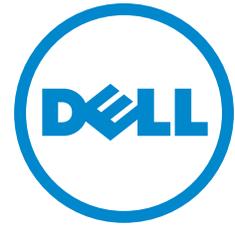


53-1002116-01
07 December 2010



Dell Converged Enhanced Ethernet

Administrator's Guide

Information in this document is subject to change without notice.

© 2010 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *Inspiron*, *Dell Precision*, *Dimension*, *OptiPlex*, *Latitude*, *PowerEdge*, *PowerVault*, *PowerApp*, and *Dell OpenManage* are trademarks of Dell Inc.; *Intel*, *Pentium*, and *Celeron* are registered trademarks of Intel Corporation in the U.S. and other countries; *Microsoft*, *Windows*, *Windows Server*, *MS-DOS* and *Windows Vista* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

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

Regulatory Model Code: M8428-k

Contents

About This Document

In this chapter	xv
How this document is organized	xv
Supported hardware and software	xvi
Document conventions	xvi
Text formatting	xvi
Command syntax conventions	xvii
Notes, cautions, and warnings	xvii
Notice to the reader	xviii
Getting technical help	xviii
Contacting Dell	xviii

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	5
Frame classification (incoming)	5
Congestion control and queuing	6
Access control	7
Access Gateway	7
Trunking	8
Flow Control	8
FCoE Initialization Protocol	8
FIP discovery	8
FIP login	9
FIP logout	10
FCoE login	10
FCoE logout	10
Logincfg	11
Name server	11
FC zoning	11
Registered State Change Notification (RSCN)	12
FCoE queuing	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	15
	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
	Internal and external 10 Gbps Ethernet interfaces syntax	20
Chapter 3	Initial FCoE and CEE Configuration	
	In this chapter	21
	Overview	21
	Configuring the FCoE interfaces	21
	Configuring FCoE VLAN	22
	Assigning FCoE map on to an interface	23
	Configuring the CEE interfaces	23
	Configuring DCBX	24
	Configuring Spanning Tree Protocol.	25
	Configuring VLAN membership	25
	Configuring protect mode	26
Chapter 4	Configuring IP static routes	
	In this chapter	27
	IP static routes overview.	27
	Configuring IP static routes	28
	Recursive IP static routes	29
	Floating IP static routes.	29
	Displaying and clearing IP static routes	29
Chapter 5	Configuring VLANs Using the CEE CLI	
	In this chapter	31
	VLAN overview.	31
	Ingress VLAN filtering	31
	VLAN configuration guidelines and restrictions	33
	Default VLAN configuration	33

VLAN configuration and management	34
Enabling and disabling an interface port	34
Configuring the MTU on an interface port	34
Creating a VLAN interface	35
Enabling STP on a VLAN	35
Disabling STP on a VLAN	35
Configuring a VLAN interface to forward FCoE traffic	36
Configuring an interface port as a Layer 2 switch port	36
Configuring an interface port as an access interface	36
Configuring an interface port as a trunk interface	37
Disabling a VLAN on a trunk interface	37
Configuring an interface port as a converged interface	37
Disabling a VLAN on a converged interface	38
Configuring protocol-based VLAN classifier rules	38
Configuring a VLAN classifier rule	39
Configuring MAC address-based VLAN classifier rules	39
Deleting a VLAN classifier rule	39
Creating a VLAN classifier group and adding rules	39
Activating a VLAN classifier group with an interface port	40
Clearing VLAN counter statistics	40
Displaying VLAN information	40
Configuring the MAC address table	40
Specifying or disabling the aging time for MAC addresses	41
Adding static addresses to the MAC address table	41

Chapter 6

Configuring STP, RSTP, and MSTP using the CEE CLI

In this chapter	43
STP overview	43
Configuring STP on Dell FCoE hardware	44
RSTP overview	45
MSTP overview	47
Configuring MSTP on Dell FCoE hardware	48
STP, RSTP, and MSTP configuration guidelines and restrictions	49
Default STP, RSTP, and MSTP configuration	50

STP, RSTP, and MSTP configuration and management	51
Enabling STP, RSTP, or MSTP	51
Disabling STP, RSTP, or MSTP	51
Shutting down STP, RSTP, or MSTP globally	51
Specifying the bridge priority.	52
Specifying the bridge forward delay	52
Specifying the bridge maximum aging time.	53
Enabling the error disable timeout timer	53
Specifying the error disable timeout interval.	53
Specifying the port-channel path cost	54
Specifying the bridge hello time (STP and RSTP).	54
Specifying the transmit hold count (RSTP and MSTP).	54
Enabling Cisco interoperability (MSTP).	55
Disabling Cisco interoperability (MSTP)	55
Mapping a VLAN to an MSTP instance	55
Specifying the maximum number of hops for a BPDU (MSTP)	56
Specifying a name for an MSTP region.	56
Specifying a revision number for an MSTP configuration	56
Flushing MAC addresses (RSTP and MSTP).	57
Clearing spanning tree counters.	57
Clearing spanning tree-detected protocols	57
Displaying STP, RSTP, and MSTP-related information	58
Configuring STP, RSTP, or MSTP on CEE interface ports	58
Enabling automatic edge detection	58
Configuring the path cost	58
Enabling a port (interface) as an edge port.	59
Enabling the guard root.	59
Specifying the MSTP hello time.	60
Specifying restrictions for an MSTP instance	60
Specifying a link type.	61
Enabling port fast (STP).	61
Specifying the port priority	61
Restricting the port from becoming a root port	62
Restricting the topology change notification	62
Enabling spanning tree	62
Disabling spanning tree.	63

Chapter 7

Configuring Link Aggregation using the CEE CLI

In this chapter	65
Link aggregation overview	65
Link Aggregation Group configuration	65
Link Aggregation Control Protocol.	68
Dynamic link aggregation	68
Static link aggregation.	68
Dell-proprietary aggregation	68
LAG distribution process	68
LACP configuration guidelines and restrictions	69
Default LACP configuration.	69

LACP configuration and management	69
Enabling LACP on a CEE interface	69
Configuring the LACP system priority	70
Configuring the LACP timeout period on a CEE interface	70
Configuring minimum links feature	70
Configuring interface tracking	71
Clearing LACP counter statistics on a LAG	71
Clearing LACP counter statistics on all LAG groups	72
Displaying LACP information	72
LACP troubleshooting tips	72

Chapter 8

Configuring LLDP using the CEE CLI

In this chapter	75
LLDP overview	75
Layer 2 topology mapping	76
DCBX overview	78
Enhanced Transmission Selection (ETS)	78
Priority Flow Control (PFC)	79
DCBX interaction with other vendor devices	79
LLDP configuration guidelines and restrictions	79
Default LLDP configuration	80
LLDP configuration and management	80
Enabling LLDP globally	80
Disabling and resetting LLDP globally	80
Configuring LLDP global command options	81
Configuring LLDP interface-level command options	84
Clearing LLDP-related information	85
Displaying LLDP-related information	85

Chapter 9

Configuring ACLs using the CEE CLI

In this chapter	87
ACL overview	87
Default ACL configuration	88
ACL configuration guidelines and restrictions	88
ACL configuration and management	88
Creating a standard MAC ACL and adding rules	88
Creating an extended MAC ACL and adding rules	89
Modifying MAC ACL rules	89
Removing a MAC ACL	90
Reordering the sequence numbers in a MAC ACL	90
Applying a MAC ACL to a CEE interface	91
Applying a MAC ACL to a VLAN interface	91

Chapter 10	Configuring QoS using the CEE CLI	
	In this chapter	93
	QoS overview	93
	Rewriting	94
	Queueing	94
	User-priority mapping.	94
	Traffic class mapping.	97
	Congestion control	100
	Tail drop	100
	Ethernet pause.	101
	Ethernet Priority Flow Control	102
	Multicast rate limiting.	103
	Scheduling.	104
	Strict priority scheduling	104
	Deficit weighted round robin scheduling	104
	Traffic class scheduling policy.	105
	Multicast queue scheduling	106
	Converged Enhanced Ethernet map configuration.	107
Chapter 11	Configuring 802.1x Port Authentication	
	In this chapter	111
	802.1x protocol overview	111
	802.1x configuration guidelines and restrictions.	111
	802.1x authentication configuration tasks.	112
	Configure authentication between the switch and CNA or NIC.	112
	Interface-specific administrative tasks for 802.1x.	112
	Configuring 802.1x on specific interface ports	113
	Configuring 802.1x timeouts on specific interface ports.	113
	Configuring 802.1x re-authentication on specific interface ports.	113
	Disabling 802.1x on specific interface ports.	114
Chapter 12	Configuring sFlow using the CEE CLI	
	In this chapter	115
	sFlow protocol overview	115
	Interface flow samples	115
	Packet counter samples	115
	Configuring the sFlow protocol globally.	116
	Interface-specific administrative tasks for sFlow	116
	Disabling sFlow on specific interfaces	116
	Configuring sFlow on specific interfaces	117

Chapter 13	Configuring Port Mirroring using the CEE CLI	
	In this chapter	119
	Port Mirroring protocol overview	119
	Port Mirroring limitations.	119
	Configuring ingress Port Mirroring	120
	Configuring egress Port Mirroring	120
	Configuring bidirectional Port Mirroring	120
	Deleting a Port Mirroring connection from a session.	121
	Deleting a Port Mirroring session.	121
Chapter 14	Configuring RMON using the CEE CLI	
	In this chapter	123
	RMON overview.	123
	RMON configuration and management.	123
	Default RMON configuration	123
	Configuring RMON settings.	123
	Configuring RMON events	124
	Configuring RMON group statistics collection	124
	Configuring RMON history collection	125
Chapter 15	Configuring IGMP	
	In this chapter	127
	About IGMP	127
	Active IGMP snooping	127
	Multicast routing	128
	Configuring IGMP	128
	Configuring IGMP snooping querier.	128
	Monitoring IGMP	129
Chapter 16	FCoE configuration using the Fabric OS CLI	
	In this chapter	131
	FCoE configuration guidelines and restrictions	131
	Managing and displaying the FCoE configuration.	132
	Enabling or disabling an FCoE port	132
	Configuring FCMAP values for a VLAN	132
	Configuring FIP multicast advertisement intervals	132
	Clearing logins	133
	Displaying FCoE configuration-related information	133

Managing and displaying the FCoE login configuration	133
Enabling or disabling FCoE login configuration management.	133
Displaying or aborting the current configuration transaction	134
Cleaning up login groups and VN_port mappings.	134
Displaying the FCoE login configuration.	135
Saving the current FCoE configuration.	135
Creating and managing the FCoE login group configuration	135
Creating an FCoE login group	135
Modifying the FCoE login group device list.	136
Deleting an FCoE login group	136
Renaming an FCoE login group.	137

Chapter 17

CEE configuration management

In this chapter	139
CEE configuration management guidelines and restrictions.	139
CEE configuration management tasks	139
Display the running configuration file.	140
Saving the running configuration file	140
Loading the startup configuration file	140
Erasing the startup configuration file.	140
Archiving the running configuration file	141
Restore an archived running configuration file	141
Archiving the startup configuration file	141
Restore an archived startup configuration file	141
Archive a startup configuration from Flash	141
Restore a startup configuration file from Flash.	142
CEE configuration management commands	142
Flash file management commands	142
Debugging and logging commands	143

Index

Figures

Figure 1	Multiple switch fabric configuration	3
Figure 2	CEE CLI command mode hierarchy	15
Figure 3	Ingress VLAN filtering	32
Figure 4	Configuring LAGs for a top-of-the-rack CEE switch—Example 1	67
Figure 5	Configuring LAGs for a top-of-the-rack CEE switch—Example 2	67
Figure 6	Queue depth	100
Figure 7	Strict priority schedule — two queues	104
Figure 8	WRR schedule — two queues	105
Figure 9	Strict priority and Weighted Round Robin scheduler	106

Tables

Table 1	FCoE terminology	1
Table 2	CEE RBAC permissions.	14
Table 3	CEE CLI command modes	16
Table 4	CEE CLI keyboard shortcuts.	17
Table 5	CEE CLI command output modifiers	19
Table 6	IP static route features and capabilities	27
Table 7	Default VLAN configuration	33
Table 8	STP versus RSTP state comparison.	45
Table 9	Default STP, RSTP, and MSTP configuration	50
Table 10	Default MSTP configuration.	50
Table 11	Default 10-Gigabit Ethernet CEE interface-specific configuration	50
Table 12	Default LACP configuration	69
Table 13	ETS priority grouping of IPC, LAN, and SAN traffic	78
Table 14	Default LLDP configuration	80
Table 15	Default MAC ACL configuration	88
Table 16	Default priority value of untrusted interfaces.	95
Table 17	IEEE 802.1Q default priority mapping.	95
Table 18	Default user priority for unicast traffic class mapping.	97
Table 19	Default user priority for multicast traffic class mapping	98
Table 20	Supported scheduling configurations.	105
Table 21	Multicast traffic class equivalence mapping	106
Table 22	Default CEE Priority Group Table configuration	108
Table 23	Default CEE priority table	108
Table 24	CEE configuration management commands	142
Table 25	CEE Flash memory file management commands.	143
Table 26	Debugging and logging commands.	143

About This Document

In this chapter

- [How this document is organized](#) xv
- [Supported hardware and software](#)..... xvi
- [Document conventions](#) xvi
- [Notice to the reader](#) xviii
- [Getting technical help](#) xviii

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 Dell FCoE hardware.
- [Chapter 2, “Using the CEE CLI,”](#) describes the Converged Enhanced Ethernet (CEE) CLI.
- [Chapter 3, “Initial FCoE and CEE Configuration,”](#) describes some basic switch configurations for command SAN and LAN environments.
- [Chapter 4, “Configuring IP static routes,”](#) describes how to configure IP static routes.
- [Chapter 5, “Configuring VLANs Using the CEE CLI,”](#) describes how to configure VLANs.
- [Chapter 6, “Configuring STP, RSTP, and MSTP using the CEE CLI,”](#) describes how to configure the Spanning Tree Protocol (STP), Rapid STP (RSTP), and Multiple STP (MSTP).
- [Chapter 7, “Configuring Link Aggregation using the CEE CLI,”](#) describes how to configure Link Aggregation and Link Aggregation Control Protocol (LACP).
- [Chapter 8, “Configuring LLDP using the CEE CLI,”](#) describes how to configure the Link Layer Discovery Protocol (LLDP) and the Data Center Bridging (DCB) Capability Exchange Protocol (DCBX).
- [Chapter 9, “Configuring ACLs using the CEE CLI,”](#) describes how to configure Access Control Lists (ACLs).
- [Chapter 10, “Configuring QoS using the CEE CLI,”](#) describes how to configure Quality of Service (QoS).
- [Chapter 11, “Configuring 802.1x Port Authentication,”](#) describes how to configure the 802.1x Port Authentication protocol.
- [Chapter 12, “Configuring sFlow using the CEE CLI,”](#) describes how to configure sFlow.

- [Chapter 13, “Configuring Port Mirroring using the CEE CLI,”](#) describes how to configure Port Mirroring.
- [Chapter 14, “Configuring RMON using the CEE CLI,”](#) describes how to configure remote monitoring (RMON).
- [Chapter 15, “Configuring IGMP,”](#) describes how to configure IGMP snooping on the Dell FCoE hardware.
- [Chapter 16, “FCoE configuration using the Fabric OS CLI,”](#) describes how to configure FCoE using the Fabric OS CLI.
- [Chapter 17, “CEE configuration management,”](#) describes how to perform the administrative tasks required by the Dell FCoE hardware.

Supported hardware and software

The following hardware platforms are supported in this release:

- Dell M8428-k

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

- Dell M8428-k
- Dell FCOE10-24 port blade
- Dell Converged 10GbE Switch Module for IBM BladeCenter
- Dell 2 port 10GbE Converged Network Adapter for IBM BladeCenter

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>code text</code>	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:

command	Commands are printed in bold.
--option, option	Command options are printed in bold.
-argument, arg	Arguments.
[]	Optional element.
<i>variable</i>	Variables are printed in italics. In the help pages, values are <u>underlined</u> or enclosed in angled brackets < >.
...	Repeat the previous element, for example “member[:member...]”
value	Fixed values following arguments are printed in plain font. For example, --show WWN
	Boolean. Elements are exclusive. Example: --show -mode egress ingress

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.

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
IBM	BladeCenter Advanced Management Module Protect Mode

Getting technical help

Dell is committed to ensuring that your investment in our products remains cost-effective. If you need assistance, or find errors in the manuals, contact Dell Technical Support.

Contacting Dell

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

NOTE

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

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

1. Visit <http://www.support.dell.com>.
2. Click your country or region at the bottom of the page. For a full listing of countries and regions, click **All**.
3. In the Support menu, click **All Support**.

Choose the method of contacting Dell that is convenient for you.

Introducing FCoE

In this chapter

- [FCoE terminology](#) 1
- [FCoE overview](#) 1
- [Layer 2 Ethernet overview](#) 3
- [FCoE Initialization Protocol](#) 8
- [FCoE queuing](#) 12

FCoE terminology

[Table 1](#) lists and describes the FCoE terminology used in this document.

TABLE 1 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

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, Dell FCoE hardware provides a unique platform that connects servers to both LAN and SAN environments.

Within this manual, any appearance of the term “Dell FCoE hardware” refers to any of Dell’s FCoE products.

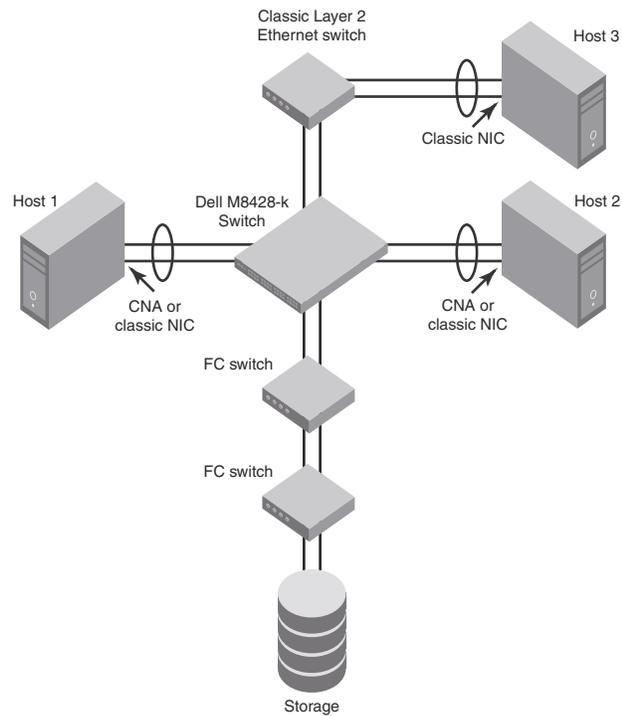
NOTE

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 Dell 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 Dell 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 Dell 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.

1 Layer 2 Ethernet overview

The Dell 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.
- 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 Dell 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 Using the CEE CLI”](#) on page 31.

Loop-free network environment

The Dell 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 using the CEE CLI”](#) on page 43.

Frame classification (incoming)

The Dell 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 Dell 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 Dell 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 using the CEE CLI”](#) on page 93.

Congestion control and queuing

The Dell 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 Dell 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).

Queuing features are described as follows:

- RED—RED increases link utilization. When multiple inbound TCP 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 using the CEE CLI”](#) on page 93.

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 using the CEE CLI”](#) on page 87.

Access Gateway

All ports on the switch come from the factory set to Access Gateway mode, with the default Access Gateway mapping. See the [“Access Gateway Administrator’s Guide”](#) for full details.

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 Dell FCoE hardware using the VLAN classifier commands. See [“Configuring an interface port as a trunk interface”](#) on page 37.

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 using the CEE CLI”](#) on page 75.

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 using the CEE CLI”](#) on page 65.

NOTE

The Dell 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 using the CEE CLI”](#) on page 93.

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 does not support FIP Keep Alive.

The Dell FCoE hardware FIP discovery phase operates as follows:

- The Dell FCoE hardware uses the FCoE Initialization Protocol (FIP). ENodes discover FCFs and initialize the FCoE connection through the FIP.
- Solicited advertisements—A typical scenario is where a Dell 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 Dell FCoE hardware that have the ENode in the login group.
- VLAN 1—The Dell 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 v6.3.1_cee 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 Dell FCoE hardware using FIP. Fabric login (FLOGI) and fabric discovery (FDISC) are accepted. Dell 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 Dell 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 Dell 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 Dell 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 Dell 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 Dell 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 Dell FCoE hardware accepts a maximum of 1024 logins per switch. Note that the Dell FCoE hardware does not reject further logins after the maximum is reached.

FIP logout

FIP logout operates as follows:

- ENodes can log out from the Dell FCoE hardware using FIP. The Dell FCoE hardware in the fabric updates the MAC address, WWN, and PID mappings upon logout. The Dell FCoE hardware also handles scenarios of implicit logout where the ENode has left the fabric without explicitly logging out.
- FIP logout (LOGO)—The Dell 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 Dell 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 Dell FCoE hardware implicitly logs out that ENode. ENode logout is updated in the VN_port table. The Dell FCoE hardware sends FCoE LOGO on behalf of the ENode.

FCoE login

The Dell FCoE hardware FCoE login operates as follows:

- ENodes can log in to the Dell FCoE hardware using FCoE encapsulated, FC Extended Link Service (ELS) frames. FLOGI and FDISC are accepted. Dell FCoE hardware in the fabric maintains the MAC address to WWN/PID mappings per login. Class 2 FLOGI is not supported.
- FCoE FLOGI—The Dell 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 Dell FCoE hardware. Requests are ignored if the FCMAP does not match. The ENode login is added to the VN_port table.
- FCoE FDISC—The Dell 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 Dell 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 Dell 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 v6.3.1_cee release.

FCoE logout

The Dell FCoE hardware FCoE logout operates as follows:

- ENodes can log out from the Dell FCoE hardware using the FCoE encapsulated, FC ELS frame. Dell FCoE hardware in the fabric updates the MAC address to WWN/PID mappings upon logout. The Dell FCoE hardware also handles scenarios of implicit logout where the ENode has left the fabric without explicitly logging out.
- FCoE LOGO—The Dell 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 Dell 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 Dell FCoE hardware logincfg mechanism operates as follows:

- The logincfg is the mechanism for controlling ENode logins per Dell FCoE hardware. Each unit of Dell 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 Dell FCoE hardware accepts valid loggingroup names and member WWNs. The Dell FCoE hardware rejects invalid entries. The Dell 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 Dell FCoE hardware does not apply the configurations to new sessions until the changes are committed. Once defined, loggingroups are committed. The Dell 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 Dell FCoE hardware name server function operates as follows:

- ENode login and logout to and from the Dell FCoE hardware updates the name server in the FC fabric. The Dell 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 Dell 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 Dell FCoE hardware FC zoning operates as follows:

- The virtual devices created by the Dell 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 v6.3.1_cee 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.

1 FCoE queuing

- 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 Dell 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 Dell 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 Dell 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 v6.3.1_cee release.

Using the CEE CLI

In this chapter

- [Management Tools](#) 13
- [CEE Command Line Interface](#)..... 13
- [Internal and external 10 Gbps Ethernet interfaces syntax](#)..... 20

Management Tools

The Dell FCoE hardware runs traditional Fabric OS (FOS) software and can be managed using the same tools traditionally used for SAN management. Using the FOS Command Line Interface (CLI), administrators have access to all commands and utilities common to other Dell switches. In addition, Fabric OS software on the Dell FCoE hardware enables Dell 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 Dell FCoE hardware introduces a new CLI 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 **fos ?** 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.

1. 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 command

Perform this task from Privileged EXEC mode.

1. 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 2](#) 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 2 CEE RBAC permissions

Root	Factory	UserID	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

NOTE

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

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.

```
switch login: userid
Password:
switch:admin> cmsh
switch#
```

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

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

NOTE

Multiple users can Telnet and issue commands using the Exec mode and the 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> cmsh
switch#
```

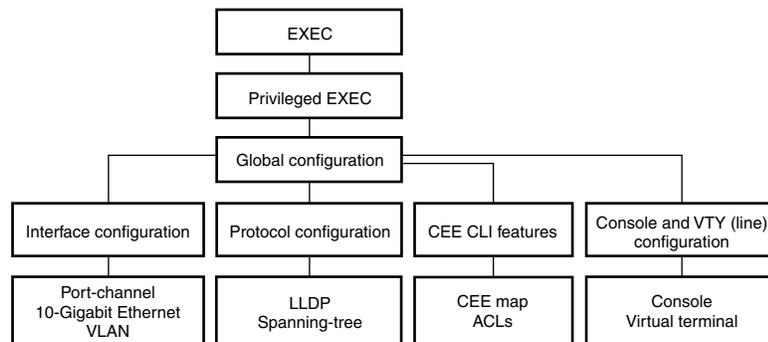
To return to the Fabric OS shell, enter the following command.

```
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 3](#) lists the CEE CLI command modes and describes how to access them.

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 3 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 interface types: <ul style="list-style-type: none"> • interface port-channel • interface intengigabitethernet • 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 protocol types: <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 3 CEE CLI command modes

Command mode	Prompt	How to access the command mode	Description
Feature configuration	CEE map: switch(config-ceemap)# Standard ACL: switch(conf-macl-std)# Extended ACL: switch(conf-macl-ext)#	From the global configuration mode, specify a CEE feature by entering one of the following feature names: <ul style="list-style-type: none"> • cee-map • mac access-list 	Access and configure CEE features.
Console and VTY (line) configuration	switch(config-line)#	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.

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 4 lists CEE CLI keyboard shortcuts.

TABLE 4 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 from all terminals.

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 the 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
switch(conf-lldp)#
```

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. This 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 5](#).

TABLE 5 CEE CLI command output modifiers

Output modifier	Description
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.
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.

Internal and external 10 Gbps Ethernet interfaces syntax

Fabric OS v6.3.1_cee uses two separate interface operands to refer to internal and external 10Gbps Ethernet ports; **intengigabitethernet** and **extengigabitethernet**.

Internal interfaces are assigned to ports 1 through 14. External interfaces are assigned to ports 15 through 22. The usual format for the 10 Gbps Ethernet interface is either **intengigabitethernet slot/port** or **extengigabitethernet slot/port**.

Initial FCoE and CEE Configuration

In this chapter

- Overview 21
- Configuring the FCoE interfaces 21
- Configuring the CEE interfaces 23
- Configuring DCBX 24
- Configuring Spanning Tree Protocol 25
- Configuring VLAN membership 25
- Configuring protect mode 26

Overview

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE) properties on the switch using CEE 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 Dell FCoE hardware. See “[CEE CLI command modes](#)” on page 15 for complete instructions on logging into the Dell FCoE hardware.

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. On assigning an FCoE map onto 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.

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

3 Configuring the FCoE interfaces

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 default CEE map (also called “default”).

The default startup-configuration file contains the following settings.

- `priority-group-table 1 weight 40 pfc`
- `priority-group-table 2 weight 60`
- `priority-table 2 2 2 1 2 2 2 2`

Configuring FCoE VLAN

By default, if there is no preexisting FCoE VLAN, one is created and assigned the VLAN ID of 1002. But if the system already contains an FCoE VLAN, then that VLAN is considered the default VLAN and is associated to the FCoE map. Only one FCoE VLAN can exist at a time.

If the system boots without any FCoE VLANs, but with VLAN 1002 assigned as a LAN VLAN, the default FCoE VLAN is automatically assigned a different VLAN ID, depending on availability.

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

NOTE

Make sure your converged mode interface is not configured to classify untagged packets to the same VLAN as the incoming VLAN-tagged packets. By configuring a converged interface to classify untagged packets (by using classifiers or the default port VLAN ID) to the same VLAN as VLAN-tagged packets coming into the interface, the FCoE hardware sends out untagged packets to the CNA. These packets may be dropped, disrupting communications.

The FCoE VLANs can be seen using `show vlan fcoe`. Alternatively, `show vlan brief` can also be used, in which FCoE VLANs are marked with an “(F)”.

NOTE

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

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

1. Create the FCoE map and enter FCoE map mode.

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

2. Configure the FCoE VLAN for the FCoE map and delete the previous FCoE VLAN.

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

The new FCoE VLAN 5 is created and the previous FCoE VLAN 1002 is deleted. This command succeeds only if VLAN 5 can be created and VLAN 1002 can be deleted.

3. Display the fcoe map to ensure the successful creation.

```
switch(conf-fcoe-map)#do show fcoe-map
fcoe-map default
  cee-map      :default
  fcoe-vlan    :6
  interface(s) :
```

Assigning FCoE map on to an interface

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

The FCoE map can be applied, irrespective of whether the interface is in 'switchport' or not. 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 to an interface, perform the following steps in global configuration mode.

1. Enter interface configuration mode.

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

2. Use the **fcoeport** command to assign the FCoE map to the interface

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

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

To enable the interface:

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

To disable the interface:

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

4. Display the fcoe map to ensure the successful creation.

```
switch(conf-fcoe-map)#do show fcoe-map
fcoe-map default
  cee-map      :default
  fcoe-vlan    :6
  interface(s) :te0/0
```

5. Repeat this procedure for any additional interfaces.

Configuring the CEE interfaces

Traffic from downstream CEE interfaces can be assigned to a VLAN using several methods:

- The VLAN tag contained in the incoming frame
- The VLAN classifiers
- The Port-VLAN ID (PVID)

Because the Ethernet uplink ports from the Dell FCoE hardware to the distribution layer switches will carry traffic for multiple VLANs, they are configured as 802.1q trunk ports.

The downstream CEE ports connected to the server CNAs are configured as access ports with a PVID of either 10 or 20. The VLAN classifier group created for the FIP and FCoE EtherTypes must be applied to the interfaces in order to place FCoE traffic on the correct VLAN. The CEE map is also applied to the interface.

If an interface is already assigned as an fcoeport, you cannot apply the CEE map to that interface. Only a single CEE map is allowed, which is created automatically with the name "default." You are not be able to delete or rename this map.

3 Configuring DCBX

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

1. Assign VLANs to the uplink Ethernet port.

NOTE

You must repeat this step for all uplink interfaces. For details, see [“Configuring an interface port as a trunk interface”](#) on page 37.

The following example assigns VLAN 10 and VLAN 20 to the uplink Ethernet port.

```
switch(config)#interface intengigabitethernet 0/1
switch(conf-if-te-0/1)#switchport
switch(conf-if-te-0/1)#switchport mode trunk
switch(conf-if-te-0/1)#switchport trunk allowed vlan add 10
switch(conf-if-te-0/1)#switchport trunk allowed vlan add 20
switch(conf-if-te-0/1)#no shutdown
```

2. Apply the VLAN classifier group to the interfaces. For details, see [“Activating a VLAN classifier group with an interface port”](#) on page 40.

The following example applies a VLAN classifier group 1 to the interfaces.

```
switch(config)#interface intengigabitethernet 0/10
switch(conf-if-te-0/1)#switchport
switch(conf-if-te-0/1)#switchport mode access
switch(conf-if-te-0/1)#switchport access vlan 10
switch(conf-if-te-0/1)#vlan classifier activate group 1 vlan 100
switch(conf-if-te-0/1)#no shutdown
```

3. Apply the CEE map to the interfaces. For details, see [“Applying a CEE provisioning map to an interface”](#) on page 110.

The following example applies the map name.

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

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

```
switch(conf-if-te-0/1)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring DCBX

DCBX (Data Center Bridging eXchange Protocol) 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 ETS and PFC—to be sent before the link is shared. DCBX parameters use a type-length-value (TLV) format. By default, DCBX is turned on, but there are two TLVs that must be enabled to support FCoE on a CEE link:

- 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 Protocol 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)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring Spanning Tree Protocol

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

It is 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(conf-rstp)#bridge-priority 61440
```

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

```
switch(conf-rstp)#exit
switch(config)#end
switch#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 Initiation Protocol (FIP) and FCoE. VLAN classifiers are rules used to dynamically classify Ethernet frames on an untagged interface to VLANs.

3 Configuring protect mode

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

1. Create the VLAN interfaces on the Dell FCoE hardware using the CEE CLI. For details, see [“Creating a VLAN interface”](#) on page 35.

The following example creates two VLAN interfaces and assigning 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 FCoE VLAN
```

2. Create VLAN rules and a VLAN classifier group for these two EtherTypes. For details, see [“Creating a VLAN classifier group and adding rules”](#) on page 39.

The following example creates VLAN rules and classifier groups.

```
switch(config)#vlan classifier rule 1 proto fip encap ethv2
switch(config)#vlan classifier rule 2 proto fcoe encap ethv2
switch(config)#vlan classifier group 1 add rule 1
switch(config)#vlan classifier group 1 add rule 2
```

3. Apply the VLAN classifier group to any CEE interface. This step is optional. For details, see [“Activating a VLAN classifier group with an interface port”](#) on page 40.
4. Enter the **copy** command to save the *running-config* file to the *startup-config* file.

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

Configuring protect mode

IBM’s Advanced Management Module (AMM) can set or control certain attributes of the switch through a backend I2C bus. The switch receives these AMM settings by reading the I2C control registers and VPD. Additionally, AMM informs the switch if there are any changes in these settings while the system is running by setting the appropriate bits in the I2C control registers and generating interrupts.

Enabling protect mode is a two-step process. First, enable protect mode on AMM. Refer to the IBM publication titled *IBM BladeCenter Advanced Management Module Protect Mode*. The second task is to enable Protect Mode on the switch by performing the following steps in global configuration mode.

1. Select an external interface to configure as an IP interface.

```
switch(config)#interface extengigabitethernet 0/16
```

2. Configure the interface as an IP interface

```
switch(conf-if-te-0/16)#ip address 10.0.0.1 5
```

3. Return to Privileged EXEC mode.

```
switch(conf-if-te-0/16)#exit
switch(config)#exit
```

4. Enable protect mode.

```
switch#protect-mode enable
```

Configuring IP static routes

In this chapter

- [IP static routes overview](#) 27
- [Configuring IP static routes](#)..... 28

IP static routes overview

Switches forward packets using routing data that is either configured manually or dynamically created using a routing protocol. Static routes define an explicit path between two networking devices. Unlike a dynamic routing protocol, static routes are not automatically updated and are manually reconfigured.

Routes generated by dynamic routing protocols cannot be redistributed into the static routing tables, but static routes can be redistributed into dynamic routing protocols. You cannot prevent routing loops from using static routes.

Static routes are useful for smaller networks with only one path to an outside network and to provide security for a larger network for certain types of traffic or links to other networks that need more control. In general, most networks use dynamic routing protocols to communicate between networking devices but may have one or two static routes configured for special cases.

Fabric OS v6.3.1_cee supports IPv4 static routes with Layer 2 and Layer 3 data forwarding for eAnvil-based platforms. The following features of IP static routes are supported:

- Static routes
- Directly attached static routes
- Recursive static routes
- Floating static routes

NOTE

Fully specified static routes are not supported by Fabric OS v6.3.1_cee.

The following table explains the maximum values supported for the following features

TABLE 6 IP static route features and capabilities

Feature	Capability
Interface with IP Address	128
ECMP	63
Members per ECMP	64
LAG	63
Members per LAG	64

TABLE 6 IP static route features and capabilities

Feature	Capability
Static Routes	Less than 2048, which are shared between IP address' and static routes)
ARPs	1024 (due to Linux IP stack limitation) eAnvil Chip capability 12288 Anvil Chip capability 16384

Configuring IP static routes

In directly attached static routes, only the output interface is specified. The destination is assumed to be directly attached to this interface, so the packet destination is used as the next-hop address.

Directly attached static routes are candidates for insertion in the IPv4 routing table only if they refer to a valid IPv4 interface; that is, an interface that is both up and has IPv4 enabled on it. This configuration is not recommended because when the next hop of a static route points to an interface, the router considers each of the hosts within the range of the route to be directly connected through that interface. With this type of configuration, a router performs Address Resolution Protocol (ARP) on the Ethernet for every destination the router finds through the default route because the router considers all of these destinations as directly connected to internal 10 Gbps interface O/1. This kind of default route, especially if it is used by a lot of packets to many different destination subnets, can cause high processor utilization and a large ARP cache.

Specifying a numerical next hop on a directly connected interface prevents the router from performing ARP on each destination address. However, if the interface with the next hop goes down and the numerical next hop is reachable through a recursive route, you should specify both the next hop IP address and the interface through which the next hop should be found.

NOTE

By default, static routes are preferred to routes learned by routing protocols. Therefore, you can configure an administrative distance with a static route if you want the static route to be overridden by dynamic routes. For example, you could have routes installed by the Open Shortest Path First (OSPF) protocol with an administrative distance of 120. To have a static route that would be overridden by an OSPF dynamic route, specify an administrative distance greater than 120.

The following example shows all destinations with address prefix 1.1.1.1/32 are directly reachable through internal 10 Gbps interface O/1, with an administrative distance of 1.

```
switch(config)#ip route 1.1.1.1/32 intengigabitethernet 0/1 1
```

NOTE

Fabric OS v6.3.1_cee allows 64 ECMP routes. Once this pool is exhausted, any additional ECMP routes are automatically rejected. If one of the existing ECMP routes is deleted, previously rejected ECMP routes must be re-added manually.

Recursive IP static routes

In a recursive IP static route, only the next hop is specified. The output interface is derived from the next hop. A recursive IP static route is valid (that is, it is a candidate for insertion in the routing table) only when the specified next hop resolves, either directly or indirectly, to a valid output interface, provided the route does not self-recurse, and the recursion depth does not exceed the maximum IPv4 forwarding recursion depth. Fabric OS v6.3.1_cee supports one level of recursion. Beyond this level of recursion the routes are held as inactive routes.

Dell does not recommend manually configuring a self-recursive IP static route, although it is not prohibited. However, a recursive IP static route that has been inserted in the routing table may become self-recursive as a result of some transient change in the network learned through a dynamic routing protocol. If this occurs, the fact that the IP static route has become self-recursive is detected and is removed from the routing table, although not from the configuration. A subsequent network change may cause the IP static route to no longer be self-recursive, in which case it is reinserted in the routing table.

The following example specifies that all destinations with address prefix 1.1.1.1/32 are reachable via the host with address 2.2.2.2.

```
switch(config)#ip route 1.1.1.1/32 2.2.2.2
```

Floating IP static routes

Floating IP static routes act as an alternative path to dynamic routes learned through configured routing protocols. A floating IP static route is configured with a higher administrative distance than the dynamic routing protocol. As a result, the dynamic route learned through the routing protocol is always used in preference to the floating IP static route. If the dynamic route learned through the routing protocol is lost, the floating IP static route is used in its place.

A floating IP static route must be configured with an administrative distance that is greater than the administrative distance of the dynamic routing protocol, because routes with smaller administrative distances are preferred. For example, if your Open Shortest Path First (OSPF) protocol is configured with an administrative distance of 120, specify an administrative distance greater than 120 for your floating IP static route.

The following example defines a floating IP static route that shows all destinations with address prefix 1.1.1.1/32 are reachable through internal 10 Gbps interface 0/1:

```
switch(config)#ip route 1.1.1.1/32 intengigabitethernet 0/1 121
```

Displaying and clearing IP static routes

To display the current IP routing information and clear the assigned settings, perform the following procedure from EXEC mode.

1. Display a summary of the IP routing table.

```
switch>show ip route summary
Route Source      Networks  Subnets  Overhead  Memory (bytes)
connected         0         1         64        160
static            1         1        128        320
internal          1         0         0         1180
Total             2         2        192        1660
```

4 Configuring IP static routes

2. Display the detailed information about all the IP static routes.

```
switch>show ip route all
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.32.144.1 to network 0.0.0.0
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
 S       10.0.0.0/8 [1/0] via 10.32.144.1
 C       10.32.144.0/20 is directly connected, GigabitEthernet4/47
 S*     0.0.0.0/0 [1/0] via 10.32.144.1
```

3. Use the **clear ip route** command to remove a static ip route.

```
switch>clear ip route 10.0.0.0/8
```

4. *Optional:* Use the **clear ip route** command to remove all static ip routes.

```
switch>clear ip route all
```

Configuring VLANs Using the CEE CLI

In this chapter

- [VLAN overview](#) 31
- [Ingress VLAN filtering](#) 31
- [VLAN configuration guidelines and restrictions](#) 33
- [Default VLAN configuration](#) 33
- [VLAN configuration and management](#) 34
- [Configuring protocol-based VLAN classifier rules](#) 38
- [Configuring the MAC address table](#) 40

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 36).

NOTE

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

Ingress VLAN filtering

A frame arriving at Dell FCoE hardware is either associated with a specific port or with a VLAN, depending 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.

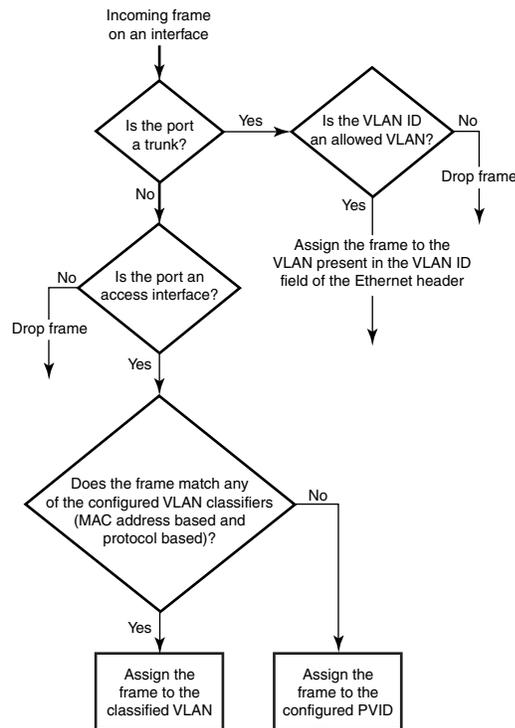
- Admit VLAN tagged and untagged frames—All tagged and untagged frames would be processed as follows:
 - All untagged frames are classified into native VLANs.
 - All frames egressing are untagged for the native VLANs.
 - Any tagged frames coming with a VLAN tag equal to the configured native VLAN are 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 3 displays the frame processing logic for an incoming frame.

FIGURE 3 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 Dell 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 Dell FCoE hardware. The aging time is between 10 and 1000000 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 overwrites an existing dynamically learned FDB entry and disables learning of the entry going forward.

NOTE

For more information on frame handling for Dell 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:

- Make sure your converged mode interface is not configured to classify untagged packets to the same VLAN as the incoming VLAN-tagged packets. By configuring a converged interface to classify untagged packets (by using classifiers or the default port VLAN ID) to the same VLAN as VLAN-tagged packets coming into the interface, the FCoE hardware sends out untagged packets to the CNA. These packets may be dropped, disrupting communications.
- 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 Dell CEE switch supports Ethernet DIX frames and 802.2 LLC SNAP encapsulated frames only.

Default VLAN configuration

[Table 7](#) lists the default VLAN configuration.

TABLE 7 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
Port assignment	VLAN 4095

VLAN configuration and management

NOTE

To see the minimum configuration required to enable FCoE on Dell FCoE hardware, refer to [Chapter 3, “Initial FCoE and CEE Configuration”](#).

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

NOTE

CEE interfaces are disabled by default.

NOTE

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

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 intengigabitethernet 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 intengigabitethernet 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 following example sets the MTU value to 4200.

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

Creating a VLAN interface

On Dell 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 1002
```

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 1002 and VLAN 55 simultaneously. In addition, VLAN 1002 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 1002
```

4. Enter the **spanning-tree shutdown** command to enable spanning tree on VLAN 1002.

```
switch(conf-if-vl-1002)#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 1002.

```
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, 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 1002
```

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 intengigabitethernet 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 intengigabitethernet 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 intengigabitethernet 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 intengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **switchport** command to configure the CEE interface as a Layer 2 switch port.

```
switch(conf-if-te-0/1)#switchport access vlan 20
```

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 intengigabitethernet 0/19
```

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

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

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

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

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

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

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

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. Any tagged frames coming with a VLAN tag equal to the configured native VLAN are dropped.

5 Configuring protocol-based VLAN classifier rules

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 intengigabitethernet 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
```

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 intengigabitethernet 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
```

Configuring protocol-based 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 39).

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:

- Ethernet hexadecimal (0x0000 through 0xffff)
- Address Resolution Protocol (ARP)
- Fibre Channel over Ethernet (FCoE)

- FCoE Initialization Protocol (FIP)
- IP version 6 (IPv6)

NOTE

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 fcoe 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)#vlan classifier group 1 delete rule 1
```

Creating a VLAN classifier group and adding rules

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

NOTE

A vlan classifier cannot be added to the FCoE VLAN.

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.

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 intengigabitethernet 0/10
```

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(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 intengigabitethernet 0/10 port-channel 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 brief
```

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. Dell 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  
intengigabitethernet 0/1 vlan 100
```

5 Configuring the MAC address table

Configuring STP, RSTP, and MSTP using the CEE CLI

In this chapter

- STP overview 43
- RSTP overview 45
- MSTP overview 47
- STP, RSTP, and MSTP configuration guidelines and restrictions 49
- Default STP, RSTP, and MSTP configuration 50
- STP, RSTP, and MSTP configuration and management 51
- Configuring STP, RSTP, or MSTP on CEE interface ports 58

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

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 59.
- PortFast BPDU guard and BPDU filter—For detailed information, see [“Enabling port fast \(STP\)”](#) on page 61.

Configuring STP on Dell FCoE hardware

The process for configuring STP on your Dell FCoE hardware is as follows.

1. Enter Global Configuration mode.
2. Enable RSTP using the global **protocol spanning-tree** command. For details, see [“Enabling STP, RSTP, or MSTP”](#) on page 51.

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

3. Designate the root switch using the **bridge-priority** command. For details, see [“Specifying the bridge priority”](#) on page 52. The range is 0 through 61440 and the priority values can be set only in increments of 4096.

```
switch(conf-stp)#bridge-priority 28582
```

4. Enable PortFast on switch ports using the **spanning-tree portfast** command. For details, see [“Enabling port fast \(STP\)”](#) on page 61. 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 intengigabitethernet 0/10
switch(conf-if-te-0/10)#spanning-tree portfast
switch(conf-if-te-0/10)#exit
switch(config)#interface intengigabitethernet 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.

5. Set the following ports to forwarding mode:
 - All ports of the root switch
 - The root port
 - The designated port

6. 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 59. Note that this step is optional.

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.

7. 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. Data received on ports that are not part of the spanning tree is blocked.

NOTE

Dell recommends leaving other STP variables at their default values.

For more information on STP, see [“STP, RSTP, and MSTP configuration and management”](#) on page 51.

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 8](#) lists the interface state changes between STP and RSTP.

TABLE 8 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 51.

Configuring RSTP on Dell FCoE hardware

The basic process for configuring RSTP on your Dell FCoE hardware is as follows.

1. Enter Global Configuration mode.
2. Enable RSTP using the global **protocol spanning-tree** command. For details, see [“Enabling STP, RSTP, or MSTP”](#) on page 51.

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

3. Designate the root switch using the **bridge-priority** command. For details, see [“Specifying the bridge priority”](#) on page 52. The range is 0 through 61440 and the priority values can be set only in increments of 4096.

```
switch(conf-stp)#bridge-priority 28582
```

4. Configure the **bridge forward delay** value. For details, see [“Specifying the bridge forward delay”](#) on page 52.

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

5. Configure the **bridge maximum aging time** value. For details, see [“Specifying the bridge maximum aging time”](#) on page 53.

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

6. Enable the **error disable timeout timer** value. For details, see [“Enabling the error disable timeout timer”](#) on page 53.

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

7. Configure the **error-disable-timeout** interval value. For details, see [“Specifying the error disable timeout interval”](#) on page 53.

8.

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

9. Configure the port-channel path cost. For details, see [“Specifying the port-channel path cost”](#) on page 54.

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

10. Configure the bridge hello time value. For details, see [“Specifying the bridge hello time \(STP and RSTP\)”](#) on page 54.

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

11. Flush the MAC addresses from the VLAN FDB. For details, see [“Flushing MAC addresses \(RSTP and MSTP\)”](#) on page 57.

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

12. Enable PortFast on switch ports using the **spanning-tree portfast** command. For details, see [“Enabling port fast \(STP\)”](#) on page 61. 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 intengigabitethernet 0/10
switch(config-if-te-0/10)#spanning-tree portfast
switch(config-if-te-0/10)#exit
switch(config)#interface intengigabitethernet 0/11
switch(config-if-te-0/11)#spanning-tree portfast
switch(config-if-te-0/11)#exit
```

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

13. Set the following ports to forwarding mode:

- All ports of the root switch
- The root port
- The designated port

For details, see [“Specifying the port priority”](#) on page 61.

14. 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 59. Note that this step is optional.

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.

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

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

Dell supports 16 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 on Dell FCoE hardware

The basic process for configuring MSTP on your Dell FCoE hardware is as follows.

1. Enter Global Configuration mode.
2. Enable MSTP using the global **protocol spanning-tree** command. For more details see [“Enabling STP, RSTP, or MSTP”](#) on page 51.

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

3. Specify the region name using the **region** *region_name* command. For more details see [“Specifying a name for an MSTP region”](#) on page 56.

```
switch(conf-mstp)#region dell1
```

4. Specify the revision number using the **revision** command. For more details see [“Specifying a revision number for an MSTP configuration”](#) on page 56.

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

5. Map a VLAN to an MSTP instance using the **instance** command. For more details see [“Mapping a VLAN to an MSTP instance”](#) on page 55.

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

6. 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 56.

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

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

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

For more information on MSTP, see [“STP, RSTP, and MSTP configuration and management”](#) on page 51.

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 16 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 9 lists the default STP, RSTP, and MSTP configuration.

TABLE 9 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 10 lists the switch defaults that apply only to MSTP configurations.

TABLE 10 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 11 lists the switch defaults for the 10-Gigabit Ethernet CEE interface-specific configuration.

TABLE 11 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

To see the minimum configuration required to enable FCoE on the Dell M8428-k switch, refer to [Chapter 3, “Initial FCoE and CEE Configuration”](#).

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 STP, RSTP, or MSTP, 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.

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 on a per-VLAN basis. 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.

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

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(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 enable the timer to bring a port out of the disabled state. 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 53.

To enable the error disable timeout timer, perform the following steps from Privileged EXEC mode. By default, the timeout feature is disabled.

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.

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(config)#transmit-holdcount 5
```

Enabling Cisco interoperability (MSTP)

In MSTP mode, use this command to enable or disable the ability of the Dell 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 Dell 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 Dell FCoE hardware to interoperate with certain legacy Cisco switches.

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

Disabling Cisco interoperability (MSTP)

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. Disable the ability of Dell 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.

```
switch(conf-mstp)#instance 5 vlan 4096
```

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 33).

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.
- Dell 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 Dell 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 Dell 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 intengigabitethernet 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.

6 Configuring STP, RSTP, or MSTP on CEE interface ports

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  
intengigabitethernet 0/1
```

Displaying STP, RSTP, and MSTP-related information

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

1. 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 Dell 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.
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 intengigabitethernet 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
```

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 61).

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 intengigabitethernet 0/1
```

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

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

6. Enter the **spanning tree** command again to assign a BPDU filter or BPDU guard.

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

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 intengigabitethernet 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 53). 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 intengigabitethernet 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 intengigabitethernet 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 5 cost 3550 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 intengigabitethernet 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 59).

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 intengigabitethernet 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.

6 Configuring STP, RSTP, or MSTP on CEE interface ports

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 intengigabitethernet 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 intengigabitethernet 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 intengigabitethernet 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

From the CEE interface, use this command to enable spanning tree on the CEE interface. By default, spanning tree is disabled.

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 intengigabitethernet 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

From the CEE interface, use this command to disable spanning tree on the CEE interface. By default, spanning tree is disabled.

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 intengigabitethernet 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)#spanning-tree shutdown
```

6 Configuring STP, RSTP, or MSTP on CEE interface ports

Configuring Link Aggregation using the CEE CLI

In this chapter

- [Link aggregation overview](#) 65
- [LACP configuration guidelines and restrictions](#) 69
- [Default LACP configuration](#) 69
- [LACP configuration and management](#) 69
- [LACP troubleshooting tips](#) 72

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 Dell FCoE hardware supports the following trunk types:

- Static, standards-based LAG.
- Dynamic, standards-based LAG using LACP.
- Static, Dell-proprietary LAG.
- Dynamic, Dell-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 Dell-proprietary LAG. Each LAG is associated with an aggregator. The aggregator manages the Ethernet frame collection and distribution functions.

7 Link aggregation overview

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 Dell 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 4](#) and [Figure 5](#) show typical IP SAN configurations using LAGs. In a data center the Dell M8428-k switch fits into the top-of-the-rack use case where all the servers in a rack are connected to the Dell M8428-k switch through Twinax copper or optical fiber cable. The database server layer connects to the top-of-the-rack Dell M8428-k switch which is located in the network access layer.

The Dell M8428-k 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 Dell M8428-k switch interoperates with all of the major Layer 2/Layer 3 aggregation routers including Foundry Networks, Cisco Systems, Brocade Communications Systems and Force10 Networks.

FIGURE 4 Configuring LAGs for a top-of-the-rack CEE switch—Example 1

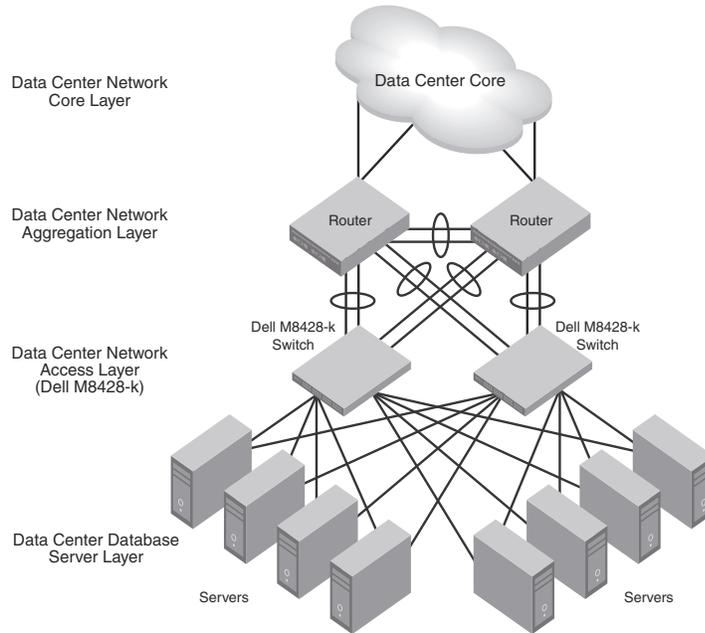
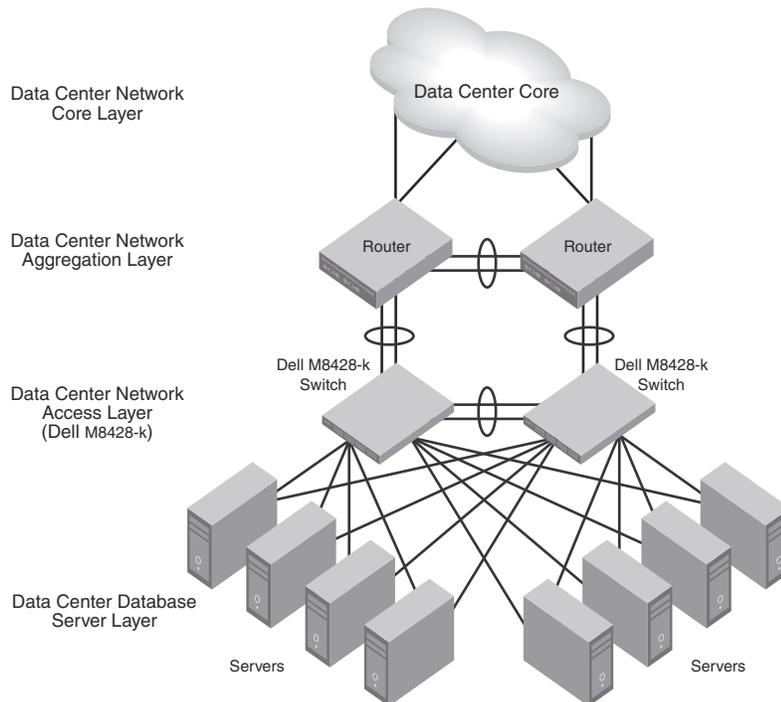


FIGURE 5 Configuring LAGs for a top-of-the-rack CEE switch—Example 2



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.

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 LACPDU to monitor the health of each member link.

Static link aggregation

In static link aggregation, links are added into a LAG without exchanging LACPDU 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.

Dell-proprietary aggregation

Dell-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 Dell 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 Dell-proprietary LAG configurations except where specifically noted otherwise.

Follow these LACP configuration guidelines and restrictions when configuring LACP:

- All ports on the Dell FCoE hardware can operate only in full-duplex mode.
- QoS—In the Fabric OS version 6.3.1_cee 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.
- Dell-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 12](#) lists the default LACP configuration.

TABLE 12 Default LACP configuration

Parameter	Default setting
System priority	32768
Port priority	32768
Timeout	Long (standard LAG) and short (Dell 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.

NOTE

To see the minimum configuration required to enable FCoE on the Dell M8428-k switch, refer to [Chapter 3, “Initial FCoE and CEE Configuration”](#).

Enabling LACP on a 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 intengigabitethernet 0/1
```

3. Enter the **no shutdown** command to enable the CEE interface.
4. Enter the **channel-group** command to configure the LACP for the CEE interface.

```
switch(conf-if)#channel-group 4 mode active type Dell
```

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 intengigabitethernet 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
```

Configuring minimum links feature

The minimum links feature allows a port-channel to operate at a certain minimum bandwidth all the time. If the bandwidth of the port-channel drops below that minimum number, then the port-channel is declared operationally DOWN even though it has operationally UP members.

For example, if a port-channel has five members, each with 10 Gigabit bandwidth and you want the port-channel to have a minimum bandwidth of 30 Gigabit, then the port-channel is declared operationally UP only when there are at least three members operationally UP in it. As soon as the operational members falls below three, the port-channel is declared operationally DOWN.

To configure this feature, perform the following steps under global configuration mode.

1. Select the interface.

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

2. Enter the port-channel sub-mode.

```
switch(conf-if-te-0/1)#interface port-channel 63
```

3. Configure the minimum links.

```
switch(conf-if-po-63)#minimum-links 3
```

Configuring interface tracking

The **track interface** command lets you track one or more external ports on your switch. Only external physical interfaces and port-channel interfaces can be tracked. The show output for internal interfaces is enhanced to include the information of which external interfaces are being tracked. Additionally, there is a command provided to enable (or disable) the tracking under internal interface sub-mode. This command is **track enable**. External interfaces can be tracked only if tracking is enabled by means of this command.

The semantics of **track interface** are as follows:

- One internal interface can track one or more external interfaces.
- Multiple internal interfaces can track the same external interface.
- If multiple external interfaces are being tracked by one internal interface, only if all of them go DOWN, the internal interface should go DOWN.
- Only the operational status of the external interface is tracked – not the admin state.
- If multiple internal interfaces are tracking a single external interface and if the external interface goes DOWN, all those internal interfaces should go DOWN.
- A maximum of eight external interfaces can be tracked from an internal interface.
- Forward referenced port-channels (non-existing port-channels) are tracked.

In the following procedure, tracking is enabled and the internal interface 0/8 is tracking external interfaces 0/15,0/16, and 0/17. Perform the following steps from global configuration mode.

1. Select the internal interface to use as the tracking monitor.

```
switch(config)#interface intengigabitethernet 0/8
```

2. Enable tracking.

```
switch(conf-if-int-0/8)#track enable
```

3. Select one or more external interfaces to be tracked.

```
switch(conf-if-int-0/8)#track interface extengigabitethernet 0/15
switch(conf-if-int-0/8)#track interface extengigabitethernet 0/16
switch(conf-if-int-0/8)#track interface extengigabitethernet 0/17
```

Clearing LACP counter statistics on a LAG

To clear LACP counter statistics, perform the following task from Privileged EXEC mode.

1. 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, enter the **clear** command to clear the LACP counter statistics for all LAG groups.

Example

```
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 Dell-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 **Dell** 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 Dell-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 **Dell** 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.

7 LACP troubleshooting tips

Configuring LLDP using the CEE CLI

In this chapter

- [LLDP overview](#) 75
- [Layer 2 topology mapping](#)..... 76
- [DCBX overview](#)..... 78
- [DCBX interaction with other vendor devices](#) 79
- [LLDP configuration guidelines and restrictions](#)..... 79
- [Default LLDP configuration](#)..... 80
- [LLDP configuration and management](#)..... 80

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

Dell's LLDP implementation supports a one-to-one connection. Each interface has one and only one neighbor.

The higher level management tools, such as Dell'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. Dell's LLDP implementation adds a proprietary Dell 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.
 - Port description 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.
- Dell 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 Dell 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 13](#) 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 13](#) are translated to priority groups in the Dell FCoE hardware.

TABLE 13 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 Dell FCoE hardware interacts with other vendor devices, the other vendor devices might not have support for the same DCBX version as the Dell FCoE hardware.

The Dell 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 Dell FCoE hardware provides the following options.

- Auto-sense (plug and play)

This is the default. The Dell 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:

- Dell's implementation of LLDP supports Dell-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 80.

Default LLDP configuration

Table 14 lists the default LLDP configuration.

TABLE 14 Default LLDP configuration

Parameter	Default setting
LLDP global state	Enabled
LLDP receive	Enabled
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
```

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 Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Disable LLDP globally.

```
switch(config)#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 Dell 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.

To specify a global system name for the Dell 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 Dell_Alpha  
Dell_Alpha(conf-lldp)#
```

Specifying an LLDP system description for the Dell FCoE hardware

NOTE

Dell 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 Dell 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 Dell 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 Dell-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)#mode 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 [port-description | system-name |
system-capabilities | system-description | management-address]
```

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 0x08.

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.

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
```

Configuring LLDP profiles

You can configure up to 64 profiles on a switch. 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.

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

2. Enter LLDP configuration mode.

3. Configure the profile name.

The following example creates the unique profile name of “UK_LLDP_IT”.

```
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)#mode tx rx
```

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.

```
switch(conf-lldp-profile-UK_LLDP_IT)#advertise optional-tlv
```

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
```

NOTE

Dell recommends against advertising dot1.tlv and dot3.tlv LLDPs if your network contains CNAs from non-Dell 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)#exit
```

```
switch(conf-lldp)#exit
```

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

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 intengigabitethernet 0/10
```

2. Apply an LLDP profile to the interface.

The following example applies the LLDP profile “network_standard” to the current 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 0x0-0xff
```

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 79. The default is to automatically detect the DCBX version.

```
switch(conf-if-te-0/10)#lldp version cee
```

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

```
switch(conf-if-te-0/10)#exit
switch(config)#end
switch#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 intengigabitethernet 0/1
```

- Clear LLDP statistics.

```
switch#clear lldp statistics intengigabitethernet 0/1
```

Displaying LLDP-related information

To display LLDP-related information, perform the following steps from Privileged EXEC mode.

1. Use the **show lldp neighbors** command to:

- Display LLDP general information.

```
switch#show lldp
```

- Display LLDP interface-related information.

```
switch#show lldp interface intengigabitethernet 0/1
```

- Display LLDP neighbor-related information.

```
switch#show lldp neighbors interface intengigabitethernet 0/1 detail
```

8 LLDP configuration and management

Configuring ACLs using the CEE CLI

In this chapter

- [ACL overview](#) 87
- [Default ACL configuration](#) 88
- [ACL configuration guidelines and restrictions](#) 88
- [ACL configuration and management](#) 88

ACL overview

NOTE

In the Dell Fabric OS release, only Layer 2 MAC access control lists (ACLs) are supported.

ACLs filter traffic for the Dell 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.

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 15 lists the default ACL configuration.

TABLE 15 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.

NOTE

To see the minimum configuration required to enable FCoE on the Dell M8428-k switch, refer to [Chapter 3, “Initial FCoE and CEE Configuration”](#).

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 90.

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
switch(conf-macl-std)#
```

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 90.

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.

The following example sets the name of the extended MAC ACL to “test_02.”

```
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)#exit
switch(config)#end
switch#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 90.

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.”

The following example deletes the extended MAC ACL named “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.

1. 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 intengigabitethernet 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
```

9 ACL configuration and management

Configuring QoS using the CEE CLI

In this chapter

- QoS overview 93
- Rewriting 94
- Queueing 94
- Congestion control..... 100
- Multicast rate limiting 103
- Scheduling 104
- Converged Enhanced Ethernet map configuration 107

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:

- Rewriting—Rewriting or marking a frame allows for overriding header fields such as the priority and VLAN ID.
- 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 16 presents the Layer 2 QoS *untrusted* user priority generation table.

TABLE 16 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 17 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 17 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.

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 intengigabitethernet 0/2`
3. Set the interface mode to 'trust'.
`switch(conf-if-te-0/2)#qos trust cos`
4. Exit the configuration mode and return to EXEC mode.
`switch(conf-if-te-0/2)#exit`
`switch(config)#end`
5. Enter the copy command to save the running-config file to the startup-config file.
`switch#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 intengigabitethernet 0/2`
3. Set the interface mode to '3'.
`switch(conf-if-te-0/2)#qos cos 3`
4. Exit the configuration mode and return to EXEC mode.
`switch(conf-if-te-0/2)#exit`
`switch(config)#end`
5. Enter the copy command to save the running-config file to the startup-config file.
`switch#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. Exit the configuration mode and return to EXEC mode.
`switch(conf-if-te-0/2)#exit`
`switch(config)#end`
4. Enter the copy command to save the running-config file to the startup-config file.
`switch#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 intengigabitethernet 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. Exit the configuration mode and return to EXEC mode.

```
switch(conf-if-te-0/2)#exit
switch(config)#end
```

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

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

Traffic class mapping

The Dell M8428-k 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 18 presents the Layer 2 default traffic class mapping supported for a COS-based user priority to conform to 802.1Q default mapping.

TABLE 18 Default user priority for unicast traffic class mapping

User priority	Traffic class
0	1
1	0
2	2

TABLE 18 Default user priority for unicast traffic class mapping (Continued)

User priority	Traffic class
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 v6.3.1_cee supports four multicast 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.

[Table 19](#) presents the Layer 2 default traffic class mapping supported for a CoS-based user priority to conform to 802.1Q default mapping.

TABLE 19 Default user priority for multicast traffic class mapping

User Priority	Traffic class
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

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
```

The following example creates CoS-to-Traffic-Class QoS map to map CoS 0 (best effort) to Traffic Class 1 and CoS 1 to below best effort Traffic Class 0, all other CoS go through unchanged. This mapping matches the default behavior recommended in IEEE 802.1Q for systems supporting 8 Traffic Classes.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#qos map cos-traffic-class test 1 0 2 3 4 5 6 7
switch(config)#end
switch#
```

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

```
switch#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 intengigabitethernet 0/2
```

3. Activate the CoS-to-Traffic-Class mapping by name.

```
switch(conf-if-te-0/2)#qos cos-traffic-class test
```

The following example activates the CoS-to-Traffic-Class QoS map on an interface.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#interface intengigabitethernet 0/2
switch(conf-if-te-0/2)#qos cos-traffic-class test
switch(conf-if-te-0/2)#exit
switch(config)#end
switch#
```

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

```
switch#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)#show qos map cos-traffic-class test
```

Congestion control

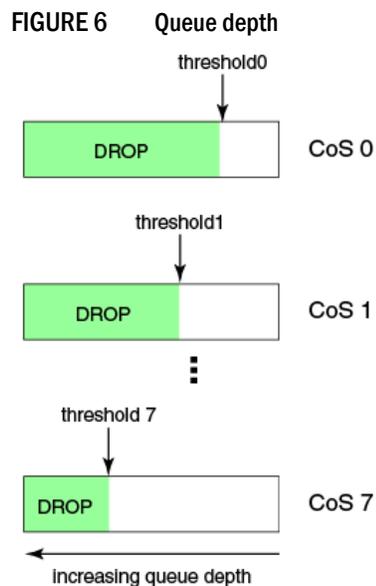
Queues can begin filling up due to a number of reasons, such as over subscription of a link or backpressure from a downstream device. Sustained, large queue buildups generally indicate congestion in the network and can affect application performance through increased queuing delays and frame loss.

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 queuing 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 6](#) describes how you can utilize this feature to ensure that 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 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 6](#) 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 1000 1000
1000 1000
```

The following example increases multicast frame expansion Tail Drop Threshold to 1000pkt for each multicast Traffic Class.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#qos rcv-queue multicast threshold 1000 1000 1000 1000 1000 1000
1000 1000
switch(config)#end
```

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

```
switch#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 intengigabitethernet 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
```

The following example enables an interface with 802.3x Pause flow control TX and RX.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#interface intengigabitethernet 0/2
switch(conf-if-te-0/2)#qos flowcontrol tx on rx on
switch(conf-if-te-0/2)#exit
switch(config)#end
```

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

```
switch#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.
- 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
```

- Specify the 10-gigabit Ethernet interface.

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

- Enable an Ethernet PFC on the interface.

```
switch(config-if-te-0/2)#qos flowcontrol pfc 3 tx on rx on
```

The following example enables an interface 802.3x Pause flow control TX and RX.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#interface intengigabitethernet 0/2
switch(config-if-te-0/2)#qos flowcontrol pfc 3 tx on rx on
switch(config-if-te-0/2)#exit
switch(config)#end
```

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

```
switch#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. 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.

- Enter global configuration mode.

```
switch#configure terminal
```

- 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
```

The following example creates a lower maximum multicast frame expansion rate to 10000pkt/s.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#qos rcv-queue multicast rate-limit 10000
switch(config)#end
```

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

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

Scheduling

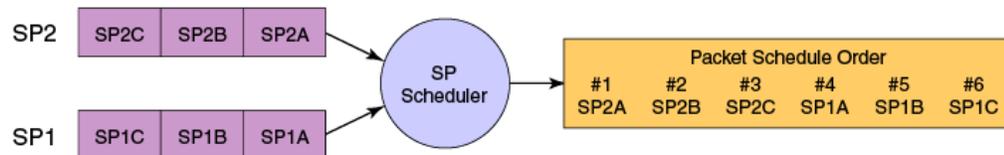
Scheduling arbitrates among multiple queues waiting to transmit a frame. The Dell M8428-k supports both Strict Priority (SP) and Deficit Weighted Round Robin (DWRR) scheduling algorithms. Also supported is the flexible selection of the number of traffic classes using SP-to-DWRR. When there are multiple 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 7 describes the frame scheduling order for an SP scheduler servicing two SP queues. The higher numbered queue, SP2, has a higher priority.

FIGURE 7 Strict priority schedule – two queues

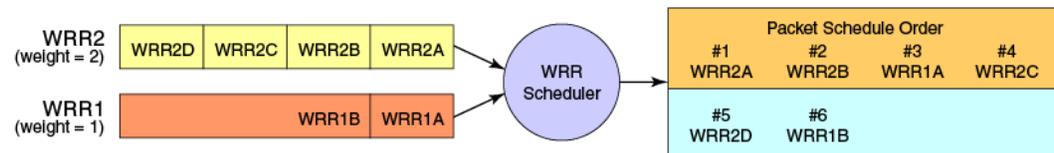


Deficit weighted round robin scheduling

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 8 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 8 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 8 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. The Dell M8428-k 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 20 describes the set of scheduling configurations supported.

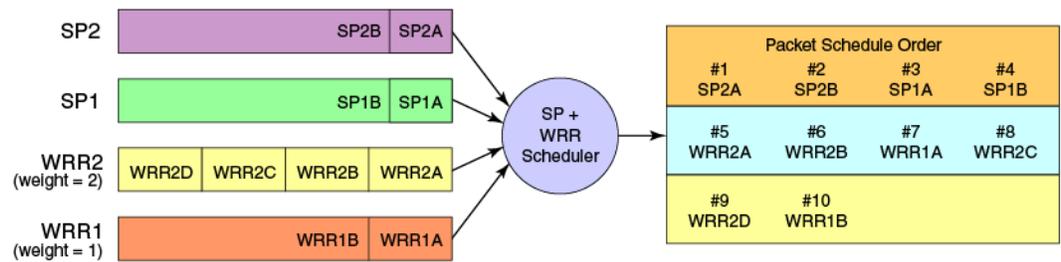
When you configure the QoS queue to use strict priority 4 (SP4), then traffic class 7 will use SP4, traffic class 6 will use SP3, and so on down the list. You use the strict priority mappings to control how the different traffic classes will be routed in the queue.

TABLE 20 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	SP2						
0	WRR1	SP1						

Figure 9 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 9 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
```

2. 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
```

The following example sets the traffic class frame scheduler for 4 Strict Priority Traffic Class and 4 DWRR Traffic Class with Traffic Class 0 getting 10 percent bandwidth, Traffic Class 1 getting 20 percent bandwidth, Traffic Class 2 getting 30 percent bandwidth, and Traffic Class 3 getting 40 percent bandwidth.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#qos queue scheduler strict-priority 4 dwrr 10 20 30 40
switch(config)#end
```

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

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

Multicast queue scheduling

The multicast traffic classes are numbered from 0 to 7; 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 21 presents the multicast traffic class equivalence mapping applied.

TABLE 21 Multicast traffic class equivalence mapping

Multicast traffic class	Equivalent unicast traffic class
0	0
1	1
2	2

TABLE 21 Multicast traffic class equivalence mapping (Continued)

Multicast traffic class	Equivalent unicast traffic class
3	3
4	4
5	5
6	6
7	7

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 20](#) on page 105 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. Because 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.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Specify the schedule to use and the traffic class to bandwidth mapping.

```
switch(config)#qos queue multicast scheduler dwrr 10 10 10 10 10 10 10 30
```

The following example sets the multicast Traffic Class frame expansion scheduler for Traffic Class 0 through 6 getting 10 percent bandwidth and Traffic Class 7 getting 30 percent bandwidth.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#qos queue multicast scheduler dwrr 10 10 10 10 10 10 10 30
switch(config)#end
```

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

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

Converged Enhanced Ethernet 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.

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 22](#) 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 22 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 7 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 23](#) details the default CEE Priority Table configuration.

TABLE 23 Default CEE priority table

CoS	PGID
0	15.7
1	15.6
2	15.5

TABLE 23 Default CEE priority table (Continued)

CoS	PGID
3	15.4
4	15.3
5	15.2
6	15.1
7	15.0

Creating a CEE map

Only a single CEE map is allowed, which is created automatically with the name “default.” You are not be able to delete or rename this map.

To create a CEE map, perform the following steps from Privileged EXEC mode.

1. Enter global configuration mode.

```
switch#configure terminal
```

2. Create a CEE map.

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

The following example creates a CEE map enter CEE-Map CLI configuration submode.

```
switch:admin>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#cee-map default
switch(config)#end
```

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

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

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. In this example ‘test’ is used.

```
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#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. Specify the name of the CEE map to define. In this example 'test' is used.

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

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#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 intengigabitethernet 0/2
```

3. Apply the CEE map on the interface.

```
switch(conf-if-te-0/2)#cee default
```

The following example applies the CEE provisioning map on an interface.

```
switch:root>cmsh
switch>enable
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#interface intengigabitethernet 0/2
switch(conf-if-te-0/2)#cee default
switch(conf-if-te-0/2)#exit
switch(config)#end
```

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

```
switch#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
```

Configuring 802.1x Port Authentication

In this chapter

- [802.1x protocol overview](#) 111
- [802.1x configuration guidelines and restrictions](#) 111
- [802.1x authentication configuration tasks](#) 112
- [Interface-specific administrative tasks for 802.1x](#) 112

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 Dell FCoE hardware.

As the authenticator, the Dell FCoE hardware prevents unauthorized network access. Upon detection of the new supplicant, the Dell 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 Dell FCoE hardware transmits an EAP-request to the supplicant, which responds with the EAP-response packet. The Dell 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 Dell 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 Dell 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)#end  
switch#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 Dell 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.

NOTE

802.1x port authentication functions only on external ports.

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 extengigabitethernet 0/16
```

3. Use the **dot1x authentication** command to enable 802.1x authentication.

```
switch(conf-if-te-0/16)#dot1x authentication
```

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

```
switch(conf-if-te-0/16)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring 802.1x timeouts on specific interface ports

NOTE

While you are free to modify the timeouts, Dell 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 extengigabitethernet 0/16
```

3. Configure the timeout interval.

```
switch(conf-if-te-0/16)#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 extengigabitethernet 0/16
```

3. Enable 802.1x authentication for the interface port.

```
switch(conf-if-te-0/16)#dot1x enable
```

4. Configure reauthentication for the interface port.

```
switch(conf-if-te-0/16)#dot1x reauthentication
```

11 Interface-specific administrative tasks for 802.1x

```
switch(conf-if-te-0/16)#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.

```
switch(config)#interface extengigabitethernet 0/16
```

3. Use the **no dot1x port-control** command to disable 802.1x Authentication.

```
switch(conf-if-te-0/16)#no dot1x authentication
```

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

```
switch(conf-if-te-0/16)#exit  
switch(config)#end  
switch#copy running-config startup-config
```

Configuring sFlow using the CEE CLI

In this chapter

- [sFlow protocol overview](#) 115
- [Configuring the sFlow protocol globally](#) 116
- [Interface-specific administrative tasks for sFlow](#) 116

sFlow protocol overview

The sFlow protocol is an industry standard technology for monitoring high-speed switched networks that can be implemented in a broad range of networking devices from Layer 2 switches to high-end core routers. The sFlow standard consists of an sFlow agent that resides on the device and sFlow collector that resides on a central server.

An sFlow agent collects statistics from the switch and forwards the data to the sFlow collector at an IP address somewhere on the network. The sFlow collector stores the sFlow datagrams from all agents on the network for processing at a later time.

The sFlow datagram provides information about the sFlow version, its originating agent's IP address, a sequence number, how many samples it contains and usually up to 10 flow samples or counter samples.

The sFlow agent uses two forms of operation:

- Time-based sampling of interface counters
- Statistical sampling of switched packets

Interface flow samples

A flow sample is based on random packets being forwarded to the sFlow collector at defined numeric intervals for a single port interface. For example, every 4,096th packet is forwarded to the sFlow collector for analysis and storage.

NOTE

This type of random sampling provides estimated flow rates, but not perfect accuracy.

Packet counter samples

A polling interval defines how often the sFlow octet and packet counter for a specific interface are sent to the collector, but the sFlow agent is free to schedule the start of the polling in order to maximize internal efficiency. If the regular schedule is chosen, each polling interval start time is chosen randomly to eliminate bottlenecks in network performance, but the polling interval remains as configured.

Configuring the sFlow protocol globally

Dell recommends that you globally configure sFlow on the switch first, then make custom alterations to specific interface ports. For details, see [“Interface-specific administrative tasks for sFlow”](#) on page 116.

NOTE

For complete information on the sFlow CLI commands for the switch, see the *Converged Enhanced Ethernet Command Reference*.

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

1. Globally enable the sFlow protocol.

```
switch(config)#sflow enable
```

2. Designate the IP address for the sFlow collector server.

```
switch(config)#sflow collector 192.10.138.176
```

3. Set the sFlow polling interval.

```
switch(config)#sflow polling-interval 135
```

4. Set the sFlow sample-rate

```
switch(config)#sflow sample-rate 4096
```

5. Use the **exit** command to return to Privileged EXEC mode.

6. Confirm the sFlow configuration status.

```
switch#show sflow
sFlow services are enabled
Global default sampling rate: 4096
Global default counter polling interval: 135
1 collectors configured
Collector IP addr: 192.10.138.176, Agent IP addr: 10.32.152.22, UDP port: 6343
77 UDP packets exported
0 UDP packets dropped
165 sFlow samples collected
```

Interface-specific administrative tasks for sFlow

Global sflow enable prepares the switch to use sFlow, as shown in [“Configuring the sFlow protocol globally”](#). However, that process does not automatically enable sFlow on individual ports. You must enable sFlow on each port. Use the administrative tasks in this section to make any necessary customizations to specific interface port settings.

Disabling sFlow on specific interfaces

By default, sFlow is disabled on all the interface ports on the switch.

NOTE

Disabling sFlow on the interface port does not completely shut down the network communication on the interface port.

To disable sFlow on a specific interface, perform the following steps in global configuration mode.

1. Enter interface mode for the port

```
switch(config)#interface intengigabitethernet 0/12
```

2. Disable sFlow on the interface.

```
switch(conf-if-te-0/12)#no sflow enable
```

3. Confirm the sFlow configuration status on the specific interface.

```
switch(conf-if)#exit
switch(config)#exit
switch#>show sflow interface intengigabitethernet 1/12
```

Configuring sFlow on specific interfaces

You can enable sFlow and configure the polling interval and sample rate for a specific interface port on the switch.

To customize sFlow on an interface, 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 extengigabitethernet 1/16
```

2. Enable sFlow for the interface.

```
switch(conf-if-te-0/1)#sflow enable interface
```

3. Configure the sFlow polling interval.

```
switch(conf-if-te-1/16)#sflow polling-interval 135
```

4. Set the sFlow sample-rate.

```
switch(conf-if-te-1/16)#sflow sample-rate 8192
```

5. Confirm the sFlow configuration status on the specific interface.

```
switch(conf-if-te-1/16)#exit
switch(config)#exit
switch#>show sflow interface extengigabitethernet 1/16
te 1/16
Configured sampling rate :8192
Actual sampling rate :8192
Counter polling interval :135
Samples rcvd from h/w :33
```

12 Interface-specific administrative tasks for sFlow

Configuring Port Mirroring using the CEE CLI

In this chapter

- [Port Mirroring protocol overview](#) 119
- [Configuring ingress Port Mirroring](#) 120
- [Configuring egress Port Mirroring](#) 120
- [Configuring bidirectional Port Mirroring](#) 120
- [Deleting a Port Mirroring connection from a session](#) 121
- [Deleting a Port Mirroring session](#) 121

Port Mirroring protocol overview

Port Mirroring is used on a network switch to send a copy of network packets seen on one switch port to a network monitoring connection on another switch port. If you are interested in listening or snooping on traffic that passes through a particular port, Port Mirroring artificially copies the packets to a port connected to your analyzer. Usually, this traffic is limited to incoming or outgoing packets, but Fabric OS v6.3.1_cee allows bidirectional traffic monitoring on the source port.

Port Mirroring limitations

The limitations of Port Mirroring connections are as follows:

- Mirror port can be any port on the same switch as the source port.
- Mirror port should not be configured to carry normal traffic.
- A port can not be mirrored to multiple locations in the same direction.
- A BP port is not supported as a source or destination mirror port, because this interface is not exposed to the user.
- Only one port per chip can be configured as a destination port for ingress mirroring.
- Only one port per chip can be configured as a destination port for egress mirroring.
- A port cannot be made a destination port for bi-directional mirroring if a different port on that chip is already configured as destination port for any type of mirroring.
- If a port is configured as a destination port of bi-directional mirroring, no other port on that chip can be made destination port for any type of mirroring.
- The destination mirror port can only handle 10G (line rate) worth of mirror traffic. If multiple ports, or both flows on same port, are mirrored to the same destination mirror port, then only 10G worth of mirror traffic is mirrored and the remaining traffic is ignored.
- If multiple ports (or both flows on same port) are mirrored to the same destination port, and two or more ports process burst traffic, the destination mirror port cannot handle all the bursts and some of the burst traffic is not mirrored.

Configuring ingress Port Mirroring

To configure Port Mirroring for incoming packets only, perform the following steps in global configuration mode.

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

2. Configure the source port and the destination port, with the rx parameter for received packets.

The destination port is always an external port.

```
switch(config-mon-sess-1)#source intengigabitethernet 0/1 destination
extengigabitethernet 0/18 direction rx
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-mon-sess-1)#description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one port. For additional ports you must create additional monitor sessions

Configuring egress Port Mirroring

To configure Port Mirroring for incoming packets only, perform the following steps in global configuration mode.

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

2. Configure the source port and the destination port, with the tx parameter for transmitted packets.

The destination port is always an external port.

```
switch(config-mon-sess-1)#source intengigabitethernet 0/1 destination
extengigabitethernet 0/18 direction tx
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-mon-sess-1)#description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one port. For additional ports you must create additional monitor sessions

Configuring bidirectional Port Mirroring

To configure Port Mirroring for packets traveling in both directions, perform the following steps in global configuration mode.

1. Open a monitor session and assign a session number

```
switch(config)# monitor session 1
```

2. Configure the source port and the destination port, with the both parameter for all packets.

The destination port is always an external port.

```
switch(config-mon-sess-1)#source intengigabitethernet 0/1 destination
extengigabitethernet 0/18 direction both
```

3. *Optional:* Use the **description** command to add a label to the monitor session.

```
switch(config-mon-sess-1)#description Hello World!
```

4. Repeat [step 1](#) and [step 2](#) as needed for additional ports.

A monitor session can have only one port. For additional ports you must create additional monitor sessions

Deleting a Port Mirroring connection from a session

To remove a single connection from a Port Mirroring session, perform the following steps in global configuration mode.

1. Display the existing configuration of the monitor session.

```
switch#show monitor session 1
```

2. Open an existing monitor session.

```
switch(config)#monitor session 1
```

3. Use the no option to delete a particular port connection.

```
switch(config-mon-sess-1)#no source intengigabitethernet 0/1 destination
extengigabitethernet 0/18 direction both
```

4. Display the monitor session again to confirm the connection deletion.

```
switch#show monitor session 1
```

5. Repeat [step 2](#) and [step 3](#) as needed for additional ports.

A monitor session can have only one port.

Deleting a Port Mirroring session

To remove a Port Mirroring session, perform the following steps in global configuration mode.

1. Display the existing configuration of the monitor session.

```
switch#show monitor session 1
```

2. Enter global configuration mode with the **enable** command.

3. Delete the existing monitor session using the no option.

```
switch(config)#no monitor session 1
```

4. Return to Privileged EXEC mode with the **exit** command.

5. Display the monitor session again to confirm the connection deletion.

```
switch#show monitor session 1
```

13 Deleting a Port Mirroring session

Configuring RMON using the CEE CLI

In this chapter

- [RMON overview](#) 123
- [RMON configuration and management](#) 123

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 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.

The following example creates an alarm that tests every sample for a rising threshold

```
switch(config)#rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 30 absolute  
rising-threshold 95 event 27 owner john_smith
```

The following example creates 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)#end  
switch#copy running-config startup-config
```

Configuring RMON events

You can add or remove an event in the RMON event table that is associated with an 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)#end  
switch#copy running-config startup-config
```

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:

- 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 intengigabitethernet 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)#exit
switch(config)#end
switch#copy running-config startup-config
```

Configuring RMON history collection

The **RMON history collection** 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. This option records periodic statistical samples from a network and stores them for later retrieval. These samples include, but aren't limited to:

- 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 intengigabitethernet 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 InTengigabitEthernet 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

Configuring IGMP

In this chapter

- About IGMP 127
- Configuring IGMP..... 128
- Configuring IGMP snooping querier 128
- Monitoring IGMP 129

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, because 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

Dell Fabric OS 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.

Use the following procedure to configure IGMP on a CEE/FCoE switch.

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. Configure a VLAN port member to be a multi-router interface.

```
switch (config)#interface vlan 25
switch (conf-if-vl-25)#ip igmp snooping mrouter interface intengigabitethernet
0/1
```

4. Repeat step 3 for each port in the VLAN, as needed.
5. Activate the default IGMP querier functionality for the VLAN.

```
switch(conf-if-vl-25)#ip igmp snooping querier enable vlan 25
```

6. *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.

Use the following procedure to configure the IGMP snooping querier.

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

2. Activate the default IGMP querier functionality for the VLAN.

```
switch(config-if-vl-25)#ip igmp snooping querier enable vlan 25
```

3. Activate IGMP querier functionality for the VLAN.

The valid range is 1 to 18000 seconds. The default is 125 seconds.

```
switch(config)#ip igmp query-interval 125
```

4. Set the last member query interval.

The valid range is 1000 to 25500 milliseconds. The default is 1000 milliseconds.

```
switch(config)#ip igmp last-member-query-interval 1000
```

5. Set the Max Response Time (MRT).

The valid range is 1 to 25 seconds. The default is 10 seconds.

```
switch(config)#ip igmp query-max-response-time 10
```

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.

Use the following procedure to monitor IGMP snooping on a CEE/FCoE switch.

1. Enter the **enable** command to access Privileged EXEC mode.

2. 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
```

3. Use the **show ip igmp statistics** command to display the IGMP statistics for a VLAN or interface.

```
switch#show ip igmp snooping statistics interface vlan 1
```

4. 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
```

5. When you have reviewed the IGMP statistics for the switch, refer to [“Configuring IGMP”](#) on page 128 or [“Configuring IGMP snooping querier”](#) on page 128 to make any needed corrections.

NOTE

Refer to the *CEE Command Reference* for additional information on IGMP CLI commands.

FCoE configuration using the Fabric OS CLI

In this chapter

- [FCoE configuration guidelines and restrictions.](#) 131
- [Managing and displaying the FCoE configuration.](#) 132
- [Managing and displaying the FCoE login configuration](#) 133
- [Creating and managing the FCoE login group configuration](#) 135

FCoE configuration guidelines and restrictions

Follow these FCoE configuration guidelines and restrictions when configuring FCoE:

- Speed negotiation—The Dell M8428-k switch supports auto-negotiated FC link speeds of 2, 4, and 8 Gbps. The Ethernet ports of the Dell M8428-k switch do not support auto-negotiation of Ethernet link speeds. The Ethernet ports only support 10-Gigabit Ethernet.
- Features that are not supported on the Dell M8428-k switch or the FCOE10-24 blade:
 - Virtual fabrics
 - Admin Domains
 - Port-based zoning
 - QoS zoning
 - Adaptive networking
 - FC-SP for the FCoE ports
 - Interop mode
 - FC routing
 - Integrated routing
 - Hot Code Load (HCL) firmware download
 - Extended fabrics
 - FICON
- 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 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.

Managing and displaying 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.

NOTE

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

Enabling or disabling an FCoE port

To enable or disable an FCoE port, perform the following tasks.

Task	Command
Enable an FCOE port.	switch:admin> fcoe --enable port
Disable an FCOE port.	switch:admin> fcoe --disable port

Configuring FCMAP values for a VLAN

NOTE

For information on FCMAPs, see [“FCoE Initialization Protocol”](#) on page 8.

If the FCMAP default value is acceptable, then it can be applied to the specified VLAN. The **fcmapunset** command is only necessary if the FCMAP value was previously set to a non-default value. For example, if you reset the default value to a value other than the default value, and then want to change the value again, you must enter the **fcmapunset** command to return the value to the default value. The **fcmapunset** command always returns the FCMAP to the default value.

To configure FCMAP values for a VLAN, perform the following tasks.

Task	Command
Configure the FCMAP values for Fabric Provided MAC Addresses (FPMA) for the specified VLANs. Syntax is as follows: <ul style="list-style-type: none"> • <i>vid</i>—Specifies the VLAN ID for which the FCMAP must be set. • <i>fcmapid</i>—Specifies the FCMAP to be set. 	switch:admin> fcoe --fcmapset -vlan vid fcmapid
Remove the FCMAP from the specified VLAN.	switch:admin> fcoe --fcmapunset -vlan vid

Configuring FIP multicast advertisement intervals

NOTE

For information on the FCoE Initialization Protocol (FIP), see [“FCoE Initialization Protocol”](#) on page 8.

To configure FIP multicast advertisement intervals, perform the following task.

Task	Command
Configure FIP multicast advertisement intervals. Syntax is as follows: <ul style="list-style-type: none"> <i>intvl</i>—Specifies the interval in seconds. The minimum interval value is 0 seconds and the maximum value is 300 seconds. A value of 0 cancels the previous advertisement interval value. 	switch:admin> fcoe -fipcfg -advintvl <i>intvl</i>

Clearing logins

To clear logins, perform the following task.

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 <i>slot/port</i>—Specifies the slot or port number. -device <i>wwn</i>—Specifies the device WWN. 	switch:admin> fcoe -resetlogin -teport <i>slot/port</i> -device <i>wwn</i>

Displaying FCoE configuration-related information

To display FCoE-related configuration information, perform the following tasks.

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 [<i>port</i>]
Display information about devices logged into a specific FCoE F_port.	switch:admin> fcoe -loginshow [<i>port</i>]
Display FIP configurations.	switch:admin> fcoe -fipcfgshow

Managing and displaying the FCoE login configuration

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

Enabling or disabling FCoE 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.

16 Managing and displaying the FCoE login configuration

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

Displaying or aborting 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

To cleanup login groups and VN_port mappings, perform the following tasks.

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
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

To display the FCoE login configuration, perform the following tasks.

Task	Command
Display the FCoE login configuration. Syntax is as follows: <ul style="list-style-type: none"> • -switch <i>swwn</i>—Displays all of the login groups for the specified switch. • -logingroup <i>lgroupname</i>—Displays the login group configuration for the specified login group. • -saved—Displays only the effective configuration. 	switch:admin> fcoelogincfg -show [-switch <i>swwn</i> -logingroup <i>lgroupname</i>] [-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

To save the current FCoE configuration, perform the following task.

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

Creating and managing the FCoE login group configuration

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

Creating an FCoE login group

The FCoE login group enables you to configure login policies.

16 Creating and managing the FCoE login group configuration

To create an FCoE login group, perform the following task.

Task	Command
Syntax is as follows: <ul style="list-style-type: none">• -create—Create a login group.• <i>lgnam</i>e—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 into the switch.	switch:admin> fcoelogingroup -create <i>lgnam</i> e -switch <i>swwn</i> -self [- allowall " <i>member</i> ; <i>member</i> ;..."]

Modifying the FCoE login group device list

To add or remove VN_port devices from the FCoE login group, perform the following tasks.

Task	Command
Add VN_port devices to the FCoE login group. Syntax is as follows: <ul style="list-style-type: none">• <i>lgnam</i>e—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 into the switch.	switch:admin> fcoelogingroup -add <i>lgnam</i> e <i>member</i> ; <i>member</i> ;...
Remove VN_port devices from the FCoE login group. Syntax is as follows: <ul style="list-style-type: none">• <i>lgnam</i>e—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.x. Only specified members are allowed to log into the switch.	switch:admin> fcoelogingroup -remove <i>lgnam</i> e <i>member</i> ; <i>member</i> ;...

Deleting an FCoE login group

To delete an FCoE login group, perform the following task.

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

Renaming an FCoE login group

To rename an FCoE login group, perform the following task.

Task	Command
Rename an FCoE login group. Syntax is as follows:	switch:admin> fcoelogingroup -rename <i>lgname</i> <i>newlgname</i>
<ul style="list-style-type: none">• <i>lgname</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.x. Only specified members are allowed to log into the switch.	

16 Creating and managing the FCoE login group configuration

CEE configuration management

In this chapter

- CEE configuration management guidelines and restrictions 139
- CEE configuration management tasks. 139
- Flash file management commands 142
- Debugging and logging commands 143

CEE configuration management 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 CEE configuration must be manually saved using the CEE CLI **write** or **copy** commands.

CEE configuration management tasks

This section describes the typical configuration management tasks you may encounter when administering the switch.

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 CEE 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.

Display the running configuration file

To display the running configuration, perform the following steps from Privileged EXEC mode.

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

```
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 steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to copy the currently running configuration to the startup configuration.

```
switch#copy running-config startup-config  
Overwrite the startup config file (y/n): y
```

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 steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to load the startup configuration.

```
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 steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the write command to erase the startup configuration file.

```
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 steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to archive the running configuration file.

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

Restore an archived running configuration file

To restore the running configuration, perform the following steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to restore the running configuration file.

```
switch#copy running-config ftp://jsmith:password@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 steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to archive the startup configuration file.

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

Restore an archived startup configuration file

To restore the startup configuration, perform the following steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to restore the startup configuration file.

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

Archive a startup configuration from Flash

This task also works for running configuration files.

To archive the startup configuration, perform the following steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **copy** command to restore the archived configuration file.

```
switch#copy startup-config flash://config_filename
```

Restore a startup configuration file from Flash

This task also works for running configuration files.

To restore the startup configuration, perform the following steps from Privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the copy command to restore the archived configuration file.

```
switch#copy flash://config_filename startup-config
```

CEE configuration management commands

Table 24 lists the common CEE configuration management commands.

TABLE 24 CEE configuration management commands

Task	Command
Write the current running configuration file to the startup configuration file. NOTE: If you enter y at the prompt, the running configuration file overwrites the startup configuration file. If you enter n at the prompt, the startup configuration file is not overwritten.	<pre>switch#copy running-config startup-config Overwrite the startup config file (y/n): y</pre>
Copy the startup configuration file to the running configuration file.	<pre>switch#copy startup-config running-config</pre>
Erase the startup configuration file. NOTE: This command does not affect the running configuration file.	<pre>switch#write erase</pre>
Copy the running configuration file to the archive folder on an FTP server.	<pre>switch#copy running-config ftp://jsmith:password@/archive/config_file]</pre>
Copy a stored startup configuration file in Flash to the running configuration.	<pre>switch#copy flash://test_filename running-config</pre>
Copy a configuration file from an FTP server to the startup configuration.	<pre>switch#copy ftp://jsmith:password@/archive/test_filename startup-config}</pre>
Display the contents of the running configuration file.	<pre>switch#show running-config</pre>

Flash file management commands

Table 25 describes the common tasks used to manage the Flash files on the switch. The *Converged Enhanced Ethernet Command Reference* contains complete information on all available CLI commands.

NOTE

Use of the flash:// prefix is optional.

TABLE 25 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 26](#) describes the tasks related to debugging and logging commands on the switch. The *Converged Enhanced Ethernet Command Reference* contains complete information on all available CLI commands.

Perform the following tasks from Privileged EXEC mode.

TABLE 26 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.

17 Debugging and logging commands

Index

Symbols

Numerics

- 8000 CEE switch
 - congestion control and queuing, 6
 - flow control, 8
 - Layer 2 Ethernet, 3
 - Layer 2 forwarding, 3
 - loop-free, 5
 - tagging, 4
 - trunking, 8
- 802.1x
 - LAG, 27, 111
 - overview, 27, 111
 - timeouts, 113

A

- Access Control Lists
 - See ACL
- access interface, configuring, 36
- access mode, 31, 36
- ACL
 - configuration guidelines and restrictions, 88
 - configuration procedures
 - applying a MAC ACL to a CEE interface, 91
 - applying a MAC ACL to a VLAN interface, 91
 - creating extended MAC ACL and adding rules, 89
 - creating standard MAC ACL and adding rules, 88
 - important notes, 88
 - modifying a MAC ACL, 89
 - removing a MAC ACL, 90
 - reordering the sequence numbers, 90
 - default configuration, 88
 - extended ACL, defined, 87
 - overview, 7, 87
 - standard ACL, defined, 87
- active IGMP, 127

- authentication server, 27, 111
- authenticator, 27, 111

B

- basic management TLV sets, 76
- bridge
 - forwarding delay, configuring for STP, RSTP, MSTP, 52
 - hello time, configuring for STP, RSTP, 54
 - maximum aging time, configuring for STP, RSTP, MSTP, 53
 - priority, configuring for STP, RSTP, MSTP, 52
 - extension TLV set, 77
 - proprietary aggregation, 68

C

- CEE interface
 - applying a MAC ACL, 91
 - configuring for STP, RSTP, MSTP, 58
 - configuring the hello time for MSTP, 60
 - disable or enable STP on the interface, 62
 - enabling and disabling, 34
 - enabling as an edge port for RSTP, MSTP, 59
 - enabling guard root for STP, RSTP, MSTP, 59
 - enabling LACP, 69
 - enabling port fast, 61
 - path cost, configuring for STP, RSTP, MSTP, 58
 - restricting the port from becoming a root port for STP, RSTP, MSTP, 62
 - restricting the topology change notification for STP, RSTP, MSTP, 62
 - spanning-tree defaults, 50
 - specifying a link type, 61
 - specifying restrictions for an MSTP instance, 60
 - specifying the port priority for STP, RSTP, MSTP, 61
- CEE map, configuring, 109
- CEE maps, verifying, 110
- CEE provisioning map, applying, 110
- Cisco interoperability, disabling for MSTP, 55
- Cisco interoperability, enabling for MSTP, 55

- classifier groups, VLAN, 39
- classifier rules, VLAN, 38
- CLI, CEE
 - accessing, 15
 - 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, 15
- command completion, CEE CLI, 19
- command modes, CEE, 15
- command output modifiers, 19
- command syntax, 18
- configuration management
 - saving changes, 14
- congestion control
 - QoS, 100
 - queuing, 6
- console interface, 15
- converged mode, 31
- counters, clearing, 40

D

- Data Center Bridging (DCB) Capability Exchange Protocol
 - See DCBX
- DCBX
 - Enhanced Transmission Selection, 78
 - interaction with other vendor devices, 79
 - overview
 - Priority Flow Control, 79
 - TLV sets, 24
- Dell FCoE hardware, 2
- document conventions, *xvi*
- dynamic link aggregation, 68

E

- EAP, 27, 111
- edge detection, configuring for STP, RSTP, MSTP, 58
- edge port, enabling a CEE interface as an edge port for
 - RSTP, MSTP, 59
- enable
 - track, 71
- Enhanced Transmission Selection
 - See ETS
- error disable timeout interval, configuring for STP, RSTP, MSTP, 53
- error disable timeout, configuring for STP, RSTP, MSTP, 53
- Ethernet, forwarding, 3
- ETS
 - overview
 - priority grouping of IPC, LAN, and SAN traffic, 78

F

- fabric OS shell, 15
- FCoE
 - configuration guidelines and restrictions, 131
 - configuration procedures
 - creating and managing the FCoE login group configuration, 135
 - managing and displaying FCoE login configuration, 133
 - managing and displaying the configuration, 132
 - Layer 2 Ethernet overview, 3
 - login, 10
 - logout, 10
 - overview, 1
 - queuing, 12
 - speed negotiation, 131
 - terminology
 - CEE, 1
 - ENode, 1
 - FCoE Forwarder (FCF), 1
 - VF_port, 1
 - VN_port, 1
 - unsupported features, 131
 - VLAN forwarding, 4
- FCoE initialization protocol
 - See FIP
- filtering VLAN ingress, 31

FIP

- FC zoning, 11
- FCoE login, 10
- FCoE logout, 10
- FIP discovery, 8
- login, 9
- logincfg, 11
- logout, 10
- name server, 11
- registered state change notification (RSCN), 12

FLOGI, 1

flow control, 8

flushing MAC addresses, 57

frame classification, incoming, 5

G

guard root, enabling on a CEE interface for STP, RSTP, MSTP, 59

H

hello time, configuring for MSTP, 60

hops, configuring for MSTP, 56

I

IEEE 802.1 organizational TLV set, 76

IEEE 802.3 organizational TLV set, 77

IGMP

- interface, 128
- interval, 128
- mrouter, 128
- MRT, 128
- passive, 127
- querier, 128
- query-interval, 128
- tcn, 128
- timer, 128
- vlan, 128

incoming frame classification, 5

ingress VLAN filtering, 31

instance

- MSTP, mapping a VLAN to, 55
- specifying restrictions for an MSTP instance, 60

interface

- track, 71

K

keyboard shortcuts, CEE CLI, 17

L

LACP

- configuration guidelines and restrictions, 69
- configuration procedures
 - clearing counters, 71
 - configuring system priority, 70
 - configuring timeout period, 70
 - displaying LACP information, 72
 - enabling on a CEE interface, 69
 - important notes, 69
- default LACP configuration, 69
- overview
- troubleshooting tips, 72

LAGs

- 802.1x, 27, 111
- distribution process, 68
- overview
- top-of-the-rack configuration, 67

Layer 2

- ACL
- Ethernet overview, 3

Layer 2 forwarding, 3

link aggregation

- Dell-proprietary, 68
- dynamic, 68
- LACP, 68
- LAG distribution process, 68
- LAGs, 65
- overview, 65
- static, 68

Link Aggregation Control Protocol

See LACP

link aggregation group

See LAGs

Link Layer Discovery Protocol

See LLDP

link type, specifying, 61

LLDP

- configuration guidelines and restrictions, 79
- configuration procedures
 - clearing LLDP-related information, 85
 - disabling LLDP globally, 80
 - displaying LLDP-related information, 85
 - enabling LLDP globally, 80
 - global command options, 81
 - important notes, 80
 - interface-level command options, 84
- DCBX overview
- default configuration, 80
- Layer 2 topology mapping, 76
- overview, 75
- TLV sets, 76

login

- FCoE, 10
- FIP, 9

logincfg, 11

logout

- FCoE, 10
- FIP, 10

loop-free network environment, 5

M

MAC addresses

- configuration guidelines and restrictions, 33
- flush from the VLAN FDB, 57

minimum links, 70

MSTP

- configuration procedures, 51
- default configuration, 50
- displaying MSTP-related information, 58
- overview, 47

MTU, configuring, 34

multicast rate limiting, QoS, 103

Multiple Spanning Tree Protocol

- See MSTP

N

name server, 11

network

- flow control, 8
- loop-free
 - STP, RSTP, MSTP, 5
- trunking, 8

O

output modifiers, CEE CLI, 19

overview

- ACL, 87
- link aggregation, 65
- MSTP, 47
- RSTP, 45
- STP, 43

P

passive IGMP, 127

path cost

- CEE interface, configuring for STP, RSTP, MSTP, 58
- port channel, configuring for STP, RSTP, MSTP, 54

PEAP, 27, 111

port configuration for STP, RSTP, MSTP, 58

port fast, enabling on a CEE interface, 61

port priority, specifying on a CEE interface for STP, RSTP, MSTP, 61

Priority Flow Control (PFC), 79

priority group table, mapping, 109

priority mapping, QoS, 94

priority-table, mapping, 110

Q

QoS

CEE QoS overview, 107

configuration procedures

- applying a CEE provisioning map, 110
- creating a CEE map, 109
- mapping a priority group table, 109
- mapping a priority-table, 110
- overview, 107
- verifying CEE maps, 110

congestion control, 100

multicast rate limiting, 103

overview, 93, 115, 119

queuing

- traffic class mapping, 97
- user-priority mapping, 94

queuing overview, 94

rewriting frame header field, 94

scheduling, 104

Quality of Service

- See QoS

querier
interval, 128
MRT, 128
VLAN, 128

queuing
congestion control, 6
FCoE, 12
QoS, 94

R

RADIUS, 27, 111
Rapid Spanning Tree Protocol
See RSTP
RBAC permissions
region name, specifying for MSTP, 56
registered state notification protocol (RSCN), 12
revision number, specifying for MSTP, 56
Role-Based Action Control
See RBAC
root port, CEE interface, restricting for STP, RSTP, MSTP, 62
RSTP
configuration guidelines and restrictions
MSTP configuration guidelines and restrictions, 49
configuration procedures, 51
default configuration, 50
displaying RSTP-related information, 58
overview, 45

S

saving configuration, 14
scheduling, QoS, 104
Spanning Tree Protocol
See STP
spanning-tree defaults, 50
speed negotiation, FC ports, 131
static link aggregation, 68
STP
configuration guidelines and restrictions, 49
configuration procedures, 51
default configuration, 50
displaying STP-related information, 58
overview, 43
supplicant, 27, 111

switch
port configuration, 36
system priority, configuring for LACP, 70

T

T11-FC-BB5, 1
technical help, xviii
telnet, 15
terminology
FCoE, 1
timeout period, configuring for LACP, 70
TLV sets
basic management TLV, 76
Dell extension TLV set, 77
configuring, 24
IEEE 802.1 organizational TLV set, 76
IEEE 802.3 organizational TLV set, 77
top-of-the-rack configuration, 67
topology change notification, CEE interface, restricting for STP, RSTP, MSTP, 62
topology mapping, LLDP, 76
track
enable, 71
interface, 71
traffic class mapping, QoS, 97
transmit hold count, configuring for RSTP, MSTP, 54
troubleshooting tips, LACP, 72
trunk interface, configuring, 36
trunk mode, 31, 36
trunking, 8

U

unsupported features, 131
user-priority mapping, QoS, 94

V

Virtual LANs
See VLAN

VLAN

- applying a MAC ACL, 91
- configuration guidelines and restrictions, 33
- configuration procedures
 - clearing VLAN counters, 40
 - configuring a CEE interface as a Layer 2 switch port, 36
 - configuring a CEE interface as an access or trunk interface, 36
 - configuring the MTU on an interface, 34
 - displaying VLAN information, 40
 - enabling and disabling a CEE interface, 34
 - important notes, 34
 - VLAN classifier groups, 39
 - VLAN classifier rules, 38
- default configuration, 33
- FDB
 - flushing, 57
 - overview, 32
- forwarding, 4
- important management notes, 34
- ingress VLAN filtering, 31
- overview, 31
- tagging, 4

Z

- zoning, FC, 11