

53-1002507-01
15 December 2011



Converged Enhanced --- Ethernet

Administrator's Guide

Supporting Fabric OS v7.0.1

BROCADE

Copyright © 2009-2011 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and AnyIO, Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

The product described by this document may contain "open source" software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 - 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1001258-01	New document	March 2009
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1001336-02	Updated for Fabric OS v6.3.0. Added new chapters for standard configurations and port authentication	November 2009
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1002507-01	Updated for Fabric OS v7.0.1.	December 2011

Title	Publication number	Summary of changes	Date
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1001761-01	Updated for Fabric OS v7.0.1. Added chapter for IGMP.	March 2010
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1002061-01	Updated to support iSCSI TLV configuration.	October 2010
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1002163-01	Updated for Fabric OS v7.0.0	April 2011
<i>Converged Enhanced Ethernet Administrator's Guide</i>	53-1002163-02	Updated Appendix A, Brocade 8000 Replacement.	June 2011

Contents

About This Document

In this chapterxvii
How this document is organizedxvii
Supported hardware and software	xviii
What's new in this document	xviii
Document conventions	xviii
Text formatting	xviii
Command syntax conventions	xix
Notes, cautions, and warnings	xix
Key terms	xx
Notice to the readerxx
Additional informationxx
Brocade resources	xx
Other industry resources	xxi
Getting technical help	xxi
Document feedbackxxii

Chapter 1

Introducing FCoE

In this chapter	1
FCoE terminology	1
FCoE overview	1
FCoE hardware	2
Layer 2 Ethernet overview	3
Layer 2 forwarding	3
VLAN tagging	4
Loop-free network environment	4
Frame classification (incoming)	5
Congestion control and queuing	5
Access control	7
Trunking	7
Flow Control	8

	FCoE Initialization Protocol	8
	FIP discovery	8
	FIP login	9
	FIP logout	9
	FCoE login	10
	FCoE logout	10
	Loginconf	10
	Name server	11
	FC zoning	11
	Registered State Change Notification (RSCN)	11
	FCoE queuing	12
	Upgrade and downgrade considerations for FCoE	12
Chapter 2	Using the CEE CLI	
	In this chapter	13
	Management Tools	13
	CEE Command Line Interface	13
	Saving your configuration changes	14
	CEE CLI RBAC permissions	14
	Accessing the CEE CLI through the console or Telnet	14
	Accessing the CEE CLI from the Fabric OS shell	15
	CEE CLI command modes	15
	CEE CLI keyboard shortcuts	17
	Using the do command as a shortcut	18
	Displaying CEE CLI commands and command syntax	18
	CEE CLI command completion	19
	CEE CLI command output modifiers	19
Chapter 3	Configuration management	
	In this chapter	21
	Configuration management tasks	21
	CEE configuration guidelines and restrictions	21
	Display the running configuration file	22
	Saving the running configuration file	22
	Loading the startup configuration file	22
	Erasing the startup configuration file	22
	Archiving the running configuration file	22
	Restore an archived running configuration file	22
	Archiving the startup configuration file	23
	Restore an archived startup configuration file	23
	Flash file management commands	23
	Debugging and logging commands	23
Chapter 4	FCoE Configuration Management	
	In this chapter	25

Managing the FCoE configuration	25
FCoE configuration guidelines and restrictions	25
Clearing logins	26
Displaying FCoE configuration-related information	26
Configuring the FCoE login information	26
Enabling or disabling login configuration management	26
Managing the current configuration transaction	27
Cleaning up login groups and VN_port mappings	27
Displaying the FCoE login configuration	27
Saving the current FCoE configuration	28
Creating an FCoE login group	28
Modifying the FCoE login group device list	29
Deleting an FCoE login group	29
Renaming an FCoE login group	29

Chapter 5 Initial FCoE and CEE Configuration

In this chapter	31
Overview	31
End to End FCoE using FC ISLs	31
FCoE Logical Port Numbering	32
FCoE port-to-front end TenGigabit port mapping	32
High availability	33
Hot plug	33
Control Processor failover	33
High availability reference	34
Configuring the FCoE interfaces	35
Configuring the FCoE VLAN	35
Assigning an FCoE map onto an interface	36
Configuring DCBX	37
Configuring Spanning Tree Protocol	38
Configuring VLAN membership	38
CEE map configuration	39
Defining a Priority Group Table map	40
Defining a Priority-Table map	41
Applying a CEE provisioning map to an interface	41
Verifying the CEE maps	41
CEE audit logs	42

Chapter 6 Configuring VLANs

In this chapter	45
VLAN overview	45
Ingress VLAN filtering	45
VLAN configuration guidelines and restrictions	47
Default VLAN configuration	47

VLAN configuration and management	48
Enabling and disabling an interface port	48
Configuring the MTU on an interface port	48
Creating a VLAN interface	48
Enabling STP on a VLAN	49
Disabling STP on a VLAN	49
Configuring a VLAN interface to forward FCoE traffic	49
Configuring an interface port as a Layer 2 switch port	50
Configuring an interface port as an access interface	50
Configuring an interface port as a trunk interface	50
Disabling a VLAN on a trunk interface	51
Configuring an interface port as a converged interface	51
Disabling a VLAN on a converged interface	52
Configuring VLAN classifier rules	52
Configuring a VLAN classifier rule	53
Configuring MAC address-based VLAN classifier rules	53
Deleting a VLAN classifier rule	53
Creating a VLAN classifier group and adding rules	53
Activating a VLAN classifier group with an interface port	53
Clearing VLAN counter statistics	54
Displaying VLAN information	54
Configuring the MAC address table	54
Specifying or disabling the aging time for MAC addresses	54
Adding static addresses to the MAC address table	55

Chapter 7

Configuring STP, RSTP, and MSTP

In this chapter	57
STP overview	57
Configuring STP	58
RSTP overview	59
MSTP overview	61
Configuring MSTP	62
STP, RSTP, and MSTP configuration guidelines and restrictions	63
Default STP, RSTP, and MSTP configuration	63

STP, RSTP, and MSTP configuration and management	64
Enabling STP, RSTP, or MSTP	64
Disabling STP, RSTP, or MSTP	65
Shutting down STP, RSTP, or MSTP globally	65
Specifying the bridge priority.	65
Specifying the bridge forward delay	66
Specifying the bridge maximum aging time.	66
Enabling the error disable timeout timer	67
Specifying the error disable timeout interval	67
Specifying the port-channel path cost	67
Specifying the bridge hello time (STP and RSTP).	68
Specifying the transmit hold count (RSTP and MSTP).	68
Enabling Cisco interoperability (MSTP).	69
Disabling Cisco interoperability (MSTP)	69
Mapping a VLAN to an MSTP instance	69
Specifying the maximum number of hops for a BPDU (MSTP)	70
Specifying a name for an MSTP region.	70
Specifying a revision number for an MSTP configuration	70
Flushing MAC addresses (RSTP and MSTP).	71
Clearing spanning tree counters.	71
Clearing spanning tree-detected protocols	71
Displaying STP, RSTP, and MSTP-related information	72
Configuring STP, RSTP, or MSTP on CEE interface ports	72
Enabling automatic edge detection	72
Configuring the path cost	72
Enabling a port (interface) as an edge port.	73
Enabling the guard root.	73
Specifying the MSTP hello time.	74
Specifying restrictions for an MSTP instance	74
Specifying a link type.	75
Enabling port fast (STP).	75
Specifying the port priority	75
Restricting the port from becoming a root port	76
Restricting the topology change notification	76
Enabling spanning tree	76
Disabling spanning tree.	77

Chapter 8

Configuring Link Aggregation

In this chapter	79
Link aggregation overview	79
Link Aggregation Group configuration	79
Link Aggregation Control Protocol.	82
Dynamic link aggregation	82
Static link aggregation.	82
Brocade-proprietary aggregation	82
LAG distribution process	83
LACP configuration guidelines and restrictions	83
Default LACP configuration.	83

	LACP configuration and management	83
	Enabling LACP on an CEE interface	83
	Enabling LACP on an interface in Layer 2 mode	84
	Configuring the LACP system priority	84
	Configuring the LACP timeout period on a CEE interface	84
	Clearing LACP counter statistics on a LAG	85
	Clearing LACP counter statistics on all LAG groups	85
	Displaying LACP information	85
	LACP troubleshooting tips	85
Chapter 9	Configuring LLDP	
	In this chapter	87
	LLDP overview	87
	Layer 2 topology mapping	88
	DCBX overview	89
	Enhanced Transmission Selection (ETS)	90
	Priority Flow Control (PFC)	90
	DCBX interaction with other vendor devices	91
	LLDP configuration guidelines and restrictions	91
	Default LLDP configuration	91
	LLDP configuration and management	92
	Enabling LLDP globally	92
	Disabling and resetting LLDP globally	92
	Configuring LLDP global command options	92
	Configuring LLDP interface-level command options	98
	Clearing LLDP-related information	98
	Displaying LLDP-related information	98
Chapter 10	Configuring ACLs	
	In this chapter	101
	ACL overview	101
	Default ACL configuration	102
	ACL configuration guidelines and restrictions	102
	ACL configuration and management	102
	Creating a standard MAC ACL and adding rules	102
	Creating an extended MAC ACL and adding rules	103
	Modifying MAC ACL rules	103
	Removing a MAC ACL	104
	Reordering the sequence numbers in a MAC ACL	104
	Applying a MAC ACL to a CEE interface	105
	Applying a MAC ACL to a VLAN interface	105
Chapter 11	Configuring QoS	
	In this chapter	107

	QoS overview	107
	Rewriting	108
	Queueing	108
	User-priority mapping.	108
	Traffic class mapping.	111
	Congestion control	113
	Tail drop	113
	Ethernet pause.	114
	Ethernet Priority Flow Control	115
	Multicast rate limiting.	116
	Scheduling.	117
	Strict priority scheduling	117
	Deficit weighted round robin scheduling	117
	Traffic class scheduling policy.	118
	Multicast queue scheduling	119
	CEE map configuration	119
Chapter 12	Configuring 802.1x Port Authentication	
	In this chapter	123
	802.1x protocol overview	123
	802.1x configuration guidelines and restrictions.	123
	802.1x authentication configuration tasks.	124
	Configure authentication between the switch and CNA or NIC.	124
	Interface-specific administrative tasks for 802.1x.	124
	Configuring 802.1x on specific interface ports	124
	Configuring 802.1x timeouts on specific interface ports	125
	Configuring 802.1x re-authentication on specific interface ports.	125
	Disabling 802.1x on specific interface ports.	125
Chapter 13	Configuring IGMP	
	In this chapter	127
	About IGMP	127
	Active IGMP snooping	127
	Multicast routing	128
	Configuring IGMP	128
	Configuring IGMP snooping querier.	129
	Configuring IGMP snooping mrouter	129
	Monitoring IGMP	130
Chapter 14	Configuring RMON	
	In this chapter	131

RMON overview131
RMON configuration and management.131
 Default RMON configuration131
 Configuring RMON group statistics collection131
 Configuring RMON settings.....132
 Configuring RMON events132
 Configuring RMON history collection133

Appendix A Brocade 8000 Replacement

Index

Figures

Figure 1	Multiple switch fabric configuration	3
Figure 2	CEE CLI command mode hierarchy	15
Figure 3	Deploying FCoE across multiple hops of FC ISLs	32
Figure 4	Ingress VLAN filtering	46
Figure 5	Configuring LAGs for a top-of-the-rack CEE switch—Example 1	81
Figure 6	Configuring LAGs for a top-of-the-rack CEE switch—Example 2	81
Figure 7	Queue depth	114
Figure 8	Strict priority schedule — two queues	117
Figure 9	WRR schedule — two queues	117
Figure 10	Strict priority and Weighted Round Robin scheduler	118

Tables

Table 1	Command syntax conventions.	xix
Table 2	FCoE terminology	1
Table 3	CEE RBAC permissions.	14
Table 4	CEE CLI command modes	16
Table 5	CEE CLI keyboard shortcuts.	17
Table 6	CEE CLI command output modifiers	19
Table 7	CEE Flash memory file management commands.	23
Table 8	Debugging and logging commands.	24
Table 9	Protocol timeout values for non-disruptive failover	33
Table 10	State preservation during CP failover	34
Table 11	Default CEE Priority Group Table configuration	39
Table 12	Default CEE priority table configuration	40
Table 13	Format of the CEE audit logs	42
Table 14	Default VLAN configuration	47
Table 15	STP versus RSTP state comparison.	59
Table 16	Default STP, RSTP, and MSTP configuration	63
Table 17	Default MSTP configuration.	64
Table 18	Default 10-Gigabit Ethernet CEE interface-specific configuration	64
Table 19	Default LACP configuration	83
Table 20	ETS priority grouping of IPC, LAN, and SAN traffic	90
Table 21	Default LLDP configuration	91
Table 22	Default MAC ACL configuration	102
Table 23	Default priority value of untrusted interfaces.	109
Table 24	IEEE 802.1Q default priority mapping.	109
Table 25	Default user priority for unicast traffic class mapping.	111
Table 26	Default user priority for multicast traffic class mapping	112
Table 27	Supported scheduling configurations	118
Table 28	Multicast traffic class equivalence mapping	119
Table 29	Default CEE Priority Group Table configuration	120
Table 30	Default CEE priority table	121

About This Document

In this chapter

• How this document is organized	xvii
• Supported hardware and software.....	xviii
• What’s new in this document.....	xviii
• Document conventions	xviii
• Notice to the reader	xx
• Additional information.....	xx
• Getting technical help.....	xxi
• Document feedback	xxii

How this document is organized

This document is organized to help you find the information that you want as quickly and easily as possible.

The document contains the following components:

- [Chapter 1, “Introducing FCoE,”](#) provides an overview of Fibre Channel over Ethernet (FCoE) on the Brocade FCoE hardware.
- [Chapter 2, “Using the CEE CLI,”](#) describes the Converged Enhanced Ethernet (CEE) CLI.
- [Chapter 3, “Configuration management,”](#) describes how to perform the administrative tasks required by the Brocade FCoE hardware.
- [Chapter 4, “FCoE Configuration Management,”](#) describes how to configure FCoE using the Fabric OS CLI.
- [Chapter 5, “Initial FCoE and CEE Configuration,”](#) describes some basic switch configurations for command SAN and LAN environments.
- [Chapter 6, “Configuring VLANs,”](#) describes how to configure VLANs.
- [Chapter 7, “Configuring STP, RSTP, and MSTP,”](#) describes how to configure the Spanning Tree Protocol (STP), Rapid STP (RSTP), and Multiple STP (MSTP).
- [Chapter 8, “Configuring Link Aggregation,”](#) describes how to configure Link Aggregation and Link Aggregation Control Protocol (LACP).
- [Chapter 9, “Configuring LLDP,”](#) describes how to configure the Link Layer Discovery Protocol (LLDP) and the Data Center Bridging (DCB) Capability Exchange Protocol (DCBX).
- [Chapter 10, “Configuring ACLs,”](#) describes how to configure Access Control Lists (ACLs).
- [Chapter 11, “Configuring QoS,”](#) describes how to configure Quality of Service (QoS).

- [Chapter 12, “Configuring 802.1x Port Authentication,”](#) describes how to configure the 802.1x Port Authentication protocol.
- [Chapter 13, “Configuring IGMP,”](#) describes how to configure IGMP snooping on the Brocade FCoE hardware.
- [Chapter 14, “Configuring RMON,”](#) describes how to configure remote monitoring (RMON).

Supported hardware and software

This document includes updated information specific to Fabric OS v7.0.1. The following hardware platforms are supported in this release of the CEE Administrator’s Guide:

- Brocade 8000

The following blades are supported by this release of the CEE Administrator’s Guide:

- Brocade FCOE10-24 blade

Within this manual, any appearance of the term “Brocade FCoE hardware” is referring to:

- Brocade 8000
- Brocade FCOE10-24 port blade

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Fabric OS v7.0.1, documenting all possible configurations and scenarios is beyond the scope of this document.

To obtain information about an OS version other than Fabric OS v7.0.1, refer to the documentation specific to that OS version.

What’s new in this document

This document has been updated for Fabric OS v7.0.1.

This document has been updated with corrections and updates for defects discovered since the release of the previous version.

For further information about new features and documentation updates for this release, refer to the release notes.

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 and operands Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
code text	Identifies CLI output Identifies command syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

Command syntax conventions

Command syntax in this manual follows these conventions:

TABLE 1 Command syntax conventions

Convention	Description
[]	Default responses to system prompts appear in square brackets.
{x y z}	A choice of required keywords appears in braces separated by vertical bars. You must select one.
screen font	Examples of information displayed on the screen.
<>	Nonprinting characters, for example passwords, appear in angle brackets.
[]	Keywords or arguments that appear within square brackets are optional.
bold face font	Commands and keywords.
<i>italic</i>	Variables for which you supply values.

Notes, cautions, and warnings

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.

ATTENTION

An Attention statement indicates potential damage to hardware or data.



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.

Key terms

For definitions specific to Brocade and Fibre Channel, see the technical glossaries on MyBrocade. See “[Brocade resources](#)” on page xx for instructions on accessing MyBrocade.

For terminology specific to this document, see “[FCoE terminology](#)” on page 1.

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

<http://www.snia.org/education/dictionary>

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.

These references are made for informational purposes only.

Corporation	Referenced Trademarks and Products
None	Not applicable

Additional information

This section lists additional Brocade and industry-specific documentation that you might find helpful.

Brocade resources

To get up-to-the-minute information, go to <http://my.brocade.com> and register at no cost for a user ID and password.

White papers, online demonstrations, and data sheets are available through the Brocade website at:

<http://www.brocade.com/products-solutions/products/index.page>

For additional Brocade documentation, visit the Brocade website:

<http://www.brocade.com>

Release notes are available on the MyBrocade website and are also bundled with the Fabric OS firmware.

Other industry resources

For additional resource information, visit the Technical Committee T11 website. This website provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, and other applications:

<http://www.t11.org>

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association website:

<http://www.fibrechannel.org>

Getting technical help

Contact your switch support supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. General Information

- Switch model
- Switch operating system version
- Software name and software version, if applicable
- Error numbers and messages received
- **supportSave** command output
- Detailed description of the problem, including the switch or fabric behavior immediately following the problem, and specific questions
- Description of any troubleshooting steps already performed and the results
- Serial console and Telnet session logs
- syslog message logs

2. Switch Serial Number

The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below:

FT00X0054E9

FT00X0054E9

The serial number label is located as follows:

- *Brocade 8000* –On the switch ID pull-out tab located inside the chassis on the port side on the left

3. World Wide Name (WWN)

Use the **licenseldShow** command to display the WWN of the chassis.

If you cannot use the **licenseldShow** command because the switch is inoperable, you can get the WWN from the same place as the serial number, except for the Brocade DCX. For the Brocade DCX, access the numbers on the WWN cards by removing the Brocade logo plate at the top of the nonport side of the chassis.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Introducing FCoE

In this chapter

- [FCoE terminology](#) 1
- [FCoE overview](#) 1
- [Layer 2 Ethernet overview](#) 3
- [FCoE Initialization Protocol](#) 8
- [FCoE queuing](#) 12
- [Upgrade and downgrade considerations for FCoE](#) 12

FCoE terminology

[Table 2](#) lists and describes the FCoE terminology used in this document.

TABLE 2 FCoE terminology

Term	Description
FCoE	Fibre Channel over Ethernet
CEE	Converged Enhanced Ethernet
VN_port	FCoE equivalent of an FC N_port
VF_port	FCoE equivalent of an FC F_port
ENode	An FCoE device that supports FCoE VN_ports (servers and target devices)
FCoE Forwarder (FCF)	An FCoE link end point that provides FC fabric services
FIP	FCoE Initialization Protocol

FCoE overview

Fibre Channel over Ethernet (FCoE) enables you to transport FC protocols and frames over Converged Enhanced Ethernet (CEE) networks. CEE is an enhanced Ethernet that enables the convergence of various applications in data centers (LAN, SAN, and HPC) onto a single interconnect technology.

FCoE provides a method of encapsulating the Fibre Channel (FC) traffic over a physical Ethernet link. FCoE frames use a unique EtherType that enables FCoE traffic and standard Ethernet traffic to be carried on the same link. FC frames are encapsulated in an Ethernet frame and sent from one FCoE-aware device across an Ethernet network to a second FCoE-aware device. The FCoE-aware devices may be FCoE end nodes (ENodes) such as servers, storage arrays, or tape drives on one end and FCoE Forwarders on the other end. FCoE Forwarders (FCFs) are switches providing FC fabric services and FCoE-to-FC bridging.

The motivation behind using CEE networks as a transport mechanism for FC arises from the desire to simplify host protocol stacks and consolidate network interfaces in data center environments. FC standards allow for building highly reliable, high-performance fabrics for shared storage, and these characteristics are what CEE brings to data centers. Therefore, it is logical to consider transporting FC protocols over a reliable CEE network in such a way that it is completely transparent to the applications. The underlying CEE fabric is highly reliable and high performing, the same as the FC SAN.

In FCoE, ENodes discover FCFs and initialize the FCoE connection through the FCoE Initialization Protocol (FIP). The FIP has a separate EtherType from FCoE. The FIP includes a discovery phase in which ENodes solicit FCFs, and FCFs respond to the solicitations with advertisements of their own. At this point, the ENodes know enough about the FCFs to log into them. The fabric login and fabric discovery (FLOGI/FDISC) for VN-to-VF port connections is also part of the FIP.

NOTE

With pre-FIP implementations, as an alternative to FIP, directly connected devices can send an FCoE-encapsulated FLOGI to the connected FCF.

FCoE hardware

At a fundamental level, FCoE is designed to enable the transport of storage and networking traffic over the same physical link. Utilizing this technology, the Brocade 8000 switch and the Brocade FCOE10-24 port blade provide a unique platform that connects servers to both LAN and SAN environments.

Within this manual, any appearance of the term “Brocade FCoE hardware” is referring to the following hardware:

- Brocade 8000 switch
- Brocade FCOE10-24 port blade

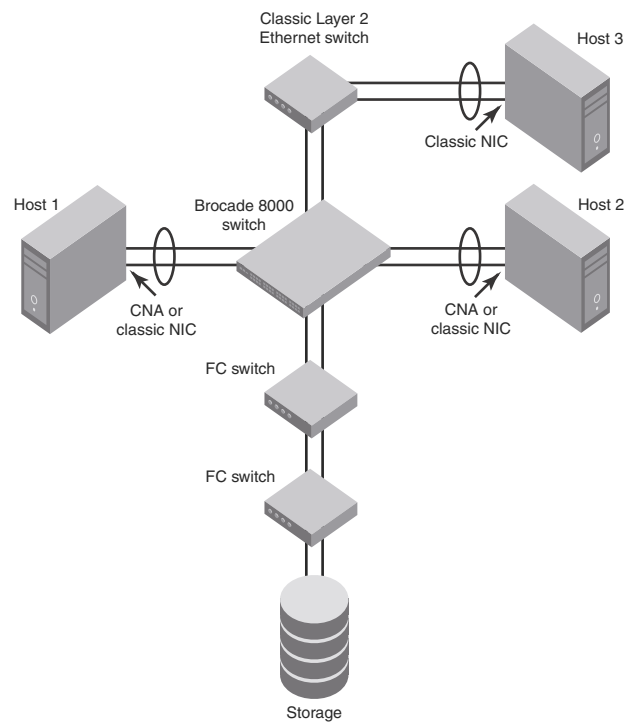
NOTE

The intermediate switching devices in the CEE network do not have to be FCoE-aware. They simply route the FCoE traffic to the FCoE device based on the Ethernet destination address in the FCoE frame.

Layer 2 Ethernet overview

The Brocade FCoE hardware contain CEE ports that support FCoE forwarding. The CEE ports are also backwards compatible and support classic Layer 2 Ethernet networks (see [Figure 1](#)). In Layer 2 Ethernet operation, a host with a Converged Network Adapter (CNA) can be directly attached to a CEE port on the Brocade FCoE hardware. Another host with a classic 10-Gigabit Ethernet NIC can be either directly attached to a CEE port, or attached to a classic Layer 2 Ethernet network which is attached to the Brocade FCoE hardware.

FIGURE 1 Multiple switch fabric configuration



Layer 2 forwarding

Layer 2 Ethernet frames are forwarded on the CEE ports. 802.1Q VLAN support is used to tag incoming frames to specific VLANs, and 802.3ac VLAN tagging support is used to accept VLAN tagged frames from external devices. The 802.1D Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) are used as the bridging protocols between Layer 2 switches.

The Brocade FCoE hardware handles Ethernet frames as follows:

- When the destination MAC address is not in the lookup table, the frame is flooded on all ports except the ingress port.
- When the destination MAC address is present in the lookup table, the frame is switched only to the correct egress port.
- When the destination MAC address is present in the lookup table, and the egress port is the same as the ingress port, the frame is dropped.

1 Layer 2 Ethernet overview

- If the Ethernet Frame Check Sequence (FCS) is incorrect, because the switch is in cut-through mode, a correctly formatted Ethernet frame is sent out with an incorrect FCS.
- If the Ethernet frame is too short, the frame is discarded and the error counter is incremented.
- If the Ethernet frame is too long, the frame is discarded and the error counter is incremented.
- Frames sent to a broadcast destination MAC address are flooded on all ports except the ingress port.
- When MAC address entries in the lookup table time out, they are removed. In this event, frame forwarding changes from unicast to flood.
- An existing MAC address entry in the lookup table is discarded when a device is moved to a new location. When a device is moved, the ingress frame from the new port causes the old lookup table entry to be discarded and the new entry inserted into the lookup table. Frame forwarding remains unicast to the new port.
- When the lookup table is full, new entries replace the oldest MAC addresses after the oldest MAC addresses age and time out. MAC addresses that still have traffic running are not timed out.

NOTE

New entries start replacing older entries when the lookup table reaches 90 percent of its 32k capacity.

VLAN tagging

The Brocade FCoE hardware handles VLAN tagging as follows:

- If the CEE port is configured to tag incoming frames with a single VLAN ID, then incoming frames that are untagged are tagged with the VLAN ID.
- If the CEE port is configured to tag incoming frames with multiple VLAN IDs, then incoming frames that are untagged are tagged with the correct VLAN ID based on the port setting.
- If the CEE port is configured to accept externally tagged frames, then incoming frames that are tagged with a VLAN ID are passed through unchanged.

NOTE

Only a single switch-wide VLAN is capable of forwarding FCoE traffic.

For detailed information on configuring VLANs, see [“Configuring VLANs”](#) on page 45.

Loop-free network environment

The Brocade FCoE hardware uses the following protocols to maintain a loop-free network environment:

- 802.1D Spanning Tree Protocol (STP)—STP is required to create a loop-free topology in the LAN.
- Rapid Spanning Tree Protocol (RSTP)—RSTP evolved from the 802.1D STP standard. RSTP provides for a faster spanning tree convergence after a topology change.
- Multiple Spanning Tree Protocol (MSTP)—MSTP defines an extension to RSTP to further develop the usefulness of VLANs. With per-VLAN MSTP, you can configure a separate spanning tree for each VLAN group. The protocol automatically blocks the links that are redundant in each spanning tree.

Using MSTP, you can create multiple loop-free active topologies on a single physical topology. These loop-free topologies are mapped to a set of configurable VLANs. This enables you to better utilize the physical resources present in the network and achieve better load balancing of VLAN traffic.

For detailed information on configuring these protocols, see [“Configuring STP, RSTP, and MSTP”](#) on page 57.

Frame classification (incoming)

The Brocade FCoE hardware is capable of classifying incoming Ethernet frames based on the following criteria:

- Port number
- Protocol
- MAC address

The classified frames can be tagged with a VLAN ID or with 802.1p Ethernet priority. The 802.1p Ethernet priority tagging is done using the Layer 2 Class of Service (CoS). The 802.1p Ethernet priority is used to tag frames in a VLAN with a Layer 2 CoS to prioritize traffic in the VLAN. The Brocade FCoE hardware also accepts frames that have been tagged by an external device.

Frame classification options are as follows:

- VLAN ID and Layer 2 CoS by physical port number—With this option, the port is set to classify incoming frames to a preset VLAN ID and the Layer 2 CoS by the physical port number on the Brocade FCoE hardware.
- VLAN ID and Layer 2 CoS by LAG virtual port number—With this option, the port is set to classify incoming frames to a preset VLAN ID and Layer 2 CoS by the Link Aggregation Group (LAG) virtual port number.
- Layer 2 CoS mutation—With this option, the port is set to change the Layer 2 CoS setting by enabling the QoS mutation feature.
- Layer 2 CoS trust—With this option, the port is set to accept the Layer 2 CoS of incoming frames by enabling the QoS trust feature.

For detailed information on configuring QoS, see [“Configuring QoS”](#) on page 107.

Congestion control and queuing

The Brocade FCoE hardware supports several congestion control and queuing strategies. As an output queue approaches congestion, Random Early Detection (RED) is used to selectively and proactively drop frames to maintain maximum link utilization. Incoming frames are classified into priority queues based on the Layer 2 CoS setting of the incoming frame, or the possible rewriting of the Layer 2 CoS field based on the settings of the CEE port or VLAN.

The Brocade FCoE hardware supports a combination of two scheduling strategies to queue frames to the egress ports; Priority queuing, which is also referred to as strict priority, and Deficit Weighted Round Robin (DWRR) queuing.

The scheduling algorithms work on the eight traffic classes as specified in 802.1Qaz Enhanced Transmission Selection (ETS).

1 Layer 2 Ethernet overview

Queuing features are described as follows:

- RED—RED increases link utilization. When multiple inbound traffic streams are switched to the same outbound port, and some traffic streams send small frames while other traffic streams send large frames, link utilization will not be able to reach 100 percent. When RED is enabled, link utilization approaches 100 percent.
- Classification—Setting user priority.
 - Inbound frames are tagged with the user priority set for the inbound port. The tag is visible when examining the frames on the outbound port. By default, all frames are tagged to priority zero.
 - Externally tagged Layer 2 frames—When the port is set to accept externally tagged Layer 2 frames, the user priority is set to the Layer 2 CoS of the inbound frames.
- Queuing
 - Input queuing—Input queuing optimizes the traffic flow in the following way. Suppose a CEE port has inbound traffic that is tagged with several priority values, and traffic from different priority settings is switched to different outbound ports. Some outbound ports are already congested with background traffic while others are uncongested. With input queuing, the traffic rate of the traffic streams switched to uncongested ports should remain high.
 - Output queuing—Output queuing optimizes the traffic flow in the following way. Suppose that several ports carry inbound traffic with different priority settings. Traffic from all ports is switched to the same outbound port. If the inbound ports have different traffic rates, some outbound priority groups will be congested while others can remain uncongested. With output queuing, the traffic rate of the traffic streams that are uncongested should remain high.
 - Multicast rate limit—A typical multicast rate limiting example is where several ports carry multicast inbound traffic that is tagged with several priority values. Traffic with different priority settings is switched to different outbound ports. The multicast rate limit is set so that the total multicast traffic rate on output ports is less than the specified set rate limit.
 - Multicast input queuing—A typical multicast input queuing example is where several ports carry multicast inbound traffic that is tagged with several priority values. Traffic with different priority settings is switched to different outbound ports. Some outbound ports are already congested with background traffic while others are uncongested. The traffic rate of the traffic streams switched to the uncongested ports should remain high. All outbound ports should carry some multicast frames from all inbound ports. This enables multicast traffic distribution relative to the set threshold values.
 - Multicast output queuing—A typical multicast output queuing example is where several ports carry multicast inbound traffic. Each port has a different priority setting. Traffic from all ports is switched to the same outbound port. If the inbound ports have varying traffic rates, some outbound priority groups will be congested while others remain uncongested. The traffic rate of the traffic streams that are uncongested remains high. The outbound ports should carry some multicast frames from all the inbound ports.
- Scheduling—A typical example of scheduling policy (using SPO and SP1 modes) is where ports 0 through 7 carry inbound traffic, each port has a unique priority level, port 0 has priority 0, port 1 has priority 1, and so on. All traffic is switched to the same outbound port. In SPO mode, all ports have DWRR scheduling; therefore, the frames-per-second (FPS) on all ports should correspond to the DWRR settings. In SP1 mode, priority 7 traffic uses SP; therefore, priority 7 can achieve a higher FPS. Frames from input ports with the same priority level should be scheduled in a round robin manner to the output port.

When setting the scheduling policy, each priority group that is using DWRR scheduling can be set to use a percentage of the total bandwidth by setting the PG_Percentage parameter.

For detailed information on configuring QoS, see [“Configuring QoS”](#) on page 107.

Access control

Access Control Lists (ACLs) are used for Layer 2 switching security. Standard ACLs inspect the source address for the inbound ports. Extended ACLs provide filtering by source and destination addresses and protocol. ACLs can be applied to the CEE ports or to VLANs.

ACLs function as follows:

- A standard Ethernet ACL configured on a physical port is used to permit or deny frames based on the source MAC address. The default is to permit all frames.
- An extended Ethernet ACL configured on a physical port is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames.
- A standard Ethernet ACL configured on a LAG virtual port is used to permit or deny frames based on the source MAC address. The default is to permit all frames. LAG ACLs apply to all ports in the LAG.
- An extended Ethernet ACL configured on a LAG virtual port is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames. LAG ACLs apply to all ports in the LAG.
- A standard Ethernet ACL configured on a VLAN is used to permit or deny frames based on the source MAC address. The default is to permit all frames. VLAN ACLs apply to the Switch Vertical Interface (SVI) for the VLAN.
- An extended Ethernet ACL configured on a VLAN is used to permit or deny frames based on the source MAC address, destination MAC address, and EtherType. The default is to permit all frames. VLAN ACLs apply to the Switch Vertical Interface (SVI) for the VLAN.

For detailed information on configuring ACLs, see [“Configuring ACLs”](#) on page 101.

Trunking

NOTE

The term “trunking” in an Ethernet network refers to the use of multiple network links (ports) in parallel to increase the link speed beyond the limits of any one single link or port, and to increase the redundancy for higher availability.

802.1ab Link Layer Discovery Protocol (LLDP) is used to detect links to connected switches or hosts. Trunks can then be configured between an adjacent switch or host and the Brocade FCoE hardware using the VLAN classifier commands. See [“Configuring an interface port as a trunk interface”](#) on page 50.

The Data Center Bridging (DCB) Capability Exchange Protocol (DCBX) extension is used to identify a CEE-capable port on an adjacent switch or host. For detailed information on configuring LLDP and DCBX, see [“Configuring LLDP”](#) on page 87.

The 802.3ad Link Aggregation Control Protocol (LACP) is used to combine multiple links to create a trunk with the combined bandwidth of all the individual links. For detailed information on configuring LACP, see [“Configuring Link Aggregation”](#) on page 79.

NOTE

The Brocade software supports a maximum 24 LAG interfaces.

Flow Control

802.3x Ethernet pause and Ethernet Priority-based Flow Control (PFC) are used to prevent dropped frames by slowing traffic at the source end of a link. When a port on a switch or host is not ready to receive more traffic from the source, perhaps due to congestion, it sends pause frames to the source to pause the traffic flow. When the congestion has been cleared, it stops requesting the source to pause traffic flow, and traffic resumes without any frame drop.

When Ethernet pause is enabled, pause frames are sent to the traffic source. Similarly, when PFC is enabled, there is no frame drop; pause frames are sent to the source switch.

For detailed information on configuring Ethernet pause and PFC, see [“Configuring QoS”](#) on page 107.

FCoE Initialization Protocol

The FCoE Initialization Protocol (FIP) discovers and initializes FCoE capable entities connected to an Ethernet cloud through a dedicated Ethertype, 0x8914, in the Ethernet frame.

FIP discovery

NOTE

This software version supports the October 8, 2008 (REV 1.03) of the ANSI FC Backbone Specification with priority-tagged FIP VLAN discovery protocol and FIP version 0. This release supports FIP Keep Alive.

The Brocade FCoE hardware FIP discovery phase operates as follows:

- The Brocade FCoE hardware uses the FCoE Initialization Protocol (FIP). ENodes discover FCFs and initialize the FCoE connection through the FIP.
- VF_port configuration—An FCoE port accepts ENode requests when it is configured as a VF_port and enabled. An FCoE port does not accept ENode requests when disabled.
- Solicited advertisements—A typical scenario is where a Brocade FCoE hardware receives a FIP solicitation from an ENode. Replies to the original FIP solicitation are sent to the MAC address embedded in the original FIP solicitation. After being accepted, the ENode is added to the VN_port table.
- Login group—When enabled, replies to solicitations are sent only by Brocade FCoE hardware that have the ENode in the login group.
- FCF forwarding—The Brocade FCoE hardware forwards FIP frames only when the VLAN is set to FCF forwarding mode.
- VLAN 1—The Brocade FCoE hardware should not forward FIP frames on VLAN 1 because it is reserved for management traffic only.
- A fabric-provided MAC address is supported. A server-provided MAC-address is not supported in the Fabric OS v7.0.1 release.

NOTE

In the fabric-provided MAC address format, VN_port MAC addresses are based on a 24-bit fabric-supplied value. The first three bytes of this value is referred to as the FCMAP. The next three bytes are the FC ID, which is assigned by the switch when the ENode logs in to the switch.

FIP login

FIP login operates as follows:

- ENodes can log in to the Brocade FCoE hardware using FIP. Fabric login (FLOGI) and fabric discovery (FDISC) are accepted. Brocade FCoE hardware in the fabric maintain the MAC address, World Wide Name (WWN), and PID mappings per login. Each ENode port should have a unique MAC address and WWN.
- FIP FLOGI—The Brocade FCoE hardware accepts the FIP FLOGI from the ENode. The FIP FLOGI acceptance (ACC) is sent to the ENode if the ENode MAC address or WWN matches the VN_port table on the Brocade FCoE hardware. The FIP FLOGI request is rejected if the ENode MAC address or WWN does not match. The ENode login is added to the VN_port table. Fabric Provided MAC addressing (FPMA) is supported.
- FIP FDISC—The Brocade FCoE hardware accepts FIP FDISC from the ENode. FIP FDISC acceptance (ACC) is sent to the ENode if the ENode MAC address or WWN matches the VN_port table on the Brocade FCoE hardware. The FIP FDISC request is rejected if the ENode MAC address or WWN does not match. The ENode login is added to the VN_port table. FPMA is supported.
- Maximum logins per VF_port—The Brocade FCoE hardware supports a maximum of 255 logins per VF_port. The VF_port rejects further logins after the maximum is reached.
- Maximum logins per switch—The Brocade FCoE hardware accepts a maximum of 1024 logins per switch. Note that the Brocade FCoE hardware rejects further logins after the maximum is reached.

FIP logout

FIP logout operates as follows:

- ENodes can log out from the Brocade FCoE hardware using FIP. The Brocade FCoE hardware in the fabric updates the MAC address, WWN, and PID mappings upon logout. The Brocade FCoE hardware also handles scenarios of implicit logout where the ENode has left the fabric without explicitly logging out.
- FIP logout (LOGO)—The Brocade FCoE hardware accepts a FIP LOGO from the ENode. The FIP LOGO ACC should be sent to the ENode if the ENode MAC address matches the VN_port table on the Brocade FCoE hardware. The LOGO is ignored (not rejected) if the ENode MAC address does not match. The ENode logout is updated in the VN_port table. FPMA is supported.
- Implicit logout—With the ENode directly connected to a CEE port, if the port that the ENode is attached to goes offline, the Brocade FCoE hardware implicitly logs out that ENode. ENode logout is updated in the VN_port table. The Brocade FCoE hardware sends clear virtual link (CVL) on behalf of the ENode.

FCoE login

The Brocade FCoE hardware FCoE login operates as follows:

- ENodes can log in to the Brocade FCoE hardware using FCoE encapsulated, FC Extended Link Service (ELS) frames. FLOGI and FDISC are accepted. Brocade FCoE hardware in the fabric maintains the MAC address to WWN/PID mappings per login. Class 2 FLOGI is not supported.
- FCoE FLOGI—The Brocade FCoE hardware accepts FCoE FLOGI from the ENode. FCoE FLOGI ACC is sent to the ENode if the FCMAP matches the VN_port table on the Brocade FCoE hardware. Requests are ignored if the FCMAP does not match. The ENode login is added to the VN_port table.
- FCoE FDISC—The Brocade FCoE hardware accepts FCoE FDISC from the ENode. FCoE FDISC ACC is sent to the ENode if the FCMAP matches the VN_port table on the Brocade FCoE hardware. The FCoE FDISC request is ignored if the FCMAP does not match. The ENode login is added to the VN_port table.
- FCMAP—The Brocade FCoE hardware accepts FCoE FLOGI from the ENode. The FCMAP determines which FCoE VLAN is accepted for the FCoE session.

NOTE

Only one FCoE VLAN is supported in the Fabric OS v7.0.1 release.

FCoE logout

The Brocade FCoE hardware FCoE logout operates as follows:

- ENodes can log out from the Brocade FCoE hardware using the FCoE encapsulated, FC ELS frame. Brocade FCoE hardware in the fabric updates the MAC address to WWN/PID mappings upon logout. The Brocade FCoE hardware also handles scenarios of implicit logout where the ENode has left the fabric without explicitly logging out.
- FCoE LOGO—The Brocade FCoE hardware accepts the FCoE LOGO from the ENode. The FCoE LOGO ACC is sent to the ENode if the ENode MAC address matches the VN_port table on the Brocade FCoE hardware. The LOGO is ignored (not rejected) if the ENode MAC address does not match. The ENode logout is updated in the VN_port table.

Logincfg

The Brocade FCoE hardware logincfg mechanism operates as follows:

- The logincfg is the mechanism for controlling ENode logins per Brocade FCoE hardware. Each unit of Brocade FCoE hardware maintains its own logincfg.
- Login configuration management is optional—when login management is disabled, the default behavior is to accept logins from any ENode.
- Loggingroup creation and deletion—The Brocade FCoE hardware accepts valid loggingroup names and member WWNs. The Brocade FCoE hardware rejects invalid entries. The Brocade FCoE hardware allows the deletion of loggingroups that are defined and committed. You can display defined and committed loggingroups. The loggingroup capability is disabled by default.
- Member add and remove—You can add valid member WWNs. Invalid WWNs are rejected. Duplicate WWNs are uniquely resolved. You can display the current view of defined loggingroups when changes are made to the configuration.

- Commit and abort—Defined loggingroup changes can be aborted with no effect on existing sessions. The Brocade FCoE hardware does not apply the configurations to new sessions until the changes are committed. Once defined, loggingroups are committed. The Brocade FCoE hardware immediately uses the new configuration.
- No traffic disruption—Changing the loggingroup without committing the changes does not affect existing sessions. After committing the changes, ENodes that were already logged in continue to function even when that member is removed from the loggingroup. New logins from the former member are rejected.

Name server

The Brocade FCoE hardware name server function operates as follows:

- ENode login and logout to and from the Brocade FCoE hardware updates the name server in the FC fabric. The Brocade FCoE hardware maintains the MAC address to WWN/PID mappings.
- ENode login and logout—When an ENode login occurs through any means (FIP FLOGI, FIP FDISC, FCoE FLOGI, or FCoE FDISC), an entry is added to the name server. When an ENode logout occurs through any means (FIP LOGO, FCoE LOGO, or implicit logout), the entry is removed from the name server.
- ENode data—The Brocade FCoE hardware maintains a VN_port table. The table tracks the ENode MAC address, FIP login parameters for each login from the same ENode, and WWN/PID mappings on the FC side. You can display the VN_port table with the **fcoe -loginshow port** command.

FC zoning

The Brocade FCoE hardware FC zoning operates as follows:

- The virtual devices created by the Brocade FCoE hardware on behalf of the ENodes are subject to FC zoning. An ENode is only allowed to access devices in the same zones. Administrative Domains (ADs) are not supported in the Fabric OS v7.0.1 release.
- ENodes can access FC devices in the same zones— FC devices that are not in the same zones cannot be accessed. Zone members can overlap in multiple zones (that is, overlapping zones). Zoning changes are immediately enabled by hardware enforced zoning.
- ENodes can access all FC devices with no zoning—ENodes can access all FC devices in the fabric when **cfgdisable** is issued and Default Zone is set to All Access Mode.
- Field replacement—When a Brocade FCoE hardware is replaced in the field, you can perform a **configdownload** on a previously saved configuration. No zoning change is required.

Registered State Change Notification (RSCN)

The Brocade FCoE hardware RSCN function operates as follows:

- RSCN events generated in the FC fabric are forwarded to the ENodes. RSCN events generated on the FCoE side are forwarded to the FC devices. CEE is not aware of RSCN events.
- Device RSCN—An RSCN is generated to all registered and affected members when an ENode either logs in or logs out of an FCF through any means. An RSCN is generated when an FC N_port device either logs in or logs out of the FC fabric.

NOTE

When transmitting an RSCN, zoning rules still apply for FCoE devices as the devices are treated as regular FC N_ports.

- VF_port RSCN—An RSCN is generated to all registered members when a VF_port goes online or offline, causing ENode or FC devices to be added or removed.
- Domain RSCN—An RSCN is generated to all registered and affected members when an FC switch port goes online or offline, causing ENode or FC devices to be added or removed. An RSCN is generated when two FC switches merge or segment, causing ENode or FC devices to be added or removed. When FC switches merge or segment, an RSCN is propagated to ENodes.
- Zoning RSCN—An RSCN is generated to all registered and affected members when a zoning exchange occurs in the FC fabric.

FCoE queuing

The QOS configuration controls the FCoE traffic distribution. Note that changing these settings requires changes on both the Brocade FCoE hardware and the CNA; therefore, the link must be taken offline and back online after a change is made. Traffic scheduler configuration changes affect FCoE traffic distribution as follows:

- Changing the priority group for a port causes the FCoE traffic distribution to update. The priority group and bandwidth are updated.
- Changing the priority table for a port causes the FCoE traffic distribution to be updated. The COS-to-priority group mapping is updated.
- Changing the class map for a port causes the FCoE traffic distribution to be updated.
- Changing the policy map for a port causes FCoE traffic distribution to be updated.
- Changing the CEE map for a port causes the FCoE traffic distribution to be updated.
- The FCMAP to VLAN mapping determines the FCoE VLAN allowed for the FCoE session. Modifying this mapping causes the existing sessions to terminate.

NOTE

Only one FCoE VLAN is supported in the Fabric OS v7.0.1 release.

Upgrade and downgrade considerations for FCoE

Firmware upgrades from Fabric OS v6.3.x or v6.4.0 to Fabric OS v7.0.0 are disruptive for FCoE traffic. When you upgrade, FCoE provisioning creates the default FCoE VLAN 1002 and the default CEE map with priority 3 with PFC enabled for FCoE traffic. No DCE ports are allowed to carry FCoE traffic by default. You are expected to enable FCoE on a DCE port using the `fcoeport` command. Once this is done, the CNA will login. See [“Assigning an FCoE map onto an interface”](#) on page 36.

Firmware downgrades from Fabric OS v7.0.0 to Fabric OS v6.3.x or v6.4.0 cause all of the FCoE provisioning configurations to be lost. To cause th CNAs to login again, you must recreate the FCoE configurations afresh using the Fabric OS v6.3.x or v6.4.0 CLIs.

Using the CEE CLI

In this chapter

- [Management Tools](#) 13
- [CEE Command Line Interface](#)..... 13

Management Tools

The Brocade FCoE hardware runs traditional Fabric OS software and can be managed using the same tools traditionally used for SAN management. Using the Fabric OS Command Line Interface (CLI), administrators have access to all commands and utilities common to other Brocade switches. In addition, the Fabric OS software enables Brocade Web Tools to support the following features for configuring and managing a Converged Ethernet Network:

- CEE interface display and configuration
- FCoE trunk display and configuration
- CEE configuration including link aggregation (LACP), Virtual LANs (VLANs), Quality of Service (QoS), and LLDP (Link Layer Discovery Protocol)/ DCBX protocol (Data Center Bridging eXchange)
- FCoE login groups

CEE Command Line Interface

The Brocade CEE CLI is designed to support the management of CEE and Layer 2 Ethernet switching functionality. The CEE CLI uses an industry-standard hierarchical shell familiar to Ethernet/IP networking administrators.

All conventional port-related Fabric OS CLI commands are only applicable to Fibre Channel. These commands have no knowledge of the Ethernet ports. The CEE features and CEE ports can only be configured through the CEE CLI interface which is accessed by entering the **cmsh** command from the Fabric OS shell.

The system starts up with the default Fabric OS configuration and the CEE startup configuration. After logging in you are in the Fabric OS shell. For information on accessing the CEE commands from the Fabric OS shell, see [“Accessing the CEE CLI from the Fabric OS shell”](#) on page 15.

Some Fabric OS commands are available in the CEE shell. Enter the **Fabric OS ?** command at the CEE CLI privileged EXEC mode command prompt to view the available Fabric OS commands. The traditional Fabric OS command help found in the Fabric OS shell is not available through the CEE shell.

NOTE

The CEE configuration is not affected by **configUpload** and **configDownload** commands entered in the Fabric OS shell.

Saving your configuration changes

Any configuration changes made to the switch are written into the running-config file. This is a dynamic file that is lost when the switch reboots. During the boot sequence, the switch resets all configuration settings to the values in the startup-config file.

To make your changes permanent, you must use either the **write memory** command or the **copy** command to commit the running-config file to the startup-config file.

Saving configuration changes with the copy command

Perform this task from privileged EXEC mode.

Enter the **copy** command to save the running-config file to the startup-config file.

```
switch#copy running-config startup-config
```

Saving configuration changes with the write memory command

Perform this task from privileged EXEC mode.

Enter the **write memory** command to save the running-config file to the startup-config file.

```
switch# write memory
Overwrite the startup config file (y/n): y
Building configuration...
```

CEE CLI RBAC permissions

Role-Based Action Control (RBAC) defines the capabilities that a user account has based on the role the account has been assigned. [Table 3](#) displays the permissions matrix for CEE. Permissions are specifically defined as follows:

- OM—When you enter the **cmsh** command, you are put directly into privileged EXEC mode.
- O—When you enter the **cmsh** command, you are limited to EXEC mode.
- N—You are not allowed access to the CEE CLI.

TABLE 3 CEE RBAC permissions

Root	Factory	Admin	User	Operator	SwitchAdmin	FabricAdmin	ZoneAdmin	BasicSwitchAdmin	SecurityAdmin
OM	OM	OM	O	N	O	OM	N	N	O

O = observe, OM = observe and modify, N = access not allowed

Accessing the CEE CLI through the console or Telnet

The procedure to access the CEE CLI is the same through either the console interface or through a Telnet session; both access methods bring you to the login prompt.

NOTE

While this example uses the **admin** role to log in to the switch, any role listed in the “[CEE CLI RBAC permissions](#)” section can be used.

```
switch login: admin
Password:
switch:admin> cms
switch#
```

To return to the Fabric OS CLI, enter the following command.

```
switch#exit
switch:admin>
```

NOTE

Multiple users can use Telnet and issue commands using EXEC mode and privileged EXEC mode.

Accessing the CEE CLI from the Fabric OS shell

To enter the CEE CLI from the Fabric OS shell, enter the following command.

```
switch:admin> cms
switch#
```

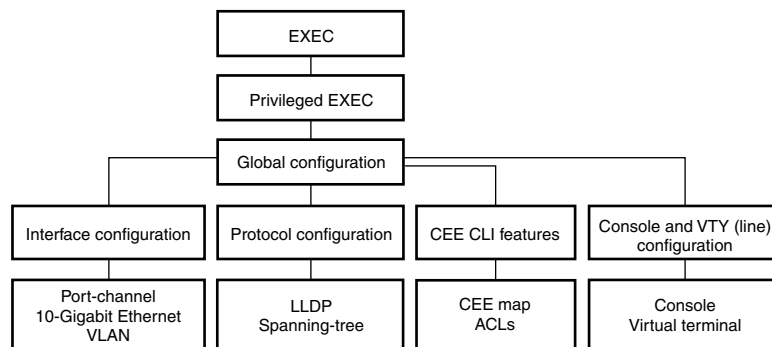
To return to the Fabric OS shell, enter the following command.

```
switch#exit
switch:admin>
```

CEE CLI command modes

[Figure 2](#) displays the CEE CLI command mode hierarchy.

FIGURE 2 CEE CLI command mode hierarchy



[Table 4](#) lists the CEE CLI command modes and describes how to access them.

2 CEE Command Line Interface

NOTE

At system startup, if you try to enter privileged EXEC mode before the system has fully booted, the following message is displayed:

```
%Info: Please wait. System configuration is being loaded.
```

After the system has fully booted, a RASlog message indicates that the CEE CLI is ready to accept configuration commands.

TABLE 4 CEE CLI command modes

Command mode	Prompt	How to access the command mode	Description
EXEC	switch>	Enter the cmsh command at the Fabric OS prompt after you have logged in as an appropriate user.	Display running system information and set terminal line parameters.
Privileged EXEC	switch#	From the EXEC mode, enter the enable command.	Display and change system parameters. Note that this is the administrative mode and also includes EXEC mode commands.
Global configuration	switch(config)#	From the EXEC mode, enter the configure terminal EXEC command.	Configure features that affect the entire switch.
Interface configuration	Port-channel: switch(conf-if-po-63)# 10-Gigabit Ethernet (CEE port): switch(conf-if-te-0/1)# VLAN: switch(conf-if-vl-1)#	From the global configuration mode, specify an interface by entering one of the following commands: <ul style="list-style-type: none"> • interface port-channel • interface tengigabitethernet • interface vlan 	Access and configure individual interfaces.
Protocol configuration	LLDP: switch(conf-lldp)# Spanning-tree: switch(conf-mstp)# switch(conf-rstp)# switch(conf-stp)#	From the global configuration mode, specify a protocol by entering one of the following commands: <ul style="list-style-type: none"> • protocol lldp • protocol spanning-tree mstp • protocol spanning-tree rstp • protocol spanning-tree stp 	Access and configure protocols.

TABLE 4 CEE CLI command modes (Continued)

Command mode	Prompt	How to access the command mode	Description
Feature configuration	CEE map: <code>switch(config-ceemap)#</code> Standard ACL: <code>switch(conf-macl-std)#</code> Extended ACL: <code>switch(conf-macl-ext)#</code>	From the global configuration mode, specify a CEE feature by entering one of the following commands: <ul style="list-style-type: none"> • cee-map • mac access-list 	Access and configure CEE features. Identify traffic based on the MAC addresses, such as VLAN IDs and different encapsulations. Standard Access Control Lists filter the traffic on a source address and block traffic close to a destination. Extended Access Control Lists block traffic based on any given packet attribute.
Console and VTY (line) configuration	<code>switch(config-line)#</code>	From the global configuration mode, configure a terminal connected through the console port by entering the line console command. Configure a terminal connected through a Telnet session by entering the line vty command.	Configure a terminal connected through the console port or a terminal connected through a Telnet session. After you apply the access list to an interface, a Virtual Teletype (VTY), or through a command using the access list keyword, it becomes effective.

NOTE

Pressing **Ctrl+Z** or entering the **end** command in any mode returns you to privileged EXEC mode. Entering **exit** in any mode returns you to the previous mode.

CEE CLI keyboard shortcuts

Table 5 lists CEE CLI keyboard shortcuts.

TABLE 5 CEE CLI keyboard shortcuts

Keystroke	Description
Ctrl+B or the left arrow key	Moves the cursor back one character.
Ctrl+F or the right arrow key	Moves the cursor forward one character.
Ctrl+A	Moves the cursor to the beginning of the command line.
Ctrl+E	Moves the cursor to the end of the command line.
Esc B	Moves the cursor back one word.
Esc F	Moves the cursor forward one word.
Ctrl+Z	Returns to privileged EXEC mode.
Ctrl+P or the up arrow key	Displays commands in the history buffer with the most recent command displayed first.
Ctrl+N or the down arrow key	Displays commands in the history buffer with the most recent command displayed last.

NOTE

In EXEC and privileged EXEC modes, use the **show history** command to list the commands most recently entered. The switch retains the history of the last 1000 commands entered.

Using the do command as a shortcut

You can use the **do** command to save time when you are working in any configuration mode and you want to run a command in EXEC or privileged EXEC mode.

For example, if you are configuring an LLDP and you want to execute a privileged EXEC mode command, such as the **dir** command, you would first have to exit the LLDP configuration mode. However, by using the **do** command with the **dir** command, you can ignore the need to change configuration modes, as shown in the following example.

```
switch(conf-lldp)#do dir
Contents of flash://
-rw-r-----      1276   Wed Feb  4 07:08:49 2009   startup_rmon_config
-rw-r-----      1276   Wed Feb  4 07:10:30 2009   rmon_config
-rw-r-----      1276   Wed Feb  4 07:12:33 2009   rmon_configuration
-rw-r-----      1276   Wed Feb  4 10:48:59 2009   starup-config
```

Displaying CEE CLI commands and command syntax

Enter a question mark (?) in any command mode to display the list of commands available in that mode.

```
switch>?
Exec commands:
enable      Turn on privileged mode command
exit        End current mode and down to previous mode
help        Description of the interactive help system
logout      Exit from the EXEC
quit        Exit current mode and down to previous mode
show        Show running system information
terminal    Set terminal line parameters
```

To display a list of commands that start with the same characters, type the characters followed by the question mark (?).

```
switch>e?
enable      Turn on privileged mode command
exit        End current mode and down to previous mode
```

To display the keywords and arguments associated with a command, enter the keyword followed by the question mark (?).

```
switch#terminal ?
length      Set number of lines on a screen
no          Negate a command or set its defaults
```

If the question mark (?) is typed within an incomplete keyword, and the keyword is the only keyword starting with those characters, the CLI displays help for that keyword only.

```
switch#show d?
dot1x      IEEE 802.1X Port-Based Access Control
<cr>
```


If the question mark (?) is typed within an incomplete keyword but the keyword matches several keywords, the CLI displays help for all the matching keywords.

```
switch#show i?
  interface  Interface status and configuration
  ip         Internet Protocol (IP)
```

The CEE CLI accepts abbreviations for commands. The following example is the abbreviation for the **show qos interface all** command.

```
switch#sh q i a
```

If the switch does not recognize a command after **Enter** is pressed, an error message displays.

```
switch#hookup
      ^
% Invalid input detected at '^' marker.
```

If an incomplete command is entered, an error message displays.

```
switch#show
% Incomplete command.
```

CEE CLI command completion

To automatically complete the spelling of commands or keywords, begin typing the command or keyword and then press **Tab**. For example, at the CLI command prompt, type **te** and press **Tab**:

```
switch#te
```

The CLI displays:

```
switch#terminal
```

If there is more than one command or keyword associated with the characters typed, the CEE CLI displays all choices. For example, at the CLI command prompt, type **show l** and press **Tab**:

```
switch#show l
```

The CLI displays:

```
switch#show l
lacp line lldp
```

CEE CLI command output modifiers

You can filter the output of the CEE CLI **show** commands using the output modifiers described in [Table 6](#).

TABLE 6 CEE CLI command output modifiers

Output modifier	Description
append	Appends the output to a file.
redirect	Redirects the command output to the specified file.
include	Displays the command output that includes the specified expression.
exclude	Displays the command output that excludes the specified expression.
append	Appends the command output to the specified file.
begin	Displays the command output that begins with the specified expression.

2 CEE Command Line Interface

TABLE 6 CEE CLI command output modifiers (Continued)

Output modifier	Description
last	Displays only the last few lines of the command output.
tee	Redirects the command output to the specified file. Note that this modifier also displays the command output.
FLASH	Redirects the output to flash memory.

Configuration management

In this chapter

- Configuration management tasks 21
- Flash file management commands 23
- Debugging and logging commands 23

Configuration management tasks

This chapter describes the typical configuration management tasks you may encounter.

The current configuration on the switch is referred to as the running configuration (running-config). The running-config file can be written to non-volatile memory to save configuration changes. Additionally, the running-config file can be saved as the startup configuration (startup-config) file. When the switch is booted, the system reads the contents of the startup-config file and applies it to the running-config.

Typical configuration management tasks are as follows:

- Saving the startup-config and running-config files to Flash.
- Uploading the startup-config and running-config files to a remote location.
- Uploading any configuration file saved and stored in Flash to a remote location.
- Downloading a configuration file from a remote location to the switch to serve as the startup-config file or the running-config file.
- Downloading a configuration file from a remote location to the switch Flash.

CEE configuration guidelines and restrictions

Follow these guidelines and restrictions when performing any CEE configuration management tasks.

- The CEE configuration database is maintained in a file separate from the Fabric OS configuration database. Note that Fabric OS configuration management remains unchanged.
- The CEE configuration is not affected by **configUpload** and **configDownload** commands entered in the Fabric OS shell.
- The configuration must be manually saved using the **write** or **copy** commands.

Display the running configuration file

To display the running configuration, perform the following task from EXEC or privileged EXEC mode.

Enter the **show running-config** command.

```
switch#show running-config
```

Saving the running configuration file

This task causes the running configuration to become the default configuration. To save the running configuration, perform the following task from privileged EXEC mode.

Enter the **copy** command to copy the currently running configuration to the startup configuration.

```
switch#copy running-config startup-config
```

Loading the startup configuration file

If you wish to reverse the changes to the running configuration, this task reloads the default startup configuration, overwriting the running configuration.

To load the default configuration, perform the following command from privileged EXEC mode.

```
switch#copy startup-config running-config
```

Erasing the startup configuration file.

NOTE

This task does not affect the running configuration file.

To erase the startup configuration, perform the following command from privileged EXEC mode.

```
switch#write erase
```

Archiving the running configuration file

This task allows you to archive the running configuration to an archive folder on an FTP site, so that it can be stored without changing the startup configuration. To archive the running configuration file, perform the following task from privileged EXEC mode.

Enter the **copy** command to archive the running configuration file.

```
switch#copy running-config  
ftp://jsmith:password@10.24.48.51/archive/config_file]
```

Restore an archived running configuration file

To restore the running configuration, perform the following task from privileged EXEC mode.

Enter the **copy** command to restore the running configuration file.

```
switch#copy running-config  
ftp://jsmith:password@10.24.48.51/archive/config_file]
```

Archiving the startup configuration file

This task allows you to archive the startup configuration to an archive folder on an FTP site. To archive the startup configuration, perform the following task from privileged EXEC mode.

Enter the **copy** command to archive the startup configuration file.

```
switch#copy startup-config
ftp://jsmith:password@10.24.48.51/archive/config_file]
```

Restore an archived startup configuration file

To restore the startup configuration, perform the following task from privileged EXEC mode.

Enter the **copy** command to restore the startup configuration file.

```
switch#copy startup-config
ftp://jsmith:password@10.24.48.51/archive/config_file]
```

Flash file management commands

[Table 7](#) describes the common tasks used to manage Flash files. The *Converged Enhanced Ethernet Command Reference* contains complete information on all available CLI commands.

NOTE

Use of the flash:// prefix is optional.

The tasks listed in [Table 7](#) are performed from privileged EXEC mode.

TABLE 7 CEE Flash memory file management commands

Task	Command
List the files in the Flash memory directory.	switch# dir
Delete a file from the Flash memory directory.	switch# delete flash://example_filename
Erase all the files in the Flash memory directory. NOTE: This command erases <i>all</i> the files in the Flash directory except the default startup configuration file which is programmed as a manufacturing default.	switch# erase flash % Warning: Erasing flash filesystem will remove all files in flash://. Continue to erase?(y/n): y
Rename a file in the Flash.	switch# rename filename new_filename
Display the contents of a file in the Flash memory directory.	switch# show file flash://example_filename

Debugging and logging commands

[Table 8](#) describes the tasks related to debugging and logging commands. The *Converged Enhanced Ethernet Command Reference* contains complete information on all available CLI commands.

3 Debugging and logging commands

The tasks listed in [Table 8](#) are performed from privileged EXEC mode.

TABLE 8 Debugging and logging commands

Task	Command
Display debugging information for CEE components.	switch# show debug
Display logging information for CEE components.	switch# show logging
Display the collection of information needed for technical support.	switch# show tech-support

NOTE: The **supportsave** command in Fabric OS includes the debugging data provided by the above commands.

FCoE Configuration Management

In this chapter

- [Managing the FCoE configuration](#) 25
- [Configuring the FCoE login information](#) 26
- [Creating an FCoE login group](#) 28

Managing the FCoE configuration

FCoE technology bridges the boundary between the SAN and LAN sections of your network. FCoE configuration tasks require mostly configuration of the interface ports on the switch.

This chapter discusses the tasks for managing your FCoE configuration, login groups, and VN_port mappings.

NOTE

To assist you in configuring FCoE, see [“FCoE Initialization Protocol”](#) on page 8.

FCoE configuration guidelines and restrictions

Follow these FCoE configuration guidelines and restrictions when configuring FCoE:

- The CEE configuration database is maintained in a file separate from the Fabric OS configuration database. Fabric OS configuration management procedures remain unchanged.
- FCoE to FCoE traffic across two FCOE10-24 blades can only reach 68% line rate using a port-based routing policy. Using an exchange based routing policy can avoid the performance drop.

Only the WWN zoning of FCoE VF ports is supported. Port-based zoning of the FCoE VF port is not supported. Additionally, inclusion of FCoE VF ports in a zone which has port-based zone members (such as zone members specified by their respective domain and index) is not supported. However there are no restrictions enforced that block you from enabling the port-based zoning. If you configure port-based zoning on a VF port, it converts all the other ports in the FCoE port group to “session based” zoning. The only way to recover the VF port from this scenario is to set all the VF ports in the FCoE port group offline, change the zoning config of VF ports to WWN-based zoning and bring them back to online.

Clearing logins

Perform the following task to clear logins.

Task	Command
Clear the logins that occurred through a front-end port or from a device specified by the ENode's VN_port WWN. Syntax is as follows: <ul style="list-style-type: none"> • -teport slot/port—Specifies the slot or port number. • -device wwn—Specifies the device WWN. 	switch:admin> fcoe -resetlogin -teport slot/port -device wwn

Displaying FCoE configuration-related information

Perform the following tasks to display FCoE-related configuration information.

Task	Command
Display the embedded FCoE port configuration. Configurations of all the ports are displayed if you do not specify a specific port.	switch:admin> fcoe -cfgshow [slot/port]
Display information about devices logged in to a specific FCoE F_port.	switch:admin> fcoe -loginshow [slot/port]
Display FIP configurations.	switch:admin> fcoe -fipcfgshow

Configuring the FCoE login information

Another important task in administrating FCoE is configuring the FCoE login information.

Enabling or disabling login configuration management

The **fcoelogincfg** command allows only configured ENode VN_ports to log in. Use the **fcoelogingroup** command to configure allowed ENode VN_ports. The default is disabled.

Disabling the **fcoelogincfg** command allows unrestricted login on ENode VN_ports.

Perform one of the following tasks to toggle the availability of FCoE login configuration management.

Task	Command
Enable the FCoE login configuration management on the switch (this is a switch-based command, not port-based).	switch:admin> fcoelogincfg -enable
Disable the FCoE login configuration management on the switch.	switch:admin> fcoelogincfg -disable

Managing the current configuration transaction

NOTE

The configuration changes created using the **fcoelogingroup** command are kept in a transaction buffer until you save the buffer using the fabric-wide **fcoelogincfg-save** command. The login configuration is saved as a transaction, and to apply it you need to specifically save it.

Perform one of the following tasks to either display or abort the current configuration transaction.

Task	Command
Display the current configuration transaction.	switch:admin> fcoelogincfg -transshow
Abort the current configuration transaction.	switch:admin> fcoelogincfg -transabort

Cleaning up login groups and VN_port mappings

Perform the following tasks to clean up login groups and VN_port mappings.

Task	Command
Perform a cleanup of all conflicting login groups and VN_port mappings from the effective configuration. This purges not only the conflicting login groups but also the non-existing switches.	switch:admin> fcoelogincfg -purge -conflicting -nonexisting
Perform a cleanup of all conflicting login groups and conflicting VN_port mappings from the effective configuration.	switch:admin> fcoelogincfg -purge -conflicting
Perform a cleanup of all login groups for non-existing switches from the effective configuration.	switch:admin> fcoelogincfg -purge -nonexisting

Displaying the FCoE login configuration

Perform the following tasks to display the FCoE login configuration.

Task	Command
Display the FCoE login configuration. Syntax is as follows: <ul style="list-style-type: none"> • -switch swwn—Displays all of the login groups for the specified switch. • -logingroup lgnam—Displays the login group configuration for the specified login group. • -saved—Displays only the effective configuration. 	switch:admin> fcoelogincfg -show [-switch swwn -logingroup lgnam] [-saved]
Display the status of the last configuration merge during the last fabric merge. This operand also displays conflicting login groups and login groups for non-existing switches.	switch:admin> fcoelogincfg -show [-mergestatus]

Saving the current FCoE configuration

Perform the following task to save the current FCoE configuration.

Task	Command
Save the current FCoE login configuration as the effective configuration fabric-wide.	switch:admin> fcoelogincfg -save

Creating an FCoE login group

The FCoE login group enables you to configure login policies.

Perform the following task to create an FCoE login group.

Task	Command
<p>Syntax is as follows:</p> <ul style="list-style-type: none"> • -create—Create a login group. • <i>lname</i>—Specify the name of the login group for this switch. The maximum length is a 64-byte string. • -switch <i>swwn</i>—Specify the WWN of the switch for which the login group is being created. • -self—Specify the WWN of the current switch. • -allowall—Allow all VN_port devices to log in to the switch. • <i>member</i>—Identify the WWN of the VN_port. The WWN must be specified in hex as xx:xx:xx:xx:xx:xx:xx:xx. Only specified members are allowed to log in to the switch. 	<pre>switch:admin> fcoelogingroup -create lname -switch swwn -self [-allowall "member; member;..."]</pre>

Modifying the FCoE login group device list

Perform the following tasks to add or remove VN_port devices from the FCoE login group.

Task	Command
Add VN_port devices to the FCoE login group. Syntax is as follows: <ul style="list-style-type: none"> • <i>lgnam</i>—Specify the name of the login group to which VN_port devices are to be added. • <i>member</i>—Identify the WWN of the VN_port. The WWN must be specified in hex as xx:xx:xx:xx:xx:xx:xx:xx. Only specified members are allowed to log in to the switch. 	switch:admin> fcoelogingroup –add <i>lgnam member; member; ...</i>
Remove VN_port devices from the FCoE login group. Syntax is as follows: <ul style="list-style-type: none"> • <i>lgnam</i>—Specify the name of the login group from which VN_port devices are to be removed. • <i>member</i>—Identify the WWN of the VN_port. The WWN must be specified in hex as xx:xx:xx:xx:xx:xx:xx:xx. Only specified members are allowed to log in to the switch. 	switch:admin> fcoelogingroup –remove <i>lgnam member; member;...</i>

Deleting an FCoE login group

Perform the following task to delete an FCoE login group.

Task	Command
Delete an FCoE login group. Syntax is as follows: <ul style="list-style-type: none"> • <i>lgnam</i>—Specify the name of the login group. 	switch:admin> fcoelogingroup –delete <i>lgnam</i>

Renaming an FCoE login group

Perform the following task to rename an FCoE login group.

Task	Command
Rename an FCoE login group. Syntax is as follows: <ul style="list-style-type: none"> • <i>lgnam</i>—Specify the name of the login group from which VN_port devices are to be removed. 	switch:admin> fcoelogingroup –rename <i>lgnam newlgnam</i>

4 Creating an FCoE login group

Initial FCoE and CEE Configuration

In this chapter

• Overview	31
• End to End FCoE using FC ISLs	31
• High availability	33
• Configuring the FCoE interfaces	35
• Configuring the FCoE VLAN	35
• Configuring DCBX	37
• Configuring Spanning Tree Protocol	38
• Configuring VLAN membership	38
• CEE map configuration	39
• CEE audit logs	42

Overview

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE) properties on the switch using CLI commands. For detailed information about all the CEE CLI commands, refer to the *Converged Enhanced Ethernet Command Reference Supporting Fabric OS v6.4.0* (53-1001762-01).

All of the CLI commands are entered using the Telnet or console interface on the Brocade FCoE hardware.

End to End FCoE using FC ISLs

Support for end-to-end FCoE is achieved through FC ISLs, allowing FCoE traffic to be carried across multiple FC ISL hops, eventually terminating in FC or FCoE storage.

The following diagram shows a deployment scenario where FCoE initiators are connected to FCoE and FC targets across multiple hops of FC ISLs.

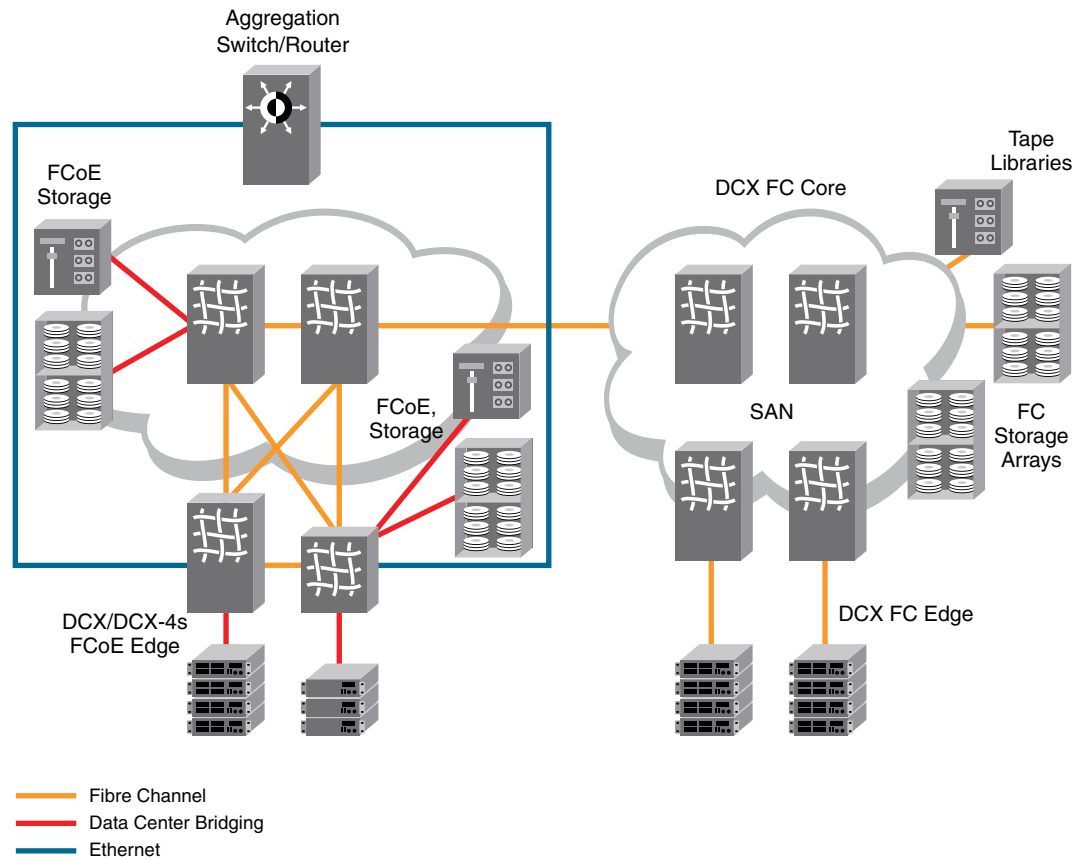


FIGURE 3 Deploying FCoE across multiple hops of FC ISLs

FCoE Logical Port Numbering

A total of 24 FCoE logical ports on FCOE10-24 belong to 6 FCoE port groups. The port number range is 0 through 23 and these belong to 6 FCoE port groups ranging from 0 through 3, 4 through 7, 8 through 11, 12 through 15, 16 through 19, and 20 through 23. For example, the FCoE ports on a 24 port blade in the 4th slot is 4/0 through 4/23.

NOTE

This grouping is relevant only for bridging bandwidth considerations. A group of four ports that is arranged in a bank has 10Gb of bridging bandwidth.

FCoE port-to-front end TenGigabit port mapping

Each of the FCoE logical ports in VF port mode is mapped to a corresponding front end TenGigabit port such that the FCoE FLOGI received through the TenGigabit port is mapped and handled at the corresponding FCoE logical port.

On an FCOE10-24 this mapping is as follows, assuming that the FCOE10-24 blade is in slot 4:

- FCoE port 4/0 maps to TE port 4/0
- FCoE port 4/1 maps to TE port 4/1
- FCoE port 4/23 maps to TE port 4/23

High availability

Fabric OS v7.0.1 ensures high availability of FCoE storage traffic going through FCOE10-24 blades on the DCX and DCX-4S platforms. If the Control Processors (CPs) on the DCX or DCX-4S are in a redundant state and the active CP fails, the standby CP takes over the control of the system. Failover from the active CP to the standby CP is non-disruptive to both FC and FCoE traffic.

Hot plug

The procedure for hot-plugging an FCOE10-24 is similar to any other blade supported on a DCX or DCX-4S. When diagnostics are running on a blade, a slotpoweroff operation is not allowed. This behavior is similar to what has been supported in the previous FOS releases.

When a core blade on a DCX or DCX-4S is removed, the traffic flowing through the ASICs in that blade is re-routed to use the ASICs in the other core blade in the chassis. This increases the congestion, but the switch continues to operate in an over-subscribed mode. When both the core blades are removed, then all the blades in the chassis are disabled.

Control Processor failover

When the active CP in a DCX or DCX-4S fails over, the FCoE storage traffic going through the FCOE10-24 blades are not disrupted. Critical protocol components are brought online quickly to ensure that PDUs are transmitted to the adjacent switches so that the associated protocol sessions are not declared dead.

ATTENTION

Upgrading your firmware from Fabric OS v6.3.x, v6.4.0, or v6.4.1 does not cause FC traffic disruption but does cause FCoE traffic disruption.

Brocade recommends that the protocol timeout values in [Table 9](#) be configured both on the local Brocade switches, as well as neighboring switches, to ensure non-disruptive HA failover.

TABLE 9 Protocol timeout values for non-disruptive failover

Protocol type	Time interval
LACP	Long timeout = 90 seconds (default value)
FIP	Hello Interval = 8 seconds (default value)
RSTP	3 Second Hello (9 second session timeout)
MSTP	3 Second Hello (9 second session timeout)

Full recovery of all services can take up to a minute. No new configuration is allowed for that period. Configuration manager generates a RASLOG message to indicate the availability of the standby CP after recovery to accept new configuration CLIs.

Every configuration is either fully committed or not, but never “half-complete.” If you had started executing the **copy startup running** command, it will be resumed after failover before any new CLI is allowed to be executed.

High availability reference

This section contains reference tables for high availability.

TABLE 10 State preservation during CP failover

Operation, data, or state	Expected behavior
CLIs that are being executed when the HA failover occurs.	All telnet/ssh sessions are disconnected and need to be reestablished after failover. Any CLI operation under progress will either be fully committed if its execution had started on active CP or not executed completely. The vty timeout configuration is preserved across failover.
Running and Startup configurations	The running and startup configuration is preserved. All other configuration files are preserved.
L2 forwarding table	This item is preserved during failover.
VLAN table and state	This item is preserved during failover.
Static Multicast MAC DA	This item is preserved during failover.
Dynamic Multicast MAC DA (learned via IGMP snooping)	This item is preserved during failover.
MAC learning process	MAC learning entries that are in the ASIC queue but yet to be processed by software is flushed.
MAC ageing timer	The timer is restarted after the failover.
Station moves during failover	Station moves are not detected.
Interface state	Preserved as long as the interfaces are either in a shut or no-shut state. Any interface which is not fully initialized is reset.
Interface counters	This item is preserved during failover.
Interface counters related to CoS, Flow control, RED Alarms and other configuration	This item is preserved during failover, however since interface counters are not preserved, the alarms are restarted upon HA failover.
xSTP Port STP state	This item is preserved as long as the ports are in a steady state, otherwise the port state machine is reset.
IGMP snooping – Dynamically learned groups	This item is preserved during failover.
Static, Dynamic Standard, and Brocade LAGs	This item is preserved during failover.
ACL policies	This item is preserved during failover, however the ACL counters are not preserved.
LLDP, neighbor information, profiles, or DCBX state machine	This item is preserved during failover.
QoS policies	This item is preserved during failover.
Port authentication state	Preserved during failover, as long as the ports are in a steady state, otherwise the port state machine are reset.
SNMP sessions	Reset during failover. SNMP manager must reconnect after the failover.

TABLE 10 State preservation during CP failover (Continued)

Operation, data, or state	Expected behavior
RMON – Alarms and other configuration	This item is preserved during failover. However, alarms timers are restarted upon HA failover.
Login, FCMAP, FIP, or VLAN configuration	This item is preserved during failover.
MIB objects and tables	This item is preserved during failover.
FCoE Initiator entries in FC login database Name server	This item is preserved during failover.

Configuring the FCoE interfaces

FCoE maps are used to configure FCoE properties on interfaces. An FCoE map is a placeholder for an FCoE VLAN and a CEE map. You can assign FCoE maps on to physical interfaces using the **fcoeport** command. Once the FCoE map is assigned onto an interface:

- The corresponding FCoE VLAN is applied to the interface.
- The corresponding CEE map is applied to the interface.
- The FCoE/FIP vlan classifiers are applied to the interface.

In short, the interface becomes capable of carrying FCoE traffic. The FCoE map can be applied on an interface only if the FCoE map is complete in all aspects. That is, it should have an FCoE VLAN and a CEE map associated with it.

Only a single FCoE map is allowed, which is created automatically with the name “default.” You are not able to delete or rename this map. By default, if there are no other conflicting configurations, the FCoE VLAN associated to the FCoE map is FCoE VLAN (1002) and the CEE map associated is the default CEE map (also called “default”).

Configuring the FCoE VLAN

The contents of the default running-configuration file are displayed with the **show running-config** command:

```
switch#show running-config
no protocol spanning-tree
cee-map default
priority-group-table 1 weight 40 pfc
priority-group-table 2 weight 60
priority-table 2 2 2 1 2 2 2
!
fcoe-map default
fcoe-vlan 1002
!
interface vlan 1
```

If the system already contains an FCoE VLAN 1002, it continues to be used as the FCoE VLAN by creating a default fcoe-map with fcoe-vlan 1002. If the pre-existing FCoE VLAN is set to something other than 1002, that VLAN is configured as a standard VLAN and new FCoE VLAN 1002 is created with the default fcoe-map.

VLAN 1002 is reserved for FCoE VLAN. If VLAN 1002 is already assigned as a LAN VLAN, it is automatically recreated as the FCoE VLAN while creating the default fcoe-map.

The FCoE VLAN can be created in the FCoE map configuration mode as well. The **fcoe-vlan** command creates the FCoE VLAN if the corresponding VLAN does not exist in the system. If a VLAN already exists with the same VLAN ID, then the command fails.

The FCoE VLANs can be seen using the **show vlan fcoe** command. Alternatively, the **show vlan brief** command can be used, in which FCoE VLANs are marked with an "(F)".

No VLAN classifiers (including FIP and FCoE classifiers) are allowed to be configured for the FCoE VLAN.

To configure the FCoE VLAN, perform the following steps in global configuration mode.

1. Activate the FCoE map configuration mode.

```
switch(config)#fcoe-map default
switch(conf-fcoe-map)#
```

2. Use the **fcoe-vlan** command to specify a new VLAN ID.

Only one FCoE VLAN can be active at a time, and this command will be successful only if the new VLAN ID can be created and the old VLAN ID can be deleted. The FCoE VLAN cannot be changed if the fcoeport profile is already applied to an interface.

```
switch(conf-fcoe-map)#fcoe-vlan 5
```

3. Use the **do show vlan fcoe** command to confirm the VLAN ID change.

```
switch(conf-fcoe-map)#do show vlan fcoe
VLAN      Name                State           Ports
(F)-FCoE                                     (u)-Untagged, (t)-Tagged
=====
5(F)      VLAN0005            ACTIVE
```

4. Use the **do show fcoe-map** command to confirm the current status of the FCoE map.

The following example shows that the FCoE VLAN has been assigned the ID of 5, but the map has not been assigned to any interfaces.

```
switch(conf-fcoe-map)#do show fcoe-map
fcoe-map default
  cee-map      : default
  fcoe-vlan    : 5
  interface(s) :
```

Assigning an FCoE map onto an interface

The FCoE map cannot be edited if it is associated with any interfaces.

The FCoE map can be applied, irrespective of whether or not the interface is in 'switchport' mode. But the FCoE map cannot be applied on an interface if the same interface already has a CEE map assigned to it.

To assign the FCoE map onto an interface, perform the following steps in global configuration mode.

1. Activate the interface configuration mode for the interface you wish to modify.

The following example activates the mode for the 10 Gigabit Ethernet interface in slot 0/port 0.

```
switch(config)#interface tengigabitethernet 0/0
```

2. Apply the current FCoE profile map to the interface using the **fcoeport** command.

```
switch(conf-if-te-0/0)# fcoeport
```

3. Return to the privileged EXEC mode using the **end** command.

```
switch(conf-if-te-0/0)#end
```

4. Confirm the changes to the interface with the **show running-config** command.

```
switch#show running-config interface tengigabitethernet 0/0
interface TenGigabitEthernet 0/0
  fcoeport
  shutdown
```

5. Use the **show fcoe-map** command to confirm the current status of the FCoE map.

The following example shows that the FCoE VLAN has been assigned the ID of 5, and the map has been assigned to the interface.

```
switch# show fcoe-map
fcoe-map default
  cee-map      : default
  fcoe-vlan    : 5
  interface(s) : te0/0
```

6. Repeat this procedure for any additional interfaces.

Configuring DCBX

Data Center Bridging eXchange Protocol (DCBX) runs on CEE links and is an extension of the Link Layer Discovery Protocol (LLDP). The primary goal of DCBX is to allow the discovery of CEE-capable hosts and switches and allow CEE-specific parameters—such as those for Enhanced Transmission Selection (ETS) and Priority Flow Control (PFC)—to be sent before the link is shared. DCBX parameters use a type-length-value (TLV) format. By default, DCBX is turned on and the two TLVs are enabled. The available TLVS are:

- dcbx-fcoe-app-tlv—IEEE Data Center Bridging eXchange FCoE Application TLV.
- dcbx-fcoe-logical-link-tlv—IEEE Data Center Bridging eXchange FCoE Logical Link TLV. The presence of this TLV declares that the FCoE part of the converged link is UP.

To configure the TLVs for DCBX, perform the following steps in global configuration mode.

1. Set the protocol type to LLDP.

```
switch(config)#protocol lldp
```

2. Activate the protocol.

```
switch(conf-lldp)#no disable
```

3. Activate the TLV formats using the **advertise** command in LLDP configuration mode.

```
switch(conf-lldp)#advertise dcbx-fcoe-app-tlv
switch(conf-lldp)#advertise dcbx-fcoe-logical-link-tlv
```

4. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(conf-lldp)#do copy running-config startup-config
```

Configuring Spanning Tree Protocol

Spanning Tree Protocol (STP) is a mechanism to detect and avoid loops in Ethernet networks by establishing a fixed path between all the switches in a LAN. The Brocade FCoE hardware supports three spanning tree variations: standard Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP).

NOTE

Configuring Spanning Tree on a port enabled for FCoE is not recommended, as it can cause login issues while the port transitions from the blocked state to the forwarding state. Either STP should be disabled, or configured in edge mode. However, Brocade strongly recommends that STP be disabled.

It is a best practice that an access layer switch does not become the root switch. Changing the bridge or STP priority helps to ensure that this does not occur. The following example performed from the CEE CLI configures the switch for RSTP and sets the bridge priority to the highest value, ensuring it will not become the root switch in an existing LAN.

To configure RSTP, perform the following steps in global configuration mode.

1. Configure the switch for RSTP.

```
switch(config)#protocol spanning-tree rstp
```

2. Set the bridge priority to the highest value so it does not become the root switch in an existing LAN.

```
switch(config-rstp)#bridge-priority 61440
```

3. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(config-rstp)#do copy running-config startup-config
```

Configuring VLAN membership

IEEE 802.1q Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow network traffic isolation into separate virtual networks, reducing the size of administrative and broadcast domains.

A VLAN contains end stations that have a common set of requirements which can be in independent physical locations. You can group end stations in a VLAN even if they are not physically located in the same LAN segment. VLANs are typically associated with IP subnets and all the end stations in a particular IP subnet belong to the same VLAN.

In addition to creating a special VLAN for FCoE traffic, VLAN classifiers are applied to incoming EtherTypes for FCoE Initialization Protocol (FIP) and FCoE. VLAN classifiers are rules used to dynamically classify Ethernet frames on an untagged interface to VLANs.

NOTE

FIP and FCoE cannot be manually configured in VLAN classifier rules. They are part of the fcoeport interface configuration and are applied automatically when an interface is configured as an fcoeport.

To configure VLAN membership, perform the following steps in global configuration mode.

1. Create the VLAN interfaces on the Brocade FCoE hardware using the CEE CLI. For details, see [“Creating a VLAN interface”](#).

The following example creates two VLAN interfaces and assigns each one to a server group.

```
switch(config)#interface vlan 10
switch-cmsh(conf-if-vl-10)#description server group 1
switch(config)#interface vlan 20
switch-cmsh(conf-if-vl-20)#description server group 2
switch(config)#interface vlan 100
switch-cmsh(conf-if-vl-100)#description server group 3
```

2. (Optional) Apply the VLAN classifier group to any CEE interface. For details, see [“Activating a VLAN classifier group with an interface port”](#).
3. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(config)#do copy running-config startup-config
```

CEE map configuration

The CEE Quality of Service (QoS) covers frame classification, priority and traffic class (queue) mapping, congestion control, and scheduling. Under the CEE Provisioning model all of these features are configured utilizing two configuration tables, the Priority Group Table and the Priority Table.

The CEE Priority Group Table defines each Priority Group ID (PGID) and its scheduling policy (Strict Priority versus Deficit Weighted Round-Robin (DWRR), DWRR weight, relative priority), and partially defines the congestion Priority Flow Control (PFC) configuration. There are 16 rows in the CEE Priority Group Table. [Table 11](#) presents the default CEE Priority Group Table configuration.

NOTE

Only a single Class-of-Service (CoS) can be mapped to a PFC-enabled priority queue. The CoS number must be identical to the priority queue number. If your configuration violates this restriction, an error message displays and the Priority Group Table is set back to the default values.

When the CEE map is applied, and the interface is connected to the CNA, only one Strict Priority PGID (PGID 15.0 to PGID 15.7) is allowed.

TABLE 11 Default CEE Priority Group Table configuration

PGID	Bandwidth%	PFC
15.0	—	N
15.1	—	N
15.2	—	N
15.3	—	N
15.4	—	N
15.5	—	N
15.6	—	N
15.7	—	N

TABLE 11 Default CEE Priority Group Table configuration (Continued)

PGID	Bandwidth%	PFC
0	0	N
1	0	N
2	0	N
3	0	N
4	0	N
5	0	N
6	0	N
7	0	N

Strict Priority versus DWRR is derived directly from the PGID value. All PGIDs with prefix 15 receive the Strict Priority scheduling policy and all PGIDs in the range 0 through 7 receive the DWRR scheduling policy. Relative priority between PGIDs is exactly the ordering of entries listed in the table, with PGID 15.0 being the highest priority and PGID 0 being the lowest priority. Congestion control configuration is partially specified by toggling the PFC column On or Off. This provides only partial configuration of congestion control because the set of priorities mapped to the PGID is not known.

The CEE Priority Table defines each CoS mapping to a PGID, and completes the PFC configuration. There are eight rows in the CEE Priority Table. [Table 12](#) details the default CEE Priority Table configuration.

TABLE 12 Default CEE priority table configuration

CoS	PGID
0	15.6
1	15.7
2	15.5
3	15.4
4	15.3
5	15.2
6	15.1
7	15.0

Defining a Priority Group Table map

To define a priority group table map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Activate the CEE map mode.

```
switch(config)#cee-map default
```

3. Define the CEE map for PGID 0.

```
switch(config-ceemap)#priority-group-table 0 weight 50 pfc
```

4. Define the CEE map for PGID 1.

```
switch(config-ceemap)#priority-group-table 1 weight 50
```

5. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(config)#do copy running-config startup-config
```

Defining a Priority-Table map

To define a Priority Table map, perform the following steps in privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Activate the CEE map mode.

```
switch(config)#cee-map default
```

3. Define the map.

```
switch(config)#priority-table 2 2 2 1 2 2 2 2
```

4. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(config)#do copy running-config startup-config
```

Applying a CEE provisioning map to an interface

To apply a CEE provisioning map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Specify the 10 Gigabit Ethernet interface. In this example, 0/2 is used.

```
switch(config)#interface tengigabitethernet 0/2
```

3. Apply the CEE map on the interface.

```
switch(conf-if-te-0/1)#cee-map default
```

4. Return to privileged EXEC mode.

```
switch(conf-if-te-0/1)#exit
```

5. Enter the **copy** command to save the **running-config** file to the **startup-config** file.

```
switch(config)#do copy running-config startup-config
```

Verifying the CEE maps

To verify the CEE map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Verify the CEE map provisioning for a specified name.

```
switch(config)#show cee maps default
```

CEE audit logs

The CEE audit logs record all changes made by authorized users. This information is always present, irrespective of the configuring entity. Because redundant configuration is not recorded, only configuration changes that have resulted in some change in the running-config file on the system are logged.

In addition to configuration commands, the **copy** and **clear** commands, which alter the system state, are logged. On the local system, the logs are maintained in a circular buffer with a capacity of 256 entries. The audit log infrastructure of the Fabric OS is used here and the capacity of 256 entries is shared between the Fabric OS and CEE audit logs.

Table 3 shows the components of the CEE audit logs.

TABLE 13 Format of the CEE audit logs

Name	Example	Description
Timestamp	2009/05/12-21:25:57	Date and time at which the log was generated
User	Root	User's login name
Role	Root	User's RBAC role
IP	10.2.2.47	IP address of the user's computer
Session	telnet	The virtual console from which the user executed the command
Hostname	EL34	Host name of the switch
Mode	Config	The configuration mode in which the command was executed
Command	rmon event 15	The CLI command that altered the configuration

To configure CEE audit logs, perform these steps from the Fabric OS command shell.

1. CEE audit logging is disabled by default. It must be enabled explicitly by using the **auditcfg** command from the Fabric OS shell. See the *Fabric OS Command Reference Manual Supporting Fabric OS v6.4.0* for details.

```
switch>auditcfg --enable
```

2. The CEE audit logs belong to the CONFIGURATION class of audit logs. The filter for this class must be enabled by using the **auditcfg** command from the Fabric OS shell. See the *Fabric OS Command Reference Manual Supporting Fabric OS v6.4.0* for details.

```
switch>auditcfg --class 3
```

3. Activate privileged EXEC mode.

```
switch>cmsh
switch#enable
switch#
```


4. Review the DCE audit logs using the **show accounting logs** command.

```
switch#show accounting logs
2009/06/24-02:15:54(UTC),root/root/NONE/console/CLI/EL34,config;rmon event 12
2009/06/24-02:21:13(UTC),root/root/NONE/console/CLI/EL34,config;rmon event 7
2009/06/24-02:21:29(UTC),root/root/NONE/console/CLI/EL34,conf-if-te-0/0;rmon
collection stats 100
```

5 CEE audit logs

Configuring VLANs

In this chapter

- [VLAN overview](#) 45
- [Ingress VLAN filtering](#) 45
- [VLAN configuration guidelines and restrictions](#) 47
- [Default VLAN configuration](#) 47
- [VLAN configuration and management](#) 48
- [Configuring VLAN classifier rules](#) 52
- [Configuring the MAC address table](#) 54

VLAN overview

IEEE 802.1Q Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow you to isolate network traffic between virtual networks and reduce the size of administrative and broadcast domains.

A VLAN contains end stations that have a common set of requirements that are independent of physical location. You can group end stations in a VLAN even if they are not physically located in the same LAN segment. VLANs are typically associated with IP subnetworks and all the end stations in a particular IP subnet belong to the same VLAN. Traffic between VLANs must be routed. VLAN membership is configurable on a per interface basis.

The VLAN used for carrying FCoE traffic needs to be explicitly designated as the FCoE VLAN. FCoE VLANs are configured through the CEE CLI (see [“Configuring a VLAN interface to forward FCoE traffic”](#) on page 49).

NOTE

Currently only one VLAN can be configured as the FCoE VLAN.

Ingress VLAN filtering

A frame arriving at Brocade FCoE hardware is either associated with a specific port or with a VLAN, based on whether the frame is tagged or untagged:

- **Admit tagged frames only**—The port the frame came in on is assigned to a single VLAN or to multiple VLANs depending on the VLAN ID in the frame’s VLAN tag. This is called trunk mode.
- **Admit untagged frames only**—These frames are assigned the port VLAN ID (PVID) assigned to the port the frame came in on. This is called access mode.

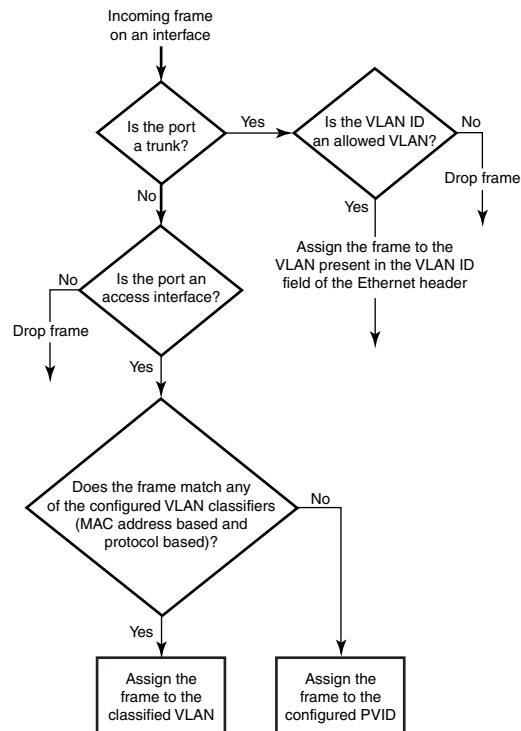
- Admit VLAN tagged and untagged frames—All tagged and untagged frames would be processed as follows:
 - All untagged frames are classified into native VLANs.
 - All frames egressing are untagged for the native VLANs.
 - Any tagged frames coming with a VLAN tag equal to the configured native VLAN are processed.
 - For ingress and egress, non-native VLAN tagged frames are processed according to the allowed VLAN user specifications. This is called converged mode.

NOTE

Ingress VLAN filtering is enabled by default on all Layer 2 interfaces. This ensures that VLANs are filtered on the incoming port (depending on the user configuration).

Figure 4 displays the frame processing logic for an incoming frame.

FIGURE 4 Ingress VLAN filtering



There are important facts you should know about Ingress VLAN filtering:

- Ingress VLAN filtering is based on port VLAN membership.
- Port VLAN membership is configured through the CEE CLI.
- Dynamic VLAN registration is not supported.
- The Brocade FCoE hardware does VLAN filtering at both the ingress and egress ports.
- The VLAN filtering behavior on logical Layer 2 interfaces such as LAG interfaces is the same as on port interfaces.
- The VLAN filtering database (FDB) determines the forwarding of an incoming frame.

Additionally, there are important facts you should know about the VLAN FDB:

- The VLAN FDB contains information that helps determine the forwarding of an arriving frame based on MAC address and VLAN ID data. The FDB contains both statically configured data and dynamic data that is learned by the switch.
- The dynamic updating of FDB entries using learning is supported (if the port state permits).
- Dynamic FDB entries are not created for multicast group addresses.
- Dynamic FDB entries are aged out based on the aging time configured per Brocade FCoE hardware. The aging time is between 10 and 100000 seconds. The default is 300 seconds.
- You can add static MAC address entries specifying a VLAN ID. Static entries are not aged out.
- A static FDB entry does not overwrite a new or pre-existing dynamic entry.
- Static entries do not disable the learning of dynamic entries.

NOTE

For more information on frame handling for Brocade FCoE hardware, see [“Layer 2 Ethernet overview”](#) on page 3.

VLAN configuration guidelines and restrictions

Follow these VLAN configuration guidelines and restrictions when configuring VLANs.

- In an active topology, MAC addresses can be learned, per VLAN, using Independent VLAN Learning (IVL) only.
- A MAC address ACL always overrides a static MAC address entry. In this case, the MAC address is the forwarding address and the forwarding entry can be overwritten by the ACL.
- The Brocade CEE switch supports Ethernet DIX frames and 802.2 LLC SNAP encapsulated frames only.

Default VLAN configuration

[Table 14](#) lists the default VLAN configuration.

TABLE 14 Default VLAN configuration

Parameter	Default setting
Default VLAN	VLAN 1
Interface VLAN assignment	All interfaces assigned to VLAN 1
VLAN state	Active
MTU size	2500 bytes

VLAN configuration and management

NOTE

You need to enter either the **copy running-config startup-config** command or the **write memory** command to save your configuration changes to Flash so that they are not lost if there is a system reload or power outage.

Enabling and disabling an interface port

CEE interfaces do not support auto-negotiation of Ethernet link speeds. The CEE interfaces only support 10-Gigabit Ethernet. CEE interfaces are disabled by default.

To enable and disable an interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **shutdown** command to toggle the availability of the interface.

To enable the CEE interface:

```
switch(conf-if-te-0/1)#no shutdown
```

To disable the CEE interface:

```
switch(conf-if-te-0/1)#shutdown
```

Configuring the MTU on an interface port

To configure the maximum transmission unit (MTU) on an interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the interface port type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the interface port.

4. Enter the **mtu** command to specify the MTU value on the interface port.

The range of valid values is from 1522 through 9208, and the default size is 2500.

```
switch(conf-if-te-0/1)#mtu 4200
```

Creating a VLAN interface

On Brocade FCoE hardware, VLANs are treated as interfaces from a configuration point of view.

By default all the CEE ports are assigned to VLAN 1 (VLAN ID equals 1). The *vlan_ID* value can be 1 through 3583. VLAN IDs 3584 through 4094 are internally-reserved VLAN IDs.

To create a VLAN interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface vlan** command to assign the VLAN interface number.

```
switch(config)#interface vlan 100
```

Enabling STP on a VLAN

Once all of the interface ports have been configured for a VLAN, you can enable spanning tree protocol (STP) for all members of the VLAN with a single command. Whichever protocol is currently selected is used by the VLAN. Only one type of STP can be active at a time.

A physical interface port can be a member of multiple VLANs. For example, a physical port can be a member of VLAN 100 and VLAN 55 simultaneously. In addition, VLAN 100 can have STP enabled and VLAN 55 can have STP disabled simultaneously.

To enable STP for a VLAN, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol spanning tree** command to select the type of STP for the VLAN.

```
switch(config)#protocol spanning tree mstp
```

3. Enter the **interface** command to select the VLAN interface number.

```
switch(config)#interface vlan 100
```

4. Enter the **no spanning-tree shutdown** command to enable spanning tree on VLAN 100. STP is active by default, but enter the command to ensure activation.

```
switch(conf-if-vl-100)#no spanning-tree shutdown
```

Disabling STP on a VLAN

Once all of the interface ports have been configured for a VLAN, you can disable STP for all members of the VLAN with a single command.

To disable STP for a VLAN, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to select the VLAN interface number.

```
switch(config)#interface vlan 55
```

3. Enter the **spanning-tree shutdown** command to disable spanning tree on VLAN 55.

```
switch(conf-if-vl-55)#spanning-tree shutdown
```

Configuring a VLAN interface to forward FCoE traffic

An FCoE Forwarder (FCF) is an FCoE device that supports FCoE VF_ports. It is the equivalent of an FC switch. A VLAN can be made FCF-capable. Only FCF-capable VLANs can carry FCoE traffic.

To configure a VLAN interface to forward FCoE traffic, see [“Configuring the FCoE VLAN”](#) on page 35.

Configuring an interface port as a Layer 2 switch port

To configure the interface as a Layer 2 switch port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```
3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **switchport** command to configure the interface as a Layer 2 switch port.
5. Enter the **do show** command to confirm the status of the CEE interface. For example

```
switch(conf-if-te-0/1)#do show interface tengigabitethernet 0/1
```
6. Enter the **do show** command to confirm the status of the CEE interface running configuration.

```
switch(conf-if-te-0/1)#do show running-config interface tengigabitethernet 0/1
```

Configuring an interface port as an access interface

Each CEE interface port supports admission policies based on whether the frames are untagged or tagged. Access mode admits only untagged and priority-tagged frames.

To configure the interface as an access interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```
3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **switchport access** command to configure the CEE interface as a Layer 2 switch port.

```
switch(conf-if-te-0/1)#switchport access vlan 20
```
5. Enter the **switchport mode** command to set the mode of the interface to access.

```
switch(conf-if-te-0/1)#switchport mode access
```

Configuring an interface port as a trunk interface

Each CEE interface port supports admission policies based on whether the frames are untagged or tagged. Trunk mode admits only VLAN-tagged frames.

To configure the interface as a trunk interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/19
```
3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **switchport mode trunk** command to place the CEE interface into trunk mode.


```
switch(conf-if-te-0/19)#switchport mode trunk
```

5. Specify whether all, one, or none of the VLAN interfaces are allowed to transmit and receive through the CEE interface. Enter the following command that is appropriate for your needs.

- This example allows the VLAN numbered as 30 to transmit/receive through the CEE interface:

```
switch(conf-if-te-0/19)#switchport trunk allowed vlan add 30
```

- To allow all VLANs to transmit/receive through the CEE interface:

```
switch(conf-if-te-0/19)#switchport trunk allowed vlan all
```

- This example allows all except VLAN 11 to transmit/receive through the CEE interface:

```
switch(conf-if-te-0/19)#switchport trunk allowed vlan except 11
```

- To allow none of the VLANs to transmit/receive through the CEE interface:

```
switch(conf-if-te-0/19)#switchport trunk allowed vlan none
```

Disabling a VLAN on a trunk interface

To disable a VLAN on a trunk interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/10
```

3. Enter the **no shutdown** command to enable the CEE interface.

4. Enter the **switchport** command to place the CEE interface into trunk mode.

```
switch(conf-if-te-0/10)#switchport mode trunk none
```

5. Enter the **switchport trunk allowed vlan remove vlan_id** command to disable a vlan on a trunk interface.

```
switch(conf-if-te-0/10)#switchport trunk allowed vlan remove 100
```

Configuring an interface port as a converged interface

Each CEE interface port supports admission policies based on whether the frames are untagged or tagged. Converged mode admits both tagged and untagged frames. Frames egressing are tagged automatically.

To configure the interface as converged interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.

4. Enter the **switchport** command to set the tagged VLAN on the interface to 100.

```
switch(conf-if-te-0/1)#switchport converged allowed vlan add 100
```

5. Enter the **switchport mode converged** command to configure the interface as converged port.

```
switch(conf-if-te-0/1)#switchport mode converged
```

6. Enter the **switchport converged allowed vlan add *vlan_id*** command to give access to individual VLANs.

```
switch(conf-if-te-0/1)#switchport converged allowed vlan add 200
switch(conf-if-te-0/1)#switchport converged allowed vlan add 300
switch(conf-if-te-0/1)#switchport converged allowed vlan add 400
```

Disabling a VLAN on a converged interface

To disable a VLAN on a converged interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/10
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **switchport** command to place the CEE interface into converged mode.

```
switch(conf-if-te-0/10)#switchport mode converged none
```

5. Enter the **switchport trunk allowed vlan remove *vlan_id*** command to disable a VLAN on a converged interface.

```
switch(conf-if-te-0/10)#switchport trunk allowed vlan remove 300
```

Configuring VLAN classifier rules

You can configure VLAN classifier rules to define specific rules for classifying frames to selected VLANs based on protocol and MAC addresses. Sets of rules can be grouped into VLAN classifier groups (see [“Creating a VLAN classifier group and adding rules”](#) on page 53).

VLAN classifier rules (1 through 256) are a set of configurable rules that reside in one of these categories:

- 802.1Q protocol-based classifier rules
- Source MAC address-based classifier rules
- Encapsulated Ethernet classifier rules

NOTE

Multiple VLAN classifier rules can be applied per interface provided the resulting VLAN IDs are unique for the different rules.

802.1Q protocol-based VLANs apply only to untagged frames, or frames with priority tagging.

With both Ethernet-II and 802.2 SNAP encapsulated frames, the following protocol types are supported by the **vlan classifier rule *number* protocol** command:

- Ethernet hexadecimal (0x0000 through 0xffff)
- Address Resolution Protocol (ARP)
- FCoE
- FIP
- IP version 4 and IP version 6 (IPv6)

For complete information on all available VLAN classifier rule options, see the *Converged Enhanced Ethernet Command Reference*.

Configuring a VLAN classifier rule

To configure a protocol-based VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **vlan classifier rule** command to configure a protocol-based VLAN classifier rule.

```
switch(config)#vlan classifier rule 1 proto ip encap ethv2
```

Configuring MAC address-based VLAN classifier rules

To configure a MAC address-based VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **vlan classifier rule** command to configure a MAC address-based VLAN classifier rule.

```
switch(config)#vlan classifier rule 5 mac 0008.744c.7fid
```

Deleting a VLAN classifier rule

VLAN classifier groups (1 through 16) can contain any number of VLAN classifier rules.

To configure a VLAN classifier group and remove a VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify a VLAN classifier group and delete a rule.

```
switch(config)#no vlan classifier rule 1
```

Creating a VLAN classifier group and adding rules

VLAN classifier groups (1 through 16) can contain any number of VLAN classifier rules.

To configure a VLAN classifier group and add a VLAN classifier rule, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Create a VLAN classifier group and add a rule.

```
switch(config)#vlan classifier group 1 add rule 1
```

Activating a VLAN classifier group with an interface port

To associate a VLAN classifier group with an interface port, perform the following steps from privileged EXEC mode.

6 Configuring the MAC address table

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.
3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **vlan classifier** command to activate and associate it with a VLAN interface (group 1 and VLAN 2 are used in this example).

```
switch(config)#interface tengigabitethernet 0/10
```

```
switch(conf-if-te-0/10)#vlan classifier activate group 1 vlan 2
```

NOTE

This example assumes that VLAN 2 was already created.

Clearing VLAN counter statistics

To clear VLAN counter statistics, perform the following steps from privileged EXEC mode.

1. Enter the **clear** command to clear the VLAN counter statistics for the specified VLAN. The `vlan_ID` value can be 1 through 3583. For example, to clear the counter for VLAN 20:

```
switch#clear counter interface vlan 20
```

Displaying VLAN information

To display VLAN information, perform the following steps from privileged EXEC mode.

1. Enter the **show interface** command to display the configuration and status of the specified interface.

```
switch#show interface tengigabitethernet 0/10 switchport
```

2. Enter the **show vlan** command to display the specified VLAN information. For example, this syntax displays the status of VLAN 20 for all interfaces, including static and dynamic:

```
switch#show vlan 20
```

Configuring the MAC address table

Each CEE port has a MAC address table. The MAC address table stores a number of unicast and multicast address entries without flooding any frames. Brocade FCoE hardware has a configurable aging timer. If a MAC address remains inactive for a specified number of seconds, it is removed from the address table. For detailed information on how the switch handles MAC addresses in a Layer 2 Ethernet environment, see [“Layer 2 Ethernet overview”](#) on page 3.

Specifying or disabling the aging time for MAC addresses

You can set the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. Static address entries are never aged or removed from the table. You can also disable the aging time. The default is 300 seconds.

NOTE

To disable the aging time for MAC addresses, enter an aging time value of 0.

To specify an aging time or disable the aging time for MAC addresses, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the appropriate command based on whether you want to specify an aging time or disable the aging time for MAC addresses:

```
switch(config)#mac-address-table aging-time 600
```

Adding static addresses to the MAC address table

To add a static address to the MAC address table, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Add the static address 0011.2222.3333 to the MAC address table with a packet received on VLAN 100:

```
switch(config)#mac-address-table static 0011.2222.3333 forward  
tengigabitethernet 0/1 vlan 100
```

6 Configuring the MAC address table

Configuring STP, RSTP, and MSTP

In this chapter

- STP overview 57
- RSTP overview 59
- MSTP overview 61
- STP, RSTP, and MSTP configuration guidelines and restrictions 63
- Default STP, RSTP, and MSTP configuration 63
- STP, RSTP, and MSTP configuration and management 64
- Configuring STP, RSTP, or MSTP on CEE interface ports 72

STP overview

The IEEE 802.1D Spanning Tree Protocol (STP) runs on bridges and switches that are 802.1D-compliant. STP prevents loops in the network by providing redundant links. If a primary link fails, the backup link is activated and network traffic is not affected. Without STP running on the switch or bridge, a link failure can result in a loop.

When the spanning tree algorithm is run, the network switches transform the real network topology into a spanning tree topology in which any LAN in the network can be reached from any other LAN through a unique path. The network switches recalculate a new spanning tree topology whenever there is a change to the network topology.

For each LAN, the switches that attach to the LAN choose a designated switch that is the closest switch to the root switch. This designated switch is responsible for forwarding all traffic to and from the LAN. The port on the designated switch that connects to the LAN is called the designated port.

The switches decide which of their ports will be part of the spanning tree. A port is included in the spanning tree if it is a root port or a designated port.

With STP, data traffic is allowed only on those ports that are part of the spanning tree topology. Ports that are not part of the spanning tree topology are automatically changed to a blocking (inactive) state. They are kept in the blocking state until there is a break in the spanning tree topology, at which time they are automatically activated to provide a new path.

The STP interface states for every Layer 2 interface running STP are as follows:

- Blocking—The interface does not forward frames.
- Listening—The interface is identified by the spanning tree as one that should participate in frame forwarding. This is a transitional state after the blocking state.
- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.

- Disabled—The interface is not participating in spanning tree because of a shutdown port, no link on the port, or no spanning tree instance running on the port.

A port participating in spanning tree moves through these states:

- From initialization to blocking state.
- From blocking to listening or disabled state.
- From listening to learning, blocking, or disabled state.
- From learning to forwarding, blocking, or disabled state.
- From forwarding to blocking or disabled state.

The following STP features are considered optional features although you might use them in your STP configuration:

- Root guard—For detailed information, see [“Enabling the guard root”](#) on page 73.
- PortFast BPDU guard and BPDU filter—For detailed information, see [“Enabling port fast \(STP\)”](#) on page 75.

Configuring STP

To configure STP, perform the following steps from global configuration mode.

1. Enable RSTP using the global **protocol spanning-tree** command. For details, see [“Enabling STP, RSTP, or MSTP”](#) on page 64.

```
switch(config)#protocol spanning-tree stp
```

2. Designate the root switch using the **bridge-priority** command. For details, see [“Specifying the bridge priority”](#) on page 65. The range is 0 through 61440 and the priority values can be set only in increments of 4096.

```
switch(conf-stp)#bridge-priority 28672
```

3. Enable PortFast on switch ports using the **spanning-tree portfast** command. For details, see [“Enabling port fast \(STP\)”](#) on page 75. Note that this step is optional.

NOTE

PortFast only needs to be enabled on ports that connect to workstations or PCs. Repeat these commands for every port connected to workstations or PCs. Do not enable PortFast on ports that connect to other switches.

```
switch(config)#interface tengigabitethernet 0/10
switch(conf-if-te-0/10)#spanning-tree portfast
switch(conf-if-te-0/10)#exit
switch(config)#interface tengigabitethernet 0/11
switch(conf-if-te-0/11)#spanning-tree portfast
switch(conf-if-te-0/11)#exit
```

Repeat these commands for every port connected to workstations or PCs.

4. *Optional:* Enable the guard root feature with the **spanning-tree guard root** command.

The guard root feature provides a way to enforce the root bridge placement in the network. For detailed information, refer to [“Enabling the guard root”](#) on page 73.

All other switch ports connect to other switches and bridges are automatically placed in blocking mode.

This does not apply to ports connected to workstations or PCs; these ports remain in the forwarding state.

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

When the spanning tree topology is completed, the network switches send and receive data only on the ports that are part of the spanning tree. Ports that receive superior BPDUs are automatically placed in blocking mode.

NOTE

Brocade recommends leaving other STP variables at their default values.

For more information on STP, see [“STP, RSTP, and MSTP configuration and management”](#) on page 64.

RSTP overview

NOTE

RSTP is designed to be compatible and interoperate with STP. However, the advantages of the RSTP fast reconvergence are lost when it interoperates with switches running STP.

The IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) standard is an evolution of the 802.1D STP standard. It provides rapid reconvergence following the failure of a switch, a switch port, or a LAN. It provides rapid reconvergence of edge ports, new root ports, and ports connected through point-to-point links.

The RSTP interface states for every Layer 2 interface running RSTP are as follows:

- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Discarding—The interface discards frames. Note that the 802.1D disabled, blocking, and listening states are merged into the RSTP discarding state. Ports in the discarding state do not take part in the active topology and do not learn MAC addresses.

[Table 15](#) lists the interface state changes between STP and RSTP.

TABLE 15 STP versus RSTP state comparison

STP interface state	RSTP interface state	Is the interface included in the active topology?	Is the interface learning MAC addresses?
Disabled	Discarding	No	No
Blocking	Discarding	No	No
Listening	Discarding	Yes	No
Learning	Learning	Yes	Yes
Forwarding	Forwarding	Yes	Yes

With RSTP, the port roles for the new interface states are also different. RSTP differentiates explicitly between the state of the port and the role it plays in the topology. RSTP uses the root port and designated port roles defined by STP, but splits the blocked port role into backup port and alternate port roles:

- Backup port—Provides a backup for the designated port and can only exist where two or more ports of the switch are connected to the same LAN; the LAN where the bridge serves as a designated switch.
- Alternate port—Serves as an alternate port for the root port providing a redundant path towards the root bridge.

Only the root port and the designated ports are part of the active topology; the alternate and backup ports do not participate in it.

When the network is stable, the root and the designated ports are in the forwarding state, while the the alternate and backup ports are in the discarding state. When there is a topology change, the new RSTP port roles allow a faster transition of an alternate port into the forwarding state.

For more information on RSTP, see [“STP, RSTP, and MSTP configuration and management”](#) on page 64.

Configuring RSTP

To configure RSTP, perform the following steps from global configuration mode.

1. Enable RSTP using the global **protocol spanning-tree** command. For details, see [“Enabling STP, RSTP, or MSTP”](#) on page 64.

```
switch(config)#protocol spanning-tree rstp
```

2. Designate the root switch using the **bridge-priority** command. For details, see [“Specifying the bridge priority”](#) on page 65. The range is 0 through 61440 and the priority values can be set only in increments of 4096.

```
switch(conf-rstp)#bridge-priority 28672
```

3. Configure the **bridge forward delay** value. For details, see [“Specifying the bridge forward delay”](#) on page 66.

```
switch(conf-rstp)#forward-delay 20
```

4. Configure the **bridge maximum aging time** value. For details, see [“Specifying the bridge maximum aging time”](#) on page 66.

```
switch(conf-rstp)#max-age 25
```

5. Enable the **error disable timeout timer** value. For details, see [“Enabling the error disable timeout timer”](#) on page 67.

```
switch(conf-rstp)#error-disable-timeout enable
```

6. Configure the **error-disable-timeout** interval value. For details, see [“Specifying the error disable timeout interval”](#) on page 67.

7. `switch(conf-rstp)#error-disable-timeout interval 60`

8. Configure the port-channel path cost method. For details, see [“Specifying the port-channel path cost”](#) on page 67.

```
switch(conf-rstp)#port-channel path-cost custom
```

9. Configure the bridge hello time value. For details, see [“Specifying the bridge hello time \(STP and RSTP\)”](#) on page 68.

```
switch(config-rstp)#hello-time 5
```

10. Set the method for flushing mac-addresses in FDB, either according to the IEEE standard or the propriety standard. For details, see [“Flushing MAC addresses \(RSTP and MSTP\)”](#) on page 71.

```
switch(config)#spanning-tree tc-flush-standard
```

11. Enable PortFast on switch ports using the **spanning-tree edgeport** command. For details, see [“Enabling port fast \(STP\)”](#) on page 75. Note that this step is optional.

NOTE

Edgeport only needs to be enabled on ports that connect to workstations or PCs. Repeat these commands for every port connected to workstations or PCs. Do not enable edgeport on ports that connect to other switches.

```
switch(config)#interface tengigabitethernet 0/10
switch(config-if-te-0/10)#spanning-tree edgeport
switch(config-if-te-0/10)#exit
switch(config)#interface tengigabitethernet 0/11
switch(config-if-te-0/11)#spanning-tree edgeport
switch(config-if-te-0/11)#exit
```

Repeat these commands for every port connected to workstations or PCs.

12. Enable the guard root feature with the **spanning-tree guard root** command. The guard root feature provides a way to enforce the root bridge placement in the network. For detailed information, refer to [“Enabling the guard root”](#) on page 73. Note that this step is optional.

Switch ports that receive superior BPDUs are automatically placed in blocking mode.

This does not apply to ports connected to workstations or PCs; these ports remain in the forwarding state.

13. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config-if-te-0/1)#do copy running-config startup-config
```

MSTP overview

The IEEE 802.1s Multiple STP (MSTP) helps create multiple loop-free active topologies on a single physical topology. MSTP enables multiple VLANs to be mapped to the same spanning tree instance (forwarding path), which reduces the number of spanning tree instances needed to support a large number of VLANs. Each MSTP instance has a spanning tree topology independent of other spanning tree instances. With MSTP you can have multiple forwarding paths for data traffic. A failure in one instance does not affect other instances. With MSTP, you are able to more effectively utilize the physical resources present in the network and achieve better load balancing of VLAN traffic.

NOTE

In MSTP mode, RSTP is automatically enabled to provide rapid convergence.

Multiple switches must be configured consistently with the same MSTP configuration to participate in multiple spanning tree instances. A group of interconnected switches that have the same MSTP configuration is called an MSTP region.

NOTE

Brocade supports 15 MSTP instances and one MSTP region.

MSTP introduces a hierarchical way of managing switch domains using regions. Switches that share common MSTP configuration attributes belong to a region. The MSTP configuration determines the MSTP region where each switch resides. The common MSTP configuration attributes are as follows:

- Alphanumeric configuration name (32 bytes)
- Configuration revision number (2 bytes)
- 4096-element table that maps each of the VLANs to an MSTP instance

Region boundaries are determined based on the above attributes. A multiple spanning tree instance is an RSTP instance that operates inside an MSTP region and determines the active topology for the set of VLANs mapping to that instance. Every region has a common internal spanning tree (CIST) that forms a single spanning tree instance that includes all the switches in the region. The difference between the CIST instance and the MSTP instance is that the CIST instance operates across the MSTP region and forms a loop-free topology across regions, while the MSTP instance operates only within a region. The CIST instance can operate using RSTP if all the switches across the regions support RSTP. However, if any of the switches operate using 802.1D STP, the CIST instance reverts to 802.1D. Each region is viewed logically as a single STP/RSTP bridge to other regions.

Configuring MSTP

To configure MSTP, perform the following steps from global configuration mode.

1. Enable MSTP using the global **protocol spanning-tree** command. For more details see [“Enabling STP, RSTP, or MSTP”](#) on page 64.

```
switch(config)#protocol spanning-tree mstp
```

2. Specify the region name using the **region** *region_name* command. For more details see [“Specifying a name for an MSTP region”](#) on page 70.

```
switch(conf-mstp)#region brocade1
```

3. Specify the revision number using the **revision** command. For more details see [“Specifying a revision number for an MSTP configuration”](#) on page 70.

```
switch(conf-mstp)#revision 1
```

4. Map a VLAN to an MSTP instance using the **instance** command. For more details see [“Mapping a VLAN to an MSTP instance”](#) on page 69.

```
switch(conf-mstp)#instance 1 vlan 2, 3  
switch(conf-mstp)#instance 2 vlan 4-6  
switch(conf-mstp)#instance 1 priority 4096
```

5. Specify the maximum hops for a BPDU to prevent the messages from looping indefinitely on the interface using the **max-hops** *hop_count* command. For more details see [“Specifying the maximum number of hops for a BPDU \(MSTP\)”](#) on page 70.

```
switch(conf-mstp)#max-hops 25
```

6. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(conf-mstp)#do copy running-config startup-config
```

For more information on MSTP, see [“STP, RSTP, and MSTP configuration and management”](#) on page 64.

STP, RSTP, and MSTP configuration guidelines and restrictions

Follow these configuration guidelines and restrictions when configuring STP, RSTP, and MSTP:

- You have to disable one form of xSTP before enabling another.
- Packet drops or packet flooding may occur if you do not enable xSTP on all devices connected on both sides of parallel links.
- LAGs are treated as normal links and by default are enabled for STP.
- You can have 15 MSTP instances and one MSTP region.
- Create VLANs before mapping them to MSTP instances.
- The MSTP force-version option is not supported.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.
- When you enable MSTP by using the global **protocol spanning-tree mstp** command, RSTP is automatically enabled.
- For two or more switches to be in the same MSTP region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same name.
- Spanning Tree topologies must not be enabled on any direct server connections to the front-end Ten Gigabit Ethernet ports that may run FCoE traffic. This may result in lost or dropped FCoE logins.

Default STP, RSTP, and MSTP configuration

[Table 16](#) lists the default STP, RSTP, and MSTP configuration.

TABLE 16 Default STP, RSTP, and MSTP configuration

Parameter	Default setting
Spanning-tree mode	By default, STP, RSTP, and MSTP are disabled
Bridge priority	32768
Bridge forward delay	15 seconds
Bridge maximum aging time	20 seconds
Error disable timeout timer	Disabled
Error disable timeout interval	300 seconds
Port-channel path cost	Standard
Bridge hello time	2 seconds
Flush MAC addresses from the VLAN FDB	Enabled

Table 17 lists the switch defaults that apply only to MSTP configurations.

TABLE 17 Default MSTP configuration

Parameter	Default setting
Cisco interoperability	Disabled
Switch priority (when mapping a VLAN to an MSTP instance)	32768
Maximum hops	20 hops
Revision number	0

Table 18 lists the switch defaults for the 10-Gigabit Ethernet CEE interface-specific configuration.

TABLE 18 Default 10-Gigabit Ethernet CEE interface-specific configuration

Parameter	Default setting
Spanning tree	Disabled on the interface
Automatic edge detection	Disabled
Path cost	2000
Edge port	Disabled
Guard root	Disabled
Hello time	2 seconds
Link type	Point-to-point
Port fast	Disabled
Port priority	128
CEE interface root port	Allow the CEE interface to become a root port.
CEE interface BPDU restriction	Restriction is disabled

STP, RSTP, and MSTP configuration and management

NOTE

You need to enter either the **copy running-config startup-config** command or the **write memory** command to save your configuration changes to Flash so that they are not lost if there is a system reload or power outage.

Enabling STP, RSTP, or MSTP

You enable STP to detect or avoid loops. STP is not required in a loop-free topology. You must turn off one form of STP before turning on another form. By default, STP, RSTP, and MSTP are not enabled.

To enable xSTP, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree rstp
```

Disabling STP, RSTP, or MSTP

NOTE

Using the **no protocol spanning-tree** command deletes the context and all the configurations defined within the context or protocol for the interface, unless the spanning tree is explicitly shutdown on the interface by the **spanning-tree shutdown** command.

To disable STP, RSTP, or MSTP, perform the following steps from privileged EXEC mode. By default, STP, RSTP, and MSTP are not enabled.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to disable STP, RSTP, or MSTP.

```
switch(config)#no protocol spanning-tree
```

Shutting down STP, RSTP, or MSTP globally

To shut down STP, RSTP, or MSTP globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **shutdown** command to globally shutdown STP, RSTP, or MSTP. The **shutdown** command below works in all three modes.

```
switch(conf-mstp)#shutdown
```

Specifying the bridge priority

In any mode (STP, RSTP, or MSTP), use this command to specify the priority of the switch. After you decide on the root switch, set the appropriate values to designate the switch as the root switch. If a switch has a bridge priority that is lower than all the other switches, the other switches automatically select the switch as the root switch.

The root switch should be centrally located and not in a “disruptive” location. Backbone switches typically serve as the root switch because they often do not connect to end stations. All other decisions in the network, such as which port to block and which port to put in forwarding mode, are made from the perspective of the root switch.

Bridge protocol data units (BPDUs) carry the information exchanged between switches. When all the switches in the network are powered up, they start the process of selecting the root switch. Each switch transmits a BPDU to directly connected switches. Each switch compares the received BPDU to the BPDU that the switch sent. In the root switch selection process, if switch 1 advertises a root ID that is a lower number than the root ID that switch 2 advertises, switch 2 stops the advertisement of its root ID, and accepts the root ID of switch 1. The switch with the lowest bridge priority becomes the root switch.

NOTE

Because each VLAN is in a separate broadcast domain, each VLAN must have its own root switch.

To specify the bridge priority, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.
3. Specify the bridge priority. The range is 0 through 61440 and the priority values can be set only in increments of 4096. The default priority is 32678.

```
switch(config)#protocol spanning-tree rstp
```

```
switch(conf-stp)#bridge-priority 20480
```

Specifying the bridge forward delay

In any mode (STP, RSTP, or MSTP), use this command to specify how long an interface remains in the listening and learning states before the interface begins forwarding all spanning tree instances.

The range is 4 through 30 seconds. The default is 15 seconds. The following relationship should be kept:

$$2 * (\text{forward_delay} - 1) \geq \text{max_age} \geq 2 * (\text{hello_time} + 1)$$

To specify the bridge forward delay, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Specify the bridge forward delay.

```
switch(conf-stp)#forward-delay 20
```

Specifying the bridge maximum aging time

In any mode (STP, RSTP, or MSTP), use this command to control the maximum length of time that passes before an interface saves its Bridge Protocol Data Unit (BPDU) configuration information.

When configuring the maximum aging time, the max-age setting must be greater than the hello-time setting. The range is 6 through 40 seconds. The default is 20 seconds. The following relationship should be kept:

$$2 * (\text{forward_delay} - 1) \geq \text{max_age} \geq 2 * (\text{hello_time} + 1)$$

To specify the bridge maximum aging time, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Specify the bridge maximum aging time.

```
switch(conf-stp)##max-age 25
```

Enabling the error disable timeout timer

In any mode (STP, RSTP, or MSTP), use this command to specify the time in seconds required for the timer to expire. When the STP BPDU guard disables a port, the port remains in the disabled state unless the port is enabled manually. This command allows you to enable the port from the disabled state. For details on configuring the error disable timeout interval, see [“Specifying the error disable timeout interval”](#) on page 67. By default, the timeout feature is disabled.

To enable the error disable timeout timer, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Enable the error disable timeout timer.

```
switch(conf-stp)#error-disable-timeout enable
```

Specifying the error disable timeout interval

In any mode (STP, RSTP, or MSTP), use this command to specify the time in seconds it takes for an interface to time out. The range is 10 through 1000000 seconds. The default is 300 seconds. By default, the timeout feature is disabled.

To specify the time in seconds it takes for an interface to time out, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Specify the time in seconds it takes for an interface to time out.

```
switch(conf-stp)#error-disable-timeout interval 60
```

Specifying the port-channel path cost

In any mode (STP, RSTP, or MSTP), use this command to specify the port-channel path cost. The default port cost is **standard**. The path cost options are:

- **custom**—Specifies that the path cost changes according to the port-channel's bandwidth.
- **standard**—Specifies that the path cost does not change according to the port-channel's bandwidth.

NOTE

For the chosen path-cost method to be activated, it should either be configured prior to creation of all port-channels or the existing port-channels must be disabled or enabled.

To specify the port-channel path cost, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Specify the port-channel path cost.

```
switch(conf-stp)#port-channel path-cost custom
```

Specifying the bridge hello time (STP and RSTP)

In STP or RSTP mode, use this command to configure the bridge hello time. The hello time determines how often the switch interface broadcasts hello Bridge Protocol Data Units (BPDUs) to other devices. The range is 1 through 10 seconds. The default is 2 seconds.

When configuring the hello-time, the max-age setting must be greater than the hello-time setting. The following relationship should be kept:

$$2 * (\text{forward_delay} - 1) \geq \text{max_age} \geq 2 * (\text{hello_time} + 1)$$

To specify the bridge hello time, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable STP, RSTP, or MSTP.

```
switch(config)#protocol spanning-tree stp
```

3. Specify the time range in seconds for the interval between the hello BPDUs sent on an interface.

```
switch(conf-stp)#hello-time 5
```

Specifying the transmit hold count (RSTP and MSTP)

In RSTP and MSTP mode, use this command to configure the BPDU burst size by specifying the transmit hold count value. The command configures the maximum number of BPDUs transmitted per second for RSTP and MSTP before pausing for 1 second. The range is 1 through 10. The default is 6 seconds.

To specify the transmit hold count, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify the transmit hold count.

```
switch(conf-rstp)#transmit-holdcount 5
```

Enabling Cisco interoperability (MSTP)

In MSTP mode, use this command to enable or disable the ability of the Brocade FCoE hardware to interoperate with certain legacy Cisco switches. If Cisco interoperability is required on any switch in the network, then all switches in the network must be compatible, and therefore enabled using this command. The default is Cisco interoperability is disabled.

NOTE

This command is necessary because the “version 3 length” field in the MSTP BPDU on some legacy Cisco switches does not conform to current standards.

To enable Brocade FCoE hardware to interoperate with certain legacy Cisco switches, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```
3. Enable the ability of Brocade FCoE hardware to interoperate with certain legacy Cisco switches.

```
switch(conf-mstp)#cisco-interoperability enable
```

Disabling Cisco interoperability (MSTP)

To disable Cisco interoperability, perform the following steps from global configuration mode.

1. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```
2. Disable the ability of Brocade FCoE hardware to interoperate with certain legacy Cisco switches.

```
switch(conf-mstp)#cisco-interoperability disable
```

Mapping a VLAN to an MSTP instance

In MSTP mode, use this command to map a VLAN to an MSTP instance. You can group a set of VLANs to an instance. This command can be used only after the VLAN is created. VLAN instance mapping is removed from the configuration if the underlying VLANs are deleted.

To map a VLAN to an MSTP instance, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```
3. Map a VLAN to an MSTP instance. The range of valid VLAN instances is from 1 through 3583 (excluding VLAN 1002).

```
switch(conf-mstp)#instance 5 vlan 300
```

Specifying the maximum number of hops for a BPDU (MSTP)

In MSTP mode, use this command to configure the maximum number of hops for a BPDU in an MSTP region. Specifying the maximum hops for a BPDU prevents the messages from looping indefinitely on the interface. When you change the number of hops, it affects all spanning tree instances. The range is 1 through 40. The default is 20 hops.

To configure the maximum number of hops for a BPDU in an MSTP region, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```

3. Enter the **max-hops** command to configure the maximum number of hops for a BPDU in an MSTP region.

```
switch(conf-mstp)#max-hops hop_count
```

Specifying a name for an MSTP region

In MSTP mode, use this command to assign a name to an MSTP region. The region name has a maximum length of 32 characters and is case-sensitive.

To assign a name to an MSTP region, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```

3. Enter the **region** command to assign a name to an MSTP region.

```
switch(conf-mstp)#region sydney
```

Specifying a revision number for an MSTP configuration

In MSTP mode, use this command to specify a revision number for an MSTP configuration. The range is 0 through 255. The default is 0.

To specify a revision number for an MSTP configuration, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **protocol** command to enable MSTP.

```
switch(config)#protocol spanning-tree mstp
```

3. Enter the **revision** command to specify a revision number for an MSTP configuration.

```
switch(conf-mstp)#revision 17
```

Flushing MAC addresses (RSTP and MSTP)

For RSTP and MSTP, use this command to flush the MAC addresses from the VLAN filtering database (FDB). The VLAN FDB determines the forwarding of an incoming frame. The VLAN FDB contains information that helps determine the forwarding of an arriving frame based on MAC address and VLAN ID data (see “[VLAN configuration guidelines and restrictions](#)” on page 47).

There are two ways to flush the MAC addresses:

- Standard method—When one port receives a BPDU frame with a topology change flag, it flushes the FDB for the other ports in the switch. If a BPDU frame with the topology change flag is received continuously, the switch continues to flush the FDB. This behavior is the default behavior.
- Brocade method—With this method, the FDB is only flushed for the first and last BPDU with a topology change flag.

Both methods flush the FDB when the switch receives BPDUs with a topology change flag, but the Brocade method causes less flushing.

To flush the MAC addresses from the VLAN FDB, perform the following steps.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the appropriate form of the **spanning-tree** command based on how you want to flush the address:
 - To flush the MAC address using the standard method:

```
switch(config)#spanning-tree tc-flush-standard
```
 - To flush the MAC addresses from the VLAN FDB using the Brocade method:

```
switch(config)#no spanning-tree tc-flush-standard
```

Clearing spanning tree counters

In privileged EXEC mode, use this command to clear spanning tree counters on all interfaces or on the specified interface.

To clear spanning tree counters, perform the following steps from privileged EXEC mode.

1. Enter the appropriate form of the **clear** command based on what you want to clear:
 - To clear all spanning tree counters on all interfaces:

```
switch#clear spanning-tree counter
```
 - To clear the spanning tree counters associated with a specific port-channel or CEE port interface:

```
switch#clear spanning-tree counter interface tengigabitethernet 0/1
```

Clearing spanning tree-detected protocols

In privileged EXEC mode, restart the protocol migration process (force the renegotiation with neighboring switches) on all interfaces or on the specified interface.

To restart the protocol migration process, perform the following tasks from privileged EXEC mode.

1. Enter the appropriate form of the **clear** command based on what you want to clear:

- To clear all spanning tree counters on all interfaces:

```
switch#clear spanning-tree detected-protocols
```

- To clear the spanning tree counters associated with a specific port-channel or CEE port interface:

```
switch#clear spanning-tree detected-protocols interface tengigabitethernet 0/1
```

Displaying STP, RSTP, and MSTP-related information

To display STP, RSTP, and MSTP-related information, perform the following task from privileged EXEC mode.

Enter the show spanning tree command to display all STP, RSTP, and MSTP-related information.

```
switch#show spanning-tree brief
```

Configuring STP, RSTP, or MSTP on CEE interface ports

This section details the commands for enabling and configuring STP, RSTP, or MSTP on individual 10-Gigabit Ethernet CEE interface ports on Brocade FCoE hardware.

Enabling automatic edge detection

From the CEE interface, use this command to automatically identify the edge port. The port can become an edge port if no BPDU is received. By default, automatic edge detection is disabled.

To enable automatic edge detection on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.

4. Enter the **spanning-tree** command to enable automatic edge detection on the CEE interface.

```
switch(config-if-te-0/1)#spanning-tree autoedge
```

Configuring the path cost

From the CEE interface, use this command to configure the path cost for spanning tree calculations. The lower the path cost means there is a greater chance of the interface becoming the root. The range is 1 through 200000000. The default path cost is 2000.

To configure the path cost for spanning tree calculations on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to configure the path cost for spanning tree calculations on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree cost cost 2190
```

Enabling a port (interface) as an edge port

From the CEE interface, use this command to enable the port as an edge port to allow the port to quickly transition to the forwarding state. To configure a port as an edge port, follow these guidelines:

- A port can become an edge port if no BPDU is received.
- When an edge port receives a BPDU, it becomes a normal spanning tree port and is no longer an edge port.
- Because ports that are directly connected to end stations cannot create bridging loops in the network, edge ports transition directly to the forwarding state and skip the listening and learning states.
- This command is only for RSTP and MSTP. Use the **spanning-tree portfast** command for STP (see [“Enabling port fast \(STP\)”](#) on page 75).

To enable the CEE interface as an edge port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to enable the CEE interface as an edge port.

```
switch(conf-if-te-0/1)#spanning-tree edgeport
```

Enabling the guard root

From the CEE interface, use this command to enable the guard root on the switch. The guard root feature provides a way to enforce the root bridge placement in the network. With the guard root enabled on an interface, the switch is able to restrict which interface is allowed to be the spanning tree root port or the path to the root for the switch. The root port provides the best path from the switch to the root switch. By default, guard root is disabled.

Guard root protects the root bridge from malicious attacks and unintentional misconfigurations where a bridge device that is not intended to be the root bridge becomes the root bridge. This causes severe bottlenecks in the data path. Guard root ensures that the port on which it is enabled is a designated port. If the guard root-enabled port receives a superior BPDU, it goes to a discarding state.

To enable the guard root on a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to enable the guard root on a CEE interface.

```
switch(config-if-te-0/1)#spanning-tree guard root
```

Specifying the MSTP hello time

From the CEE interface, use this command to set the time interval between BPDUs sent by the root switch. Changing the **hello-time** affects all spanning tree instances.

The **max-age** setting must be greater than the **hello-time** setting (see [“Specifying the bridge maximum aging time”](#) on page 66). The range is 1 through 10 seconds. The default is 2 seconds.

To specify the MSTP hello time on a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to specify the hello time on a CEE interface.

```
switch(config-if-te-0/1)#spanning-tree hello-time 5
```

Specifying restrictions for an MSTP instance

From the CEE interface, use this command to specify restrictions on the interface for an MSTP instance.

To specify restrictions for an MSTP instance on a CEE interface, perform the following steps.

1. Enter the **configure terminal** command to access global configuration mode from privileged EXEC mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.

4. Enter the **spanning-tree** command to specify the restrictions for an MSTP instance on a CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree instance 1 restricted-role
switch(conf-if-te-0/1)#spanning-tree instance 1 restricted-tcn
```

Specifying a link type

From the CEE interface, use this command to specify a link type. Specifying the **point-to-point** keyword enables rapid spanning tree transitions to the forwarding state. Specifying the **shared** keyword disables spanning tree rapid transitions. The default setting is point-to-point.

To specify a link type on a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to specify the link type on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree link-type shared
```

Enabling port fast (STP)

From the CEE interface, use this command to enable port fast on an interface to allow the interface to quickly transition to the forwarding state. Port fast immediately puts the interface into the forwarding state without having to wait for the standard forward time.

NOTE

If you enable the **portfast bpduguard** option on an interface and the interface receives a BPDU, the software disables the interface and puts the interface in the ERR_DISABLE state.

Use the **spanning-tree edgeport** command for MSTP and RSTP (see [“Enabling a port \(interface\) as an edge port”](#) on page 73).

To enable port fast on the CEE interface for STP, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to enable port fast on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree portfast
```

Specifying the port priority

From the CEE interface, use this command to specify the port priority. The range is 0 through 240 in increments of 16. The default is 128.

To specify the port priority on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to specify the port priority on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree priority 32
```

Restricting the port from becoming a root port

From the CEE interface, use this command to restrict a port from becoming a root port. The default is to allow the CEE interface to become a root port.

To restrict the CEE interface from becoming a root port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to restrict the CEE interface from becoming a root port.

```
switch(conf-if-te-0/1)#spanning-tree restricted-role
```

Restricting the topology change notification

From the CEE interface, use this command to restrict the topology change notification BPDUs sent on the interface. By default, the restriction is disabled.

To restrict the topology change notification BPDUs sent on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to restrict the topology change notification BPDUs sent on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree restricted-tcn
```

Enabling spanning tree

By default, spanning tree is enabled on all switch ports when spanning tree is enabled globally. If it is shutdown previously on a particular interface, it is saved across all modes of spanning-tree (STP/RSTP/MSTP).

To enable spanning tree on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to enable spanning tree on the CEE interface.

```
switch(conf-if-te-0/1)#no spanning-tree shutdown
```

Disabling spanning tree

By default, spanning tree is enabled on all switch ports when spanning tree is enabled globally.

To enable spanning tree on the CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **spanning-tree** command to disable spanning tree on the CEE interface.

```
switch(conf-if-te-0/1)#spanning-tree shutdown
```

7 Configuring STP, RSTP, or MSTP on CEE interface ports

Configuring Link Aggregation

In this chapter

- [Link aggregation overview](#) 79
- [LACP configuration guidelines and restrictions](#) 83
- [Default LACP configuration](#) 83
- [LACP configuration and management](#) 83
- [LACP troubleshooting tips](#) 85

Link aggregation overview

Link aggregation allows you to bundle multiple physical Ethernet links to form a single logical trunk providing enhanced performance and redundancy. The aggregated trunk is referred to as a Link Aggregation Group (LAG). The LAG is viewed as a single link by connected devices, the spanning tree protocol, IEEE 802.1Q VLANs, and so on. When one physical link in the LAG fails, the other links stay up and there is no disruption to traffic.

To configure links to form a LAG, the physical links must be the same speed and all links must go to the same neighboring device. Link aggregation can be done by manually configuring the LAG or by dynamically configuring the LAG using the IEEE 802.3ad Link Aggregation Control Protocol (LACP).

NOTE

The LAG or LAG interface is also referred to as a port-channel.

The benefits of link aggregation are summarized as follows:

- Increased bandwidth. The logical bandwidth can be dynamically changed as the demand changes.
- Increased availability.
- Load sharing.
- Rapid configuration and reconfiguration.

The Brocade FCoE hardware supports the following trunk types:

- Static, standards-based LAG.
- Dynamic, standards-based LAG using LACP.
- Static, Brocade-proprietary LAG.
- Dynamic, Brocade-proprietary LAG using proprietary enhancements to LACP.

Link Aggregation Group configuration

You can configure a maximum of 24 Link Aggregation Groups (LAG) with up to 16 links per standard LAG and four links per Brocade-proprietary LAG. Each LAG is associated with an aggregator. The aggregator manages the Ethernet frame collection and distribution functions.

On each port, link aggregation control:

- Maintains configuration information to control port aggregation.
- Exchanges configuration information with other devices to form LAGs.
- Attaches ports to and detaches ports from the aggregator when they join or leave a LAG.
- Enables or disables an aggregator's frame collection and distribution functions.

Each link in the Brocade FCoE hardware can be associated with a LAG; a link cannot be associated with more than one LAG. The process of adding and removing links to and from a LAG is controlled either statically, dynamically, or through LACP.

Each LAG consists of the following components:

- A MAC address that is different from the MAC addresses of the LAG's individual member links.
- An interface index for each link to identify the link to neighboring devices.
- An administrative key for each link. Only links having the same administrative key value can be aggregated into a LAG. On each link configured to use LACP, LACP automatically configures an administrative key value equal to the port-channel identification number.

[Figure 5](#) and [Figure 6](#) show typical IP SAN configurations using LAGs. In a data center the Brocade 8000 switch fits into the top-of-the-rack use case where all the servers in a rack are connected to the Brocade 8000 switch through Twinax copper or optical fiber cable. The database server layer connects to the top-of-the-rack Brocade 8000 switch which is located in the network access layer.

The Brocade 8000 switch connects to Layer 2/Layer 3 aggregation routers which provide access into the existing LAN. This connectivity is formed in a standard V-design or square-design. Both designs use the LAG as the uplink to provide redundancy and improved bandwidth.

The Brocade 8000 switch interoperates with all of the major Layer 2/Layer 3 aggregation routers including Foundry Networks, Cisco Systems, and Force10 Networks.

FIGURE 5 Configuring LAGs for a top-of-the-rack CEE switch—Example 1

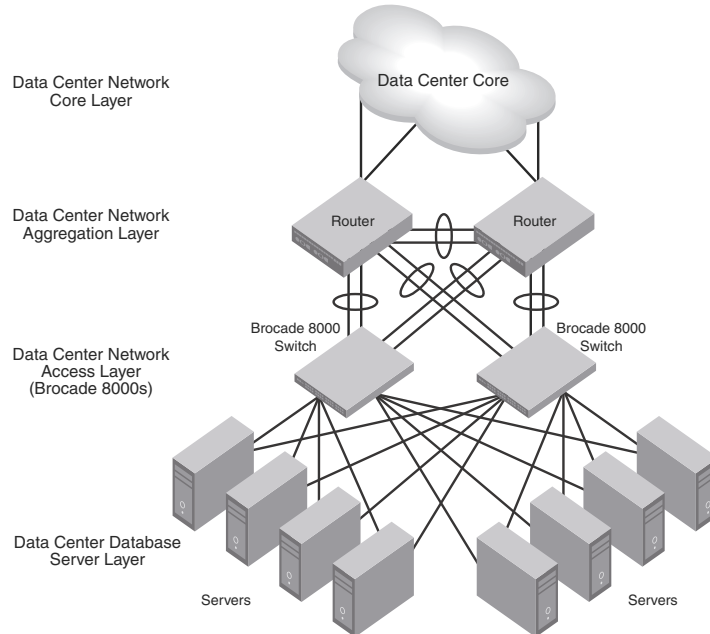
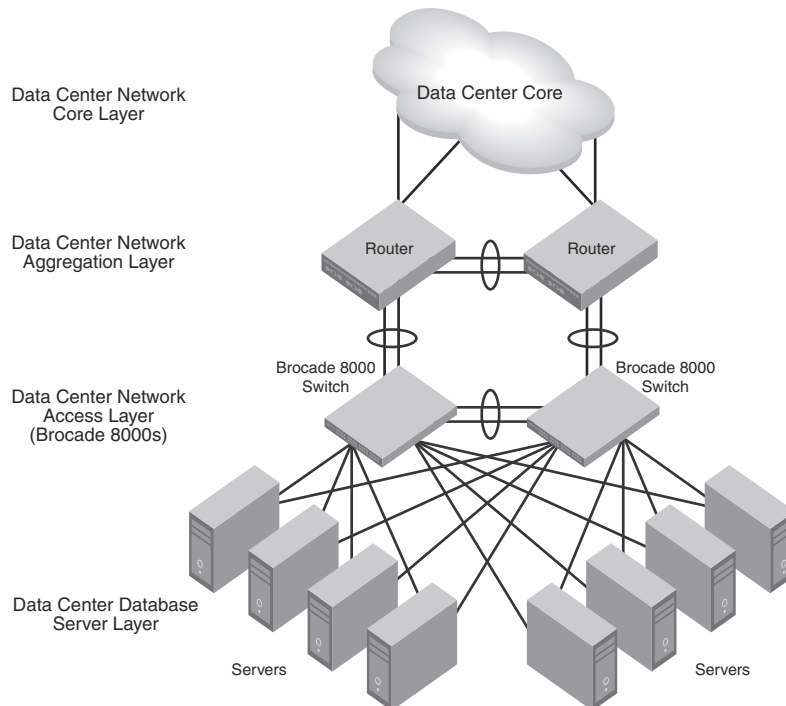


FIGURE 6 Configuring LAGs for a top-of-the-rack CEE switch—Example 2



Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standards-based protocol that allows two partner systems to dynamically negotiate attributes of physical links between them to form logical trunks. LACP determines whether a link can be aggregated into a LAG. If a link can be aggregated into a LAG, LACP puts the link into the LAG. All links in a LAG inherit the same administrative characteristics. LACP operates in two modes:

- **Passive mode**—LACP responds to Link Aggregation Control Protocol Data Units (LACPDU) initiated by its partner system but does not initiate the LACPDU exchange.
- **Active mode**—LACP initiates the LACPDU exchange regardless of whether the partner system sends LACPDU.

There are multiple ASIC limitations related to LAGs:

- Multicast traffic uses the primary link of standard LAG only.
- If more than 10Gb of traffic hashes to a single link of a standard LAG, the traffic is dropped (even though the other links are not used.)
- Brocade trunk can handle only up to 15G traffic coming from another single chip. If traffic comes from multiple chips then it can use the full bandwidth. If the traffic comes from same chip on which the Brocade trunk is present, then it can also use the full bandwidth.

Dynamic link aggregation

Dynamic link aggregation uses LACP to negotiate which links can be added and removed from a LAG. Typically, two partner systems sharing multiple physical Ethernet links can aggregate a number of those physical links using LACP. LACP creates a LAG on both partner systems and identifies the LAG by the LAG ID. All links with the same administrative key and all links that are connected to the same partner switch become members of the LAG. LACP continuously exchanges LACPDUs to monitor the health of each member link.

Static link aggregation

In static link aggregation, links are added into a LAG without exchanging LACPDUs between the partner systems. The distribution and collection of frames on static links is determined by the operational status and administrative state of the link.

Brocade-proprietary aggregation

Brocade-proprietary aggregation is similar to standards-based link aggregation but differs in how the traffic is distributed. It also has additional rules that member links must meet before they are aggregated:

- The most important rule requires that there is not a significant difference in the length of the fiber between the member links, and that all member links are part of the same port-group. The ports that belong to port-group 1, port-group 2, and port-group 3 are te0/0 to te0/7, te0/8 to te0/15, and te0/16 to te0/23, respectively.
- A maximum of four Brocade LAGs can be created per port-group.

LAG distribution process

The LAG aggregator is associated with the collection and distribution of Ethernet frames. The collection and distribution process is required to guarantee the following:

- Inserting and capturing control PDUs.
- Restricting the traffic of a given conversation to a specific link.
- Load balancing between individual links.
- Handling dynamic changes in LAG membership.

LACP configuration guidelines and restrictions

This section applies to standards-based and Brocade-proprietary LAG configurations except where specifically noted otherwise.

Follow these LACP configuration guidelines and restrictions when configuring LACP:

- All ports on the Brocade FCoE hardware can operate only in full-duplex mode.
- QoS—In the Fabric OS version 7.0.1 release, QoS commands for a LAG need to be specified on each LAG member link, instead of on the logical LAG interface (port-group). Additionally, the QoS commands specified on each LAG member link need to be the same on each link.
- Brocade-proprietary LAGs only—All LAG member links need to be part of the same port-group.
- Switchport interfaces—Interfaces configured as “switchport” interfaces cannot be aggregated into a LAG. However, a LAG can be configured as a switchport.

Default LACP configuration

Table 19 lists the default LACP configuration.

TABLE 19 Default LACP configuration

Parameter	Default setting
System priority	32768
Port priority	32768
Timeout	Long (standard LAG) and short (Brocade LAG)

LACP configuration and management

You need to enter either the **copy running-config startup-config** command or the **write memory** command to save your configuration changes to Flash memory so that they are not lost if there is a system reload or power outage.

Enabling LACP on an CEE interface

To add additional interfaces to an existing LAG, repeat this procedure using the same LAG group number for the new interfaces.

To enable LACP on a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.

```
switch(config-if)#no shutdown
```

4. Enter the **channel-group** command to configure the LACP for the CEE interface.

```
switch(conf-if)#channel-group 4 mode active type brocade
```

Enabling LACP on an interface in Layer 2 mode

To enable LACP on a Layer 2 interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the interface port-channel.

```
switch(config)#interface port-channel 4
```

3. Enter the **switchport** command to set the interface into Layer 2 mode.

```
switch(conf-if-po-4)#switchport
```

4. Enter the **switchport mode** command to set the mode.

```
switch(conf-if-po-4)#switchport mode access
```

5. Enter the **no shutdown** command to enable the interface.

```
switch(config-if)#no shutdown
```

Configuring the LACP system priority

You configure an LACP system priority on each switch running LACP. LACP uses the system priority with the switch MAC address to form the system ID and also during negotiation with other switches.

The system priority value must be a number in the range of 1 through 65535. The higher the number, the lower the priority. The default priority is 32768.

To configure the global LACP system priority, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Specify the LACP system priority.

```
switch(config)#lACP system-priority 25000
```

Configuring the LACP timeout period on a CEE interface

The LACP timeout period indicates how long LACP waits before timing out the neighboring device. The **short** timeout period is 3 seconds and the **long** timeout period is 90 seconds. The default is **long**.

To configure the LACP timeout period on a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Specify the LACP timeout period for the CEE interface.

```
switch(conf-if-te-0/1)#lACP timeout short
```

Clearing LACP counter statistics on a LAG

To clear LACP counter statistics, perform the following task from privileged EXEC mode.

Enter the **clear** command to clear the LACP counter statistics for the specified LAG group number.

```
switch#clear lACP 42 counters
```

Clearing LACP counter statistics on all LAG groups

To clear LACP counter statistics, perform the following task from privileged EXEC mode.

Enter the **clear** command to clear the LACP counter statistics for all LAG groups.

```
switch#clear lACP counters
```

Displaying LACP information

Use the **show** command to display LACP statistics and configuration information. See the *Converged Enhanced Ethernet Command Reference* for information.

LACP troubleshooting tips

To troubleshoot problems with your LACP configuration, use the following troubleshooting tips.

If a standard IEEE 802.3ad-based dynamic trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **standard** for the trunk type.
- Make sure that both ends of the link *are not* configured for **passive** mode. They must be configured as either **active/active**, **active/passive**, or **passive/active**.
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.
- Make sure that the links that are part of the LAG are connected to the same neighboring switch.
- Make sure that the system ID of the switches connected by the link is unique. This can be verified by entering the **show lACP sys-id** command on both switches.

- Make sure that LACPDUs are being received and transmitted on both ends of the link and that there are no error PDUs. This can be verified by entering the **show lacp counters** *port-channel-num* command and looking at the receive mode (rx) and transmit mode (tx) statistics. The statistics should be incrementing and should not be at zero or a fixed value. If the PDU rx count is not incrementing, check the interface for possible CRC errors by entering the **show interface** *link-name* command on the neighboring switch. If the PDU tx count is not incrementing, check the operational status of the link by entering the **show interface** *link-name* command and verifying that the interface status is “up.”

If a Brocade-based dynamic trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **Brocade** for trunk type.
- Make sure that both ends of the link are *not* configured for **passive** mode. They must be configured as either **active/active**, **active/passive**, or **passive/active**.
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.
- Make sure that the links that are part of the LAG are connected to the same neighboring switch.
- Make sure that the system ID of the switches connected by the link is unique. This can be verified by entering the **show lacp sys-id** command on both switches.
- Make sure that LACPDUs are being received and transmitted on both ends of the link and there are no error PDUs. This can be verified by entering the **show lacp** *port-channel-num* **counters** command and looking at the rx and tx statistics. The statistics should be incrementing and should not be at zero or a fixed value. If the PDU rx count is not incrementing, check the interface for possible CRC errors by entering the **show interface** *link-name* command on the neighboring switch.
- Make sure that the fiber length of the link has a deskew value of 7 microseconds. If it does not, the link will not be able to join the LAG and the following RASLOG message is generated:

```
Deskew calculation failed for link <link-name>.
```

When a link has this problem, the **show port-channel** command displays the following:

```
Mux machine state : Deskew not OK.
```

If a Brocade-based static trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **Brocade** for trunk type and verify that the mode is “on.”
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.

If a standards-based static trunk is configured on a link and the link is not able to join the LAG:

- Make sure that both ends of the link are configured as **standard** for trunk type and verify that the mode is “on.”
- Make sure that the port-channel interface is in the administrative “up” state by ensuring that the **no shutdown** command was entered on the interface on both ends of the link.

Configuring LLDP

In this chapter

- [LLDP overview](#) 87
- [Layer 2 topology mapping](#)..... 88
- [DCBX overview](#)..... 89
- [DCBX interaction with other vendor devices](#) 91
- [LLDP configuration guidelines and restrictions](#)..... 91
- [Default LLDP configuration](#)..... 91
- [LLDP configuration and management](#)..... 92

LLDP overview

The IEEE 802.1AB Link Layer Discovery Protocol (LLDP) enhances the ability of network management tools to discover and maintain accurate network topologies and simplify LAN troubleshooting in multi-vendor environments. To efficiently and effectively operate the various devices in a LAN you must ensure the correct and valid configuration of the protocols and applications that are enabled on these devices. With Layer 2 networks expanding dramatically, it is difficult for a network administrator to statically monitor and configure each device in the network.

Using LLDP, network devices such as routers and switches advertise information about themselves to other network devices and store the information they discover. Details such as device configuration, device capabilities, and device identification are advertised. LLDP defines the following:

- A common set of advertisement messages.
- A protocol for transmitting the advertisements.
- A method for storing the information contained in received advertisements.

NOTE

LLDP runs over the data-link layer which allows two devices running different network layer protocols to learn about each other.

LLDP information is transmitted periodically and stored for a finite period. Every time a device receives an LLDP advertisement frame, it stores the information and initializes a timer. If the timer reaches the time to live (TTL) value, the LLDP device deletes the stored information ensuring that only valid and current LLDP information is stored in network devices and is available to network management systems.

Layer 2 topology mapping

The LLDP protocol lets network management systems accurately discover and model Layer 2 network topologies. As LLDP devices transmit and receive advertisements, the devices store information they discover about their neighbors. Advertisement data such as a neighbor's management address, device type, and port identification is useful in determining what neighboring devices are in the network.

NOTE

Brocade's LLDP implementation supports a one-to-one connection. Each interface has one and only one neighbor.

The higher level management tools, such as Brocade's DCFM, can query the LLDP information to draw Layer 2 physical topologies. The management tools can continue to query a neighboring device through the device's management address provided in the LLDP information exchange. As this process is repeated, the complete Layer 2 topology is mapped.

In LLDP the link discovery is achieved through the exchange of link-level information between two link partners. The link-level information is refreshed periodically to reflect any dynamic changes in link-level parameters. The basic format for exchanging information in LLDP is in the form of a type, length, value (TLV) field.

LLDP keeps a database for both local and remote configurations. The LLDP standard currently supports three categories of TLVs. Brocade's LLDP implementation adds a proprietary Brocade extension TLV set. The four TLV sets are described as follows:

- Basic management TLV set. This set provides information to map the Layer 2 topology and includes the following TLVs:
 - Chassis ID TLV—Provides the ID for the switch or router where the port resides. This is a mandatory TLV.
 - Optional TLV—Provides a description of the port in an alphanumeric format. If the LAN device supports RFC-2863, the port description TLV value equals the "ifDescr" object. This is a mandatory TLV.
 - System name TLV—Provides the system-assigned name in an alphanumeric format. If the LAN device supports RFC-3418, the system name TLV value equals the "sysName" object. This is an optional TLV.
 - System description TLV—Provides a description of the network entity in an alphanumeric format. This includes system name, hardware version, operating system, and supported networking software. If the LAN device supports RFC-3418, the value equals the "sysDescr" object. This is an optional TLV.
 - System capabilities TLV—Indicates the primary functions of the device and whether these functions are enabled in the device. The capabilities are indicated by two octets. The first octet indicates Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station, respectively. The second octet is reserved. This is an optional TLV.
 - Management address TLV—Indicates the addresses of the local switch. Remote switches can use this address to obtain information related to the local switch. This is an optional TLV.

- IEEE 802.1 organizational TLV set. This set provides information to detect mismatched settings between local and remote devices. A trap or event can be reported once a mismatch is detected. This is an optional TLV. This set includes the following TLVs:
 - Port VLANID TLV—Indicates the port VLAN ID (PVID) that is associated with an untagged or priority tagged data frame received on the VLAN port.
 - PPVLAN ID TLV—Indicates the port- and protocol-based VLAN ID (PPVID) that is associated with an untagged or priority tagged data frame received on the VLAN port. The TLV supports a “flags” field that indicates whether the port is capable of supporting port- and protocol-based VLANs (PPVLANs) and whether one or more PPVLANs are enabled. The number of PPVLAN ID TLVs in a Link Layer Discovery Protocol Data Unit (LLDPDU) corresponds to the number of the PPVLANs enabled on the port.
 - VLAN name TLV—Indicates the assigned name of any VLAN on the device. If the LAN device supports RFC-2674, the value equals the “dot1QVLANStaticName” object. The number of VLAN name TLVs in an LLDPDU corresponds to the number of VLANs enabled on the port.
 - Protocol identity TLV—Indicates the set of protocols that are accessible at the device's port. The protocol identity field in the TLV contains a number of octets after the Layer 2 address that can enable the receiving device to recognize the protocol. For example, a device that wishes to advertise the spanning tree protocol includes at least eight octets: 802.3 length (two octets), LLC addresses (two octets), 802.3 control (one octet), protocol ID (two octets), and the protocol version (one octet).
- IEEE 802.3 organizational TLV set. This is an optional TLV set. This set includes the following TLVs:
 - MAC/PHY configuration/status TLV—Indicates duplex and bit rate capabilities and the current duplex and bit rate settings of the local interface. It also indicates whether the current settings were configured through auto-negotiation or through manual configuration.
 - Power through media dependent interface (MDI) TLV—Indicates the power capabilities of the LAN device.
 - Link aggregation TLV—Indicates whether the link (associated with the port on which the LLDPDU is transmitted) can be aggregated. It also indicates whether the link is currently aggregated and provides the aggregated port identifier if the link is aggregated.
 - Maximum Ethernet frame size TLV—Indicates the maximum frame size capability of the device's MAC and PHY implementation.
- Brocade extension TLV set. This set is used to identify vendor-specific information. This set includes the following TLVs:
 - Link Vendor/Version TLV—Indicates the vendor for the switch, host, or router where the port resides.
 - Primitive supported/version TLV—Indicates where the link-level primitives are supported (if supported) and the link-level primitive version.

DCBX overview

Storage traffic requires a lossless communication which is provided by CEE. The Data Center Bridging (DCB) Capability Exchange Protocol (DCBX) is used to exchange CEE-related parameters with neighbors to achieve more efficient scheduling and a priority-based flow control for link traffic.

DCBX uses LLDP to exchange parameters between two link peers; DCBX is built on the LLDP infrastructure for the exchange of information. DCBX-exchanged parameters are packaged into organizationally specific TLVs. The DCBX protocol requires an acknowledgement from the other side of the link, therefore LLDP is turned on in both transmit and receive directions. DCBX requires version number checking for both control TLVs and feature TLVs.

DCBX interacts with other protocols and features as follows:

- LLDP—LLDP is run in parallel with other Layer 2 protocols such as RSTP and LACP. DCBX is built on the LLDP infrastructure to communicate capabilities supported between link partners. The DCBX protocol and feature TLVs are treated as a superset of the LLDP standard.
- QoS management—DCBX capabilities exchanged with a link partner are passed down to the QoS management entity to set up the Brocade FCoE hardware to control the scheduling and priority-based flow control in the hardware.

The DCBX standard is subdivided into two features sets:

- “Enhanced Transmission Selection (ETS)”
- “Priority Flow Control (PFC)”

Enhanced Transmission Selection (ETS)

In a converged network, different traffic types affect the network bandwidth differently. The purpose of ETS is to allocate bandwidth based on the different priority settings of the converged traffic. For example, Inter-process communications (IPC) traffic can use as much bandwidth as needed and there is no bandwidth check; LAN and SAN traffic share the remaining bandwidth. [Table 20](#) displays three traffic groups: IPC, LAN, and SAN. ETS allocates the bandwidth based on traffic type and also assigns a priority to the three traffic types as follows: Priority 7 traffic is mapped to priority group 0 which does not get a bandwidth check, priority 2 and priority 3 are mapped to priority group 1, priorities 6, 5, 4, 1 and 0 are mapped to priority group 2.

The priority settings shown in [Table 20](#) are translated to priority groups in the Brocade FCoE hardware.

TABLE 20 ETS priority grouping of IPC, LAN, and SAN traffic

Priority	Priority group	Bandwidth check
7	0	No
6	2	Yes
5	2	Yes
4	2	Yes
3	1	Yes
2	1	Yes
1	2	Yes
0	2	Yes

Priority Flow Control (PFC)

With PFC, it is important to provide lossless frame delivery for certain traffic classes while maintaining existing LAN behavior for other traffic classes on the converged link. This differs from the traditional 802.3 PAUSE type of flow control where the pause affects all traffic on an interface.

PFC is defined by a one-byte bitmap. Each bit position stands for a user priority. If a bit is set, the flow control is enabled in both directions (Rx and Tx).

DCBX interaction with other vendor devices

When the Brocade FCoE hardware interacts with other vendor devices, the other vendor devices might not have support for the same DCBX version as the Brocade FCoE hardware.

The Brocade FCoE hardware supports two DCBX versions:

- CEE version (1.0.1)—Based on the CEE standard.
- Pre-CEE version.

To accommodate the different DCBX versions, the Brocade FCoE hardware provides the following options.

- Auto-sense (plug and play)
This is the default. The Brocade FCoE hardware detects the version used by the link neighbor and automatically switches between the CEE version and the pre-CEE version.
- CEE version
Forces the use of the CEE version for the link (auto-sense is off).
- Pre-CEE version
Forces the use of the pre-CEE version for the link (auto-sense is off).

LLDP configuration guidelines and restrictions

Follow these LLDP configuration guidelines and restrictions when configuring LLDP:

- Brocade's implementation of LLDP supports Brocade-specific TLV exchange in addition to the standard LLDP information.
- Mandatory TLVs are always advertised.
- The exchange of LLDP link-level parameters is transparent to the other Layer 2 protocols. The LLDP link-level parameters are reported by LLDP to other interested protocols.

NOTE

DCBX configuration simply involves configuring DCBX-related TLVs to be advertised. Detailed information is provided in the [“LLDP configuration and management”](#) on page 92.

Default LLDP configuration

[Table 21](#) lists the default LLDP configuration.

TABLE 21 Default LLDP configuration

Parameter	Default setting
LLDP global state	Enabled
LLDP receive	Enabled

TABLE 21 Default LLDP configuration (Continued)

Parameter	Default setting
LLDP transmit	Enabled
Transmission frequency of LLDP updates	30 seconds
Hold time for receiving devices before discarding	120 seconds
DCBX-related TLVs to be advertised	dcbx-tlv

LLDP configuration and management

NOTE

You need to enter either the **copy running-config startup-config** command or the **write memory** command to save your configuration changes to Flash so that they are not lost if there is a system reload or power outage.

Enabling LLDP globally

The **protocol lldp** command enables LLDP globally on all interfaces unless it has been specifically disabled on an interface. LLDP is globally enabled by default.

To enable LLDP globally, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Use the **no disable** command to enable LLDP.

```
switch(conf-lldp)#no disable
```

Disabling and resetting LLDP globally

The **protocol lldp** command returns all configuration settings made using the **protocol lldp** commands to their default settings. LLDP is globally enabled by default.

To disable and reset LLDP globally, perform the following steps from global configuration mode.

1. Enter the **protocol lldp** command to enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

2. Disable the LLDP protocol.

```
switch(conf-lldp)#no protocol lldp
```

Configuring LLDP global command options

After entering the **protocol lldp** command from global configuration mode, you are in LLDP configuration mode which is designated with the `switch(conf-lldp)#` prompt. Using the keywords in this mode, you can set non-default parameter values that apply globally to all interfaces.

Specifying a system name for the Brocade FCoE hardware

The global system name for LLDP is useful for differentiating between switches. By default, the “host-name” from the chassis/entity MIB is used. By specifying a descriptive system name, you will find it easier to configure the switch for LLDP.

The name used in the **system-name** command must begin with a letter, and can consist of letters, digits, hyphens, and underscore characters. Spaces are prohibited. Special characters are not supported, and cause the name to truncate.

To specify a global system name for the Brocade FCoE hardware, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Specify an LLDP system name for the CEE switch.

```
switch(conf-lldp)#system-name Brocade_Alpha
Brocade_Alpha(conf-lldp)#
```

Specifying an LLDP system description for the Brocade FCoE hardware

NOTE

Brocade recommends you use the operating system version for the description or use the description from the chassis/entity MIB.

To specify an LLDP system description for the Brocade FCoE hardware, perform the following steps from privileged EXEC mode. The system description is seen by neighboring switches.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Specify a system description for the Brocade FCoE hardware.

```
switch(conf-lldp)#system-description IT_1.6.2_LLDP_01
```

Specifying a user description for LLDP

To specify a user description for LLDP, perform the following steps from privileged EXEC mode. This description is for network administrative purposes and is not seen by neighboring switches.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Specify a user description for LLDP.

```
switch(conf-lldp)#description Brocade-LLDP-installed-july-25
```

Enabling and disabling the receiving and transmitting of LLDP frames

By default both transmit and receive for LLDP frames is enabled. To enable or disable the receiving (rx) and transmitting (tx) of LLDP frames, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mode** command to:
 - Enable only receiving of LLDP frames:
`switch(conf-lldp)#mode rx`
 - Enable only transmitting of LLDP frames:
`switch(conf-lldp)#mode tx`
 - Disable all LLDP frame transmissions
`switch(conf-lldp)#no mode`

Configuring the transmit frequency of LLDP frames

To configure the transmit frequency of LLDP frames, perform the following steps from privileged EXEC mode. The default is 30 seconds.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.
`switch(config)#protocol lldp`
3. Configure the transmit frequency of LLDP frames.
`switch(conf-lldp)#hello 45`

Configuring the hold time for receiving devices

To configure the hold time for receiving devices, perform the following steps from privileged EXEC mode. This configures the number of consecutive LLDP hello packets that can be missed before declaring the neighbor information as invalid. The default is 4.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.
`switch(config)#protocol lldp`
3. Configure the hold time for receiving devices.
`switch(conf-lldp)#multiplier 6`

Advertising the optional LLDP TLVs

NOTE

If the **advertise optional-tlv** command is entered without keywords, all optional LLDP TLVs are advertised.

To advertise the optional LLDP TLVs, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Advertise the optional LLDP TLVs.

```
switch(conf-lldp)#advertise optional-tlv management-address port-description
system-capabilities system-name system-description
```

Configuring the advertisement of LLDP DCBX-related TLVs

NOTE

By default, the dcbx-tlv is advertised; the dot1-tlv, dot3-tlv, dcbx-fcoe-app-tlv, and dcbx-fcoe-logical-link-tlv are not advertised.

To configure the LLDP DCBX-related TLVs to be advertised, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Advertise the LLDP DCBX-related TLVs using these commands:

- `switch(conf-lldp)#advertise dcbx-fcoe-app-tlv`
- `switch(conf-lldp)#advertise dcbx-fcoe-logical-link-tlv`
- `switch(conf-lldp)#advertise dcbx-tlv`
- `switch(conf-lldp)#advertise dot1-tlv`
- `switch(conf-lldp)#advertise dot3-tlv`

Configuring FCoE priority bits

The FCoE priority bit setting is a bitmap setting where each bit position stands for a priority. When you set a bit for a particular priority, that priority setting is applied to the FCoE traffic (that is, the incoming FCoE traffic will have that priority). The default value is 0.

NOTE

FCoE traffic is only supported on the priority level that also has flow control enabled. This means that the final advertised FCoE priority consists of the configured FCoE priority setting and the per-priority flow control setting.

Even though setting multiple bits is allowed (exercising the full range of values), it doesn't make sense to set more than one bit, because adapters don't support multiple priorities for FCoE.

To configure the FCoE priority bits, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Configure the FCoE priority bits.

```
switch(conf-lldp)#lldp fcoe-priority-bits 3
```

4. Advertise the TLV.

```
switch(conf-lldp)#advertise dcbx-fcoe-app-tlv
```

Configuring iSCSI priority bits

The iSCSI priority bit setting is a bitmap setting where each bit position stands for a priority. When you set a bit for a particular priority, that priority setting is applied to the iSCSI traffic (that is, the incoming iSCSI traffic will have that priority). The default value is 4.

The iSCSI TLV is used only to advertise the iSCSI traffic configuration parameters to the attached CEE enabled servers and targets. No verification or enforcement of the usage of the advertised parameters by the iSCSI server or target is done by the switch.

To configure the iSCSI priority bits, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Configure the iSCSI priority bits.

```
switch(conf-lldp)#iscsi-priority-bits list 2 3
```

4. Advertise the TLV.

```
switch (conf-lldp)#advertise dcbx-iscsi-app-tlv
```

Configuring LLDP profiles

You can configure up to 64 profiles on a switch, but the best practice is to limit this to one profile for each port. Using the **no profile name** command deletes the entire profile.

To configure LLDP profiles, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter LLDP configuration mode.

```
switch(config)#protocol lldp
```

3. Configure the profile name.

```
switch(conf-lldp)#profile UK_LLDP_IT
```

4. Specify a description for the profile.

```
switch(conf-lldp-profile-UK_LLDP_IT)#description standard_profile_by_Jane
```

5. Enable the transmitting and receiving of LLDP frames.

```
switch(conf-lldp-profile-UK_LLDP_IT)#no mode
```

6. Configure the transmission frequency of LLDP updates.

```
switch(conf-lldp-profile-UK_LLDP_IT)#hello 10
```

7. Configure the hold time for receiving devices.

```
switch(conf-lldp-profile-UK_LLDP_IT)#multiplier 2
```

8. Advertise the optional LLDP TLVs.

9. Advertise the LLDP DCBX-related TLVs.

```
switch(conf-lldp-profile-UK_LLDP_IT)#advertise dot1-tlv
switch(conf-lldp-profile-UK_LLDP_IT)#advertise dot3-tlv
switch(conf-lldp-profile-UK_LLDP_IT)#advertise advertise dcbx-tlv
switch(conf-lldp-profile-UK_LLDP_IT)#advertise dcbx-fcoe-logical-link-tlv
switch(conf-lldp-profile-UK_LLDP_IT)#advertise dcbx-fcoe-app-tlv
switch(conf-lldp-profile-UK_LLDP_IT)#advertise dcbx-iscsi-app-tlv
```

NOTE

Brocade recommends against advertising dot1.tlv and dot3.tlv LLDPs if your network contains CNAs from non-Brocade vendors. This configuration may cause functionality problems.

10. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(conf-lldp-profile-UK_LLDP_IT)#do copy running-config startup-config
```

Configuring the iSCSI profile

You can configure an iSCSI profile to be applied to individual interfaces. However, the priority bit must be set manually for each interface. Using the **no profile name** command deletes the entire profile.

To configure iSCSI profiles, perform the following steps from privileged EXEC mode.

1. Use the **cee-map** command to create the cee profile, if it has not already been created.

```
switch(config)#cee-map
switch(conf-ceemap)#priority-group-table 1 weight 50
switch(conf-ceemap)#priority-group-table 2 weight 30 pfc
switch(conf-ceemap)#priority-group-table 3 weight 20 pfc
switch(conf-ceemap)#priority-table 1 1 1 1 2 3 1 1
```

2. Enter LLDP configuration mode.

```
switch(conf-ceemap)#protocol lldp
```

3. Create an LLDP profile for iSCSI.

```
switch(conf-lldp)#profile iscsi_config
```

4. Advertise the iSCSI TLV.

```
switch(conf-lldp-profile-iscsi_config)#advertise dcbx-iscsi-app-tlv
```

5. Enter configuration mode for the specific interface.

```
switch (conf-lldp-profile-iscsi_config)#interface te 0/1
```

6. Apply the CEE Provisioning map to the interface.

```
switch(conf-if-te-0/1)#cee default
```

7. Apply the LLDP profile you created for iSCSI.

```
switch(conf-if-te-0/1)#lldp profile iscsi_config
```

8. Set the iSCSI priority bits for the interface.

```
switch(conf-if-te-0/1)#iscsi-priority-bits list 6
```

9. Repeat steps 5 through 7 for additional interfaces.

```
switch(conf-if-te-0/1)#interface te 0/7
switch(conf-if-te-0/7)#cee default
switch(conf-if-te-0/7)#lldp profile iscsi_config
switch(conf-if-te-0/7)#iscsi-priority-bits list 6
```

Configuring LLDP interface-level command options

Only one LLDP profile can be assigned to an interface. If you do not use the **lldp profile** option at the interface level, the global configuration is used on the interface. If there are no global configuration values defined, the global default values are used.

To configure LLDP interface-level command options, perform the following steps from privileged EXEC mode.

1. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/10
```

2. Apply an LLDP profile to the interface.

```
switch(conf-if-te-0/10)#lldp profile network_standard
```

3. Configure the FCoE priority bits for an interface. The value is specified as 0x0-0xff.

```
switch(conf-if-te-0/10)#fcoe-priority-bits 6
```

4. Configure the DCBX version for an interface for CEE. For detailed information on these version command keywords, see [“DCBX interaction with other vendor devices”](#) on page 91. The default option is to automatically detect the DCBX version.

```
switch(conf-if-te-0/10)#lldp dcbx-version cee
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(conf-if-te-0/10)#do copy running-config startup-config
```

Clearing LLDP-related information

To clear LLDP-related information, perform the following steps from privileged EXEC mode.

1. Use the **clear** command to:

- Clear LLDP neighbor information.

```
switch#clear lldp neighbors tengigabitethernet 0/1
```

- Clear LLDP statistics.

```
switch#clear lldp statistics tengigabitethernet 0/1
```

- Clear all LLDP statistics.

```
switch#clear lldp statistics
```

- Clear all LLDP neighbors.

```
switch#clear lldp neighbors
```

Displaying LLDP-related information

To display LLDP-related information, perform the following steps from EXEC or privileged EXEC mode.

1. Use the **show lldp neighbors** command to:

- Display LLDP general information.

```
switch#show lldp
```

- Display LLDP neighbor-related information.


```
switch#show lldp neighbors
```

- Display LLDP neighbor-related detail information.

```
switch#show lldp neighbors detail
```

- Display LLDP interface-related information.

```
switch#show lldp interface tengigabitethernet 0/1
```

- Display LLDP neighbor-related information for a particular interface.

```
switch#show lldp neighbors interface tengigabitethernet 0/1 detail
```

- Display all LLDP statistics.

```
switch#show lldp statistics and show lldp statistics interface  
tengigabitethernet 0/1
```

9 LLDP configuration and management

Configuring ACLs

In this chapter

- [ACL overview](#) 101
- [Default ACL configuration](#) 102
- [ACL configuration guidelines and restrictions](#) 102
- [ACL configuration and management](#) 102

ACL overview

ACLs filter traffic for the Brocade FCoE hardware and permit or deny *incoming* frames from passing through interfaces that have the ACLs applied to them. You can apply ACLs on VLANs and on Layer 2 interfaces. Each ACL is a unique collection of permit and deny statements (rules) that apply to frames. When a frame is received on an interface, the switch compares the fields in the frame against any ACLs applied to the interface to verify that the frame has the required permissions to be forwarded. The switch compares the frame, sequentially, against each rule in the ACL and either forwards the frame or drops the frame.

NOTE

In the Brocade Fabric OS v7.0.1 release, only Layer 2 MAC access control lists (ACLs) are supported.

The switch examines ACLs associated with options configured on a given interface. As frames enter the switch on an interface, ACLs associated with all inbound options configured on that interface are examined. With MAC ACLs you can identify and filter traffic based on the MAC address, and EtherType.

The primary benefits of ACLs are as follows:

- Provide a measure of security.
- Save network resources by reducing traffic.
- Block unwanted traffic or users.
- Reduce the chance of denial of service (DOS) attacks.

There are two types of MAC ACLs:

- Standard ACLs—Permit and deny traffic according to the source MAC address in the incoming frame. Use standard MAC ACLs if you only need to filter traffic based on source addresses.
- Extended ACLs—Permit and deny traffic according to the source and destination MAC addresses in the incoming frame, as well as EtherType.

MAC ACLs are supported on the following interface types:

- Physical interfaces
- Logical interfaces (LAGs)
- VLANs

Default ACL configuration

Table 22 lists the default ACL configuration.

TABLE 22 Default MAC ACL configuration

Parameter	Default setting
MAC ACLs	By default, no MAC ACLs are configured.

ACL configuration guidelines and restrictions

Follow these ACL configuration guidelines and restrictions when configuring ACLs:

- The order of the rules in an ACL is critical. The first rule that matches the traffic stops further processing of the frames.
- Standard ACLs and extended ACLs cannot have the same name.

ACL configuration and management

You need to enter either the **copy running-config startup-config** command or the **write memory** command to save your configuration changes to Flash so that they are not lost if there is a system reload or power outage.

Creating a standard MAC ACL and adding rules

NOTE

You can use the **resequence** command to change all the sequence numbers assigned to the rules in a MAC ACL. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 104.

To create a standard MAC ACL and add rules, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Create a standard MAC ACL and enter ACL configuration mode.

In this example, the name of the standard MAC ACL is “test_01.”

```
switch(config)#mac access-list standard test_01
```

3. Enter the **deny** command to create a rule in the MAC ACL to drop traffic with the source MAC address.

```
switch(conf-macl-std)#deny 0022.3333.4444 count
```

4. Enter the **permit** command to create a rule in the MAC ACL to permit traffic with the source MAC address.

```
switch(conf-macl-std)#permit 0022.5555.3333 count
```

5. Use the **seq** command to create MAC ACL rules in a specific sequence.

```
switch(conf-macl-std)#seq 100 deny 0011.2222.3333 count
switch(conf-macl-std)#seq 1000 permit 0022.1111.2222 count
```

Creating an extended MAC ACL and adding rules

NOTE

You can use the **resequence** command to change all the sequence numbers assigned to the rules in a MAC ACL. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 104.

The MAC ACL name length is limited to 64 characters.

To create an extended MAC ACL and add rules, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

2. Create an extended MAC ACL and enter ACL configuration mode.

```
switch(config)#mac access-list extended test_02
```

3. Create a rule in the MAC ACL to **permit** traffic with the source MAC address and the destination MAC address.

```
switch(conf-macl-ext)#permit 0022.3333.4444 0022.3333.5555
```

4. Use the **seq** command to insert the rule anywhere in the MAC ACL.

```
switch(conf-macl-std)#seq 5 permit 0022.3333.4444 0022.3333.5555
```

5. Enter the copy command to save the running-config file to the startup-config file.

```
switch(conf-macl-std)#do copy running-config startup-config
```

Modifying MAC ACL rules

You cannot modify the existing rules of a MAC ACL. However, you can remove the rule and then recreate it with the desired changes.

If you need to add more rules between existing rules than the current sequence numbering allows, you can use the **resequence** command to reassign sequence numbers. For detailed information, see [“Reordering the sequence numbers in a MAC ACL”](#) on page 104.

Use a sequence number to specify the rule you wish to modify. Without a sequence number, a new rule is added to the end of the list, and the existing rule is unchanged.

NOTE

Using the **permit** and **deny** keywords, you can create many different rules. The examples in this section provide the basic knowledge needed to modify MAC ACLs.

NOTE

This example assumes that test_02 contains an existing rule number 100 with the “deny any any” options.

To modify a MAC ACL, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mac** command to specify the ACL called test_02 for modification.

```
switch(config)#mac access-list extended test_02
```
3. Enter the **no seq** command to delete the existing rule 100.

```
switch (config)#no seq 100
```
4. Enter the **seq** command to re create rule number 100 by recreating it with new parameters.

```
switch(conf-macl-ext)#seq 100 permit any any
```

Removing a MAC ACL

To remove a MAC ACL, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **mac** command to specify and delete the ACL that you want to remove. In this example, the extended MAC ACL name is “test_02.”

```
switch(config)#no mac access-list extended test_02
```

Reordering the sequence numbers in a MAC ACL

You can reorder the sequence numbers assigned to rules in a MAC ACL. Reordering the sequence numbers is useful when you need to insert rules into an ACL and there are not enough available sequence numbers.

The first rule receives the number specified by the starting-sequence number that you specify. Each subsequent rule receives a number larger than the preceding rule. The difference in numbers is determined by the increment number that you specify. The starting-sequence number and the increment number must be in the range of 1 through 65535.

For example, in the task listed below the **resequence** command assigns a sequence number of 50 to the rule named test_02, then the second rule has a sequence number of 55 and the third rule a has a sequence number of 60.

To reorder the rules in a MAC ACL, perform the following task from privileged EXEC mode.

Enter the **resequence** command to assign sequence numbers to the rules contained in the MAC ACL.

```
switch#resequence access-list mac test_02 50 5
```

Applying a MAC ACL to a CEE interface

Ensure that the ACL that you want to apply exists and is configured to filter traffic in the manner that you need for this CEE interface. An ACL does not take effect until it is expressly applied to an interface using the **access-group** command. Frames can be filtered as they enter an interface (ingress direction).

To apply a MAC ACL to a CEE interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enter the **switchport** command to configure the interface as a Layer 2 switch port.
4. Enter the **mac-access-group** command to specify the MAC ACL that is to be applied to the Layer 2 CEE interface in the ingress direction.

```
switch(conf-if-te-0/1)#mac access-group test_02 in
```

Applying a MAC ACL to a VLAN interface

Ensure that the ACL that you want to apply exists and is configured to filter traffic in the manner that you need for this VLAN interface. An ACL does not take effect until it is expressly applied to an interface using the **access-group** command. Frames can be filtered as they enter an interface (ingress direction).

To apply a MAC ACL to a VLAN interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **interface** command to apply the VLAN interface to the MAC ACL.

```
switch(config)#interface vlan 50
```

3. Enter the **mac-access-group** command to specify the MAC ACL that is to be applied to the VLAN interface in the ingress direction.

```
switch(conf-if-vl-82)# mac access-group test_02 in
```

10 ACL configuration and management

Configuring QoS

In this chapter

- QoS overview 107
- Rewriting 108
- Queueing 108
- Congestion control..... 113
- Multicast rate limiting 116
- Scheduling 117
- CEE map configuration 119

QoS overview

Quality of Service (QoS) provides you with the capability to control how the traffic is moved from switch to switch. In a network that has different types of traffic with different needs (CoS), the goal of QoS is to provide each traffic type with a virtual pipe. FCoE uses traffic class mapping, scheduling, and flow control to provide quality of service.

Traffic running through the switches can be classified as either multicast traffic or unicast traffic. Multicast traffic has a single source but multiple destinations. Unicast traffic has a single source with a single destination. With all this traffic going through inbound and outbound ports, QoS can be set based on egress port and priority level of the CoS.

QoS can also be set on interfaces where the end-station knows how to mark traffic with QoS and it lies with the same trusted interfaces. An untrusted interface is when the end-station is untrusted and is at the administrative boundaries.

The QoS features are:

- Queueing—Queueing provides temporary storage for frames while waiting for transmission. Queues are selected based on ingress ports, egress ports, and configured user priority level.
- Congestion control—When queues begin filling up and all buffering is exhausted, frames are dropped. This has a detrimental effect on application throughput. Congestion control techniques are used to reduce the risk of queue overruns without adversely affecting network throughput. Congestion control features include IEEE 802.3x Ethernet Pause, Tail Drop, and Ethernet Priority Flow Control (PFC).
- Multicast rate limiting—Many multicast applications cannot be adapted for congestion control techniques and the replication of frames by switching devices can exacerbate this problem. Multicast rate limiting controls frame replication to minimize the impact of multicast traffic.

- Scheduling—When multiple queues are active and contending for output on a common physical port the scheduling algorithm selects the order the queues are serviced. Scheduling algorithms include Strict Priority (SP) and Deficit Weighted Round Robin (DWRR) queueing. The scheduler supports a hybrid policy combining SP and DWRR servicing. Under a hybrid scheduler configuration, the highest priority queues are serviced by SP while lower priority queues share the remaining bandwidth using the DWRR service.
- Converged Enhanced Ethernet—CEE describes an enhanced Ethernet that will enable convergence of various applications in data centers (LAN, SAN, and IPC) onto a single interconnect technology.

Rewriting

Rewriting a frame header field is typically performed by an edge device. Rewriting occurs on frames as they enter or exit a network because the neighboring device is untrusted, unable to mark the frame, or is using a different QoS mapping.

The frame rewriting rules set the Ethernet CoS and VLAN ID fields. Egress Ethernet CoS rewriting is based on the user-priority mapping derived for each frame as described later in the queueing section.

Queueing

Queue selection begins by mapping an incoming frame to a configured user priority, then each user-priority mapping is assigned to one of the switch's eight unicast traffic class queues or one of the four multicast traffic class queues.

NOTE

You need to enter the **copy running-config startup-config** command to save your configuration changes to NVRAM so that they are not lost if there is a system reload or power outage.

User-priority mapping

There are several ways an incoming frame can be mapped into a user-priority. If the neighboring devices are untrusted or unable to properly set QoS, then the interface is considered untrusted. All traffic must be user-priority mapped using explicit policies for the interface to be trusted; if it is not mapped in this way, the IEEE 802.1Q default-priority mapping is used. If an interface is trusted to have QoS set then the CoS header field can be interpreted.

NOTE

The user priority mapping described in this section applies to both unicast and multicast traffic.

Default user-priority mappings for untrusted interfaces

When Layer 2 QoS trust is set to *untrusted* then the default is to map all Layer 2 switched traffic to the port default user priority value of 0 (best effort), unless configured to a different value.

Table 23 presents the Layer 2 QoS *untrusted* user priority generation table.

TABLE 23 Default priority value of untrusted interfaces

Incoming CoS	User Priority
0	port <user priority> (default 0)
1	port <user priority> (default 0)
2	port <user priority> (default 0)
3	port <user priority> (default 0)
4	port <user priority> (default 0)
5	port <user priority> (default 0)
6	port <user priority> (default 0)
7	port <user priority> (default 0)

NOTE

Non-tagged Ethernet frames are interpreted as incoming CoS value of 0 (zero).

You can override the default user-priority mapping by applying explicit user-priority mappings.

When neighboring devices are trusted and able to properly set QoS then Layer 2 QoS trust can be set to *COS* and the IEEE 802.1Q default-priority mapping is applied.

Table 24 presents the Layer 2 CoS user priority generation table conforming to 802.1Q default mapping. You can override this default user priority table per port if you want to change (mutate) the CoS value.

TABLE 24 IEEE 802.1Q default priority mapping

Incoming CoS	User Priority
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

Configuring the QoS trust mode

The QoS trust mode controls user priority mapping of incoming traffic. The Class of Service (CoS) mode sets the user priority based on the incoming CoS value. If the incoming packet is not priority tagged, then fallback is to the Interface Default CoS value.

NOTE

When a CEE map is applied on an interface, the **qos trust** command is not allowed. The CEE map always puts the interface in the CoS trust mode.

11 Queueing

To configure the QoS trust mode, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Specify the 10-gigabit Ethernet interface.
`switch(config)#interface tengigabitethernet 0/2`
3. Set the interface mode to 'trust'.
`switch(conf-if-te-0/2)#qos trust cos`
4. Enter the copy command to save the running-config file to the startup-config file.
`switch(conf-if-te-0/2)#do copy running-config startup-config`

Configuring user-priority mappings

To configure user-priority mappings, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Specify the 10-gigabit Ethernet interface.
`switch(config)#interface tengigabitethernet 0/2`
3. Set the interface mode to '3'.
`switch(conf-if-te-0/2)#qos cos 3`
4. Enter the copy command to save the running-config file to the startup-config file.
`switch(conf-if-te-0/2)#do copy running-config startup-config`

Creating a CoS-to-CoS mutation QoS map

To create a CoS-to-CoS mutation, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Create the CoS-to-CoS mutation QoS map name. In this example 'test' is used.
`switch(config)#qos map cos-mutation test 0 1 2 3 5 4 6 7`
3. Enter the copy command to save the running-config file to the startup-config file.
`switch(conf-if-te-0/2)#do copy running-config startup-config`

Applying a CoS-to-CoS mutation QoS map

To apply a CoS-to-CoS mutation QoS map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Specify the 10-gigabit Ethernet interface.
`switch(config)#interface tengigabitethernet 0/2`

3. Activate or apply changes made to the CoS-to-CoS mutation QoS map name. In this example 'test' is used.

```
switch(conf-if-te-0/2)#qos map cos-mutation test
```

4. Specify the trust mode for incoming traffic.

Use this command to specify the interface ingress QoS trust mode, which controls user priority mapping of incoming traffic. The untrusted mode overrides all incoming priority markings with the Interface Default CoS. The CoS mode sets the user priority based on the incoming CoS value, if the incoming packet is not priority tagged, then fallback is to the Interface Default CoS value.

```
switch(conf-if-te-0/2)#qos trust cos
```

5. Enter the copy command to save the running-config file to the startup-config file.

```
switch(conf-if-te-0/2)#do copy running-config startup-config
```

Traffic class mapping

Fabric OS v7.0.1 supports eight unicast traffic classes for isolation and to control servicing for different priorities of application data. Traffic classes are numbered from 0 through 7, with higher values designating higher priority.

The traffic class mapping stage provides some flexibility in queue selection:

- The mapping may be many-to-one, such as mapping one byte user priority (256 values) to eight traffic classes.
- There may be a non-linear ordering between the user priorities and traffic classes.

Unicast traffic

[Table 25](#) presents the Layer 2 default traffic class mapping supported for a COS-based user priority to conform to 802.1Q default mapping.

TABLE 25 Default user priority for unicast traffic class mapping

User priority	Traffic class
0	1
1	0
2	2
3	3
4	4
5	5
6	6
7	7

You are allowed to override these default traffic class mappings per port. Once the traffic class mapping has been resolved it is applied consistently across any queueing incurred on the ingress and the egress ports.

Multicast traffic

Fabric OS v7.0.1 supports four multicast traffic classes for isolation and to control servicing for different priorities of application data. Traffic classes are numbered from 0 through 3, with higher values designating higher priority. The traffic class mapping stage provides some flexibility in queue selection.

Table 26 presents the Layer 2 default traffic class mapping supported for a COS-based user priority to conform to 802.1Q default mapping.

TABLE 26 Default user priority for multicast traffic class mapping

User Priority	Traffic class
0	0
1	0
2	1
3	1
4	2
5	2
6	3
7	3

Once the traffic class mapping has been resolved for ingress traffic, it is applied consistently across all queueing incurred on the ingress and egress ports.

Mapping CoS-to-Traffic-Class

To map a CoS-to-Traffic-Class, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```
2. Create the CoS-Traffic-Class mapping by specifying a name and the mapping.

```
switch(config)#qos map cos-traffic-class test 1 0 2 3 4 5 6 7
```
3. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Activating a mapping CoS-to-Traffic-Class

To activate a CoS-to-traffic class mapping, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```
2. Specify the 10-gigabit Ethernet interface.

```
switch(config)#interface tengigabitethernet 0/2
```
3. Activate the CoS-to-Traffic-Class mapping by name.

```
switch(conf-if-te-0/2)#qos cos-traffic-class test
```

4. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Verifying a mapping CoS-to-Traffic-Class

To verify a CoS-to-Traffic-Class mapping, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Verify the CoS-Traffic-Class mapping specifying a name and the mapping.

```
switch(config)#do show qos map cos-traffic-class test
```

Congestion control

Queues can begin filling up due to a number of reasons, such as over subscription of a link or backpressure from a downstream device. Sustained, large queue buildups generally indicate congestion in the network and can affect application performance through increased queueing delays and frame loss.

Congestion control covers features that define how the system responds when congestion occurs or active measures taken to prevent the network from entering a congested state.


Tail drop

Tail drop queueing is the most basic form of congestion control. Frames are queued in FIFO order and queue buildup can continue until all buffer memory is exhausted. This is the default behavior when no additional QoS has been configured.

The basic tail drop algorithm does not have any knowledge of multiple priorities and per traffic class drop thresholds can be associated with a queue to address this. When the queue depth breaches a threshold, then any frame arriving with the associated priority value will be dropped. [Figure 7](#) describes how you can utilize this feature to ensure that lower priority traffic cannot totally

11 Congestion control

consume the full buffer memory. Thresholds can also be used to bound the maximum queuing delay for each traffic class. Additionally if the sum of the thresholds for a port is set below 100 percent of the buffer memory, then you can also ensure that a single port does not monopolize the entire shared memory pool.



The tail drop algorithm can be extended to support per priority drop thresholds. When the ingress port CoS queue depth breaches a threshold, then any frame arriving with the associated priority value will be dropped. Figure 7 describes how you can utilize this feature to ensure lower priority traffic cannot totally consume the full buffer memory. Thresholds can also be used to bound the maximum queuing delay for each traffic class. Additionally if the sum of the thresholds for a port is set below 100 percent of the buffer memory then you can also ensure that a single CoS does not monopolize the entire shared memory pool allocated to the port.

Changing the Tail Drop threshold

To change the Tail Drop threshold, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Change the Tail Drop threshold for each multicast traffic class. In this example, 1000pkt is used.

```
switch(config)#qos rcv-queue multicast threshold 1000 1000 1000 1000
```

3. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Ethernet pause

Ethernet Pause is an IEEE 802.3 standard mechanism for back pressuring a neighboring device. Pause messages are sent by utilizing the optional MAC control sublayer. A Pause frame contains a 2-byte pause number, which states the length of the pause in units of 512 bit times. When a device receives a Pause frame, it must stop sending any data on the interface for the specified length of

time, once it completes transmission of any frame in progress. You can use this feature to reduce Ethernet frame losses by using a standardized mechanism. However the Pause mechanism does not have the ability to selectively back pressure data sources multiple hops away, or exert any control per VLAN or per priority, so it is disruptive to all traffic on the link.

Ethernet Pause includes the following features:

- All configuration parameters can be specified independently per interface.
- Pause On/Off can be specified independently for TX and RX directions. No support is provided for auto-negotiation.
- Pause generation is based on input (receive) queueing. Queue levels are tracked per input port. You can change the high-water and low-water threshold for each input port. When the instantaneous queue depth crosses the high-water mark then a Pause is generated. If any additional frames are received and the queue length is still above the low-water mark then additional Pauses are generated. Once the queue length drops below the low-water mark then Pause generation ceases.
- A Pause that is received and processed halts transmission of the output queues associated with the port for the duration specified in the Pause frame.

Enabling Ethernet Pause

To enable Ethernet Pause, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Specify the 10-gigabit Ethernet interface.

```
switch(config)#interface tengigabitethernet 0/2
```

3. Enable Ethernet Pause on the interface for both TX and RX traffic.

```
switch(conf-if-te-0/2)#qos flowcontrol tx on rx on
```

4. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Ethernet Priority Flow Control

Ethernet Priority Flow Control (PFC) is a basic extension of the Ethernet Pause. The Pause MAC control message is extended with eight 2-byte pause numbers and a bitmask to indicate which values are valid. Each pause number is interpreted identically to the base Pause protocol; however each is applied to the corresponding Ethernet priority / class level. For example, the Pause number zero applies to priority zero, Pause number one applies to priority one, and so on. This addresses one shortcoming of the Ethernet Pause mechanism, which is disruptive to all traffic on the link. However, it still suffers from the other Ethernet Pause limitations.

Ethernet Priority Flow Control includes the following features:

- Everything operates exactly as in Ethernet Pause described above except there are eight high-water and low-water thresholds for each input port. This means queue levels are tracked per input port plus priority.
- Pause On/Off can be specified independently for TX and RX directions per priority.
- Pause time programmed into Ethernet MAC is a single value covering all priorities.

11 Multicast rate limiting

- Both ends of a link must be configured identically for Ethernet Pause or Ethernet Priority Flow Control because they are incompatible.

Enabling an Ethernet PFC

To enable Ethernet PFC, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Specify the 10-gigabit Ethernet interface.

```
switch(config)#interface tengigabitethernet 0/2
```

3. Enable trust mode on the interface.

```
switch(conf-if-te-0/2)#qos trust cos
```

4. Enable an Ethernet PFC on the interface.

```
switch(conf-if-te-0/2)#qos flowcontrol pfc 3 tx on rx on
```

5. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Multicast rate limiting

Multicast rate limiting provides a mechanism to control multicast frame replication and cap the effect of multicast traffic.

Multicast rate limit is applied to the output of each multicast receive queue. Rate limits apply equally to ingress receive queueing (first level expansion) and egress receive queueing (second level expansion) since the same physical receive queues are utilized. You can set policies to limit the maximum multicast frame rate differently for each traffic class level and cap the total multicast egress rate out of the system.

Multicast rate limiting includes the following features:

- All configuration parameters are applied globally. Multicast rate limits are applied to multicast receive queues as frame replications are placed into the multicast expansion queues. The same physical queues are used for both ingress receive queues and egress receive queues so rate limits are applied to both ingress and egress queueing.
- Four explicit multicast rate limit values are supported, one for each traffic class. The rate limit values represent the maximum multicast expansion rate in packets per second (PPS).

Creating a receive queue multicast rate-limit

To create the receive queue multicast rate-limit, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Create a lower maximum multicast frame expansion rate. In this example, the rate is to 10000 PPS.

```
switch(config)#qos rcv-queue multicast rate-limit 10000 burst 5000
```

3. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Scheduling


Scheduling arbitrates among multiple queues waiting to transmit a frame. Fabric OS v7.0.1 supports both Strict Priority (SP) and Deficit Weighted Round Robin (DWRR) scheduling algorithms. Also supported is the flexible selection of the number of traffic classes using SP-to-DWRR. When there are multiple queues for the same traffic class, then scheduling takes these equal priority queues into consideration.

Strict priority scheduling

Strict priority scheduling is used to facilitate support for latency-sensitive traffic. A strict priority scheduler drains all frames queued in the highest priority queue before continuing on to service lower priority traffic classes. A danger with this type of service is that a queue can potentially starve out lower priority traffic classes.

Figure 8 describes the frame scheduling order for an SP scheduler servicing two SP queues. The higher numbered queue, SP2, has a higher priority.

FIGURE 8 Strict priority schedule – two queues




Deficit weighted round robin scheduling

Weighted Round Robin (WRR) scheduling is used to facilitate controlled sharing of the network bandwidth. WRR assigns a weight to each queue; that value is then used to determine the amount of bandwidth allocated to the queue. The round robin aspect of the scheduling allows each queue to be serviced in a set ordering, sending a limited amount of data before moving onto the next queue and cycling back to the highest priority queue after the lowest priority is serviced.

Figure 9 describes the frame scheduling order for a WRR scheduler servicing two WRR queues. The higher numbered queue is considered higher priority (WRR2) and the weights indicate the network bandwidth should be allocated in a 2:1 ratio between the two queues. In Figure 9 WRR2 should receive 66 percent of bandwidth and WRR1 receives 33 percent. The WRR scheduler tracks the extra bandwidth used and subtracts it from the bandwidth allocation for the next cycle through the queues. In this way, the bandwidth utilization statistically matches the queue weights over longer time periods.

FIGURE 9 WRR schedule – two queues



Deficit Weighted Round Robin (DWRR) is an improved version of WRR. DWRR remembers the excess used when a queue goes over its bandwidth allocation and reduces the queue's bandwidth allocation in the subsequent rounds. This way the actual bandwidth usage is closer to the defined level when compared to WRR.

Traffic class scheduling policy

The traffic classes are numbered from 0 to 7; higher numbered traffic classes are considered higher priority. Fabric OS v7.0.1 provides full flexibility in controlling the number of SP-to-WRR queues. The number of SP queues is specified in N (SP1 through 8), then the highest priority traffic classes are configured for SP service and the remaining eight are WRR serviced. [Table 27](#) describes the set of scheduling configurations supported.

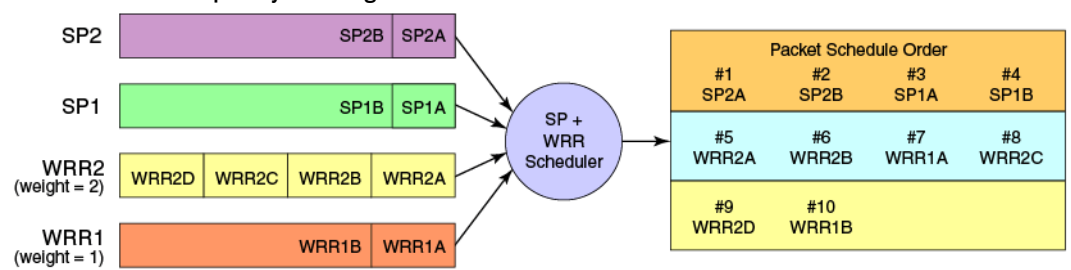
When you configure the QoS queue to use strict priority 4 (SP4), then traffic class 7 will use SP4, traffic class 6 will use SP3, and so on down the list. You use the strict priority mappings to control how the different traffic classes will be routed in the queue.

TABLE 27 Supported scheduling configurations

Traffic Class	SP0	SP1	SP2	SP3	SP4	SP5	SP6	SP8
7	WRR8	SP1	SP2	SP3	SP4	SP5	SP6	SP8
6	WRR7	WRR7	SP1	SP2	SP3	SP4	SP5	SP7
5	WRR6	WRR6	WRR6	SP1	SP2	SP3	SP4	SP6
4	WRR5	WRR5	WRR5	WRR5	SP1	SP2	SP3	SP5
3	WRR4	WRR4	WRR4	WRR4	WRR4	SP1	SP2	SP4
2	WRR3	WRR3	WRR3	WRR3	WRR3	WRR3	SP1	SP3
1	WRR2	WRR2	WRR2	WRR2	WRR2	WRR2	WRR2	SP2
0	WRR1	WRR1	WRR1	WRR1	WRR1	WRR1	WRR1	SP1

[Figure 10](#) shows that extending the frame scheduler to a hybrid SP+WRR system is fairly straightforward. All SP queues are considered strictly higher priority than WRR so they are serviced first. Once all SP queues are drained, then the normal WRR scheduling behavior is applied to the non-empty WRR queues.

FIGURE 10 Strict priority and Weighted Round Robin scheduler



Scheduling the QoS queue

To specify the schedule to use, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

- Specify the schedule to use and the traffic class to bandwidth mapping.

```
switch(config)#qos queue scheduler strict-priority 4 dwrr 10 20 30 40
```

- Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Multicast queue scheduling

The multicast traffic classes are numbered from 0 to 3; higher numbered traffic classes are considered higher priority. A fixed mapping from multicast traffic class to equivalent unicast traffic class is applied to select the queue scheduling behavior. [Table 28](#) presents the multicast traffic class equivalence mapping applied.

TABLE 28 Multicast traffic class equivalence mapping

Multicast traffic class	Equivalent unicast traffic class
3	6
2	4
1	2
0	0

Once the multicast traffic class equivalence mapping has been applied, then scheduling and any scheduler configuration are inherited from the equivalent unicast traffic class. See [Table 27](#) on page 118 for details on exact mapping equivalencies.

Unicast ingress and egress queueing utilizes a hybrid scheduler that simultaneously supports SP+WRR service and multiple physical queues with the same service level. Multicast adds additional multicast expansion queues. Since multicast traffic classes are equivalent to unicast service levels, they're treated exactly as their equivalent unicast service policies.

Scheduling the QoS multicast queue

To schedule the QoS multicast queue, perform the following steps from privileged EXEC mode.

- Enter global configuration mode.

```
switch#configure terminal
```

- Specify the schedule to use and the traffic class to bandwidth mapping.

```
switch(config)#qos queue multicast scheduler dwrr 10 20 30 40
```

- Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

CEE map configuration

The CEE QoS covers frame classification, priority and traffic class (queue) mapping, congestion control, and scheduling. Under the CEE Provisioning model all of these features are configured utilizing two configuration tables, Priority Group Table and Priority Table.

11 CEE map configuration

CEE Priority Group Table defines each Priority Group ID (PGID) and its scheduling policy (Strict Priority versus DWRR, DWRR weight, relative priority), and partially defines the congestion control (PFC) configuration. There are 16 rows in the CEE Priority Group Table. [Table 29](#) presents the default CEE Priority Group Table configuration.

NOTE

Only a single CoS can be mapped to a PFC-enabled priority queue. The CoS number must be identical to the priority queue number. If your configuration violates this restriction an error message displays and the Priority Group Table is set back to the default values.

When the CEE map is applied, and the interface is connected to the CNA, only one strict priority PGID (PGID 15.0 to PGID 15.7) is allowed.

TABLE 29 Default CEE Priority Group Table configuration

PGID	Bandwidth%	PFC
15.0	—	N
15.1	—	N
15.2	—	N
15.3	—	N
15.4	—	N
15.5	—	N
15.6	—	N
15.7	—	N
0	0	N
1	0	N
2	0	N
3	0	N
4	0	N
5	0	N
6	0	N
7	0	N

Strict Priority versus DWRR is derived directly from the PGID value. All PGIDs with prefix 15 receive Strict Priority scheduling policy and all PGIDs in the range 0 through 7 receive DWRR scheduling policy. Relative priority between Priority Group is exactly the ordering of entries listed in the table, with PGID 15.0 being highest priority and PGID 0 being lowest priority. Congestion control configuration is partially specified by toggling the PFC column On or Off. This provides only partial configuration of congestion control because the set of priorities mapped to the Priority Group is not known, which leads into the CEE Priority Table.

CEE Priority Table defines each CoS mapping to Priority Group, and completes PFC configuration. There are eight rows in the CEE Priority Table. [Table 30](#) details the default CEE Priority Table configuration.

TABLE 30 Default CEE priority table

CoS	PGID
0	15.6
1	15.7
2	15.5
3	15.4
4	15.3
5	15.2
6	15.1
7	15.0

Defining a priority group table

To define a priority group table map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Specify the name of the CEE map to define.
`switch(config)#cee-map default`
3. Define the CEE map for PGID 0.
`switch(config-ceemap)#priority-group-table 0 weight 50 pfc`
4. Define the CEE map for PGID 1.
`switch(config-ceemap)#priority-group-table 1 weight 50`
5. Enter the copy command to save the running-config file to the startup-config file.
`switch(config)#do copy running-config startup-config`

Defining a priority-table map

To define a priority-table map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.
`switch#configure terminal`
2. Activate the CEE map mode.
`switch(config)#cee-map default`
3. Define the map.
`switch(config-ceemap)#priority-table 1 1 1 0 1 1 1 15.0`
4. Enter the copy command to save the running-config file to the startup-config file.
`switch(config)#do copy running-config startup-config`

11 CEE map configuration

Applying a CEE provisioning map to an interface

To apply a CEE provisioning map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Specify the 10-gigabit Ethernet interface. In this example, 0/2 is used.

```
switch(config)#interface tengigabitethernet 0/2
```

3. Apply the CEE map on the interface.

```
switch(conf-if-te-0/2)#cee default
```

4. Enter the copy command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Verifying the CEE maps

To verify the CEE map, perform the following steps from privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Verify the CEE map provisioning for a specified name.

```
switch(config)#do show cee maps name
```


Configuring 802.1x Port Authentication

In this chapter

- [802.1x protocol overview](#) 123
- [802.1x configuration guidelines and restrictions](#) 123
- [802.1x authentication configuration tasks](#) 124
- [Interface-specific administrative tasks for 802.1x](#) 124

802.1x protocol overview

The 802.1x protocol defines a port-based authentication algorithm involving network data communication between client-based supplicant software, an authentication database on a server, and the authenticator device. In this situation the authenticator device is the Brocade FCoE hardware.

As the authenticator, the Brocade FCoE hardware prevents unauthorized network access. Upon detection of the new supplicant, the Brocade FCoE hardware enables the port and marks it “unauthorized”. In this state, only 802.1x traffic is allowed. All other traffic, such as DHCP and HTTP, is blocked. The Brocade FCoE hardware transmits an EAP-request to the supplicant, which responds with the EAP-response packet. The Brocade FCoE hardware, which then forwards the EAP-response packet to the RADIUS authentication server. If the credentials are validated by the RADIUS server database, the supplicant may access the protected network resources.

NOTE

802.1x port authentication is not supported by LAG (Link Aggregation Group) or interfaces that participate in a LAG.

NOTE

The EAP-MD5, EAP-TLS, EAP-TTLS and PEAP-v0 protocols are supported by the RADIUS server and are transparent to the authenticator switch.

When the supplicant logs off, it sends an EAP-logoff message to the Brocade FCoE hardware which then sets the port back to the “unauthorized” state.

802.1x configuration guidelines and restrictions

Follow these 802.1x configuration guidelines and restrictions when configuring 802.1x:

- If you globally disable 802.1x, then all interface ports with 802.1x authentication enabled automatically switch to force-authorized port-control mode.

802.1x authentication configuration tasks

The tasks in this section describe the common 802.1x operations that you will need to perform. For a complete description of all the available 802.1x CLI commands for the Brocade FCoE hardware, see the *Converged Enhanced Ethernet Command Reference*.

Configure authentication between the switch and CNA or NIC

For complete information on the **aaaConfig** command, see the *Fabric OS Command Reference* and the *Fabric OS Administrator's Guide*.

NOTE

The **aaaConfig** command attempts to connect to the first RADIUS server. If the RADIUS server is not reachable, the next RADIUS server is contacted. However, if the RADIUS server is contacted and the authentication fails, the authentication process does not check for the next server in the sequence.

To configure authentication, perform the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Add the RADIUS to the switch as the authentication server. This Fabric OS CLI command moves the new RADIUS server to the top of the access list.

```
switch:admin> aaaconfig --add 10.2.2.147 -conf radius 1
```

3. Enter global configuration mode.

```
switch:admin>cmsh
switch#configure t
```

4. Enable 802.1x authentication globally

```
switch(config)#dot1x enable
```

5. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config)#do copy running-config startup-config
```

Interface-specific administrative tasks for 802.1x

It is essential to configure the 802.1x port authentication protocol globally on the Brocade FCoE hardware, and then enable 802.1x and make customized changes for each interface port. Since 802.1x was enabled and configured in “[802.1x authentication configuration tasks](#)”, use the administrative tasks in this section to make any necessary customizations to specific interface port settings.

Configuring 802.1x on specific interface ports

To configure 802.1x port authentication on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.

2. Use the **interface** command to select the interface port to modify.

```
switch(config)#interface tengigabitethernet 1/12
```

3. Use the **dot1x authentication** command to enable 802.1x authentication.

```
switch(config-if-te-1/12)#dot1x authentication
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(config-if-te-1/12)#do copy running-config startup-config
```

Configuring 802.1x timeouts on specific interface ports

NOTE

While you are free to modify the timeouts, Brocade recommends that you leave timeouts set to their default values.

To configure 802.1x timeout attributes on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

```
switch(config)#interface tengigabitethernet 1/12
```

3. Configure the timeout interval.

```
switch(config-if-te-1/12)#dot1x timeout supp-timeout 40
```

Configuring 802.1x re-authentication on specific interface ports

To configure 802.1x port re-authentication on a specific interface port, perform the following steps from privileged EXEC mode. Repeat this task for each interface port you wish to modify.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

```
switch(config)#interface tengigabitethernet 1/12
```

3. Enable 802.1x authentication for the interface port.

```
switch(config-if-te-1/12)#dot1x authentication
```

4. Configure reauthentication for the interface port.

```
switch(config-if-te-1/12)#dot1x reauthentication  
switch(config-if-te-1/12)#dot1x timeout re-authperiod 4000
```

Disabling 802.1x on specific interface ports

To disable 802.1x authentication on a specific interface port, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Use the **interface** command to select the interface port to modify.

12 Interface-specific administrative tasks for 802.1x

```
switch(config)#interface tengigabitethernet 1/12
```

3. Use the **no dot1x port-control** command to disable 802.1x Authentication.

```
switch(conf-if-te-1/12)#no dot1x authentication
```

4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

```
switch(conf-if-te-1/12)#do copy running-config startup-config
```

Configuring IGMP

In this chapter

- About IGMP 127
- Configuring IGMP..... 128
- Configuring IGMP snooping querier 129
- Configuring IGMP snooping mrouter 129
- Monitoring IGMP 130

About IGMP

Multicast Control packet and Data Forwarding through a Layer-2 switch configured with VLANs is most easily achieved by Layer-2 forwarding of received Multicast Packets on all the member ports of the VLAN interfaces. However, this simple approach is not bandwidth efficient, since only a subset of member ports may be connected to devices interested in receiving those Multicast packets. In the worst case scenario the data would get forwarded to all port members of a VLAN with a large number of member ports (for example, all 24 ports), even if only a single VLAN member is interested in receiving the data. Such scenarios can lead to loss of throughput for a switch that gets hit by a high rate of Multicast Data Traffic.

Internet Group Management Protocol (IGMP) snooping is a mechanism by which a Layer-2 switch can effectively address this issue of inefficient Multicast Forwarding to VLAN port members. Snooping involves “learning” forwarding states for Multicast Data traffic on VLAN port members from the IGMP control (Join/Leave) packets received on them. The Layer-2 switch also provides for a way to configure forwarding states statically through the CLI.

NOTE

Brocade Fabric OS 7.0.1 supports IGMPv1 and IGMPv2.

Active IGMP snooping

IGMP snooping is normally passive by nature, as it simply monitors IGMP traffic without filtering. However, active IGMP snooping actively filters IGMP packets to reduce load on the multicast router. Upstream traffic is filtered so that only the minimal quantity of information is sent. The switch ensures the router only has a single entry for the VLAN, regardless of the number of active listeners downstream.

In active IGMP snooping, the router only knows about the most recent member of the VLAN. If there are two active listeners in a VLAN and the original member drops from the VLAN, the switch determines that the router does not need this information as the status of the VLAN remains unchanged. However the next time there is a routine query from the router, the switch will forward the reply from the remaining host to prevent the router from assuming there are no active listeners.

Multicast routing

Multicast routers use IGMP to learn which groups have members on each of their attached physical networks. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

NOTE

“Multicast group memberships” means that at least one member of a multicast group on a given attached network is available.

There are two ways that hosts join multicast routing groups:

- Send an unsolicited IGMP join request
- Send an IGMP join request as a response to a general query from a multicast router

In response to the request, the switch creates an entry in its Layer 2 forwarding table for that VLAN. When other hosts send join requests for the same multicast, the switch adds them to the existing table entry. Only one entry is created per VLAN in the Layer 2 forwarding table for each multicast group.

IGMP snooping suppresses all but one of the host join messages per multicast group and forwards this one join message to the multicast router. The switch forwards multicast traffic for the specified multicast group to the interfaces where the join messages were received.

Configuring IGMP

By default, IGMP snooping is globally disabled on all VLAN interfaces. Refer to the *CEE Command Reference* for complete information about the commands in this section.

To configure IGMP, use the following procedure in privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **ip igmp snooping enable** command to enable IGMP for all interfaces.

This command ensures that IGMP snooping is active on all interfaces.

```
switch(config)#ip igmp snooping enable
```

3. Enter VLAN interface configuration mode.

```
switch(config)#interface vlan 25
```

4. Configure a VLAN port member to be a multi-router interface.

```
switch(conf-if-vl-25)#ip igmp snooping mrouter interface tengigabitethernet 0/1
```

5. Repeat step 4 for each port in the VLAN, as needed.

6. Return to global configuration mode.

```
switch(conf-if-vl-25)#exit
```

7. Enable snooping fast-leave.

```
switch(config)#ip igmp snooping fast-leave
```

8. Activate the default IGMP querier functionality for the VLAN.

```
switch(conf-if-vl-25)#ip igmp snooping querier enable
```

9. *Optional:* Activate the IGMP querier functionality with additional features.

Configuring IGMP snooping querier

If your multicast traffic is not routed because Protocol-Independent Multicast (PIM) and IGMP are not configured, use the IGMP snooping querier in a VLAN.

IGMP snooping querier sends out IGMP queries to trigger IGMP responses from switches that wish to receive IP multicast traffic. IGMP snooping listens for these responses to map the appropriate forwarding addresses.

Refer to the *CEE Command Reference* for complete information about the commands in this section.

To configure the IGMP snooping querier, use the following procedure from privileged EXEC mode.

1. Enter interface configuration mode.

```
switch(config)#interface vlan 25
```

2. Activate the default IGMP querier functionality for the VLAN.

```
switch(config-if-vl-25)#ip igmp snooping querier enable
```

3. Return to privileged EXEC mode.

```
switch(config-if-vl-25)#exit
```

4. Set the IGMP querier interval for the VLAN.

The valid range is from 1 through 18000 seconds. The default value is 125 seconds.

```
switch(config)#ip igmp query-interval 125
```

5. Set the last member query interval.

The valid range is from 100 through 25500 milliseconds. The default value is 1000 milliseconds.

```
switch(config)#ip igmp last-member-query-interval 2000
```

6. Set the Max Response Time (MRT).

The valid range is from 1 through 25 seconds. The default value is 10 seconds.

```
switch(config)#ip igmp query-max-response-time 10
```

Configuring IGMP snooping mrouter

To configure IGMP snooping mrouter, perform this task in global configuration mode.

1. Enter interface configuration mode for the VLAN port member.

```
switch(config)#interface vlan 10
```

2. Configure a the port member to be a multicast router interface.

```
switch(config-if-vl-10)#ip igmp snooping mrouter interface tengigabitethernet 0/1
```

NOTE

Use the **no ip igmp snooping mrouter** command to return to the default setting.

- Return to global configuration mode.

```
switch(conf-if-vl-25)#exit
```

- Configure the mrouter timeout value.

```
switch(config)#ip igmp snooping mrouter-timeout 500
```

Monitoring IGMP

Monitoring the performance of your IGMP traffic allows you to diagnose any potential issues on your switch. This helps you utilize bandwidth more efficiently by setting the switch to forward IP multicast traffic only to connected hosts that request multicast traffic.

Refer to the *CEE Command Reference* for complete information about the commands in this section.

To monitor IGMP snooping, perform the following procedure in EXEC or privileged EXEC mode.

- Enter the **show ip igmp groups** command to display all information on IGMP multicast groups for the switch.

Use this command to display the IGMP database, including configured entries for either all groups on all interfaces, or all groups on specific interfaces, or specific groups on specific interfaces.

```
switch#show ip igmp groups
```

- Use the **show ip igmp statistics** command to display the IGMP statistics for a VLAN or interface.

```
switch#show ip igmp statistics interface vlan 1
```

- Use the **show ip igmp mrouter** to display multicast router (mrouter) port related information for all VLANs, or a specific VLAN.

```
switch#show ip igmp snooping mrouter
```

- or -

```
switch#show ip igmp snooping mrouter interface vlan 1
```

- Use the **show ip igmp interface** command to display information relating to the IGMP configuration on the specified VLAN interface.

```
switch#show ip igmp interface vlan 100
```

- When you have reviewed the IGMP statistics for the switch, refer to [“Configuring IGMP”](#) on page 128 or [“Configuring IGMP snooping querier”](#) on page 129 to make any needed corrections.

NOTE

Refer to the *CEE Command Reference* for additional information on IGMP CLI commands.

Configuring RMON

In this chapter

- [RMON overview](#) 131
- [RMON configuration and management](#) 131

RMON overview

Remote monitoring (RMON) is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. The RMON specification defines a set of statistics and functions that can be exchanged between RMON-compliant console managers and network probes. As such, RMON provides you with comprehensive network-fault diagnosis, planning, and performance-tuning information.

RMON configuration and management

Alarms and events are configurable RMON parameters:

- **Alarms**—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- **Events**—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

Default RMON configuration

By default, no RMON alarms and events are configured and RMON collection statistics are not enabled.

Configuring RMON group statistics collection

You can collect RMON group statistics on an interface. RMON alarms and events must be configured for you to display collection statistics. By default, RMON group statistics are not enabled. The statistics are measured by the probe for each monitored interface on the switch. These statistics include, but aren't limited to the following items:

- Packets dropped and sent
- Bytes sent (octets)

- Broadcast packets
- Multicast packets
- CRC errors

To collect RMON group statistics on an interface and enable collection history, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
`switch#configure terminal`
2. Enter the **interface** command to specify the CEE interface type and slot/port number.
`switch(config)#interface tengigabitethernet 0/1`
3. Enable the CEE interface.
`switch(conf-if-te-0/1)#no shutdown`
4. Configure RMON group statistics on the interface.
`switch(conf-if-te-0/1)#rmon collection stats 200 owner john_smith`
5. Enter the **copy** command to save the running-config file to the startup-config file.
`switch(conf-if-te-0/1)#do copy running-config startup-config`

Configuring RMON settings

To configure RMON alarms and events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
`switch#configure terminal`
2. Configure the RMON alarms.

Example of an alarm that tests every sample for a rising threshold

```
switch(config)#rmon alarm 1 1.3.6.1.2.1.16.1.1.1.5.65535 interval 10 absolute  
rising-threshold 10 event 1 falling-threshold 10 event 20 owner john_smith
```

Example of an alarm that tests the delta between samples for a falling threshold

```
switch(config)#rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 10 delta  
falling-threshold 65 event 42 owner john_smith
```

3. Enter the **copy** command to save the running-config file to the startup-config file.
`switch(config)#do copy running-config startup-config`

Configuring RMON events

You can add or remove an event in the RMON event table that is associated with a RMON alarm number.

To configure RMON events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
`switch#configure terminal`

2. Configure the RMON event.

```
switch(config)#rmon event 27 description Rising_Threshold log owner john_smith
trap syslog
```

3. Enter the **copy** command to save the running-config file to the startup-config file.

```
switch(config)#do copy running-config startup-config
```

Configuring RMON history collection

The **RMON collection history** command collects the periodic statistical samples of RMON group statistics on an interface for later retrieval. By default, the collection of RMON history is not enabled. The periodic statistical samples include, but are not limited to the following items:

- Sample period
- Number of samples
- Items sampled

To enable collection history, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch#configure terminal
```

2. Enter the **interface** command to specify the CEE interface type and slot/port number.

```
switch(config)#interface tengigabitethernet 0/1
```

3. Enable the CEE interface.

```
switch(conf-if-te-0/1)#no shutdown
```

4. Activate RMON history collection for the interface.

```
switch(conf-if-te-0/1)#rmon collection history 1 buckets 20 interval 10 owner
john_smith
```

5. Return to privileged EXEC mode.

```
switch(conf-if-te-0/1)#exit
switch(config)#end
```

6. Display the RMON history.

```
switch#show rmon history
RMON history control entry 1
  interface: ifIndex.1745682445 TengigabitEthernet 0/13
  buckets requested: 20
  buckets granted: 20
  sampling interval: 10
  Owner: jsmith
```

7. Enter the **copy** command to save the running-config file to the startup-config file.

```
switch#copy running-config startup-config
```

14 RMON configuration and management

Brocade 8000 Replacement

It is important that the new Brocade 8000 has the identical setup to the original Brocade 8000. This process updates the firmware and configuration files on the new Brocade 8000 to match the original Brocade 8000, then you move the data cables from the original switch to the new switch.

Refer to the *CEE Command Reference* and *Fabric OS Command Reference* for complete information on the commands referenced in this procedure.

NOTE

This procedure assumes that the Brocade 8000 being replaced is still functional on the network. However, if original Brocade 8000 is non-functional, the information from the original switch must have been previously gathered.

If the original Brocade 8000 is still active on the network, you must assign the replacement Brocade 8000 a temporary IP address, resulting in both the original and replacement Brocade 8000s being active on the network at the same time.

NOTE

Before you begin this procedure, obtain all non-factory installed licenses on the original switch that are required on the replacement switch.

To replace a Brocade 8000, perform the following steps.

1. Use the **firmwreshow** command to determine the firmware version of a given switch. Refer to the *Fabric OS Administrator's Guide* for details on installing and maintaining firmware.

```
new_switch:admin> firmwreshow
Appl      Primary/Secondary Versions
-----
FOS       v6.3.2b
          v6.3.2b
```

```
original_switch:admin> firmwreshow
Appl      Primary/Secondary Versions
-----
FOS       v6.4.0
          v6.4.0
```

2. Upgrade the firmware on the new switch to match the firmware of the original switch.
3. Use the **ipaddrset** command to set the IP address on the new switch to be on the same subnet as the original switch.

```
original_switch:admin> ipaddrshow
SWITCH
Ethernet IP Address [10.24.51.23]:
Ethernet Subnetmask: 255.255.252.0
Gateway IP Address: 10.24.48.1
DHCP: Off
IPv6 Autoconfiguration Enabled: No
Local IPv6 Addresses:
static fec0:60:69bc:60:260:69ff:fed0:107/64 preferred
```

IPv6 Gateways:

```
new_switch:admin> ipaddrset
Ethernet IP Address [10.77.77.77]: 10.24.51.25
Ethernet Subnetmask [255.255.255.0]:255.255.252.0
Gateway IP Address [10.77.77.1]:10.24.48.1
DHCP [Off]:
```

4. Use the **chassisname** command to learn the chassis name of the original switch.

```
original_switch:admin> chassisname
ESNSVT_8000
```

5. Use the **chassisname** command to set the chassis name on the new switch to match that of the original switch.

```
new_switch:admin> chassisname "ESNSVT_8000"
ESNSVT_8000
```

6. Use the **licenseshow** command on both switches to ensure you have the correct licenses available for the new switch.

```
original_switch:admin> licenseshow
by9RzReccyczAS0c:
  Fabric license
SeR9Qz9yRzTvSAR3:
  Ports on Demand license - additional 8 port upgrade license
SeR9Qz9yRzXfSARr:
  Performance Monitor license
SeR9Qz9yRzbfSARv:
  Trunking license
SeR9Qz9yRzVfSARp:
  Fabric Watch license
R9A4fF9ZfttXDWgXM9HX3gRHa4rG9NRaBAGXD:
  Enhanced Group Management license
fLTJNCKgmGY3rB743A34GaYCR4YZfgYMB7P9M:
  FCoE Base license
Z3LE7M3MNHQPQ4JDHfgESH9T9ZT77HATrBJf4A:
  Enhanced Group Management license

new_switch:admin> licenseshow
RFFZf49B3MLmQJWEYCRRN4rZNEJt9aJWBJaJM:
  Enhanced Group Management license
SMtS7FRfTLYYQWEPXt4MLTaPKGT7Q7YKBJEPA:
  FCoE Base license
```

NOTE

If your original switch contains licenses for features that are not native to the Brocade 8000, contact your Support Representative for assistance in transferring these licenses. Be sure to have the WWN for both switches. See [“Getting technical help”](#) on page xxi for details.

7. Use the **licenseadd** command on the new switch to install any transferred licenses from the original switch.

```
new_switch:admin> licenseadd
Usage: licenseadd "license-key"
```

8. Take the new switch offline.

You will be uploading the configuration of the original switch.

```
new_switch:admin> switchdisable
```

9. Use the **configupload** command to save the SAN configuration settings of the original switch.

```
original_switch:admin> configupload
Protocol (scp, ftp, local) [ftp]:
Server Name or IP Address [host]: 10.243.80.42
User Name [user]: anonymous
Path/Filename [<home dir>/config.txt]: /8000_configs_dir/MP_8000B_config.txt
Section (all|chassis|switch [all]):
```

configUpload complete: All selected config parameters are uploaded

10. Use the **configdownload** command to download the SAN configuration settings to the new switch.

```
new_switch:admin> configdownload
Protocol (scp, ftp, local) [ftp]:
Server Name or IP Address [host]: 10.243.80.42
User Name [user]: anonymous
Path/Filename [<home dir>/config.txt]: /8000_configs_dir/MP_8000B_config.txt
Section (all|chassis|switch [all]):
```

*** CAUTION ***

This command is used to download a backed-up configuration for a specific switch. If using a file from a different switch, this file's configuration settings will override any current switch settings. Downloading a configuration file, which was uploaded from a different type of switch, may cause this switch to fail. A switch reboot might be required for some parameter changes to take effect.

configDownload operation may take several minutes to complete for large files.

Do you want to continue [y/n]: **y**
Password:

Doing configDownload on switch ...

Activating configDownload without disabling switch:

NOTE: Not all configuration parameters will be downloaded with switch online

configDownload complete: All selected config parameters are downloaded

NOTE

The **configdownload** command updates the switch with the correct Domain ID. However, if the original switch is dead and no configuration upload file from it exists, the one parameter you'll want to setup properly on the replacement switch is its Domain ID. To configure a switch's Domain ID, when the switch is a disabled state, enter the **configure** command. When prompted "Fabric parameters (yes, y, no, n): [no]", type in **yes** and you'll be prompted with the existing Domain ID as follows "Domain: (1..239) [7]". At this point, enter the value of the original switch's Domain ID.

11. Use the **cfgactvshow** command on both switches to ensure the effective zoning configurations match:

```
original_switch:admin> cfgactvshow
```

```
Effective Configuration:
  cfg: host1_fcoe
  zone: w2k3_cna1:
                                          21:01:00:00:c8:98:ab:78
.....output truncated ....
```

```
new_switch:admin> cfgactvshow
```

```
Effective Configuration:
  cfg: host1_fcoe
  zone: w2k3_cna1:
                                          21:01:00:00:c8:98:ab:78
.....output truncated ....
```

12. Bring the new switch back online.

```
new_switch:admin> switchenable
```

13. Copy the ethernet settings of the original switch, using **copy startup-config** and **copy running-config**, to an FTP device.

```
original_switch:admin#cmsh
original_switch#copy startup-config ftp://admin@10.10.10.10/file1
Source password[]?
Building configuration...
original_switch#copy running-config ftp://admin@10.10.10.10/file2
Source password[]?
Building configuration...
```

14. Upload the **startup-config** file to the new switch.

```
new_switch:admin#cmsh
new_switch#copy ftp://admin@10.10.10.10/file1 startup-config
```

15. Upload the **running-config** file to the new switch.

```
new_switch#copy ftp://admin@10.10.10.10/file2 running-config
Source password[]?
Building configuration...
```

NOTE

In the examples above, the IP address of the FTP server is listed as '10.10.10.10' – you must enter the correct IP address of your FTP server.

16. To verify the Ethernet configuration has been downloaded successfully to the replacement switch, use the **show running-config** command on both switches to ensure the running configurations match.

```
original_switch#show running-config
!
Protocol spanning-tree rstp
!
vlan classifier rule 1 proto fcoe encap ethv2
vlan classifier rule 2 proto fip encap ethv2
vlan classifier group 1 add rule 1
vlan classifier group 1 add rule 2
!
.....output truncated .....

original_switch#exit

new_switch#show running-config
```



```

!
Protocol spanning-tree rstp
!
vlan classifier rule 1 proto fcoe encap ethv2
vlan classifier rule 2 proto fip encap ethv2
vlan classifier group 1 add rule 1
vlan classifier group 1 add rule 2
!
.....output truncated .....

new_switch#exit

```

17. Use the **switchshow** command to learn the switch name of the original switch.

```

original_switch:admin> switchname
original_switch

```

18. Use the **switchname** command to set the switch name on the new switch to match that of the original switch.

The new switch name is not reflected until the next user login.

```

new_switch:admin> switchname "original_switch"
original_switch

```

19. Remove the Ethernet cable from the original switch.
20. If a temporary IP address was being used on the new switch during the replacement process, use the **ipaddrset** command to set the replacement Brocade 8000's IP address to match the original Brocade 8000's IP address.

```

new_switch:admin> ipaddrset
Ethernet IP Address [10.24.51.25]:10.24.51.23
Ethernet Subnetmask [255.255.252.0]:
Gateway IP Address [10.24.48.1]:
DHCP [Off]:

```

21. One at a time, transfer the data cables from the original switch to the new switch.

NOTE

Wait for the lights on the new port to match the lights on the original port before proceeding to the next cable.

22. When all the cables have been moved, remove the original switch from the rack.

A Brocade 8000 Replacement

Index

Symbols

Numerics

- 802.1x
 - LAG, 123
 - overview, 123
 - timeouts, 125

A

Access Control Lists

See ACL

access interface, configuring, 50

access mode, 45, 50

ACL

configuration guidelines and restrictions, 102

configuration procedures

applying a MAC ACL to a CEE interface, 105

applying a MAC ACL to a VLAN interface, 105

creating extended MAC ACL and adding rules, 103

creating standard MAC ACL and adding rules, 102

important notes, 102

modifying a MAC ACL, 103

removing a MAC ACL, 104

reordering the sequence numbers, 104

default configuration, 102

extended ACL, defined, 101

overview, 7, 101

standard ACL, defined, 101

active IGMP, 127

audit log, 42

authentication server, 123

authenticator, 123

B

basic management TLV sets, 88

bridge

forwarding delay, configuring for STP, RSTP, MSTP, 66

hello time, configuring for STP, RSTP, 68

maximum aging time, configuring for STP, RSTP, MSTP, 66

priority, configuring for STP, RSTP, MSTP, 65

Brocade

Brocade Connect, xx

extension TLV set, 89

proprietary aggregation, 82

website, xx

Brocade FCoE hardware, 2

C

CEE audit log, 42

CEE interface

applying a MAC ACL, 105

configuring for STP, RSTP, MSTP, 72

configuring the hello time for MSTP, 74

disable or enable STP on the interface, 76

enabling and disabling, 48

enabling as an edge port for RSTP, MSTP, 73

enabling guard root for STP, RSTP, MSTP, 73

enabling LACP, 83

enabling port fast, 75

path cost, configuring for STP, RSTP, MSTP, 72

restricting the port from becoming a root port for STP, RSTP, MSTP, 76

restricting the topology change notification for STP, RSTP, MSTP, 76

spanning-tree defaults, 64

specifying a link type, 75

specifying restrictions for an MSTP instance, 74

specifying the port priority for STP, RSTP, MSTP, 75

CEE map, 39

CEE maps, verifying, 41, 122

CEE provisioning map, applying, 41, 122

- CEE switch
 - congestion control and queuing, 5
 - flow control, 8
 - Layer 2 Ethernet, 3
 - Layer 2 forwarding, 3
 - loop-free, 4
 - tagging, 4
 - trunking, 7
- Cisco interoperability, disabling for MSTP, 69
- Cisco interoperability, enabling for MSTP, 69
- classifier groups, VLAN, 53
- classifier rules, VLAN, 52
- CLI, CEE
 - accessing, 14
 - command completion, 19
 - command modes, 15
 - console and VTY (line) configuration, 17
 - EXEC, 16
 - feature configuration, 17
 - global configuration, 16
 - interface configuration, 16
 - Privileged EXEC, 16
 - protocol configuration, 16
 - command syntax, 18
 - configuration guidelines and restrictions, 13
 - displaying commands, 18
 - keyboard shortcuts, 17
 - output modifiers, 19
 - RBAC permissions, 14
- cmsh command, 14
- command completion, CEE CLI, 19
- command modes, CEE, 15
- command output modifiers, 19
- command syntax, 18
- configuration management
 - saving changes, 14
- configure
 - FCoE VLAN, 35
 - VLAN membership, 38
- congestion control
 - QoS, 113
 - queuing, 5
- console interface, 14
- control processors, 33
- converged mode, 45
- counters, clearing, 54

D

- Data Center Bridging (DCB) Capability Exchange Protocol
 - See DCBX
- DCBX
 - Enhanced Transmission Selection, 90
 - interaction with other vendor devices, 91
 - overview
 - Priority Flow Control, 90
 - TLV sets, 37
- document conventions, *xviii*
- DWRR, 39
- dynamic link aggregation, 82

E

- EAP, 123
- edge detection, configuring for STP, RSTP, MSTP, 72
- edge port, enabling a CEE interface as an edge port for
 - RSTP, MSTP, 73
- Enhanced Transmission Selection
 - See ETS
- error disable timeout interval, configuring for STP, RSTP,
 - MSTP, 67
- error disable timeout, configuring for STP, RSTP, MSTP, 67
- Ethernet, forwarding, 3
- ETS
 - overview
 - priority grouping of IPC, LAN, and SAN traffic, 90

F

- fabric OS shell, 15
- failover, 33
- FCoE
 - configuration guidelines and restrictions, 25
 - configuration procedures
 - creating and managing the FCoE login group configuration, 28
 - managing and displaying FCoE login

- configuration, 26
- Layer 2 Ethernet overview, 3
- login, 10
- logout, 10
- overview, 1
- queuing, 12
- speed negotiation, 25
- terminology
 - CEE, 1
 - ENode, 1
 - FCoE Forwarder (FCF), 1
 - VF_port, 1
 - VN_port, 1
- unsupported features, 25
- VLAN forwarding, 4
- FCoE initialization protocol
 - See FIP
- feedback, *xxii*
- Fibre Channel Association, *xxi*
- filtering VLAN ingress, 45
- FIP
 - FC zoning, 11
 - FCoE login, 10
 - FCoE logout, 10
 - FIP discovery, 8
 - login, 9
 - logincfg, 10
 - logout, 9
 - name server, 11
 - registered state change notification (RSCN), 11
- FLOGI, 1
- flow control, 8
- flushing MAC addresses, 71
- frame classification, incoming, 5

G

- glossary, *xx*
- guard root, enabling on a CEE interface for STP, RSTP, MSTP, 73

H

- hello time, configuring for MSTP, 74
- high availability, 33
- hops, configuring for MSTP, 70

I

- IEEE 802.1 organizational TLV set, 89
- IEEE 802.3 organizational TLV set, 89
- IGMP
 - interface, 128
 - interval, 129
 - mrouter, 128
 - MRT, 129
 - passive, 127
 - querier, 129
 - query-interval, 128
 - tcn, 128
 - timer, 128
 - vlan, 128
- incoming frame classification, 5
- ingress VLAN filtering, 45
- instance
 - MSTP, mapping a VLAN to, 69
 - specifying restrictions for an MSTP instance, 74
- interface, 35, 36
- iSCSI, 96

K

- key terms, *xx*
- keyboard shortcuts, CEE CLI, 17

L

- LACP
 - configuration guidelines and restrictions, 83
 - configuration procedures
 - clearing counters, 85
 - configuring system priority, 84
 - configuring timeout period, 84
 - displaying LACP information, 85
 - enabling on a CEE interface, 83
 - important notes, 83
 - default LACP configuration, 83
 - overview
 - troubleshooting tips, 85
- LAGs
 - 802.1x, 123
 - distribution process, 83
 - overview
 - top-of-the-rack configuration, 81

- Layer 2
 - ACL
 - Ethernet overview, 3
- Layer 2 forwarding, 3
- link aggregation
 - Brocade-proprietary, 82
 - dynamic, 82
 - LACP, 82
 - LAG distribution process, 83
 - LAGs, 79
 - overview, 79
 - static, 82
- Link Aggregation Control Protocol
 - See LACP
- link aggregation group
 - See LAGs
- Link Layer Discovery Protocol
 - See LLDP
- link type, specifying, 75
- LLDP
 - configuration guidelines and restrictions, 91
 - configuration procedures
 - clearing LLDP-related information, 98
 - disabling LLDP globally, 92
 - displaying LLDP-related information, 98
 - enabling LLDP globally, 92
 - global command options, 92
 - important notes, 92
 - interface-level command options, 98
 - DCBX overview
 - default configuration, 91
 - Layer 2 topology mapping, 88
 - overview, 87
 - TLV sets, 88
- login
 - FCoE, 10
 - FIP, 9
- logincfg, 10
- logout
 - FCoE, 10
 - FIP, 9
- loop-free network environment, 4

M

- MAC addresses
 - configuration guidelines and restrictions, 47
 - flush from the VLAN FDB, 71
- map, 36

- MSTP, 38
 - configuration procedures, 64
 - default configuration, 63
 - displaying MSTP-related information, 72
 - overview, 61
- MTU, configuring, 48
- multicast rate limiting, QoS, 116
- Multiple Spanning Tree Protocol
 - See MSTP

N

- name server, 11
- network
 - flow control, 8
 - loop-free
 - STP, RSTP, MSTP, 4
 - trunking, 7

O

- output modifiers, CEE CLI, 19
- overview
 - ACL, 101
 - link aggregation, 79
 - MSTP, 61
 - RSTP, 59
 - STP, 57

P

- passive IGMP, 127
- path cost
 - CEE interface, configuring for STP, RSTP, MSTP, 72
 - port channel, configuring for STP, RSTP, MSTP, 67
- PEAP, 123
- PFC, 39
- PGID, 39
- port configuration for STP, RSTP, MSTP, 72
- port fast, enabling on a CEE interface, 75
- port priority, specifying on a CEE interface for STP, RSTP, MSTP, 75
- Priority Flow Control (PFC), 90
- priority group table, mapping, 40, 121
- priority mapping, QoS, 108
- priority-table, mapping, 41, 121

Q

QoS

- CEE QoS overview, 39, 119
- configuration procedures
 - applying a CEE provisioning map, 41, 122
 - mapping a priority group table, 40, 121
 - mapping a priority-table, 41, 121
 - overview, 39, 119
 - verifying CEE maps, 41, 122
- congestion control, 113
- multicast rate limiting, 116
- overview, 107
- queuing
 - traffic class mapping, 111
 - user-priority mapping, 108
- queuing overview, 108
- rewriting frame header field, 108
- scheduling, 117

Quality of Service

See QoS

querier

- interval, 129
- MRT, 129
- VLAN, 129

queuing

- congestion control, 5
- FCoE, 12
- QoS, 108

R

RADIUS, 123

Rapid Spanning Tree Protocol

See RSTP

RBAC permissions

- region name, specifying for MSTP, 70
- registered state notification protocol (RSCN), 11
- revision number, specifying for MSTP, 70

Role-Based Action Control

See RBAC

root port, CEE interface, restricting for STP, RSTP, MSTP, 76

RSTP, 38

- configuration guidelines and restrictions
 - MSTP configuration guidelines and restrictions,

63

- configuration procedures, 64
- default configuration, 63
- displaying RSTP-related information, 72
- overview, 59

S

saving configuration, 14

scheduling, QoS, 117

Spanning Tree Protocol

See STP

spanning-tree defaults, 64

speed negotiation, FC ports, 25

standby, 33

static link aggregation, 82

STP, 38

configuration guidelines and restrictions, 63

configuration procedures, 64

default configuration, 63

displaying STP-related information, 72

overview, 57

supplicant, 123

switch

port configuration, 50

serial number, *xxi*

system priority, configuring for LACP, 84

T

T11-FC-BB5, 1

technical help, *xxi*

telnet, 14

terminology

document, *xx*

FCoE, 1

timeout period, configuring for LACP, 84

TLV sets

basic management TLV, 88

Brocade extension TLV set, 89

configuring, 37

IEEE 802.1 organizational TLV set, 89

IEEE 802.3 organizational TLV set, 89

top-of-the-rack configuration, 81

topology change notification, CEE interface, restricting for STP, RSTP, MSTP, 76

topology mapping, LLDP, 88

traffic class mapping, QoS, 111

- transmit hold count, configuring for RSTP, MSTP, 68
- troubleshooting tips, LACP, 85
- trunk interface, configuring, 50
- trunk mode, 45, 50
- trunking, 7

U

- unsupported features, 25
- user-priority mapping, QoS, 108

V

Virtual LANs

- See VLAN

VLAN

- applying a MAC ACL, 105
- configuration guidelines and restrictions, 47
- configuration procedures
 - clearing VLAN counters, 54
 - configuring a CEE interface as a Layer 2 switch port, 50
 - configuring a CEE interface as an access or trunk interface, 50
 - configuring the MTU on an interface, 48
 - displaying VLAN information, 54
 - enabling and disabling a CEE interface, 48
 - important notes, 48
 - VLAN classifier groups, 53
 - VLAN classifier rules, 52
- default configuration, 47
- FDB
 - flushing, 71
 - overview, 46
- forwarding, 4
- important management notes, 48
- ingress VLAN filtering, 45
- overview, 45
- tagging, 4

W

- website, Brocade, xx

Z

- zoning, FC, 11