syntax: `ip tcp burst-normal <value> burst-max <value> lockup <seconds>`

NOTE

This command is available at the global CONFIG level on both Chassis devices and Stackable devices. On Chassis devices, this command is available at the Interface level as well. This command is supported on Ethernet and Layer 3 ATM interfaces.

The PowerConnect B-Series TI24X device supports the following burst-normal, burst-max, and lockup values.

The **burst-normal** value ranges from 30 through 16000000.

The **burst-max** value ranges from 30 through 16000000.

The **lockup** value ranges from 1 through 10000.

The number of incoming TCP SYN packets per second are measured and compared to the threshold values as follows:

- If the number of TCP SYN packets exceeds the **burst-normal** value, the excess TCP SYN packets are dropped.
- If the number of TCP SYN packets exceeds the **burst-max** value, all TCP SYN packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

In the example above, if the number of TCP SYN packets received per second exceeds 30, the excess packets are dropped. If the number of TCP SYN packets received per second exceeds 100, the device drops all TCP SYN packets for the next 300 seconds (five minutes).

Configuration notes

Consider the following statements when DoS attack protection is implemented at port level or Virtual Interface (VE) level.

- The ACL based ingress rate-limiting for TCP flow on a port is not accurate if TCP DoS attack protection is enabled on the same port. Non-TCP flows are not affected.
- TCP DoS attack protection considers packet marked as drop by port-based ingress rate limiting. In this case, even if the port-based ingress rate-limiting reduces the packet per byte rate, DoS attack is still detected by using actual ingress packet per byte rate on a port.

TCP security enhancement

TCP security enhancement improves upon the handling of TCP inbound segments. This enhancement eliminates or minimizes the possibility of a TCP reset attack, in which a perpetrator attempts to prematurely terminate an active TCP session, and a data injection attack, wherein an attacker injects or manipulates data in a TCP connection.

In both cases, the attack is blind, meaning the perpetrator does not have visibility into the content of the data stream between two devices, but blindly injects traffic. Also, the attacker does not see the direct effect, the continuing communications between the devices and the impact of the injected packet, but may see the indirect impact of a terminated or corrupted session.

The TCP security enhancement prevents and protects against the following three types of attacks:

- Blind TCP reset attack using the reset (RST) bit.
- Blind TCP reset attack using the synchronization (SYN) bit

- Blind TCP packet injection attack

The TCP security enhancement is automatically enabled.

Protecting against a blind TCP reset attack using the RST bit

In a blind TCP reset attack using the RST bit, a perpetrator attempts to guess the RST segments in order to prematurely terminate an active TCP session.

To prevent a user from using the RST bit to reset a TCP connection, the RST bit is subject to the following rules when receiving TCP segments:

- If the RST bit is set and the sequence number is outside the expected window, the device silently drops the segment.
- If the RST bit is exactly the next expected sequence number, the device resets the connection.
- If the RST bit is set and the sequence number does not exactly match the next expected sequence value, but is within the acceptable window, the device sends an acknowledgement.

Protecting against a blind TCP reset attack using the SYN bit

In a blind TCP reset attack, a perpetrator attempts to guess the SYN bits to prematurely terminate an active TCP session.

To prevent a user from using the SYN bit to tear down a TCP connection, the SYN bit is subject to the following rules when receiving TCP segments:

- If the SYN bit is set and the sequence number is outside the expected window, the device sends an acknowledgement (ACK) back to the peer.
- If the SYN bit is set and the sequence number is an exact match to the next expected sequence, the device sends an ACK segment to the peer. Before sending the ACK segment, the software subtracts one from the value being acknowledged.
- If the SYN bit is set and the sequence number is acceptable, the device sends an acknowledgement (ACK) segment to the peer.

Protecting against a blind injection attack

In a blind TCP injection attack, a perpetrator tries to inject or manipulate data in a TCP connection.

To reduce the chances of a blind injection attack, an additional check on all incoming TCP segments is performed.

Displaying statistics about packets dropped because of DoS attacks

To display information about ICMP and TCP SYN packets dropped because burst thresholds were exceeded, enter the following command.

31 Protecting against TCP SYN attacks

```
PowerConnect# show statistics dos-attack
----- Local Attack Statistics -----
ICMP Drop Count      ICMP Block Count      SYN Drop Count      SYN Block Count
-----
0                    0                    0                    0
-----
----- Transit Attack Statistics -----
Port  ICMP Drop Count      ICMP Block Count      SYN Drop Count      SYN Block Count
-----
11    0                    0                    0                    0
```

Syntax: show statistics dos-attack

To clear statistics about ICMP and TCP SYN packets dropped because burst thresholds were exceeded, enter the following command.

```
PowerConnect# clear statistics dos-attack
```

Syntax: clear statistics dos-attack

Displaying statistics about packets dropped because of DoS attacks in PowerConnect devices

To display information about ICMP and TCP SYN packets dropped because burst thresholds were exceeded, enter the following command.

```
PowerConnect# show statistics dos-attack
----- Local Attack Statistics -----
          ICMP                      TCP-SYN
-----
Dropped pkts Blocked pkts Lockup Count  Dropped pkts Blocked pkts Lockup Count
-----
0            0            0            0            0            0
-----
----- Transit Attack Statistics -----
          ICMP                      TCP-SYN
-----
Port/VE Dropped pkts Blocked pkts Lockup Count  Dropped pkts Blocked pkts Lockup Count
-----
25      76            1754            2            0            0            0
```

Syntax: show statistics dos-attack

[Table 164](#) defines the fields in the output command.

TABLE 164 Output of show statistics dos-attack command

Table 0.10:

This field..	Displays
Dropped pkts	Displays the number of packets dropped due to exceeding the configured burst-normal rate.
Blocked pkt	Displays the number of packets dropped due to exceeding the configured burst-max rate.
Lockup Count	Displays the number of times a port or VE is locked down for ICMP or TCP-SYN packets.

Securing SNMP Access

SNMP overview

SNMP is a set of protocols for managing complex networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

Chapter 26, “Securing Access to Management Functions” introduced a few methods used to secure SNMP access. They included the following:

- “Using ACLs to restrict SNMP access” on page 858
- “Restricting SNMP access to a specific IP address” on page 860
- “Restricting SNMP access to a specific VLAN” on page 862
- “Disabling SNMP access” on page 865

This chapter presents additional methods for securing SNMP access to devices. It contains the following sections:

- “Establishing SNMP community strings”
- “Using the user-based security modelSNMP version 3 (RFC 2570 through 2575) introduces a User-Based Security model (RFC 2574) for authentication and privacy services.”
- “SNMP v3 Configuration examples”
- “SNMP version 3 traps”
- “Displaying SNMP Information”
- “SNMP v3 Configuration examples”

Restricting SNMP access using ACL, VLAN, or a specific IP address constitute the first level of defense when the packet arrives at a device. The next level uses one of the following methods:

- Community string match In SNMP versions 1 and 2
- User-based model in SNMP version 3

SNMP views are incorporated in community strings and the user-based model.

Establishing SNMP community strings

SNMP versions 1 and 2 use community strings to restrict SNMP access.

- The default read-only community string is “public”.
- There is no default read-write community string. You first must configure a read-write community string using the CLI. Then you can log on using “set” as the user name and the read-write community string you configure as the password.

You can configure as many additional read-only and read-write community strings as you need. The number of strings you can configure depends on the memory on the device. There is no practical limit.

NOTE

If you delete the startup-config file, the device automatically re-adds the default “public” read-only community string the next time you load the software.

Encryption of SNMP community strings

The software automatically encrypts SNMP community strings. Users with read-only access or who do not have access to management functions in the CLI cannot display the strings. For users with read-write access, the strings are encrypted in the CLI.

Encryption is enabled by default. You can disable encryption for individual strings or trap receivers if desired. Refer to the next section for information about encryption.

Adding an SNMP community string

When you add a community string, you can specify whether the string is encrypted or clear. By default, the string is encrypted.

To add an encrypted community string, enter commands such as the following.

```
PowerConnect(config)# snmp-server community private rw
PowerConnect(config)# write memory
```

Syntax: `snmp-server community [0 | 1] <string>`
`ro | rw [view <viewname>] [<standard-ACL-name> | <standard-ACL-id>]`

The `<string>` parameter specifies the community string name. The string can be up to 32 characters long.

The `ro` | `rw` parameter specifies whether the string is **read-only (ro)** or **read-write (rw)**.

The `0` | `1` parameter affects encryption for display of the string in the running-config and the startup-config file. Encryption is enabled by default. When encryption is enabled, the community string is encrypted in the CLI regardless of the access level you are using.

The encryption option can be omitted (the default) or can be one of the following:

- **0** – Disables encryption for the community string you specify with the command. The community string is shown as clear text in the running-config and the startup-config file. Use this option if you do not want the display of the community string to be encrypted.
- **1** – Assumes that the community string you enter is encrypted, and decrypts the value before using it.

NOTE

If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter **0** or **1**. Instead, omit the encryption option and allow the software to use the default behavior.

NOTE

If you specify encryption option **1**, the software assumes that you are entering the encrypted form of the community string. In this case, the software decrypts the community string you enter before using the value for authentication. If you accidentally enter option **1** followed by the clear-text version of the community string, authentication will fail because the value used by the software will not match the value you intended to use.

The command in the example above adds the read-write SNMP community string “private”. When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server community 1 <encrypted-string> rw
```

To add a non-encrypted community string, you must explicitly specify that you do not want the software to encrypt the string. Here is an example.

```
PowerConnect(config)# snmp-server community 0 private rw
PowerConnect(config)# write memory
```

The command in this example adds the string “private” in the clear, which means the string is displayed in the clear. When you save the new community string to the startup-config file, the software adds the following command to the file.

```
snmp-server community 0 private rw
```

The **view** <viewname> parameter is optional. It allows you to associate a view to the members of this community string. Enter up to 32 alphanumeric characters. If no view is specified, access to the full MIB is granted. The view that you want must exist before you can associate it to a community string. Here is an example of how to use the view parameter in the community string command.

```
PowerConnect(config)# snmp-s community myread ro view sysview
```

The command in this example associates the view “sysview” to the community string named “myread”. The community string has read-only access to “sysview”. For information on how to create views, refer to [“SNMP v3 Configuration examples”](#) on page 1023.

The <standard-ACL-name> | <standard-ACL-id> parameter is optional. It allows you to specify which ACL group will be used to filter incoming SNMP packets. You can enter either the ACL name or its ID. Here are some examples.

```
PowerConnect(config)# snmp-s community myread ro view sysview 2
PowerConnect(config)# snmp-s community myread ro view sysview myACL
```

The command in the first example indicates that ACL group 2 will filter incoming SNMP packets; whereas, the command in the second example uses the ACL group called “myACL” to filter incoming packets. Refer to [“Using ACLs to restrict SNMP access”](#) on page 858 for more information.

NOTE

To make configuration changes, including changes involving SNMP community strings, you must first configure a read-write community string using the CLI. Alternatively, you must configure another authentication method and log on to the CLI using a valid password for that method.

Displaying the SNMP community strings

To display the configured community strings, enter the following command at any CLI level.

```
PowerConnect# show snmp server
Contact: Marshall
Location: Copy Center
Community(ro): public
Community(rw): private
Traps
    Cold start: Enable
    Link up: Enable
    Link down: Enable
    Authentication: Enable
    Locked address violation: Enable
    Power supply failure: Enable
    Fan failure: Enable
    Temperature warning: Enable
    STP new root: Enable
    STP topology change: Enable
    ospf: Enable

Total Trap-Receiver Entries: 4
Trap-Receiver IP Address      Community
1          207.95.6.211
2          207.95.5.21
```

Syntax: show snmp server

NOTE

If display of the strings is encrypted, the strings are not displayed. Encryption is enabled by default.

Using the user-based security model SNMP version 3 (RFC 2570 through 2575) introduces a User-Based Security model (RFC 2574) for authentication and privacy services.

SNMP version 1 and version 2 use community strings to authenticate SNMP access to management modules. This method can still be used for authentication. In SNMP version 3, the User-Based Security model of SNMP can be used to secure against the following threats:

- Modification of information
- Masquerading the identity of an authorized entity
- Message stream modification
- Disclosure of information

SNMP version 3 also supports View-Based Access Control Mechanism (RFC 2575) to control access at the PDU level. It defines mechanisms for determining whether or not access to a managed object in a local MIB by a remote principal should be allowed. (refer to [“SNMP v3 Configuration examples”](#) on page 1023.)

Configuring your NMS

In order to use the SNMP version 3 features.

1. Make sure that your Network Manager System (NMS) supports SNMP version 3.
2. Configure your NMS agent with the necessary users.
3. Configure the SNMP version 3 features in devices.

Configuring SNMP version 3

Follow the steps given below to configure SNMP version 3 on devices.

1. Enter an engine ID for the management module using the **snmp-server engineid** command if you will not use the default engine ID. Refer to “[Defining the engine id](#)” on page 1015.
2. Create views that will be assigned to SNMP user groups using the **snmp-server view** command. Refer to “[SNMP v3 Configuration examples](#)” on page 1023 for details.
3. Create ACL groups that will be assigned to SNMP user groups using the **access-list** command.
4. Create user groups using the **snmp-server group** command. Refer to “[Defining an SNMP group](#)” on page 1016.
5. Create user accounts and associate these accounts to user groups using the **snmp-server user** command. Refer to “[Defining an SNMP user account](#)” on page 1017.

If SNMP version 3 is not configured, then community strings by default are used to authenticate access.

Defining the engine id

A default engine ID is generated during system start up. To determine what the default engine ID of the device is, enter the **show snmp engineid** command and find the following line:

```
Local SNMP Engine ID: 800007c70300e05290ab60
```

See the section “[Displaying the Engine ID](#)” on page 1022 for details.

The default engine ID guarantees the uniqueness of the engine ID for SNMP version 3. If you want to change the default engine ID, enter a command such as the following.

```
PowerConnect(config)# snmp-server engineid local 800007c70300e05290ab60
```

Syntax: [no] **snmp-server engineid local** <hex-string>

The **local** parameter indicates that engine ID to be entered is the ID of this device, representing an SNMP management entity.

NOTE

Each user localized key depends on the SNMP server engine ID, so all users need to be reconfigured whenever the SNMP server engine ID changes.

NOTE

Since the current implementation of SNMP version 3 does not support Notification, remote engine IDs cannot be configured at this time.

The <hex-string> variable consists of 11 octets, entered as hexadecimal values. There are two hexadecimal characters in each octet. There should be an even number of hexadecimal characters in an engine ID.

The default engine ID has a maximum of 11 octets:

- Octets 1 through 4 represent the agent's SNMP management private enterprise number as assigned by the Internet Assigned Numbers Authority (IANA). The most significant bit of Octet 1 is "1". For example, "000007c7" is the ID for Dell Communications, Inc. in hexadecimal. With Octet 1 always equal to "1", the first four octets in the default engine ID is always "800007c7" (which is 1991 in decimal).

- Octet 5 is always 03 in hexadecimal and indicates that the next set of values represent a MAC address.
- Octets 6 through 11 form the MAC address of the lowest port in the management module.

NOTE

Engine ID must be a unique number among the various SNMP engines in the management domain. Using the default engine ID ensures the uniqueness of the numbers.

Defining an SNMP group

SNMP groups map SNMP users to SNMP views. For each SNMP group, you can configure a read view, a write view, or both. Users who are mapped to a group will use its views for access control.

To configure an SNMP user group, enter a command such as the following.

```
PowerConnect(config)# snmp-server group admin v3 auth read all write all
```

Syntax: `[no] snmp-server group <groupname> v1 | v2 | v3 auth | noauth | priv [access <standard-ACL-id>] [read <viewstring> | write <viewstring>]`

NOTE

This command is not used for SNMP version 1 and SNMP version 2. In these versions, groups and group views are created internally using community strings. (refer to [“Establishing SNMP community strings”](#) on page 1011.) When a community string is created, two groups are created, based on the community string name. One group is for SNMP version 1 packets, while the other is for SNMP version 2 packets.

The **group** `<groupname>` parameter defines the name of the SNMP group to be created.

The **v1**, **v2**, or **v3** parameter indicates which version of SNMP is used. In most cases, you will be using v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If **auth** is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The **access** `<standard-ACL-id>` parameter is optional. It allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **read** `<viewstring>` | **write** `<viewstring>` parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The `<viewstring>` variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

The value of `<viewstring>` is defined using the **snmp-server view** command. The SNMP agent comes with the "all" default view, which provides access to the entire MIB; however, it must be specified when creating the group. The "all" view also allows SNMP version 3 to be backwards compatible with SNMP version 1 and version 2.

NOTE

If you will be using a view other than the "all" view, that view must be configured before creating the user group. Refer to the section “[SNMP v3 Configuration examples](#)” on page 1023, especially for details on the include | exclude parameters.

Defining an SNMP user account

The **snmp-server user** command does the following:

- Creates an SNMP user.
- Defines the group to which the user will be associated.
- Defines the type of authentication to be used for SNMP access by this user.
- Specifies one of the following encryption types used to encrypt the privacy password:
 - Data Encryption Standard (DES) – A symmetric-key algorithm that uses a 56-bit key.
 - Advanced Encryption Standard (AES) – The 128-bit encryption standard adopted by the U.S. government. This standard is a symmetric cipher algorithm chosen by the National Institute of Standards and Technology (NIST) as the replacement for DES.

Here is an example of how to create an SNMP User account.

```
PowerConnect(config)# snmp-s user bob admin v3 access 2 auth md5 bobmd5 priv des bobdes
```

The CLI for creating SNMP version 3 users has been updated as follows.

Syntax: [no] **snmp-server user** <name> <groupname> **v3**
 [[**access** <standard-ACL-id>]
 [[**encrypted**] [**auth md5** <md5-password> | **sha** <sha-password>]
 [**priv** [**encrypted**] **des** <des-password-key> | **aes** <aes-password-key>]]]

The <name> parameter defines the SNMP user name or security name used to access the management module.

The <groupname> parameter identifies the SNMP group to which this user is associated or mapped. All users must be mapped to an SNMP group. Groups are defined using the **snmp-server group** command.

NOTE

The SNMP group to which the user account will be mapped should be configured before creating the user accounts; otherwise, the group will be created without any views. Also, ACL groups must be configured before configuring user accounts.

The **v3** parameter is required.

The **access** <standard-ACL-id> parameter is optional. It indicates that incoming SNMP packets are filtered based on the ACL attached to the user account.

NOTE

The ACL specified in a user account overrides the ACL assigned to the group to which the user is mapped. If no ACL is entered for the user account, then the ACL configured for the group will be used to filter packets.

The **encrypted** parameter means that the MD5 or SHA password will be a digest value. MD5 has 16 octets in the digest. SHA has 20. The digest string has to be entered as a hexadecimal string. In this case, the agent need not generate any explicit digest. If the **encrypted** parameter is not used, the user is expected to enter the authentication password string for MD5 or SHA. The agent will convert the password string to a digest, as described in RFC 2574.

The **auth md5 | sha** parameter is optional. It defines the type of encryption that the user must have to be authenticated. Choose between MD5 or SHA encryption. MD5 and SHA are two authentication protocols used in SNMP version 3.

The `<md5-password>` and `<sha-password>` define the password the user must use to be authenticated. These password must have a minimum of 8 characters. If the encrypted parameter is used, then the digest has 16 octets for MD5 or 20 octets for SHA.

NOTE

Once a password string is entered, the generated configuration displays the digest (for security reasons), not the actual password.

The **priv [encrypted]** parameter is optional after you enter the md5 or sha password. The **priv** parameter specifies the encryption type (DES or AES) used to encrypt the privacy password. If the **encrypted** keyword is used, do the following:

- If DES is the privacy protocol to be used, enter **des** followed by a 16-octet DES key in hexadecimal format for the `<des-password-key>`. If you include the encrypted keyword, enter a password string of at least 8 characters.
- If AES is the privacy protocol to be used, enter **aes** followed by the AES password key. For a small password key, enter 12 characters. For a big password key, enter 16 characters. If you include the encrypted keyword, enter a password string containing 32 hexadecimal characters.

Defining SNMP views

SNMP views are named groups of MIB objects that can be associated with user accounts to allow limited access for viewing and modification of SNMP statistics and system configuration. SNMP views can also be used with other commands that take SNMP views as an argument. SNMP views reference MIB objects using object names, numbers, wildcards, or a combination of the three. The numbers represent the hierarchical location of the object in the MIB tree. You can reference individual objects in the MIB tree or a subset of objects from the MIB tree.

To configure the number of SNMP views available on the device, enter the following command.

```
PowerConnect(config)# system-max view 15
```

Syntax: `system-max view <number-of-views>`

This command specifies the maximum number of SNMPv2 and v3 views that can be configured on a device. The number of views can be from 10 – 65536. The default is 10 views.

To add an SNMP view, enter one of the following commands.

```
PowerConnect(config)# snmp-server view Maynes system included
PowerConnect(config)# snmp-server view Maynes system.2 excluded
PowerConnect(config)# snmp-server view Maynes 2.3.*.6 included
PowerConnect(config)# write mem
```


NOTE

The **snmp-server view** command supports the MIB objects as defined in RFC 1445.

Syntax: [no] snmp-server view <name> <mib_tree> included | excluded

The <name> parameter can be any alphanumeric name you choose to identify the view. The names cannot contain spaces.

The <mib_tree> parameter is the name of the MIB object or family. MIB objects and MIB sub-trees can be identified by a name or by the numbers called Object Identifiers (OIDs) that represent the position of the object or sub-tree in the MIB hierarchy. You can use a wildcard (*) in the numbers to specify a sub-tree family.

The **included** | **excluded** parameter specifies whether the MIB objects identified by the <mib_family> parameter are included in the view or excluded from the view.

NOTE

All MIB objects are automatically excluded from any view unless they are explicitly included; therefore, when creating views using the **snmp-server view** command, indicate which portion of the MIB you want users to access.

For example, you may want to assign the view called "admin" a community string or user group. The "admin" view will allow access to the Dell MIBs objects that begin with the 1.3.6.1.4.1.1991 object identifier. Enter the following command.

```
PowerConnect(config)# snmp-server view admin 1.3.6.1.4.1.1991 included
```

You can exclude portions of the MIB within an inclusion scope. For example, if you want to exclude the snAgentSys objects, which begin with 1.3.6.1.4.1.1991.1.1.2 object identifier from the admin view, enter a second command such as the following.

```
PowerConnect(config)# snmp-server view admin 1.3.6.1.4.1.1991.1.1.2 excluded
```

NOTE

Note that the exclusion is within the scope of the inclusion.

To delete a view, use the no parameter before the command.

SNMP version 3 traps

Devices support SNMP notifications in SMIV2 format. This allows notifications to be encrypted and sent to the target hosts in a secure manner.

Defining an SNMP group and specifying which view is notified of traps

The SNMP group command allows configuration of a viewname for notification purpose, similar to the read and write view. The default viewname is "all", which allows access to the entire MIB.

To configure an SNMP user group, first configure SNMP v3 views using the **snmp-server view** command. Refer to "[SNMP v3 Configuration examples](#)" on page 1023. Then enter a command such as the following.

```
PowerConnect(config)# snmp-server group admin v3 auth read all write all
notify all
```

Syntax: [no] snmp-server group <groupname>
 v1 | v2 | v3
 auth | noauth | priv
 [access <standard-ACL-id>] [read <viewstring> | write <viewstring> | notify <viewstring>]

The **group** <groupname> parameter defines the name of the SNMP group to be created.

The **v1**, **v2**, or **v3** parameter indicates which version of SNMP to use. In most cases, you will use v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If auth is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The **access** <standard-ACL-id> parameter is optional. It allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **read** <viewstring> | **write** <viewstring> parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The **notify** view allows administrators to restrict the scope of varbind objects that will be part of the notification. All of the varbinds need to be in the included view for the notification to be created.

The <viewstring> variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

Defining the UDP port for SNMP v3 traps The SNMP host command enhancements allow configuration of notifications in SMIv2 format, with or without encryption, in addition to the previously supported SMIv1 trap format.

You can define a port that receives the SNMP v3 traps by entering a command such as the following.

```
PowerConnect(config)# snmp-server host 192.168.4.11 version v3 auth security-name
port 1 [no] snmp-server host <ip-addr> | <ipv6-addr> version [ v1 | v2c
<community-string> | v3 auth | noauth | priv <security-name>] [port
<trap-UDP-port-number>]
```

The <ip-addr> parameter specifies the IP address of the host that will receive the trap.

For **version**, indicate one of the following

For SNMP version 1, enter **v1** and the name of the community string (<community-string>). This string is encrypted within the system.

NOTE

The options "v2c" and "v3" are new in this release. If the configured version is v2c, then the notification is sent out in SMIv2 format, using the community string, but in cleartext mode. To send the SMIv2 notification in SNMPv3 packet format, configure v3 with auth or privacy parameters, or both, by specifying a security name. The actual authorization and privacy values are obtained from the security name.

For SNMP version 2c, enter **v2** and the name of the community string. This string is encrypted within the system.

For SNMP version 3, enter one of the following depending on the authorization required for the host:

- **v3 auth** <security-name>: Allow only authenticated packets.
- **v3 no auth** <security-name>: Allow all packets.
- **v3 priv** <security-name>: A password is required

For **port** <trap-UDP-port-number>, specify the UDP port number on the host that will receive the trap.

Trap MIB changes

To support the SNMP V3 trap feature, the Dell Enterprise Trap MIB was rewritten in SMIv2 format, as follows:

- The MIB name was changed from FOUNDRY-SN-TRAP-MIB to FOUNDRY-SN-NOTIFICATION-MIB
- Individual notifications were changed to NOTIFICATION-TYPE instead of TRAP-TYPE.
- As per the SMIv2 format, each notification has an OID associated with it. The root node of the notification is snTraps (OID enterprise.foundry.0). For example, OID for snTrapRunningConfigChanged is {snTraps.73}. Earlier, each trap had a trap ID associated with it, as per the SMIv1 format.

Backward compatibility with SMIv1 trap format

The device will continue to support creation of traps in SMIv1 format, as before. To allow the device to send notifications in SMIv2 format, configure the device as described above. The default mode is still the original SMIv1 format.

Restricting SNMP Access to an IPv6 node

You can restrict SNMP access so that the device (including Brocade Network Advisor) can only be accessed by the IPv6 host address that you specify. To do so, enter a command such as the following.

```
PowerConnect(config)# snmp-client ipv6 2001:efff:89::23
```

Syntax: **snmp-client ipv6** <ipv6-address>

The <ipv6-address> must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specifying an IPv6 host as an SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter a command such as the following.

```
PowerConnect(config)# snmp-server host ipv6 2001:efff:89::13
```

Syntax: **snmp-server host ipv6** <ipv6-address>

The <ipv6-address> must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Displaying SNMP Information

This section lists the commands for viewing SNMP-related information.

Displaying the Engine ID

To display the engine ID of a management module, enter a command such as the following.

```
PowerConnect# show snmp engineid
Local SNMP Engine ID: 800007c70300e05290ab60
Engine Boots: 3
Engine time: 5
```

Syntax: show snmp engineid

The engine ID identifies the source or destination of the packet.

The engine boots represents the number of times that the SNMP engine reinitialized itself with the same engine ID. If the engineID is modified, the boot count is reset to 0.

The engine time represents the current time with the SNMP agent.

Displaying SNMP groups

To display the definition of an SNMP group, enter a command such as the following.

```
PowerConnect# show snmp group
groupname = exceptifgrp
security model = v3
security level = authNoPriv
ACL id = 2
readview = exceptif
writeview = <none>
```

Syntax: show snmp group

The value for security level can be one of the following.

Table 0.11:

Security level	Authentication
<none>	If the security model shows v1 or v2, then security level is blank. User names are not used to authenticate users; community strings are used instead.
noauthNoPriv	Displays if the security model shows v3 and user authentication is by user name only.
authNoPriv	Displays if the security model shows v3 and user authentication is by user name and the MD5 or SHA algorithm.

Displaying user information

To display the definition of an SNMP user account, enter a command such as the following.

```
PowerConnect# show snmp user
username = bob
ACL id = 2
group = admin
security model = v3
group ACL id = 0
authtype = md5
authkey = 3aca18d90b8d172760e2dd2e8f59b7fe
privtype = des, privkey = 1088359afb3701730173a6332d406eec
engine ID= 800007c70300e052ab0000
```

Syntax: show snmp user

Interpreting varbinds in report packets

If an SNMP version 3 request packet is to be rejected by an SNMP agent, the agent sends a report packet that contains one or more varbinds. The varbinds contain additional information, showing the cause of failures. An SNMP manager application decodes the description from the varbind. The following table presents a list of varbinds supported by the SNMP agent.

Table 0.12:

Varbind object Identifier	Description
1. 3. 6. 1. 6. 3. 11. 2. 1. 3. 0	Unknown packet data unit.
1. 3. 6. 1. 6. 3. 12. 1. 5. 0	The value of the varbind shows the engine ID that needs to be used in the snmp-server engineid command
1. 3. 6. 1. 6. 3. 15. 1. 1. 1. 0	Unsupported security level.
1. 3. 6. 1. 6. 3. 15. 1. 1. 2. 0	Not in time packet.
1. 3. 6. 1. 6. 3. 15. 1. 1. 3. 0	Unknown user name. This varbind may also be generated: <ul style="list-style-type: none"> • If the configured ACL for this user filters out this packet. • If the group associated with the user is unknown.
1. 3. 6. 1. 6. 3. 15. 1. 1. 4. 0	Unknown engine ID. The value of this varbind would be the correct authoritative engineID that should be used.
1. 3. 6. 1. 6. 3. 15. 1. 1. 5. 0	Wrong digest.
1. 3. 6. 1. 6. 3. 15. 1. 1. 6. 0	Decryption error.

SNMP v3 Configuration examples

The following sections present examples of how to configure SNMP v3.

Simple SNMP v3 configuration

```
PowerConnect(config)# snmp-s group admingrp v3 priv read all write all notify all
PowerConnect(config)# snmp-s user adminuser admingrp v3 auth md5 <auth password>
priv <privacy password>
PowerConnect(config)# snmp-s host <dest-ip> version v3 privacy adminuser
```

More detailed SNMP v3 configuration

```
PowerConnect(config)# snmp-server view internet internet included
PowerConnect(config)# snmp-server view system system included
PowerConnect(config)# snmp-server community ..... ro
PowerConnect(config)# snmp-server community ..... rw
PowerConnect(config)# snmp-server contact isc-operations
PowerConnect(config)# snmp-server location sdh-pillbox
PowerConnect(config)# snmp-server host 128.91.255.32 .....
PowerConnect(config)# snmp-server group ops v3 priv read internet write system
PowerConnect(config)# snmp-server group admin v3 priv read internet write internet
PowerConnect(config)# snmp-server group restricted v3 priv read internet
PowerConnect(config)# snmp-server user ops ops v3 encrypted auth md5
ab8e9cd6d46e7a270b8c9549d92a069 priv encrypted des
0e1b153303b6188089411447dbc32de
PowerConnect(config)# snmp-server user admin admin v3 encrypted auth md5
0d8a2123f91bfbd8695fef16a6f4207b priv encrypted des
18e0cf359fce4fcd60df19c2b6515448
PowerConnect(config)# snmp-server user restricted restricted v3 encrypted auth
md5 261fd8f56a3ad51c8bcecle4609f54dc priv encrypted des
d32e66152f89de9b2e0cb17a65595f43
```

Enabling the Foundry Discovery Protocol and Reading Cisco Discovery Protocol Packets

33

Using FDP

The Foundry Discovery Protocol (FDP) enables Dell devices to advertise themselves to other devices on the network. When you enable FDP on a device, the device periodically advertises information including the following:

- Hostname (device ID)
- Product platform and capability
- Software version
- VLAN and Layer 3 protocol address information for the port sending the update. IP, IPX, and AppleTalk Layer 3 information is supported.

A device running FDP sends FDP updates on Layer 2 to MAC address 01-E0-52-CC-CC-CC. Other devices listening on that address receive the updates and can display the information in the updates. Devices can send and receive FDP updates on Ethernet interfaces.

FDP is disabled by default.

NOTE

If FDP is not enabled on a device that receives an FDP update or the device is running a software release that does not support FDP, the update passes through the device at Layer 2.

Configuring FDP

The following sections describe how to enable FDP and how to change the FDP update and hold timers.

Enabling FDP globally

To enable a device to globally send FDP packets, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# fdp run
```

Syntax: [no] fdp run

The feature is disabled by default.

Enabling FDP at the interface level

By default, FDP is enabled at the interface level after FDP is enabled on the device.

When FDP is enabled globally, you can disable and re-enable FDP on individual ports.

Disable FDP by entering commands such as the following:

```
PowerConnect(config)# int e 1
```

```
PowerConnect(config-if-1)# no fdp enable
```

Enable or re-enable FDP by entering commands such as the following:

```
PowerConnect(config-if-1)# fdp enable
```

Syntax: [no] fdp enable

Changing the FDP update timer

By default, a device enabled for FDP sends an FDP update every 60 seconds. You can change the update timer to a value from 5 – 900 seconds.

To change the FDP update timer, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# fdp timer 120
```

Syntax: [no] fdp timer <secs>

The <secs> parameter specifies the number of seconds between updates and can be from 5 – 900 seconds. The default is 60 seconds.

Changing the FDP hold time

By default, a device that receives an FDP update holds the information until one of the following events occurs:

- The device receives a new update.
- 180 seconds have passed since receipt of the last update. This is the hold time.

Once either of these events occurs, the device discards the update.

To change the FDP hold time, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)# fdp holdtime 360
```

Syntax: [no] fdp holdtime <secs>

The <secs> parameter specifies the number of seconds a device that receives an FDP update can hold the update before discarding it. You can specify from 10 – 255 seconds. The default is 180 seconds.

Displaying FDP information

You can display the following FDP information:

- FDP entries for Dell neighbors
- Individual FDP entries
- FDP information for an interface on the device you are managing
- FDP packet statistics

NOTE

If the Dell device has intercepted CDP updates, then the CDP information is also displayed.

Displaying neighbor information

To display a summary list of all the neighbors that have sent FDP updates to this device, enter the following command.

```
PowerConnectA# show fdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a CDP device
```

```

   Device ID           Local Int   Holdtm  Capability Platform      Port ID
   -----
PowerConnectB         Eth 9      178    Router PowerConnect Rou Eth 9
```

Syntax: `show fdp neighbor [ethernet<portnum>] [detail]`

The `ethernet<portnum>` parameter lists the information for updates received on the specified port.

The `detail` parameter lists detailed information for each device.

The `show fdp neighbor` command, without optional parameters, displays the following information.

TABLE 165 Summary FDP and CDP neighbor information

This line...	Displays...
Device ID	The hostname of the neighbor.
Local Int	The interface on which this device received an FDP or CDP update for the neighbor.
Holdtm	The maximum number of seconds this device can keep the information received in the update before discarding it.
Capability	The role the neighbor is capable of playing in the network.
Platform	The product platform of the neighbor.
Port ID	The interface through which the neighbor sent the update.

To display detailed information, enter the following command.

```
PowerConnectA# show fdp neighbor detail
Device ID: FCX648 Switch
           configured as tag-type8100
Entry address(es):
  IP address: 10.20.64.230
Platform: FCX648 Switch, Capabilities: Switch
Interface: ethernet5
Port ID (outgoing port): ethernet4 is TAGGED in following VLAN(s):
  202 203 204 205 1702 1703 1704 1705 1706 1707 1708 1709 1710 1711 1712 1713
  1714 1715 1716 1717 1718 1719 1720 2702 2703 2704 2705 2706 2707 2708 2709
  2710 2711 2712 2713 2714 2715 2716 2717 2718 2719 2720
Holdtime : 124 seconds
Brocade Communications Systems, Inc. Stacking System FCX648, IronWare Version
07.3.00a031T7f1 Compiled on Dec 02 2010 at 07:20:46 labeled as FCXS07300a031
The show fdp neighbor detail command displays the following information.
```

TABLE 166 Detailed FDP and CDP neighbor information

This line...	Displays...
Device ID	The hostname of the neighbor. In addition, this line lists the VLAN memberships and other VLAN information for the neighbor port that sent the update to this device.
Entry address(es)	The Layer 3 protocol addresses configured on the neighbor port that sent the update to this device. If the neighbor is a Layer 2 Switch, this field lists the management IP address.
Platform	The product platform of the neighbor.
Capabilities	The role the neighbor is capable of playing in the network.
Interface	The interface on which this device received an FDP or CDP update for the neighbor.
Port ID	The interface through which the neighbor sent the update.
Holdtime	The maximum number of seconds this device can keep the information received in the update before discarding it.
Version	The software version running on the neighbor.

Displaying FDP entries

To display the detailed neighbor information for a specific device, enter a command such as the following.

```
PowerConnectA# show fdp entry "FESX424 Switch"
Device ID: FESX424 Switch
           configured as tag-type8100
Entry address(es):
  IP address: 10.20.64.244
Platform: FESX424 Switch, Capabilities: Switch
Interface: ethernet25
Port ID (outgoing port): ethernet4 is TAGGED in following VLAN(s):
 302 303 304 305 1802 1803 1804 1805 1806 1807 1808 1809 1810 1811 1812 1813
 1814 1815 1816 1817 1818 1819 1820 2802 2803 2804 2805 2806 2807 2808 2809
 2810 2811 2812 2813 2814 2815 2816 2817 2818 2819 2820
Holdtime : 144 seconds
Brocade Communications Systems, Inc. FESX424-PREM-PoE, IronWare Version
07.2.00T3e1 Compiled on Sep 20 2010 at 23:06:03 labeled as SXS07200
```

Syntax: `show fdp entry * | <device-id>`

The `* | <device-id>` parameter specifies the device ID. If you enter `*`, the detailed updates for all neighbor devices are displayed. If you enter a specific device ID, the update for that device is displayed. For information about the display, refer to [Table 166](#).

Displaying FDP information for an interface

To display FDP information for an interface, enter a command such as the following.

```
PowerConnectA# show fdp interface ethernet 3
FastEthernet 3 is up, line protocol is up
  Encapsulation ethernet
  Sending FDP packets every 5 seconds
  Holdtime is 180 seconds
```

This example shows information for Ethernet port 3. The port sends FDP updates every 5 seconds. Neighbors that receive the updates can hold them for up to 180 seconds before discarding them.

Syntax: `show fdp interface [ethernet <portnum>]`

The `ethernet <portnum>` parameter lists the information only for the specified interface.

Displaying FDP and CDP statistics

To display FDP and CDP packet statistics, enter the following command.

```
PowerConnectA# show fdp traffic
CDP/FDP counters:
  Total packets output: 6, Input: 5
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
  Internal errors: 0
```

Syntax: `show fdp traffic`

Clearing FDP and CDP information

You can clear the following FDP and CDP information:

- Information received in FDP and CDP updates
- FDP and CDP statistics

The same commands clear information for both FDP and CDP.

Clearing FDP and CDP neighbor information

To clear the information received in FDP and CDP updates from neighboring devices, enter the following command.

```
PowerConnect# clear fdp table
```

Syntax: `clear fdp table`

NOTE

This command clears all the updates for FDP and CDP.

Clearing FDP and CDP statistics

To clear FDP and CDP statistics, enter the following command.

```
PowerConnect# clear fdp counters
```

Syntax: `clear fdp counters`

Reading CDP packets

Cisco Discovery Protocol (CDP) packets are used by Cisco devices to advertise themselves to other Cisco devices. By default, devices forward these packets without examining their contents. You can configure a device to intercept and display the contents of CDP packets. This feature is useful for learning device and interface information for Cisco devices in the network.

Devices support intercepting and interpreting CDP version 1 and version 2 packets.

NOTE

The Dell device can interpret only the information fields that are common to both CDP version 1 and CDP version 2.

NOTE

When you enable interception of CDP packets, the Dell device drops the packets. As a result, Cisco devices will no longer receive the packets.

Enabling interception of CDP packets globally

To enable the device to intercept and display CDP packets, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)# cdp run
```

Syntax: [no] cdp run

The feature is disabled by default.

Enabling interception of CDP packets on an interface

You can disable and enable CDP at the interface level.

You can enter commands such as the following.

```
PowerConnect(config)# int e 1  
PowerConnect(config-if-1)# cdp enable
```

Syntax: [no] cdp enable

By default, the feature is enabled on an interface once CDP is enabled on the device.

Displaying CDP information

You can display the following CDP information:

- Cisco neighbors
- CDP entries for all Cisco neighbors or a specific neighbor
- CDP packet statistics

Displaying neighbors

To display the Cisco neighbors the Dell device has learned from CDP packets, enter the following command.

```
PowerConnect# show fdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a Cisco device
```

Device ID	Local Int	Holdtm	Capability	Platform	Port ID
(*)Router	Eth 1	124	R	cisco RSP4	

FastEthernet0

To display detailed information for the neighbors, enter the following command.

```
PowerConnect# show fdp neighbors detail
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1, Port ID (outgoing port): FastEthernet0
Holdtime : 150 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

To display information about a neighbor attached to a specific port, enter a command such as the following.

```
PowerConnect# show fdp neighbors ethernet 1
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1, Port ID (outgoing port): FastEthernet0
Holdtime : 127 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Syntax: show fdp neighbors [detail | ethernet <portnum>]

Displaying CDP entries

To display CDP entries for all neighbors, enter the following command.

```
PowerConnect# show fdp entry *
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1, Port ID (outgoing port): FastEthernet0
Holdtime : 124 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fcl)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

To display CDP entries for a specific device, specify the device ID. Here is an example.

```
PowerConnect# show fdp entry Router1
Device ID: Router1
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1, Port ID (outgoing port): FastEthernet0
Holdtime : 156 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fcl)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Syntax: `show fdp entry * | <device-id>`

Displaying CDP statistics

To display CDP packet statistics, enter the following command.

```
PowerConnect# show fdp traffic
CDP counters:
  Total packets output: 0, Input: 3
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
```

Syntax: `show fdp traffic`

Clearing CDP information

You can clear the following CDP information:

- Cisco Neighbor information
- CDP statistics

To clear the Cisco neighbor information, enter the following command.

```
PowerConnect# clear fdp table
```

Syntax: `clear fdp table`

To clear CDP statistics, enter the following command.

```
PowerConnect# clear fdp counters
```

Syntax: clear fdp counters

33 Reading CDP packets

Using Syslog

This chapter describes how to display Syslog messages and how to configure the Syslog facility, and lists the Syslog messages that devices can display during standard operation.

NOTE

This chapter does not list Syslog messages that can be displayed when a debug option is enabled.

Overview

Device software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer.

You also can specify the IP address or host name of up to six Syslog servers. When you specify a Syslog server, the device writes the messages both to the system log and to the Syslog server.

Using a Syslog server ensures that the messages remain available even after a system reload. The local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the Syslog server remain on the server.

The Syslog service on a Syslog server receives logging messages from applications on the local host or from devices such as a Layer 2 Switch or Layer 3 Switch. Syslog adds a time stamp to each received message and directs messages to a log file. Most Unix workstations come with Syslog configured. Some third party vendor products also provide Syslog running on NT.

Syslog uses UDP port 514 and each Syslog message thus is sent with destination port 514. Each Syslog message is one line with Syslog message format. The message is embedded in the text portion of the Syslog format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfield order is insensitive except that the text subfield should be the last field in the message. All the subfields are optional.

Displaying Syslog messages

To display the Syslog messages in the device local buffer, enter the **show logging** command at any level of the CLI. The following shows an example display output.

```
PowerConnect#show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

For information about the Syslog configuration information, time stamps, and dynamic and static buffers, refer to [“Displaying the Syslog configuration”](#) on page 1037.

Enabling real-time display of Syslog messages

By default, to view Syslog messages generated by a device, you need to display the Syslog buffer or the log on a Syslog server used by the device.

You can enable real-time display of Syslog messages on the management console. When you enable this feature, the software displays a Syslog message on the management console when the message is generated. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI.

```
PowerConnect(config)#logging console
```

Syntax: [no] logging console

This command enables the real-time display of Syslog messages on the serial console. You can enter this command from the serial console or a Telnet or SSH session.

Enabling real-time display for a Telnet or SSH session

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session.

```
telnet@PowerConnect#terminal monitor
Syslog trace was turned ON
```

Syntax: terminal monitor

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@PowerConnect#terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed.

```
telnet@PowerConnect#terminal monitor
Syslog trace was turned ON
SYSLOG: <9>PowerConnect, Power supply 2, power supply on left connector, failed

SYSLOG: <14>PowerConnect, Interface ethernet 6, state down

SYSLOG: <14>PowerConnect, Interface ethernet 2, state up
```

Show log on all terminals

It permits any terminal logged on to a switch to receive real-time Syslog messages when the **terminal monitor** command is issued.

Configuring the Syslog service

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a Syslog server. You can configure the device to use up to six Syslog servers. (Use of a Syslog server is optional. The system can hold up to 1000 Syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local Syslog buffer can hold.
- Display the Syslog configuration.
- Clear the local Syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- By default, up to 50 messages are retained in the local Syslog buffer. This can be changed.
- No Syslog server is specified.

Displaying the Syslog configuration

To display the Syslog parameters currently in effect on a device, enter the following command from any level of the CLI.

```
PowerConnect#show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Syntax: show logging

The Syslog display shows the following configuration information, in the rows above the log entries themselves.

TABLE 167 CLI display of Syslog buffer configuration

This field...	Displays...
Syslog logging	The state (enabled or disabled) of the Syslog buffer.
messages dropped	The number of Syslog messages dropped due to user-configured filters. By default, the software logs messages for all Syslog levels. You can disable individual Syslog levels, in which case the software filters out messages at those levels. Refer to “Disabling logging of a message level” on page 1042. Each time the software filters out a Syslog message, this counter is incremented.
flushes	The number of times the Syslog buffer has been cleared by the clear logging command. Refer to “Clearing the Syslog messages from the local buffer” on page 1044.
overruns	The number of times the dynamic log buffer has filled up and been cleared to hold new entries. For example, if the buffer is set for 100 entries, the 101st entry causes an overrun. After that, the 201st entry causes a second overrun.
level	The message levels that are enabled. Each letter represents a message type and is identified by the key (level code) below the value. If you disable logging of a message level, the code for that level is not listed.
messages logged	The total number of messages that have been logged since the software was loaded.
level code	The message levels represented by the one-letter codes.

Static and dynamic buffers

The software provides two buffers:

- Static – logs power supply failures, fan failures, and temperature warning or shutdown messages
- Dynamic – logs all other message types

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (50 by default), then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
PowerConnect#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
```

Static Log Buffer:

```
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
```

Dynamic Log Buffer (50 entries):

```
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Notice that the static buffer contains two separate messages for fan failures. Each message of each type has its own buffer. Thus, if you replace fan 1 but for some reason that fan also fails, the software replaces the first message about the failure of fan 1 with the newer message. The software does not overwrite the message for fan 2, unless the software sends a newer message for fan 2.

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command at the Privileged EXEC level.

```
PowerConnect#clear logging dynamic-buffer
```

Syntax: clear logging [dynamic-buffer | static-buffer]

You can specify **dynamic-buffer** to clear the dynamic buffer or **static-buffer** to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Time stamps

The contents of the time stamp differ depending on whether you have set the time and date on the onboard system clock:

- If you have set the time and date on the onboard system clock, the date and time are shown in the following format.

mm dd hh:mm:ss

where

- *mm* – abbreviation for the name of the month
- *dd* – day

- *hh* – hours
- *mm* – minutes
- *ss* – seconds

For example, “Oct 15 17:38:03” means October 15 at 5:38 PM and 3 seconds.

- If you have not set the time and date on the onboard system clock, the time stamp shows the amount of time that has passed since the device was booted, in the following format.

`<num>d<num>h<num>m<num>s`

where

- `<num>d` – day
- `<num>h` – hours
- `<num>m` – minutes
- `<num>s` – seconds

For example, “188d1h01m00s” means the device had been running for 188 days, 11 hours, one minute, and zero seconds when the Syslog entry with this time stamp was generated.

Example of Syslog messages on a device with the onboard clock set

The example shows the format of messages on a device where the onboard system clock has been set. Each time stamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
PowerConnect#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 38 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
```

```
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
```

```
Dynamic Log Buffer (50 entries):
Oct 15 17:38:03:warning:list 101 denied tcp 209.157.22.191(0)(Ethernet 18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
Oct 15 07:03:30:warning:list 101 denied tcp 209.157.22.26(0)(Ethernet 18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
Oct 15 06:58:30:warning:list 101 denied tcp 209.157.22.198(0)(Ethernet 18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
```

Example of Syslog messages on a device with the onboard clock not set

The example shows the format of messages on a device where the onboard system clock is not set. Each time stamp shows the amount of time the device had been running when the message was generated. For example, the most recent message, at the top of the list of messages, was generated when the device had been running for 21 days, seven hours, two minutes, and 40 seconds.

```
PowerConnect#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 38 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):
21d07h02m40s:warning:list 101 denied tcp 209.157.22.191(0)(Ethernet18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
19d07h03m30s:warning:list 101 denied tcp 209.157.22.26(0)(Ethernet 18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
17d06h58m30s:warning:list 101 denied tcp 209.157.22.198(0)(Ethernet18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
```

Disabling or re-enabling Syslog

Syslog is enabled by default. To disable it, enter the following command at the global CONFIG level.

```
PowerConnect(config)#no logging on
```

Syntax: [no] logging on [*udp-port*]

The *udp-port* parameter specifies the application port used for the Syslog facility. The default is 514.

To re-enable logging, enter the following command.

```
PowerConnect(config)#logging on
```

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No Syslog server is specified.

Specifying a Syslog server

To specify a Syslog server, enter a command such as the following.

```
PowerConnect(config)#logging host 10.0.0.99
```

Syntax: logging host *ip-addr* | *server-name*

Specifying an additional Syslog server

To specify an additional Syslog server, enter the **logging host** *ip-addr* command again, as in the following example. You can specify up to six Syslog servers.

```
PowerConnect(config)#logging host 10.0.0.99
```

Syntax: logging host *ip-addr* | *server-name*

Disabling logging of a message level

To change the message level, disable logging of specific message levels. You must disable the message levels on an individual basis.

For example, to disable logging of debugging and informational messages, enter the following commands.

```
PowerConnect(config)#no logging buffered debugging
PowerConnect(config)#no logging buffered informational
```

Syntax: [no] logging buffered <level> | <num-entries>

The <level> parameter can have one of the following values:

- alerts
- critical
- debugging
- emergencies
- errors
- informational
- notifications
- warnings

The commands in the example above change the log level to notification messages or higher. The software will not log informational or debugging messages. The changed message level also applies to the Syslog servers.

Changing the number of entries the local buffer can hold

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store. For example.

```
PowerConnect(config)#logging buffered 100
PowerConnect(config)#write mem
PowerConnect(config)#exit
PowerConnect#reload
```

Syntax: logging buffered <num>

The default number of messages is 50.

Configuration notes

- If you decrease the size of the buffer, the software clears the buffer before placing the change into effect.

Changing the log facility

The Syslog daemon on the Syslog server uses a facility to determine where to log the messages from the device. The default facility for messages the device sends to the Syslog server is “user”. You can change the facility using the following command.

NOTE

You can specify only one facility. If you configure the device to use two Syslog servers, the device uses the same facility on both servers.

```
PowerConnect(config)#logging facility local0
```

Syntax: `logging facility <facility-name>`

The <facility-name> can be one of the following:

- kern – kernel messages
- user – random user-level messages
- mail – mail system
- daemon – system daemons
- auth – security or authorization messages
- syslog – messages generated internally by Syslog
- lpr – line printer subsystem
- news – netnews subsystem
- uucp – uucp subsystem
- sys9 – cron/at subsystem
- sys10 – reserved for system use
- sys11 – reserved for system use
- sys12 – reserved for system use
- sys13 – reserved for system use
- sys14 – reserved for system use
- cron – cron/at subsystem
- local0 – reserved for local use
- local1 – reserved for local use
- local2 – reserved for local use
- local3 – reserved for local use
- local4 – reserved for local use
- local5 – reserved for local use
- local6 – reserved for local use
- local7 – reserved for local use

Displaying Interface names in Syslog messages

By default, an interface port number are displayed when you display Syslog messages. If you want to display the name of the interface instead of its number, enter the following command:

```
PowerConnect(config)# ip show-portname
```

This command is applied globally to all interfaces on Layer 2 Switches and Layer 3 Switches.

Syntax: `[no] ip show-portname`

When you display the messages in the Syslog, you see the interface name under the Dynamic Log Buffer section. The actual interface number is appended to the interface name. For example, if the interface name is "lab" and its port number is "2", you see "lab2" displayed as in the example below:

```
PowerConnect# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet Lab2, state up
Dec 15 18:45:15:I:Warm start
```

Displaying TCP or UDP port numbers in Syslog messages

The command **ip show-service-number-in-log** allows you to change the display of TCP or UDP application information from the TCP or UDP well-known port name to the TCP or UDP port number. For example, when this command is in effect, the device will display **http** (the well-known port name) instead of **80** (the port number) in the output of show commands, and other commands that contain application port information. By default, devices display TCP or UDP application information in named notation.

To display TCP or UDP port numbers instead of their names, enter the following command.

```
PowerConnect(config)#ip show-service-number-in-log
```

Syntax: [no] ip show-service-number-in-log

Clearing the Syslog messages from the local buffer

To clear the Syslog messages stored in the local buffer of the device, enter the following command.

```
PowerConnect#clear logging
```

Syntax: clear logging

Syslog messages

[Table 168](#) lists all of the Syslog messages. Note that some of the messages apply only to Layer 3 Switches. The messages are listed by message level, in the following order, then by message type:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

TABLE 168 Dell Syslog messages

Message level	Message	Explanation
Alert	<num-modules> modules and 1 power supply, need more power supply!!	Indicates that the chassis needs more power supplies to run the modules in the chassis. The <num-modules> parameter indicates the number of modules in the chassis.
Alert	Fan <num>, <location>, failed	A fan has failed. The <num> is the fan number. The <location> describes where the failed fan is in the chassis.
Alert	ISIS MEMORY USE EXCEEDED	IS-IS is requesting more memory than is available.
Alert	MAC Authentication failed for <mac-address> on <portnum>	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the VLAN returned in the RADIUS Access-Accept message did not refer to a valid VLAN or VLAN ID on the device. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (Invalid User)	RADIUS authentication failed for the specified <mac-address> on the specified <portnum> because the MAC address sent to the RADIUS server was not found in the RADIUS server users database.
Alert	MAC Authentication failed for <mac-address> on <portnum> (No VLAN Info received from RADIUS server)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, dynamic VLAN assignment was enabled for the port, but the RADIUS Access-Accept message did not include VLAN information. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (Port is already in another radius given vlan)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the RADIUS Access-Accept message specified a VLAN ID, although the port had previously been moved to a different RADIUS-assigned VLAN. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (RADIUS given vlan does not exist)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the RADIUS Access-Accept message specified a VLAN that does not exist in the configuration. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (RADIUS given VLAN does not match with TAGGED vlan)	Multi-device port authentication failed for the <mac-address> on a tagged port because the packet with this MAC address as the source was tagged with a VLAN ID different from the RADIUS-supplied VLAN ID.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Alert	Management module at state changed from <module-state> to <module-state>.	Indicates a state change in a management module. The <module-state> can be one of the following: <ul style="list-style-type: none"> • active • standby • crashed • coming-up • unknown
Alert	OSPF LSA Overflow, LSA Type = <lsa-type>	Indicates an LSA database overflow. The <lsa-type> parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following: <ul style="list-style-type: none"> • 1 – Router • 2 – Network • 3 – Summary • 4 – Summary • 5 – External
Alert	OSPF Memory Overflow	OSPF has run out of memory.
Alert	Power supply <num>, <location>, failed	A power supply has failed. The <num> is the power supply number. The <location> describes where the failed power supply is in the chassis.
Alert	System: No Free Tcam Entry available. System will be unstable	The limit for the TCAM routing entries has been reached. You must reboot the device.
Alert	System: Temperature is over shutdown level, system is going to be reset in <num> seconds	The chassis temperature has risen above shutdown level. The system will be shut down in the amount of time indicated.
Alert	Temperature <degrees> C degrees, warning level <warn-degrees> C degrees, shutdown level <shutdown-degrees> C degrees	Indicates an over temperature condition on the active module. The <degrees> value indicates the temperature of the module. The <warn-degrees> value is the warning threshold temperature configured for the module. The <shutdown-degrees> value is the shutdown temperature configured for the module.
Critical	Authentication shut down <portnum> due to DOS attack	Denial of Service (DoS) attack protection was enabled for multi-device port authentication on the specified <portnum>, and the per-second rate of RADIUS authentication attempts for the port exceeded the configured limit. The device considers this to be a DoS attack and disables the port.
Debug	BGP4: Not enough memory available to run BGP4	The device could not start the BGP4 routing protocol because there is not enough memory available.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Debug	DOT1X: Not enough memory	There is not enough system memory for 802.1X authentication to take place. Contact Dell Technical Support.
Error	No of prefixes received from BGP peer <ip-addr> exceeds maximum prefix-limit...shutdown	The Layer 3 Switch has received more than the specified maximum number of prefixes from the neighbor, and the Layer 3 Switch is therefore shutting down its BGP4 session with the neighbor.
Information	IPv6: IPv6 protocol disabled on the device from <session-id>	IPv6 protocol was disabled on the device during the specified session.
Information	IPv6: IPv6 protocol enabled on the device from <session-id>	IPv6 protocol was enabled on the device during the specified session.
Information	MAC Filter applied to port <port-id> by <username> from <session-id> (filter id=<filter-ids>)	Indicates a MAC filter was applied to the specified port by the specified user during the specified session. <session-id> can be console, telnet, ssh, or snmp. <filter-ids> is a list of the MAC filters that were applied.
Information	MAC Filter removed from port <port-id> by <username> from <session-id> (filter id=<filter-ids>)	Indicates a MAC filter was removed from the specified port by the specified user during the specified session. <session-id> can be console, telnet, ssh, or snmp. <filter-ids> is a list of the MAC filters that were removed.
Information	Security: Password has been changed for user <username> from <session-id>	Password of the specified user has been changed during the specified session ID or type. <session-id> can be console, telnet, ssh, or snmp.
Informational	<device-name> : Logical link on interface ethernet <port#> is down.	The specified ports were logically brought down while singleton was configured on the port.
Informational	<device-name>: Logical link on interface ethernet <port#> is up.	The specified ports were logically brought up while singleton was configured on the port.
Informational	<user-name> login to PRIVILEGED mode	A user has logged into the Privileged EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> login to USER EXEC mode	A user has logged into the USER EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> logout from PRIVILEGED mode	A user has logged out of Privileged EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> logout from USER EXEC mode	A user has logged out of the USER EXEC mode of the CLI. The <user-name> is the user name.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	ACL <ACL id> added deleted modified from console telnet ssh snmp session	A user created, modified, deleted, or applied an ACL through a SNMP, console, SSH, or Telnet session.
Informational	Bridge is new root, vlan <vlan-id>, root ID <root-id>	A Spanning Tree Protocol (STP) topology change has occurred, resulting in the device becoming the root bridge. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <root-id> is the STP bridge root ID.
Informational	Bridge root changed, vlan <vlan-id>, new root ID <string>, root interface <portnum>	A Spanning Tree Protocol (STP) topology change has occurred. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <root-id> is the STP bridge root ID. The <portnum> is the number of the port connected to the new root bridge.
Informational	Bridge topology change, vlan <vlan-id>, interface <portnum>, changed state to <stp-state>	A Spanning Tree Protocol (STP) topology change has occurred on a port. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <portnum> is the port number. The <stp-state> is the new STP state and can be one of the following: <ul style="list-style-type: none"> • disabled • blocking • listening • learning • forwarding • unknown
Informational	Cold start	The device has been powered on.
Informational	DOT1X : port <portnum> - mac <mac address> Cannot apply an ACL or MAC filter on a port member of a VE (virtual interface)	The RADIUS server returned an IP ACL or MAC address filter, but the port is a member of a virtual interface (VE).
Informational	DOT1X : port <portnum> - mac <mac address> cannot remove inbound ACL	An error occurred while removing the inbound ACL.
Informational	DOT1X : port <portnum> - mac <mac address> Downloading a MAC filter, but MAC filter have no effect on router port	The RADIUS server returned a MAC address filter, but the <portnum> is a router port (it has one or more IP addresses).
Informational	DOT1X : port <portnum> - mac <mac address> Downloading an IP ACL, but IP ACL have no effect on a switch port	The RADIUS server returned an IP ACL, but the <portnum> is a switch port (no IP address).
Informational	DOT1X : port <portnum> - mac <mac address> Error - could not add all MAC filters	The device was unable to implement the MAC address filters returned by the RADIUS server.
Informational	DOT1X : port <portnum> - mac <mac address> Invalid MAC filter ID - this ID doesn't exist	The MAC address filter ID returned by the RADIUS server does not exist in the configuration.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	DOT1X : port <portnum> - mac <mac address> Invalid MAC filter ID - this ID is user defined and cannot be used	The port was assigned a MAC address filter ID that had been dynamically created by another user.
Informational	DOT1X : port <portnum> - mac <mac address> is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication failed for the Client with the specified <mac address> on the specified <portnum> either due to insufficient system resources on the device, or due to invalid IP ACL or MAC address filter information returned by the RADIUS server.
Informational	DOT1X : port <portnum> - mac <mac address> Port is already bound with MAC filter	The RADIUS server returned a MAC address filter, but a MAC address filter had already been applied to the port.
Informational	DOT1X : port <portnum> - mac <mac address> This device doesn't support ACL with MAC Filtering on the same port	The RADIUS server returned a MAC address filter while an IP ACL was applied to the port, or returned an IP ACL while a MAC address filter was applied to the port.
Informational	DOT1X Port <portnum> is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: <ul style="list-style-type: none"> • Insufficient system resources were available on the device to apply an IP ACL or MAC address filter to the port • Invalid information was received from the RADIUS server (for example, the Filter-ID attribute did not refer to an existing IP ACL or MAC address filter)
Informational	DOT1X: Port <portnum> currently used vlan-id changes to <vlan-id> due to dot1x-RADIUS vlan assignment	A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by <vlan-id>.
Informational	DOT1X: Port <portnum> currently used vlan-id is set back to port default vlan-id <vlan-id>	The user connected to <portnum> has disconnected, causing the port to be moved back into its default VLAN, <vlan-id>.
Informational	DOT1X: Port <portnum>, AuthControlledPortStatus change: authorized	The status of the interface controlled port has changed from unauthorized to authorized.
Informational	DOT1X: Port <portnum>, AuthControlledPortStatus change: unauthorized	The status of the interface controlled port has changed from authorized to unauthorized.
Informational	Enable super port-config read-only password deleted added modified from console telnet ssh snmp OR Line password deleted added modified from console telnet ssh snmp	A user created, re-configured, or deleted an Enable or Line password through the SNMP, console, SSH, or Telnet session.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	ERR_DISABLE: Interface ethernet <port-number>, err-disable recovery timeout	Errdisable recovery timer expired and the port has been reenabled.
Informational	ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout	If the wait time (port is down and is waiting to come up) expires and the port is brought up the following message is displayed.
Informational	ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold; port in err-disable state	The threshold for the number of times that a port link toggles from “up” to “down” and “down” to “up” has been exceeded.
Informational	Interface <portnum>, line protocol down	The line protocol on a port has gone down. The <portnum> is the port number.
Informational	Interface <portnum>, line protocol up	The line protocol on a port has come up. The <portnum> is the port number.
Informational	Interface <portnum>, state down	A port has gone down. The <portnum> is the port number.
Informational	Interface <portnum>, state up	A port has come up. The <portnum> is the port number.
Informational	MAC Based Vlan Disabled on port <port id>	A MAC Based VLAN has been disabled on a port
Informational	MAC Based Vlan Enabled on port <port id>	A MAC Based VLAN has been enabled on a port.
Informational	MAC Filter added deleted modified from console telnet ssh snmp session filter id = <MAC filter ID>, src mac = <Source MAC address> any, dst mac = <Destination MAC address> any	A user created, modified, deleted, or applied this MAC filter through the SNMP, console, SSH, or Telnet session.
Informational	MSTP: BPDU-guard interface ethernet <port-number> detect (Received BPDU), putting into err-disable state.	BPDU guard violation occurred in MSTP.
Informational	OPTICAL MONITORING: port <port-number> is not capable.	The optical transceiver is qualified by Dell, but the transceiver does not support digital optical performance monitoring.
Informational	Port <p> priority changed to <n>	A port priority has changed.
Informational	Port <portnum>, srcip-security max-ipaddr-per-int reached.Last IP=<ipaddr>	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Informational	Port <portnum>, srcip-security max-ipaddr-per-int reached.Last IP=<ipaddr>	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Informational	Security: console login by <username> to USER PRIVILEGE EXEC mode	The specified user logged into the device console into the specified EXEC mode.
Informational	Security: console logout by <username>	The specified user logged out of the device console.
Informational	Security: telnet SSH login by <username> from src IP <ip-address>, src MAC <mac-address> to USER PRIVILEGE EXEC mode	The specified user logged into the device using Telnet or SSH from either or both the specified IP address and MAC address. The user logged into the specified EXEC mode.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	Security: telnet SSH logout by <username> from src IP <ip-address>, src MAC <mac-address> to USER PRIVILEGE EXEC mode	The specified user logged out of the device. The user was using Telnet or SSH to access the device from either or both the specified IP address and MAC address. The user logged out of the specified EXEC mode.
Informational	SNMP read-only community read-write community contact location user group view engineId trap [host] [<value-str>] deleted added modified from console telnet ssh snmp session	A user made SNMP configuration changes through the SNMP, console, SSH, or Telnet session. [<value-str>] does not appear in the message if SNMP community or engineId is specified.
Informational	SNMP Auth. failure, intruder IP: <ip-addr>	A user has tried to open a management session with the device using an invalid SNMP community string. The <ip-addr> is the IP address of the host that sent the invalid community string.
Informational	SSH telnet server enabled disabled from console telnet ssh snmp session [by user <username>]	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable/disable configuration through the SNMP, console, SSH, or Telnet session.
Informational	startup-config was changed or startup-config was changed by <user-name>	A configuration change was saved to the startup-config file. The <user-name> is the user ID, if they entered a user ID to log in.
Informational	STP: Root Guard Port <port-number>, VLAN <vlan-ID> consistent (Timeout).	Root guard unblocks a port.
Informational	STP: Root Guard Port <port-number>, VLAN <vlan-ID> inconsistent (Received superior BPDU).	Root guard blocked a port.
Informational	STP: VLAN <vlan id> BPDU-Guard on Port <port id> triggered (Received BPDU), putting into err-disable state	The BPDU guard feature has detected an incoming BPDU on {vlan-id, port-id}
Informational	STP: VLAN <vlan id> Root-Protect Port <port id>, Consistent (Timeout)	The root protect feature goes back to the consistent state.
Informational	STP: VLAN <vlan id> Root-Protect Port <port id>, Inconsistent (Received superior BPDU)	The root protect feature has detected a superior BPDU and goes into the inconsistent state on {vlan-id, port-id}.
Informational	STP: VLAN <vlan-id> BPDU-guard port <port-number> detect (Received BPDU), putting into err-disable state	STP placed a port into an errdisable state for BPDU guard.
Informational	STP: VLAN 1 BPDU-guard port <port-number> detect (Received BPDU), putting into err-disable state.	BPDU guard violation in occurred in STP or RSTP.
Informational	Syslog server <IP-address> deleted added modified from console telnet ssh snmp OR Syslog operation enabled disabled from console telnet ssh snmp	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation through the SNMP, console, SSH, or Telnet session.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	SYSTEM: Optic is not Dell-qualified (<port-number>)	Dell does not support the optical transceiver.
Informational	System: Fan <fan id> (from left when facing right side), ok	The fan status has changed from fail to normal.
Informational	System: Fan speed changed automatically to <fan speed>	The system automatically changed the fan speed to the speed specified in this message.
Informational	System: No free TCAM entry. System will be unstable	There are no TCAM entries available.
Informational	System: Static Mac entry with Mac Address <mac-address> is added from the <unit>/<port> to <unit>/<port> on VLANs <vlan-id> to <vlan-id>	A MAC address is added to a range of interfaces, which are members of the specified VLAN range.
Informational	System: Static Mac entry with Mac Address <mac-address> is added to ethe <unit>/<port> to <unit>/<port> on <vlan-id>	A MAC address is added to a range of interfaces, which are members of the specified VLAN.
Informational	System: Static Mac entry with Mac Address <mac-address> is added to portnumber <unit>/<port> on VLAN <vlan-id>	A MAC address is added to an interface and the interface is a member of the specified VLAN.
Informational	System: Static Mac entry with Mac Address <mac-address> is deleted from ethe <unit>/<port> to <unit>/<port> on <vlan-id>	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN.
Informational	System: Static Mac entry with Mac Address <mac-address> is deleted from ethe <unit>/<port> to <unit>/<port> on VLANs <vlan-id> to <vlan-id>	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN range.
Informational	System: Static Mac entry with Mac Address <mac-address> is deleted from portnumber <unit>/<port> on <vlan-id>	A MAC address is deleted from an interface and the interface is a member of the specified VLAN.
Informational	System: Static Mac entry with Mac Address <mac-address> is deleted from portnumber <unit>/<port> on VLANs <vlan-id> to <vlan-id>	A MAC address is deleted from an interface and the interface is a member of the specified VLAN range.
Informational	telnet SSH access [by <username>] from src IP <source ip address>, src MAC <source MAC address> rejected, <n> attempts	There were failed SSH, or Telnet login access attempts from the specified source IP and MAC address. <ul style="list-style-type: none"> [by <user> <username>] does not appear if telnet or SSH clients are specified. <n> is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Informational	Trunk group (<ports>) created by 802.3ad link-aggregation module.	802.3ad link aggregation is configured on the device, and the feature has dynamically created a trunk group (aggregate link). The <ports> is a list of the ports that were aggregated to make the trunk group.
Informational	user <username> added deleted modified from console telnet ssh snmp	A user created, modified, or deleted a local user account through the SNMP, console, SSH, or Telnet session.
Informational	vlan <vlan id> added deleted modified from console telnet ssh snmp session	A user created, modified, or deleted a VLAN through the SNMP, console, SSH, or Telnet session.
Informational	Warm start	The system software (flash code) has been reloaded.
Informational	vlan <vlan-id> Bridge is RootBridge <mac-address> (MgmtPriChg)	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Informational	vlan <vlan-id> Bridge is RootBridge <mac-address> (MsgAgeExpiry)	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Informational	vlan <vlan-id> interface <portnum> Bridge TC Event (DOT1wTransition)	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.
Informational	vlan <vlan-id> interface <portnum> STP state -> <state> (DOT1wTransition)	802.1W changed the state of a port to a new state: forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Informational	vlan <vlan-id> New RootBridge <mac-address> RootPort <portnum> (BpduRcvd)	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Informational	vlan <vlan-id> New RootPort <portnum> (RootSelection)	802.1W changed the port role to Root port, using the root selection computation.
Notification	ACL exceed max DMA L4 cam resource, using flow based ACL instead	The port does not have enough Layer 4 CAM entries for the ACL. To correct this condition, allocate more Layer 4 CAM entries. To allocate more Layer 4 CAM entries, enter the following command at the CLI configuration level for the interface: ip access-group max-l4-cam <num>
Notification	ACL insufficient L4 cam resource, using flow based ACL instead	The port does not have a large enough CAM partition for the ACLs
Notification	ACL insufficient L4 session resource, using flow based ACL instead	The device does not have enough Layer 4 session entries. To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command at the global CONFIG level of the CLI interface: system-max session-limit <num>

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	ACL port fragment packet inspect rate <rate> exceeded on port <portnum>	The fragment rate allowed on an individual interface has been exceeded. The <rate> indicates the maximum rate allowed. The <portnum> indicates the port. This message can occur if fragment throttling is enabled.
Notification	ACL system fragment packet inspect rate <rate> exceeded	The fragment rate allowed on the device has been exceeded. The <rate> indicates the maximum rate allowed. This message can occur if fragment throttling is enabled.
Notification	Authentication Disabled on <portnum>	The multi-device port authentication feature was disabled on the on the specified <portnum>.
Notification	Authentication Enabled on <portnum>	The multi-device port authentication feature was enabled on the on the specified <portnum>.
Notification	BGP Peer <ip-addr> DOWN (IDLE)	Indicates that a BGP4 neighbor has gone down. The <ip-addr> is the IP address of the neighbor BGP4 interface with the device.
Notification	BGP Peer <ip-addr> UP (ESTABLISHED)	Indicates that a BGP4 neighbor has come up. The <ip-addr> is the IP address of the neighbor BGP4 interface with the PowerConnect device.
Notification	DOT1X issues software but not physical port down indication of Port <portnum> to other software applications	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Notification	DOT1X issues software but not physical port up indication of Port <portnum> to other software applications	The device has indicated that the specified port has been authenticated, but the actual port may not be active.
Notification	DOT1X: Port <port_id> Mac <mac_address> -user <user_id> - RADIUS timeout for authentication	The RADIUS session has timed out for this 802.1x port.
Notification	ISIS ENTERED INTO OVERLOAD STATE	The Layer 3 Switch has set the overload bit to on (1), indicating that the Layer 3 Switch IS-IS resources are overloaded.
Notification	ISIS EXITING FROM OVERLOAD STATE	The Layer 3 Switch has set the overload bit to off (0), indicating that the Layer 3 Switch IS-IS resources are no longer overloaded.
Notification	ISIS L1 ADJACENCY DOWN <system-id> on circuit <circuit-id>	The Layer 3 Switch adjacency with this Level-1 IS has gone down. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	ISIS L1 ADJACENCY UP <system-id> on circuit <circuit-id>	The Layer 3 Switch adjacency with this Level-1 IS has come up. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	ISIS L2 ADJACENCY DOWN <system-id> on circuit <circuit-id>	The Layer 3 Switch adjacency with this Level-2 IS has gone down. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	ISIS L2 ADJACENCY UP <system-id> on circuit <circuit-id>	The Layer 3 Switch adjacency with this Level-2 IS has come up. The <system-id> is the system ID of the IS. The <circuit-id> is the ID of the circuit over which the adjacency was established.
Notification	Local ICMP exceeds <burst-max> burst packets, stopping for <lockup> seconds!!	The number of ICMP packets exceeds the <burst-max> threshold set by the ip icmp burst command. The device may be the victim of a Denial of Service (DoS) attack. All ICMP packets will be dropped for the number of seconds specified by the <lockup> value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	Local TCP exceeds <burst-max> burst packets, stopping for <lockup> seconds!!	The number of TCP SYN packets exceeds the <burst-max> threshold set by the ip tcp burst command. The PowerConnect device may be the victim of a TCP SYN DoS attack. All TCP SYN packets will be dropped for the number of seconds specified by the <lockup> value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	Local TCP exceeds <num> burst packets, stopping for <num> seconds!!	Threshold parameters for local TCP traffic on the device have been configured, and the maximum burst size for TCP packets has been exceeded. The first <num> is the maximum burst size (maximum number of packets allowed). The second <num> is the number of seconds during which additional TCP packets will be blocked on the device. NOTE: This message can occur in response to an attempted TCP SYN attack.
Notification	MAC Authentication RADIUS timeout for <mac_address> on port <port_id>	The RADIUS session has timed out for the MAC address for this port.
Notification	MAC Authentication succeeded for <mac-address> on <portnum>	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF interface state changed, rid <router-id>, intf addr <ip-addr>, state <ospf-state>	<p>Indicates that the state of an OSPF interface has changed.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the interface IP address.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown
Notification	OSPF intf authen failure, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF interface authentication failure has occurred.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf config error, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF interface configuration error has occurred.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the error packet.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPF intf rcvd bad pkt, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, pkt type <pkt-type>	<p>Indicates that an OSPF interface received a bad packet.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf rcvd bad pkt: Bad Checksum, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet that had an invalid checksum.</p> <p>The rid <ip-addr> is the router ID.</p> <p>The intf addr <ip-addr> is the IP address of the interface that received the packet.</p> <p>The pkt size <num> is the number of bytes in the packet.</p> <p>The checksum <num> is the checksum value for the packet.</p> <p>The pkt src addr <ip-addr> is the IP address of the neighbor that sent the packet.</p> <p>The pkt type <type> is the OSPF packet type and can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state acknowledgement • unknown (indicates an invalid packet type)
Notification	OSPF intf rcvd bad pkt: Bad Packet type, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet with an invalid type.</p> <p>The parameters are the same as for the Bad Checksum message. The pkt type <type> value is "unknown", indicating that the packet type is invalid.</p>
Notification	OSPF intf rcvd bad pkt: Invalid packet size, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet with an invalid packet size.</p> <p>The parameters are the same as for the Bad Checksum message.</p>
Notification	OSPF intf rcvd bad pkt: Unable to find associated neighbor, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The neighbor IP address in the packet is not in the list of OSPF neighbors in the device.</p> <p>The parameters are the same as for the Bad Checksum message.</p>

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf retransmit, rid <router-id>, intf addr <ip-addr>, nbr rid <nbr-router-id>, pkt type is <pkt-type>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <nbr-router-id> is the router ID of the neighbor router.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>
Notification	OSPF LSDB approaching overflow, rid <router-id>, limit <num>	<p>The software is close to an LSDB condition.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <num> is the number of LSAs.</p>
Notification	OSPF LSDB overflow, rid <router-id>, limit <num>	<p>A Link State Database Overflow (LSDB) condition has occurred.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <num> is the number of LSAs.</p>
Notification	OSPF max age LSA, rid <router-id>, area <area-id>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An LSA has reached its maximum age.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <area-id> is the OSPF area.</p> <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF nbr state changed, rid <router-id>, nbr addr <ip-addr>, nbr rid <nbr-router-id>, state <ospf-state>	<p>Indicates that the state of an OSPF neighbor has changed.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the neighbor.</p> <p>The <nbr-router-id> is the router ID of the neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	OSPF originate LSA, rid <router-id>, area <area-id>, LSA type <lsa-type>, LSA id <lsa-id>, LSA router id <lsa-router-id>	<p>An OSPF interface has originated an LSA.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <area-id> is the OSPF area.</p> <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf authen failure, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF virtual routing interface authentication failure has occurred.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf config error, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF virtual routing interface configuration error has occurred. The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the error packet.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPF virtual intf rcvd bad pkt, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, pkt type <pkt-type>	<p>Indicates that an OSPF interface received a bad packet. The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the device received the authentication failure.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf retransmit, rid <router-id>, intf addr <ip-addr>, nbr rid <nbr-router-id>, pkt type is <pkt-type>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An OSPF interface on the device has retransmitted a Link State Advertisement (LSA).</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the interface on the device.</p> <p>The <nbr-router-id> is the router ID of the neighbor router.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>
Notification	OSPF virtual intf state changed, rid <router-id>, area <area-id>, nbr <ip-addr>, state <ospf-state>	<p>Indicates that the state of an OSPF virtual routing interface has changed.</p> <p>The <router-id> is the router ID of the router the interface is on.</p> <p>The <area-id> is the area the interface is in.</p> <p>The <ip-addr> is the IP address of the OSPF neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual nbr state changed, rid <router-id>, nbr addr <ip-addr>, nbr rid <nbr-router-id>, state <ospf-state>	<p>Indicates that the state of an OSPF virtual neighbor has changed.</p> <p>The <router-id> is the router ID of the device.</p> <p>The <ip-addr> is the IP address of the neighbor.</p> <p>The <nbr-router-id> is the router ID of the neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	Transit ICMP in interface <portnum> exceeds <num> burst packets, stopping for <num> seconds!!	<p>Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded.</p> <p>The <portnum> is the port number.</p> <p>The first <num> is the maximum burst size (maximum number of packets allowed).</p> <p>The second <num> is the number of seconds during which additional ICMP packets will be blocked on the interface.</p> <p>NOTE: This message can occur in response to an attempted Smurf attack.</p>
Notification	Transit TCP in interface <portnum> exceeds <num> burst packets, stopping for <num> seconds!	<p>Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded.</p> <p>The <portnum> is the port number.</p> <p>The first <num> is the maximum burst size (maximum number of packets allowed).</p> <p>The second <num> is the number of seconds during which additional TCP packets will be blocked on the interface.</p> <p>NOTE: This message can occur in response to an attempted TCP SYN attack.</p>

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Notification	VRRP intf state changed, intf <portnum>, vrid <virtual-router-id>, state <vrrp-state>	A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) interface. The <portnum> is the port. The <virtual-router-id> is the virtual router ID (VRID) configured on the interface. The <vrrp-state> can be one of the following: <ul style="list-style-type: none"> • init • master • backup • unknown
Warning	DOT1X security violation at port <portnum>, malicious mac address detected: <mac-address>	A security violation was encountered at the specified port number.
Warning	Dup IP <ip-addr> detected, sent from MAC <mac-addr> interface <portnum>	Indicates that the device received a packet from another device on the network with an IP address that is also configured on the device. The <ip-addr> is the duplicate IP address. The <mac-addr> is the MAC address of the device with the duplicate IP address. The <portnum> is the port that received the packet with the duplicate IP address. The address is the packet source IP address.
Warning	IGMP/MLD no hardware vidx, broadcast to the entire vlan. rated limited number	IGMP or MLD snooping has run out of hardware application VLANs. There are 4096 application VLANs per device. Traffic streams for snooping entries without an application VLAN are switched to the entire VLAN and to the CPU to be dropped. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number on non-printed warnings.
Warning	IGMP/MLD: <vlanId>(<portId>) is V1 but rcvd V2 from nbr <ipAddr>	Port has received a query with a MLD version that does not match the port MLD version. This message is rated-limited to appear a maximum of once every 10 hours.
Warning	Latched low RX Power TX Power TX Bias Current Supply Voltage Temperature warning alarm warning, port <port-number>	The optical transceiver on the given port has risen above or fallen below the alarm or warning threshold.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Warning	list <ACL-num> denied <ip-proto> <src-ip-addr> (<src-tcp/udp-port>) (Ethernet <portnum> <mac-addr>) -> <dst-ip-addr> (<dst-tcp/udp-port>), 1 event(s)	Indicates that an Access Control List (ACL) denied (dropped) packets. The <ACL-num> indicates the ACL number. Numbers 1 – 99 indicate standard ACLs. Numbers 100 – 199 indicate extended ACLs. The <ip-proto> indicates the IP protocol of the denied packets. The <src-ip-addr> is the source IP address of the denied packets. The <src-tcp/udp-port> is the source TCP or UDP port, if applicable, of the denied packets. The <portnum> indicates the port number on which the packet was denied. The <mac-addr> indicates the source MAC address of the denied packets. The <dst-ip-addr> indicates the destination IP address of the denied packets. The <dst-tcp/udp-port> indicates the destination TCP or UDP port number, if applicable, of the denied packets.
Warning	Locked address violation at interface e<portnum>, address <mac-address>	Indicates that a port on which you have configured a lock-address filter received a packet that was dropped because the packet source MAC address did not match an address learned by the port before the lock took effect. The e<portnum> is the port number. The <mac-address> is the MAC address that was denied by the address lock. Assuming that you configured the port to learn only the addresses that have valid access to the port, this message indicates a security violation.
Warning	mac filter group denied packets on port <portnum> src macaddr <mac-addr>, <num> packets	Indicates that a Layer 2 MAC filter group configured on a port has denied packets. The <portnum> is the port on which the packets were denied. The <mac-addr> is the source MAC address of the denied packets. The <num> indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.
Warning	multicast no software resource: resource-name, rate limited number	IGMP or MLD snooping has run out of software resources. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number of non-printed warnings.

TABLE 168 Dell Syslog messages (Continued)

Message level	Message	Explanation
Warning	No global IP! cannot send IGMP msg.	The device is configured for ip multicast active but there is no configured IP address and the device cannot send out IGMP queries.
Warning	No of prefixes received from BGP peer <ip-addr> exceeds warning limit <num>	The Layer 3 Switch has received more than the allowed percentage of prefixes from the neighbor. The <ip-addr> is the IP address of the neighbor. The <num> is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the Layer 3 Switch receives a 76th prefix from the neighbor.
Warning	NTP server <ip-addr> failed to respond	Indicates that a Simple Network Time Protocol (SNTP) server did not respond to the device query for the current time. The <ip-addr> indicates the IP address of the SNTP server.
Warning	rip filter list <list-num> <direction> V1 V2 denied <ip-addr>, <num> packets	Indicates that a RIP route filter denied (dropped) packets. The <list-num> is the ID of the filter list. The <direction> indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following: <ul style="list-style-type: none"> • in • out The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2). The <ip-addr> indicates the network number in the denied updates. The <num> indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.
Warning	Temperature is over warning level.	The chassis temperature has risen above the warning level.

34 Syslog messages

Network Monitoring

This appendix describes the remote monitoring features available on devices.

Basic management

The following sections contain procedures for basic system management tasks.

Viewing system information

You can access software and hardware specifics for a Layer 2 Switch or Layer 3 Switch. For software specifics, refer to [“Determining the software versions installed and running on a device”](#) on page 39.

To view the software and hardware details for the system, enter the **show version** command. The following shows an example output.

```
PowerConnect# show version
SW: Version 4.2.00b Copyright (c) 1996-2010 Brocade Communications Systems, Inc.
    Compiled on Dec 02 2010 at 08:07:06 labeled as TIR07202b071
    (6092645 bytes) from Secondary TIR07202b071
    Compressed Boot-Monitor Image size = 373767, Version:04.1.00T205 (grz04100)
HW: Stackable TurboIron-X24
=====
    Serial #: BFF2342E00X
    P-ASIC 0: type B820, rev 01 subrev 00
=====
    833 MHz Power PC processor 8541 (version 32/0020) 66 MHz bus
    512 KB boot flash memory
    31744 KB code flash memory
    512 MB DRAM
The system uptime is 5 minutes 34 seconds
The system : started=warm start   reloaded=by "reload"
```

Syntax: show version

Viewing configuration information

You can view a variety of configuration details and statistics with the **show** option. The **show** option provides a convenient way to check configuration changes before saving them to flash.

The show options available will vary for Layer 2 Switches and Layer 3 Switches and by configuration level.

To determine the available show commands for the system or a specific level of the CLI, enter the following command.

```
PowerConnect#show ?
```

Syntax: `show <option>`

You also can enter “show” at the command prompt, then press the TAB key.

Viewing port statistics

Port statistics are polled by default every 10 seconds.

You can view statistics for ports by entering the following **show** commands:

- show interfaces
- show configuration
- show statistics

To display the statistics, enter a command such as the following.

```
PowerConnect# show statistics ethernet 3
Port Link State Dupl Speed Trunk Tag Prior MAC Name
3 Up Forward Half 100M None No level0 00e0.5200.0102

Port 3 Counters:
      InOctets          3200          OutOctets          256
      InPkts            50           OutPkts            4
InBroadcastPkts      0       OutBroadcastPkts    3
InMulticastPkts     48       OutMulticastPkts    0
  InUnicastPkts      2       OutUnicastPkts     1
      InBadPkts        0
      InFragments      0
      InDiscards        0           OutErrors          0
      CRC              0           Collisions         0
      InErrors         0       LateCollisions     0
      InGiantPkts      0
      InShortPkts      0
      InJabber         0
InFlowCtrlPkts      0       OutFlowCtrlPkts    0
      InBitsPerSec     264       OutBitsPerSec      16
      InPktsPerSec     0       OutPktsPerSec      0
      InUtilization    0.00%    OutUtilization     0.00%
```

Syntax: `show statistics [ethernet <portnum>]`

The *portnum* parameter is a valid port number.

[Table 169](#) lists the statistics displayed in the output of the **show statistics** command.

TABLE 169 Port statistics

This line...	Displays...
Port configuration	
Port	The port number.
Link	The link state.
State	The STP state.
Dupl	The mode (full-duplex or half-duplex).
Speed	The port speed (10M, 100M, or 1000M).

TABLE 169 Port statistics (Continued)

This line...	Displays...
Trunk	The trunk group number, if the port is a member of a trunk group.
Tag	Whether the port is a tagged member of a VLAN.
Priori	The QoS forwarding priority of the port (level0 – level7).
MAC	The MAC address of the port.
Name	The name of the port, if you assigned a name.
Statistics	
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets sent.
InPkts	The total number of packets received. The count includes rejected and local packets that are not sent to the switching core for transmission.
OutPkts	The total number of good packets sent. The count includes unicast, multicast, and broadcast packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets sent.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets sent.
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets sent.
InBadPkts	The total number of packets received for which one of the following is true: <ul style="list-style-type: none"> • The CRC was invalid. • The packet was oversized. • Jabbers: The packets were longer than 1518 octets and had a bad FCS. • Fragments: The packets were less than 64 octets long and had a bad FCS. • The packet was undersized (short).
InFragments	The total number of packets received for which both of the following was true: <ul style="list-style-type: none"> • The length was less than 64 bytes. • The CRC was invalid.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.
OutErrors	The total number of packets with internal transmit errors such as TX underruns.
CRC	The total number of packets received for which all of the following was true: <ul style="list-style-type: none"> • The data length was between 64 bytes and the maximum allowable frame size. • No Collision or Late Collision was detected. • The CRC was invalid.
Collisions	The total number of packets received in which a Collision event was detected.
InErrors	The total number of packets received that had Alignment errors or phy errors.
LateCollisions	The total number of packets received in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.

TABLE 169 Port statistics (Continued)

This line...	Displays...
InGiantPkts	The total number of packets for which all of the following was true: <ul style="list-style-type: none"> • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. <p>NOTE: Packets are counted for this statistic regardless of whether the CRC is valid or invalid.</p>
InShortPkts	The total number of packets received for which all of the following was true: <ul style="list-style-type: none"> • The data length was less than 64 bytes. • No Rx Error was detected. • No Collision or Late Collision was detected. <p>NOTE: Packets are counted for this statistic regardless of whether the CRC is valid or invalid.</p>
InJabber	The total number of packets received for which all of the following was true: <ul style="list-style-type: none"> • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. • The CRC was invalid.
InFlowCtrlPkts	The total number of flow control packets received.
OutFlowCtrlPkts	The total number of flow control packets transmitted.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits sent per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets sent per second.
InUtilization	The percentage of the port bandwidth used by received traffic.
OutUtilization	The percentage of the port bandwidth used by sent traffic.

Viewing STP statistics

You can view a summary of STP statistics for Layer 2 Switches and Layer 3 Switches. STP statistics are by default polled every 10 seconds.

To view spanning tree statistics, enter the **show span** command. To view STP statistics for a VLAN, enter the **span vlan** command.

Clearing statistics

You can clear statistics for many parameters with the **clear** option.

To determine the available **clear** commands for the system, enter the following command.

```
PowerConnect#clear ?
```

Syntax: **clear** <option>

You also can enter “clear” at the command prompt, then press the TAB key.

NOTE

Clear commands are found at the Privileged EXEC level.

Traffic counters for outbound traffic

You can configure traffic counters (also called transmit counters) that enable the device to count the following packet types on a port or port region:

- broadcast packets
- multicast packets
- unicast packets
- dropped packets due to congestion and egress filtering

Depending on the parameters specified with the traffic counter configuration, traffic counters record the number of outbound packets from any combination of the following sources:

- a specific port or all ports in a specific port region
- a specific VLAN or all VLANs
- a specific 802.1p priority queue or all priority queues

Configuration notes

Consider the following rules when configuring traffic counters for outbound traffic.

- This feature is supported in the Layer 2, base Layer 3, and full Layer 3 codes.
- This feature applies to physical ports only, including 10 Gbps Ethernet ports and trunk ports. It does not apply to virtual interfaces.
- Once the enhanced traffic counters are read using the **show transmit-counter values** command, the counters are cleared (reset to zero).
- For each port region, you can enable a maximum of two traffic counters, regardless of whether traffic counters are enabled on individual ports or on all ports in the port region.
- Traffic counters increase for bridged filtered outbound traffic when any of the following conditions occur:
 - The port is disabled or the link is down.
 - The port or port region does not belong to the VLAN specified in the transmit counter configuration.
 - A Layer 2 protocol (e.g., spanning tree) has the port in a Blocked state.
 - The source port needs to be suppressed for multi-target packets.
 - The priority queue specified in the traffic counter is not allowed for some other reason.
 - Unknown unicast and unregistered multicast packets are filtered.

Configuration syntax

This section provides the syntax and configuration examples for enhanced traffic counters.

Example

To configure traffic counters for outbound traffic on a specific port, enter a command such as the following.

```
PowerConnect(config)#transmit-counter 4 port 18 only vlan 1 prio 7 enable
```

The above command creates and enables traffic counter 4 on port 18. The device will count the number of packets sent out on port 18 that are in VLAN 1 and have a priority queue of 7.

Example

To configure traffic counters for outbound traffic in a specific port region, enter a command such as the following.

```
PowerConnect(config)#transmit-counter 1 port 1 region vlan all prio all enable
```

The above command creates and enables traffic counter 1 on all ports that are in the same port region as port 1. The device will count the number of packets transmitted in this port region that belong to any VLAN and have any assigned priority queue.

Syntax: `[no] transmit-counter <counter-ID> port <port-num> only | region vlan <vlan-ID> | all priority <priority-queue> | all enable`

Enter the **no** form of the command to remove the outbound traffic counter.

The `<counter-ID>` parameter identifies the traffic counter. You can configure up to 64 traffic counters. Enter a number from 1 – 64.

The `<port-num>` parameter is the port number to which enhanced traffic counters will apply. Enter the port number followed by **only** to apply the enhanced traffic counter to a specific port, or enter the port number followed by **region** to apply the enhanced traffic counter to all of the ports in the port region.

The `<vlan-ID>` parameter identifies the VLAN ID for which outbound traffic will be counted. Enter a number from 0 – 4095 or enter **all** to indicate all VLANs.

The `<priority-queue>` parameter identifies the 802.1p priority queue for which traffic will be counted. Enter a number from 0 – 7 or enter **all** to indicate all priority queues.

Displaying enhanced traffic counter profiles

To display the details of the traffic counters configured on your device, enter the **show transmit-counter profiles** command. The following shows an example output.

```
PowerConnect#show transmit-counter profiles
Tx Counter      Port(s)      Vlan Id      Priority      Device      Set
      1         1 - 12         All          All          Dev 0      Set0
      4              18           1            7          Dev 1      Set0
     10        13 - 24        100          All          Dev 1      Set1
```

Displaying enhanced traffic counter statistics

To display the traffic counters for outbound traffic, enter the **show transmit-counter profiles** command.

NOTE

Once the enhanced traffic counters are displayed, the counters are cleared (reset to zero).

The following shows an example output.


```
PowerConnect#show transmit-counter values 1
```

```
Transmit Queue Counter Values for Counter 1:
```

```
Transmitted Frames:
```

```
Known Unicast           : 17204
Multicast & Unknown Unicast : 2797
Broadcast               : 5
```

```
Dropped Frames:
```

```
Bridge Egress Filtered  : 2
Congestion Drops        : 0
```

```
PowerConnect#show transmit-counter values 4
```

```
Transmit Queue Counter Values for Counter 4:
```

```
Transmitted Frames:
```

```
Known Unicast           : 124
Multicast & Unknown Unicast : 2752
Broadcast               : 0
```

```
Dropped Frames:
```

```
Bridge Egress Filtered  : 37
Congestion Drops        : 0
```

Syntax: `show transmit-counter values <number>`

where <number> identifies a valid enhanced traffic counter and is a value from 1 – 64.

TABLE 170 Outbound traffic counter statistics

This line...	Displays...
Transmitted frames	
Known Unicast	The number of known unicast packets transmitted.
Multicast & Unknown Unicast	The number of multicast and unknown unicast packets transmitted.
Broadcast	The number of broadcast packets transmitted.
Dropped Frames	
Bridge Egress Filtered	The number of bridged outbound packets that were filtered and dropped. This number includes the number of packets that were dropped because of any one of the following conditions: <ul style="list-style-type: none"> • The port was disabled or the link was down. • The port or port region does not belong to the VLAN specified in the transmit counter configuration. • A Layer 2 protocol (e.g., spanning tree) had the port in a Blocked state. • The source port was suppressed for multi-target packets. • The priority queue specified in the traffic counter was not allowed for some other reason. • Unknown unicast and unregistered multicast packets were filtered.
Congestion Drops	The number of outbound packets that were dropped because of traffic congestion.

RMON support

The RMON agent supports the following groups. The group numbers come from the RMON specification (RFC 1757):

- Statistics (RMON Group 1)
- History (RMON Group 2)
- Alarms (RMON Group 3)
- Events (RMON Group 9)

The CLI allows you to make configuration changes to the control data for these groups, but you need a separate RMON application to view and display the data graphically.

Maximum number of entries allowed in the RMON control table

You can specify the maximum number of entries allowed in the RMON control table, including alarms, history, and events. The default number of RMON entries allowed in the RMON control table is 2048 on the PowerConnect B-Series TI24X devices. The maximum number of RMON entries supported is 32768.

To set the maximum number of allowable entries to 3000 in the RMON history table, enter commands such as the following.

```
PowerConnect(config)#system-max rmon-entries 3000
PowerConnect(config)#write mem
PowerConnect(config)#exit
PowerConnect#reload
```

NOTE

You must save the change to the startup-config file and reload or reboot. The change does not take effect until you reload or reboot.

Syntax: `system-max rmon-entries <value>`

where *<value>* can be:

- 1536 – 32768 for PowerConnect B-Series TI24X devices

Statistics (RMON group 1)

Count information on multicast and broadcast packets, total packets sent, undersized and oversized packets, CRC alignment errors, jabbers, collision, fragments and dropped events is collected for each port on a Layer 2 Switch or Layer 3 Switch.

No configuration is required to activate collection of statistics for the Layer 2 Switch or Layer 3 Switch. This activity is by default automatically activated at system start-up.

You can view a textual summary of the statistics for all ports by entering the following CLI command.

```
PowerConnect#show rmon statistics
Ethernet statistics 1 is active, owned by monitor
Interface 1 (ifIndex 1) counters
    Octets 0
    Drop events 0
    Broadcast pkts 0
    CRC alignment errors 0
    Oversize pkts 0
    Jabbers 0
    64 octets pkts 0
    128 to 255 octets pkts 0
    512 to 1023 octets pkts 0
    Packets 0
    Multicast pkts 0
    Undersize pkts 0
    Fragments 0
    Collisions 0
    65 to 127 octets pkts 0
    256 to 511 octets pkts 0
    1024 to 1518 octets pkts 0
```

Syntax: show rmon statistics<portnum>

The <portnum> parameter specifies the port number. You can use the physical port number or the SNMP port number. The SNMP numbers of the ports start at 1 and increase sequentially. This command shows the following information.

TABLE 171 Export configuration and statistics

This line...	Displays...
Octets	The total number of octets of data received on the network. This number includes octets in bad packets. This number does not include framing bits but does include Frame Check Sequence (FCS) octets.
Drop events	Indicates an overrun at the port. The port logic could not receive the traffic at full line rate and had to drop some packets as a result. The counter indicates the total number of events in which packets were dropped by the RMON probe due to lack of resources. This number is not necessarily the number of packets dropped, but is the number of times an overrun condition has been detected.
Packets	The total number of packets received. This number includes bad packets, broadcast packets, and multicast packets.
Broadcast pkts	The total number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
Multicast pkts	The total number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
CRC alignment errors	The total number of packets received that were from 64 – 1518 octets long, but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). The packet length does not include framing bits but does include FCS octets.
Undersize pkts	The total number of packets received that were less than 64 octets long and were otherwise well formed. This number does not include framing bits but does include FCS octets.
Fragments	The total number of packets received that were less than 64 octets long and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). It is normal for this counter to increment, since it counts both runts (which are normal occurrences due to collisions) and noise hits. This number does not include framing bits but does include FCS octets.

TABLE 171 Export configuration and statistics (Continued)

This line...	Displays...
Oversize packets	The total number of packets received that were longer than 1518 octets and were otherwise well formed. This number does not include framing bits but does include FCS octets.
Jabbers	The total number of packets received that were longer than 1518 octets and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (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 where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms. This number does not include framing bits but does include FCS octets.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
64 octets pkts	The total number of packets received that were 64 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
65 to 127 octets pkts	The total number of packets received that were 65 – 127 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
128 to 255 octets pkts	The total number of packets received that were 128 – 255 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
256 to 511 octets pkts	The total number of packets received that were 256 – 511 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
512 to 1023 octets pkts	The total number of packets received that were 512 – 1023 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
1024 to 1518 octets pkts	The total number of packets received that were 1024 – 1518 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.

History (RMON group 2)

All active ports by default will generate two history control data entries per active Layer 2 Switch port or Layer 3 Switch interface. An active port is defined as one with a link up. If the link goes down the two entries are automatically deleted.

Two history entries are generated for each device:

- A sampling of statistics every 30 seconds
- A sampling of statistics every 30 minutes

The history data can be accessed and displayed using any of the popular RMON applications

A sample RMON history command and its syntax is shown below.

```
PowerConnect(config)#rmon history 1 interface 1 buckets 10 interval 10 owner nyc02
```

Syntax: `rmon history <entry-number> interface <portnum> buckets <number> interval <sampling-interval> owner <text-string>`

You can modify the sampling interval and the bucket (number of entries saved before overwrite) using the CLI. In the above example, owner refers to the RMON station that will request the information.

NOTE

To review the control data entry for each port or interface, enter the **show rmon history** command.

Alarm (RMON group 3)

Alarm is designed to monitor configured thresholds for any SNMP integer, time tick, gauge or counter MIB object. Using the CLI, you can define what MIB objects are monitored, the type of thresholds that are monitored (falling, rising or both), the value of those thresholds, and the sample type (absolute or delta).

An alarm event is reported each time that a threshold is exceeded. The alarm entry also indicates the action (event) to be taken if the threshold be exceeded.

A sample CLI alarm entry and its syntax is shown below.

```
PowerConnect(config)#rmon alarm 1 ifInOctets.6 10 delta rising-threshold 100 1
falling threshold 50 1 owner nyc02
```

Syntax: **rmon alarm** <entry-number> <MIB-object.interface-num> <sampling-time>
 <sample-type>
 <threshold-type> <threshold-value> <event-number> <threshold-type> <threshold-value>
 <event-number>
owner <text-string>

Event (RMON group 9)

There are two elements to the Event Group—the **event control table** and the **event log table**.

The event control table defines the action to be taken when an alarm is reported. Defined events can be found by entering the CLI command, show event. The Event Log Table collects and stores reported events for retrieval by an RMON application.

A sample entry and syntax of the event control table is shown below.

```
PowerConnect(config)#rmon event 1 description 'testing a longer string'
log-and-trap public owner nyc02
```

Syntax: **rmon event** <event-entry> **description** <text-string> **log | trap | log-and-trap** **owner**
 <rmon-station>

sFlow

sFlow is a standards-based protocol that allows network traffic to be sampled at a user-defined rate for the purpose of monitoring traffic flow patterns and identifying packet transfer rates on user-specified interfaces.

When sFlow is enabled on a Layer 2 or Layer 3 switch, the system performs the following sFlow-related tasks:

- Samples traffic flows by copying packet header information

- Identifies ingress and egress interfaces for the sampled flows
- Combines sFlow samples into UDP packets and forwards them to the sFlow collectors for analysis
- Forwards byte and packet count data, or counter samples, to sFlow collectors

sFlow is described in RFC 3176, “InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks”.

NOTE

When sFlow is enabled on a PowerConnect B-Series TI24X switch QoS will support 7 priority queues rather than 8. This is because QoS queue 1 is reserved for sFlow and does not get used for other types of traffic. Any non-sFlow packets assigned to QoS queue 1 will be redirected to QoS 0.

sFlow support for IPv6 packets

The Dell implementation of sFlow features support IPv6 packets. This support includes extended router information and extended gateway information in the sampled packet. Note that sFlow support for IPv6 packets exists only on devices running software that supports IPv6.

The configuration procedures for this feature are the same as for IPv4, except where the collector is a link-local address on a Layer 3 switch. For details refer to “[Specifying the collector](#)” on page 1082.

Extended router information

IPv6 sFlow sampled packets include the following extended router information:

- IP address of the next hop router
- Outgoing VLAN ID
- Source IP address prefix length
- Destination IP address prefix length

Note that in IPv6 devices, the prefix lengths of the source and destination IP addresses are collected if BGP is configured and the route lookup is completed. In IPv4 devices, this information is collected only if BGP is configured on the devices.

Extended gateway information

If BGP is enabled, extended gateway information is included in IPv6 sFlow sampled packets, including the following BGP information about a packet destination route:

- The autonomous system (AS) number for the router
- The source IP AS of the route
- The source peer AS for the route
- The AS patch to the destination

NOTE

AS communities and local preferences are not included in the sampled packets.

To obtain extended gateway information, use “struct extended_gateway” as described in RFC 3176.

IPv6 packet sampling

IPv6 sampling is performed by the packet processor. The system uses the sampling rate setting to selectively mark the monitoring bit in the header of an incoming packet. Marked packets tell the CPU that the packets are subject to sFlow sampling.

Configuration considerations

This section lists the sFlow configuration considerations on devices.

Hardware support

- Devices support sFlow packet sampling of inbound traffic only. These devices do not sample outbound packets. However, devices support byte and packet count statistics for both traffic directions.
- sFlow is supported on all Ethernet ports (10/100, Gbps, and 10 Gbps)

CPU utilization

Enabling sFlow may cause a slight and noticeable increase of up to 20% in CPU utilization. In typical scenarios, this is normal behavior for sFlow, and does not affect the functionality of other features on the switch.

Source address

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the IP address of the device that sent the data:

- On a Layer 2 Switch, `agent_address` is the Layer 2 Switch management IP address. You must configure the management IP address in order to export sFlow data from the device. If the switch has both an IPv4 and IPv6 address, the `agent_address` is the IPv4 address. If the switch has an IPv6 address only, the `agent_address` is the global IPv6 address.
- On a Layer 3 Switch with IPv6 interfaces only, sFlow looks for an IPv6 address in the following order, and uses the first address found:
 - The first IPv6 address on the lowest-numbered loopback interface
 - The first IPv6 address on the lowest-numbered VE interface
 - The first IPv6 address on any interface
- On a Layer 3 Switch with both IPv4 and IPv6 interfaces, or with IPv4 interfaces only, sFlow looks for an IP address in the following order, and uses the first address found:
 - The IPv4 router ID configured by the `ip router-id` command
 - The first IPv4 address on the lowest-numbered loopback interface
 - The first IPv4 address on the lowest-numbered virtual interface
 - The first IPv4 address on any interface

NOTE

The device uses the router ID only if the device also has an IP interface with the same address. Router ID is not supported on IPv6 devices.

NOTE

If an IP address is not already configured when you enable sFlow, the feature uses the source address 0.0.0.0. To display the `agent_address`, enable sFlow, then enter the **show sflow** command. Refer to “[Enabling sFlow forwarding](#)” on page 1086 and “[Displaying sFlow information](#)” on page 1087.

Sampling rate

The **sampling rate** is the average ratio of the number of packets incoming on an sFlow enabled port, to the number of flow samples taken from those packets. sFlow sampling can affect performance in some configurations.

Note that on the PowerConnect B-Series TI24X devices, the configured sampling rate and the actual rate may not be the same, depending on the configured sampling rate. The configured sampling rate adjusts to the closest value configurable on the device.

Configuring and enabling sFlow

To configure sFlow, perform the following tasks:

- Specify collector information. The collector is the external device to which you are exporting the sFlow data. You can specify up to four collectors.
- Optional – Change the polling interval
- Optional – Change the sampling rate
- Enable sFlow globally
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports

NOTE

If you change the router ID or other IP address value that sFlow uses for its `agent_address`, you need to disable and then re-enable sFlow to cause the feature to use the new source address.

Specifying the collector

sFlow exports traffic statistics to an external collector. You can specify up to four collectors. You can specify more than one collector with the same IP address if the UDP port numbers are unique. You can have up to four unique combinations of IP addresses and UDP port numbers.

IPv4 devices

To specify an sFlow collector on an IPv4 device, enter a command such as the following.

```
PowerConnect(config)#sflow destination 10.10.10.1
```

This command specifies a collector with IPv4 address 10.10.10.1, listening for sFlow data on UDP port 6343.

Syntax: `[no] sflow destination <ip-addr> [<dest-udp-port>]`

The `<ip-addr>` parameter specifies the IP address of the collector.

The `<dest-udp-port>` parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343.

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the device that sent the data. Refer to “[Source address](#)” on page 1081.

IPv6 devices

To specify an sFlow collector on an IPv6 device, enter a command such as the following.

```
PowerConnect(config)#sflow destination ipv6 2003:0:0::0b:02a
```

This command specifies a collector with IPv6 address 2003:0::0b:02a, listening for sFlow data on UDP port 6343.

Syntax: `[no] sflow destination ipv6 <ip-addr> [<dest-udp-port>]`

The `<ip-addr>` parameter specifies the IP address of the collector.

The `<dest-udp-port>` parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343.

If the IPv6 address you specify is a link-local address on a Layer 3 switch, you must also specify the **outgoing-interface ethernet** `<port-num>` or the **ve** `<port-num>`. This identifies the outgoing interface through which the sampled packets will be sent.

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the device that sent the data. Refer to “[Source address](#)” on page 1081.

Changing the polling interval

The polling interval defines how often sFlow byte and packet counter data for a port are sent to the sFlow collectors. If multiple ports are enabled for sFlow, the device staggers transmission of the counter data to smooth performance. For example, if sFlow is enabled on two ports and the polling interval is 20 seconds, the device sends counter data every ten seconds. The counter data for one of the ports are sent after ten seconds, and counter data for the other port are sent after an additional ten seconds. Ten seconds later, new counter data for the first port are sent. Similarly, if sFlow is enabled on five ports and the polling interval is 20 seconds, the device sends counter data every four seconds.

The default polling interval is 20 seconds. You can change the interval to a value from 1 to any higher value. The interval value applies to all interfaces on which sFlow is enabled. If you set the polling interval to 0, counter data sampling is disabled.

To change the polling interval, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#sflow polling-interval 30
```

Syntax: `[no] sflow polling-interval <secs>`

The `<secs>` parameter specifies the interval and can be from 1 to any higher value. The default is 20 seconds. If you specify 0, counter data sampling is disabled.

Changing the sampling rate

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets.

You can change the default (global) sampling rate. You also can change the rate on an individual port, overriding the default sampling rate of 512. With a sampling rate of 512, on average, one in every 512 packets forwarded on an interface is sampled.

Configuration considerations

The sampling rate is a fraction in the form $1/N$, meaning that, on average, one out of every N packets will be sampled. The **sflow sample** command at the global level or port level specifies N , the denominator of the fraction. Thus a higher number for the denominator means a lower sampling rate since fewer packets are sampled. Likewise, a lower number for the denominator means a higher sampling rate because more packets are sampled. For example, if you change the denominator from 512 to 128, the sampling rate increases because four times as many packets will be sampled.

NOTE

Dell recommends that you do not change the denominator to a value lower than the default. Sampling requires CPU resources. Using a low denominator for the sampling rate can cause high CPU utilization.

Configured rate and actual rate

When you enter a sampling rate value, this value is the **configured rate**. The software rounds the value you enter to the next higher odd power of 2 to obtain the **actual rate**. This value becomes the actual sampling rate. For example, if the configured sampling rate is 1000, then the actual rate is 2048 and 1 in 2048 packets are sampled by the hardware.

Change to global rate

If you change the global sampling rate, the change is applied to all sFlow-enabled ports **except** those ports on which you have already explicitly set the sampling rate. For example, suppose that sFlow is enabled on ports 1, 2, and 3. If you configure the sampling rate on port 1 but leave the other two ports using the default rate, then a change to the global sampling rate applies to ports 2 and 3 but not port 1. sFlow assumes that you want to continue using the sampling rate you explicitly configured on an individual port even if you globally change the sampling rate for the other ports.

Module rate

While different ports on a module may be configured to have different sampling rates, the hardware for the module will be programmed to take samples at a single rate (the module sampling rate). The module sampling rate will be the highest sampling rate (i.e. lowest number) configured for any of the ports on the module.

When ports on a given module are configured with different sampling rates, the CPU discards some of the samples supplied by the hardware for ports with configured sampling rates which are lower than the module sampling rate. This is referred to as subsampling, and the ratio between the port sampling rate and the module sampling rate is known as the subsampling factor. For example, if sFlow enabled on ports 2 and 8, and port 2 is using the default sampling rate of 512, and port 8 is configured explicitly for a rate of 2048, then the module sampling rate will be 512 because this is this highest port sampling rate (lowest number). The subsampling factor for port 2 will be 1, meaning that every sample taken by the hardware will be exported, while the subsampling factor for port 8 will be 4, meaning that one out of every four samples taken by the hardware will be exported. Whether a port's sampling rate is configured explicitly, or whether it uses the global default setting, has no effect on the calculations.

You do not need to perform any of these calculations to change a sampling rate. For simplicity, the syntax information in this section lists the valid sampling rates. In addition, the software will round the value you enter up to the nearest value listed. You can display the rates you entered (the configured rates) as well as the rates rounded up to by the software (the actual rates) for the default sampling rate, module rates, and all sFlow-enabled ports by entering the **show sflow** command. Refer to “[Displaying sFlow information](#)” on page 1087.

Sampling rate for new ports

When you enable sFlow on a port, the port's sampling rate is set to the global default sampling rate. This also applies to ports on which you disable and then re-enable sFlow. The port does not retain the sampling rate it had when you disabled sFlow on the port, even if you had explicitly set the sampling rate on the port.

Changing the default sampling rate

To change the default (global) sampling rate, enter a command such as the following at the global CONFIG level of the CLI.

```
PowerConnect(config)#sflow sample 2048
```

Syntax: [no] sflow sample <num>

The <num> parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter to the next higher odd power of 2. This value becomes the actual default sampling rate and is one of the following:

- 2
- 8
- 32
- 128
- 512
- 2048
- 4096
- 8192
- 32768
- 131072
- 524288
- 2097152
- 8388608
- 33554432
- 134217728
- 536870912
- 2147483648

For example, if the configured sampling rate is 1000, then the actual rate is 2048 and 1 in 2048 packets are sampled by the hardware.

Changing the sampling rate of a module

You cannot change a module sampling rate directly. You can change a module sampling rate only by changing the sampling rate of a port on that module.

Changing the sampling rate on a port

You can configure an individual port to use a different sampling rate than the global default sampling rate. This is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To change the sampling rate on an individual port, enter a command such as the following at the configuration level for the port.

```
PowerConnect(config-if-1)#sflow sample 8192
```

Syntax: [no] sflow sample <num>

The <num> parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter up to the next odd power of 2. The actual sampling rate becomes one of the values listed in [“Changing the default sampling rate”](#).

Changing the sampling rate for a trunk port

You can configure an individual static trunk port to use a different sampling rate than the global default sampling rate. This feature is also supported on LACP trunk ports is not supported on LACP trunk ports on PowerConnect B-Series TI24X devices. This feature is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To change the sampling rate on an individual trunk port, enter commands such as the following.

```
PowerConnect(config)#trunk e 1 to 8
PowerConnect(config-trunk-1-8)#config-trunk-ind
PowerConnect(config-trunk-1-8)#sflow-subsampling e2 8192
```

Syntax: [no] sflow-subsampling ethernet<portnum> <num>

OR

Syntax: [no] sflow sample ethernet <portnum> <num>

The <num> parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter up to the next odd power of 2. The actual sampling rate becomes one of the values listed in [“Changing the default sampling rate”](#).

Enabling sFlow forwarding

sFlow exports data only for the interfaces on which you enable sFlow forwarding. You can enable sFlow forwarding on Ethernet interfaces.

To enable sFlow forwarding, perform the following:

- Globally enable the sFlow feature
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports

NOTE

Before you enable sFlow, make sure the device has an IP address that sFlow can use as its source address. Refer to [“Source address”](#) on page 1081 for the source address requirements.

NOTE

When you enable sFlow forwarding on an 802.1X-enabled interface, the samples taken from the interface include the username used to obtain access to either or both the inbound and outbound ports, if that information is available. For information about 802.1X, refer to [Chapter 28, “Configuring 802.1X Port Security”](#).

Command syntax

This section shows how to enable sFlow forwarding.

Globally enabling sFlow forwarding

To enable sFlow forwarding, you must first enable it on a global basis, then on individual interfaces or trunk ports, or both.

To globally enable sFlow forwarding, enter the following command.

```
PowerConnect(config)#sflow enable
```

You can now enable sFlow forwarding on individual ports as described in the next two sections.

Syntax: [no] sflow enable

Enabling sFlow forwarding on individual interfaces

To enable sFlow forwarding enter commands such as the following.

```
PowerConnect(config)#sflow enable
PowerConnect(config)#interface ethernet 1 to 8
PowerConnect(config-mif-1-8)#sflow-forwarding
```

These commands globally enable sFlow, then enable sFlow forwarding on Ethernet ports 1 – 8. You must use both the **sflow enable** and **sflow-forwarding** commands to enable the feature.

Syntax: [no] sflow enable

Syntax: [no] sflow-forwarding

Displaying sFlow information

To display sFlow configuration information and statistics, enter the following command at any level of the CLI.

```
PowerConnect#show sflow
sFlow services are enabled.
sFlow agent IP address: 123.123.123.1
4 collector destinations configured:
Collector IP 192.168.4.204, UDP 6343
Collector IP 192.168.4.200, UDP 6333
Collector IP 192.168.4.202, UDP 6355
Collector IP 192.168.4.203, UDP 6565
Polling interval is 0 seconds.
Configured default sampling rate: 1 per 512 packets
Actual default sampling rate: 1 per 512 packets
10552 UDP packets exported
24127 sFlow samples collected.
sFlow ports: ethe 2 to 12 ethe 15 ethe 25 to 26 ethe 1 ethe 10 to
20 ethe 1 ethe 4
```

Port Sampling Rates

```
-----
Port 4, configured rate=512, actual rate=512, Subsampling factor=1
Port 1, configured rate=512, actual rate=512, Subsampling factor=1
Port 20, configured rate=3000, actual rate=8192, Subsampling factor=16
Port 19, configured rate=512, actual rate=512, Subsampling factor=1
Port 18, configured rate=512, actual rate=512, Subsampling factor=1
Port 17, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 16, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 15, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 14, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 13, configured rate=512, actual rate=512, Subsampling factor=1
Port 12, configured rate=512, actual rate=512, Subsampling factor=1
Port 11, configured rate=512, actual rate=512, Subsampling factor=1
Port 10, configured rate=512, actual rate=512, Subsampling factor=1
Port 1, configured rate=10000, actual rate=32768, Subsampling factor=1
Port 26, configured rate=512, actual rate=512, Subsampling factor=1
Port 25, configured rate=512, actual rate=512, Subsampling factor=1
Port 15, configured rate=512, actual rate=512, Subsampling factor=1
Port 12, configured rate=512, actual rate=512, Subsampling factor=1
```

...continued on next page...

...continued from previous page...

Port 11, configured rate=512, actual rate=512, Subsampling factor=1
 Port 10, configured rate=512, actual rate=512, Subsampling factor=1
 Port 9, configured rate=512, actual rate=512, Subsampling factor=1
 Port 8, configured rate=512, actual rate=512, Subsampling factor=1
 Port 7, configured rate=1000, actual rate=2048, Subsampling factor=4
 Port 6, configured rate=512, actual rate=512, Subsampling factor=1
 Port 5, configured rate=512, actual rate=512, Subsampling factor=1
 Port 4, configured rate=512, actual rate=512, Subsampling factor=1
 Port 3, configured rate=512, actual rate=512, Subsampling factor=1
 Port 2, configured rate=1000, actual rate=2048, Subsampling factor=4

Syntax: show sflow

This command shows the following information.

TABLE 172 sFlow information

This field...	Displays...
sFlow services	The feature state, which can be one of the following: <ul style="list-style-type: none"> • disabled • enabled
sFlow agent IP address	The IP address that sFlow is using in the agent_address field of packets sent to the collectors. Refer to " Source address " on page 1081.
Collector	The collector information. The following information is displayed for each collector: <ul style="list-style-type: none"> • IP address • UDP port If more than one collector is configured, the line above the collectors indicates how many have been configured.
Polling interval	The port counter polling interval.
Configured default sampling rate	The configured global sampling rate. If you changed the global sampling rate, the value you entered is shown here. The actual rate calculated by the software based on the value you entered is listed on the next line, "Actual default sampling rate".
Actual default sampling rate	The actual default sampling rate.
UDP packets exported	The number of sFlow export packets the device has sent. NOTE: Each UDP packet can contain multiple samples.
sFlow samples collected	The number of sampled packets that have been sent to the collectors.
sFlow ports	The ports on which you enabled sFlow.
Module Sampling Rates	The configured and actual sampling rates for each module. If a module does not have any sFlow-enabled ports, the rates are listed as 0.
Port Sampling Rates	The configured and actual sampling rates for each sFlow-enabled port. The Subsampling factor indicates how many times the sampling rate of the port's module is multiplied to achieve the port's sampling rate. Because of the way the actual sampling rates are computed, the Subsampling factors are always whole numbers.

Clearing sFlow statistics

To clear the UDP packet and sFlow sample counters in the **show sflow** display, enter the following command.

A Configuring a utilization list for an uplink port

```
PowerConnect#clear statistics
```

Syntax: clear statistics

This command clears the values in the following fields of the **show sflow** display:

- UDP packets exported
- sFlow samples collected

NOTE

This command also clears the statistics counters used by other features.

Configuring a utilization list for an uplink port

You can configure uplink utilization lists that display the percentage of a given uplink port bandwidth that is used by a specific list of downlink ports. The percentages are based on 30-second intervals of RMON packet statistics for the ports. Both transmit and receive traffic is counted in each percentage.

NOTE

This feature is intended for ISP or collocation environments in which downlink ports are dedicated to various customers' traffic and are isolated from one another. If traffic regularly passes between the downlink ports, the information displayed by the utilization lists does not provide a clear depiction of traffic exchanged by the downlink ports and the uplink port.

Each uplink utilization list consists of the following:

- Utilization list number (1, 2, 3, or 4)
- One or more uplink ports
- One or more downlink ports

Each list displays the uplink port and the percentage of that port bandwidth that was utilized by the downlink ports over the most recent 30-second interval.

You can configure up to four bandwidth utilization lists.

Command syntax

To configure an uplink utilization list, enter commands such as the following. The commands in this example configure a link utilization list with port 1 as the uplink port and ports 2 and 3 as the downlink ports.

```
PowerConnect(config)#relative-utilization 1 uplink eth 1 downlink eth 2 to 3
PowerConnect(config)#write memory
```

Syntax: [no] **relative-utilization** <num> **uplink ethernet** <portnum> [**to** <portnum>] **downlink ethernet** <portnum> [**to**<portnum>]

The <num> parameter specifies the list number. You can configure up to four lists. Specify a number from 1 - 4.

The **uplink ethernet** parameters and the port numbers you specify after the parameters indicate the uplink ports.

The **downlink ethernet** parameters and the port numbers you specify after the parameters indicate the downlink ports.

Displaying utilization percentages for an uplink

After you configure an uplink utilization list, you can display the list to observe the percentage of the uplink bandwidth that each of the downlink ports used during the most recent 30-second port statistics interval. The number of packets sent and received between the two ports is listed, as well as the ratio of each individual downlink port packets relative to the total number of packets on the uplink.

To display an uplink utilization list, enter a command such as the following at any level of the CLI.

```
PowerConnect#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
    2:60    3:40
```

In this example, ports 2 and 3 are sending traffic to port 1. Port 2 and port 3 are isolated (not shared by multiple clients) and typically do not exchange traffic with other ports except for the uplink port, 1.

Syntax: `show relative-utilization <num>`

The `<num>` parameter specifies the list number.

NOTE

The example above represents a pure configuration in which traffic is exchanged only by ports 2 and 1, and by ports 3 and 1. For this reason, the percentages for the two downlink ports equal 100%. In some cases, the percentages do not always equal 100%. This is true in cases where the ports exchange some traffic with other ports in the system or when the downlink ports are configured together in a port-based VLAN.

In the following example, ports 2 and 3 are in the same port-based VLAN.

```
PowerConnect#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
    2:100   3:100
```

Here is another example showing different data for the same link utilization list. In this example, port 2 is connected to a hub and is sending traffic to port 1. Port 3 is unconnected.

```
PowerConnect#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 2996
packet count ratio (%)
    2:100   3:---
```

A Configuring a utilization list for an uplink port

Software Specifications

IEEE compliance

Dell devices support the following standards.

TABLE 173 IEEE compliance

Standard	Description	PowerConnect B-Series T124X
802.1AB	Station and Media Access Control Connectivity Discovery Also supports TIA-1057, Telecommunications – IP Telephony Infrastructure -- Link Layer Discovery Protocol (LLDP) for Media Endpoint Devices	Yes
802.1d	Ethernet Bridging	Yes
802.1D	MAC Bridges	Yes
802.1p	Mapping to Priority Queue	Yes
802.1p/q	VLAN Tagging	Yes
802.1s	Multiple Spanning Tree	Yes
802.1w	Rapid Spanning Tree	Yes
802.3	10Base-T	Yes
802.3	MAU MIB (RFC 2239)	Yes
802.3ab	1000Base-T	Yes
802.3ad	Link Aggregation (Dynamic and Static) and Trunk Groups	Yes
802.3ae	10-Gigabit Ethernet	Yes
802.3u	100Base-TX, 100Base-FX, 100Base_LX	100Base_TX only
802.3z	1000Base-SX, 1000Base-LX	Yes
802.3x	Flow Control	Yes

RFC support

The following table lists the RFCs supported by Dell devices.

NOTE

Some devices support only a subset of the RFCs. For example, Layer 2 Switches do not support router-specific RFCs. For a list of features supported on your device, refer to the data sheet or the software release notes for the version of software running on your device.

TABLE 174 Dell RFC support

RFC number	Protocol or Standard	PowerConnect B-Series T124X
768	User Datagram Protocol (UDP)	Yes
783	Trivial File Transfer Protocol (TFTP)	Yes
791	Internet Protocol (IP)	Yes
792	Internet Control Message Protocol (ICMP)	Yes
793	Transmission Control Protocol (TCP)	Yes
826	Ethernet Address Resolution Protocol (ARP)	Yes
854, 855, and 857	Telnet	Yes
894	IP over Ethernet frames	Yes
903	Reverse ARP (RARP)	Yes
906	Bootstrap loading using TFTP	Yes
919	Broadcast Internet datagrams	Yes
920	Domain requirements	Yes
922	Broadcast Internet datagrams in the presence of subnets	Yes
950	Internet standard subnetting procedure	Yes
951	Bootstrap Protocol (BootP)	Yes
1027	Proxy ARP	Yes
1042	IP datagrams over IEEE 802 networks (for Ethernet)	Yes
1112	Internet Gateway Management Protocol (IGMP) version 1	Yes
1155	Structure and Identification of Management Information (SMI)	Yes
1157	Simple Network Management Protocol (SNMP) version 1	Yes
1191	Path MTU Discovery	No
1212	Concise MIB Definitions	Yes
1213	MIB II Definitions	Yes
1215	SNMP generic traps	Yes
1340	Assigned numbers (where applicable)	Yes
1398	Ethernet-Like MIB	Yes
1492	An Access Control Protocol, Sometimes Called TACACS	Yes
1493	Bridge MIB (excluding filtering of objects)	Yes
1516	Repeater MIB	Yes

TABLE 174 Dell RFC support (Continued)

RFC number	Protocol or Standard	PowerConnect B-Series T124X
1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy	Yes
1541	Dynamic Host Configuration Protocol (DHCP)	Yes
1542	BootP Extensions	Yes
1573	SNMP MIB II	Yes
1591	Domain Name System (DNS) Structure and Delegation	Yes
1643	Ethernet Interface MIB	Yes
1757	Remote Monitoring (RMON) groups 1, 2, 3, 9	Yes
1905	Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)	Yes
1906	Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2)	Yes
1965	Autonomous System Configurations for BGP4	No
2011	SNMPv2 Management Information Base for the Internet Protocol using SMIv2	Yes
2012	SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2	Yes
2013	SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2	Yes
2030	SNTP	Yes
2138	Remote Authentication Dial In User Server (RADIUS)	Yes
2139	RADIUS Accounting	Yes
2236	Internet Gateway Management Protocol (IGMP) version 2	Yes
2239	802.3 Medium Attachment Units (MAUs) using SMIv2	Yes
2336	IGMP version 2	Yes
2482	Language Tagging in Unicode Plain Text	Yes
2544	Benchmarking Methodology for Network Interconnect Devices	Yes
2570	Introduction to version 3 of the Internet-standard Network Management Framework	Yes

TABLE 174 Dell RFC support (Continued)

RFC number	Protocol or Standard	PowerConnect B-Series T124X
2571	An Architecture of Describing SNMP Management Frameworks	Yes
2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)	Yes
2573	SNMP version 3 Applications	Yes
2574	User-based Security (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)	Yes
2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)	Yes
2576	Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework	Yes
2578	Structure of Management Information Version 2 (SMIv2)	Yes
2579	Textual Conventions for SMIv2	Yes
2580	Conformance Statements for SMIv2	Yes
2665	Ethernet Like MIB (incorporates RFC 1398)	Yes
2674	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions	Yes
2865	Remote Authentication Dial In User Service (RADIUS)	Yes
2866	RADIUS Accounting	Yes
2869	RADIUS Extensions	Yes
2889	Benchmarking Methodology for LAN Switching Devices	Yes
2933	Internet Group Management Protocol MIB	Yes
3176	InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks	Yes
3376	Internet Gateway Management Protocol (IGMP) version 3	Yes
3411	Simple Network Management Protocol (SNMP) Management Frameworks	Yes
3412	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP V3)	Yes
3413	Simple Network Management Protocol (SNMP) Applications	Yes

TABLE 174 Dell RFC support (Continued)

RFC number	Protocol or Standard	PowerConnect B-Series T124X
3414	User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMP V3)	Yes
3415	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)	Yes
3416	Version 2 of the Protocol Operations for the SNMP	Yes
3418	Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)	Yes
3584	Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework	Yes
3918	Benchmarking Methodology for IP Multicast	Yes
4188	Definitions of Managed Objects for Bridges	Yes
4251	The Secure Shell (SSH) Protocol Architecture	Yes
4252	The Secure Shell (SSH) Authentication Protocol	Yes
4253	The Secure Shell (SSH) Transport Protocol	Yes
4254	The Secure Shell (SSH) Connection Protocol	Yes
4330	Simple Network Time Protocol (SNTP) version 4	Yes
	Authentication, Authorization, and Accounting (AAA)	Yes
	Bi-level access mode (standard and EXEC level)	Yes
	DNS Client	Yes
	IGMP Proxy	Yes
	IGMP Snooping (versions 1, 2, and 3)	Yes
	Integrated standard-based Command Line Interface (CLI)	Yes
	MRP	Yes
	Protection of Denial of Service Attack, such as TCP SYN or Smurf Attacks	Yes
	PVST/PVST+/PVRST	Yes
	RMON and Windows NT	Yes
	Secure Copy (SCP)	Yes
	SSH V 1.5	Yes

B Internet drafts

TABLE 174 Dell RFC support (Continued)

RFC number	Protocol or Standard	PowerConnect B-Series T124X
	SSH V 2	Yes
	SNMP V1, V2c, and V3	Yes
	TACACS/TACACS+	Yes
	TELNET and SSH V1	Yes
	UDLD	Yes
	Username or Password (challenge and response)	Yes

Internet drafts

In addition to the RFCs listed in [“RFC support”](#) on page 1093, Dell devices support the following Internet drafts:

- draft-ietf-magma-igmp-proxy.txt
- TACACS+ Protocol version 1.78