

**Dell PowerEdge FN I/O Aggregator
Configuration Guide
9.7(0.0)**



Notes, Cautions, and Warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your computer.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Copyright © 2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

2015 - 02

Rev. A01

Contents

1 About this Guide.....	13
Audience.....	13
Conventions.....	13
Information Symbols.....	14
Related Documents.....	14
2 Before You Start.....	15
IOA Operational Modes.....	15
Standalone mode.....	15
VLT mode.....	15
Programmable MUX mode.....	15
Stacking mode.....	16
Default Settings.....	16
Other Auto-Configured Settings.....	16
Data Center Bridging Support.....	17
FCoE Connectivity and FIP Snooping.....	17
iSCSI Operation.....	17
Link Aggregation.....	18
Link Tracking.....	18
Configuring VLANs.....	18
Uplink LAG.....	18
Server-Facing LAGs.....	19
Where to Go From Here.....	19
3 Configuration Fundamentals.....	20
Accessing the Command Line.....	20
CLI Modes.....	20
Navigating CLI Modes.....	21
The do Command.....	22
Undoing Commands.....	23
Obtaining Help.....	23
Entering and Editing Commands.....	24
Command History.....	25
Filtering show Command Outputs.....	25
Multiple Users in Configuration Mode.....	26
4 Data Center Bridging (DCB).....	27
Supported Modes.....	27

Ethernet Enhancements in Data Center Bridging.....	27
Priority-Based Flow Control.....	28
Configuring Priority-Based Flow Control.....	29
Enhanced Transmission Selection.....	31
Configuring Enhanced Transmission Selection.....	33
Creating a DCB Map.....	33
Important Points to Remember.....	34
Applying a DCB Map on Server-Facing Ethernet Ports.....	34
Data Center Bridging Exchange Protocol (DCBx).....	34
Data Center Bridging in a Traffic Flow.....	35
Data Center Bridging: Auto-DCB-Enable Mode.....	35
QoS dot1p Traffic Classification and Queue Assignment.....	37
How Priority-Based Flow Control is Implemented.....	38
How Enhanced Transmission Selection is Implemented.....	39
ETS Operation with DCBx.....	40
DCBx Operation.....	40
DCBx Operation.....	40
DCBx Port Roles.....	41
DCB Configuration Exchange.....	42
Configuration Source Election.....	42
Propagation of DCB Information.....	43
Auto-Detection of the DCBx Version.....	43
DCBx Example.....	44
DCBx Prerequisites and Restrictions.....	45
DCBx Error Messages.....	46
Debugging DCBx on an Interface.....	46
Verifying the DCB Configuration.....	47
Hierarchical Scheduling in ETS Output Policies.....	57
Troubleshooting PFC, ETS, and DCBx Operation.....	57

5 Dynamic Host Configuration Protocol (DHCP)..... 60

Supported Modes.....	60
Assigning an IP Address using DHCP.....	60
Debugging DHCP Client Operation.....	62
DHCP Client.....	65
How DHCP Client is Implemented.....	65
DHCP Client on a Management Interface.....	66
DHCP Client on a VLAN.....	66
DHCP Packet Format and Options.....	66
Option 82.....	68
Releasing and Renewing DHCP-based IP Addresses.....	69
Viewing DHCP Statistics and Lease Information.....	69

6 FIP Snooping.....	71
Supported Modes.....	71
Fibre Channel over Ethernet.....	71
Ensuring Robustness in a Converged Ethernet Network.....	71
FIP Snooping on Ethernet Bridges.....	73
How FIP Snooping is Implemented.....	74
FIP Snooping on VLANs.....	75
FC-MAP Value.....	75
Bridge-to-FCF Links.....	75
Impact on other Software Features.....	75
FIP Snooping Prerequisites.....	75
FIP Snooping Restrictions.....	76
Configuring FIP Snooping.....	76
Displaying FIP Snooping Information.....	77
FIP Snooping Example.....	83
Debugging FIP Snooping	84
7 Internet Group Management Protocol (IGMP).....	85
IGMP Overview.....	85
IGMP Version 2.....	85
Joining a Multicast Group.....	86
Leaving a Multicast Group.....	86
IGMP Version 3.....	86
Joining and Filtering Groups and Sources.....	87
Leaving and Staying in Groups.....	88
IGMP Snooping.....	89
How IGMP Snooping is Implemented on an Aggregator.....	89
Disabling Multicast Flooding.....	90
Displaying IGMP Information.....	90
8 Interfaces.....	92
Basic Interface Configuration.....	92
Advanced Interface Configuration.....	92
Interface Auto-Configuration.....	92
Interface Types.....	93
Viewing Interface Information.....	93
Disabling and Re-enabling a Physical Interface.....	95
Layer 2 Mode.....	95
Management Interfaces.....	96
Accessing an Aggregator.....	96
Configuring a Management Interface.....	96

Configuring a Static Route for a Management Interface.....	97
VLAN Membership.....	98
Default VLAN	98
Port-Based VLANs.....	98
VLANs and Port Tagging.....	99
Configuring VLAN Membership.....	99
Displaying VLAN Membership.....	100
Adding an Interface to a Tagged VLAN.....	101
Adding an Interface to an Untagged VLAN.....	101
VLAN Configuration on Physical Ports and Port-Channels.....	102
Port Channel Interfaces.....	103
Port Channel Definitions and Standards.....	103
Port Channel Benefits.....	104
Port Channel Implementation.....	104
10GbE Interface in Port Channels.....	104
Uplink Port Channel: VLAN Membership.....	105
Server-Facing Port Channel: VLAN Membership.....	105
Displaying Port Channel Information.....	105
Interface Range.....	106
Bulk Configuration Examples.....	107
Monitor and Maintain Interfaces.....	108
Maintenance Using TDR.....	109
Flow Control Using Ethernet Pause Frames.....	109
Enabling Pause Frames.....	110
MTU Size.....	111
Auto-Negotiation on Ethernet Interfaces.....	112
Setting Auto-Negotiation Options.....	113
Viewing Interface Information.....	114
Clearing Interface Counters.....	115
Fibre Channel Interface.....	116
Configuring Fibre Channel Interfaces.....	116
Enabling Fibre Channel Capability.....	116
Configuring Fibre Channel Interfaces.....	116
9 iSCSI Optimization.....	117
Supported Modes.....	117
iSCSI Optimization Overview.....	117
Monitoring iSCSI Traffic Flows.....	119
Information Monitored in iSCSI Traffic Flows.....	119
Synchronizing iSCSI Sessions Learned on VLT-Lags with VLT-Peer.....	119
iSCSI Optimization: Operation.....	119
Configuring iSCSI Optimization.....	120

Displaying iSCSI Optimization Information.....	122
10 Isolated Networks for Aggregators.....	124
Configuring and Verifying Isolated Network Settings.....	124
11 Link Aggregation.....	125
Supported Modes.....	125
How the LACP is Implemented on an Aggregator.....	125
Uplink LAG.....	126
Server-Facing LAGs.....	126
LACP Modes.....	126
Auto-Configured LACP Timeout.....	126
Link Aggregation Control Protocol (LACP).....	127
Configuration Tasks for Port Channel Interfaces.....	127
Creating a Port Channel.....	127
Adding a Physical Interface to a Port Channel.....	127
Reassigning an Interface to a New Port Channel.....	130
Configuring the Minimum Oper Up Links in a Port Channel.....	130
Configuring VLAN Tags for Member Interfaces.....	131
Deleting or Disabling a Port Channel.....	131
Configuring Auto LAG.....	132
Configuring the Minimum Number of Links to be Up for Uplink LAGs to be Active.....	134
Optimizing Traffic Disruption Over LAG Interfaces On IOA Switches in VLT Mode.....	135
Preserving LAG and Port Channel Settings in Nonvolatile Storage.....	135
Enabling the Verification of Member Links Utilization in a LAG Bundle.....	135
Monitoring the Member Links of a LAG Bundle.....	136
Verifying LACP Operation and LAG Configuration.....	137
Multiple Uplink LAGs with 10G Member Ports.....	139
.....	139
12 Layer 2.....	141
Supported Modes.....	141
Managing the MAC Address Table.....	141
Clearing the MAC Address Entries.....	141
Displaying the MAC Address Table.....	142
Network Interface Controller (NIC) Teaming.....	142
MAC Address Station Move.....	143
MAC Move Optimization.....	144
13 Link Layer Discovery Protocol (LLDP).....	145
Supported Modes.....	145
Protocol Data Units.....	145

Configure LLDP.....	147
Related Configuration Tasks.....	147
Important Points to Remember.....	147
CONFIGURATION versus INTERFACE Configurations.....	147
Enabling LLDP.....	148
Disabling and Undoing LLDP.....	148
Advertising TLVs.....	148
Optional TLVs.....	150
Management TLVs.....	150
IEEE Organizationally Specific TLVs.....	150
LLDP-MED Capabilities TLV.....	152
LLDP-MED Network Policies TLV.....	153
Extended Power via MDI TLV.....	154
LLDP Operation.....	155
Viewing the LLDP Configuration.....	155
Viewing Information Advertised by Adjacent LLDP Agents.....	156
Configuring LLDPDU Intervals.....	157
Configuring a Time to Live.....	157
Clearing LLDP Counters.....	158
Debugging LLDP.....	158
Relevant Management Objects.....	159
14 Port Monitoring.....	165
Supported Modes.....	165
Configuring Port Monitoring.....	165
Important Points to Remember.....	166
Port Monitoring.....	167
15 Security.....	169
Supported Modes.....	169
Understanding Banner Settings.....	169
Accessing the I/O Aggregator Using the CMC Console Only.....	169
AAA Accounting.....	170
Configuration Task List for AAA Accounting.....	170
AAA Authentication.....	172
Configuration Task List for AAA Authentication.....	172
AAA Authorization.....	175
Privilege Levels Overview.....	175
Configuration Task List for Privilege Levels.....	176
RADIUS.....	180
RADIUS Authentication.....	180
Configuration Task List for RADIUS.....	180

TACACS+.....	183
Configuration Task List for TACACS+.....	183
TACACS+ Remote Authentication.....	185
Enabling SCP and SSH.....	186
Using SCP with SSH to Copy a Software Image.....	187
Secure Shell Authentication.....	187
Telnet.....	187
VTY Line and Access-Class Configuration.....	188
VTY Line Local Authentication and Authorization.....	188
VTY Line Remote Authentication and Authorization.....	189
16 Simple Network Management Protocol (SNMP).....	190
Supported Modes.....	190
Implementation Information.....	190
Configuring the Simple Network Management Protocol.....	190
Important Points to Remember.....	190
Setting up SNMP.....	191
Creating a Community.....	191
Reading Managed Object Values.....	191
Displaying the Ports in a VLAN using SNMP.....	192
Fetching Dynamic MAC Entries using SNMP.....	194
Deriving Interface Indices.....	195
Monitor Port-Channels.....	196
Entity MIBS.....	197
Example of Sample Entity MIBS outputs.....	197
SNMP Traps for Link Status.....	198
Standard VLAN MIB.....	198
Enhancements.....	198
Fetching the Switchport Configuration and the Logical Interface Configuration	199
MIB Support to Display the Available Memory Size on Flash.....	200
Viewing the Available Flash Memory Size.....	200
MIB Support to Display the Software Core Files Generated by the System.....	200
Viewing the Software Core Files Generated by the System.....	201
17 Stacking.....	202
Supported Modes.....	202
Configuring a Switch Stack.....	202
Stacking Prerequisites.....	202
Master Selection Criteria.....	203
Configuring Priority and stack-group.....	203
Cabling the Switch Stack.....	204
Accessing the CLI.....	204

Configuring and Bringing Up a Stack.....	204
Adding a Stack Unit.....	205
Resetting a Unit on a Stack.....	206
Removing an Aggregator from a Stack.....	206
Merging Two Operational Stacks.....	207
Verifying a Stack Configuration.....	207
Using Show Commands.....	207
Troubleshooting a Switch Stack.....	208
Failure Scenarios.....	208
Upgrading a Switch Stack.....	210
Upgrading a Single Stack Unit.....	211
18 Broadcast Storm Control.....	213
Supported Modes.....	213
Disabling Broadcast Storm Control.....	213
Displaying Broadcast-Storm Control Status.....	213
Configuring Storm Control.....	213
19 System Time and Date.....	214
Supported Modes.....	214
Setting the Time for the Software Clock.....	214
Setting the Timezone.....	214
Setting Daylight Savings Time.....	215
Setting Daylight Saving Time Once.....	215
Setting Recurring Daylight Saving Time.....	216
20 Uplink Failure Detection (UFD).....	218
Supported Modes.....	218
Feature Description.....	218
How Uplink Failure Detection Works.....	219
UFD and NIC Teaming.....	221
Important Points to Remember.....	221
Uplink Failure Detection (SMUX mode).....	222
Configuring Uplink Failure Detection (PMUX mode).....	222
Clearing a UFD-Disabled Interface (in PMUX mode).....	224
Displaying Uplink Failure Detection.....	226
Sample Configuration: Uplink Failure Detection.....	228
21 PMUX Mode of the IO Aggregator.....	230
I/O Aggregator (IOA) Programmable MUX (PMUX) Mode.....	230
Configuring and Changing to PMUX Mode.....	230
Configuring the Commands without a Separate User Account.....	231

Virtual Link Trunking (VLT).....	231
Overview.....	231
Setting up VLT.....	232
VLT Terminology.....	233
Configure Virtual Link Trunking.....	233
Verifying a VLT Configuration.....	238
VLT Sample Configurations.....	241
Troubleshooting VLT.....	243
22 NPIV Proxy Gateway.....	245
NPIV Proxy Gateway Configuration.....	245
NPIV Proxy Gateway Operations and Capabilities.....	245
NPIV Proxy Gateway Operation	245
NPIV Proxy Gateway: Protocol Services.....	246
NPIV Proxy Gateway Functionality.....	246
NPIV Proxy Gateway: Terms and Definitions.....	247
Configuring an NPIV Proxy Gateway.....	249
Enabling Fibre Channel Capability on the Switch.....	250
Creating a DCB Map	250
Applying a DCB Map on Server-facing Ethernet Ports	251
Creating an FCoE VLAN.....	252
Creating an FCoE Map	252
Applying an FCoE Map on Server-facing Ethernet Ports.....	253
Applying an FCoE Map on Fabric-facing FC Ports.....	254
Sample Configuration.....	255
Displaying NPIV Proxy Gateway Information.....	256
show interfaces status Command Example.....	256
show fcoe-map Command Examples	257
show qos dcb-map Command Examples	258
show npiv devices brief Command Example.....	259
show npiv devices Command Example	260
show fc switch Command Example	261
Displaying NPIV Proxy Gateway Information.....	261
show interfaces status Command Example.....	262
show fcoe-map Command Examples	263
show qos dcb-map Command Examples	264
show npiv devices brief Command Example.....	265
show npiv devices Command Example	265
show fc switch Command Example	267
23 Upgrade Procedures.....	268
Get Help with Upgrades.....	268


24 Debugging and Diagnostics.....	269
Supported Modes.....	269
Debugging Aggregator Operation.....	269
All interfaces on the Aggregator are operationally down.....	269
Broadcast, unknown multicast, and DLF packets switched at a very low rate.....	270
Flooded packets on all VLANs are received on a server.....	270
Software show Commands.....	271
Offline Diagnostics.....	272
Important Points to Remember.....	272
Running Offline Diagnostics.....	272
Trace Logs.....	273
Auto Save on Crash or Rollover.....	274
Using the Show Hardware Commands.....	274
Environmental Monitoring.....	276
Recognize an Over-Temperature Condition.....	276
Troubleshoot an Over-Temperature Condition.....	277
Recognize an Under-Voltage Condition.....	277
Troubleshoot an Under-Voltage Condition.....	278
Buffer Tuning.....	278
Deciding to Tune Buffers.....	280
Sample Buffer Profile Configuration.....	283
Troubleshooting Packet Loss.....	284
Displaying Drop Counters.....	284
Dataplane Statistics.....	285
Displaying Stack Port Statistics.....	286
Restoring the Factory Default Settings.....	287
Important Points to Remember.....	287
25 Standards Compliance.....	288
IEEE Compliance.....	288
RFC and I-D Compliance.....	288
General Internet Protocols.....	289
General IPv4 Protocols.....	289
Network Management.....	290
MIB Location.....	292

About this Guide

This guide describes the supported protocols and software features, and provides configuration instructions and examples, for the Dell Networking FN I/O Aggregator running Dell Networking OS version 9.6(0.0).

The I/O Aggregator is installed in a Dell PowerEdge FX2 server chassis. For information about how to install and perform the initial switch configuration, refer to the *Getting Started Guides* on the Dell Support website at <http://www.dell.com/support/manuals>

Though this guide contains information about protocols, it is not intended to be a complete reference. This guide is a reference for configuring protocols on Dell Networking systems. For complete information about protocols, refer to other documentation, including IETF requests for comment (RFCs). The instructions in this guide cite relevant RFCs, and Standards Compliance contains a complete list of the supported RFCs and management information base files (MIBs).

 **NOTE:** You can perform some of the configuration tasks described in this document by using either the Dell command line or the chassis management controller (CMC) graphical interface. Tasks supported by the CMC interface are shown with the CMC icon: CMC

Audience

This document is intended for system administrators who are responsible for configuring and maintaining networks and assumes knowledge in Layer 2 and Layer 3 networking technologies.

Conventions

This guide uses the following conventions to describe command syntax.

Keyword	Keywords are in Courier (a monospaced font) and must be entered in the CLI as listed.
<i>parameter</i>	Parameters are in italics and require a number or word to be entered in the CLI.
{X}	Keywords and parameters within braces must be entered in the CLI.
[X]	Keywords and parameters within brackets are optional.
x y	Keywords and parameters separated by a bar require you to choose one option.
x y	Keywords and parameters separated by a double bar allows you to choose any or all of the options.

Information Symbols

This book uses the following information symbols.



NOTE: The Note icon signals important operational information.



CAUTION: The Caution icon signals information about situations that could result in equipment damage or loss of data.



WARNING: The Warning icon signals information about hardware handling that could result in injury.

* (Exception). This symbol is a note associated with additional text on the page that is marked with an asterisk.

Related Documents

For more information about the Dell PowerEdge FN I/O Aggregator, refer to the following documents:

- *Dell PowerEdge FN I/O Aggregator Command Line Reference Guide*
- *Dell PowerEdge FN I/O Aggregator Getting Started Guide*
- *Release Notes for the Dell PowerEdge FN I/O Aggregator*

Before You Start

To install the Aggregator in a Dell PowerEdge FX2 server chassis, use the instructions in the Dell PowerEdge FN I/O Aggregator Getting Started Guide that is shipped with the product. The I/O Aggregator (also known as Aggregator) installs with zero-touch configuration. After you power it on, an Aggregator boots up with default settings and auto-configures with software features enabled. This chapter describes the default settings and software features that are automatically configured at startup. To reconfigure the Aggregator for customized network operation, use the tasks described in the other chapters.

IOA Operational Modes

IOA supports four operational modes. Select the operational mode that meets your deployment needs. To enable a new operational mode, reload the switch.

Standalone mode

stack-unit *unit* iom-mode standalone

CONFIGURATION mode

```
Dell(conf)#stack-unit 0 iom-mode standalone
```

This is the default mode for IOA. It is a fully automated zero-touch mode that allows you to configure VLAN memberships. (Supported in CMC)

VLT mode

stack-unit *unit* iom-mode vlt

CONFIGURATION mode

```
Dell(conf)#stack-unit 0 iom-mode vlt
```

Select this mode to multi-home server interfaces to different IOA modules. This is a low-touch mode where all configuration except VLAN membership is automated. To enable VLAN, you must configure the VLANs at the server port level. In this mode, port 9 link, which is associated with LAG-127, is dedicated to VLT interconnect.

Programmable MUX mode

stack-unit *unit* iom-mode programmable-mux

CONFIGURATION mode

```
Dell(conf)#stack-unit 0 iom-mode programmable-mux
```

Select this mode to configure PMUX mode CLI commands.

For more information on the PMUX mode, refer to [PMUX Mode of the IO Aggregator](#).

Stacking mode

stack-unit *unit* **iom-mode** **stack**

CONFIGURATION mode

```
Dell(conf)#stack-unit 0 iom-mode stack
```

Select this mode to configure Stacking mode CLI commands.


For more information on the Stacking mode, refer to [Stacking](#).

Default Settings

The I/O Aggregator provides zero-touch configuration with the following default configuration settings:

- default user name (**root**)
- password (**calvin**)
- VLAN (vlan1) and IP address for in-band management (**DHCP**)
- IP address for out-of-band (OOB) management (**DHCP**)
- read-only SNMP community name (**public**)
- broadcast storm control (**enabled in Standalone mode and disabled in VLT mode**)
- IGMP multicast flooding (**enabled**)
- VLAN configuration (**in Standalone mode, all ports belong to all VLANs**)

You can change any of these default settings using the CLI. Refer to the appropriate chapter for details.

 **NOTE:** You can also change many of the default settings using the chassis management controller (CMC) interface. For information about how to access the CMC to configure the aggregator, refer to the *Dell Chassis Management Controller (CMC) User's Guide* on the Dell Support website at <http://support.dell.com/>

Other Auto-Configured Settings

After the Aggregator powers on, it auto-configures and is operational with software features enabled, including:

- Ports: Ports are administratively up and auto-configured to operate as hybrid ports to transmit tagged and untagged VLAN traffic.
 - Ports from 1 to 8 are internal server-facing ports.
 - Ports from 9 to 12 are external ports.

For more information about how ports are numbered, refer to Port Numbering.

- Link aggregation: All uplink ports are configured in a single LAG (LAG 128).
- VLANs: All ports are configured as members of all (4094) VLANs. All VLANs are up and can send or receive layer 2 traffic. For more information, refer to VLAN Membership.
- Data center bridging capability exchange protocol (DCBx): Server-facing ports auto-configure in auto-downstream port roles; uplink ports auto-configure in auto-upstream port roles.

- Fibre Channel over Ethernet (FCoE) connectivity and FCoE initiation protocol (FIP) snooping: The uplink port channel (LAG 128) is enabled to operate in Fibre channel forwarder (FCF) port mode.
- Link layer discovery protocol (LLDP): Enabled on all ports to advertise management TLV and system name with neighboring devices.
- Internet small computer system interface (iSCSI) optimization.
- Internet group management protocol (IGMP) snooping.
- Jumbo frames: Ports are set to a maximum MTU of 12,000 bytes by default.
- Link tracking: Uplink-state group 1 is automatically configured. In uplink state-group 1, server-facing ports auto-configure as downstream interfaces; the uplink port-channel (LAG 128) auto-configures as an upstream interface. Server-facing links are auto-configured to be brought up only if the uplink port-channel is up.
- In VLT mode, port 9 is automatically configured as VLT interconnect ports. VLT domain configuration is automatic. This includes peer-link, configured MAC, backup link and setting every port channel as VLT port-channel.

Data Center Bridging Support

To eliminate packet loss and provision links with required bandwidth, Data Center Bridging (DCB) enhancements for data center networks are supported.

The aggregator provides zero-touch configuration for DCB. The aggregator auto-configures DCBX port roles as follows:

- Server-facing ports are configured as auto-downstream interfaces.
- Uplink ports are configured as auto-upstream interfaces.

In operation, DCBx auto-configures uplink ports to match the DCB configuration in the ToR switches to which they connect.

The Aggregator supports DCB only in standalone mode.

FCoE Connectivity and FIP Snooping

Many data centers use Fiber Channel (FC) in storage area networks (SANs). Fiber Channel over Ethernet (FCoE) encapsulates Fiber Channel frames over Ethernet networks.

On an Aggregator, the internal ports support FCoE connectivity and connects to the converged network adapter (CNA) in servers. FCoE allows Fiber Channel to use 10-Gigabit Ethernet networks while preserving the Fiber Channel protocol.

The Aggregator also provides zero-touch configuration for FCoE connectivity. The Aggregator auto-configures to match the FCoE settings used in the switches to which it connects through its uplink ports.

FIP snooping is automatically configured on an Aggregator. The auto-configured port channel (LAG 128) operates in FCF port mode.

iSCSI Operation

Support for iSCSI traffic is turned on by default when the Aggregator powers up. No configuration is required.

When an aggregator powers up, it monitors known TCP ports for iSCSI storage devices on all interfaces. When a session is detected, an entry is created and monitored as long as the session is active.

The Aggregator also detects iSCSI storage devices on all interfaces and autoconfigures to optimize performance. Performance optimization operations, such as Jumbo frame size support and disabling storm control on interfaces connected to an iSCSI equallogic (EQL) storage device, are applied automatically.


Link Aggregation

All uplink ports are configured in a single LAG (LAG 128). Server-facing ports are auto-configured as part of link aggregation groups if the corresponding server is configured for LACP-based network interface controller (NIC) teaming. Static LAGs are not supported.

 **NOTE:** The recommended LACP timeout is Long-Timeout mode.

Link Tracking

By default, all server-facing ports are tracked by the operational status of the uplink LAG. If the uplink LAG goes down, the aggregator loses its connectivity and is no longer operational; all server-facing ports are brought down after the specified defer-timer interval, which is 10 seconds by default. If you have configured VLAN, you can reduce the defer time by changing the defer-timer value or remove it by using the `no defer-timer` command.

 **NOTE:** If installed servers do not have connectivity to a switch, check the Link Status LED of uplink ports on the aggregator. If all LEDs are on, to ensure the LACP is correctly configured, check the LACP configuration on the ToR switch that is connected to the aggregator .

Configuring VLANs

By default, in Standalone mode, all aggregator ports belong to all 4094 VLANs and are members of untagged VLAN 1. To configure only the required VLANs on a port, use the CLI or CMC interface.

You can configure VLANs only on server ports. The uplink LAG will automatically get the VLANs, based on the server ports VLAN configuration.

When you configure VLANs on server-facing interfaces (ports from 1 to 8), you can assign VLANs to a port or a range of ports by entering the `vlan tagged` or `vlan untagged` commands in Interface Configuration mode; for example:

```
Dell(conf)# interface range tengigabitethernet 0/2 - 4
Dell(conf-if-range-te-0/2-4)# vlan tagged 5,7,10-12
Dell(conf-if-range-te-0/2-4)# vlan untagged 3
```

Uplink LAG


The tagged VLAN membership of the uplink LAG is automatically configured based on the VLAN configuration of all server-facing ports (ports from 1 to 8).

The untagged VLAN used for the uplink LAG is always the default VLAN.

Server-Facing LAGs

The tagged VLAN membership of a server-facing LAG is automatically configured based on the server-facing ports that are members of the LAG.

The untagged VLAN of a server-facing LAG is configured based on the untagged VLAN to which the lowest numbered server-facing port in the LAG belongs.

 **NOTE:** Dell Networking recommends configuring the same VLAN membership on all LAG member ports.

Where to Go From Here

You can customize the Aggregator for use in your data center network as necessary. To perform additional switch configuration, do one of the following:

- For remote out-of-band management, enter the OOB management interface IP address into a Telnet or SSH client and log in to the switch using the user ID and password to access the CLI.
- For local management using the CLI, use the attached console connection.
- For remote in-band management from a network management station, enter the IP address of the default VLAN and log in to the switch to access the CLI.

In case of a Dell upgrade, you can check to see that an Aggregator is running the latest Dell version by entering the `show version` command. To download Dell version, go to <http://support.dell.com>


For detailed information about how to reconfigure specific software settings, refer to the appropriate chapter.

Configuration Fundamentals

The Dell Networking Operating System (OS) command line interface (CLI) is a text-based interface you can use to configure interfaces and protocols.

The CLI is structured in modes for security and management purposes. Different sets of commands are available in each mode, and you can limit user access to modes using privilege levels.

In Dell Networking OS, after you enable a command, it is entered into the running configuration file. You can view the current configuration for the whole system or for a particular CLI mode. To save the current configuration, copy the running configuration to another location. For more information, refer to Save the Running-Configuration.

 **NOTE:** You can use the chassis management controller (CMC) out-of-band management interface to access and manage an Aggregator using the Dell Networking OS command-line reference. For more information about how to access the CMC to configure an Aggregator, refer to the Dell Chassis Management Controller (CMC) User's Guide on the Dell Support website at <http://support.dell.com/support/edocs/systems/pem/en/index.htm>.

Accessing the Command Line

Access the command line through a serial console port or a Telnet session (Logging into the System using Telnet). When the system successfully boots, enter the command line in EXEC mode.

Logging into the System using Telnet

```
telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: username
Password:
Dell>
```

CLI Modes

Different sets of commands are available in each mode.

A command found in one mode cannot be executed from another mode (except for EXEC mode commands with a preceding `do` command (refer to the `do` Command section)).


The Dell Networking OS CLI is divided into three major mode levels:

- EXEC mode is the default mode and has a privilege level of 1, which is the most restricted level. Only a limited selection of commands is available, notably the `show` commands, which allow you to view system information.

- EXEC Privilege mode has commands to view configurations, clear counters, manage configuration files, run diagnostics, and enable or disable debug operations. The privilege level is 15, which is unrestricted. You can configure a password for this mode.
- CONFIGURATION mode allows you to configure security features, time settings, set logging and SNMP functions, and set line cards on the system.

Beneath CONFIGURATION mode are submodes that apply to interfaces, protocols, and features. The following example shows the submode command structure. Two sub-CONFIGURATION modes are important when configuring the chassis for the first time:

- INTERFACE submode is the mode in which you configure Layer 2 protocols and IP services specific to an interface. An interface can be physical (10 Gigabit Ethernet) or logical (port channel, or virtual local area network [VLAN]).
- LINE submode is the mode in which you to configure the console and virtual terminal lines.

 **NOTE:** At any time, entering a question mark (?) displays the available command options. For example, when you are in CONFIGURATION mode, entering the question mark first lists all available commands, including the possible submodes.

The CLI modes are:

```
EXEC
  EXEC Privilege
  CONFIGURATION
    INTERFACE
      10 GIGABIT ETHERNET
      INTERFACE RANGE
      MANAGEMENT ETHERNET
    LINE
      CONSOLE
      VIRTUAL TERMINAL
      MONITOR SESSION
```

Navigating CLI Modes

The Dell prompt changes to indicate the CLI mode.

The following table lists the CLI mode, its prompt, and information about how to access and exit the CLI mode. Move linearly through the command modes, except for the `end` command which takes you directly to EXEC Privilege mode and the `exit` command which moves you up one command mode level.



 **NOTE:** Sub-CONFIGURATION modes all have the letters “conf” in the prompt with more modifiers to identify the mode and slot/port information.

Table 1. Dell Command Modes

CLI Command Mode	Prompt	Access Command
EXEC	Dell>	Access the router through the console or Telnet.
EXEC Privilege	Dell#	<ul style="list-style-type: none"> • From EXEC mode, enter the <code>enable</code> command. • From any other mode, use the <code>end</code> command.

CLI Command Mode	Prompt	Access Command
CONFIGURATION	Dell (conf) #	<ul style="list-style-type: none"> From EXEC privilege mode, enter the <code>configure</code> command. From every mode except EXEC and EXEC Privilege, enter the <code>exit</code> command.
10 Gigabit Ethernet Interface	Dell (conf-if-te-0/1) #	<code>interface</code> (INTERFACE modes)
Interface Range	Dell (conf-if-range) #	<code>interface</code> (INTERFACE modes)
Management Ethernet Interface	Dell (conf-if-ma-0/0) #	<code>interface</code> (INTERFACE modes)
MONITOR SESSION	Dell (conf-mon-sess) #	<code>monitor session</code>
CONSOLE	Dell (config-line-console) #	<code>line</code> (LINE Modes)
VIRTUAL TERMINAL	Dell (config-line-vty) #	<code>line</code> (LINE Modes)

 **NOTE:** Access all of the following modes from CONFIGURATION mode.

The following example shows how to change the command mode from CONFIGURATION mode to INTERFACE configuration mode.

Example of Changing Command Modes

```
Dell(conf)#interface tengigabitethernet 0/2
Dell(conf-if-te-0/2)#
```

The do Command

You can enter an EXEC mode command from any CONFIGURATION mode (CONFIGURATION, INTERFACE, and so on.) without having to return to EXEC mode by preceding the EXEC mode command with the `do` command.

The following example shows the output of the `do` command.

```
Dell(conf)#do show system brief

Stack MAC : 00:1e:c9:de:03:7b


-- Stack Info --
Unit  UnitType      Status      ReqTyp      CurTyp      Version
-----
-----
0      Management      online      PE-FN-410S-IOA  PE-FN-410S-IOA  1-0 (0-1864)
12
1      Member          not present
2      Member          not present
3      Member          not present
4      Member          not present
5      Member          not present
```

Dell#

Undoing Commands

When you enter a command, the command line is added to the running configuration file (running-config).

To disable a command and remove it from the running-config, enter the `no` command, then the original command. For example, to delete an IP address configured on an interface, use the `no ip address ip-address` command.

 **NOTE:** Use the help or ? command as described in [Obtaining Help](#).

Example of Viewing Disabled Commands

```
Dell(conf)# interface managementethernet 0/0
Dell(conf-if-ma-0/0)# ip address 192.168.5.6/16
Dell(conf-if-ma-0/0)#
Dell(conf-if-ma-0/0)#
Dell(conf-if-ma-0/0)#show config
!
interface ManagementEthernet 0/0
ip address 192.168.5.6/16
no shutdown
Dell(conf-if-ma-0/0)#
Dell(conf-if-ma-0/0)# no ip address
Dell(conf-if-ma-0/0)#
Dell(conf-if-ma-0/0)# show config
!
interface ManagementEthernet 0/0
  no ip address
  no shutdown
Dell(conf-if-ma-0/0)#
```

Obtaining Help

Obtain a list of keywords and a brief functional description of those keywords at any CLI mode using the ? or help command:

- To list the keywords available in the current mode, enter ? at the prompt or after a keyword.
- Enter ? after a prompt lists all of the available keywords. The output of this command is the same for the help command.

```
Dell#?
start          Start Shell
capture        Capture Packet
cd             Change current directory
clear          Reset functions
clock          Manage the system clock
configure      Configuring from terminal
copy           Copy from one file to another
--More--
```

- Enter ? after a partial keyword lists all of the keywords that begin with the specified letters.

```
Dell(conf)#cl?
clock
Dell(conf)#cl
```

- Enter [space]? after a keyword lists all of the keywords that can follow the specified keyword.

```
Dell (conf) #clock ?
summer-time      Configure summer (daylight savings) time
timezone         Configure time zone
Dell (conf) #clock
```

Entering and Editing Commands

Notes for entering commands.

- The CLI is not case-sensitive.
- You can enter partial CLI keywords.
 - Enter the minimum number of letters to uniquely identify a command. For example, you cannot enter `cl` as a partial keyword because both the `clock` and `class-map` commands begin with the letters “cl.” You can enter `clo`, however, as a partial keyword because only one command begins with those three letters.
- The TAB key auto-completes keywords in commands. Enter the minimum number of letters to uniquely identify a command.
- The UP and DOWN arrow keys display previously entered commands (refer to [Command History](#)).
- The BACKSPACE and DELETE keys erase the previous letter.
- Key combinations are available to move quickly across the command line. The following table describes these short-cut key combinations.

Short-Cut Key Combination	Action
CNTL-A	Moves the cursor to the beginning of the command line.
CNTL-B	Moves the cursor back one character.
CNTL-D	Deletes character at cursor.
CNTL-E	Moves the cursor to the end of the line.
CNTL-F	Moves the cursor forward one character.
CNTL-I	Completes a keyword.
CNTL-K	Deletes all characters from the cursor to the end of the command line.
CNTL-L	Re-enters the previous command.
CNTL-N	Return to more recent commands in the history buffer after recalling commands with CTRL-P or the UP arrow key.
CNTL-P	Recalls commands, beginning with the last command.
CNTL-U	Deletes the line.
CNTL-W	Deletes the previous word.
CNTL-X	Deletes the line.
CNTL-Z	Ends continuous scrolling of command outputs.
Esc B	Moves the cursor back one word.
Esc F	Moves the cursor forward one word.

Short-Cut Key Combination	Action
Esc D	Deletes all characters from the cursor to the end of the word.

Command History

Dell Networking OS maintains a history of previously-entered commands for each mode. For example:

- When you are in EXEC mode, the UP and DOWN arrow keys display the previously-entered EXEC mode commands.
- When you are in CONFIGURATION mode, the UP or DOWN arrows keys recall the previously-entered CONFIGURATION mode commands.

Filtering show Command Outputs

Filter the output of a `show` command to display specific information by adding `| [except | find | grep | no-more | save] specified_text` after the command.


The variable `specified_text` is the text for which you are filtering and it IS case sensitive unless you use the `ignore-case` sub-option.

Starting with Dell Networking OS version 7.8.1.0, the `grep` command accepts an `ignore-case` sub-option that forces the search to case-insensitive. For example, the commands:

- `show run | grep Ethernet` returns a search result with instances containing a capitalized "Ethernet," such as `interface TenGigabitEthernet 0/1`.
- `show run | grep ethernet` does not return that search result because it only searches for instances containing a non-capitalized "ethernet."
- `show run | grep Ethernet ignore-case` returns instances containing both "Ethernet" and "ethernet."

The `grep` command displays only the lines containing specified text. The following example shows this command used in combination with the `show linecard all` command.

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | grep 0
stack unit 0 stack-port all
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
```

 **NOTE:** Dell accepts a space or no space before and after the pipe. To filter a phrase with spaces, underscores, or ranges, enclose the phrase with double quotation marks.

The `except` keyword displays text that does not match the specified text. The following example shows this command used in combination with the `show linecard all` command.

Example of the `except` Keyword

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | except 0
Admin mode is On
```

```
Admin is enabled
Local is enabled
Link Delay 65535 pause quantum
Dell(conf)#
```


The `find` keyword displays the output of the `show` command beginning from the first occurrence of specified text. The following example shows this command used in combination with the `show linecard all` command.

Example of the `find` Keyword

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | find 0
stack unit 0 stack-port all
Admin mode is On
Admin is enabled
Local is enabled
Link Delay 65535 pause quantum
0 Pause Tx pkts, 0 Pause Rx pkts
Dell(conf)#
```

The `no-more` command displays the output all at once rather than one screen at a time. This is similar to the `terminal length` command except that the `no-more` option affects the output of the specified command only.

The `save` command copies the output to a file for future reference.

 **NOTE:** You can filter a single command output multiple times. The `save` option must be the last option entered. For example: `Dell# command | grep regular-expression | except regular-expression | grep other-regular-expression | find regular-expression | save.`

Multiple Users in Configuration Mode

Dell notifies all users when there are multiple users logged in to CONFIGURATION mode.

A warning message indicates the username, type of connection (console or VTY), and in the case of a VTY connection, the IP address of the terminal on which the connection was established. For example:

- On the system that telnets into the switch, this message appears:
% Warning: The following users are currently configuring the system:
User "<username>" on line console0
- On the system that is connected over the console, this message appears:
% Warning: User "<username>" on line vty0 "10.11.130.2" is in configuration mode

If either of these messages appears, Dell Networking recommends coordinating with the users listed in the message so that you do not unintentionally overwrite each other's configuration changes.

Data Center Bridging (DCB)

On an I/O Aggregator, data center bridging (DCB) features are auto-configured in standalone mode. You can display information on DCB operation by using **show** commands.

 **NOTE:** DCB features are not supported on an Aggregator in stacking mode.

Supported Modes

Standalone, Stacking, PMUX, VLT

Ethernet Enhancements in Data Center Bridging

DCB refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic. Through network consolidation, DCB results in reduced operational cost, simplified management, and easy scalability by avoiding the need to deploy separate application-specific networks. For example, instead of deploying an Ethernet network for LAN traffic, additional storage area networks (SANs) to ensure lossless fibre-channel traffic, and a separate InfiniBand network for high-performance inter-processor computing within server clusters, only one DCB-enabled network is required in a data center. The Dell Networking switches that support a unified fabric and consolidate multiple network infrastructures use a single input/output (I/O) device called a converged network adapter (CNA).


A CNA is a computer input/output device that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC). Multiple adapters on different devices for several traffic types are no longer required.

Data center bridging satisfies the needs of the following types of data center traffic in a unified fabric:

- LAN traffic consists of a large number of flows that are generally insensitive to latency requirements, while certain applications, such as streaming video, are more sensitive to latency. Ethernet functions as a best-effort network that may drop packets in case of network congestion. IP networks rely on transport protocols (for example, TCP) for reliable data transmission with the associated cost of greater processing overhead and performance impact.
- Storage traffic based on Fibre Channel media uses the SCSI protocol for data transfer. This traffic typically consists of large data packets with a payload of 2K bytes that cannot recover from frame loss. To successfully transport storage traffic, data center Ethernet must provide no-drop service with lossless links.
- Servers use InterProcess Communication (IPC) traffic within high-performance computing clusters to share information. Server traffic is extremely sensitive to latency requirements.

To ensure lossless delivery and latency-sensitive scheduling of storage and service traffic and I/O convergence of LAN, storage, and server traffic over a unified fabric, IEEE data center bridging adds the following extensions to a classical Ethernet network:

- 802.1Qbb - Priority-based Flow Control (PFC)
- 802.1Qaz - Enhanced Transmission Selection (ETS)
- 802.1Qau - Congestion Notification
- Data Center Bridging Exchange (DCBx) protocol

 **NOTE:** In Dell Networking OS version 9.4.0.x, only the PFC, ETS, and DCBx features are supported in data center bridging.

Priority-Based Flow Control

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion.

When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device. In this way, PFC ensures that large amounts of queued LAN traffic do not cause storage traffic to be dropped, and that storage traffic does not result in high latency for high-performance computing (HPC) traffic between servers.

PFC enhances the existing 802.3x pause and 802.1p priority capabilities to enable flow control based on 802.1p priorities (classes of service). Instead of stopping all traffic on a link (as performed by the traditional Ethernet pause mechanism), PFC pauses traffic on a link according to the 802.1p priority set on a traffic type. You can create lossless flows for storage and server traffic while allowing for loss in case of LAN traffic congestion on the same physical interface.

The following illustration shows how PFC handles traffic congestion by pausing the transmission of incoming traffic with dot1p priority 3.




Figure 1. Priority-Based Flow Control

In the system, PFC is implemented as follows:

- PFC is supported on specified 802.1p priority traffic (dot1p 0 to 7) and is configured per interface. However, only two lossless queues are supported on an interface: one for Fibre Channel over Ethernet (FCoE) converged traffic and one for Internet Small Computer System Interface (iSCSI) storage traffic. Configure the same lossless queues on all ports.
- A dynamic threshold handles intermittent traffic bursts and varies based on the number of PFC priorities contending for buffers, while a static threshold places an upper limit on the transmit time of a queue after receiving a message to pause a specified priority. PFC traffic is paused only after surpassing both static and dynamic thresholds for the priority specified for the port.
- By default, PFC is enabled when you enabled DCB. When you enable DCB globally, you cannot simultaneously enable TX and RX on the interface for flow control and link-level flow control is disabled.

- Buffer space is allocated and de-allocated only when you configure a PFC priority on the port.
- PFC delay constraints place an upper limit on the transmit time of a queue after receiving a message to pause a specified priority.
- By default, PFC is enabled on an interface with no dot1p priorities configured. You can configure the PFC priorities if the switch negotiates with a remote peer using DCBX. During DCBX negotiation with a remote peer:
 - DCBx communicates with the remote peer by link layer discovery protocol (LLDP) type, length, value (TLV) to determine current policies, such as PFC support and enhanced transmission selection (ETS) BW allocation.
 - If the negotiation succeeds and the port is in DCBX Willing mode to receive a peer configuration, PFC parameters from the peer are used to configured PFC priorities on the port. If you enable the link-level flow control mechanism on the interface, DCBX negotiation with a peer is not performed.
 - If the negotiation fails and PFC is enabled on the port, any user-configured PFC input policies are applied. If no PFC dcb-map has been previously applied, the PFC default setting is used (no priorities configured). If you do not enable PFC on an interface, you can enable the 802.3x link-level pause function. By default, the link-level pause is disabled, when you disable DCBx and PFC. If no PFC dcb-map has been applied on the interface, the default PFC settings are used.
- PFC supports buffering to receive data that continues to arrive on an interface while the remote system reacts to the PFC operation.
- PFC uses the DCB MIB IEEE802.1azd2.5 and the PFC MIB IEEE802.1bb-d2.2.

If DCBx negotiation is not successful (for example, due to a version or TLV mismatch), DCBx is disabled and you cannot enable PFC or ETS.

Configuring Priority-Based Flow Control

PFC provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default when you enable DCB.

As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (Class of Service (CoS) values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that is to be stopped. Data Center Bridging Exchange protocol (DCBx) provides the link-level exchange of PFC parameters between peer devices. PFC allows network administrators to create zero-loss links for Storage Area Network (SAN) traffic that requires no-drop service, while retaining packet-drop congestion management for Local Area Network (LAN) traffic.

To ensure complete no-drop service, apply the same dcb-map on all PFC and ETS enabled interfaces.

1. Create a DCB map to apply priority based flow control or enhanced transmission selection for specified priority groups and priorities.

CONFIGURATION mode

```
dcb-map map-name
```

The maximum is 32 alphanumeric characters.

2. Configure the priority group with PGID, bandwidth percentage or strict priority for ETS and PFC mode.

DCB-MAP mode

```
priority-group pg_num [bandwidth percentage | strict-priority] pfc [on | off]
```

pg_num range is from 0 to 7.

bandwidth percentage range is from 1 to 100.

Either strict-priority or bandwidth percentage can be set for ETS on the priority group.

PFC can either be enabled or disabled for the priority group.

3. Configure the priorities to priority group.

DCB-MAP mode

```
priority-pgid <pgid> <pgid> <pgid> <pgid> <pgid> <pgid> <pgid> <pgid>
```

pgid range is from 0 to 7.

Configure priority to priority group mapping from priority 0 to priority 7 in order.

4. Exit the DCB MAP configuration mode.

DCB-MAP mode

```
exit
```

5. Enter interface configuration mode.

CONFIGURATION mode

```
interface type slot/port
```

6. Apply the dcb-map with PFC and ETS configurations to both ingress and egress interfaces.

INTERFACE mode

```
dcb-map map-name
```

7. Repeat steps 1 to 6 on all PFC and ETS enabled interfaces to ensure lossless traffic service.



NOTE: All these configurations are available only in PMUX mode and you cannot perform these configurations in Standalone mode.

Dell Networking OS Behavior: As soon as you apply a DCB MAP with PFC enabled on an interface, DCBx starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE, and CIN versions of PFC Type, Length, Value (TLV) are supported. DCBx also validates PFC configurations that are received in TLVs from peer devices.

By applying a DCB MAP with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues.

To remove a DCB MAP, including the PFC and ETS configurations it contains, use the `no dcb-map map-name` command in INTERFACE Configuration mode.

You can enable any number of 802.1p priorities for PFC. Queues to which PFC priority traffic is mapped are lossless by default. Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure the lossless queues for no-drop priorities in a PFC dcb-map and reapply the policy to an interface.

To apply PFC, a PFC peer must support the configured priority traffic (as detected by DCBx).

To honor a PFC pause frame multiplied by the number of PFC-enabled ingress ports, the minimum link delay must be greater than the round-trip transmission time the peer requires.

If you apply `dcb-map` with PFC disabled (`no pfc mode on`):

- You can enable link-level flow control on the interface. To delete the `dcb-map`, first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is by default PFC-enabled.
- PFC still allows you to configure lossless queues on a port to ensure no-drop handling of lossless traffic.


 **NOTE:** You cannot enable PFC and link-level flow control at the same time on an interface.

When you apply a `dcb-map` to an interface, an error message displays if:

- The PFC dot1p priorities result in more than two lossless port queues globally on the switch.
- Link-level flow control is already enabled. You cannot enable PFC and link-level flow control at the same time on an interface.
- In a switch stack, configure all stacked ports with the same PFC configuration.

A DCB MAP for PFC applied to an interface may become invalid if you reconfigure dot1p-queue mapping. This situation occurs when the new dot1p-queue assignment exceeds the maximum number (2) of lossless queues supported globally on the switch. In this case, all PFC configurations received from PFC-enabled peers are removed and resynchronized with the peer devices.

Traffic may be interrupted when you reconfigure PFC no-drop priorities in a `dcb-map` or reapply the `dcb-map` to an interface.

 **NOTE:** All these configurations are available only in PMUX mode and you cannot perform these configurations in Standalone mode.

Enhanced Transmission Selection

Enhanced transmission selection (ETS) supports optimized bandwidth allocation between traffic types in multiprotocol (Ethernet, FCoE, SCSI) links.

ETS allows you to divide traffic according to its 802.1p priority into different priority groups (traffic classes) and configure bandwidth allocation and queue scheduling for each group to ensure that each traffic type is correctly prioritized and receives its required bandwidth. For example, you can prioritize low-latency storage or server cluster traffic in a traffic class to receive more bandwidth and restrict best-effort LAN traffic assigned to a different traffic class.

Although you can configure strict-priority queue scheduling for a priority group, ETS introduces flexibility that allows the bandwidth allocated to each priority group to be dynamically managed according to the amount of LAN, storage, and server traffic in a flow. Unused bandwidth is dynamically allocated to prioritized priority groups. Traffic is queued according to its 802.1p priority assignment, while flexible bandwidth allocation and the configured queue-scheduling for a priority group is supported.

The following figure shows how ETS allows you to allocate bandwidth when different traffic types are classed according to 802.1p priority and mapped to priority groups.




Figure 2. Enhanced Transmission Selection

The following table lists the traffic groupings ETS uses to select multiprotocol traffic for transmission.

Table 2. ETS Traffic Groupings

Traffic Groupings	Description
Priority group	A group of 802.1p priorities used for bandwidth allocation and queue scheduling. All 802.1p priority traffic in a group must have the same traffic handling requirements for latency and frame loss.
Group ID	A 4-bit identifier assigned to each priority group. The range is from 0 to 7.
Group bandwidth	Percentage of available bandwidth allocated to a priority group.
Group transmission selection algorithm (TSA)	Type of queue scheduling a priority group uses.

In the Dell Networking OS, ETS is implemented as follows:

- ETS supports groups of 802.1p priorities that have:
 - PFC enabled or disabled
 - No bandwidth limit or no ETS processing
- Bandwidth allocated by the ETS algorithm is made available after strict-priority groups are serviced. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups so that the sum of the bandwidth use is 100%. If priority group bandwidth use exceeds 100%, all configured priority group bandwidth is decremented based on the configured percentage ratio until all priority group bandwidth use is 100%. If priority group bandwidth usage is less than or equal to 100% and any default priority groups exist, a minimum of 1% bandwidth use is assigned by decreasing 1% of bandwidth from the other priority groups until priority group bandwidth use is 100%.
- For ETS traffic selection, an algorithm is applied to priority groups using:
 - Strict priority shaping

- ETS shaping
- (Credit-based shaping is not supported)
- ETS uses the DCB MIB IEEE 802.1azd2.5.

Configuring Enhanced Transmission Selection

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic.

Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.


For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link by:

- Allocating a guaranteed share of bandwidth to each priority group.
- Allowing each group to exceed its minimum guaranteed bandwidth if another group is not fully using its allotted bandwidth.

To configure ETS and apply an ETS dcb-map to an interface, you must follow the steps described in [Configuring Priority-Based Flow Control](#).

Creating a DCB Map

Configure the priority-based flow control (PFC) and enhanced traffic selection (ETS) settings in a DCB map before you can apply them on downstream server-facing ports.

 **NOTE:** This feature is supported only on PMUX mode.

Task	Command	Command Mode
Create a DCB map to specify the PFC and ETS settings for groups of dot1p priorities.	<code>dcb-map name</code>	CONFIGURATION
Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group or whether priority group traffic should be handled with strict priority scheduling.	<code>priority-group group_num {bandwidth percentage strict-priority} pfc {on off}</code>	DCB MAP
Specify the priority group ID number to handle VLAN traffic for each dot1p class-of-service: 0 through 7. Leave a space between each priority group number.	<code>priority-pgid dot1p0_group-num dot1p1_group-num dot1p2_group-num dot1p3_group-num dot1p4_group-num dot1p5_group-num dot1p6_group-num dot1p7_group-num</code>	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (the `no priority pgid` command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.
- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings and apply the new map to the interfaces to override the previous DCB map settings. Then delete the original dot1p priority-priority group mapping.

Applying a DCB Map on Server-Facing Ethernet Ports

You can apply a DCB map only on a physical Ethernet interface and can apply only one DCB map per interface.

Task	Command	Command Mode
Enter Interface Configuration mode on a server-facing port to apply a DCB map.	<pre>interface {tengigabitEthernet slot/ port fortygigabitEthernet slot/port}</pre>	CONFIGURATION
Apply the DCB map on an Ethernet port. Repeat this step to apply a DCB map to more than one port.	<pre>dcb-map name</pre>	INTERFACE

Data Center Bridging Exchange Protocol (DCBx)

The data center bridging exchange (DCBx) protocol is disabled by default on any switch on which PFC or ETS are enabled.

DCBx allows a switch to automatically discover DCB-enabled peers and exchange configuration information. PFC and ETS use DCBx to exchange and negotiate parameters with peer devices. DCBx capabilities include:

- Discovery of DCB capabilities on peer-device connections.
- Determination of possible mismatch in DCB configuration on a peer link.
- Configuration of a peer device over a DCB link.

DCBx requires the link layer discovery protocol (LLDP) to provide the path to exchange DCB parameters with peer devices. Exchanged parameters are sent in organizationally specific TLVs in LLDP data units. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#). The following LLDP TLVs are supported for DCB parameter exchange:

PFC parameters	PFC Configuration TLV and Application Priority Configuration TLV.
ETS parameters	ETS Configuration TLV and ETS Recommendation TLV.

Data Center Bridging in a Traffic Flow

The following figure shows how DCB handles a traffic flow on an interface.

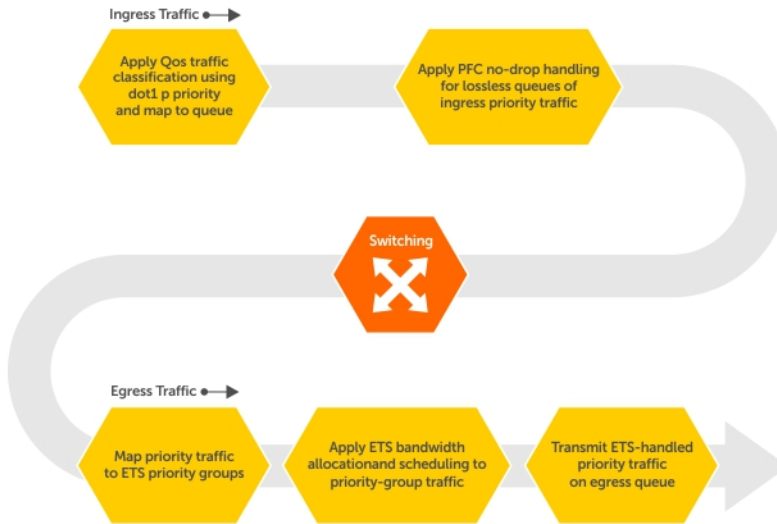


Figure 3. DCB PFC and ETS Traffic Handling

Data Center Bridging: Auto-DCB-Enable Mode

On an Aggregator in standalone or VLT modes, the default mode of operation for data center bridging on Ethernet ports is auto-DCB-enable mode. In this mode, Aggregator ports detect whether peer devices support CEE or not, and enable ETS and PFC or link-level flow control accordingly:

- Interfaces come up with DCB disabled and link-level flow control enabled to control data transmission between the Aggregator and other network devices (see Flow Control Using Ethernet Pause Frames). When DCB is disabled on an interface, PFC, and ETS are also disabled.
- When DCBx protocol packets are received, interfaces automatically enable DCB and disable link-level flow control.

DCB is required for PFC, ETS, DCBx, and FCoE initialization protocol (FIP) snooping to operate.

NOTE: Normally, interfaces do not flap when DCB is automatically enabled.

DCB processes VLAN-tagged packets and dot1p priority values. Untagged packets are treated with a dot1p priority of 0.

For DCB to operate effectively, ingress traffic is classified according to its dot1p priority so that it maps to different data queues. The dot1p-queue assignments used on an Aggregator are shown in Table 6-1 in dcb enable auto-detect on-next-reload Command Example QoS dot1p Traffic Classification and Queue Assignment.

When DCB is Disabled (Default) By default, Aggregator interfaces operate with DCB disabled and link-level flow control enabled. When an interface comes up, it is automatically configured with:

- Flow control enabled on input interfaces.
- A DCB-MAP policy is applied with PFC disabled.

The following example shows a default interface configuration with DCB disabled and link-level flow control enabled.

show interfaces Command Example: DCB disabled and Flow Control enabled

```
Dell#show running-config interface te 0/4
!
interface TenGigabitEthernet 0/4
  mtu 12000
  portmode hybrid
  switchport
  auto vlan
  flowcontrol rx on tx off
  dcb-map DCB_MAP_PFC_OFF
!
  protocol lldp
    advertise management-tlv management-address system-name
    dcbx port-role auto-downstream
  no shutdown
Dell#
```

When DCB is Enabled When an interface receives a DCBx protocol packet, it automatically enables DCB and disables link-level flow control. The dcb-map and flow control configurations are removed as shown in the following example.

show interfaces Command Example: DCB enabled and Flow Control disabled

```
Dell#show running-config interface te 0/3
!
interface TenGigabitEthernet 0/3
  mtu 12000
  portmode hybrid
  switchport
  auto vlan
!
  protocol lldp
    advertise management-tlv management-address system-name
    dcbx port-role auto-downstream
  no shutdown
Dell#
```

When no DCBx TLVs are received on a DCB-enabled interface for 180 seconds, DCB is automatically disabled and flow control is re-enabled.

Lossless Traffic Handling In auto-DCB-enable mode, Aggregator ports operate with the auto-detection of DCBx traffic. At any moment, some ports may operate with link-level flow control while others operate with DCB-based PFC enabled.

As a result, lossless traffic is ensured only if traffic ingresses on a PFC-enabled port and egresses on another PFC-enabled port.

Lossless traffic is not guaranteed when it is transmitted on a PFC-enabled port and received on a link-level flow control-enabled port, or transmitted on a link-level flow control-enabled port and received on a PFC-enabled port.

Enabling DCB on Next Reload To configure the Aggregator so that all interfaces come up with DCB enabled and flow control disabled, use the `dcb enable on-next-reload` command. Internal PFC buffers are automatically configured.

Task	Command	Command Mode
Globally enable DCB on all interfaces after next switch reload.	<code>dcb enable on-next-reload</code>	CONFIGURATION

To reconfigure the Aggregator so that all interfaces come up with DCB disabled and link-level flow control enabled, use the `no dcb enable on-next-reload` command. PFC buffer memory is automatically freed.

Enabling Auto-DCB-Enable Mode on Next Reload To configure the Aggregator so that all interfaces come up in auto-DCB-enable mode with DCB disabled and flow control enabled, use the `dcb enable auto-detect on-next-reload` command.

Task	Command	Command Mode
Globally enable auto-detection of DCBx and auto-enabling of DCB on all interfaces after switch reload.	<code>dcb enable auto-detect on-next-reload</code>	CONFIGURATION

Enabling DCB To configure the Aggregator so that all interfaces are DCB enabled and flow control disabled, use the `dcb enable` command.

Disabling DCB To configure the Aggregator so that all interfaces are DCB disabled and flow control enabled, use the `no dcb enable` command.

dcb enable auto-detect on-next-reload Command Example

```
Dell#dcb enable auto-detect on-next-reload
```

QoS dot1p Traffic Classification and Queue Assignment

DCB supports PFC, ETS, and DCBx to handle converged Ethernet traffic that is assigned to an egress queue according to the following QoS methods:

- Honor dot1p** dot1p priorities in ingress traffic are used at the port or global switch level.
- Layer 2 class maps** dot1p priorities are used to classify traffic in a class map and apply a service policy to an ingress port to map traffic to egress queues.



NOTE: Dell Networking does not recommend mapping all ingress traffic to a single queue when using PFC and ETS. However, Dell Networking does recommend using Ingress traffic classification using the `service-class dynamic dot1p` command (honor dot1p) on all DCB-enabled interfaces. If you use L2 class maps to map dot1p priority traffic to egress queues, take into account the default dot1p-queue assignments in the following table and the maximum number of two lossless queues supported on a port.

Although the system allows you to change the default dot1p priority-queue assignments, DCB policies applied to an interface may become invalid if you reconfigure dot1p-queue mapping. If the configured dcb-map remains valid, the change in the dot1p-queue assignment is allowed. For DCB ETS enabled interfaces, traffic destined to queue that is not mapped to any dot1p priority are dropped.

dot1p Value in the Incoming Frame	Egress Queue Assignment
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

How Priority-Based Flow Control is Implemented

Priority-based flow control provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default. As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (CoS values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that needs to be stopped. DCBx provides the link-level exchange of PFC parameters between peer devices. PFC creates zero-loss links for SAN traffic that requires no-drop service, while at the same time retaining packet-drop congestion management for LAN traffic.

PFC is implemented on an Aggregator as follows:

- If DCB is enabled, as soon as a dcb-map with PFC is applied on an interface, DCBx starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE and CIN versions of PFC TLV are supported. DCBx also validates PFC configurations received in TLVs from peer devices.
- To achieve complete lossless handling of traffic, enable PFC operation on ingress port traffic and on all DCB egress port traffic.
- All 802.1p priorities are enabled for PFC. Queues to which PFC priority traffic is mapped are lossless by default. Traffic may be interrupted due to an interface flap (going down and coming up).
- For PFC to be applied on an Aggregator port, the auto-configured priority traffic must be supported by a PFC peer (as detected by DCBx).
- A dcb-map for PFC applied to an interface may become invalid if dot1p-queue mapping is reconfigured. This situation occurs when the new dot1p-queue assignment exceeds the maximum


number (2) of lossless queues supported globally on the switch. In this case, all PFC configurations received from PFC-enabled peers are removed and re-synchronized with the peer devices.

- Dell Networking OS does not support MACsec Bypass Capability (MBC).

How Enhanced Transmission Selection is Implemented

Enhanced transmission selection (ETS) provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic. Different traffic types have different service needs. Using ETS, groups within an 802.1p priority class are auto-configured to provide different treatment for traffic with different bandwidth, latency, and best-effort needs.

For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link.

 **NOTE:** The IEEE 802.1Qaz, CEE, and CIN versions of ETS are supported.

ETS is implemented on an Aggregator as follows:

- Traffic in priority groups is assigned to strict-queue or WERR scheduling in a dcb-map and is managed using the ETS bandwidth-assignment algorithm. Dell Networking OS de-queues all frames of strict-priority traffic before servicing any other queues. A queue with strict-priority traffic can starve other queues in the same port.
- ETS-assigned bandwidth allocation and scheduling apply only to data queues, not to control queues.
- Dell Networking OS supports hierarchical scheduling on an interface. Dell Networking OS control traffic is redirected to control queues as higher priority traffic with strict priority scheduling. After control queues drain out, the remaining data traffic is scheduled to queues according to the bandwidth and scheduler configuration in the dcb-map. The available bandwidth calculated by the ETS algorithm is equal to the link bandwidth after scheduling non-ETS higher-priority traffic.
- By default, equal bandwidth is assigned to each port queue and each dot1p priority in a priority group.
- By default, equal bandwidth is assigned to each priority group in the dcb-map applied to an egress port. The sum of auto-configured bandwidth allocation to dot1p priority traffic in all ETS priority groups is 100%.
- dot1p priority traffic on the switch is scheduled according to the default dot1p-queue mapping. dot1p priorities within the same queue should have the same traffic properties and scheduling method.
- A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share the same latency and loss requirements. All 802.1p priorities mapped to the same queue should be in the same priority group.
 - By default:
 - * All 802.1p priorities are grouped in priority group 0.
 - * 100% of the port bandwidth is assigned to priority group 0. The complete bandwidth is equally assigned to each priority class so that each class has 12 to 13%.
 - The maximum number of priority groups supported in ETS output policies on an interface is equal to the number of data queues (4) on the port. The 802.1p priorities in a priority group can map to multiple queues.
- A dcb-map is created to associate a priority group with a dcb-map with scheduling and bandwidth configuration, and applied on egress ports.
 - The ETS configuration associated with 802.1p priority traffic in a dcb-map is used in DCBx negotiation with ETS peers.
 - When a dcb-map is applied to an interface, ETS-configured scheduling and bandwidth allocation take precedence over any auto-configured settings in the QoS output policies.

- ETS is enabled by default with the default ETS configuration applied (all dot1p priorities in the same group with equal bandwidth allocation).

ETS Operation with DCBx

In DCBx negotiation with peer ETS devices, ETS configuration is handled as follows:

- ETS TLVs are supported in DCBx versions CIN, CEE, and IEEE2.5.
- ETS operational parameters are determined by the DCBX port-role configurations.
- ETS configurations received from TLVs from a peer are validated.
- In case of a hardware limitation or TLV error, the DCBx operation on an ETS port goes down.
- ETS operates with legacy DCBx versions as follows:
 - In the CEE version, the priority group/traffic class group (TCG) ID 15 represents a non-ETS priority group. Any priority group configured with a scheduler type is treated as a strict-priority group and is given the priority-group (TCG) ID 15.
 - The CIN version supports two types of strict-priority scheduling:
 - * Group strict priority: Allows a single priority flow in a priority group to increase its bandwidth usage to the bandwidth total of the priority group. A single flow in a group can use all the bandwidth allocated to the group.
 - * Link strict priority: Allows a flow in any priority group to increase to the maximum link bandwidth.

CIN supports only the default dot1p priority-queue assignment in a priority group.

DCBx Operation

The data center bridging exchange protocol (DCBx) is used by DCB devices to exchange configuration information with directly connected peers using the link layer discovery protocol (LLDP) protocol. DCBx can detect the misconfiguration of a peer DCB device, and optionally, configure peer DCB devices with DCB feature settings to ensure consistent operation in a data center network.

DCBx is a prerequisite for using DCB features, such as priority-based flow control (PFC) and enhanced traffic selection (ETS), to exchange link-level configurations in a converged Ethernet environment. DCBx is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices are DCBx-enabled (DCBx is enabled end-to-end).

The following versions of DCBx are supported on an Aggregator: CIN, CEE, and IEEE2.5.

DCBx requires the LLDP to be enabled on all DCB devices.

DCBx Operation

DCBx performs the following operations:

- Discovers DCB configuration (such as PFC and ETS) in a peer device.
- Detects DCB mis-configuration in a peer device; that is, when DCB features are not compatibly configured on a peer device and the local switch. Mis-configuration detection is feature-specific because some DCB features support asymmetric configuration.
- Reconfigures a peer device with the DCB configuration from its configuration source if the peer device is willing to accept configuration.
- Accepts the DCB configuration from a peer if a DCBx port is in “willing” mode to accept a peer’s DCB settings and then internally propagates the received DCB configuration to its peer ports.

DCBx Port Roles

The following DCBx port roles are auto-configured on an Aggregator to propagate DCB configurations learned from peer DCBx devices internally to other switch ports:

Auto-upstream

The port advertises its own configuration to DCBx peers and receives its configuration from DCBx peers (ToR or FCF device). The port also propagates its configuration to other ports on the switch.

The first auto-upstream that is capable of receiving a peer configuration is elected as the *configuration source*. The elected configuration source then internally propagates the configuration to other auto-upstream and auto-downstream ports. A port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-upstream port (besides the configuration source) receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBx peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The configuration received from a DCBx peer or from an internally propagated configuration is not stored in the switch's running configuration.

On a DCBx port in an auto-upstream role, the PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Auto-downstream


The port advertises its own configuration to DCBx peers but is *not willing* to receive remote peer configuration. The port always accepts internally propagated configurations from a configuration source. An auto-downstream port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-downstream port receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBx peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The internally propagated configuration is not stored in the switch's running configuration. On a DCBx port in an auto-downstream role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

Default DCBx port role: Uplink ports are auto-configured in an auto-upstream role. Server-facing ports are auto-configured in an auto-downstream role.

 **NOTE:** You can change the port roles only in the PMUX mode. Use the following command to change the port roles:

```
dcbx port-role {auto-downstream | auto-upstream | config-source | manual}
```

`manual` is the default port role.

 **NOTE:** On a DCBx port, application priority TLV advertisements are handled as follows:

- The application priority TLV is transmitted only if the priorities in the advertisement match the configured PFC priorities on the port.
- On auto-upstream and auto-downstream ports:
 - If a configuration source is elected, the ports send an application priority TLV based on the application priority TLV received on the configuration-source port. When an application priority TLV is received on the configuration-source port, the auto-upstream and auto-downstream ports use the internally propagated PFC priorities to match against the received application priority. Otherwise, these ports use their locally configured PFC priorities in application priority TLVs.
 - If no configuration source is configured, auto-upstream and auto-downstream ports check to see that the locally configured PFC priorities match the priorities in a received application priority TLV.
- On manual ports, an application priority TLV is advertised only if the priorities in the TLV match the PFC priorities configured on the port.

DCB Configuration Exchange

On an Aggregator, the DCBx protocol supports the exchange and propagation of configuration information for the following DCB features.

- Enhanced transmission selection (ETS)
- Priority-based flow control (PFC)

DCBx uses the following methods to exchange DCB configuration parameters:

Asymmetric DCB parameters are exchanged between a DCBx-enabled port and a peer port without requiring that a peer port and the local port use the same configured values for the configurations to be compatible. For example, ETS uses an asymmetric exchange of parameters between DCBx peers.

Symmetric DCB parameters are exchanged between a DCBx-enabled port and a peer port but requires that each configured parameter value be the same for the configurations in order to be compatible. For example, PFC uses an symmetric exchange of parameters between DCBx peers.

Configuration Source Election

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port first checks to see if there is an active configuration source on the switch.

- If a configuration source already exists, the received peer configuration is checked against the local port configuration. If the received configuration is compatible, the DCBx marks the port as DCBx-enabled. If the configuration received from the peer is not compatible, a warning message is logged and the DCBx frame error counter is incremented. Although DCBx is operationally disabled, the port

keeps the peer link up and continues to exchange DCBx packets. If a compatible peer configuration is later received, DCBx is enabled on the port.


- If there is no configuration source, a port may elect itself as the configuration source. A port may become the configuration source if the following conditions exist:
 - No other port is the configuration source.
 - The port role is auto-upstream.
 - The port is enabled with link up and DCBx enabled.
 - The port has performed a DCBx exchange with a DCBx peer.
 - The switch is capable of supporting the received DCB configuration values through either a symmetric or asymmetric parameter exchange.

A newly elected configuration source propagates configuration changes received from a peer to the other auto-configuration ports. Ports receiving auto-configuration information from the configuration source ignore their current settings and use the configuration source information.

Propagation of DCB Information

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port acts as a DCBx client and checks if a DCBx configuration source exists on the switch.

- If a configuration source is found, the received configuration is checked against the currently configured values that are internally propagated by the configuration source. If the local configuration is compatible with the received configuration, the port is enabled for DCBx operation and synchronization.
- If the configuration received from the peer is not compatible with the internally propagated configuration used by the configuration source, the port is disabled as a client for DCBx operation and synchronization and a syslog error message is generated. The port keeps the peer link up and continues to exchange DCBx packets. If a compatible configuration is later received from the peer, the port is enabled for DCBx operation.

 **NOTE:** When a configuration source is elected, all auto-upstream ports other than the configuration source are marked as *willing disabled*. The internally propagated DCB configuration is refreshed on all auto-configuration ports and each port may begin configuration negotiation with a DCBx peer again.

Auto-Detection of the DCBx Version

The Aggregator operates in auto-detection mode so that a DCBx port automatically detects the DCBx version on a peer port. Legacy CIN and CEE versions are supported in addition to the standard IEEE version 2.5 DCBx.

A DCBx port detects a peer version after receiving a valid frame for that version. The local DCBx port reconfigures to operate with the peer version and maintains the peer version on the link until one of the following conditions occurs:

- The switch reboots.
- The link is reset (goes down and up).
- The peer times out.
- Multiple peers are detected on the link.

DCBx operations on a port are performed according to the auto-configured DCBx version, including fast and slow transmit timers and message formats. If a DCBx frame with a different version is received, a

syslog message is generated and the peer version is recorded in the peer status table. If the frame cannot be processed, it is discarded and the discard counter is incremented.

DCBx Example

The following figure shows how DCBx is used on an Aggregator installed in a Dell PowerEdge FX2 server chassis in which servers are also installed.

The Aggregator ports are numbered 1 to 12. Ports 1 to 8 are internal server-facing interfaces. Ports 9 to 12 are uplink ports. The uplink ports on the base module (ports 9 to 12) are used for uplinks configured as DCBx auto-upstream ports. The Aggregator is connected to third-party, top-of-rack (ToR) switches through the uplinks. The ToR switches are part of a Fibre Channel storage network.

The internal ports (ports 1 to 8) connected to the 10GbE backplane are configured as auto-downstream ports.

On the Aggregator, PFC and ETS use DCBx to exchange link-level configuration with DCBx peer devices.

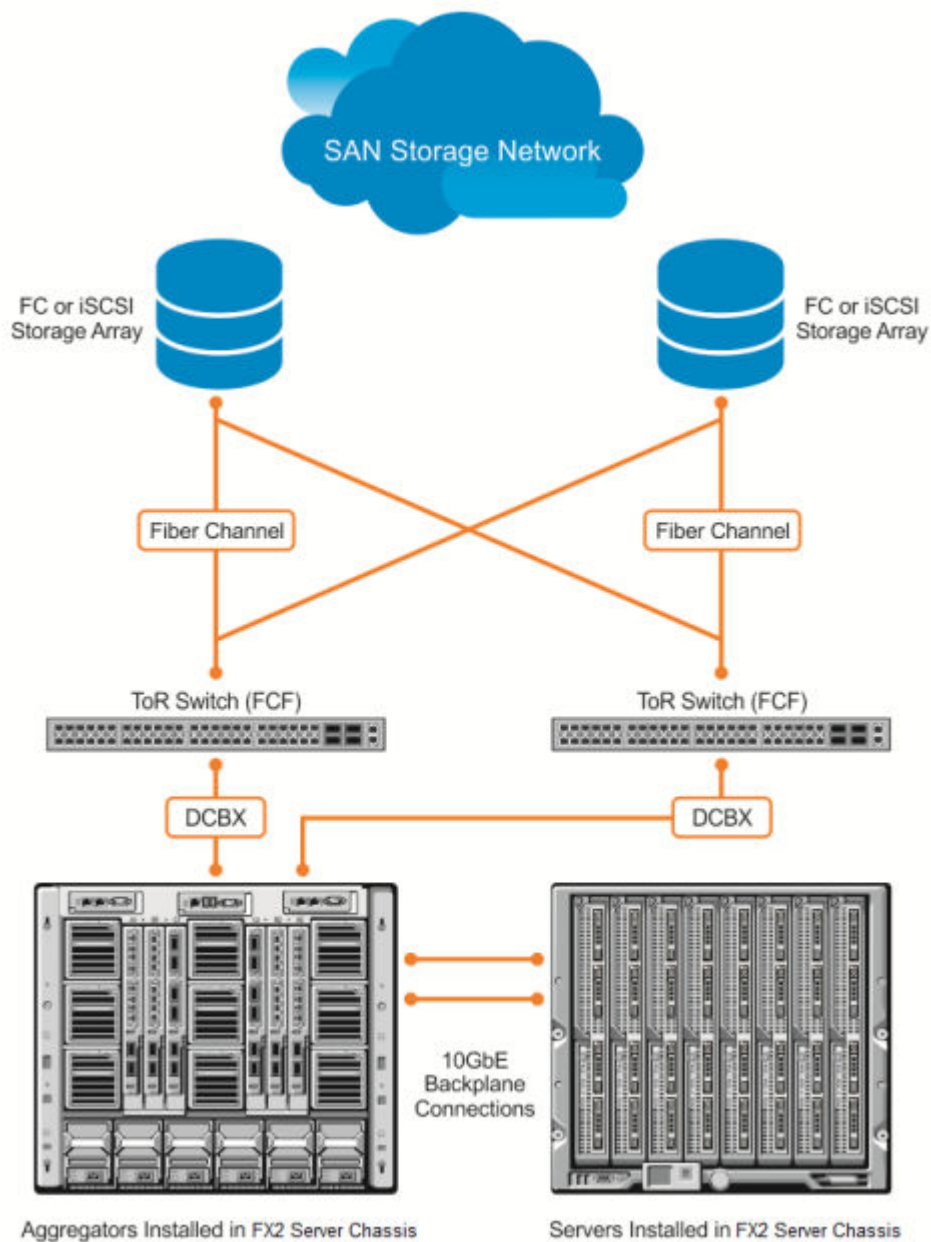


Figure 4. DCBx Sample Topology

DCBx Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure DCBx operation on a port:

- DCBx requires LLDP in both send (TX) and receive (RX) modes to be enabled on a port interface. If multiple DCBx peer ports are detected on a local DCBx interface, LLDP is shut down.

- The CIN version of DCBx supports only PFC, ETS, and FCOE; it does not support iSCSI, backward congestion management (BCN), logical link down (LLD), and network interface virtualization (NIV).

DCBx Error Messages

The following syslog messages appear when an error in DCBx operation occurs.

LLDP_MULTIPLE_PEER_DETECTED: DCBx is operationally disabled after detecting more than one DCBx peer on the port interface.

LLDP_PEER_AGE_OUT: DCBx is disabled as a result of LLDP timing out on a DCBx peer interface.

DSM_DCBx_PEER_VERSION_CONFLICT: A local port expected to receive the IEEE, CIN, or CEE version in a DCBx TLV from a remote peer but received a different, conflicting DCBx version.

DSM_DCBx_PFC_PARAMETERS_MATCH and DSM_DCBx_PFC_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) PFC configuration from a peer.

DSM_DCBx_ETS_PARAMETERS_MATCH and DSM_DCBx_ETS_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) ETS configuration from a peer.

LLDP_UNRECOGNISED_DCBx_TLV_RECEIVED: A local DCBx port received an unrecognized DCBx TLV from a peer.

Debugging DCBx on an Interface

To enable DCBx debug traces for all or a specific control paths, use the following command.

- Enable DCBx debugging.
EXEC PRIVILEGE mode

```
debug dcbx {all | auto-detect-timer | config-exchng | fail | mgmt | resource
| sem | tlv}
```

- all: enables all DCBx debugging operations.
- auto-detect-timer: enables traces for DCBx auto-detect timers.
- config-exchng: enables traces for DCBx configuration exchanges.
- fail: enables traces for DCBx failures.
- mgmt: enables traces for DCBx management frames.
- resource: enables traces for DCBx system resource frames.
- sem: enables traces for the DCBx state machine.
- tlv: enables traces for DCBx TLVs.

Verifying the DCB Configuration

To display DCB configurations, use the following `show` commands.

Table 3. Displaying DCB Configurations

Command	Output
<code>show dcb [stack-unit unit-number]</code>	Displays the data center bridging status, number of PFC-enabled ports, and number of PFC-enabled queues. On the master switch in a stack, you can specify a stack-unit number. The range is from 0 to 5.
<code>show interface port-type slot/port pfc statistics</code>	Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.
<code>show interface port-type slot/port pfc {summary detail}</code>	Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay. To clear PFC TLV counters, use the <code>clear pfc counters {stack-unit unit-number tengigabitethernet slot/port}</code> command.
<code>show interface port-type slot/port ets {summary detail}</code>	Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation. To clear ETS TLV counters, enter the <code>clear ets counters stack-unit unit-number</code> command.

Example of the `show dcb` Command

```
Dell(conf)#show dcb stack-unit 0 port-set 0

DCB Status: Enabled, PFC Queue Count: 2

stack-unit Total Buffer PFC Total Buffer PFC Shared Buffer PFC Available Buffer
      PP          (KB)          (KB)          (KB)          (KB)
-----
--
0      0      3822          1912          832          450
Dell(conf)#
```

Example of the `show interface pfc statistics` Command

```
Dell#show interfaces tengigabitethernet 0/3 pfc statistics
Interface TenGigabitEthernet 0/3

Priority Rx XOFF Frames Rx Total Frames Tx Total Frames
-----
0      0      0      0      0
1      0      0      0      0
2      0      0      0      0
3      0      0      0      0
4      0      0      0      0
5      0      0      0      0
```

```

6          0          0          0
7          0          0          0

```

Example of the show interfaces pfc summary Command

```

Dell# show interfaces tengigabitethernet 0/4 pfc summary
Interface TenGigabitEthernet 0/4
  Admin mode is on
  Admin is enabled
  Remote is enabled, Priority list is 4
  Remote Willing Status is enabled
  Local is enabled
  Oper status is Recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quantams
  Application Priority TLV Parameters :
  -----
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8

```

```

Dell# show interfaces tengigabitethernet 0/4 pfc detail
Interface TenGigabitEthernet 0/4
  Admin mode is on
  Admin is enabled
  Remote is enabled
  Remote Willing Status is enabled
  Local is enabled
  Oper status is recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quanta
  Application Priority TLV Parameters :
  -----
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8

```

```

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 0 Pause Rx
pkts
2 Input Appln Priority TLV pkts, 0 Output Appln Priority TLV pkts, 0 Error
Appln Priority TLV Pkts

```

The following table describes the show interface pfc summary command fields.

Table 4. show interface pfc summary Command Description

Fields	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on; Admin is enabled	PFC Admin mode is on or off with a list of the configured PFC priorities . When PFC admin mode

Fields	Description
Remote is enabled; Priority list Remote Willing Status is enabled	is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration takes effect. The admin operational status for a DCBx exchange of PFC configuration is enabled or disabled. Operational status (enabled or disabled) of peer device for DCBx exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBx exchange (Willing bit received in PFC TLV): enabled or disabled.
Local is enabled	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities.
Operational status (local port)	Port state for current operational PFC configuration: <ul style="list-style-type: none"> • Init: Local PFC configuration parameters were exchanged with peer. • Recommend: Remote PFC configuration parameters were received from peer. • Internally propagated: PFC configuration parameters were received from configuration source.
PFC DCBx Oper status	Operational status for exchange of PFC configuration on local port: match (up) or mismatch (down).
Reason	Reason displayed when the DCBx operational status for PFC on a port is down.
State Machine Type	Type of state machine used for DCBx exchanges of PFC parameters: <ul style="list-style-type: none"> • Feature: for legacy DCBx versions • Symmetric: for an IEEE version
TLV Tx Status	Status of PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: ISCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap used by local DCBx port in FCoE advertisements in application priority TLVs.

Fields	Description
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap used by local DCBx port in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Priority bitmap received from the remote DCBx port in FCoE advertisements in application priority TLVs.
Application Priority TLV: Remote ISCSI Priority Map	Priority bitmap received from the remote DCBx port in iSCSI advertisements in application priority TLVs.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received.
Input Appln Priority TLV pkts	Number of Application Priority TLVs received.
Output Appln Priority TLV pkts	Number of Application Priority TLVs transmitted.
Error Appln Priority TLV pkts	Number of Application Priority error packets received.

Example of the show interface ets summary Command

```

Dell# show interfaces te 0/1 ets summary
Interface TenGigabitEthernet 0/1
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
-----
Admin is enabled
TC-grp    Priority#    Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%           ETS
2         0%           ETS
3         0%           ETS
4         0%           ETS
5         0%           ETS
6         0%           ETS
7         0%           ETS

Priority#  Bandwidth  TSA
0         13%       ETS
1         13%       ETS
2         13%       ETS
3         13%       ETS
4         12%       ETS
5         12%       ETS
6         12%       ETS
7         12%       ETS
Remote Parameters:

```

```

-----
Remote is disabled

Local Parameters :
-----
Local is enabled
TC-grp    Priority#    Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%              ETS
2         0%              ETS
3         0%              ETS
4         0%              ETS
5         0%              ETS
6         0%              ETS
7         0%              ETS

Priority#    Bandwidth    TSA
0           13%         ETS
1           13%         ETS
2           13%         ETS
3           13%         ETS
4           12%         ETS
5           12%         ETS
6           12%         ETS
7           12%         ETS

Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled

```

Example of the show interface ets detail Command

```

Dell# show interfaces tengigabitethernet 0/4 ets detail
Interface TenGigabitEthernet 0/4
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
-----
Admin is enabled
TC-grp    Priority#    Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%              ETS
2         0%              ETS
3         0%              ETS
4         0%              ETS
5         0%              ETS
6         0%              ETS
7         0%              ETS

Remote Parameters:
-----
Remote is disabled

Local Parameters :
-----
Local is enabled
PG-grp    Priority#    Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%              ETS
2         0%              ETS
3         0%              ETS

```

```

4          0%          ETS
5          0%          ETS
6          0%          ETS
7          0%          ETS

```

```

Oper status is init
ETS DCBX Oper status is Down
Reason: Port Shutdown
State Machine Type is Asymmetric
Conf TLV Tx Status is enabled
Reco TLV Tx Status is enabled

```

```

0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
0 Input Reco TLV Pkts, 0 Output Reco TLV Pkts, 0 Error Reco TLV Pkts

```

The following table describes the `show interface ets detail` command fields.

Table 5. show interface ets detail Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBx TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including Admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If the ETS Admin mode is enabled on the remote port for DCBx exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including Admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: <ul style="list-style-type: none"> • Init: Local ETS configuration parameters were exchanged with peer. • Recommend: Remote ETS configuration parameters were received from peer.

Field	Description
	<ul style="list-style-type: none"> Internally propagated: ETS configuration parameters were received from configuration source.
ETS DCBx Oper status	Operational status of ETS configuration on local port: match or mismatch.
Reason	Reason displayed when the DCBx operational status for ETS on a port is down.
State Machine Type	Type of state machine used for DCBx exchanges of ETS parameters: <ul style="list-style-type: none"> Feature: for legacy DCBx versions Asymmetric: for an IEEE version
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
Reco TLV Tx Status	Status of ETS Recommendation TLV advertisements: enabled or disabled.
Input Conf TLV pkts, Output Conf TLV pkts, Error Conf TLV pkts	Number of ETS Configuration TLVs received and transmitted, and number of ETS Error Configuration TLVs received.
Input Reco TLV pkts, Output Reco TLV pkts, Error Reco TLV pkts	Number of ETS Recommendation TLVs received and transmitted, and number of ETS Error Recommendation TLVs received.

Example of the show stack-unit all stack-ports all pfc details Command

```
Dell# show stack-unit all stack-ports all pfc details

stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

Example of the show stack-unit all stack-ports all ets details Command

```
Dell# show stack-unit all stack-ports all ets details
Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1

Admin mode is on
```

Admin Parameters:

```
-----  
Admin is enabled  
TC-grp    Priority#          Bandwidth    TSA  
-----  
0          0,1,2,3,4,5,6,7    100%        ETS  
1          -                    -           -  
2          -                    -           -  
3          -                    -           -  
4          -                    -           -  
5          -                    -           -  
6          -                    -           -  
7          -                    -           -  
8          -                    -           -
```

Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on

Admin Parameters:

```
-----  
Admin is enabled  
TC-grp    Priority#          Bandwidth    TSA  
-----  
0          0,1,2,3,4,5,6,7    100%        ETS  
1          -                    -           -  
2          -                    -           -  
3          -                    -           -  
4          -                    -           -  
5          -                    -           -  
6          -                    -           -  
7          -                    -           -  
8          -                    -           -
```

Example of the show interface DCBx detail Command

```
Dell# show interface tengigabitethernet 0/4 dcbx detail  
Dell#show interface te 0/4 dcbx detail
```

```
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled  
R-ETS Recommendation TLV enabled         r-ETS Recommendation TLV disabled  
P-PFC Configuration TLV enabled         p-PFC Configuration TLV disabled  
F-Application priority for FCOE enabled  f-Application Priority for FCOE  
disabled  
I-Application priority for iSCSI enabled i-Application Priority for iSCSI  
disabled
```

--

```
Interface TenGigabitEthernet 0/4  
  Remote Mac Address 00:00:00:00:00:11  
  Port Role is Auto-Upstream  
  DCBX Operational Status is Enabled  
  Is Configuration Source? TRUE
```

```
Local DCBX Compatibility mode is CEE  
Local DCBX Configured mode is CEE  
Peer Operating version is CEE  
Local DCBX TLVs Transmitted: ErPfi
```

Local DCBX Status

```

-----
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 2
Acknowledgment Number: 2
Protocol State: In-Sync

Peer DCBX Status:
-----
DCBX Operational Version is 0
DCBX Max Version Supported is 255
Sequence Number: 2
Acknowledgment Number: 2
2 Input PFC TLV pkts, 3 Output PFC TLV pkts, 0 Error PFC pkts, 0 PFC Pause Tx
pkts,
0 Pause Rx pkts
2 Input PG TLV Pkts, 3 Output PG TLV Pkts, 0 Error PG TLV Pkts
2 Input Appln Priority TLV pkts, 0 Output Appln Priority TLV pkts, 0 Error
Appln Priority
TLV Pkts
Total DCBX Frames transmitted 27
Total DCBX Frames received 6
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0

```

The following table describes the show interface DCBx detail command fields.

Table 6. show interface DCBx detail Command Description

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured DCBx port role: auto-upstream or auto-downstream.
DCBx Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBx operational status is the combination of PFC and ETS operational status.
Configuration Source	Specifies whether the port serves as the DCBx configuration source on the switch: true (yes) or false (no).
Local DCBx Compatibility mode	DCBx version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only received a DCBx version supported on the remote peer.
Local DCBx Configured mode	DCBx version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBx version received from a peer).
Peer Operating version	DCBx version that the peer uses to exchange DCB parameters.

Field	Description
Local DCBx TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs.
Local DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs.
Local DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBx Status: Protocol State	Current operational state of DCBx protocol: ACK or IN-SYNC.
Peer DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs received from peer device.
Peer DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs received from peer device.
Peer DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs received from peer device.
Peer DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from peer device.
Total DCBx Frames transmitted	Number of DCBx frames sent from local port.
Total DCBx Frames received	Number of DCBx frames received from remote peer port.
Total DCBx Frame errors	Number of DCBx frames with errors received.
Total DCBx Frames unrecognized	Number of unrecognizable DCBx frames received.
PFC TLV Statistics: Input PFC TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output PFC TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error PFC pkts	Number of PFC error packets received.
PFC TLV Statistics: PFC Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: PFC Pause Rx pkts	Number of PFC pause frames received.
PG TLV Statistics: Input PG TLV Pkts	Number of PG TLVs received.
PG TLV Statistics: Output PG TLV Pkts	Number of PG TLVs transmitted.
PG TLV Statistics: Error PG TLV Pkts	Number of PG error packets received.
Application Priority TLV Statistics: Input Appln Priority TLV pkts	Number of Application TLVs received.
Application Priority TLV Statistics: Output Appln Priority TLV pkts	Number of Application TLVs transmitted.
Application Priority TLV Statistics: Error Appln Priority TLV Pkts	Number of Application TLV error packets received.

Hierarchical Scheduling in ETS Output Policies

ETS supports up to three levels of hierarchical scheduling.

For example, you can apply ETS output policies with the following configurations:

- Priority group 1** Assigns traffic to one priority queue with 20% of the link bandwidth and strict-priority scheduling.
- Priority group 2** Assigns traffic to one priority queue with 30% of the link bandwidth.
- Priority group 3** Assigns traffic to two priority queues with 50% of the link bandwidth and strict-priority scheduling.

In this example, the configured ETS bandwidth allocation and scheduler behavior is as follows:

- Unused bandwidth usage:** Normally, if there is no traffic or unused bandwidth for a priority group, the bandwidth allocated to the group is distributed to the other priority groups according to the bandwidth percentage allocated to each group. However, when three priority groups with different bandwidth allocations are used on an interface:
 - If priority group 3 has free bandwidth, it is distributed as follows: 20% of the free bandwidth to priority group 1 and 30% of the free bandwidth to priority group 2.
 - If priority group 1 or 2 has free bandwidth, (20 + 30)% of the free bandwidth is distributed to priority group 3. Priority groups 1 and 2 retain whatever free bandwidth remains up to the (20+ 30)%.
- Strict-priority groups:** If two priority groups have strict-priority scheduling, traffic assigned from the priority group with the higher priority-queue number is scheduled first. However, when three priority groups are used and two groups have strict-priority scheduling (such as groups 1 and 3 in the example), the strict priority group whose traffic is mapped to one queue takes precedence over the strict priority group whose traffic is mapped to two queues.

Therefore, in this example, scheduling traffic to priority group 1 (mapped to one strict-priority queue) takes precedence over scheduling traffic to priority group 3 (mapped to two strict-priority queues).

Troubleshooting PFC, ETS, and DCBx Operation

In the `show interfaces pfc | ets | dcbx` output, the DCBx operational status may be down for any of the reasons described in the following table.

When DCBx is down, the following values display in the `show` output field for DCBx Oper status:

- PFC DCBx Oper status: Down
- ETS DCBx Oper status: Down
- DCBx Oper status: Disabled.

Reason	Description
Port Shutdown	Port is shut down. All other reasons for DCBx inoperation, if any, are ignored.

Reason	Description
LLDP Rx/Tx is disabled	LLDP is disabled (Admin Mode set to rx or tx only) globally or on the interface.
Waiting for Peer	Waiting for peer or detected peer connection has aged out.
Multiple Peer Detected	Multiple peer connections detected on the interface.
Version Conflict	DCBx version on peer version is different than the local or globally configured DCBx version.
Config-source Down	Although DCBx parameters match in auto-upstream or auto-downstream port, the configuration source (elected dynamically or CLI-configured) is down.
Unrecognized TLV received	Invalid TLV length TLV has been received. In the case of an invalid PFC or ETS TLV, the error displays in the <code>show interfaces pfc ets</code> output. The <code>show interfaces dcbx</code> output displays PFC or ETS as down.
Admin Mode OFF (Local)	Local Admin Mode is disabled.
Admin Mode OFF (Remote)	Remote Admin Mode is disabled.
Waiting for ACK from Peer	For a legacy DCBx version, a peer has not acknowledged the reception of a sent packet. This reason displays only when a remote peer is willing to receive a DCB configuration.
Error Bit set	For a legacy DCBx version, a peer has sent packets with an error bit set. This reason displays only when a remote peer is willing to receive a DCB configuration.
Enabled with ETS Mismatch (<code>show interfaces dcbx</code> output)	DCBx is enabled but an ETS validation failure error has occurred.
PFC is down (<code>show interfaces pfc</code> output)	One of the following PFC-specific errors has occurred: <ul style="list-style-type: none"> • No MBC support. • Configured PFC priorities exceed maximum PFC capability limit. • New dot1p-to-queue mapping violates the allowed system limit for PFC Enable status per priority
ETS is down (<code>show interfaces ets</code> output)	One of the following ETS-specific errors occurred in ETS validation: <ul style="list-style-type: none"> • Unsupported PGID • A priority group exceeds the maximum number of supported priorities. • COSQ is mapped to more than one priority group. • COSQ is mapped to more than one priority group. - Invalid or unsupported transmission selection algorithm (TSA). • Bandwidth is configured for an unconfigured priority group. • Total ETS bandwidth mapped in priority groups is not equal to 100%. • Priorities mapped to a queue use different TSAs.

Reason**Description**

- Total bandwidth assigned to priorities in one or more priority groups is not equal to 100%.

Or one of the following ETS failure errors occurred:


- Incompatible priority group ID (PGID).
- Incompatible bandwidth (BW) allocation.
- Incompatible TSA.
- Incompatible TC BW.
- Incompatible TC TSA.

Dynamic Host Configuration Protocol (DHCP)

The Aggregator is auto-configured to operate as a dynamic host configuration protocol (DHCP) client. The DHCP server, DHCP relay agent, and secure DHCP features are not supported. The DHCP is an application layer protocol that dynamically assigns IP addresses and other configuration parameters to network end-stations (hosts) based on configuration policies determined by network administrators. DHCP relieves network administrators of manually configuring hosts, which can be a tedious and error-prone process when hosts often join, leave, and change locations on the network and it reclaims IP addresses that are no longer in use to prevent address exhaustion.

DHCP is based on a client-server model. A host discovers the DHCP server and requests an IP address, and the server either leases or permanently assigns one. There are three types of devices that are involved in DHCP negotiation:

DHCP Server	This is a network device offering configuration parameters to the client.
DHCP Client	This is a network device requesting configuration parameters from the server.
Relay Agent	This is an intermediary network device that passes DHCP messages between the client and server when the server is not on the same subnet as the host.

 **NOTE:** The DHCP server and relay agent features are not supported on an Aggregator.

Supported Modes

Stacking, PMUX, Standalone, VLT

Assigning an IP Address using DHCP

The following section describes DHCP and the client in a network.

When a client joins a network:

1. The client initially broadcasts a **DHCPDISCOVER** message on the subnet to discover available DHCP servers. This message includes the parameters that the client requires and might include suggested values for those parameters.
2. Servers unicast or broadcast a **DHCPOFFER** message in response to the DHCPDISCOVER that offers to the client values for the requested parameters. Multiple servers might respond to a single DHCPDISCOVER; the client might wait a period of time and then act on the most preferred offer.
3. The client broadcasts a **DHCPREQUEST** message in response to the offer, requesting the offered values.
4. After receiving a DHCPREQUEST, the server binds the clients' unique identifier (the hardware address plus IP address) to the accepted configuration parameters and stores the data in a database called a

binding table. The server then broadcasts a **DHCPACK** message, which signals to the client that it may begin using the assigned parameters.

There are additional messages that are used in case the DHCP negotiation deviates from the process previously described and shown in the illustration below.

- | | |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DHCPDECLINE | A client sends this message to the server in response to a DHCPACK if the configuration parameters are unacceptable; for example, if the offered address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER. |
| DHCPINFORM | A client uses this message to request configuration parameters when it assigned an IP address manually rather than with DHCP. The server responds by unicast. |
| DHCPNAK | A server sends this message to the client if it is not able to fulfill a DHCPREQUEST; for example, if the requested address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER. |
| DHCPRELEASE | A DHCP client sends this message when it is stopped forcefully to return its IP address to the server. |

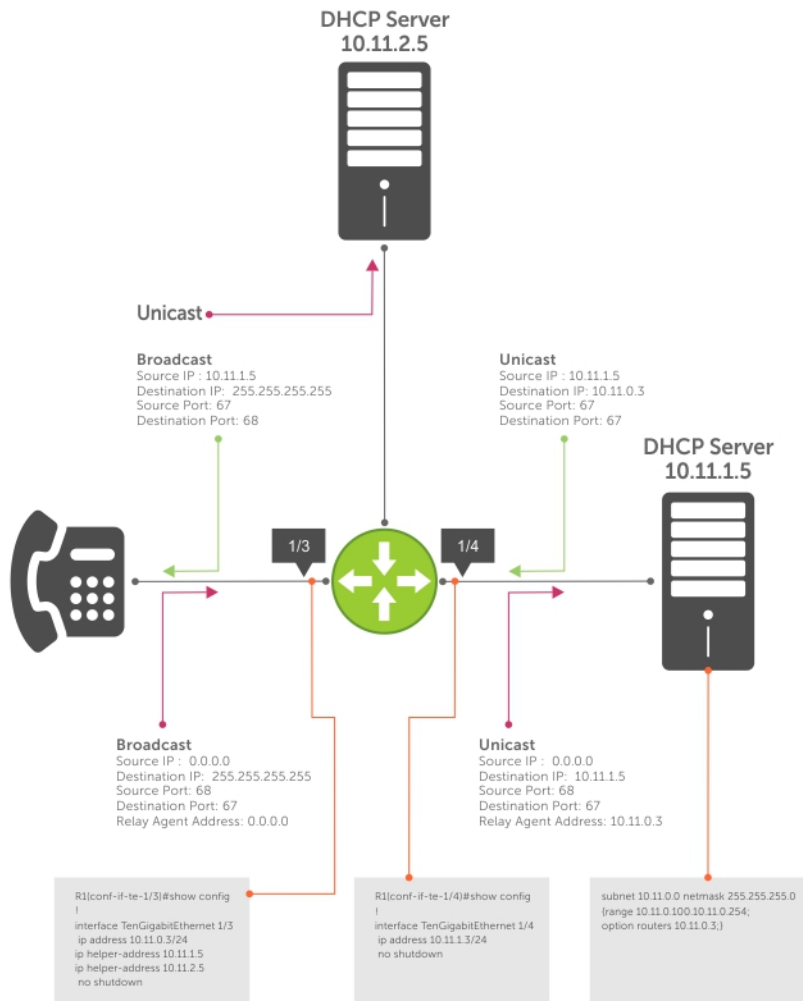


Figure 5. Assigning Network Parameters using DHCP

Dell Networking OS Behavior: DHCP is implemented in Dell Networking OS based on RFC 2131 and 3046.

Debugging DHCP Client Operation

To enable debug messages for DHCP client operation, enter the following **debug** commands:

- Enable the display of log messages for all DHCP packets sent and received on DHCP client interfaces.
EXEC Privilege

[no] debug ip dhcp client packets [interface type slot/port]
- Enable the display of log messages for the following events on DHCP client interfaces: IP address acquisition, IP address release, Renewal of IP address and lease time, and Release of an IP address.
EXEC Privilege

```
[no] debug ip dhcp client events [interface type slot/port]
```

The following example shows the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface.

DHCP Client: Debug Messages Logged during DHCP Client Enabling/Disabling

```
Dell (conf-if-Ma-0/0)# ip address dhcp
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :DHCP ENABLE
CMD
Received in state START
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :Transitioned
to state SELECTING
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: DHCP
DISCOVER sent in
Interface
Ma 0/0
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: Received
DHCP OFFER packet in
Interface Ma 0/0 with Lease-ip:10.16.134.250, Mask:255.255.0.0, Server-Id:
10.16.134.249
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :Transitioned
to state REQUESTING
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:DHCP REQUEST
sent in Interface
Ma 0/0
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:Received
DHCP ACK packet in
Interface
Ma 0/0 with Lease-IP:10.16.134.250, Mask:255.255.0.0,DHCP REQUEST sent in
Interface Ma 0/0
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :Transitioned
to state BOUND
```

```
Dell(conf-if-ma-0/0)# no ip address
Dell(conf-if-ma-0/0)#1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG:
DHCLIENT_DBG_EVT: Interface
Ma 0/0 :DHCP_DISABLE CMD Received in state SELECTING
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :Transitioned
to state START
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :DHCP DISABLED
CMD
sent to FTOS in state START
```

```
Dell# release dhcp int Ma 0/0
Dell#1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :DHCP
RELEASE
CMD Received in state BOUND
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: DHCP
RELEASE sent in Interface
Ma 0/0
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
```

```

Ma 0/0 :Transitioned
to state STOPPED
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :DHCP IP
RELEASED
CMD sent to FTOS in state STOPPED

```

```

Dell# renew dhcp int Ma 0/0
Dell#1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :DHCP
RENEW
CMD Received in state STOPPED
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface
Ma 0/0 :Transitioned
to state SELECTING
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: DHCP
DISCOVER sent in Interface
Ma 0/0
1w2d23h: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: Received
DHCP OFFER packet in
Interface Ma 0/0 with Lease-Ip:10.16.134.250, Mask:255.255.0.0, Server-Id:
10.16.134.249

```

The following example shows the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface when you release and renew a DHCP client.

DHCP Client: Debug Messages Logged during DHCP Client Release/Renew

```

Dell# release dhcp interface managementethernet 0/0
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :
DHCP RELEASE CMD Received in state BOUND
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in Interface Ma 0/0
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :
Transitioned to state STOPPED
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :
DHCP IP RELEASED CMD sent to FTOS in state STOPPED

```

```

Dell# renew dhcp interface tengigabitethernet 0/1
Dell#May 27 15:55:28: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG:
DHCLIENT_DBG_EVT: Interface Ma 0/0 :
DHCP RENEW CMD Received in state STOPPED
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Ma 0/0 :
Transitioned to state SELECTING
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in Interface Ma 0/0
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCP OFFER packet in Interface Ma 0/0 with Lease-Ip:10.16.134.250,
Mask:255.255.0.0, Server-Id:10.16.134.249

```


DHCP Client

An Aggregator is auto-configured to operate as a DHCP client. The DHCP client functionality is enabled only on the default VLAN and the management interface.

A DHCP client is a network device that requests an IP address and configuration parameters from a DHCP server. On an Aggregator, the DHCP client functionality is implemented as follows:

- The public out-of-band management (OOB) interface and default VLAN 1 are configured, by default, as a DHCP client to acquire a dynamic IP address from a DHCP server.

You can override the DHCP-assigned address on the OOB management interface by manually configuring an IP address using the CLI or CMC interface. If no user-configured IP address exists for the OOB interface exists and if the OOB IP address is not in the startup configuration, the Aggregator will automatically obtain it using DHCP.

You can also manually configure an IP address for the VLAN 1 default management interface using the CLI. If no user-configured IP address exists for the default VLAN management interface exists and if the default VLAN IP address is not in the startup configuration, the Aggregator will automatically obtain it using DHCP.

- The default VLAN 1 with all ports configured as members is the only L3 interface on the Aggregator. When the default management VLAN has a DHCP-assigned address and you reconfigure the default VLAN ID number, the Aggregator:
 - Sends a DHCP release to the DHCP server to release the IP address.
 - Sends a DHCP request to obtain a new IP address. The IP address assigned by the DHCP server is used for the new default management VLAN.


How DHCP Client is Implemented

The Aggregator is enabled by default to receive DHCP server-assigned dynamic IP addresses on an interface. This setting persists after a switch reboot. If you enter the `shutdown` command on the interface, DHCP transactions are stopped and the dynamically-acquired IP address is saved. Use the `show interface type slot/port` command to display the dynamic IP address and DHCP as the mode of IP address assignment. If you later enter the `no shutdown` command and the lease timer for the dynamic IP address has expired, the IP address is unconfigured and the interface tries to acquire a new dynamic address from DHCP server.

If you later enter the `no shutdown` command and the lease timer for the dynamic IP address has expired, the IP address is released.

When you enter the `release dhcp` command, although the IP address that was dynamically-acquired from a DHCP server is released from an interface, the ability to acquire a new DHCP server-assigned address remains in the running configuration for the interface. To acquire a new IP address, enter either the `renew dhcp` command at the EXEC privilege level or the `ip address dhcp` command at the interface configuration level.

If you enter `renew dhcp` command on an interface already configured with a dynamic IP address, the lease time of the dynamically acquired IP address is renewed.

 **Important:** To verify the currently configured dynamic IP address on an interface, enter the `show ip dhcp lease` command. The `show running-configuration` command output only displays `ip address dhcp`; the currently assigned dynamic IP address is not displayed.

DHCP Client on a Management Interface

These conditions apply when you enable a management interface to operate as a DHCP client.

- The management default route is added with the gateway as the router IP address received in the DHCP ACK packet. It is required to send and receive traffic to and from other subnets on the external network. The route is added irrespective of whether the DHCP client and server are on same or different subnets. The management default route is deleted if the management IP address is released like other DHCP client management routes.
- `ip route for 0.0.0.0` takes precedence if it is present or added later.
- Management routes added by a DHCP client display with Route Source as **DHCP** in the `show ip management route` and `show ip management-route dynamic` command output.
- Management routes added by DHCP are automatically reinstalled if you configure a static IP route with the `ip route` command that replaces a management route added by the DHCP client. If you remove the statically configured IP route using the `no ip route` command, the management route is reinstalled. Manually delete management routes added by the DHCP client.
- To reinstall management routes added by the DHCP client that is removed or replaced by the same statically configured management routes, release the DHCP IP address and renew it on the management interface.
- Management routes added by the DHCP client have higher precedence over the same statically configured management route. Static routes are not removed from the running configuration if a dynamically acquired management route added by the DHCP client overwrites a static management route.
- Management routes added by the DHCP client are not added to the running configuration.

 **NOTE:** Management routes added by the DHCP client include the specific routes to reach a DHCP server in a different subnet and the management route.

DHCP Client on a VLAN

The following conditions apply on a VLAN that operates as a DHCP client:

- The default VLAN 1 with all ports auto-configured as members is the only L3 interface on the Aggregator.
- When the default management VLAN has a DHCP-assigned address and you reconfigure the default VLAN ID number, the Aggregator:
 - Sends a DHCP release to the DHCP server to release the IP address.
 - Sends a DHCP request to obtain a new IP address. The IP address assigned by the DHCP server is used for the new default management VLAN.

DHCP Packet Format and Options

DHCP uses the user datagram protocol (UDP) as its transport protocol.

The server listens on port 67 and transmits to port 68; the client listens on port 68 and transmits to port 67. The configuration parameters are carried as options in the DHCP packet in Type, Length, Value (TLV) format; many options are specified in RFC 2132. To limit the number of parameters that servers must

provide, hosts specify the parameters that they require, and the server sends only those parameters. Some common options are shown in the following illustration.




Figure 6. DHCP packet Format

The following table lists common DHCP options.

Option	Number and Description
Subnet Mask	Option 1 Specifies the client's subnet mask.
Router	Option 3 Specifies the router IP addresses that may serve as the client's default gateway.
Domain Name Server	Option 6 Specifies the domain name servers (DNSs) that are available to the client.
Domain Name	Option 15 Specifies the domain name that clients should use when resolving hostnames via DNS.
IP Address Lease Time	Option 51 Specifies the amount of time that the client is allowed to use an assigned IP address.
DHCP Message Type	Option 53 <ul style="list-style-type: none"> • 1: DHCPDISCOVER • 2: DHCPOFFER • 3: DHCPREQUEST • 4: DHCPDECLINE • 5: DHCPACK • 6: DHCPNACK • 7: DHCPRELEASE • 8: DHCPINFORM

Option	Number and Description
Parameter Request List	Option 55 Clients use this option to tell the server which parameters it requires. It is a series of octets where each octet is DHCP option code.
Renewal Time	Option 58 Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with the <i>original</i> server.
Rebinding Time	Option 59 Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with <i>any</i> server, if the original server does not respond.
End	Option 255 Signals the last option in the DHCP packet.

Option 82

RFC 3046 (the relay agent information option, or Option 82) is used for class-based IP address assignment.

The code for the relay agent information option is 82, and is comprised of two sub-options, circuit ID and remote ID.

Circuit ID	This is the interface on which the client-originated message is received.
Remote ID	This identifies the host from which the message is received. The value of this sub-option is the MAC address of the relay agent that adds Option 82.

The DHCP relay agent inserts Option 82 before forwarding DHCP packets to the server. The server can use this information to:

- track the number of address requests per relay agent. Restricting the number of addresses available per relay agent can harden a server against address exhaustion attacks.
- associate client MAC addresses with a relay agent to prevent offering an IP address to a client spoofing the same MAC address on a different relay agent.
- assign IP addresses according to the relay agent. This prevents generating DHCP offers in response to requests from an unauthorized relay agent.

The server echoes the option back to the relay agent in its response, and the relay agent can use the information in the option to forward a reply out the interface on which the request was received rather than flooding it on the entire VLAN.

The relay agent strips Option 82 from DHCP responses before forwarding them to the client.

To insert Option 82 into DHCP packets, follow this step.

- Insert Option 82 into DHCP packets.

CONFIGURATION mode

```
int ma 0/0
ip add dhcp relay information-option remote-id
```

For routers between the relay agent and the DHCP server, enter the `trust-downstream` option.

Releasing and Renewing DHCP-based IP Addresses

On an Aggregator configured as a DHCP client, you can release a dynamically-assigned IP address without removing the DHCP client operation on the interface.

To manually acquire a new IP address from the DHCP server, use the following command.

- Release a dynamically-acquired IP address while retaining the DHCP client configuration on the interface.

EXEC Privilege mode

```
release dhcp interface type slot/port
```

- Acquire a new IP address with renewed lease time from a DHCP server.

EXEC Privilege mode

```
renew dhcp interface type slot/port
```

Viewing DHCP Statistics and Lease Information

To display DHCP client information, enter the following **show** commands:

- Display statistics about DHCP client interfaces.

EXEC Privilege

```
show ip dhcp client statistics interface type slot/port
```

- Clear DHCP client statistics on a specified or on all interfaces.

EXEC Privilege

```
clear ip dhcp client statistics {all | interface type slot/port}
```

- Display lease information about the dynamic IP address currently assigned to a DHCP client interface.

EXEC Privilege

```
show ip dhcp lease [interface type slot/port]
```

View the statistics about DHCP client interfaces with the **show ip dhcp client statistics** command and the lease information about the dynamic IP address currently assigned to a DHCP client interface with the **show ip dhcp lease** command.

Example of the show ip dhcp client statistics Command

```
Dell#show ip dhcp client statistics interface managementethernet 0/0
Interface Name      Ma 0/0
Message            Received
DHCPOFFER          0
DHCPACK             0
DHCPNAK            0
Message            Sent
DHCPDISCOVER       1626
DHCPREQUEST        0
```

```

DHCPDECLINE          0
DHCPRELEASE          0
DHCPREBIND          0
DHCPRENEW            0
DHCPINFORM           0
Dell#

```

Example of the show ip dhcp lease Command

```

Dell# show ip dhcp
Interface Lease-IP Def-Router ServerId State Lease Obtnd At Lease
Expires At
=====
===== Ma 0/0
0.0.0.0/0 0.0.0.0 0.0.0.0 INIT -----NA----- -----NA-----

Vl 1 10.1.1.254/24 0.0.0.0 10.1.1.1 BOUND 08-26-2011 04:33:39
08-27-2011 04:33:39

Renew Time Rebind Time
=====
-----NA----- -----NA-----

08-26-2011 16:21:50 08-27-2011 01:33:39

```

FIP Snooping

This chapter describes about the FIP snooping concepts and configuration procedures.

Supported Modes

Standalone, PMUX, VLT

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) provides a converged Ethernet network that allows the combination of storage-area network (SAN) and LAN traffic on a Layer 2 link by encapsulating Fibre Channel data into Ethernet frames.

FCoE works with Ethernet enhancements provided in data center bridging (DCB) to support lossless (no-drop) SAN and LAN traffic. In addition, DCB provides flexible bandwidth sharing for different traffic types, such as LAN and SAN, according to 802.1p priority classes of service. For more information, refer to the Data Center Bridging (DCB) chapter.

Ensuring Robustness in a Converged Ethernet Network

Fibre Channel networks used for SAN traffic employ switches that operate as trusted devices. End devices log into the switch to which they are attached in order to communicate with the other end devices attached to the Fibre Channel network. Because Fibre Channel links are point-to-point, a Fibre Channel switch controls all storage traffic that an end device sends and receives over the network. As a result, the switch can enforce zoning configurations, ensure that end devices use their assigned addresses, and secure the network from unauthorized access and denial-of-service attacks.

To ensure similar Fibre Channel robustness and security with FCoE in an Ethernet cloud network, the Fibre Channel over Ethernet initialization protocol (FIP) establishes virtual point-to-point links between FCoE end-devices (server ENodes and target storage devices) and FCoE forwarders (FCFs) over transit FCoE-enabled bridges.

Ethernet bridges commonly provide access control list (ACLs) that can emulate a point-to-point link by providing the traffic enforcement required to create a Fibre Channel-level of robustness. In addition, FIP serves as a Layer 2 protocol to:

- Operate between FCoE end-devices and FCFs over intermediate Ethernet bridges to prevent unauthorized access to the network and achieve the required security.
- Allow transit Ethernet bridges to efficiently monitor FIP frames passing between FCoE end-devices and an FCF, and use the FIP snooping data to dynamically configure ACLs on the bridge to only permit traffic authorized by the FCF.

FIP enables FCoE devices to discover one another, initialize and maintain virtual links over an Ethernet network, and access storage devices in a storage area network. FIP satisfies the Fibre Channel

requirement for point-to-point connections by creating a unique virtual link for each connection between an FCoE end-device and an FCF via a transit switch.

FIP provides a functionality for discovering and logging in to an FCF. After discovering and logging in, FIP allows FCoE traffic to be sent and received between FCoE end-devices (ENodes) and the FCF. FIP uses its own EtherType and frame format. The below illustration about FIP discovery, depicts the communication that occurs between an ENode server and an FCoE switch (FCF).

FIP performs the following functions:

- FIP virtual local area network (VLAN) discovery: FCoE devices (ENodes) discover the FCoE VLANs on which to transmit and receive FIP and FCoE traffic.
- FIP discovery: FCoE end-devices and FCFs are automatically discovered.
- Initialization: FCoE devices perform fabric login (FLOGI) and fabric discovery (FDISC) to create a virtual link with an FCoE switch.
- Maintenance: A valid virtual link between an FCoE device and an FCoE switch is maintained and the link termination logout (LOGO) functions properly.

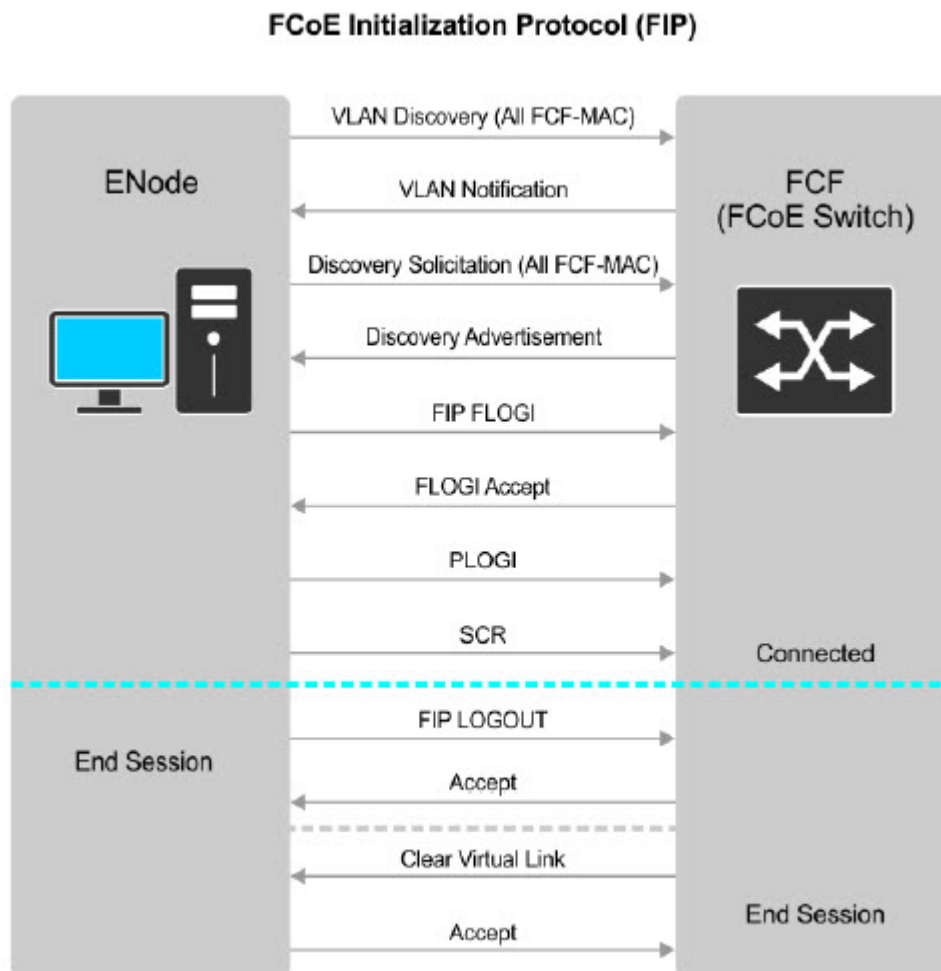


Figure 7. FIP discovery and login between an ENode and an FCF

FIP Snooping on Ethernet Bridges

In a converged Ethernet network, intermediate Ethernet bridges can snoop on FIP packets during the login process on an FCF. Then, using ACLs, a transit bridge can permit only authorized FCoE traffic to be transmitted between an FCoE end-device and an FCF. An Ethernet bridge that provides these functions is called a FIP snooping bridge (FSB).

On a FIP snooping bridge, ACLs are created dynamically as FIP login frames are processed. The ACLs are installed on switch ports configured for the following port modes:

- ENode mode for server-facing ports
- FCF mode for a trusted port directly connected to an FCF

You must enable FIP snooping on an Aggregator and configure the FIP snooping parameters. When you enable FIP snooping, all ports on the switch by default become ENode ports.

Dynamic ACL generation on an Aggregator operating as a FIP snooping bridge functions as follows:

- Global ACLs are applied on server-facing ENode ports.
- Port-based ACLs are applied on ports directly connected to an FCF and on server-facing ENode ports.
- Port-based ACLs take precedence over global ACLs.
- FCoE-generated ACLs take precedence over user-configured ACLs. A user-configured ACL entry cannot deny FCoE and FIP snooping frames.

The below illustration depicts an Aggregator used as a FIP snooping bridge in a converged Ethernet network. The ToR switch operates as an FCF for FCoE traffic. Converged LAN and SAN traffic is transmitted between the ToR switch and an Aggregator. The Aggregator operates as a lossless FIP snooping bridge to transparently forward FCoE frames between the ENode servers and the FCF switch.

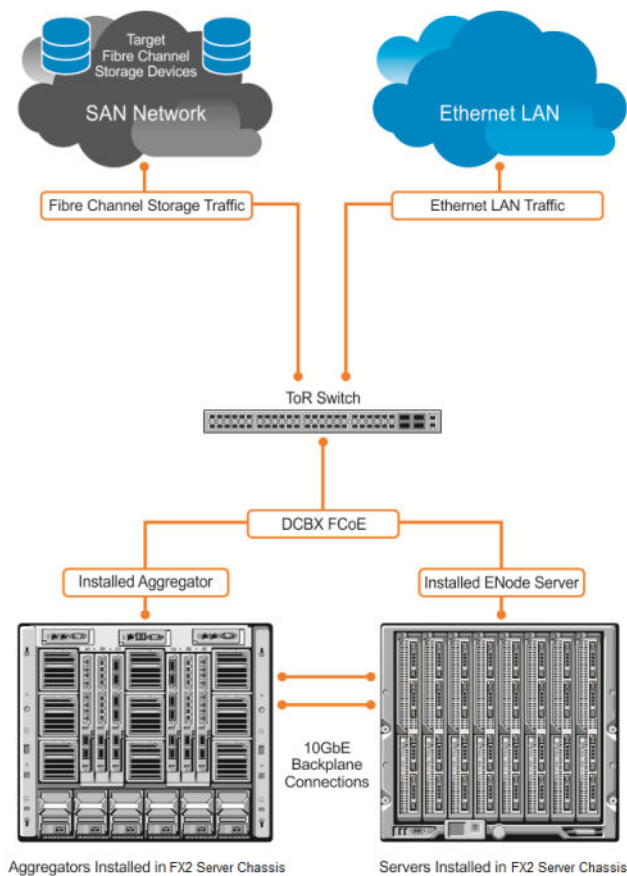


Figure 8. FIP Snooping on an Aggregator

The following sections describes how to configure the FIP snooping feature on a switch that functions as a FIP snooping bridge so that it can perform the following functions:

- Performs FIP snooping (allowing and parsing FIP frames) globally on all VLANs or on a per-VLAN basis.
- Set the FCoE MAC address prefix (FC-MAP) value used by an FCF to assign a MAC address to an ECoE end-device (server ENode or storage device) after a server successfully logs in.
- Set the FCF mode to provide additional port security on ports that are directly connected to an FCF.
- Check FIP snooping-enabled VLANs to ensure that they are operationally active.
- Process FIP VLAN discovery requests and responses, advertisements, solicitations, FLOGI/FDISC requests and responses, FLOGO requests and responses, keep-alive packets, and clear virtual-link messages.

How FIP Snooping is Implemented

As soon as the Aggregator is activated in an Dell PowerEdge FX2 server chassis as a switch-bridge, existing VLAN—specific and FIP snooping auto-configurations are applied. The Aggregator snoops FIP packets on VLANs enabled for FIP snooping and allows legitimate sessions. By default, all FCoE and FIP frames are dropped unless specifically permitted by existing FIP snooping-generated ACLs.

FIP Snooping on VLANs

FIP snooping is enabled globally on an Aggregator on all VLANs:

- FIP frames are allowed to pass through the switch on the enabled VLANs and are processed to generate FIP snooping ACLs.
- FCoE traffic is allowed on VLANs only after a successful virtual-link initialization (fabric login FLOGI) between an ENode and an FCF. All other FCoE traffic is dropped.
- At least one interface is auto-configured for FCF (FIP snooping bridge — FCF) mode on a FIP snooping-enabled VLAN. Multiple FCF trusted interfaces are auto-configured in a VLAN.
- A maximum of eight VLANs are supported for FIP snooping on an Aggregator. FIP snooping processes FIP packets in traffic only from the first eight incoming VLANs.

FC-MAP Value

The FC-MAP value that is applied globally by the Aggregator on all FCoE VLANs to authorize FCoE traffic is auto-configured.

The FC-MAP value is used to check the FC-MAP value for the MAC address assigned to ENodes in incoming FCoE frames. If the FC-MAP values does not match, FCoE frames are dropped. A session between an ENode and an FCF is established by the switch —bridge only when the FC-MAP value on the FCF matches the FC-MAP value on the FIP snooping bridge.

Bridge-to-FCF Links

A port directly connected to an FCF is auto-configured in FCF mode. Initially, all FCoE traffic is blocked; only FIP frames are allowed to pass.

FCoE traffic is allowed on the port only after a successful FLOGI request/response and confirmed use of the configured FC-MAP value for the VLAN.

Impact on other Software Features

FIP snooping affects other software features on an Aggregator as follows:

- MAC address learning: MAC address learning is not performed on FIP and FCoE frames, which are denied by ACLs dynamically created by FIP snooping in server-facing ports in ENode mode.
- MTU auto-configuration: MTU size is set to mini-jumbo (2500 bytes) when a port is in Switchport mode, the FIP snooping feature is enabled on the switch, and the FIP snooping is enabled on all or individual VLANs.
- Link aggregation group (LAG): FIP snooping is supported on port channels on ports on which PFC mode is on (PFC is operationally up).

FIP Snooping Prerequisites

On an Aggregator, FIP snooping requires the following conditions:

- A FIP snooping bridge requires DCBX and PFC to be enabled on the switch for lossless Ethernet connections (refer to **Data Center Bridging (DCB)**). Dell recommends that you also enable ETS; ETS is recommended but not required. DCBX and PFC mode are auto-configured on Aggregator ports and FIP snooping is operational on the port. If the PFC parameters in a DCBX exchange with a peer are not synchronized, FIP and FCoE frames are dropped on the port.
- VLAN membership:
 - The Aggregator auto-configures the VLANs which handle FCoE traffic. You can reconfigure VLAN membership on a port (`vlan tagged` command).

- Each FIP snooping port is auto-configured to operate in Hybrid mode so that it accepts both tagged and untagged VLAN frames.
- Tagged VLAN membership is auto-configured on each FIP snooping port that sends and receives FCoE traffic and has links with an FCF, ENode server or another FIP snooping bridge.
- The default VLAN membership of the port should continue to operate with untagged frames. FIP snooping is not supported on a port that is configured for non-default untagged VLAN membership.

FIP Snooping Restrictions

The following restrictions apply to FIP snooping on an Aggregator:

- The maximum number of FCoE VLANs supported on the Aggregator is eight.
- The maximum number of FIP snooping sessions supported per ENode server is 32. To increase the maximum number of sessions to 64, use the `fip-snooping max-sessions-per-enodemac` command. This is configurable only in PMUX mode.
- In a full FCoE N port ID virtualization (NPIV) configuration, 16 sessions (one FLOGI + 15 NPIV sessions) are supported per ENode. In an FCoE NPV configuration, only one session is supported per ENode.
- The maximum number of FCFs supported per FIP snooping-enabled VLAN is 12.
- Links to other FIP snooping bridges on a FIP snooping-enabled port (bridge-to-bridge links) are not supported on the Aggregator.

Configuring FIP Snooping

FIP snooping is auto-configured on an Aggregator in standalone mode. You can display information on FIP snooping operation and statistics by entering `show` commands. You can enable FIP snooping globally on all FCoE VLANs on a switch or on an individual FCoE VLAN.

By default, FIP snooping is disabled.

To enable FCoE transit on the switch and configure the FCoE transit parameters on ports, follow these steps.

1. Enable the FCoE transit feature on a switch.

CONFIGURATION mode.

```
feature fip-snooping
```

2. Enable FIP snooping on all VLANs or on a specified VLAN.

CONFIGURATION mode or VLAN INTERFACE mode.

```
fip-snooping enable
```

By default, FIP snooping is disabled on all VLANs.

3. Configure the FC-MAP value used by FIP snooping on all VLANs.

CONFIGURATION VLAN or INTERFACE mode

```
fip-snooping fc-map fc-map-value
```

The default is 0x0EFC00.

The valid values are from 0EFC00 to 0EFCFF.

4. Enter interface configuration mode to configure the port for FIP snooping links.


CONFIGURATION mode


```
interface port-type slot/port
```

By default, a port is configured for bridge-to-ENode links.

5. Configure the port for bridge-to-FCF links.
INTERFACE or CONFIGURATION mode

```
fip-snooping port-mode fcf
```

 **NOTE:** All these configurations are available only in PMUX mode.

-  **NOTE:** To disable the FIP snooping feature or FIP snooping on VLANs, use the no version of a command; for example, no feature fip-snooping or no fip-snooping enable.

Displaying FIP Snooping Information

Use the show commands from the table below, to display information on FIP snooping.

Command	Output
<pre>show fip-snooping sessions [interface vlan vlan-id]</pre>	Displays information on FIP-snooped sessions on all VLANs or a specified VLAN, including the ENode interface and MAC address, the FCF interface and MAC address, VLAN ID, FCoE MAC address and FCoE session ID number (FC-ID), worldwide node name (WWNN) and the worldwide port name (WWPN). Information on NPIV sessions is also displayed.
<pre>show fip-snooping config</pre>	Displays the FIP snooping status and configured FC-MAP values.
<pre>show fip-snooping enode [enode- mac-address]</pre>	Displays information on the ENodes in FIP-snooped sessions, including the ENode interface and MAC address, FCF MAC address, VLAN ID and FC-ID.
<pre>show fip-snooping fcf [fcf-mac- address]</pre>	Displays information on the FCFs in FIP-snooped sessions, including the FCF interface and MAC address, FCF interface, VLAN ID, FC-MAP value, FKA advertisement period, and number of ENodes connected.
<pre>clear fip-snooping database interface vlan vlan-id {fcoe-mac- address enode-mac-address fcf- mac-address}</pre>	Clears FIP snooping information on a VLAN for a specified FCoE MAC address, ENode MAC address, or FCF MAC address, and removes the corresponding ACLs generated by FIP snooping.
<pre>show fip-snooping statistics [interface vlan vlan-id interface port-type port/slot interface port-channel port- channel-number]</pre>	Displays statistics on the FIP packets snooped on all interfaces, including VLANs, physical ports, and port channels.

```
clear fip-snooping statistics
[interface vlan vlan-id |
interface port-type port/slot |
interface port-channel port-
channel-number]
```

Clears the statistics on the FIP packets snooped on all VLANs, a specified VLAN, or a specified port interface.

```
show fip-snooping system
```

Display information on the status of FIP snooping on the switch (enabled or disabled), including the number of FCoE VLANs, FCFs, ENodes, and currently active sessions.

```
show fip-snooping vlan
```

Display information on the FCoE VLANs on which FIP snooping is enabled.

show fip-snooping sessions **Command Example**

```
Dell#show fip-snooping sessions
Enode MAC          Enode Intf      FCF MAC        FCF Intf VLAN FCoE MAC          FC-ID
00:0e:1e:0c:54:a6 Te 0/1  00:05:73:f2:4f:ae Po128 100 0e:fc:00:9a:00:27 9a:
00:27 20:01:00:0e:1e:0c:54:a6
00:0e:1e:06:01:5e Te 0/3  00:05:73:f2:4f:af Po128 100 0e:fc:00:9a:01:18 9a:
01:18 20:01:00:0e:1e:06:01:5
Port WWNN
20:00:00:0e:1e:0c:54:a6
20:00:00:0e:1e:0c:54:a6
```

show fip-snooping sessions **Command Description**

Field	Description
ENode MAC	MAC address of the ENode.
ENode Interface	Slot/ port number of the interface connected to the ENode.
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/ port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FCoE MAC	MAC address of the FCoE session assigned by the FCF.
FC-ID	Fibre Channel ID assigned by the FCF.
Port WWPN	Worldwide port name of the CNA port.
Port WWNN	Worldwide node name of the CNA port.

show fip-snooping config **Command Example**

```
Dell#show fip-snooping config
FIP Snooping Feature enabled Status: Enabled
FIP Snooping Global enabled Status: Enabled
Global FC-MAP Value: 0X0EFC00
Maximum Sessions Per Enode Mac: 32
Dell#
```

show fip-snooping enode **Command Example**

```
Dell# show fip-snooping enode
ENode MAC          ENode Interface      FCF MAC
VLAN              FC-ID                -----
-----
d4:ae:52:1b:e3:cd  Te 0/1              54:7f:ee:37:34:40
100                62:00:11
```

show fip-snooping enode **Command Description**

Field	Description
ENode MAC	MAC address of the ENode.
ENode Interface	Slot/ port number of the interface connected to the ENode.
FCF MAC	MAC address of the FCF.
VLAN	VLAN ID number used by the session.
FC-ID	Fibre Channel session ID assigned by the FCF.

show fip-snooping fcf **Command Example**

```
Dell# show fip-snooping fcf
FCF MAC          FCF Interface      VLAN      FC-MAP      FKA_ADV_PERIOD
No. of Enodes
-----
54:7f:ee:37:34:40  Po 22              100       0e:fc:00    4000
2
```

show fip-snooping fcf **Command Description**

Field	Description
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FC-MAP	FC-Map value advertised by the FCF.
ENode Interface	Slot/ number of the interface connected to the ENode.
FKA_ADV_PERIOD	Period of time (in milliseconds) during which FIP keep-alive advertisements are transmitted.
No of ENodes	Number of ENodes connected to the FCF.

FC-ID

Fibre Channel session ID assigned by the FCF.

show fip-snooping statistics (VLAN and port) Command Example

```
Dell# show fip-snooping statistics interface vlan 100
Number of Vlan Requests :0
Number of Vlan Notifications :0
Number of Multicast Discovery Solicits :2
Number of Unicast Discovery Solicits :0
Number of FLOGI :2
Number of FDISC :16
Number of FLOGO :0
Number of Enode Keep Alive :9021
Number of VN Port Keep Alive :3349
Number of Multicast Discovery Advertisement :4437
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts :2
Number of FLOGI Rejects :0
Number of FDISC Accepts :16
Number of FDISC Rejects :0
Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0
Dell(conf)#
```

```
Dell# show fip-snooping statistics int tengigabitethernet 0/1
Number of Vlan Requests :1
Number of Vlan Notifications :0
Number of Multicast Discovery Solicits :1
Number of Unicast Discovery Solicits :0
Number of FLOGI :1
Number of FDISC :16
Number of FLOGO :0
Number of Enode Keep Alive :4416
Number of VN Port Keep Alive :3136
Number of Multicast Discovery Advertisement :0
Number of Unicast Discovery Advertisement :0
Number of FLOGI Accepts :0
Number of FLOGI Rejects :0
Number of FDISC Accepts :0
Number of FDISC Rejects :0
Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0
```

show fip-snooping statistics (port channel) Command Example

```
Dell# show fip-snooping statistics interface port-channel 22
Number of Vlan Requests :0
Number of Vlan Notifications :2
Number of Multicast Discovery Solicits :0
Number of Unicast Discovery Solicits :0
Number of FLOGI :0
Number of FDISC :0
Number of FLOGO :0
Number of Enode Keep Alive :0
```



```

Number of VN Port Keep Alive :0
Number of Multicast Discovery Advertisement :4451
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts :2
Number of FLOGI Rejects :0
Number of FDISC Accepts :16
Number of FDISC Rejects :0
Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0

```


show fip-snooping statistics **Command Description**

Field	Description
Number of Vlan Requests	Number of FIP-snooped VLAN request frames received on the interface.
Number of VLAN Notifications	Number of FIP-snooped VLAN notification frames received on the interface.
Number of Multicast Discovery Solicits	Number of FIP-snooped multicast discovery solicit frames received on the interface.
Number of Unicast Discovery Solicits	Number of FIP-snooped unicast discovery solicit frames received on the interface.
Number of FLOGI	Number of FIP-snooped FLOGI request frames received on the interface.
Number of FDISC	Number of FIP-snooped FDISC request frames received on the interface.
Number of FLOGO	Number of FIP-snooped FLOGO frames received on the interface.
Number of ENode Keep Alives	Number of FIP-snooped ENode keep-alive frames received on the interface.
Number of VN Port Keep Alives	Number of FIP-snooped VN port keep-alive frames received on the interface.
Number of Multicast Discovery Advertisements	Number of FIP-snooped multicast discovery advertisements received on the interface.
Number of Unicast Discovery Advertisements	Number of FIP-snooped unicast discovery advertisements received on the interface.
Number of FLOGI Accepts	Number of FIP FLOGI accept frames received on the interface.
Number of FLOGI Rejects	Number of FIP FLOGI reject frames received on the interface.
Number of FDISC Accepts	Number of FIP FDISC accept frames received on the interface.

Number of FDISC Rejects	Number of FIP FDISC reject frames received on the interface.
Number of FLOGO Accepts	Number of FIP FLOGO accept frames received on the interface.
Number of FLOGO Rejects	Number of FIP FLOGO reject frames received on the interface.
Number of CVLs	Number of FIP clear virtual link frames received on the interface.
Number of FCF Discovery Timeouts	Number of FCF discovery timeouts that occurred on the interface.
Number of VN Port Session Timeouts	Number of VN port session timeouts that occurred on the interface.
Number of Session failures due to Hardware Config	Number of session failures due to hardware configuration that occurred on the interface.

`show fip-snooping system` **Command Example**


```
Dell# show fip-snooping system
Global Mode           : Enabled
FCOE VLAN List (Operational) : 1, 100
FCFs                  : 1
Enodes                : 2
Sessions              : 17
```

 **NOTE:** NPIV sessions are included in the number of FIP-snooped sessions displayed.

`show fip-snooping vlan` **Command Example**

```
Dell# show fip-snooping vlan
* = Default VLAN

VLAN    FC-MAP          FCFs    Enodes    Sessions
----    -
*1      -                -       -         -
100     0X0EFC00         1       2         17
```

 **NOTE:** NPIV sessions are included in the number of FIP-snooped sessions displayed.

FIP Snooping Example

The following figure shows an Aggregator used as a FIP snooping bridge for FCoE traffic between an ENode (server blade) and an FCF (ToR switch). The ToR switch operates as an FCF and FCoE gateway.

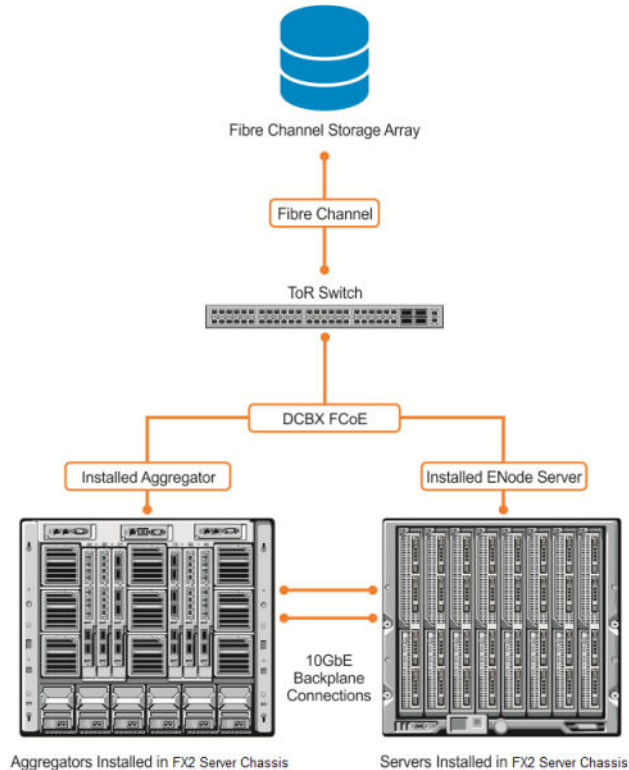


Figure 9. FIP Snooping on an Aggregator

In the above figure, DCBX and PFC are enabled on the Aggregator (FIP snooping bridge) and on the FCF ToR switch. On the FIP snooping bridge, DCBX is configured as follows:

- A server-facing port is configured for DCBX in an auto-downstream role.
- An FCF-facing port is configured for DCBX in an auto-upstream or configuration-source role.

The DCBX configuration on the FCF-facing port is detected by the server-facing port and the DCB PFC configuration on both ports is synchronized. For more information about how to configure DCBX and PFC on a port, refer to [FIP Snooping](#)

After FIP packets are exchanged between the ENode and the switch, a FIP snooping session is established. ACLS are dynamically generated for FIP snooping on the FIP snooping bridge/switch.

Debugging FIP Snooping

To enable debug messages for FIP snooping events, enter the `debug fip-snooping` command..

Task	Command	Command Mode
Enable FIP snooping debugging on for all or a specified event type, where: <code>all</code> enables all debugging options. <code>acl</code> enables debugging only for ACL-specific events. <code>error</code> enables debugging only for error conditions. <code>ifm</code> enables debugging only for IFM events. <code>info</code> enables debugging only for information events. <code>ipc</code> enables debugging only for IPC events. <code>rx</code> enables debugging only for incoming packet traffic.	<code>debug fip-snooping [all acl error ifm info ipc rx]</code>	EXEC PRIVILEGE

To turn off debugging event messages, enter the `no debug fip-snooping` command.

Internet Group Management Protocol (IGMP)

On an Aggregator, IGMP snooping is auto-configured. You can display information on IGMP by using `show ip igmp` command.

Multicast is based on identifying many hosts by a single destination IP address. Hosts represented by the same IP address are a multicast group. The internet group management protocol (IGMP) is a Layer 3 multicast protocol that hosts use to join or leave a multicast group. Multicast routing protocols (such as protocol-independent multicast [PIM]) use the information in IGMP messages to discover which groups are active and to populate the multicast routing table.

This chapter contains the following sections:

- [IGMP Overview](#)
- [IGMP Snooping](#)

IGMP Overview

IGMP has three versions. Version 3 obsoletes and is backwards-compatible with version 2; version 2 obsoletes version 1.

IGMP Version 2

IGMP version 2 improves upon version 1 by specifying IGMP Leave messages, which allows hosts to notify routers that they no longer care about traffic for a particular group. Leave messages reduce the amount of time that the router takes to stop forwarding traffic for a group to a subnet (leave latency) after the last host leaves the group. In version 1 hosts quietly leave groups, and the router waits for a query response timer several times the value of the query interval to expire before it stops forwarding traffic.

To receive multicast traffic from a particular source, a host must join the multicast group to which the source is sending traffic. A host that is a member of a group is called a "receiver." A host may join many groups, and may join or leave any group at any time. A host joins and leaves a multicast group by sending an IGMP message to its IGMP querier. The querier is the router that surveys a subnet for multicast receivers and processes survey responses to populate the multicast routing table.

IGMP messages are encapsulated in IP packets which is as illustrated below:

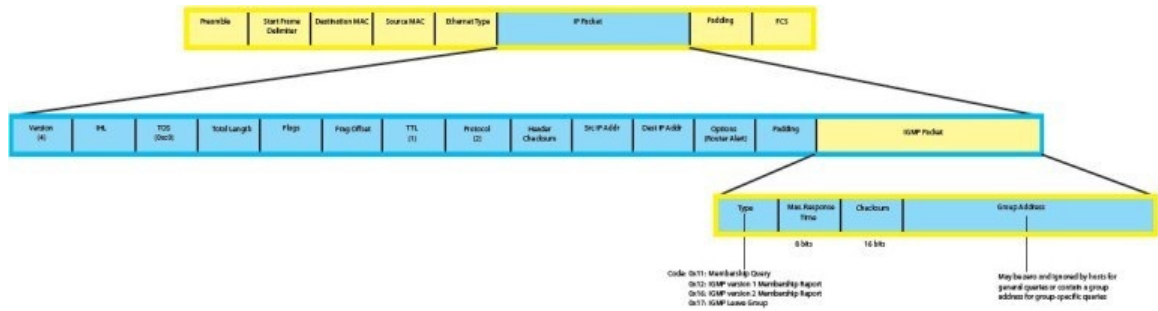


Figure 10. IGMP Version 2 Packet Format

Joining a Multicast Group

There are two ways that a host may join a multicast group: it may respond to a general query from its querier, or it may send an unsolicited report to its querier.

- Responding to an IGMP Query.
 - One router on a subnet is elected as the querier. The querier periodically multicasts (to all-multicast-systems address 224.0.0.1) a general query to all hosts on the subnet.
 - A host that wants to join a multicast group responds with an IGMP membership report that contains the multicast address of the group it wants to join (the packet is addressed to the same group). If multiple hosts want to join the same multicast group, only the report from the first host to respond reaches the querier, and the remaining hosts suppress their responses (for how the delay timer mechanism works, refer to [IGMP Snooping](#)).
 - The querier receives the report for a group and adds the group to the list of multicast groups associated with its outgoing port to the subnet. Multicast traffic for the group is then forwarded to that subnet.
- Sending an Unsolicited IGMP Report.
 - A host does not have to wait for a general query to join a group. It may send an unsolicited IGMP membership report, also called an IGMP Join message, to the querier.

Leaving a Multicast Group

- A host sends a membership report of type 0x17 (IGMP Leave message) to the all routers multicast address 224.0.0.2 when it no longer cares about multicast traffic for a particular group.
- The querier sends a group-specific query to determine whether there are any remaining hosts in the group. There must be at least one receiver in a group on a subnet for a router to forward multicast traffic for that group to the subnet.
- Any remaining hosts respond to the query according to the delay timer mechanism (refer to [IGMP Snooping](#)). If no hosts respond (because there are none remaining in the group), the querier waits a specified period and sends another query. If it still receives no response, the querier removes the group from the list associated with forwarding port and stops forwarding traffic for that group to the subnet.

IGMP Version 3

Conceptually, IGMP version 3 behaves the same as version 2. However, there are differences:

- Version 3 adds the ability to filter by multicast source, which helps the multicast routing protocols avoid forwarding traffic to subnets where there are no interested receivers.

- To enable filtering, routers must keep track of more state information, that is, the list of sources that must be filtered. An additional query type, the group-and-source-specific query, keeps track of state changes, while the group-specific and general queries still refresh existing state.
- Reporting is more efficient and robust. Hosts do not suppress query responses (non-suppression helps track state and enables the immediate-leave and IGMP snooping features), state-change reports are retransmitted to insure delivery, and a single membership report bundles multiple statements from a single host, rather than sending an individual packet for each statement.

To accommodate these protocol enhancements, the IGMP version 3 packet structure is different from version 2. Queries (shown below in query packet format) are still sent to the all-systems address 224.0.0.1, but reports (shown below in report packet format) are sent to all the IGMP version 3 – capable multicast routers address 224.0.0.22.

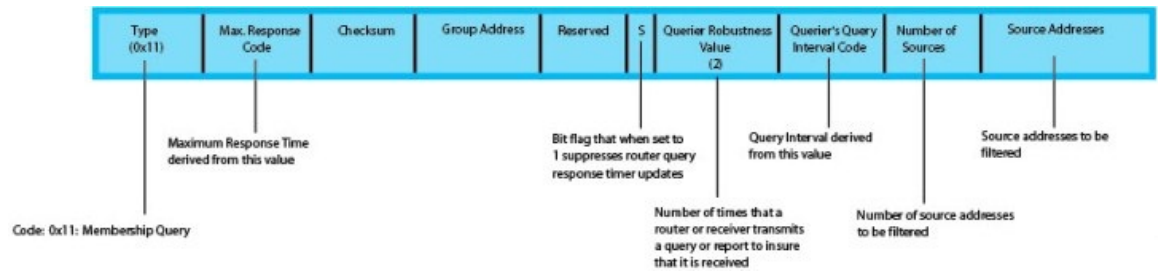


Figure 11. IGMP version 3 Membership Query Packet Format

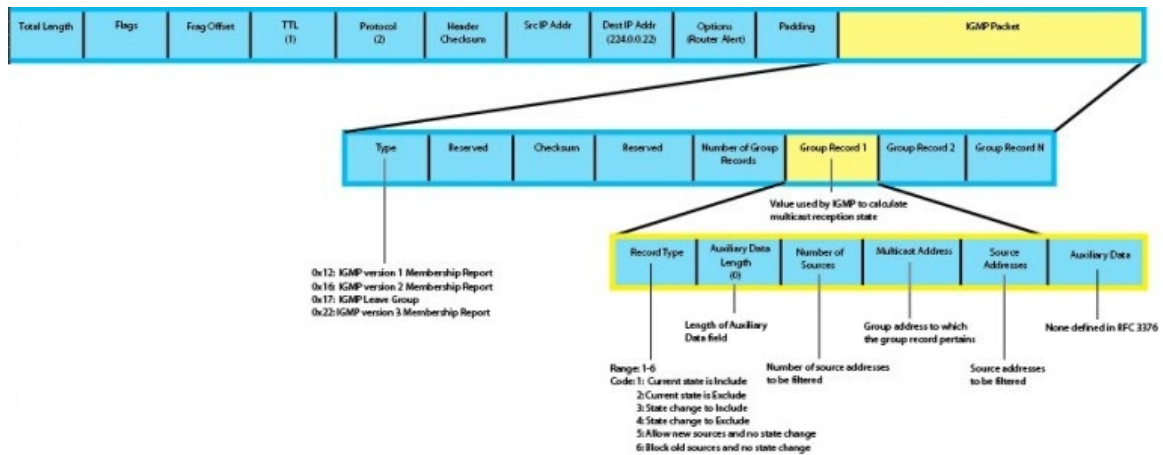


Figure 12. IGMP version 3 Membership Report Packet Format

Joining and Filtering Groups and Sources

The below illustration shows how multicast routers maintain the group and source information from unsolicited reports.

- The first unsolicited report from the host indicates that it wants to receive traffic for group 224.1.1.1.
- The host's second report indicates that it is only interested in traffic from group 224.1.1.1, source 10.11.1.1. Include messages prevent traffic from all other sources in the group from reaching the subnet, so before recording this request, the querier sends a group-and-source query to verify that there are no hosts interested in any other sources. The multicast router must satisfy all hosts if they have conflicting requests. For example, if another host on the subnet is interested in traffic from 10.11.1.3, the router cannot record the include request. There are no other interested hosts, so the request is recorded. At this point, the multicast routing protocol prunes the tree to all but the specified sources.

- The host's third message indicates that it is only interested in traffic from sources 10.11.1.1 and 10.11.1.2. Because this request again prevents all other sources from reaching the subnet, the router sends another group-and-source query so that it can satisfy all other hosts. There are no other interested hosts, so the request is recorded.

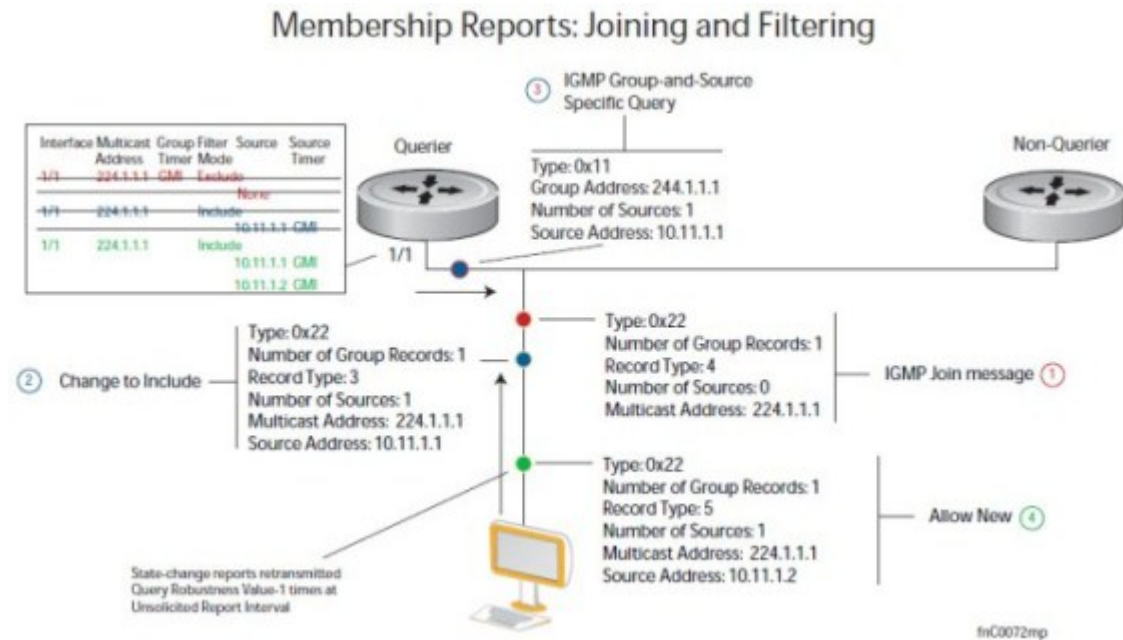


Figure 13. IGMP Membership Reports: Joining and Filtering

Leaving and Staying in Groups

The below illustration shows how multicast routers track and refreshes the state change in response to group-and-specific and general queries.

- Host 1 sends a message indicating it is leaving group 224.1.1.1 and that the included filter for 10.11.1.1 and 10.11.1.2 are no longer necessary.
- The querier, before making any state changes, sends a group-and-source query to see if any other host is interested in these two sources; queries for state-changes are retransmitted multiple times. If any are interested, they respond with their current state information and the querier refreshes the relevant state information.
- Separately in the below figure, the querier sends a general query to 224.0.0.1.
- Host 2 responds to the periodic general query so the querier refreshes the state information for that group.

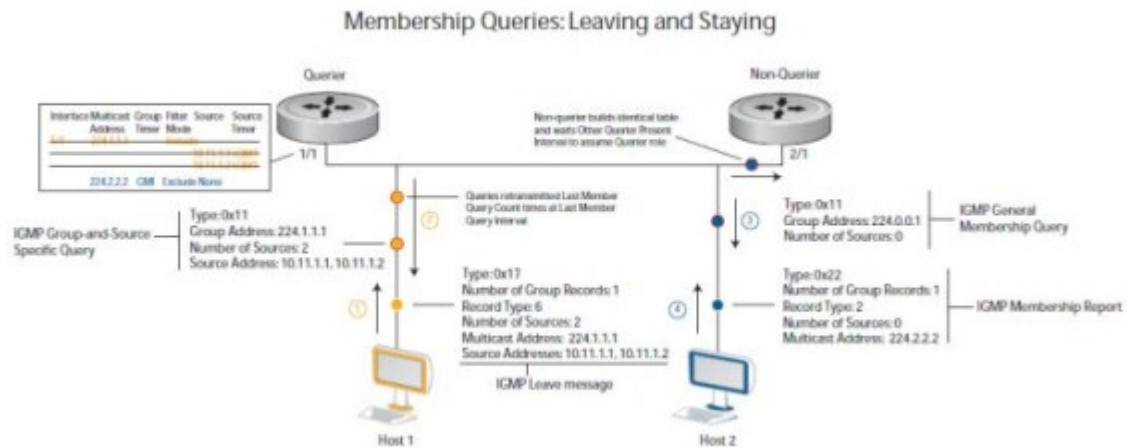


Figure 14. IGMP Membership Queries: Leaving and Staying in Groups

IGMP Snooping

IGMP snooping is auto-configured on an Aggregator.

Multicast packets are addressed with multicast MAC addresses, which represents a group of devices rather than one unique device. Switches forward multicast frames out of all ports in a VLAN by default, even if there are only a small number of interested hosts, resulting in a waste of bandwidth. IGMP snooping enables switches to use information in IGMP packets to generate a forwarding table that associate ports with multicast groups, so that the received multicast frames are forwarded only to interested receivers.

How IGMP Snooping is Implemented on an Aggregator

- IGMP snooping is enabled by default on the switch.
- Dell Networking OS supports version 1, version 2 and version 3 hosts.
- Dell Networking OS — IGMP snooping is based on the IP multicast address (not on the Layer 2 multicast MAC address). IGMP snooping entries are stored in the Layer 3 flow table instead of in the Layer 2 forwarding information base (FIB).
- Dell Networking OS — IGMP snooping is based on draft-ietf-magma-snoop-10.
- A maximum of 2k groups and 4k virtual local area networks (VLAN) are supported.
- IGMP snooping is not supported on the default VLAN interface.
- Flooding of unregistered multicast traffic is enabled by default.
- Queries are not accepted from the server side ports and are only accepted from the uplink LAG.
- Reports and Leaves are flooded by default to the uplink LAG irrespective of whether it is an mrouter port or not.

Disabling Multicast Flooding

If the switch receives a multicast packet that has an IP address of a group it has not learned (unregistered frame), the switch floods that packet out of all ports on the VLAN. To disable multicast flooding on all VLAN ports, enter the **no ip igmp snooping flood** command in global configuration mode.

When multicast flooding is disabled, unregistered multicast data traffic is forwarded to only multicast router ports on all VLANs. If there is no multicast router port in a VLAN, unregistered multicast data traffic is dropped.

Displaying IGMP Information

Use the `show` commands from the below table, to display information on IGMP. If you specify a group address or interface:

- Enter a group address in dotted decimal format; for example, 225.0.0.0.
- Enter an interface in one of the following formats: `tengigabitethernet slot/port`, `port-channel port-channel-number`, or `vlan vlan-number`.

Displaying IGMP Information

Command	Output
<code>show ip igmp snooping groups [group-address [detail] detail interface [group-address [detail]]]</code>	Displays information on IGMP groups.
<code>show ip igmp snooping interface [interface]</code>	Displays IGMP information on IGMP-enabled interfaces.
<code>show ip igmp snooping mrouter [vlan vlan-number]</code>	Displays information on IGMP-enabled multicast router (mrouter) interfaces.
<code>clear ip igmp snooping groups [group-address interface]</code>	Clears IGMP information for group addresses and IGMP-enabled interfaces.

show ip igmp snooping groups **Command Example**

```
Dell# show ip igmp snooping groups
Total Number of Groups: 2
IGMP Connected Group Membership
Group Address      Interface          Mode      Uptime      Expires
Last Reporter
226.0.0.1          Vlan 1500         INCLUDE   00:00:19    Never
1.1.1.2
226.0.0.1          Vlan 1600         INCLUDE   00:00:02    Never
1.1.1.2
Dell#show ip igmp snooping groups detail

Interface          Vlan 1500
Group              226.0.0.1
Uptime             00:00:21
Expires            Never
Router mode        INCLUDE
Last reporter      1.1.1.2
Last reporter mode INCLUDE
Last report received IS_INCL
```

```

Group source list
Source address          Uptime    Expires
1.1.1.2                 00:00:21  00:01:48
  Member Ports: Po 1

Interface              Vlan 1600
Group                  226.0.0.1
Uptime                 00:00:04
Expires                Never
Router mode            INCLUDE
Last reporter          1.1.1.2
Last reporter mode     INCLUDE
Last report received   IS_INCL
Group source list
Source address          Uptime    Expires
1.1.1.2                 00:00:04  00:02:05
  Member Ports: Po 1
Dell#

```

show ip igmp snooping interface **Command Example**

```

Dell# show ip igmp snooping interface

Vlan 2 is up, line protocol is down
  Inbound IGMP access group is not set
  Interface IGMP group join rate limit is not set
  IGMP snooping is enabled on interface
  IGMP Snooping query interval is 60 seconds
  IGMP Snooping querier timeout is 125 seconds
  IGMP Snooping last member query response interval is 1000 ms
  IGMP snooping fast-leave is disabled on this interface
  IGMP snooping querier is disabled on this interface
Vlan 3 is up, line protocol is down
  Inbound IGMP access group is not set
  Interface IGMP group join rate limit is not set
  IGMP snooping is enabled on interface
  IGMP Snooping query interval is 60 seconds
  IGMP Snooping querier timeout is 125 seconds
  IGMP Snooping last member query response interval is 1000 ms
  IGMP snooping fast-leave is disabled on this interface
  IGMP snooping querier is disabled on this interface
--More--

```

show ip igmp snooping mrouter **Command Example**

```

Dell# show ip igmp snooping mrouter
Interface Router Ports
Vlan 1000 Po 128
Dell#

```

Interfaces

This chapter describes TenGigabit Ethernet interface types, both physical and logical, and how to configure them with the Dell Networking Operating Software (OS).

Basic Interface Configuration

- [Interface Auto-Configuration](#)
- [Interface Types](#)
- [Viewing Interface Information](#)
- [Disabling and Re-enabling a Physical Interface](#)
- [Layer 2 Mode](#)
- [Management Interfaces](#)
- [VLAN Membership](#)
- [Port Channel Interfaces](#)

Advanced Interface Configuration

- [Monitor and Maintain Interfaces](#)
- [Flow Control Using Ethernet Pause Frames](#)
- [MTU Size](#)
- [Auto-Negotiation on Ethernet Interfaces](#)
- [Viewing Interface Information](#)

Interface Auto-Configuration

An Aggregator auto-configures interfaces as follows:

- Aggregator ports are numbered 1 to 12. Ports 1 to 8 are internal server-facing interfaces. Ports 9 to 12 are uplink ports.
- All 10GbE uplink interfaces belong to the same 10GbE link aggregation group (LAG).
 - The tagged Virtual Local Area Network (VLAN) membership of the uplink LAG is automatically configured based on the VLAN configuration of all server-facing ports (ports 1 to 8). The untagged VLAN used for the uplink LAG is always the default VLAN 1.
 - The tagged VLAN membership of a server-facing LAG is automatically configured based on the server-facing ports that are members of the LAG. The untagged VLAN of a server-facing LAG is auto-configured based on the untagged VLAN to which the lowest numbered server-facing port in the LAG belongs.

- All interfaces are auto-configured as members of all (4094) VLANs and untagged VLAN 1. All VLANs are up and can send or receive layer 2 traffic. You can use the Command Line Interface (CLI) or CMC interface to configure only the required VLANs on a port interface.

Interface Types


The following interface types are supported on an Aggregator.


Interface Type	Supported Modes	Default Mode	Requires Creation	Default State
Physical	L2	10GbE uplink	No	No Shutdown (enabled)
Management	L3	L3	No	No Shutdown (enabled)
Port Channel	L2	L2	No	L2 - No Shutdown (enabled)
Default VLAN	L2 and L3	L2 and L3 (VLAN 1)	No	L2 - No Shutdown (enabled)L3 - No Shutdown (enabled)
Non-default VLANs (VLANs 2 - 4094)	L2	L2 and L3	Yes	L2 - No Shutdown (enabled)L3 - No Shutdown (enabled)

Viewing Interface Information

To view interface status and auto-configured parameters use show commands.

The `show interfaces` command in EXEC mode lists all configurable interfaces on the chassis and has options to display the interface status, IP and MAC addresses, and multiple counters for the amount and type of traffic passing through the interface. If you configure a port channel interface, the `show interfaces` command lists the interfaces configured in the port channel.

 **NOTE:** To end output from the system, such as the output from the `show interfaces` command, enter CTRL+C and the Dell Networking Operating System (OS) returns to the command prompt.

 **NOTE:** The CLI output may be incorrectly displayed as 0 (zero) for the Rx/Tx power values. Perform an simple network management protocol (SNMP) query to obtain the correct power information.

The following example shows the configuration and status information for one interface.

```
Dell#show interface tengig 0/7
TenGigabitEthernet 0/7 is up, line protocol is down
Hardware is DellEth, address is 00:1e:c9:de:04:9c
  Current address is 00:1e:c9:de:04:9c
Server Port AdminState is N/A
Pluggable media not present
Interface index is 5313282
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :001ec9de049c
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Flowcontrol rx on tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 05:30:18
Queueing strategy: fifo
Input Statistics:
```

```

    0 packets, 0 bytes
    0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    0 Multicasts, 0 Broadcasts
    0 runts, 0 giants, 0 throttles
    0 CRC, 0 overrun, 0 discarded
Output Statistics:
    0 packets, 0 bytes, 0 underruns
    0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    0 Multicasts, 0 Broadcasts, 0 Unicasts
    0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
    Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
    Output 00.00 Mbits/sec,         0 packets/sec, 0.00% of line-rate
Time since last interface status change: 05:29:16

```

To view only configured interfaces use the `show interfaces configured` command in EXEC Privilege mode.

To determine which physical interfaces are available, use the `show running-config` command in EXEC mode. This command displays all physical interfaces available on the switch, which is as shown in the following example.

```

Dell#show running-config
Current Configuration ...
! Version 1-0(0-4)
! Last configuration change at Tue Mar 25 04:50:51 2014 by default
!
boot system stack-unit 0 primary tftp://10.11.8.12/dv-ci-stomp-tc-1-a1
!
redundancy auto-synchronize full
redundancy disable-auto-reboot stack-unit
!
logging coredump stack-unit all
!
hostname Dell
!
enable password 7 b125455cf679b208e79b910e85789edf
!
username admin password 7 1d28e9f33f99cf5c
!
enable restricted 7 e0f78bb16023c392
!
vlt domain 1
!
stack-unit 0 provision PE-FN-410S-IOA
--More--

```

Disabling and Re-enabling a Physical Interface

By default, all port interfaces on an Aggregator are operationally enabled (no shutdown) to send and receive Layer 2 traffic. You can reconfigure a physical interface to shut it down by entering the `shutdown` command. To re-enable the interface, enter the `no shutdown` command.

Step	Command Syntax	Command Mode	Purpose
1.	<code>interface <i>interface</i></code>	CONFIGURATION	Enter the keyword <code>interface</code> followed by the type of interface and <code>slot/port</code> information: <ul style="list-style-type: none">For a 10GbE interface, enter the keyword <code>TenGigabitEthernet</code> followed by the <code>slot/port</code> numbers; for example, <code>interface tengigabitethernet 0/5</code>.For the management interface on a stack-unit, enter the keyword <code>ManagementEthernet</code> followed by the <code>slot/port</code> numbers; for example, <code>interface managementethernet 0/0</code>.
2.	<code>shutdown</code>	INTERFACE	Enter the <code>shutdown</code> command to disable the interface.

To confirm that the interface is enabled, use the `show config` command in INTERFACE mode.

To leave INTERFACE mode, use the `exit` command or `end` command.

You cannot delete a physical interface.

The management IP address on the D-fabric provides a dedicated management access to the system.


The switch interfaces support Layer 2 traffic over the 10-Gigabit Ethernet interfaces. These interfaces can also become part of virtual interfaces such as VLANs or port channels.

For more information about VLANs, refer to VLANs and Port Tagging. For more information about port channels, refer to Port Channel Interfaces.

Dell Networking OS Behavior: The Aggregator uses a single MAC address for all physical interfaces.

Layer 2 Mode

On an Aggregator, physical interfaces, port channels, and VLANs auto-configure to operate in Layer 2 mode. Following example demonstrates about the basic configurations found in Layer 2 interface.

 **NOTE:** Layer 3 (Network) mode is not supported on Aggregator physical interfaces, port channels and VLANs. Only management interfaces operate in Layer 3 mode.

```
Dell(conf-if-te-0/1)#show config
!
interface TenGigabitEthernet 0/1
  mtu 12000
  portmode hybrid
  switchport
  auto vlan
!
  protocol lldp
```

```

    advertise management-tlv system-name
    dcbx port-role auto-downstream
    no shutdown
Dell (conf-if-te-0/1) #

```

To view the interfaces in Layer 2 mode, use the `show interfaces switchport` command in EXEC mode.

Management Interfaces

An Aggregator auto-configures with a DHCP-based IP address for in-band management on VLAN 1 and remote out-of-band (OOB) management.

The Aggregator management interface has both a public IP and private IP address on the internal Fabric D interface. The public IP address is exposed to the outside world for WebGUI configurations/WSMAN and other proprietary traffic. You can statically configure the public IP address or obtain the IP address dynamically using the dynamic host configuration protocol (DHCP).

Accessing an Aggregator

You can access the Aggregator using:

- Internal RS-232 using the chassis management controller (CMC). Telnet into CMC and do a connect –b switch-id to get console access to the corresponding Aggregator.
- External serial port with a universal serial bus (USB) connector (front panel): connect using the Aggregator front panel USB serial line to get console access (Labeled as USB B).
- Telnet/ssh using the public IP interface on the fabric D interface.
- CMC through the private IP interface on the fabric D interface.

The Aggregator supports the management ethernet interface as well as the standard interface on any front-end port. You can use either method to connect to the system.

Configuring a Management Interface

On the Aggregator, the dedicated management interface provides management access to the system. You can configure this interface with Dell Networking OS, but the configuration options on this interface are limited. You cannot configure gateway addresses and IP addresses if it appears in the main routing table of Dell Networking OS. In addition, the proxy address resolution protocol (ARP) is not supported on this interface.

For additional management access, the Aggregator supports the default VLAN (VLAN 1) L3 interface in addition to the public fabric D management interface. You can assign the IP address for the VLAN 1 default management interface using the setup wizard or through the CLI.

If you do not configure the default VLAN 1 in the startup configuration using the wizard or CLI, by default, the VLAN 1 management interface gets its IP address using DHCP.

To configure a management interface, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface Managementethernet <i>interface</i>	CONFIGURATION	Enter the slot and the port (0).

Slot range: 0-0

To configure an IP address on a management interface, use either of the following commands in MANAGEMENT INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>ip address <i>ip-address mask</i></code>	INTERFACE	Configure an IP address and mask on the interface. <ul style="list-style-type: none">• <i>ip-address mask</i>: enter an address in dotted-decimal format (A.B.C.D), the mask must be in /prefix format (/x)
<code>ip address dhcp</code>	INTERFACE	Acquire an IP address from the DHCP server.

To access the management interface from another LAN, you must configure the management route command to point to the management interface.

There is only one management interface for the whole stack.

To display the routing table for a given port, use the `show ip route` command from EXEC Privilege mode.

Configuring a Static Route for a Management Interface

When an IP address used by a protocol and a static management route exists for the same prefix, the protocol route takes precedence over the static management route.

To configure a static route for the management port, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>management route <i>ip-address mask {forwarding-router-address ManagementEthernet slot/port}</i></code>	CONFIGURATION	Assign a static route to point to the management interface or forwarding router.

To view the configured static routes for the management port, use the `show ip management-route` command in EXEC privilege mode.

```
Dell#show ip management-route all
```

Destination	Gateway	State	Route Source
-----	-----	-----	-----
10.11.0.0/16	ManagementEthernet 0/0	Connected	Connected

```
Dell#
```


VLAN Membership

A virtual LAN (VLANs) is a logical broadcast domain or logical grouping of interfaces in a LAN in which all data received is kept locally and broadcast to all members of the group. In Layer 2 mode, VLANs move traffic at wire speed and can span multiple devices. Dell Networking OS supports up to 4093 port-based VLANs and one default VLAN, as specified in IEEE 802.1Q.

VLAN provide the following benefits:

- Improved security because you can isolate groups of users into different VLANs.
- Ability to create one VLAN across multiple devices.

On an Aggregator in standalone mode, all ports are configured by default as members of all (4094) VLANs, including the default VLAN. All VLANs operate in Layer 2 mode. You can reconfigure the VLAN membership for individual ports by using the `vlan tagged` or `vlan untagged` commands in INTERFACE configuration mode (Configuring VLAN Membership). Physical Interfaces and port channels can be members of VLANs.

 **NOTE:** You can assign a static IP address to default VLAN 1 using the `ip address` command. To assign a different VLAN ID to the default VLAN, use the `default vlan-id vlan-id` command.

Following table lists out the VLAN defaults in Dell Networking OS:

Feature	Default
Mode	Layer 2 (no IP address is assigned)
Default VLAN ID	VLAN 1

Default VLAN

When an Aggregator boots up, all interfaces are up in Layer 2 mode and placed in the default VLAN as untagged interfaces. Only untagged interfaces can belong to the default VLAN.

By default, VLAN 1 is the default VLAN. To change the default VLAN ID, use the `default vlan-id <1-4094>` command in CONFIGURATION mode. You cannot delete the default VLAN.

Port-Based VLANs

Port-based VLANs are a broadcast domain defined by different ports or interfaces. In Dell Networking OS, a port-based VLAN can contain interfaces from different stack units within the chassis. Dell Networking OS supports 4094 port-based VLANs.

Port-based VLANs offer increased security for traffic, conserve bandwidth, and allow switch segmentation. Interfaces in different VLANs do not communicate with each other, adding some security to the traffic on those interfaces. Different VLANs can communicate between each other by means of IP routing. Because traffic is only broadcast or flooded to the interfaces within a VLAN, the VLAN conserves bandwidth. Finally, you can have multiple VLANs configured on one switch, thus segmenting the device

Interfaces within a port-based VLAN must be in Layer 2 mode and can be tagged or untagged in the VLAN ID.

VLANs and Port Tagging

To add an interface to a VLAN, it must be in Layer 2 mode. After you place an interface in Layer 2 mode, it is automatically placed in the default VLAN. Dell Networking OS supports IEEE 802.1Q tagging at the interface level to filter traffic. When you enable tagging, a tag header is added to the frame after the destination and source MAC addresses. The information that is preserved as the frame moves through the network. The below figure shows the structure of a frame with a tag header. The VLAN ID is inserted in the tag header.




Figure 15. Tagged Frame Format

The tag header contains some key information used by Dell Networking OS:

- The VLAN protocol identifier identifies the frame as tagged according to the IEEE 802.1Q specifications (2 bytes).
- Tag control information (TCI) includes the VLAN ID (2 bytes total). The VLAN ID can have 4,096 values, but two are reserved.

NOTE: The insertion of the tag header into the Ethernet frame increases the size of the frame to more than the 1518 bytes specified in the IEEE 802.3 standard. Some devices that are not compliant with IEEE 802.3 may not support the larger frame size.

Information contained in the tag header allows the system to prioritize traffic and to forward information to ports associated with a specific VLAN ID. Tagged interfaces can belong to multiple VLANs, while untagged interfaces can belong only to one VLAN.

Configuring VLAN Membership

By default, all Aggregator ports are member of all (4094) VLANs, including the default untagged VLAN 1. You can use the CLI or CMC interface to reconfigure VLANs only on server-facing interfaces (1–8) so that an interface has membership only in specified VLANs.

To assign an Aggregator interface in Layer 2 mode to a specified group of VLANs, use the `vlan tagged` and `vlan untagged` commands. To view which interfaces are tagged or untagged and to which VLAN they belong, use the `show vlan` command (Displaying VLAN Membership).

To reconfigure an interface as a member of only specified tagged VLANs, enter the `vlan tagged` command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>vlan tagged {vlan-id }</code>	INTERFACE	Add the interface as a tagged member of one or more VLANs, where:

vlan-id specifies a tagged VLAN number. Range: 2-4094

To reconfigure an interface as a member of only specified untagged VLANs, enter the `vlan untagged` command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>vlan untagged {vlan-id}</code>	INTERFACE	Add the interface as an untagged member of one or more VLANs, where: <i>vlan-id</i> specifies an untagged VLAN number. Range: 2-4094

If you configure additional VLAN membership and save it to the startup configuration, the new VLAN configuration takes place immediately.

Dell Networking OS Behavior: When two or more server-facing ports with VLAN membership are configured in a LAG based on the NIC teaming configuration in connected servers learned via LACP, the resulting LAG is a tagged member of all the configured VLANs and an untagged member of the VLAN to which the port with the lowest port ID belongs. For example, if port 0/3 is an untagged member of VLAN 2 and port 0/4 is an untagged member of VLAN 3, the resulting LAG consisting of the two ports is an untagged member of VLAN 2 and a tagged member of VLAN 3.

Displaying VLAN Membership

To view the configured VLANs, enter the `show vlan` command in EXEC privilege mode:

```
Dell#show vlan
```

```
Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P - Primary, C - Community, I - Isolated
```

```
       O - Openflow
```

```
Q: U - Untagged, T - Tagged
```


```
    x - Dot1x untagged, X - Dot1x tagged
```

```
    o - OpenFlow untagged, O - OpenFlow tagged
```

```
    G - GVRP tagged, M - Vlan-stack, H - VSN tagged
```

```
    i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged
```

```
      NUM      Status      Description                Q Ports
*       1      Inactive
Dell#
```

 **NOTE:** A VLAN is active only if the VLAN contains interfaces and those interfaces are operationally up. In the above example, VLAN 1 is inactive because it does not contain any interfaces. The other VLANs listed contain enabled interfaces and are active. In a VLAN, the `shutdown` command stops Layer 3 (routed) traffic only. Layer 2 traffic continues to pass through the VLAN. If the VLAN is not a routed VLAN (that is, configured with an IP address), the `shutdown` command has no effect on VLAN traffic.

Adding an Interface to a Tagged VLAN


The following example shows you how to add a tagged interface (Te 0/2) to the VLANs.

Enter the `vlan tagged` command to add interface Te 0/2 to VLANs 2 - 4, which is as shown below.

Enter the `show config` command to verify that interface Te 0/2 is a tagged member of the VLANs.

```
Dell(conf-if-te-0/2)#vlan tagged ?
VLAN-RANGE          Comma/Hyphen separated VLAN ID set
Dell(conf-if-te-0/2)#vlan tagged 2,3-4
Dell(conf-if-te-0/2)#show config
!
interface TenGigabitEthernet 0/2
mtu 12000
vlan tagged 2-4
!
port-channel-protocol LACP
port-channel 1 mode active
!
protocol lldp
advertise management-tlv system-name
dcbx port-role auto-downstream
no shutdown
Dell(conf-if-te-0/2)#
```

Except for hybrid ports, only a tagged interface can be a member of multiple VLANs. You can assign hybrid ports to two VLANs if the port is untagged in one VLAN and tagged in all others.

 **NOTE:** When you remove a tagged interface from a VLAN (using the `no vlan tagged` command), it remains tagged only if it is a tagged interface in another VLAN. If you remove the tagged interface from the only VLAN to which it belongs, the interface is placed in the default VLAN as an untagged interface.


Adding an Interface to an Untagged VLAN

To move an untagged interfaces from the default VLAN to another VLAN, use the `vlan untagged` command as shown in the below figure.

```
Dell(conf)# interface tengigabit 0/2
Dell(conf-if-te-0/2)#vlan untagged ?
<1-4094> Untagged VLAN id
Dell(conf-if-te-0/2)#
Dell(conf-if-te-0/2)#vlan untagged 4094
Dell(conf-if-te-0/2)#show config
!
interface TenGigabitEthernet 0/2
mtu 12000
vlan untagged 4094
!
port-channel-protocol LACP
port-channel 1 mode active
!
protocol lldp
advertise management-tlv system-name
dcbx port-role auto-downstream
no shutdown
Dell(conf-if-te-0/2)#
```

VLAN Configuration on Physical Ports and Port-Channels

Unlike other Dell Networking OS platforms, IOA allows VLAN configurations on port and port-channel levels. This allows you to assign VLANs to a port/port-channel.

 **NOTE:** In PMUX mode, in order to avoid loops, only disjoint VLANs are allowed between the uplink ports/uplink LAGs and uplink-to-uplink switching is disabled.

1. Initialize the port with configurations such as admin up, portmode, and switchport.

```
Dell#configure
Dell(conf)#int tengigabitethernet 0/1
Dell(conf-if-te-0/1)#no shutdown
Dell(conf-if-te-0/1)#portmode hybrid
Dell(conf-if-te-0/1)#switchport
```

2. Configure the tagged VLANs 10 through 15 and untagged VLAN 20 on this port.

```
Dell(conf-if-te-0/1)#vlan tagged 10-15
Dell(conf-if-te-0/1)#vlan untagged 20
Dell(conf-if-te-0/1)#
```

3. Show the running configurations on this port.

```
Dell(conf-if-te-0/1)#show config
!
interface TenGigabitEthernet 0/1
portmode hybrid
switchport
vlan tagged 10-15
vlan untagged 20
no shutdown
Dell(conf-if-te-0/1)#end
Dell#
```

4. Initialize the port-channel with configurations such as admin up, portmode, and switchport.

```
Dell#configure
Dell(conf)#int port-channel 128
Dell(conf-if-po-128)#portmode hybrid
Dell(conf-if-po-128)#switchport
```

5. Configure the tagged VLANs 10 through 15 and untagged VLAN 20 on this port-channel.

```
Dell(conf-if-po-128)#vlan tagged 10-15
Dell(conf-if-po-128)#
Dell(conf-if-po-128)#vlan untagged 20
```

6. Show the running configurations on this port-channel.

```
Dell(conf-if-po-128)#show config
!
interface Port-channel 128
portmode hybrid
switchport
vlan tagged 10-15
vlan untagged 20
shutdown
Dell(conf-if-po-128)#end
Dell#
```

7. Show the VLAN configurations.

```
Dell#show vlan
Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port
```

Mirroring VLANs, P - Primary, C - Community, I - Isolated
 O - Openflow
 Q: U - Untagged, T - Tagged
 x - Dot1x untagged, X - Dot1x tagged
 o - OpenFlow untagged, O - OpenFlow tagged
 G - GVRP tagged, M - Vlan-stack, H - VSN tagged
 i - Internal untagged, I - Internal tagged, v - VLT
 untagged, V - VLT tagged

NUM	Status	Description	Q	Ports
* 1	Active		U	Te 0/3
10	Active		T	Pol28 (Te 0/4-5) T Te 0/1
11	Active		T	Pol28 (Te 0/4-5)
12	Active		T	Pol28 (Te 0/4-5) T Te 0/1
13	Active		T	Pol28 (Te 0/4-5) T Te 0/1
14	Active		T	Pol28 (Te 0/4-5) T Te 0/1
15	Active		T	Pol28 (Te 0/4-5) T Te 0/1
20	Active		U	Pol28 (Te 0/4-5) U Te 0/1

Dell#

You can remove the inactive VLANs that have no member ports using the following command:

```
Dell#configure
Dell(conf)#no interface vlan vlan-id

vlan-id – Inactive VLAN with no member ports
```

You can remove the tagged VLANs using the `no vlan tagged vlan-range` command. You can remove the untagged VLANs using the `no vlan untagged` command in the physical port/port-channel.

Port Channel Interfaces

On an Aggregator, port channels are auto-configured as follows:

- All 10GbE uplink interfaces (ports 9 to 12) are auto-configured to belong to the same 10GbE port channel (LAG 128).
- Server-facing interfaces (ports 1 to 8) auto-configure in LAGs (1 to 127) according to the NIC teaming configuration on the connected servers.
- In VLT mode, LAG-127 is reserved for VLTi.

Port channel interfaces support link aggregation, as described in IEEE Standard 802.3ad.

 **NOTE:** A port channel may also be referred to as a *link aggregation group* (LAG).

Port Channel Definitions and Standards

Link aggregation is defined by IEEE 802.3ad as a method of grouping multiple physical interfaces into a single logical interface—a link aggregation group (LAG) or port channel. A LAG is “a group of links that appear to a MAC client as if they were a single link” according to IEEE 802.3ad. In Dell Networking OS, a LAG is referred to as a port channel interface.

A port channel provides redundancy by aggregating physical interfaces into one logical interface. If one physical interface goes down in the port channel, another physical interface carries the traffic.

Port Channel Benefits

A port channel interface provides many benefits, including easy management, link redundancy, and sharing.

Port channels are transparent to network configurations and can be modified and managed as one interface.

With this feature, you can create larger-capacity interfaces by utilizing a group of lower-speed links. For example, you can build a 40-Gigabit interface by aggregating four 10-Gigabit Ethernet interfaces together. If one of the four interfaces fails, traffic is redistributed across the three remaining interfaces.

Port Channel Implementation

An Aggregator supports only port channels that are dynamically configured using the link aggregation control protocol (LACP). For more information, refer to Link Aggregation. Statically-configured port channels are not supported.

The table below lists out the number of port channels per platform.

Platform	Port-channels	Members/Channel
FN I/O Aggregator	128	4

As soon as a port channel is auto-configured, the Dell Networking OS treats it like a physical interface. For example, IEEE 802.1Q tagging is maintained while the physical interface is in the port channel.

Member ports of a LAG are added and programmed into hardware in a predictable order based on the port ID, instead of in the order in which the ports come up. With this implementation, load balancing yields predictable results across switch resets and chassis reloads.

A physical interface can belong to only one port channel at a time.

Each port channel must contain interfaces of the same interface type/speed.

Port channels can contain a mix of 1000 or 10000 Mbps Ethernet interfaces . The interface speed (100, 1000, or 10000 Mbps) used by the port channel is determined by the first port channel member that is physically up. Dell Networking OS disables the interfaces that do not match the interface speed set by the first channel member. That first interface may be the first interface that is physically brought up or was physically operating when interfaces were added to the port channel. For example, if the first operational interface in the port channel is a TenGigabit Ethernet interface, all interfaces at 1000 Mbps are kept up, and all 100/1000/10000 interfaces that are not set to 1000 Mbps speed or auto negotiate are disabled.

10GbE Interface in Port Channels

When TenGigabitEthernet interfaces are added to a port channel, the interfaces must share a common speed. When interfaces have a configured speed different from the port channel speed, the software disables those interfaces.

The common speed is determined when the port channel is first enabled. At that time, the software checks the first interface listed in the port channel configuration. If that interface is enabled, its speed

configuration becomes the common speed of the port channel. If the other interfaces configured in that port channel are configured with a different speed, Dell Networking OS disables them.

For example, if four interfaces (TenGig 0/1, 0/2, 0/3 and 0/4) in which TenGig 0/1 and TenGig 0/2 are set to speed 1000 Mb/s and the TenGig 0/3 and TenGig 0/4 are set to 10000 Mb/s, with all interfaces enabled, and you add them to a port channel by entering `channel-member tengigabitethernet 0/1-4` while in port channel interface mode, and the Dell Networking OS determines if the first interface specified (TenGig 0/1) is up. After it is up, the common speed of the port channel is 1000 Mb/s. Dell Networking OS disables those interfaces configured with speed 10000 Mb/s or whose speed is 10000 Mb/s as a result of auto-negotiation. The `channel-member` command is available only in PMUX mode.

In this example, you can change the common speed of the port channel by changing its configuration so the first enabled interface referenced in the port level is a 1000 Mb/s speed interface. You can also change the common speed of the port channel by setting the speed of the TenGig 0/1 interface to 1000 Mb/s.

Uplink Port Channel: VLAN Membership

The tagged VLAN membership of the uplink LAG is automatically configured based on the VLAN configuration of all server-facing ports (ports 1 to 8).

The untagged VLAN used for the uplink LAG is always the default VLAN 1.

Server-Facing Port Channel: VLAN Membership

The tagged VLAN membership of a server-facing LAG is automatically configured based on the server-facing ports that are members of the LAG.

The untagged VLAN of a server-facing LAG is auto-configured based on the untagged VLAN to which the lowest numbered server-facing port in the LAG belongs.

Displaying Port Channel Information

To view the port channel's status and channel members in a tabular format, use the `show interfaces port-channel brief` command in EXEC Privilege mode.

```
Dell#sh int port-channel brief
Codes: L - LACP Port-channel
       O - OpenFlow Controller Port-channel

   LAG  Mode  Status      Uptime      Ports
L   1    L2    up          00:00:19    Te 0/7      (Up)
                   Te 0/8      (Up)
L  128  L2    up          00:00:36    Te 0/9      (Up)
                   Te 0/10     (Up)
                   Te 0/11     (Up)

Dell#
```

To display detailed information on a port channel, enter the `show interfaces port-channel` command in EXEC Privilege mode. The below example shows the port channel's mode (L2 for Layer 2, L3 for Layer 3, and L2L3 for a Layer 2 port channel assigned to a routed VLAN), the status, and the number of interfaces belonging to the port channel.

In this example, the Port-channel 1 is a dynamically created port channel based on the NIC teaming configuration in connected servers learned via LACP. Also, the Port-channel 128 is the default port channel to which all the uplink ports are assigned by default.


```
Dell#show interfaces port-channel
Port-channel 128 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:1e:c9:de:04:9c, Current address is 00:1e:c9:de:04:9c
Interface index is 1107492992
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :001ec9de049c
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed 30000 Mbit
Members in this channel: Te 0/9(U) Te 0/10(U) Te 0/11(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 04:44:48
Queueing strategy: fifo
Input Statistics:
    10063 packets, 749248 bytes
    8419 64-byte pkts, 0 over 64-byte pkts, 1644 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    10063 Multicasts, 0 Broadcasts
    0 runs, 0 giants, 0 throttles
    0 CRC, 0 overrun, 0 discarded
Output Statistics:
    61970 packets, 7743149 bytes, 0 underruns
    0 64-byte pkts, 12741 over 64-byte pkts, 48946 over 127-byte pkts
    283 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    61687 Multicasts, 283 Broadcasts, 0 Unicasts
    0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
    Input 00.00 Mbits/sec,          1 packets/sec, 0.00% of line-rate
    Output 00.00 Mbits/sec,         4 packets/sec, 0.00% of line-rate
Time since last interface status change: 04:43:55
Dell#
```


Interface Range

An interface range is a set of interfaces to which other commands may be applied, and may be created if there is at least one valid interface within the range. Bulk configuration excludes from configuring any non-existing interfaces from an interface range. A default VLAN may be configured only if the interface range being configured consists of only VLAN ports.

The `interface range` command allows you to create an interface range allowing other commands to be applied to that range of interfaces.

The interface range prompt offers the interface (with slot and port information) for valid interfaces. The maximum size of an interface range prompt is 12. If the prompt size exceeds this maximum, it displays (...) at the end of the output.

 **NOTE:** Non-existing interfaces are excluded from interface range prompt.

 **NOTE:** When creating an interface range, interfaces appear in the order they were entered and are not sorted.

To display all interfaces that have been validated under the interface range context, use the `show range` in Interface Range mode.

To display the running configuration only for interfaces that are part of interface range, use the `show configuration` command in Interface Range mode.

Bulk Configuration Examples

The following are examples of using the interface range command for bulk configuration:

- [Create a Single-Range](#)
- [Create a Multiple-Range](#)
- [Exclude a Smaller Port Range](#)
- [Overlap Port Ranges](#)
- [Commas](#)

Create a Single-Range

Creating a Single-Range Bulk Configuration

```
Dell(conf)# interface range tengigabitethernet 0/1 - 5
Dell(conf-if-range-te-0/1-5)# no shutdown
Dell(conf-if-range-te-0/1-5)#
```

Create a Multiple-Range

Creating a Multiple-Range Prompt

```
Dell(conf)#interface range tengigabitethernet 0/5 - 10 , tengigabitethernet
0/1 , vlan 1
Dell(conf-if-range-te-0/5-10,te-0/1,vl-1)#
```

Exclude a Smaller Port Range

If the interface range has multiple port ranges, the smaller port range is excluded from the prompt.

Interface Range Prompt Excluding a Smaller Port Range

```
Dell(conf)#interface range tengigabitethernet 0/1 - 2 , tengigab 0/1 - 7
Dell(conf-if-range-te-0/1-7)#
```

Overlap Port Ranges

If overlapping port ranges are specified, the port range is extended to the smallest start port number and largest end port number.

Interface Range Prompt Including Overlapping Port Ranges

```
Dell(conf)#inte ra tengig 0/1 - 3 , tengig 0/1 - 7
Dell(conf-if-range-te-0/1-7)#
```

Commas

The example below shows how to use commas to add different interface types to the range, enabling all Ten Gigabit Ethernet interfaces in the range 0/1 to 0/5.

Multiple-Range Bulk Configuration Ten-Gigabit Ethernet and Ten-Gigabit Ethernet

```
Dell(conf-if)# interface range tengigabitethernet 0/1 - 2 , tengigabitethernet
0/1 - 5
```

```
Dell(conf-if-range-te-0/1-5)# no shutdown
Dell(conf-if-range-te-0/1-5)#
```

Monitor and Maintain Interfaces

You can display interface statistics with the `monitor interface` command. This command displays an ongoing list of the interface status (up/down), number of packets, traffic statistics, etc.

Command Syntax	Command Mode	Purpose
<code>monitor interface <i>interface</i></code>	EXEC Privilege	View interface statistics. Enter the type of interface and slot/port information: <ul style="list-style-type: none"> For a 10GbE interface, enter the keyword <code>TenGigabitEthernet</code> followed by the <i>slot/port</i> numbers; for example, interface tengigabitethernet 0/7.

The information displays in a continuous run, refreshes every two seconds by default (Refer `monitor interface` command example below). Use the following keys to manage the output.

m - Change mode	c - Clear screen
l - Page up	a - Page down
T - Increase refresh interval (by 1 second)	t - Decrease refresh interval (by 1 second)
q - Quit	

monitor interface command example

```
Dell#monitor interface tengig 0/1
```

```
Dell Networking OS uptime is 1 day(s), 4 hour(s), 31 minute(s)
Monitor time: 00:00:00 Refresh Intvl.: 2s
```

```
Interface: Te 0/1, Disabled, Link is Down, Linespeed is 1000 Mbit
```

```
Traffic statistics:
  Input bytes:          Current      Rate      Delta
  Output bytes:        0          0 Bps     0
  Input packets:       0          0 pps     0
  Output packets:     0          0 pps     0
    64B packets:      0          0 pps     0
  Over 64B packets:   0          0 pps     0
  Over 127B packets:  0          0 pps     0
  Over 255B packets:  0          0 pps     0
  Over 511B packets:  0          0 pps     0
  Over 1023B packets: 0          0 pps     0
Error statistics:
  Input underruns:    0          0 pps     0
  Input giants:      0          0 pps     0
  Input throttles:   0          0 pps     0
  Input CRC:         0          0 pps     0
  Input IP checksum: 0          0 pps     0
  Input overrun:     0          0 pps     0
```

```

Output underruns:          0          0 pps          0
Output throttles:         0          0 pps          0


m - Change mode           c - Clear screen
l - Page up               a - Page down
T - Increase refresh interval
q - Quit                  t - Decrease refresh interval

```

Maintenance Using TDR

The time domain reflectometer (TDR) is supported on all Dell Networking switch/routers. TDR is an assistance tool to resolve link issues that helps detect obvious open or short conditions within any of the four copper pairs. TDR sends a signal onto the physical cable and examines the reflection of the signal that returns. By examining the reflection, TDR is able to indicate whether there is a cable fault (when the cable is broken, becomes unterminated, or if a transceiver is unplugged).

TDR is useful for troubleshooting an interface that is not establishing a link, that is, when the link is flapping or not coming up. Do not use TDR on an interface that is passing traffic. When a TDR test is run on a physical cable, it is important to shut down the port on the far end of the cable. Otherwise, it may lead to incorrect test results.

 **NOTE:** TDR is an intrusive test. Do not run TDR on a link that is up and passing traffic.

To test the condition of cables on 100/1000/10000 BASE-T modules, follow the below steps using the `tdr-cable-test` command.

Step	Command Syntax	Command Mode	Usage
1.	<code>tdr-cable-test tengigabitethernet <slot>/<port></code>	EXEC Privilege	To test for cable faults on the TenGigabitEthernet cable. <ul style="list-style-type: none"> Between two ports, you must not start the test on both ends of the cable. Enable the interface before starting the test. The port must be enabled to run the test or the test prints an error message.
2.	<code>show tdr tengigabitethernet <slot>/<port></code>	EXEC Privilege	Displays TDR test results.

Flow Control Using Ethernet Pause Frames

An Aggregator auto-configures to operate in auto-DCB-enable mode (Refer to Data Center Bridging: Auto-DCB-Enable Mode). In this mode, Aggregator ports detect whether peer devices support converged enhanced Ethernet (CEE) or not, and enable DCBX and PFC or link-level flow control accordingly:

- Interfaces come up with DCB disabled and link-level flow control enabled to control data transmission between the Aggregator and other network devices.
- When DCB is disabled on an interface, PFC, ETS, and DCBX are also disabled.
- When DCBX protocol packets are received, interfaces automatically enable DCB and disable link level flow control.

- DCB is required for PFC, ETS, DCBX, and FCoE initialization protocol (FIP) snooping to operate.

Link-level flow control uses Ethernet pause frames to signal the other end of the connection to pause data transmission for a certain amount of time as specified in the frame. Ethernet pause frames allow for a temporary stop in data transmission. A situation may arise where a sending device may transmit data faster than a destination device can accept it. The destination sends a pause frame back to the source, stopping the sender's transmission for a period of time.

The globally assigned 48-bit Multicast address 01-80-C2-00-00-01 is used to send and receive pause frames. To allow full duplex flow control, stations implementing the pause operation instruct the MAC to enable reception of frames with a destination address equal to this multicast address.

The pause frame is defined by IEEE 802.3x and uses MAC Control frames to carry the pause commands. Ethernet pause frames are supported on full duplex only. The only configuration applicable to half duplex ports is `rx off tx off`.

Note that if a port is over-subscribed, Ethernet Pause Frame flow control does not ensure no loss behavior.

The following error message appears when trying to enable flow control when half duplex is already configured:

```
Can't configure flowcontrol when half duplex is configure, config ignored.
```

The following error message appears when trying to enable half duplex and flow control configuration is on:


```
Can't configure half duplex when flowcontrol is on, config ignored.
```

Enabling Pause Frames

Enable Ethernet pause frames flow control on all ports on a chassis. If not, the system may exhibit unpredictable behavior.

 **NOTE:** If you disable `rx flow control`, Dell Networking recommends rebooting the system.

The flow control sender and receiver must be on the same port-pipe. Flow control is not supported across different port-pipes. (also refer to [iSCSI Optimization: Operation](#)).

 **NOTE:** After you disable DCB, if link-level flow control is not automatically enabled on an interface, to enable flow control, manually shut down the interface (`shutdown` command) and re-enable it (`no shutdown` command).


To enable pause frames, use the following command.

- Control how the system responds to and generates 802.3x pause frames on 10 and 40 Gig ports.
INTERFACE mode

```
flowcontrol rx [off | on] tx [off | on] [negotiate]
```

- `rx on`: enter the keywords `rx on` to process the received flow control frames on this port.
- `rx off`: enter the keywords `rx off` to ignore the received flow control frames on this port.
- `tx on`: enter the keywords `tx on` to send control frames from this port to the connected device when a higher rate of traffic is received.

- `tx off`: enter the keywords `tx off` so that flow control frames are not sent from this port to the connected device when a higher rate of traffic is received.
- `negotiate`: enable pause-negotiation with the egress port of the peer device. If the `negotiate` command is not used, pause-negotiation is disabled.

 **NOTE:** The default is `rx off`.

MTU Size

The Aggregator auto-configures interfaces to use a maximum MTU size of 12,000 bytes.

If a packet includes a Layer 2 header, the difference in bytes between the link MTU and IP MTU must be enough to include the Layer 2 header. For example, for VLAN packets, if the MTU is 1400, the link MTU must be no less than 1422.

`1400-byte IP MTU + 22-byte L2 header = 1422-byte link MTU`

The MTU range is 592-12000, with a default of 1554.

The table below lists out the various Layer 2 overheads found in Dell Networking OS and the number of bytes.

Difference between Link MTU and IP MTU

Layer 2 Overhead	Difference between Link MTU and IP MTU
Ethernet (untagged)	18 bytes
VLAN Tag	4 bytes
Untagged Packet with VLAN-Stack Header	22 bytes
Tagged Packet with VLAN-Stack Header	26 bytes

Link MTU and IP MTU considerations for port channels and VLANs are as follows.

Port Channels:

- All members must have the same link MTU value and the same IP MTU value.
- The port channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the channel members.

For example, if the members have a link MTU of 2100 and an IP MTU 2000, the port channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.

VLANs:

- All members of a VLAN must have the same IP MTU value.
- Members can have different link MTU values. Tagged members must have a link MTU 4 bytes higher than untagged members to account for the packet tag.
- The VLAN link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the VLAN members.


For example, the VLAN contains tagged members with a link MTU of 1522 and an IP MTU of 1500 and untagged members with a link MTU of 1518 and an IP MTU of 1500. The VLAN's Link MTU cannot be higher than 1518 bytes and its IP MTU cannot be higher than 1500 bytes.

Auto-Negotiation on Ethernet Interfaces

Setting Speed and Duplex Mode of Ethernet Interfaces

By default, auto-negotiation of speed and duplex mode is enabled on 10GbE Ethernet interface on an Aggregator.


The local interface and the directly connected remote interface must have the same setting. Auto-negotiation is the easiest way to accomplish these settings, as long as the remote interface is capable of auto-negotiation.

 **NOTE:** As a best practice, Dell Networking recommends keeping auto-negotiation enabled. Auto-negotiation should only be disabled on switch ports that attach to devices not capable of supporting negotiation or where connectivity issues arise from interoperability issues.

The negotiation auto command is tied to the speed command. Auto-negotiation is always enabled when the speed command is set to 1000 or auto. In Dell Networking OS, the speed 1000 command is an exact equivalent of speed auto 1000 in IOS.

To discover whether the remote and local interface require manual speed synchronization, and to manually synchronize them if necessary, follow these steps.

Step	Task	Command Syntax	Command Mode
1.	Determine the local interface status.	show interfaces [<i>interface</i>] status	EXEC Privilege
2.	Determine the remote interface status.	[Use the command on the remote system that is equivalent to the above command.]	EXEC EXEC Privilege
3.	Access CONFIGURATION mode.	config	EXEC Privilege
4.	Access the port.	interface <i>interface slot/port</i>	CONFIGURATION
5.	Set the local port speed.	speed {100 1000 10000 auto}	INTERFACE
6.	Optionally, set full- or half-duplex.	duplex {half full}	INTERFACE
7.	Disable auto-negotiation on the port. If the speed is set to 1000, you do not need to disable auto-negotiation	no negotiation auto	INTERFACE
8.	Verify configuration changes.	show config	INTERFACE

 **NOTE:** The show interfaces status command displays link status, but not administrative status. For link and administrative status, use the show ip interface [*interface* | brief] [*configuration*] command.

show interface status Command Example:

```
Dell# show interfaces status
Port      Description      Status Speed      Duplex Vlan
Te 0/1                    Up      10000 Mbit Full    1-4094
Te 0/2                    Down    Auto      Auto    1-1001,1003-4094
Te 0/3                    Up      10000 Mbit Full    1-1001,1003-4094
Te 0/4                    Down    Auto      Auto    1-1001,1003-4094
Te 0/5                    Up      10000 Mbit Full    1-4094
Te 0/6                    Up      10000 Mbit Full    1-4094
Te 0/7                    Up      10000 Mbit Full    1-4094
Te 0/8 toB300            Down    Auto      Auto    1-1001,1003-4094
Fc 0/9                    Up      8000 Mbit Full    --
Fc 0/10                   Up      8000 Mbit Full    --
Te 0/11                   Down    Auto      Auto    --
Te 0/12                   Down    Auto      Auto    --
```

In the above example, several ports display "Auto" in the speed field, including port 0/1. Now, in the below example, the speed of port 0/1 is set to 100 Mb and then its auto-negotiation is disabled.

Setting Port Speed Example

```
Dell#configure
Dell(conf)#interface tengig 0/1
Dell(conf-if-te-0/1)#speed 1000

Dell(conf-if-te-0/1)#no negotiation auto
Dell(conf-if-te-0/1)#show config
!
interface TenGigabitEthernet 0/1
speed 1000
no shutdown
```

Setting Auto-Negotiation Options

Auto-Negotiation, Speed, and Duplex Settings on Different Optics:

Command	Mode	10GbaseT module	10G SFP+ optics	1G SFP optics	Copper SFP - 1000baseT	Comments
speed 100	interface-config mode	Supported	Not supported(Error message is thrown) (% Error: Speed 100 not supported on this interface, config ignored Te 0/49)	Not supported(Error message is thrown) (% Error: Speed 100 not supported on this interface, config ignored Te 0/49)	% Error: Speed 100 not supported on this interface,	

speed auto	interface-config mode	Supported	Supported	Supported	Supported	Error messages not thrown wherever it says not supported
speed 1000	interface-config mode	Supported	Supported	Supported	Supported	
speed 10000	interface-config mode	Supported	Supported	Not Supported	Not supported	Error messages not thrown wherever it says not supported
negotiation auto	interface-config mode	Supported	Not supported(<i>Should some error message be thrown?</i>)	Not supported	Not supported	Error messages not thrown wherever it says not supported
duplex half	interface-config mode	Supported	CLI not available	CLI not available	Invalid Input error- CLI not available	
duplex full	interface-config mode	Supported	CLI not available	CLI not available	Invalid Input error-CLI not available	

Viewing Interface Information

Displaying Non-Default Configurations.

The `show [ip | running-config] interfaces configured` command allows you to display only interfaces that have non-default configurations are displayed.

The below example illustrates the possible `show` commands that have the available `configured` keyword.

```
Dell#show interfaces configured
Dell#show interfaces tengigabitEthernet 0 configured
Dell#show ip interface configured
Dell#show ip interface tengigabitEthernet 1 configured
Dell#show ip interface brief configured
Dell#show running-config interfaces configured
Dell#show running-config interface tengigabitEthernet 1 configured
```

In EXEC mode, `show interfaces switchport` command displays only interfaces in Layer 2 mode and their relevant configuration information. The `show interfaces switchport` command displays the

interface, whether the interface supports IEEE 802.1Q tagging or not, and the VLANs to which the interface belongs.

show interfaces switchport Command Example:

```
Dell#show interfaces switchport
```

```
Codes:  U - Untagged, T - Tagged
        x - Dot1x untagged, X - Dot1x tagged
        G - GVRP tagged, M - Trunk, H - VSN tagged
        i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT
tagged
```

```
Name: TenGigabitEthernet 0/1
802.1QTagged: False
Vlan membership:
Q      Vlans
U      1
```

```
Name: TenGigabitEthernet 0/2
802.1QTagged: False
Vlan membership:
Q      Vlans
U      1
```

```
Name: TenGigabitEthernet 0/3
802.1QTagged: False
Vlan membership:
Q      Vlans
U      1
```

```
Name: TenGigabitEthernet 0/4
802.1QTagged: False
Vlan membership:
Q      Vlans
U      1
```

```
--More--
```

Clearing Interface Counters

The counters in the `show interfaces` command are reset by the `clear counters` command. This command does not clear the counters captured by any SNMP program.

To clear the counters, use the following command in EXEC Privilege mode:

Command Syntax	Command Mode	Purpose
<code>clear counters [interface]</code>	EXEC Privilege	Clear the counters used in the show interface commands for all VLANs, and physical interfaces or selected ones.

Without an interface specified, the command clears all interface counters.

- (OPTIONAL) Enter the following interface keywords and slot/port or number information:
- For a Port Channel interface, enter the keyword `port-channel` followed by a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port numbers.
- For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.

When you enter this command, you must confirm that you want Dell Networking OS to clear the interface counters for the interface (refer to the below clearing interface example).

Clearing an Interface:

```
Dell#clear counters tengig 0/1
Clear counters on TenGigabitEthernet 0/1 [confirm]
Dell#
```

Fibre Channel Interface

The FN 2210S functions as a converged enhanced Ethernet (CEE) switch that supports both LAN and storage area network (SAN) traffic using the Fibre Channel protocol.

To access a SAN fabric, use a Fibre Channel (FC) module installed in the aggregator.

Configuring Fibre Channel Interfaces

To configure FC interfaces on the FN 2210S aggregator, follow these steps.

1. Enable Fibre Channel capability on the switch.
2. Configure Fibre Channel interfaces.
3. Verify the global Fibre Channel configuration on the switch and on individual interfaces.

Enabling Fibre Channel Capability

To enable Fibre Channel capability, follow this step.

Enable Fibre Channel capability.

CONFIGURATION mode

```
feature fc
```

Configuring Fibre Channel Interfaces

To configure a Fibre Channel interface, follow these steps.

Convert the interfaces 9 and 10 from FC to Ethernet mode.

CONFIGURATION mode

```
stack-unit unit number port-group 0 portmode ethernet
```

iSCSI Optimization

An Aggregator enables internet small computer system interface (iSCSI) optimization with default iSCSI parameter settings(Default iSCSI Optimization Values) and is auto-provisioned to support:

[iSCSI Optimization: Operation](#)

To display information on iSCSI configuration and sessions, use `show` commands.

iSCSI optimization enables quality-of-service (QoS) treatment for iSCSI traffic.

Supported Modes

Stacking, VLT

iSCSI Optimization Overview

iSCSI is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices and initiators in a storage area network (SAN).

iSCSI optimization enables the network switch to auto-detect Dell's iSCSI storage arrays and triggers self-configuration of several key network configurations that enables optimization of the network for better storage traffic throughput.

iSCSI optimization provides a means of monitoring iSCSI sessions and applying QoS policies on iSCSI traffic. When enabled, iSCSI optimization allows a switch to monitor (snoop) the establishment and termination of iSCSI connections. The switch uses the snooped information to detect iSCSI sessions and connections established through the switch.

iSCSI optimization allows you to reduce deployment time and management complexity in data centers. In a data center network, Dell EqualLogic and Compellent iSCSI storage arrays are connected to a converged Ethernet network using the data center bridging exchange protocol (DCBx) through Ethernet switches.

iSCSI session monitoring over virtual link trunking (VLT) synchronizes the iSCSI session information between the VLT peers, allowing session information to be available in both VLT peers.

iSCSI optimization functions as follows:

- Manual configuration to detect Compellent storage arrays where auto-detection is not supported.
- Jumbo frames — Ports are set to a maximum transmission unit (MTU) of 12,000 bytes.
- If you configured flow-control, iSCSI uses the current configuration. If you did not configure flow-control, iSCSI auto-configures flow control.
- iSCSI monitoring sessions — the switch monitors and tracks active iSCSI sessions in connections on the switch, including port information and iSCSI session information.

- iSCSI QoS — A user-configured iSCSI class of service (CoS) profile is applied to all iSCSI traffic. Classifier rules are used to direct the iSCSI data traffic to queues that can be given preferential QoS treatment over other data passing through the switch. Preferential treatment helps to avoid session interruptions during times of congestion that would otherwise cause dropped iSCSI packets.
- iSCSI DCBx TLVs are supported.

The following figure shows iSCSI optimization between servers in a server enclosure and a storage array in which an Aggregator connects installed servers (iSCSI initiators) to a storage array (iSCSI targets) in a SAN network. iSCSI optimization running on the Aggregator is configured to use dot1p priority-queue assignments to ensure that iSCSI traffic in these sessions receives priority treatment when forwarded on Aggregator hardware.




Figure 16. iSCSI Optimization Example

Monitoring iSCSI Traffic Flows

The switch snoops iSCSI session-establishment and termination packets by installing classifier rules that trap iSCSI protocol packets to the CPU for examination.

Devices that initiate iSCSI sessions usually use well-known TCP ports 3260 or 860 to contact targets. When you enable iSCSI optimization, by default the switch identifies IP packets to or from these ports as iSCSI traffic.

You can configure the switch to monitor traffic for additional port numbers or a combination of port number and target IP address, and you can remove the well-known port numbers from monitoring.

Information Monitored in iSCSI Traffic Flows

iSCSI optimization examines the following data in packets and uses the data to track the session and create the classifier entries that enable QoS treatment:

- Initiator's IP Address
- Target's IP Address
- ISID (Initiator defined session identifier)
- Initiator's IQN (iSCSI qualified name)
- Target's IQN
- Initiator's TCP Port
- Target's TCP Port

If no iSCSI traffic is detected for a session during a user-configurable aging period, the session data clears.

Synchronizing iSCSI Sessions Learned on VLT-Lags with VLT-Peer

The following behavior occurs during synchronization of iSCSI sessions.

- If the iSCSI login request packet is received on a port belonging to a VLT lag, the information is synced to the VLT peer and the connection is associated with this interface.
- Additional updates to connections (including aging updates) that are learnt on VLT lag members are synced to the peer.
- When receiving an iSCSI login request on a non-VLT interface followed by a response from a VLT interface, the session is not synced since it is initially learnt on a non-VLT interface through the request packet.
- The peer generates a new connection log that sees the login response packet. If the login response packet uses the ICL path, it is seen by both the peers, which in turn generate logs for this connection.

iSCSI Optimization: Operation

When the Aggregator auto-configures with iSCSI enabled, the following occurs:

- Link-level flow control is enabled on PFC disabled interfaces.
- iSCSI session snooping is enabled.

- iSCSI LLDP monitoring starts to automatically detect EqualLogic arrays.

iSCSI optimization requires LLDP to be enabled. LLDP is enabled by default when an Aggregator auto-configures.

The following message displays when you enable iSCSI on a switch and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_ENABLE: iSCSI has been enabled causing flow control to be enabled on all interfaces. EQL detection and enabling iscsi profile-compellent on an interface may cause some automatic configurations to occur like jumbo frames on all ports and no storm control and spanning tree port-fast on the port of detection.
```

Configuring iSCSI Optimization

To configure iSCSI optimization, use the following commands.

1. **For a non-DCB environment:** Enable iSCSI.

```
CONFIGURATION mode
```

```
iscsi enable
```

2. (Optional) Configure the iSCSI target ports and optionally the IP addresses on which iSCSI communication is monitored.

```
CONFIGURATION mode
```

```
[no] iscsi target port tcp-port-1 [tcp-port-2...tcp-port-16] [ip-address address]
```

- *tcp-port-n* is the TCP port number or a list of TCP port numbers on which the iSCSI target listens to requests. You can configure up to 16 target TCP ports on the switch in one command or multiple commands. The default is **860, 3260**.

Separate port numbers with a comma. If multiple IP addresses are mapped to a single TCP port, use the `no iscsi target port tcp-port-n` command to remove all IP addresses assigned to the TCP number.

To delete a specific IP address from the TCP port, use the `no iscsi target port tcp-port-n ip-address address` command to specify the address to be deleted.

- *ip-address* specifies the IP address of the iSCSI target. When you enter the `no` form of the command, and the TCP port you want to delete is one bound to a specific IP address, include the IP address value in the command.

If multiple IP addresses are mapped to a single TCP port, use the `no iscsi target port` command to remove all IP addresses assigned to the TCP port number.

To remove a single IP address from the TCP port, use the `no iscsi target port ip-address` command.

3. (Optional) Set the QoS policy that is applied to the iSCSI flows.

```
CONFIGURATION mode
```

```
[no] iscsi cos {enable | disable | dot1p vlan-priority-value [remark] | dscp dscp-value [remark]}
```


- `enable`: enables the application of preferential QoS treatment to iSCSI traffic so that iSCSI packets are scheduled in the switch with a dot1p priority 4 regardless of the VLAN priority tag in the packet. The default is: iSCSI packets are handled with dotp1 priority 4 without `remark`.
 - `disable`: disables the application of preferential QoS treatment to iSCSI frames.
 - `dot1p vlan-priority-value`: specifies the virtual local area network (VLAN) priority tag assigned to incoming packets in an iSCSI session. The range is from 0 to 7. The default is: the dot1p value in ingress iSCSI frames is not changed and the same priority is used in iSCSI TLV advertisements if you do not enter the `iscsi priority-bits` command (Step 10).
 - `dscp dscp-value`: specifies the DSCP value assigned to incoming packets in an iSCSI session. The range is from 0 to 63. The default is: the DSCP value in ingress packets is not changed.
 - `remark`: marks incoming iSCSI packets with the configured dot1p or DSCP value when they egress the switch. The default is: the dot1 and DSCP values in egress packets are not changed.
4. (Optional) Set the aging time for iSCSI session monitoring.

CONFIGURATION mode

```
[no] iscsi aging time time.
```

The range is from 5 to 43,200 minutes.

The default is **10 minutes**.

5. (Optional) Configures DCBX to send iSCSI TLV advertisements.
LLDP CONFIGURATION mode or INTERFACE LLDP CONFIGURATION mode

```
[no] advertise dcbx-app-tlv iscsi.
```

You can send iSCSI TLVs either globally or on a specified interface. The interface configuration takes priority over global configuration.

The default is **Enabled**.

6. (Optional) Configures the advertised priority bitmap in iSCSI application TLVs.
LLDP CONFIGURATION mode

```
[no] iscsi priority-bits.
```

The default is **4** (0x10 in the bitmap).

7. (Optional) Enter interface configuration mode to configure the auto-detection of Dell Compellent disk arrays.

CONFIGURATION mode

```
interface port-type slot/port
```

8. (Optional) Configures the auto-detection of Compellent arrays on a port.
INTERFACE mode

```
[no] iscsi profile-compellent.
```

The default is: Compellent disk arrays are not detected.



NOTE: All these configurations are available only in PMUX mode.

Displaying iSCSI Optimization Information

To display information on iSCSI optimization, use the show commands detailed in the below table:

Table 7. Displaying iSCSI Optimization Information

Command	Output
show iscsi	Displays the currently configured iSCSI settings.
show iscsi sessions	Displays information on active iSCSI sessions on the switch that have been established since the last reload.
show iscsi sessions detailed [session <i>isid</i>]	Displays detailed information on active iSCSI sessions on the switch. To display detailed information on specified iSCSI session, enter the session's iSCSI ID.
show run iscsi	Displays all globally-configured non-default iSCSI settings in the current Dell Networking OS session.

show iscsi Command Example

```
Dell# show iscsi
iSCSI is enabled
iSCSI session monitoring is enabled
iSCSI COS : dot1p is 4 no-remark
Session aging time: 10
Maximum number of connections is 256
-----
iSCSI Targets and TCP Ports:
-----
TCP Port Target IP Address
3260
860
```

show iscsi sessions Command Example

```
Dell# show iscsi sessions
Session 0:
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0e70c2002-10a0018426a48c94-iom010
Initiator: iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000
Session 1:
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0f60c2002-0360018428d48c94-iom011
Initiator: iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000.
```

show iscsi sessions detailed Command Example

```
Dell# show iscsi sessions detailed
Session 0 :
-----
Target:iqn.2010-11.com.ixia:ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-2c
```

```

Up Time:00:00:01:28(DD:HH:MM:SS)
Time for aging out:00:00:09:34(DD:HH:MM:SS)
ISID:806978696102
Initiator Initiator Target Target Connection
IP Address TCP Port IP Address TCPPort ID
10.10.0.44 33345 10.10.0.101 3260 0
Session 1 :

```

```

-----
Target:iqn.2010-11.com.ixia.ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-35
Up Time:00:00:01:22(DD:HH:MM:SS)
Time for aging out:00:00:09:31(DD:HH:MM:SS)
ISID:806978696102
Initiator      Initiator      Target      Target      Connection ID
IP Address     TCP Port      IP Address  TCPPort
10.10.0.53    33432         10.10.0.101 3260        0

```

Isolated Networks for Aggregators

An Isolated Network is an environment in which servers can only communicate with the uplink interfaces and not with each other even though they are part of same VLAN. If the servers in the same chassis need to communicate with each other, it requires a non-isolated network connectivity between them or it needs to be routed in the TOR.

Isolated Networks can be enabled on per VLAN basis. If a VLAN is set to be isolated, all the packets of originating from the server ports for that VLAN (Isolated Network) will be redirected to uplink LAG, including the packets destined for the server ports on the same blade.

ToR applies required ACLs and other necessary actions before sending the packet to destination. If the packet is destined to server on the same IOA blade, it is routed back on the uplink lag where it was received. Traffic that hits at the uplink ports are regularly switched based on the L2 MAC lookup. Unknown Unicast and Multicast packets from Uplink Port towards server port on an isolated network enabled VLAN, is dropped.

The isolated network feature is supported only in the standalone mode.

Isolated network is currently not supported in the following modes:

- VLT mode
- Stacking mode
- PMUX mode

 **NOTE:** Isolated Networks is not enabled for FCOE VLANs and on default VLAN. It can be managed via CLI or AFM. For more information, refer to AFM user manual.

Configuring and Verifying Isolated Network Settings

Enable the isolated-network functionality for a particular VLAN or a set of VLANs using below command:

```
Dell(conf)#io-aggregator isolated-network vlan <vlan-range>
```

To disable the isolated-network functionality, use the no form of command.

```
Dell(conf)#no io-aggregator isolated-network vlan <vlan-range>
```

To display the VLANs that are configured to be part of an isolated network on the Aggregator, use the below command.

```
Dell#show io-aggregator isolated-networks
Isolated Network Enabled VLANs : 5-10
```

Link Aggregation

Unlike IOA Automated modes (Standalone and VLT modes), the IOA Programmable MUX (PMUX) can support multiple uplink LAGs. You can provision multiple uplink LAGs. The I/O Aggregator auto-configures with link aggregation groups (LAGs) as follows:

- All uplink ports are automatically configured in a single port channel (LAG 128).
- Server-facing LAGs are automatically configured if you configure server for link aggregation control protocol (LACP)-based NIC teaming (Network Interface Controller (NIC) Teaming).

No manual configuration is required to configure Aggregator ports in the uplink or a server-facing LAG.

 **NOTE:** Static LAGs are not supported on the SMUX Aggregator.


 **NOTE:** In order to avoid loops, only disjoint VLANs are allowed between the uplink ports/uplink LAGs and uplink-to-uplink switching is disabled.

Supported Modes

Standalone, VLT, PMUX, Stacking

How the LACP is Implemented on an Aggregator


The LACP provides a means for two systems (also called partner systems) to exchange information through dynamic negotiations to aggregate two or more ports with common physical characteristics to form a link aggregation group.

 **NOTE:** A link aggregation group is referred to as a *port channel* by the Dell Networking OS.

A LAG provides both load-sharing and port redundancy across stack units. An Aggregator supports LACP for auto-configuring dynamic LAGs. Use CLI commands to display LACP information, clear port-channel counters, and debug LACP operation for auto-configured LAG on an Aggregator.

The Dell Networking OS implementation of LACP is based on the standards specified in the IEEE 802.3: “Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.”

LACP functions by constantly exchanging custom MAC protocol data units (PDUs) across local area network (LAN) Ethernet links. The protocol packets are only exchanged between ports that you configure as LACP-capable.

 **NOTE:** In Standalone, VLT, and Stacking modes, you can configure a maximum of 16 members in port-channel 128. In PMUX mode, you can have multiple port-channels with up to 16 members per channel.

Uplink LAG

When the Aggregator power is on, all uplink ports are configured in a single LAG (LAG 128).

Server-Facing LAGs

Server-facing ports are configured as individual ports by default. If you configure a server NIC in standalone, stacking, or VLT mode for LACP-based NIC teaming, server-facing ports are automatically configured as part of dynamic LAGs. The LAG range 1 to 127 is reserved for server-facing LAGs.

After the Aggregator receives LACPDU from server-facing ports, the information embedded in the LACPDU (remote-system ID and port key) is used to form a server-facing LAG. The LAG/port-channel number is assigned based on the first available number in the range from 1 to 127. For each unique remote system-id and port-key combination, a new LAG is formed and the port automatically becomes a member of the LAG.

All ports with the same combination of system ID and port key automatically become members of the same LAG. Ports are automatically removed from the LAG if the NIC teaming configuration on a server-facing port changes or if the port goes operationally down. Also, a server-facing LAG is removed when the last port member is removed from the LAG.

The benefit of supporting a dynamic LAG is that the Aggregator's server-facing ports can toggle between participating in the LAG or acting as individual ports based on the dynamic information exchanged with a server NIC. LACP supports the exchange of messages on a link to allow their LACP instances to:

- Reach agreement on the identity of the LAG to which the link belongs.
- Attach the link to that LAG.
- Enable the transmission and reception functions in an orderly manner.
- Detach the link from the LAG if one of the partner stops responding.

LACP Modes

The Aggregator supports only LACP active mode as the default mode of operation. In active mode, a port interface is considered to be not part of a LAG but rather in an active negotiating state.

A port in active mode automatically initiates negotiations with other ports by sending LACP packets. If you configure server-facing ports for LACP-based NIC teaming, LACP negotiations take place to aggregate the port in a dynamic LAG. If you do not configure server-facing ports for LACP-based NIC teaming, a port is treated as an individual port in active negotiating state.


Auto-Configured LACP Timeout

LACP PDUs are exchanged between port channel (LAG) interfaces to maintain LACP sessions. LACP PDUs are transmitted at a slow or fast transmission rate, depending on the LACP timeout value configured on the partner system.

The timeout value is the amount of time that a LAG interface waits for a PDU from the partner system before bringing the LACP session down. The default timeout is long-timeout (30 seconds) and is not user-configurable on the Aggregator.

Link Aggregation Control Protocol (LACP)

The commands for Dell Networks's implementation of the link aggregation control protocol (LACP) for creating dynamic link aggregation groups (LAGs) — known as port-channels in the Dell Networking OS — are provided in the following sections.

 **NOTE:** For static LAG commands, refer to the [Interfaces](#) chapter, based on the standards specified in the IEEE 802.3 Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.

Configuration Tasks for Port Channel Interfaces

To configure a port channel (LAG), use the commands similar to those found in physical interfaces. By default, no port channels are configured in the startup configuration. In VLT mode, port channel configurations are allowed in the startup configuration.

These are the mandatory and optional configuration tasks:

- [Creating a Port Channel](#) (mandatory)
- [Adding a Physical Interface to a Port Channel](#) (mandatory)
- [Reassigning an Interface to a New Port Channel](#) (optional)
- [Configuring the Minimum Oper Up Links in a Port Channel](#) (optional)
- [Configuring VLAN Tags for Member Interfaces](#) (optional)
- [Deleting or Disabling a Port Channel](#) (optional)

Creating a Port Channel

You can create up to 128 port channels with four port members per group on the Aggregator. To configure a port channel, use the following commands.

1. Create a port channel.
CONFIGURATION mode

```
interface port-channel id-number
```

2. Ensure that the port channel is active.
INTERFACE PORT-CHANNEL mode


```
no shutdown
```

After you enable the port channel, you can place it in Layer 3 mode. To configure an IP address to place the port channel in Layer 3 mode, use the `switchport` command.

You can configure a port channel as you would a physical interface by enabling or configuring protocols or assigning access control lists.

Adding a Physical Interface to a Port Channel

The physical interfaces in a port channel can be on any line card in the chassis, but must be the same physical type.

 **NOTE:** Port channels can contain a mix of Gigabit Ethernet and 10/100/1000 Ethernet interfaces, but Dell Networking OS disables the interfaces that are not the same speed of the first channel member in the port channel.

You can add any physical interface to a port channel if the interface configuration is minimal. You can configure only the following commands on an interface if it is a member of a port channel:

- `description`
- `shutdown/no shutdown`
- `mtu`
- `ip mtu` (if the interface is on a Jumbo-enabled by default)

 **NOTE:**

A logical port channel interface cannot have flow control. Flow control can only be present on the physical interfaces if they are part of a port channel.

To view the interface's configuration, enter INTERFACE mode for that interface and use the `show config` command or from EXEC Privilege mode, use the `show running-config interface interface` command.

When an interface is added to a port channel, Dell Networking OS recalculates the hash algorithm.

To add a physical interface to a port, use the following commands.

1. Add the interface to a port channel.

```
INTERFACE PORT-CHANNEL mode
```

```
channel-member interface
```

This command is applicable only in PMUX mode.

The `interface` variable is the physical interface type and slot/port information.

2. Double check that the interface was added to the port channel.

```
INTERFACE PORT-CHANNEL mode
```

```
show config
```

To view the port channel's status and channel members in a tabular format, use the `show interfaces port-channel brief` command in EXEC Privilege mode, as shown in the following example.

Example of the `show interfaces port-channel brief` Command

```
Dell#sh int port-channel brief
Codes: L - LACP Port-channel
       O - OpenFlow Controller Port-channel

   LAG  Mode  Status      Uptime      Ports
L   1    L2    up          00:00:19    Te 0/7      (Up)
                      Te 0/8      (Up)
L  128  L2    up          00:00:36    Te 0/9      (Up)
                      Te 0/10     (Up)
                      Te 0/11     (Up)

Dell#
```

The following example shows the port channel's mode (L2 for Layer 2 and L3 for Layer 3 and L2L3 for a Layer 2-port channel assigned to a routed VLAN), the status, and the number of interfaces belonging to the port channel.

Example of the show interface port-channel Command

```
Dell#show interface port-channel 1
Port-channel 1 is down, line protocol is down
Hardware address is 00:1e:c9:de:04:9c, Current address is 00:1e:c9:de:04:9c
Interface index is 1107492865
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :001ec9de049c
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Members in this channel:
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 02:57:05
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,        0 packets/sec, 0.00% of line-rate
Time since last interface status change: 02:57:05

Dell#
```

When more than one interface is added to a Layer 2-port channel, Dell Networking OS selects one of the active interfaces in the port channel to be the primary port. The primary port replies to flooding and sends protocol data units (PDUs). An asterisk in the show interfaces port-channel brief command indicates the primary port.

As soon as a physical interface is added to a port channel, the properties of the port channel determine the properties of the physical interface. The configuration and status of the port channel are also applied to the physical interfaces within the port channel. For example, if the port channel is in Layer 2 mode, you cannot add an IP address or a static MAC address to an interface that is part of that port channel.

Example of Error Due to an Attempt to Configure an Interface that is Part of a Port Channel

```
Dell(conf)#int port-channel 128
Dell(conf-if-po-128)#show config
!
interface Port-channel 128
mtu 12000
portmode hybrid
switchport
fip-snooping port-mode fcf
no shutdown
link-bundle-monitor enable
Dell(conf-if-po-128)#
```

Reassigning an Interface to a New Port Channel

An interface can be a member of only one port channel. If the interface is a member of a port channel, remove it from the first port channel and then add it to the second port channel. Each time you add or remove a channel member from a port channel, Dell Networking OS recalculates the hash algorithm for the port channel.

To reassign an interface to a new port channel, use the following commands.

1. Remove the interface from the first port channel.
INTERFACE PORT-CHANNEL mode

`no channel-member interface`
2. Change to the second port channel INTERFACE mode.
INTERFACE PORT-CHANNEL mode

`interface port-channel id number`

This command is applicable only in PMUX mode.
3. Add the interface to the second port channel.
INTERFACE PORT-CHANNEL mode

`channel-member interface`

Example of Moving an Interface to a New Port Channel

The following example shows moving the TenGigabitEthernet 0/8 interface from port channel 4 to port channel 3.

```
Dell(conf-if-po-4)#show config
!
interface Port-channel 4
channel-member TenGigabitEthernet 0/8
no shutdown
Dell(conf-if-po-4)#no chann tengi 0/8
Dell(conf-if-po-4)#int port 3
Dell(conf-if-po-3)#channel tengi 0/8
Dell(conf-if-po-3)#sho conf
!
interface Port-channel 3
channel-member TenGigabitEthernet 0/8
shutdown
Dell(conf-if-po-3)#
```

Configuring the Minimum Oper Up Links in a Port Channel

You can configure the minimum links in a port channel (LAG) that must be in “oper up” status to consider the port channel to be in “oper up” status.

To set the “oper up” status of your links, use the following command.

- Enter the number of links in a LAG that must be in “oper up” status.
INTERFACE mode

`minimum-links number`

The default is **1**.

Example of Configuring the Minimum Oper Up Links in a Port Channel

```
Dell#config t
Dell(conf)#int po 128
Dell(conf-if-po-128)#minimum-links 5
Dell(conf-if-po-128)#
```

Configuring VLAN Tags for Member Interfaces

To configure and verify VLAN tags for individual members of a port channel, perform the following:

1. Configure VLAN membership on individual ports

INTERFACE mode

```
Dell(conf-if-te-0/2)#vlan tagged 2,3-4
```

2. Use the `switchport` command in INTERFACE mode to enable Layer 2 data transmissions through an individual interface

INTERFACE mode

```
Dell(conf-if-te-0/2)#switchport
```

This `switchport` configuration is allowed only in PMUX mode. In all other modes, it is automatically configured.

3. Verify the manually configured VLAN membership (**show interfaces switchport** *interface* command).

EXEC mode

```
Dell(conf)# interface tengigabitethernet 0/1
Dell(conf-if-te-0/1)#switchport
Dell(conf-if-te-0/1)# vlan tagged 2-5,100,4010
Dell#show interfaces switchport te 0/1
```

```
Codes:  U - Untagged, T - Tagged
         x - Dot1x untagged, X - Dot1x tagged
         G - GVRP tagged, M - Trunk, H - VSN tagged
         i - Internal untagged, I - Internal tagged, v - VLT untagged, V -
VLT tagged
```

```
Name: TenGigabitEthernet 0/1
802.1QTagged: True
Vlan membership:
Q      Vlans
T      2-5,100,4010
```

```
Dell#
```

Deleting or Disabling a Port Channel

To delete or disable a port channel, use the following commands.

- Delete a port channel.

CONFIGURATION mode

```
no interface portchannel channel-number
```

- Disable a port channel.

shutdown

When you disable a port channel, all interfaces within the port channel are operationally down also.

Configuring Auto LAG

You can enable or disable auto LAG on the server-facing interfaces. By default, auto LAG is enabled. This functionality is supported on the Aggregator in Standalone, Stacking, and VLT modes.

To configure auto LAG, use the following commands:

1. Enable the auto LAG on all the server ports.

CONFIGURATION mode

```
io-aggregator auto-lag enable
```

```
Dell(config)# io-aggregator auto-lag enable
```

To disable the auto LAG on all the server ports, use the `no io-aggregator auto-lag enable` command. When disabled, all the server ports associated in a LAG are removed and the LAG itself gets removed. Any LACPDUs received on the server ports are discarded.

In VLT mode, the global auto LAG is automatically synced to the peer VLT through ICL message.

2. Enable the auto LAG on a specific server port.

Interface Configuration mode

```
auto-lag enable
```

```
Dell(conf-if-te-0/1)# auto-lag enable
```

To disable the auto LAG, use the `no auto-lag enable` command. When disabled, the server port is removed from the LAG and if the server port is the last member of the LAG, the LAG itself gets removed. Any LACPDUs received on the server port are discarded.

In VLT mode, the interface level auto LAG configuration is not synced to the peer. Only the global auto LAG is synced to the peer.

3. View the auto LAG configurations.

EXEC mode

```
show io-aggregator auto-lag status
```

```
Dell# show io-aggregator auto-lag status
```

```
Auto LAG creation on server port(s) is enabled
```

For the interface level auto LAG configurations, use the `show interface` command.

```
Dell(conf-if-te-0/1)#no auto-lag enable
Dell(conf-if-te-0/1)#
Dell(conf-if-te-0/1)#do show interface tengigabitethernet 0/1
TenGigabitEthernet 0/1 is up, line protocol is down(error-disabled[UFD])
Hardware is DellEth, address is f8:b1:56:07:1d:8e
    Current address is f8:b1:56:07:1d:8e
```

```

Server Port AdminState is Up
Pluggable media not present
Interface index is 15274753
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :f8b156071d8e
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed auto
Auto-lag is disabled
Flowcontrol rx on tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:12:53
Queueing strategy: fifo
Input Statistics:
    0 packets, 0 bytes
    0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    0 Multicasts, 0 Broadcasts
    0 runts, 0 giants, 0 throttles
    0 CRC, 0 overrun, 0 discarded
Output Statistics:
    0 packets, 0 bytes, 0 underruns
    0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    0 Multicasts, 0 Broadcasts, 0 Unicasts
    0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
    Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
    Output 00.00 Mbits/sec,         0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:11:36

Dell(conf-if-te-0/1)#

```

Sample Configuration

```

Dell# config terminal
Dell(conf)# no io-aggregator auto-lag enable
Dell(conf)# end
Dell# show io-aggregator auto-lag status

Auto LAG creation on server port(s) is disabled

Dell#

Dell# config terminal
Dell(config)# interface tengigabitethernet 0/1
Dell(config-if-te-0/1)# no auto-lag enable
Dell(config-if-te-0/1)# show config
!
interface TenGigabitEthernet 0/1
mtu 12000
portmode hybrid
switchport
no auto-lag enable
!
protocol lldp
  advertise management-tlv management-address system-name
  dcbx port-role auto-downstream
no shutdown
Dell#

```

Configuring the Minimum Number of Links to be Up for Uplink LAGs to be Active

You can activate the LAG bundle for uplink interfaces or ports (the uplink port-channel is LAG 128) on the I/O Aggregator only when a minimum number of member interfaces of the LAG bundle is up. For example, based on your network deployment, you may want the uplink LAG bundle to be activated only if a certain number of member interface links is also in the up state. If you enable this setting, the uplink LAG bundle is brought up only when the specified minimum number of links are up and the LAG bundle is moved to the down state when the number of active links in the LAG becomes less than the specified number of interfaces. By default, the uplink LAG 128 interface is activated when at least one member interface is up.

To configure the minimum number of member links that must be up for a LAG bundle to be fully up, perform the following steps:

Specify the minimum number of member interfaces of the uplink LAG 128 bundle that must be up for the LAG bundle to be brought up. The default minimum number of member links that must be active for the uplink LAG to be active is 1. Enter the `minimum-links` number command in the Port Channel Interface 128 Configuration mode to specify this value.

```
Dell(conf)#interface port-channel 128
```

```
Dell(conf-if-po-128)#minimum-links 4
```

Use the `show interfaces port-channel` command to view information regarding the configured LAG or port channel settings. The **Minimum number of links to bring Port-channel up is** field in the output of this command displays the configured number of active links for the LAG to be enabled.

```
Dell#show interfaces port-channel 128
Port-channel 128 is up, line protocol is down(minimum links not up)
Created by LACP protocol
Hardware address is 00:01:02:03:04:05, Current address is 00:01:02:03:04:05
Interface index is 1107492992
Minimum number of links to bring Port-channel up is 4
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :000102030405
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed auto
Members in this channel:
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 05:22:24
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
```

```
Output 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
Time since last interface status change: 05:22:28
```

Optimizing Traffic Disruption Over LAG Interfaces On IOA Switches in VLT Mode

When you use the `write memory` command while an Aggregator operates in VLT mode, the VLT LAG configurations are saved in nonvolatile storage (NVS).

By restoring the settings saved in NVS, the VLT ports come up quicker on the primary VLT node and traffic disruption is reduced. The delay in restoring the VLT LAG parameters is reduced (90 seconds by default) on the secondary VLT peer node before it becomes operational. This makes sure that the configuration settings of the primary VLT node are synchronized with the secondary VLT peer node before the secondary VLT mode is up. The traffic outage is less than 200 milliseconds during the restart or switchover of the VLT peer nodes from primary to secondary.

Preserving LAG and Port Channel Settings in Nonvolatile Storage

Use the `write memory` command on an I/O Aggregator, which operates in either standalone or PMUX modes, to save the LAG port channel configuration parameters. This behavior enables the port channels to be brought up because the configured interface attributes are available in the system database during the booting of the device. With the reduction in time for the port channels to become active after the switch is booted, the loss in the number of packets that are serviced by these interfaces is minimized.

Enabling the Verification of Member Links Utilization in a LAG Bundle

To examine the working efficiency of the LAG bundle interfaces, perform the following steps:

1. The functionality to detect the working efficiency of the LAG bundle interfaces is automatically activated on all the port channels, except the port channel that is configured as a VLT interconnect link, during the booting of the switch.
2. Use the `show link-bundle-distribution port-channel interface-number` command to display the traffic handling and utilization of the member interfaces of the port channel. The following sample output is displayed when you enter this `show` command.

EXEC

```
Dell#show link-bundle-distribution port-channel
```

```
Dell#show link-bundle-distribution port-channel 1
```

```
Link-bundle trigger threshold - 60
```

```
LAG bundle - 1          Utilization[In Percent] - 0          Alarm State - Inactive
```

Interface	Line Protocol	Utilization[In Percent]
Te 0/5	Up	0
Te 0/7	Up	0

Monitoring the Member Links of a LAG Bundle

You can examine and view the operating efficiency and the traffic-handling capacity of member interfaces of a LAG or port channel bundle. This method of analyzing and tracking the number of packets processed by the member interfaces helps you manage and distribute the packets that are handled by the LAG bundle. The functionality to detect the working efficiency of the LAG bundle interfaces is automatically activated on all the port channels, except the port channel that is configured as a VLT interconnect link, during the booting of the switch. This functionality is supported on I/O Aggregators in PMUX, Standalone, and VLT modes. By default, this capability is enabled on all of the port channels set up on the switch.

You can use the `show link-bundle-distribution port-channel interface-number` command to display the traffic-handling and utilization of the member interfaces of the port channel. The following table describes the output fields of this show command.

Table 8. Output Field Descriptions for show link-bundle-distribution port-channel Command

Field	Description
Link-bundle trigger threshold	Threshold value that is the checkpoint, exceeding which the link bundle is marked as being overutilized and alarm is generated
LAG bundle number	Number of the LAG bundle
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the port channel
Alarm State	Indicates whether an alarm is generated if over-utilization of the port channel occurred. The value, Active, is displayed for this field.
Interface	Slot and port number, and the type of the member interface of the port channel
Line Protocol	Indicates whether the interface is administratively up or down
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the particular member interface

You can also use the `show running-configuration interface port-channel` command in EXEC Privilege mode to view whether the mechanism to evaluate the utilization of the member interfaces of the LAG bundle is enabled. The following sample output illustrates the portion of this show command:

```
Dell#show running-config int port-channel 1
!
interface Port-channel 1
  mtu 12000
  portmode hybrid
  switchport
  vlt-peer-lag port-channel 1
  no shutdown
  link-bundle-monitor enable
```


Verifying LACP Operation and LAG Configuration

To verify the operational status and configuration of a dynamically created LAG, and LACP operation on a LAG on an Aggregator, enter the `show interfaces port-channel port-channel-number` and `show lacp port-channel-number` commands.

The `show` outputs in this section for uplink LAG 128 and server-facing LAG 1 refer to the *LACP Operation on an Aggregator* figure

```
show interfaces port-channel 128
```

Command Example

```
Dell# show interfaces port-channel 128
```

```
Port-channel 128 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:01:e8:e1:e1:c1, Current address is 00:01:e8:e1:e1:c1
Interface index is 1107755136
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag1280001e8e1e1c1
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed 40000 Mbit
Members in this channel: Te0/9 Te0/10 Te 0/11 Te0/12
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:11:50
Queueing strategy: fifo
Input Statistics:
    182 packets, 17408 bytes
    92 64-byte pkts, 0 over 64-byte pkts, 90 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    182 Multicasts, 0 Broadcasts
    0 runs, 0 giants, 0 throttles
    0 CRC, 0 overrun, 0 discarded
Output Statistics:
    2999 packets, 383916 bytes, 0 underruns
    5 64-byte pkts, 214 over 64-byte pkts, 2727 over 127-byte pkts
    53 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    2904 Multicasts, 95 Broadcasts, 0 Unicasts
    0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
    Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
    Output 00.00 Mbits/sec,        4 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:11:42
```

```
show lacp 128 Command Example
```

```
Dell# show lacp 128
```

```
Port-channel 128 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e8e1.e1c3
Partner System ID: Priority 32768, Address 0001.e88b.253d
Actor Admin Key 128, Oper Key 128, Partner Oper Key 128, VLT Peer Oper Key 128
LACP LAG 128 is an aggregatable link
LACP LAG 128 is a normal LAG
```

A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC

I - Collection enabled, J - Collection disabled, K - Distribution enabled
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state

Port Te 0/9 is enabled, LACP is enabled and mode is lacp
Port State: Bundle

Actor Admin: State ADEHJLMP Key 128 Priority 32768
Oper: State ADEGIKNP Key 128 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 128 Priority 32768

Port Te 0/10 is enabled, LACP is enabled and mode is lacp
Port State: Bundle

Actor Admin: State ADEHJLMP Key 128 Priority 32768
Oper: State ADEGIKNP Key 128 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 128 Priority 32768

Port Te 0/11 is enabled, LACP is enabled and mode is lacp
Port State: Bundle

Actor Admin: State ADEHJLMP Key 128 Priority 32768
Oper: State ADEGIKNP Key 128 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 128 Priority 32768

show interfaces port-channel 1 **Command Example**

Dell# show interfaces port-channel 1

Port-channel 1 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:01:e8:e1:e1:c1, Current address is 00:01:e8:e1:e1:c1
Interface index is 1107755009
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag10001e8e1e1c1
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed 10000 Mbit
Members in this channel: Te 0/12(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:12:41
Queueing strategy: fifo
Input Statistics:
112 packets, 18161 bytes
0 64-byte pkts, 46 over 64-byte pkts, 37 over 127-byte pkts
29 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
59 Multicasts, 53 Broadcasts
0 runts, 0 giants, 0 throttles
0 CRC, 0 overrun, 0 discarded
Output Statistics:
135 packets, 19315 bytes, 0 underruns
0 64-byte pkts, 79 over 64-byte pkts, 32 over 127-byte pkts
24 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
93 Multicasts, 42 Broadcasts, 0 Unicasts
0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:12:38

show lacp 1 **Command Example**

```
Dell# show lacp 1
```

```
Port-channel 1 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e8e1.e1c3
Partner System ID: Priority 65535, Address 24b6.fd87.d8ac
Actor Admin Key 1, Oper Key 1, Partner Oper Key 33, VLT Peer Oper Key 1
LACP LAG 1 is an aggregatable link
LACP LAG 1 is a normal LAG

A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC
I - Collection enabled, J - Collection disabled, K - Distribution enabled
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state
```

```
Port Te 0/12 is enabled, LACP is enabled and mode is lacp
```

```
Port State: Bundle
```

```
Actor Admin: State ADEHJLMP Key 1 Priority 32768
Oper: State ADEGIKPN Key 1 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ADEGIKPN Key 33 Priority 255
```

Multiple Uplink LAGs with 10G Member Ports

The following sample commands configure multiple dynamic uplink LAGs with 10G member ports based on LACP.

1. Bring up all the ports.

```
Dell#configure
Dell(conf)#int range tengigabitethernet 0/1 - 12
Dell(conf-if-range-te-0/1-12)#no shutdown
```

2. Associate the member ports into LAG-10 and 11.

```
Dell#configure
Dell(conf)#int range tengigabitethernet 0/4 - 5
Dell(conf-if-range-te-0/4-5)#port-channel-protocol lacp
Dell(conf-if-range-te-0/4-5-lacp)#port-channel 10 mode active
Dell(conf-if-range-te-0/4-5-lacp)#end
Dell#
```

```
Dell#configure
Dell(conf)#int tengigabitethernet 0/6
Dell(conf-if-te-0/6)#port-channel-protocol lacp
Dell(conf-if-te-0/6-lacp)#port-channel 11 mode active
Dell(conf-if-te-0/6-lacp)#end
Dell#
```

3. Show the LAG configurations and operational status.

```
Dell#show interface port-channel brief Codes: L - LACP Port-channel
Dell#show interface port-channel brief Codes: L - LACP Port-channel
O - OpenFlow Controller Port-channel LAG Mode Status Uptime Ports
L 10 L3 up 00:01:00 Te 0/4 (Up) Te 0/5 (Up)
L 11 L3 up 00:00:01 Te 0/6 (Up)
Dell#
```

4. Configure the port mode, VLAN, and so forth on the port-channel.

```
Dell#configure
Dell(conf)#int port-channel 10
Dell(conf-if-po-10)#portmode hybrid
Dell(conf-if-po-10)#switchport
Dell(conf-if-po-10)#vlan tagged 1000
Dell(conf-if-po-10)#link-bundle-monitor enable

Dell#configure
Dell(conf)#int port-channel 11
Dell(conf-if-po-11)#portmode hybrid
Dell(conf-if-po-11)#switchport
Dell(conf-if-po-11)#vlan tagged 1000
% Error: Same VLAN cannot be added to more than one uplink port/LAG.
Dell(conf-if-po-11)#vlan tagged 1001
Dell(conf-if-po-11)#link-bundle-monitor enable

Dell#show vlan
Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port
Mirroring VLANs, P - Primary, C - Community, I - Isolated
       O - Openflow
Q: U - Untagged, T - Tagged
    x - Dot1x untagged, X - Dot1x tagged
    o - OpenFlow untagged, O - OpenFlow tagged
    G - GVRP tagged, M - Vlan-stack, H - VSN tagged
    i - Internal untagged, I - Internal tagged, v - VLT
untagged, V - VLT tagged

      NUM      Status      Description      Q      Ports
* 1          Active
                                U Po10(Te 0/4-5)
                                U Po11(Te 0/6)
  1000 Active
                                T Po10(Te 0/4-5)
  1001 Active
                                T Po11(Te 0/6)
Dell#
```

5. Show LAG member ports utilization.


```
Dell#show link-bundle-distribution
Link-bundle trigger threshold - 60
LAG bundle - 10 Utilization[In Percent] - 0 Alarm State - Inactive
Interface      Line Protocol      Utilization[In Percent]
Te 0/4          Up                  0
Te 0/5          Up                  0
LAG bundle - 11 Utilization[In Percent] - 0 Alarm State - Inactive
Interface      Line Protocol      Utilization[In Percent]
Te 0/6          Up                  0
Dell#
```

Layer 2

The Aggregator supports CLI commands to manage the MAC address table:

- [Clearing the MAC Address Entries](#)
- [Displaying the MAC Address Table](#)

The Aggregator auto-configures with support for Network Interface Controller (NIC) Teaming.

 **NOTE:** On an Aggregator, all ports are configured by default as members of all (4094) VLANs, including the default VLAN. All VLANs operate in Layer 2 mode. You can reconfigure the VLAN membership for individual ports by using the `vlan tagged` or `vlan untagged` commands in INTERFACE configuration mode. See [VLAN Membership](#) for more information.

Supported Modes

Standalone, PMUX, VLT, Stacking

Managing the MAC Address Table

On an Aggregator in VLT and PMUX modes, you can manage the MAC address table by:

- [Clearing the MAC Address Entries](#)
- [Displaying the MAC Address Table](#)

In the Standalone mode, use the `show cam mac stack-unit 0 port-set 0` command to view the mac-addresses.

The Aggregator auto-configures with support for Network Interface Controller (NIC) Teaming.

Clearing the MAC Address Entries

Learned MAC addresses are entered in the table as dynamic entries, which means that they are subject to aging. For any dynamic entry, if no packet arrives on the switch with the MAC address as the source or destination address within the timer period, the address is removed from the table. The default aging time is 1800 seconds in PMUX mode and 300 seconds in Standalone and VLT modes.

You can manually clear the MAC address table of dynamic entries.

To clear a MAC address table, use the following command:

1. Clear a MAC address table of dynamic entries. EXEC Privilege mode


```
clear mac-address-table dynamic {all | interface {tengigabitethernet <0-5> |  
SLOT/PORT} }
```

- `all`: deletes all dynamic entries.
- `interface`: deletes all entries for the specified interface.

Displaying the MAC Address Table

To display the MAC address table, use the following command.

- Display the contents of the MAC address table.
EXEC Privilege mode

 **NOTE:** This command is available only in PMUX mode.

```
show mac-address-table [address | aging-time [vlan vlan-id]| count | dynamic  
| interface | static | vlan]
```

- `address`: displays the specified entry.
- `aging-time`: displays the configured aging-time.
- `count`: displays the number of dynamic and static entries for all VLANs, and the total number of entries.
- `dynamic`: displays only dynamic entries.
- `interface`: displays only entries for the specified interface.
- `static`: displays only static entries.
- `vlan`: displays only entries for the specified VLAN.

Network Interface Controller (NIC) Teaming

NIC teaming is a feature that allows multiple network interface cards in a server to be represented by one MAC address and one IP address in order to provide transparent redundancy, balancing, and to fully utilize network adapter resources.

Support for NIC teaming is auto-configured on the Aggregator, including support for:

- [MAC Address Station Move](#)
- [MAC Move Optimization](#)

The below fig shows a topology where two NICs have been teamed together. In this case, if the primary NIC fails, traffic switches to the secondary NIC, because they are represented by the same set of addresses.

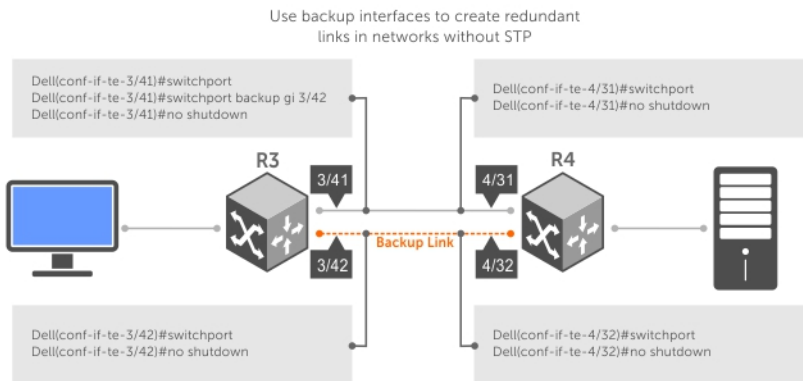



Figure 17. Redundant NOCs with NIC Teaming

MAC Address Station Move

When you use NIC teaming, consider that the server MAC address is originally learned on Port 0/1 of the switch (see figure below). If the NIC fails, the same MAC address is learned on Port 0/5 of the switch. The MAC address is disassociated with one port and re-associated with another in the ARP table; in other words, the ARP entry is "moved". The Aggregator is auto-configured to support MAC Address station moves.




Figure 18. MAC Address Station Move

MAC Move Optimization

Station-move detection takes 5000ms because this is the interval at which the detection algorithm runs.

Link Layer Discovery Protocol (LLDP)

Link layer discovery protocol (LLDP) advertises connectivity and management from the local station to the adjacent stations on an IEEE 802 LAN.

LLDP facilitates multi-vendor interoperability by using standard management tools to discover and make available a physical topology for network management. The Dell Networking operating software implementation of LLDP is based on IEEE standard 801.1ab.

The starting point for using LLDP is invoking LLDP with the protocol `lldp` command in either CONFIGURATION or INTERFACE mode.

The information LLDP distributes is stored by its recipients in a standard management information base (MIB). You can access the information by a network management system through a management protocol such as simple network management protocol (SNMP).

An Aggregator auto-configures to support the link layer discovery protocol (LLDP) for the auto-discovery of network devices. You can use CLI commands to display acquired LLDP information, clear LLDP counters, and debug LACP operation.

Supported Modes

Standalone, PMUX, VLT, Stacking

Protocol Data Units

Configuration information is exchanged in the form of type, length, value (TLV) segments. The below figure shows the chassis ID TLV.

- **Type** — Indicates the type of field that a part of the message represents.
- **Length** — Indicates the size of the value field (in byte).
- **Value** — Indicates the data for this part of the message.

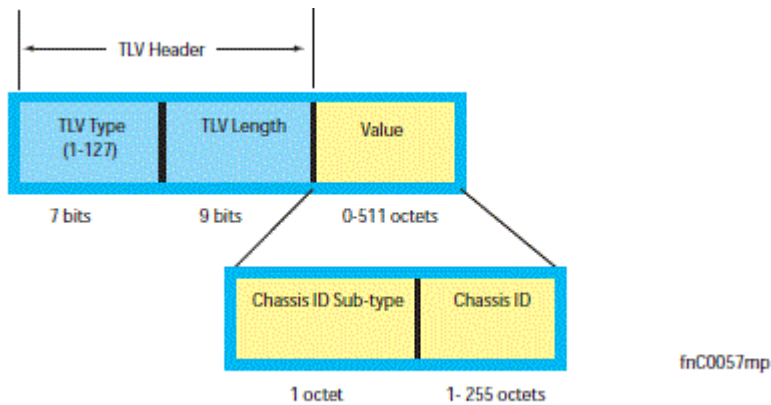


Figure 19. Type, Length, Value (TLV) Segment

TLVs are encapsulated in a frame called an LLDP data unit (LLDPDU), which is transmitted from one LLDP-enabled device to its LLDP-enabled neighbors. LLDP is a one-way protocol. LLDP-enabled devices (LLDP agents) can transmit and/or receive advertisements, but they cannot solicit and do not respond to advertisements.

There are five types of TLVs (as shown in the below table). All types are mandatory in the construction of an LLDPDU except Optional TLVs. You can configure the inclusion of individual Optional TLVs.

Type, Length, Value (TLV) Types

Type	TLV	Description
0	End of LLDPDU	Marks the end of an LLDPDU.
1	Chassis ID	The Chassis ID TLV is a mandatory TLV that identifies the chassis containing the IEEE 802 LAN station associated with the transmitting LLDP agent.
2	Port ID	The Port ID TLV is a mandatory TLV that identifies the port component of the MSAP identifier associated with the transmitting LLDP agent.
3	Time to Live	The Time To Live TLV indicates the number of seconds that the recipient LLDP agent considers the information associated with this MSAP identifier to be valid.
—	Optional	Includes sub-types of TLVs that advertise specific configuration information. These sub-types are Management TLVs, IEEE 802.1, IEEE 802.3, and TIA-1057 Organizationally Specific TLVs.




Figure 20. LLDPDU Frame

Configure LLDP

Configuring LLDP is a two-step process.

1. Enable LLDP globally.
2. Advertise TLVs out of an interface.

Related Configuration Tasks

- [Viewing the LLDP Configuration](#)
- [Viewing Information Advertised by Adjacent LLDP Agents](#)
- [Configuring LLDPDU Intervals](#)
- [Configuring a Time to Live](#)
- [Debugging LLDP](#)

Important Points to Remember

- LLDP is enabled by default.
- Dell Networking systems support up to eight neighbors per interface.
- Dell Networking systems support a maximum of 8000 total neighbors per system. If the number of interfaces multiplied by eight exceeds the maximum, the system does not configure more than 8000.
- INTERFACE level configurations override all CONFIGURATION level configurations.
- LLDP is not hitless.

CONFIGURATION versus INTERFACE Configurations

All LLDP configuration commands are available in PROTOCOL LLDP mode, which is a sub-mode of the CONFIGURATION mode and INTERFACE mode.

- Configurations made at the CONFIGURATION level are global; that is, they affect all interfaces on the system.
- Configurations made at the INTERFACE level affect only the specific interface; they override CONFIGURATION level configurations.

Example of the `protocol lldp` Command (CONFIGURATION Level)

```
Dell(conf)#protocol lldp
Dell(conf-lldp)#?
advertise          Advertise TLVs
disable           Disable LLDP protocol globally
end              Exit from configuration mode
```

```

exit          Exit from LLDP configuration mode
hello        LLDP hello configuration
mode         LLDP mode configuration (default = rx and tx)
multiplier   LLDP multiplier configuration
no           Negate a command or set its defaults
show         Show LLDP configuration

```

```

Dell(conf-lldp)#exit
Dell(conf)#interface tengigabitethernet 0/3
Dell(conf-if-te-0/3)#protocol lldp
Dell(conf-if-te-0/3-lldp)#?
advertise    Advertise TLVs
disable      Disable LLDP protocol on this interface
end          Exit from configuration mode
exit         Exit from LLDP configuration mode
hello        LLDP hello configuration
mode         LLDP mode configuration (default = rx and tx)
multiplier   LLDP multiplier configuration
no           Negate a command or set its defaults
show         Show LLDP configuration
Dell(conf-if-te-0/3-lldp)#

```

Enabling LLDP

LLDP is enabled by default. Enable and disable LLDP globally or per interface. If you enable LLDP globally, all UP interfaces send periodic LLDPDUs.

To enable LLDP, use the following command.

1. Enter Protocol LLDP mode.
CONFIGURATION or INTERFACE mode

```
protocol lldp
```

2. Enable LLDP.
PROTOCOL LLDP mode

```
no disable
```

Disabling and Undoing LLDP

To disable or undo LLDP, use the following command.

- Disable LLDP globally or for an interface.
disable

To undo an LLDP configuration, precede the relevant command with the keyword `no`.

Advertising TLVs

You can configure the system to advertise TLVs out of all interfaces or out of specific interfaces.

- If you configure the system globally, all interfaces send LLDPDUs with the specified TLVs.
- If you configure an interface, only the interface sends LLDPDUs with the specified TLVs.
- If you configure LLDP both globally and at interface level, the interface level configuration overrides the global configuration.

To advertise TLVs, use the following commands.

1. Enter LLDP mode.
CONFIGURATION or INTERFACE mode

```
protocol lldp
```

2. Advertise one or more TLVs.
PROTOCOL LLDP mode

```
advertise {dcbx-appln-tlv | dcbx-tlv | dot3-tlv | interface-port-desc |  
management-tlv | med }
```

Include the keyword for each TLV you want to advertise.

- For management TLVs: `system-capabilities`, `system-description`.
- For 802.1 TLVs: `port-protocol-vlan-id`, `port-vlan-id`.
- For 802.3 TLVs: `max-frame-size`.
- For TIA-1057 TLVs:
 - `guest-voice`
 - `guest-voice-signaling`
 - `location-identification`
 - `power-via-mdi`
 - `softphone-voice`
 - `streaming-video`
 - `video-conferencing`
 - `video-signaling`
 - `voice`
 - `voice-signaling`

In the following example, LLDP is enabled globally. R1 and R2 are transmitting periodic LLDPDUs that contain management, 802.1, and 802.3 TLVs.

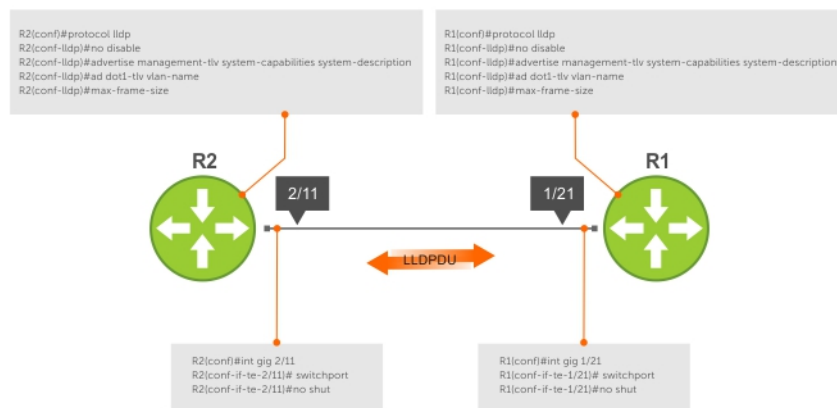


Figure 21. Configuring LLDP

Optional TLVs

The Dell Networking Operating System (OS) supports the following optional TLVs: Management TLVs, IEEE 802.1 and 802.3 organizationally specific TLVs, and TIA-1057 organizationally specific TLVs.

Management TLVs

A management TLV is an optional TLVs sub-type. This kind of TLV contains essential management information about the sender.

Organizationally Specific TLVs

A professional organization or a vendor can define organizationally specific TLVs. They have two mandatory fields (as shown in the following illustration) in addition to the basic TLV fields.

- Organizationally Unique Identifier (OUI)—a unique number assigned by the IEEE to an organization or vendor.
- OUI Sub-type—These sub-types indicate the kind of information in the following data field. The sub-types are determined by the owner of the OUI.



Figure 22. Organizationally Specific TLVs

IEEE Organizationally Specific TLVs

Eight TLV types have been defined by the IEEE 802.1 and 802.3 working groups as a basic part of LLDP; the IEEE OUI is 00-80-C2. You can configure the Dell Networking system to advertise any or all of these TLVs.

Table 9. Optional TLV Types

Type	TLV	Description
Optional TLVs		
4	Port description	A user-defined alphanumeric string that describes the port. The Dell Networking OS does not currently support this TLV.
5	System name	A user-defined alphanumeric string that identifies the system.
6	System description	A user-defined alphanumeric string that identifies the system.
7	System capabilities	Identifies the chassis as one or more of the following: repeater, bridge, WLAN Access Point,

Type	TLV	Description
8	Management address	Router, Telephone, DOCSIS cable device, end station only, or other. Indicates the network address of the management interface. The Dell Networking OS does not currently support this TLV.
IEEE 802.1 Organizationally Specific TLVs		
127	Port-VLAN ID	On Dell Networking systems, indicates the untagged VLAN to which a port belongs.
127	Port and Protocol VLAN ID	On Dell Networking systems, indicates the tagged VLAN to which a port belongs (and the untagged VLAN to which a port belongs if the port is in Hybrid mode).
127	VLAN Name	Indicates the user-defined alphanumeric string that identifies the VLAN.
127	Protocol Identity	Indicates the protocols that the port can process. The Dell Networking OS does not currently support this TLV.
IEEE 802.3 Organizationally Specific TLVs		
127	MAC/PHY Configuration/Status	Indicates the capability and current setting of the duplex status and bit rate, and whether the current settings are the result of auto-negotiation. This TLV is not available in the Dell Networking OS implementation of LLDP, but is available and mandatory (non-configurable) in the LLDP-MED implementation.
127	Power via MDI	Dell Networking supports the LLDP-MED protocol, which recommends that Power via MDI TLV be not implemented, and therefore Dell Networking implements Extended Power via MDI TLV only.

Type	TLV	Description
127	Link Aggregation	Indicates whether the link is capable of being aggregated, whether it is currently in a LAG, and the port identification of the LAG. The Dell Networking OS does not currently support this TLV.
127	Maximum Frame Size	Indicates the maximum frame size capability of the MAC and PHY.

LLDP-MED Capabilities TLV

The LLDP-MED capabilities TLV communicates the types of TLVs that the endpoint device and the network connectivity device support. LLDP-MED network connectivity devices must transmit the Network Policies TLV.

- The value of the LLDP-MED capabilities field in the TLV is a 2-octet bitmap, each bit represents an LLDP-MED capability (as shown in the following table).
- The possible values of the LLDP-MED device type are shown in the following. The Dell Networking system is a network connectivity device, which is Type 4.

When you enable LLDP-MED in Dell Networking OS (using the `advertise med` command), the system begins transmitting this TLV.




Figure 23. LLDP-MED Capabilities TLV

Table 10. Dell Networking OS LLDP-MED Capabilities

Bit Position	TLV	Dell Networking OS Support
0	LLDP-MED Capabilities	Yes
1	Network Policy	Yes
2	Location Identification	Yes
3	Extended Power via MDI-PSE	Yes
4	Extended Power via MDI-PD	No
5	Inventory	No
6–15	reserved	No

Table 11. LLDP-MED Device Types

Value	Device Type
0	Type Not Defined
1	Endpoint Class 1
2	Endpoint Class 2
3	Endpoint Class 3
4	Network Connectivity
5–255	Reserved

LLDP-MED Network Policies TLV

A network policy in the context of LLDP-MED is a device's VLAN configuration and associated Layer 2 and Layer 3 configurations.

LLDP-MED network policies TLV include:

- VLAN ID
- VLAN tagged or untagged status
- Layer 2 priority
- DSCP value

An integer represents the application type (the Type integer shown in the following table), which indicates a device function for which a unique network policy is defined. An individual LLDP-MED network policy TLV is generated for each application type that you specify with the CLI (XXAdvertising TLVs).


 **NOTE:** As shown in the following table, signaling is a series of control packets that are exchanged between an endpoint device and a network connectivity device to establish and maintain a connection. These signal packets might require a different network policy than the media packets for which a connection is made. In this case, configure the signaling application.

Table 12. Network Policy Applications

Type	Application	Description
0	Reserved	—
1	Voice	Specify this application type for dedicated IP telephony handsets and other appliances supporting interactive voice services.
2	Voice Signaling	Specify this application type only if voice control packets use a separate network policy than voice data.
3	Guest Voice	Specify this application type for a separate limited voice service for guest users with their own IP telephony handsets and other appliances supporting interactive voice services.

Type	Application	Description
4	Guest Voice Signaling	Specify this application type only if guest voice control packets use a separate network policy than voice data.
5	Softphone Voice	Specify this application type only if guest voice control packets use a separate network policy than voice data.
6	Video Conferencing	Specify this application type for dedicated video conferencing and other similar appliances supporting real-time interactive video.
7	Streaming Video	Specify this application type for dedicated video conferencing and other similar appliances supporting real-time interactive video.
8	Video Signaling	Specify this application type only if video control packets use a separate network policy than video data.
9–255	Reserved	—




Figure 24. LLDP-MED Policies TLV

Extended Power via MDI TLV

The extended power via MDI TLV enables advanced PoE management between LLDP-MED endpoints and network connectivity devices.

Advertise the extended power via MDI on all ports that are connected to an 802.3af powered, LLDP-MED endpoint device.

- **Power Type** — there are two possible power types: power source entity (PSE) or power device (PD). The Dell Networking system is a PSE, which corresponds to a value of 0, based on the TIA-1057 specification.
- **Power Source** — there are two possible power sources: primary and backup. The Dell Networking system is a primary power source, which corresponds to a value of 1, based on the TIA-1057 specification.
- **Power Priority** — there are three possible priorities: Low, High, and Critical. On Dell Networking systems, the default power priority is **High**, which corresponds to a value of 2 based on the TIA-1057 specification. You can configure a different power priority through the CLI. Dell Networking also honors the power priority value the powered device sends; however, the CLI configuration takes precedence.
- **Power Value** — Dell Networking advertises the maximum amount of power that can be supplied on the port. By default the power is **15.4W**, which corresponds to a power value of 130, based on the TIA-1057 specification. You can advertise a different power value using the `max-milliwatts` option with the `power inline auto | static` command. Dell Networking also honors the power value (power requirement) the powered device sends when the port is configured for `power inline auto`.




Figure 25. Extended Power via MDI TLV

LLDP Operation

On an Aggregator, LLDP operates as follows:

- LLDP is enabled by default.
- LLDPDUs are transmitted and received by default. LLDPDUs are transmitted periodically. The default interval is 30 seconds.
- LLDPDU information received from a neighbor expires after the default Time to Live (TTL) value: 120 seconds.
- Dell Networking OS supports up to eight neighbors per interface.
- Dell Networking OS supports a maximum of 8000 total neighbors per system. If the number of interfaces multiplied by eight exceeds the maximum, the system does not configure more than 8000.
- LLDP is not hitless.

Viewing the LLDP Configuration

To view the LLDP configuration, use the following command.

- Display the LLDP configuration.
CONFIGURATION or INTERFACE mode

```
show config
```

Example of Viewing LLDP Global Configurations

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  hello 10
  no disable
R1(conf-lldp)#
```

Example of Viewing LLDP Interface Configurations

```
R1(conf-lldp)#exit
R1(conf)#interface tengigabitethernet 0/3
R1(conf-if-te-0/3)#show config
!
interface tengigabitEthernet 0/3
  switchport
  no shutdown
R1(conf-if-te-0/3)#protocol lldp
R1(conf-if-te-0/3-lldp)#show config
!
```

```
protocol lldp
R1(conf-if-te-0/3-lldp)#
```

Viewing Information Advertised by Adjacent LLDP Agents

To view brief information about adjacent devices or to view all the information that neighbors are advertising, use the following commands.

- Display brief information about adjacent devices.
`show lldp neighbors`
- Display all of the information that neighbors are advertising.
`show lldp neighbors detail`

Example of Viewing Brief Information Advertised by Neighbors

```
Dell(conf-if-te-0/3-lldp)#end
Dell(conf-if-te-0/3)#do show lldp neighbors
Loc PortID    Rem Host Name    Rem Port Id      Rem Chassis Id
-----
Te 0/2        -                TenGigabitEthernet 0/7    00:01:e8:06:95:3e
Te 0/3        -                TenGigabitEthernet 0/8    00:01:e8:09:c2:4a
```

Example of Viewing Details Advertised by Neighbors

```
Dell#show lldp neighbors detail
=====
Local Interface Te 0/2 has 1 neighbor
  Total Frames Out: 6547
  Total Frames In: 4136
  Total Neighbor information Age outs: 0
  Total Frames Discarded: 0
  Total In Error Frames: 0
  Total Unrecognized TLVs: 0
  Total TLVs Discarded: 0
  Next packet will be sent after 7 seconds
  The neighbors are given below:
-----

Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:01:e8:06:95:3e
Remote Port Subtype: Interface name (5)
Remote Port ID: TenGigabitEthernet 0/7
Local Port ID: TenGigabitEthernet 0/2
Locally assigned remote Neighbor Index: 4
Remote TTL: 120
Information valid for next 120 seconds
Time since last information change of this neighbor: 01:50:16
Remote MTU: 1554
Remote System Desc: Dell Real Time Operating System Software
  Dell Operating System Version: 2.0.
  Dell Application Software Version: 9-4(0-180).
  Copyright (c) 1999-2014 by Dell Inc. All Rights Reserved.
Existing System Capabilities: Repeater Bridge Router
Enabled System Capabilities: Repeater Bridge Router
Remote Port Vlan ID: 1
Port and Protocol Vlan ID: 1, Capability: Supported, Status: Enabled
-----
=====
```

Configuring LLDPDU Intervals

LLDPDU s are transmitted periodically; the default interval is **30 seconds**.

To configure LLDPDU intervals, use the following command.

- Configure a non-default transmit interval.
CONFIGURATION mode or INTERFACE mode

```
hello
```

Example of Viewing LLDPDU Intervals

```
Dell#conf
Dell(conf)#protocol lldp
Dell(conf-lldp)#show config
!
protocol lldp
Dell(conf-lldp)#hello ?
<5-180>                Hello interval in seconds (default=30)
Dell(conf-lldp)#hello 10
Dell(conf-lldp)#show config
!
protocol lldp
hello 10
Dell(conf-lldp)#
Dell(conf-lldp)#no hello
Dell(conf-lldp)#show config
!
protocol lldp
Dell(conf-lldp)#
```

Configuring a Time to Live

The information received from a neighbor expires after a specific amount of time (measured in seconds) called a time to live (TTL).

The TTL is the product of the LLDPDU transmit interval (hello) and an integer called a multiplier. The default multiplier is **4**, which results in a default TTL of 120 seconds.

- Adjust the TTL value.
CONFIGURATION mode or INTERFACE mode.

```
multiplier
```

- Return to the default multiplier value.
CONFIGURATION mode or INTERFACE mode.

```
no multiplier
```

Example of the multiplier Command to Configure Time to Live

```
R1(conf-lldp)#show config
!
protocol lldp
advertise dot1-tlv port-protocol-vlan-id port-vlan-id
advertise dot3-tlv max-frame-size
advertise management-tlv system-capabilities system-description
```

```

no disable
R1(conf-lldp)#multiplier ?
<2-10>          Multiplier (default=4)
R1(conf-lldp)#multiplier 5
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  multiplier 5
  no disable
R1(conf-lldp)#no multiplier
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#

```

Clearing LLDP Counters

You can clear LLDP statistics that are maintained on an Aggregator for LLDP counters for frames transmitted to and received from neighboring devices on all or a specified physical interface.

To clear LLDP counters, enter the `clear lldp counters` command.

Command Syntax	Command Mode	Purpose
<code>clear lldp counters [interface]</code>	EXEC Privilege	Clear counters for LLDP frames sent to and received from neighboring devices on all Aggregator interfaces or on a specified interface.
		<i>interface</i> specifies a 10GbE uplink port in the format: tenGigabitEthernet slot/port .

Debugging LLDP

You can view the TLVs that your system is sending and receiving.

To view the TLVs, use the following commands.

- View a readable version of the TLVs.


```
debug lldp brief
```
- View a readable version of the TLVs plus a hexadecimal version of the entire LLDPDU.


```
debug lldp detail
```

```

Dell# debug lldp interface tengigabitethernet 1/2 packet detail tx
Dell#1w1d19h : Transmit timer blew off for local interface Gi 1/2
1w1d19h : Forming LLDP pkt to send out of interface Gi 1/2
1w1d19h : TLV: Chassis ID, Len: 7, Subtype: Mac address (4), Value: 00:01:e8:0d:b6:d6
1w1d19h : TLV: Port ID, Len: 20, Subtype: Interface name (5), Value: TenGigabitEthernet 1/2
1w1d19h : TLV: TTL, Len: 2, Value: 120
1w1d19h : TLV: SYS_DESC, Len: 207, Value: Dell Networks Real Time Operating System Software. Dell
Operating System Version: 1.0. Dell Application Software Version: E_MAIN4.7.5.276. Copyright (c)1999-Build
Time: Fri Oct 26 12:22:22 PDT 2007
1w1d19h : TLV: SYSTEM CAPAB, Len: 4, Value: Existing: Repeater Bridge Router, Enabled: Repeater Bridge Router
1w1d19h : TLV: ENDOFPDU, Len: 0
1w1d19h : Sending LLDP pkt out of Gi 1/2 of length 270
1w1d19h : Packet dump:
1w1d19h : 01 80 c2 00 00 0e 00 01 e8 0d b7 3b 81 00 00 00
1w1d19h : 88 cc 02 07 04 00 01 e8 0d b6 d6 04 14 05 47 69
1w1d19h : 67 61 62 69 74 45 74 68 65 72 6e 65 74 20 31 2f
1w1d19h : 32 06 02 00 78 0c cf 46 6f 72 63 65 31 30 20 4e
1w1d19h : 65 74 77 6f 72 6b 73 20 52 65 61 6c 20 54 69 6d
1w1d19h : 65 20 4f 70 65 72 61 74 69 6e 67 20 53 79 73 74
1w1d19h : 65 6d 20 53 6f 66 74 77 61 72 65 2e 20 46 6f 72
1w1d19h : 63 65 31 30 20 4f 70 65 72 61 74 69 6e 67 20 53
1w1d19h : 79 73 74 65 6d 20 56 65 72 73 69 6f 6e 3a 20 31
1w1d19h : 2e 30 2e 20 46 6f 72 63 65 31 30 20 41 70 70 6c
1w1d19h : 69 63 61 74 69 6f 6e 20 53 6f 66 74 77 61 72 65
1w1d19h : 20 56 65 72 73 69 6f 6e 3a 20 45 5f 4d 41 49 4e
1w1d19h : 34 2e 37 2e 35 2e 32 37 36 2e 20 43 6f 70 79 72
1w1d19h : 69 67 68 74 20 28 63 29 20 31 39 39 39 2d 42 75
1w1d19h : 69 6c 64 20 54 69 6d 65 3a 20 46 72 69 20 4f 63
1w1d19h : 74 20 32 36 20 31 32 3a 32 32 3a 32 32 20 50 44
1w1d19h : 54 20 32 30 30 37 0e 04 00 16 00 16 00 00
1w1d19h : LLDP frame sent out successfully of Gi 1/2
1w1d19h : Started Transmit timer for Loc interface Gi 1/2 for time 30 sec

```

Figure 26. The debug lldp detail Command — LLDPDU Packet Dissection

Relevant Management Objects

Dell Networkings OS supports all IEEE 802.1AB MIB objects. The following tables list the objects associated with:

- received and transmitted TLVs
- the LLDP configuration on the local agent
- IEEE 802.1AB Organizationally Specific TLVs
- received and transmitted LLDP-MED TLVs

Table 13. LLDP Configuration MIB Objects

MIB Object Category	LLDP Variable	LLDP MIB Object	Description
LLDP Configuration	adminStatus	lldpPortConfigAdminStatus	Whether you enable the local LLDP agent for transmit, receive, or both.
	msgTxHold	lldpMessageTxHoldMultiplier	Multiplier value.

MIB Object Category	LLDP Variable	LLDP MIB Object	Description
	msgTxInterval	lldpMessageTxInterval	Transmit Interval value.
	rxInfoTTL	lldpRxInfoTTL	Time to live for received TLVs.
	txInfoTTL	lldpTxInfoTTL	Time to live for transmitted TLVs.
Basic TLV Selection	mibBasicTLVsTxEnable	lldpPortConfigTLVsTxEnable	Indicates which management TLVs are enabled for system ports.
	mibMgmtAddrInstanceTxEnable	lldpManAddrPortsTxEnable	The management addresses defined for the system and the ports through which they are enabled for transmission.
LLDP Statistics	statsAgeoutsTotal	lldpStatsRxPortAgeoutsTotal	Total number of times that a neighbor's information is deleted on the local system due to an rxInfoTTL timer expiration.
	statsFramesDiscardedTotal	lldpStatsRxPortFramesDiscardedTotal	Total number of LLDP frames received then discarded.
	statsFramesInErrorsTotal	lldpStatsRxPortFramesErrors	Total number of LLDP frames received on a port with errors.
	statsFramesInTotal	lldpStatsRxPortFramesTotal	Total number of LLDP frames received through the port.
	statsFramesOutTotal	lldpStatsTxPortFramesTotal	Total number of LLDP frames transmitted through the port.
	statsTLVsDiscardedTotal	lldpStatsRxPortTLVsDiscardedTotal	Total number of TLVs received then discarded.
	statsTLVsUnrecognizedTotal	lldpStatsRxPortTLVsUnrecognizedTotal	Total number of all TLVs the local agent does not recognize.

Table 14. LLDP System MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
1	Chassis ID	chassis ID subtype	Local	lldpLocChassisIdSub type
			Remote	lldpRemChassisIdSu btype
		chassid ID	Local	lldpLocChassisId
			Remote	lldpRemChassisId
2	Port ID	port subtype	Local	lldpLocPortIdSubtyp e
			Remote	lldpRemPortIdSubty pe
		port ID	Local	lldpLocPortId
			Remote	lldpRemPortId
4	Port Description	port description	Local	lldpLocPortDesc
			Remote	lldpRemPortDesc
5	System Name	system name	Local	lldpLocSysName
			Remote	lldpRemSysName
6	System Description	system description	Local	lldpLocSysDesc
			Remote	lldpRemSysDesc
7	System Capabilities	system capabilities	Local	lldpLocSysCapSupp orted
			Remote	lldpRemSysCapSupp orted
8	Management Address	enabled capabilities	Local	lldpLocSysCapEnabl ed
			Remote	lldpRemSysCapEnab led
		management address length	Local	lldpLocManAddrLen
			Remote	lldpRemManAddrLen
		management address subtype	Local	lldpLocManAddrSubt ype
			Remote	lldpRemManAddrSu btype
		management address	Local	lldpLocManAddr
			Remote	lldpRemManAddr

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
		interface numbering subtype	Local	lldpLocManAddrIfSubtype
			Remote	lldpRemManAddrIfSubtype
		interface number	Local	lldpLocManAddrIfId
			Remote	lldpRemManAddrIfId
		OID	Local	lldpLocManAddrOID
			Remote	lldpRemManAddrOID

Table 15. LLDP 802.1 Organizationally specific TLV MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
127	Port-VLAN ID	PVID	Local	lldpXdot1LocPortVlanId
			Remote	lldpXdot1RemPortVlanId
127	Port and Protocol VLAN ID	port and protocol VLAN supported	Local	lldpXdot1LocProtoVlanSupported
			Remote	lldpXdot1RemProtoVlanSupported
		port and protocol VLAN enabled	Local	lldpXdot1LocProtoVlanEnabled
			Remote	lldpXdot1RemProtoVlanEnabled
		PPVID	Local	lldpXdot1LocProtoVlanId
			Remote	lldpXdot1RemProtoVlanId
127	VLAN Name	VID	Local	lldpXdot1LocVlanId
			Remote	lldpXdot1RemVlanId
		VLAN name length	Local	lldpXdot1LocVlanName
			Remote	lldpXdot1RemVlanName
		VLAN name	Local	lldpXdot1LocVlanName
			Remote	lldpXdot1RemVlanName

Table 16. LLDP-MED System MIB Objects

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
1	LLDP-MED Capabilities	LLDP-MED Capabilities	Local	lldpXMedPortCapSupported
			Remote	lldpXMedPortConfigTLVsTxEnable
		LLDP-MED Class Type	Local	lldpXMedLocDeviceClass
			Remote	lldpXMedRemDeviceClass
2	Network Policy	Application Type	Local	lldpXMedLocMediaPolicyAppType
			Remote	lldpXMedRemMediaPolicyAppType
		Unknown Policy Flag	Local	lldpXMedLocMediaPolicyUnknown
			Remote	lldpXMedLocMediaPolicyUnknown
		Tagged Flag	Local	lldpXMedLocMediaPolicyTagged
			Remote	lldpXMedLocMediaPolicyTagged
		VLAN ID	Local	lldpXMedLocMediaPolicyVlanID
			Remote	lldpXMedRemMediaPolicyVlanID
		L2 Priority	Local	lldpXMedLocMediaPolicyPriority
			Remote	lldpXMedRemMediaPolicyPriority
		DSCP Value	Local	lldpXMedLocMediaPolicyDscp
			Remote	lldpXMedRemMediaPolicyDscp

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object	
3	Location Identifier	Location Data Format	Local	IldpXMedLocLocationSubtype	
			Remote	IldpXMedRemLocationSubtype	
		Location ID Data	Local	IldpXMedLocLocationInfo	
			Remote	IldpXMedRemLocationInfo	
4	Extended Power via MDI	Power Device Type	Local	IldpXMedLocXPoEDeviceType	
			Remote	IldpXMedRemXPoEDeviceType	
			Power Source	Local	IldpXMedLocXPoEPSPowerSource
				Remote	IldpXMedRemXPoEPSPowerSource
		Power Priority	Local		IldpXMedLocXPoEPDPowerSource
					IldpXMedLocXPoEPDPowerSource
			Remote		IldpXMedRemXPoEPSEPowerSource
					IldpXMedRemXPoEPDPowerSource
		Power Value	Local		IldpXMedLocXPoEPDPowerPriority
					IldpXMedLocXPoEPSPortPDPriority
			Remote		IldpXMedRemXPoEPSEPowerPriority
					IldpXMedRemXPoEPDPowerPriority
	Local		IldpXMedLocXPoEPSPortPowerAv		
			IldpXMedLocXPoEPDPowerReq		
	Remote		IldpXMedRemXPoEPSEPowerAv		
			IldpXMedRemXPoEPDPowerReq		

Port Monitoring

The Aggregator supports user-configured port monitoring. See *Configuring Port Monitoring* for the configuration commands to use.

Port monitoring copies all incoming or outgoing packets on one port and forwards (mirrors) them to another port. The source port is the monitored port (MD) and the destination port is the monitoring port (MG).

Supported Modes

Standalone, PMUX, VLT, Stacking

Configuring Port Monitoring

To configure port monitoring, use the following commands.

1. Verify that the intended monitoring port has no configuration other than no shutdown, as shown in the following example.

EXEC Privilege mode

```
show interface
```

```
Dell(conf-if-te-0/1)#show config
!
interface TenGigabitEthernet 0/1
no shutdown
```

2. Create a monitoring session using the command `monitor session` from CONFIGURATION mode, as shown in the following example.

CONFIGURATION mode

```
monitor session
```

```
Dell(conf)#monitor session 1
Dell(conf-mon-sess-1)#source tengig 0/1 destination tengig
0/8 direction both
```

3. Specify the source and destination port and direction of traffic, as shown in the following example.

MONITOR SESSION mode

```
source
```

```
Dell(conf)#monitor session 1
Dell(conf-mon-sess-1)#source tengig 0/1 destination tengig
0/8 direction both
```

NOTE: By default, all uplink ports are assigned to port-channel (LAG) 128 and the destination port in a port monitoring session must be an uplink port. When you configure the destination port using the `source` command, the destination port is removed from LAG 128. To display the uplink ports currently assigned to LAG 128, enter the `show lag 128` command.

Example of Viewing Port Monitoring Configuration

To display information on currently configured port-monitoring sessions, use the `show monitor session` command from EXEC Privilege mode.

```
Dell(conf)# monitor session 0
Dell(conf-mon-sess-0)# source tengig 0/1 dest tengig 0/8 direction rx
Dell(conf-mon-sess-0)#exit
Dell(conf)# do show monitor session 0
```

SessionID	Source	Destination	Direction	Mode	Type
0	TenGig 0/1	TenGig 0/8	rx	interface	Port-based

```
Dell(conf)#
```

In the following example, the host and server are exchanging traffic which passes through the uplink interface 0/1. Port 0/1 is the monitored port and port 0/8 is the destination port, which is configured to only monitor traffic received on tengigabitethernet 0/1 (host-originated traffic).




Figure 27. Port Monitoring Example

Important Points to Remember


- Port monitoring is supported on physical ports only; virtual local area network (VLAN) and port-channel interfaces do not support port monitoring.

- The monitored (the source, [MD]) and monitoring ports (the destination, [MG]) must be on the same switch.
- The monitored (source) interface must be a server-facing interface in the format *slot/port*, where the valid slot numbers are 0 and server-facing port numbers are from 1 to 8.
- The destination interface must be an uplink port (ports 9 to 12).
- In general, a monitoring port should have no ip address and no shutdown as the only configuration; the Dell Networking OS permits a limited set of commands for monitoring ports. You can display these commands using the ? command.
- A monitoring port may not be a member of a VLAN.
- There may only be one destination port in a monitoring session.
- A source port (MD) can only be monitored by one destination port (MG). If you try to assign a monitored port to more than one monitoring port, the following message displays:

```
Dell(conf)#mon ses 1
Dell(conf-mon-sess-1)#source tengig 0/1 destination tengig 0/8 direction both
Dell(conf-mon-sess-1)#do show monitor session
```

SessionID	Source	Destination	Direction	Mode	Type
1	TenGig 0/1	TenGig 0/8	both	interface	Port-based

```
Dell(conf-mon-sess-1)#mon ses 2
Dell(conf-mon-sess-2)#source tengig 0/1 destination tengig 0/8 direction both
% Error: MD port is already being monitored.
```

-  **NOTE:** There is no limit to the number of monitoring sessions per system, provided that there are only four destination ports per port-pipe. If each monitoring session has a unique destination port, the maximum number of session is four per port-pipe.

Port Monitoring

The Aggregator supports multiple source-destination statements in a monitor session, but there may only be one destination port in a monitoring session.

There may only be one destination port in a monitoring session (% Error: Only one MG port is allowed in a session.).

The number of source ports the Dell Networking OS allows within a port-pipe is equal to the number of physical ports in the port-pipe (n). Multiple source ports may have up to four different destination ports (Exceeding max MG ports for this MD port pipe.).

In the following examples, ports 0/1, 0/2, 0/3, and 0/4 all belong to the same port-pipe. These ports mirror traffic to four different destinations (0/9, 0/10, 0/11, and 0/12).

Example of Number of Monitoring Ports

```
Dell#show mon session
SessionID Source Destination Direction Mode Type
-----
0 TenGig 0/1 TenGig 0/9 rx interface Port-based
10 TenGig 0/2 TenGig 0/10 rx interface Port-based
20 TenGig 0/3 TenGig 0/11 rx interface Port-based
30 TenGig 0/4 TenGig 0/12 rx interface Port-based
Dell(conf)#
```

A source port may only be monitored by one destination port, but a destination port may monitor more than one source port.

Dell Networking OS Behavior: All monitored frames are tagged if the configured monitoring direction is transmit (TX), regardless of whether the monitored port (MD) is a Layer 2 or Layer 3 port.

- If the MD port is a Layer 2 port, the frames are tagged with the VLAN ID of the VLAN to which the MD belongs.
- If the MD port is a Layer 3 port, the frames are tagged with VLAN ID 4095.
- If the MD port is in a Layer 3 VLAN, the frames are tagged with the respective Layer 3 VLAN ID.

For example, in the configuration *source tengig 0/1 destination tengig 0/9 direction tx*, if the source port 0/1 is an untagged member of any VLAN, all monitored frames that the destination port 0/9 receives are tagged with the VLAN ID of the source port.

Security

Security features are supported on the I/O Aggregator.

This chapter describes several ways to provide access security to the Dell Networking system.

For details about all the commands described in this chapter, refer to the *Security* chapter in the *Dell PowerEdge FN I/O Aggregator Command Line Reference Guide*.

Supported Modes

Standalone, PMUX, VLT, Stacking

Understanding Banner Settings

This functionality is supported on the Aggregator.

A banner is a note that is displayed when you log in to the system, depending on the privilege level and the command mode into which the you log in. You can specify different banners to be displayed as the message-of-the-day (MOTD), as the opening quote in EXEC mode, or as the beginning message in EXEC Privilege mode. Setting up a banner enables you to display any important information or group-level notification that needs to be communicated to all the users of the system.

A login banner message is displayed only in EXEC Privilege mode after entering the enable command followed by the password. These banners are not displayed to users in EXEC mode. When you connect to a system, the message-of-the-day (MOTD) banner is displayed first, followed by the login banner and prompts. After you log in to the system with valid authentication credentials, the EXEC banner is shown.

You can use the MOTD banner to alert users of critical upcoming events so that they can plan and schedule their accessibility to the device. You can modify the banner messages depending on the requirements or conditions.

Accessing the I/O Aggregator Using the CMC Console Only

This functionality is supported on the Aggregator.

You can enable the option to access and administer an Aggregator only using the chassis management controller (CMC) interface, and prevent the usage of the CLI interface of the device to configure and monitor settings. You can configure the restrict-access session command to disable access of the Aggregator using a Telnet or SSH session; the device is accessible only using the CMC GUI. You can use the no version of this command to reactivate the Telnet or SSH session capability for the device. Use the

show restrict-access command to view whether the access to a device using Telnet or SSH is disabled or not.

AAA Accounting

Accounting, authentication, and authorization (AAA) accounting is part of the AAA security model.

For details about commands related to AAA security, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

AAA accounting enables tracking of services that users are accessing and the amount of network resources being consumed by those services. When you enable AAA accounting, the network server reports user activity to the security server in the form of accounting records. Each accounting record comprises accounting attribute/value (AV) pairs and is stored on the access control server.

As with authentication and authorization, you must configure AAA accounting by defining a named list of accounting methods and then applying that list to various virtual terminal line (VTY) lines.

Configuration Task List for AAA Accounting

The following sections present the AAA accounting configuration tasks.

- [Enabling AAA Accounting](#) (mandatory)
- [Suppressing AAA Accounting for Null Username Sessions](#) (optional)
- [Configuring Accounting of EXEC and Privilege-Level Command Usage](#) (optional)
- [Configuring AAA Accounting for Terminal Lines](#) (optional)
- [Monitoring AAA Accounting](#) (optional)

Enabling AAA Accounting

The `aaa accounting` command allows you to create a record for any or all of the accounting functions monitored.

To enable AAA accounting, use the following command.

- Enable AAA accounting and create a record for monitoring the accounting function.

CONFIGURATION mode

```
aaa accounting {commands | exec | suppress | system level} {default | name}
{start-stop | wait-start | stop-only} {tacacs+}
```

The variables are:

- `system`: sends accounting information of any other AAA configuration.
- `exec`: sends accounting information when a user has logged in to EXEC mode.
- `command level`: sends accounting of commands executed at the specified privilege level.
- `suppress`: Do not generate accounting records for a specific type of user.
- `default | name`: enter the name of a list of accounting methods.
- `start-stop`: use for more accounting information, to send a start-accounting notice at the beginning of the requested event and a stop-accounting notice at the end.
- `wait-start`: ensures that the TACACS+ security server acknowledges the start notice before granting the user's process request.
- `stop-only`: use for minimal accounting; instructs the TACACS+ server to send a stop record accounting notice at the end of the requested user process.

- tacacs+: designate the security service. Currently, Dell Networking OS supports only TACACS+.

Suppressing AAA Accounting for Null Username Sessions

When you activate AAA accounting, the Dell Networking OS software issues accounting records for all users on the system, including users whose username string is NULL because of protocol translation. An example of this is a user who comes in on a line where the AAA authentication `login method-list none` command is applied. To prevent accounting records from being generated for sessions that do not have usernames associated with them, use the following command.

- Prevent accounting records from being generated for users whose username string is NULL.

CONFIGURATION mode

```
aaa accounting suppress null-username
```

Configuring Accounting of EXEC and Privilege-Level Command Usage

The network access server monitors the accounting functions defined in the TACACS+ attribute/value (AV) pairs.

- Configure AAA accounting to monitor accounting functions defined in TACACS+.

CONFIGURATION mode

```
aaa accounting system default start-stop tacacs+
```

```
aaa accounting command 15 default start-stop tacacs+
```

System accounting can use only the default method list.

Example of Configuring AAA Accounting to Track EXEC and EXEC Privilege Level Command Use

In the following sample configuration, AAA accounting is set to track all usage of EXEC commands and commands on privilege level 15.

```
Dell(conf)#aaa accounting exec default start-stop tacacs+
Dell(conf)#aaa accounting command 15 default start-stop tacacs+
```

Configuring AAA Accounting for Terminal Lines

To enable AAA accounting with a named method list for a specific terminal line (where `com15` and `execAcct` are the method list names), use the following commands.

- Configure AAA accounting for terminal lines.

CONFIG-LINE-VTY mode

```
accounting commands 15 com15
```

```
accounting exec execAcct
```

Example of Enabling AAA Accounting with a Named Method List

```
Dell(config-line-vty)# accounting commands 15 com15
Dell(config-line-vty)# accounting exec execAcct
```

Monitoring AAA Accounting

Dell Networking OS does not support periodic interim accounting because the `periodic` command can cause heavy congestion when many users are logged in to the network.

No specific `show` command exists for TACACS+ accounting.

To obtain accounting records displaying information about users currently logged in, use the following command.

- Step through all active sessions and print all the accounting records for the actively accounted functions.

CONFIGURATION mode or EXEC Privilege mode

```
show accounting
```

Example of the `show accounting` Command for AAA Accounting


```
Dell#show accounting
Active accounted actions on tty2, User admin Priv 1
  Task ID 1, EXEC Accounting record, 00:00:39 Elapsed, service=shell
Active accounted actions on tty3, User admin Priv 1
  Task ID 2, EXEC Accounting record, 00:00:26 Elapsed, service=shell
Dell#
```

AAA Authentication

Dell Networking OS supports a distributed client/server system implemented through authentication, authorization, and accounting (AAA) to help secure networks against unauthorized access.

In the Dell Networking implementation, the Dell Networking system acts as a RADIUS or TACACS+ client and sends authentication requests to a central remote authentication dial-in service (RADIUS) or Terminal access controller access control system plus (TACACS+) server that contains all user authentication and network service access information.

Dell Networking uses local usernames/passwords (stored on the Dell Networking system) or AAA for login authentication. With AAA, you can specify the security protocol or mechanism for different login methods and different users. In Dell Networking OS, AAA uses a list of authentication methods, called method lists, to define the types of authentication and the sequence in which they are applied. You can define a method list or use the default method list. User-defined method lists take precedence over the default method list.

 **NOTE:** If a console user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server if the privilege level is configured for that user in RADIUS, whether you configure RADIUS authorization.

Configuration Task List for AAA Authentication

The following sections provide the configuration tasks.

- [Configure Login Authentication for Terminal Lines](#)
- [Configuring AAA Authentication Login Methods](#)
- [Enabling AAA Authentication](#)

For a complete list of all commands related to login authentication, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Configure Login Authentication for Terminal Lines

You can assign up to five authentication methods to a method list. Dell Networking OS evaluates the methods in the order in which you enter them in each list.

If the first method list does not respond or returns an error, Dell Networking OS applies the next method list until the user either passes or fails the authentication. If the user fails a method list, Dell Networking OS does not apply the next method list.

Configuring AAA Authentication Login Methods

To configure an authentication method and method list, use the following commands.

Dell Networking OS Behavior: If you use a method list on the console port in which RADIUS or TACACS is the last authentication method, and the server is not reachable, Dell Networking OS allows access even though the username and password credentials cannot be verified. Only the console port behaves this way, and does so to ensure that users are not locked out of the system if network-wide issue prevents access to these servers.

1. Define an authentication method-list (*method-list-name*) or specify the default.

CONFIGURATION mode

```
aaa authentication login {method-list-name | default} method1 [... method4]
```

The default method-list is applied to all terminal lines.

Possible methods are:

- `enable`: use the password you defined using the `enable secret` or `enable password` command in CONFIGURATION mode.
- `line`: use the password you defined using the `password` command in LINE mode.
- `local`: use the username/password database defined in the local configuration.
- `none`: no authentication.
- `radius`: use the RADIUS servers configured with the `radius-server host` command.
- `tacacs+`: use the TACACS+ servers configured with the `tacacs-server host` command.

2. Enter LINE mode.

CONFIGURATION mode


```
line {aux 0 | console 0 | vty number [... end-number]}
```

3. Assign a *method-list-name* or the default list to the terminal line.

LINE mode

```
login authentication {method-list-name | default}
```

To view the configuration, use the `show config` command in LINE mode or the `show running-config` in EXEC Privilege mode.

 **NOTE:** Dell Networking recommends using the `none` method only as a backup. This method does not authenticate users. The `none` and `enable` methods do not work with secure shell (SSH).

You can create multiple method lists and assign them to different terminal lines.

Enabling AAA Authentication

To enable AAA authentication, use the following command.

- Enable AAA authentication.
CONFIGURATION mode

```
aaa authentication enable {method-list-name | default} method1 [... method4]
```

- *default*: uses the listed authentication methods that follow this argument as the default list of methods when a user logs in.
- *method-list-name*: character string used to name the list of enable authentication methods activated when a user logs in.
- *method1 [... method4]*: any of the following: RADIUS, TACACS, enable, line, none.

If you do not set the default list, only the local enable is checked. This setting has the same effect as issuing an `aaa authentication enable default enable` command.

Enabling AAA Authentication — RADIUS

To enable authentication from the RADIUS server, and use TACACS as a backup, use the following commands.

1. Enable RADIUS and set up TACACS as backup.

CONFIGURATION mode

```
aaa authentication enable default radius tacacs
```

2. Establish a host address and password.

CONFIGURATION mode

```
radius-server host x.x.x.x key some-password
```

3. Establish a host address and password.

CONFIGURATION mode

```
tacacs-server host x.x.x.x key some-password
```

To get enable authentication from the RADIUS server and use TACACS as a backup, issue the following commands.

Example of Enabling Authentication from the RADIUS Server

```
Dell(config)# aaa authentication enable default radius tacacs
Radius and TACACS server has to be properly setup for this.
Dell(config)# radius-server host x.x.x.x key <some-password>
Dell(config)# tacacs-server host x.x.x.x key <some-password>
```

To use local authentication for `enable secret` on the console, while using remote authentication on VTY lines, issue the following commands.

Example of Enabling Local Authentication for the Console and Remote Authentication for VTY Lines

```
Dell(config)# aaa authentication enable mymethodlist radius tacacs
Dell(config)# line vty 0 9
Dell(config-line-vty)# enable authentication mymethodlist
```

Server-Side Configuration

- **TACACS+** — When using TACACS+, Dell Networking OS sends an initial packet with service type SVC_ENABLE, and then sends a second packet with just the password. The TACACS server must have an entry for username \$enable\$.
- **RADIUS** — When using RADIUS authentication, Dell Networking OS sends an authentication packet with the following:
Username: \$enab15\$
Password: <password-entered-by-user>

Therefore, the RADIUS server must have an entry for this username.

AAA Authorization

The Dell Networking OS enables AAA new-model by default.

You can set authorization to be either `local` or `remote`. Different combinations of authentication and authorization yield different results. By default, the system sets both to **local**.

Privilege Levels Overview

Limiting access to the system is one method of protecting the system and your network. However, at times, you might need to allow others access to the router and you can limit that access to a subset of commands. In the Dell Networking OS, you can configure a privilege level for users who need limited access to the system.

Every command in the Dell Networking OS is assigned a privilege level of 0, 1, or 15. You can configure up to 16 privilege levels. The Dell Networking OS is pre-configured with three privilege levels and you can configure 13 more. The three pre-configured levels are:

- **Privilege level 1** — is the default level for EXEC mode. At this level, you can interact with the router, for example, view some `show` commands and Telnet and ping to test connectivity, but you cannot configure the router. This level is often called the “user” level. One of the commands available in Privilege level 1 is the `enable` command, which you can use to enter a specific privilege level.
- **Privilege level 0** — contains only the `end`, `enable`, and `disable` commands.
- **Privilege level 15** — the default level for the `enable` command, is the highest level. In this level you can access any command in the Dell Networking OS.

Privilege levels 2 through 14 are not configured and you can customize them for different users and access.

After you configure other privilege levels, enter those levels by adding the level parameter after the `enable` command or by configuring a user name or password that corresponds to the privilege level. For more information about configuring user names, refer to [Configuring a Username and Password](#).

By default, commands are assigned to different privilege levels. You can access those commands only if you have access to that privilege level. For example, to reach the `protocol spanning-tree` command,

log in to the router, enter the `enable` command for privilege level 15 (this privilege level is the default level for the command) and then enter CONFIGURATION mode.

You can configure passwords to control access to the box and assign different privilege levels to users. The Dell Networking OS supports the use of passwords when you log in to the system and when you enter the `enable` command. If you move between privilege levels, you are prompted for a password if you move to a higher privilege level.

Configuration Task List for Privilege Levels

The following list has the configuration tasks for privilege levels and passwords.

- [Configuring a Username and Password](#) (mandatory)
- [Configuring the Enable Password Command](#) (mandatory)
- [Configuring Custom Privilege Levels](#) (mandatory)
- [Specifying LINE Mode Password and Privilege](#) (optional)
- [Enabling and Disabling Privilege Levels](#) (optional)

For a complete listing of all commands related to privilege levels and passwords, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Configuring a Username and Password

In the Dell Networking OS, you can assign a specific username to limit user access to the system. To configure a username and password, use the following command.

- Assign a user name and password.

CONFIGURATION mode

```
username name [access-class access-list-name] [nopassword | password
[encryption-type] password] [privilege level] [secret]
```

Configure the optional and required parameters:

- `name`: Enter a text string up to 63 characters long.
- `access-class access-list-name`: Restrict access by access-class.
- `nopassword`: Require password for the user to login.
- `encryption-type`: Enter 0 for plain text or 7 for encrypted text.
- `password`: Enter a string. Specify the password for the user.
- `privilege level`: The range is from 0 to 15.
- `secret`: Specify the secret for the user.

To view username, use the `show users` command in EXEC Privilege mode.

Configuring the Enable Password Command

To configure the Dell Networking OS, use the `enable` command to enter EXEC Privilege level 15. After entering the command, the system requests that you enter a password.

Privilege levels are not assigned to passwords, rather passwords are assigned to a privilege level. You can always change a password for any privilege level. To change to a different privilege level, enter the `enable` command, then the privilege level. If you do not enter a privilege level, the default level **15** is assumed.

To configure a password for a specific privilege level, use the following command.

- Configure a password for a privilege level.

CONFIGURATION mode

```
enable password [level level] [encryption-mode] password
```

Configure the optional and required parameters:

- *level level*: Specify a level from 0 to 15. Level 15 includes all levels.
- *encryption-type*: Enter 0 for plain text or 7 for encrypted text.
- *password*: Enter a string.

To change only the password for the `enable` command, configure only the *password* parameter.

To view the configuration for the `enable secret` command, use the `show running-config` command in EXEC Privilege mode.

In custom-configured privilege levels, the `enable` command is always available. No matter what privilege level you entered, you can enter the `enable 15` command to access and configure all CLIs.

Configuring Custom Privilege Levels

In addition to assigning privilege levels to the user, you can configure the privilege levels of commands so that they are visible in different privilege levels.

Within the Dell Networking OS, commands have certain privilege levels. With the `privilege` command, you can change the default level or you can reset their privilege level back to the default. Assign the launch keyword (for example, `configure`) for the keyword's command mode.

To assign commands and passwords to a custom privilege level, use the following commands. You must be in privilege level 15.

1. Assign a user name and password.

CONFIGURATION mode

```
username name [access-class access-list-name] [privilege level] [nopassword  
| password [encryption-type] password] [secret]
```

Configure the optional and required parameters:

- *name*: enter a text string (up to 63 characters).
- *access-class access-list-name*: enter the name of a configured IP ACL.
- *privilege level*: the range is from 0 to 15.
- *nopassword*: do not require the user to enter a password.
- *encryption-type*: enter 0 for plain text or 7 for encrypted text.
- *password*: enter a text string.
- *secret*: specify the secret for the user.

2. Configure a password for privilege level.

CONFIGURATION mode

```
enable password [level level] [encryption-mode] password
```

Configure the optional and required parameters:

- *level level*: specify a level from 0 to 15. Level 15 includes all levels.
- *encryption-type*: enter 0 for plain text or 7 for encrypted text.

- *password*: enter a text string up to 32 characters long.

To change only the password for the `enable` command, configure only the `password` parameter.

3. Configure level and commands for a mode or reset a command's level.

CONFIGURATION mode

```
privilege mode {level level command | reset command}
```

Configure the following required and optional parameters:

- *mode*: enter a keyword for the modes (`exec`, `configure`, `interface`, `line`, `route-map`, or `router`)
- *level level*: the range is from 0 to 15. Levels 0, 1, and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.
- *command*: an Dell CLI keyword (up to five keywords allowed).
- *reset*: return the command to its default privilege mode.

To view the configuration, use the `show running-config` command in EXEC Privilege mode.

The following example shows a configuration to allow a user *john* to view only EXEC mode commands and all `snmp-server` commands. Because the `snmp-server` commands are *enable* level commands and, by default, found in CONFIGURATION mode, also assign the launch command for CONFIGURATION mode, `configure`, to the same privilege level as the `snmp-server` commands.

Line 1: The user *john* is assigned privilege level 8 and assigned a password.

Line 2: All other users are assigned a password to access privilege level 8.

Line 3: The `configure` command is assigned to privilege level 8 because it needs to reach CONFIGURATION mode where the `snmp-server` commands are located.

Line 4: The `snmp-server` commands, in CONFIGURATION mode, are assigned to privilege level 8.

Example of Configuring a Custom Privilege Level

```
Dell(conf)#username john privilege 8 password john
Dell(conf)#enable password level 8 notjohn
Dell(conf)#privilege exec level 8 configure
Dell(conf)#privilege config level 8 snmp-server
Dell(conf)#end
Dell#show running-config
Current Configuration ...
!
hostname FTOS
!
enable password level 8 notjohn
enable password FTOS
!
username admin password 0 admin
username john password 0 john privilege 8
!
```

The following example shows the Telnet session for user *john*. The `show privilege` command output confirms that *john* is in privilege level 8. In EXEC Privilege mode, *john* can access only the commands listed. In CONFIGURATION mode, *john* can access only the `snmp-server` commands.

Example of Privilege Level Login and Available Commands

```
apollo% telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: john
Password:
Dell#show priv
Current privilege level is 8
Dell#?
configure      Configuring from terminal
disable        Turn off privileged commands
enable         Turn on privileged commands
exit           Exit from the EXEC
no             Negate a command
show           Show running system information
terminal       Set terminal line parameters
traceroute     Trace route to destination
Dell#confi
Dell(conf)#?
end            Exit from Configuration mode
```

Specifying LINE Mode Password and Privilege

You can specify a password authentication of all users on different terminal lines. The user's privilege level is the same as the privilege level assigned to the terminal line, unless a more specific privilege level is assigned to the user.

To specify a password for the terminal line, use the following commands.

- Configure a custom privilege level for the terminal lines.

LINE mode

```
privilege level level
```

- *level level*: The range is from 0 to 15. Levels 0, 1, and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.

- Specify either a plain text or encrypted password.

LINE mode

```
password [encryption-type] password
```

Configure the following optional and required parameters:

- *encryption-type*: Enter 0 for plain text or 7 for encrypted text.
- *password*: Enter a text string up to 25 characters long.

To view the password configured for a terminal, use the `show config` command in LINE mode.

Enabling and Disabling Privilege Levels

To enable and disable privilege levels, use the following commands.

- Set a user's security level.

EXEC Privilege mode

```
enable or enable privilege-level
```

If you do not enter a privilege level, the system sets it to 15 by default.

- Move to a lower privilege level.

EXEC Privilege mode

```
disable level-number
```

- *level-number*: The level-number you wish to set.

If you enter `disable` without a level-number, your security level is 1.

RADIUS

Remote authentication dial-in user service (RADIUS) is a distributed client/server protocol.

This protocol transmits authentication, authorization, and configuration information between a central RADIUS server and a RADIUS client (the Dell Networking system). The system sends user information to the RADIUS server and requests authentication of the user and password. The RADIUS server returns one of the following responses:

- **Access-Accept** — the RADIUS server authenticates the user.
- **Access-Reject** — the RADIUS server does not authenticate the user.

If an error occurs in the transmission or reception of RADIUS packets, you can view the error by enabling the `debug radius` command.

Transactions between the RADIUS server and the client are encrypted (the users' passwords are not sent in plain text). RADIUS uses UDP as the transport protocol between the RADIUS server host and the client.

For more information about RADIUS, refer to RFC 2865, *Remote Authentication Dial-in User Service*.

RADIUS Authentication

Dell Networking OS supports RADIUS for user authentication (text password) at login and can be specified as one of the login authentication methods in the `aaa authentication login` command.

Idle Time

Every session line has its own idle-time. If the idle-time value is not changed, the default value of **30 minutes** is used.

RADIUS specifies idle-time allow for a user during a session before timeout. When a user logs in, the lower of the two idle-time values (configured or default) is used. The idle-time value is updated if both of the following happens:

- The administrator changes the idle-time of the line on which the user has logged in.
- The idle-time is lower than the RADIUS-returned idle-time.

Configuration Task List for RADIUS


To authenticate users using RADIUS, you must specify at least one RADIUS server so that the system can communicate with and configure RADIUS as one of your authentication methods.

The following list includes the configuration tasks for RADIUS.

- [Defining a AAA Method List to be Used for RADIUS](#) (mandatory)
- [Applying the Method List to Terminal Lines](#) (mandatory except when using default lists)

- [Specifying a RADIUS Server Host](#) (mandatory)
- [Setting Global Communication Parameters for all RADIUS Server Hosts](#) (optional)
- [Monitoring RADIUS](#) (optional)

For a complete listing of all Dell Networking OS commands related to RADIUS, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

 **NOTE:** RADIUS authentication and authorization are done in a single step. Hence, authorization cannot be used independent of authentication. However, if you have configured RADIUS authorization and have not configured authentication, a message is logged stating this. During authorization, the next method in the list (if present) is used, or if another method is not present, an error is reported.

To view the configuration, use the `show config` in LINE mode or the `show running-config` command in EXEC Privilege mode.

Defining a AAA Method List to be Used for RADIUS

To configure RADIUS to authenticate or authorize users on the system, create a AAA method list. Default method lists do not need to be explicitly applied to the line, so they are not mandatory.

To create a method list, use the following commands.

- Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the RADIUS authentication method.

CONFIGURATION mode

```
aaa authentication login method-list-name radius
```

- Create a method list with RADIUS and TACACS+ as authorization methods.

CONFIGURATION mode

```
aaa authorization exec {method-list-name | default} radius tacacs+
```

Typical order of methods: RADIUS, TACACS+, Local, None.

If RADIUS denies authorization, the session ends (RADIUS must not be the last method specified).

Applying the Method List to Terminal Lines

To enable RADIUS AAA login authentication for a method list, apply it to a terminal line.

To configure a terminal line for RADIUS authentication and authorization, use the following commands.

- Enter LINE mode.

CONFIGURATION mode

```
line {aux 0 | console 0 | vty number [end-number]}
```

- Enable AAA login authentication for the specified RADIUS method list.

LINE mode

```
login authentication {method-list-name | default}
```

This procedure is mandatory if you are not using default lists.

- To use the method list.

CONFIGURATION mode

```
authorization exec methodlist
```

Specifying a RADIUS Server Host

When configuring a RADIUS server host, you can set different communication parameters, such as the UDP port, the key password, the number of retries, and the timeout.

To specify a RADIUS server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the RADIUS server host.

CONFIGURATION mode

```
radius-server host {hostname | ip-address} [auth-port port-number]
[retransmit retries] [timeout seconds] [key [encryption-type] key]
```

Configure the optional communication parameters for the specific host:

- `auth-port port-number`: the range is from 0 to 65535. Enter a UDP port number. The default is **1812**.
- `retransmit retries`: the range is from 0 to 100. Default is **3**.
- `timeout seconds`: the range is from 0 to 1000. Default is **5 seconds**.
- `key [encryption-type] key`: enter 0 for plain text or 7 for encrypted text, and a string for the key. The key can be up to 42 characters long. This key must match the key configured on the RADIUS server host.

If you do not configure these optional parameters, the global default values for all RADIUS host are applied.

To specify multiple RADIUS server hosts, configure the `radius-server host` command multiple times. If you configure multiple RADIUS server hosts, Dell Networking OS attempts to connect with them in the order in which they were configured. When Dell Networking OS attempts to authenticate a user, the software connects with the RADIUS server hosts one at a time, until a RADIUS server host responds with an accept or reject response.

If you want to change an optional parameter setting for a specific host, use the `radius-server host` command. To change the global communication settings to all RADIUS server hosts, refer to [Setting Global Communication Parameters for all RADIUS Server Hosts](#).

To view the RADIUS configuration, use the `show running-config radius` command in EXEC Privilege mode.

To delete a RADIUS server host, use the `no radius-server host {hostname | ip-address}` command.

Setting Global Communication Parameters for all RADIUS Server Hosts

You can configure global communication parameters (`auth-port`, `key`, `retransmit`, and `timeout` parameters) and specific host communication parameters on the same system.

However, if you configure both global and specific host parameters, the specific host parameters override the global parameters for that RADIUS server host.

To set global communication parameters for all RADIUS server hosts, use the following commands.

- Set a time interval after which a RADIUS host server is declared dead.

CONFIGURATION mode

```
radius-server deadtime seconds
```

- *seconds*: the range is from 0 to 2147483647. The default is **0 seconds**.
- Configure a key for all RADIUS communications between the system and RADIUS server hosts.
CONFIGURATION mode

```
radius-server key [encryption-type] key
```

- *encryption-type*: enter 7 to encrypt the password. Enter 0 to keep the password as plain text.
- *key*: enter a string. The key can be up to 42 characters long. You cannot use spaces in the key.
- Configure the number of times Dell Networking OS retransmits RADIUS requests.
CONFIGURATION mode

```
radius-server retransmit retries
```

- *retries*: the range is from 0 to 100. Default is **3 retries**.
- Configure the time interval the system waits for a RADIUS server host response.
CONFIGURATION mode

```
radius-server timeout seconds
```

- *seconds*: the range is from 0 to 1000. Default is **5 seconds**.

To view the configuration of RADIUS communication parameters, use the `show running-config` command in EXEC Privilege mode.

Monitoring RADIUS

To view information on RADIUS transactions, use the following command.

- View RADIUS transactions to troubleshoot problems.
EXEC Privilege mode

```
debug radius
```

TACACS+

Dell Networking OS supports terminal access controller access control system (TACACS+ client, including support for login authentication.

Configuration Task List for TACACS+

The following list includes the configuration task for TACACS+ functions.

- [Choosing TACACS+ as the Authentication Method](#)
- [Monitoring TACACS+](#)
- [TACACS+ Remote Authentication](#)
- [Specifying a TACACS+ Server Host](#)

For a complete listing of all commands related to TACACS+, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Choosing TACACS+ as the Authentication Method

One of the login authentication methods available is TACACS+ and the user's name and password are sent for authentication to the TACACS hosts specified.

To use TACACS+ to authenticate users, specify at least one TACACS+ server for the system to communicate with and configure TACACS+ as one of your authentication methods.

To select TACACS+ as the login authentication method, use the following commands.

1. Configure a TACACS+ server host.

CONFIGURATION mode

```
tacacs-server host {ip-address | host}
```

Enter the IP address or host name of the TACACS+ server.

Use this command multiple times to configure multiple TACACS+ server hosts.

2. Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the TACACS+ authentication method.

CONFIGURATION mode

```
aaa authentication login {method-list-name | default} tacacs+ [...method3]
```

The TACACS+ method must not be the last method specified.

3. Enter LINE mode.

CONFIGURATION mode

```
line {aux 0 | console 0 | vty number [end-number]}
```

4. Assign the *method-list* to the terminal line.

LINE mode

```
login authentication {method-list-name | default}
```

Example of a Failed Authentication

To view the configuration, use the `show config` in LINE mode or the `show running-config tacacs +` command in EXEC Privilege mode.

If authentication fails using the primary method, Dell Networking OS employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, Dell Networking OS proceeds to the next authentication method. In the following example, the TACACS+ is incorrect, but the user is still authenticated by the secondary method.

First bold line: Server key purposely changed to incorrect value.

Second bold line: User authenticated using the secondary method.

```
Dell(conf)#
Dell(conf)#do show run aaa
!
aaa authentication enable default tacacs+ enable
aaa authentication enable LOCAL enable tacacs+
aaa authentication login default tacacs+ local
aaa authentication login LOCAL local tacacs+
aaa authorization exec default tacacs+ none
aaa authorization commands 1 default tacacs+ none
```



```

aaa authorization commands 15 default tacacs+ none
aaa accounting exec default start-stop tacacs+
aaa accounting commands 1 default start-stop tacacs+
aaa accounting commands 15 default start-stop tacacs+
Dell(conf)#
Dell(conf)#do show run tacacs+
!
tacacs-server key 7 d05206c308f4d35b
tacacs-server host 10.10.10.10 timeout 1
Dell(conf)#tacacs-server key angeline
Dell(conf)##RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user admin on
vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 ( 10.11.9.209 )
%RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated for user admin on line
vty0 (10.11.9.209)
Dell(conf)#username angeline password angeline
Dell(conf)##RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user angeline
on vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 ( 10.11.9.209 )

```

Monitoring TACACS+

To view information on TACACS+ transactions, use the following command.

- View TACACS+ transactions to troubleshoot problems.
EXEC Privilege mode

```
debug tacacs+
```

TACACS+ Remote Authentication

When configuring a TACACS+ server host, you can set different communication parameters, such as the key password.

Example of Specifying a TACACS+ Server Host

```

Dell(conf)#
Dell(conf)#aaa authentication login tacacsmethod tacacs+
Dell(conf)#aaa authentication exec tacacsauthorization tacacs+
Dell(conf)#tacacs-server host 25.1.1.2 key Force
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication tacacsmethod
Dell(config-line-vty)#end

```

Specifying a TACACS+ Server Host

To specify a TACACS+ server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the TACACS+ server host.
CONFIGURATION mode

```
tacacs-server host {hostname | ip-address} [port port-number] [timeout
seconds] [key key]
```

Configure the optional communication parameters for the specific host:

- `port port-number`: the range is from 0 to 65535. Enter a TCP port number. The default is **49**.
- `timeout seconds`: the range is from 0 to 1000. Default is **10 seconds**.
- `key key`: enter a string for the key. The key can be up to 42 characters long. This key must match a key configured on the TACACS+ server host. This parameter must be the last parameter you configure.

If you do not configure these optional parameters, the default global values are applied.

Example of Connecting with a TACACS+ Server Host

To specify multiple TACACS+ server hosts, configure the `tacacs-server host` command multiple times. If you configure multiple TACACS+ server hosts, Dell Networking OS attempts to connect with them in the order in which they were configured.

To view the TACACS+ configuration, use the `show running-config tacacs+` command in EXEC Privilege mode.


To delete a TACACS+ server host, use the `no tacacs-server host {hostname | ip-address}` command.

```
freebsd2# telnet 2200:2200:2200:2200:2200::2202
Trying 2200:2200:2200:2200:2200::2202...
Connected to 2200:2200:2200:2200:2200::2202.
Escape character is '^]'.
Login: admin
Password:
Dell#
Dell#
```

Enabling SCP and SSH

Secure shell (SSH) is a protocol for secure remote login and other secure network services over an insecure network. Dell Networking OS is compatible with SSH versions 1.5 and 2, in both the client and server modes. SSH sessions are encrypted and use authentication. SSH is enabled by default. For details about the command syntax, refer to the *Security* chapter in the *Dell Networking OS Command Line Interface Reference Guide*.

Dell Networking OS SCP, which is a remote file copy program that works with SSH.

 **NOTE:** The Windows-based WinSCP client software is not supported for secure copying between a PC and a Dell Networking OS-based system. Unix-based SCP client software is supported.

To use the SSH client, use the following command.

- Open an SSH connection and specify the hostname, username, port number, encryption cipher, HMAC algorithm and version of the SSH client.

EXEC Privilege mode

```
ssh {hostname} [-l username | -p port-number | -v {1 | 2}] -c encryption
cipher | -m HMAC algorithm
```

hostname is the IP address or host name of the remote device. Enter an IPv4 or IPv6 address in dotted decimal format (A.B.C.D).

- SSH V2 is enabled by default on all the modes.
- Display SSH connection information.

EXEC Privilege mode

```
show ip ssh
```

Specifying an SSH Version

The following example uses the `ip ssh server version 2` command to enable SSH version 2 and the `show ip ssh` command to confirm the setting.

```
Dell(conf)#ip ssh server version 2
Dell(conf)#do show ip ssh
SSH server          : enabled.
SSH server version  : v2.
SSH server vrf      : default.
SSH server ciphers  : 3des-cbc,aes128-cbc,aes192-cbc,aes256-cbc,aes128-
ctr,aes192-ctr,aes256-ctr.
SSH server macs     : hmac-md5,hmac-md5-96,hmac-sha1,hmac-sha1-96,hmac-
sha2-256,hmac-sha2-256-96.
SSH server kex algorithms : diffie-hellman-group-exchange-sha1,diffie-hellman-
group1-sha1,diffie-hellman-group14-sha1.
Password Authentication : enabled.
Hostbased Authentication : disabled.
RSA Authentication    : disabled.
      Vty          Encryption          HMAC          Remote IP
Dell(conf)#
```

To disable SSH server functions, use the `no ip ssh server enable` command.

Using SCP with SSH to Copy a Software Image

To use secure copy (SCP) to copy a software image through an SSH connection from one switch to another, use the following commands.

```
On the chassis, invoke SCP.
CONFIGURATION mode
```

```
copy scp: flash:
```

Example of Using SCP to Copy from an SSH Server on Another Switch

The following example shows the use of SCP and SSH to copy a software image from one switch running SSH server on UDP port 99 to the local switch.

```
Dell#copy scp: flash:
Address or name of remote host []: 10.10.10.1
Port number of the server [22]: 99
Source file name []: test.cfg
User name to login remote host: admin
Password to login remote host:
```

Secure Shell Authentication

Secure Shell (SSH) is enabled by default using the SSH Password Authentication method.

Telnet

To use Telnet with SSH, first enable SSH, as previously described.

By default, the Telnet daemon is enabled. If you want to disable the Telnet daemon, use the following command, or disable Telnet in the startup config. To enable or disable the Telnet daemon, use the `[no] ip telnet server enable` command.

Example of Using Telnet for Remote Login

```
Dell(conf)#ip telnet server enable
Dell(conf)#no ip telnet server enable
```

VTY Line and Access-Class Configuration

Various methods are available to restrict VTY access in Dell Networking OS. These depend on which authentication scheme you use – line, local, or remote.

Table 17. VTY Access

Authentication Method	VTY access-class support?	Username access-class support?	Remote authorization support?
Line	YES	NO	NO
Local	NO	YES	NO
TACACS+	YES	NO	YES (with Dell Networking OS version 5.2.1.0 and later)
RADIUS	YES	NO	YES (with Dell Networking OS version 6.1.1.0 and later)

Dell Networking OS provides several ways to configure access classes for VTY lines, including:

- [VTY Line Local Authentication and Authorization](#)
- [VTY Line Remote Authentication and Authorization](#)

VTY Line Local Authentication and Authorization

Dell Networking OS retrieves the access class from the local database.


To use this feature:

1. Create a username.
2. Enter a password.
3. Assign an access class.
4. Enter a privilege level.

You can assign line authentication on a per-VTY basis; it is a simple password authentication, using an access-class as authorization.

Configure local authentication globally and configure access classes on a per-user basis.

Dell Networking OS can assign different access classes to different users by username. Until users attempt to log in, Dell Networking OS does not know if they will be assigned a VTY line. This means that incoming users always see a login prompt even if you have excluded them from the VTY line with a deny-all access class. After users identify themselves, Dell Networking OS retrieves the access class from the local database and applies it. (Dell Networking OS then can close the connection if a user is denied access.)

 **NOTE:** If a VTY user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server only if you configure RADIUS authentication.

The following example shows how to allow or deny a Telnet connection to a user. Users see a login prompt even if they cannot log in. No access class is configured for the VTY line. It defaults from the local database.

Example of Configuring VTY Authorization Based on Access Class Retrieved from a Local Database (Per User)

```
Dell(conf)#user gooduser password abc privilege 10 access-class permitall
Dell(conf)#user baduser password abc privilege 10 access-class denyall
Dell(conf)#
Dell(conf)#aaa authentication login localmethod local
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication localmethod
Dell(config-line-vty)#end
```

VTY Line Remote Authentication and Authorization

Dell Networking OS retrieves the access class from the VTY line.

The Dell Networking OS takes the access class from the VTY line and applies it to ALL users. Dell Networking OS does not need to know the identity of the incoming user and can immediately apply the access class. If the authentication method is RADIUS, TACACS+, or line, and you have configured an access class for the VTY line, Dell Networking OS immediately applies it. If the access-class is set to deny all or deny for the incoming subnet, Dell Networking OS closes the connection without displaying the login prompt. The following example shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt. The example uses TACACS+ as the authentication mechanism.


Example of Configuring VTY Authorization Based on Access Class Retrieved from the Line (Per Network Address)

```
Dell(conf)#ip access-list standard deny10
Dell(conf-ext-nacl)#permit 10.0.0.0/8
Dell(conf-ext-nacl)#deny any
Dell(conf)#
Dell(conf)#aaa authentication login tacacsmethod tacacs+
Dell(conf)#tacacs-server host 256.1.1.2 key Force10
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication tacacsmethod
Dell(config-line-vty)#
Dell(config-line-vty)#access-class deny10
Dell(config-line-vty)#end
(same applies for radius and line authentication)
```

Simple Network Management Protocol (SNMP)

Network management stations use SNMP to retrieve or alter management data from network elements. A datum of management information is called a managed object; the value of a managed object can be static or variable. Network elements store managed objects in a database called a management information base (MIB).

MIBs are hierarchically structured and use object identifiers to address managed objects, but managed objects also have a textual name called an *object descriptor*.

 **NOTE:** An I/O Aggregator supports standard and private SNMP MIBs, including Get operations in supported MIBs.


Supported Modes

Standalone, PMUX, VLT, Stacking

Implementation Information

The Dell Networking OS supports SNMP version 1 as defined by RFC 1155, 1157, and 1212, SNMP version 2c as defined by RFC 1901.

Configuring the Simple Network Management Protocol

 **NOTE:** The configurations in this chapter use a UNIX environment with net-snmp version 5.4. This is only one of many RFC-compliant SNMP utilities you can use to manage the Aggregator using SNMP. Also, these configurations use SNMP version 2c.

Configuring SNMP version 1 or version 2 requires only a single step:

1. Create a community.

 **NOTE:** IOA supports only Read-only mode.

Important Points to Remember

- Typically, 5-second timeout and 3-second retry values on an SNMP server are sufficient for both local area network (LAN) and wide area network (WAN) applications. If you experience a timeout with these values, increase the timeout value to greater than 3 seconds, and increase the retry value to greater than 2 seconds on your SNMP server.

Setting up SNMP

Dell Networking OS supports SNMP version 1 and version 2 which are community-based security models. The primary difference between the two versions is that version 2 supports two additional protocol operations (informs operation and **snmpgetbulk** query) and one additional object (counter64 object).

Creating a Community

For SNMPv1 and SNMPv2, create a community to enable the community-based security in the Dell Networking OS.

The management station generates requests to either retrieve or alter the value of a management object and is called the *SNMP manager*. A network element that processes SNMP requests is called an *SNMP agent*. An *SNMP community* is a group of SNMP agents and managers that are allowed to interact. Communities are necessary to secure communication between SNMP managers and agents; SNMP agents do not respond to requests from management stations that are not part of the community.

The Dell Networking OS enables SNMP automatically when you create an SNMP community and displays the following message. You must specify whether members of the community may retrieve values in Read-Only mode. Read-write access is not supported.

```
22:31:23: %RPM1-P:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP
WARM_START.
```

To create an SNMP community:

- Choose a name for the community.
CONFIGURATION mode

```
snmp-server community name ro
```

Example of Creating an SNMP Community

To view your SNMP configuration, use the `show running-config snmp` command from EXEC Privilege mode.

```
Dell(conf)#snmp-server community my-snmp-community ro
22:31:23: %RPM1-P:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP
WARM_START.
Dell#show running-config snmp
!
snmp-server community mycommunity ro
Dell#
```

Reading Managed Object Values

You may only retrieve (read) managed object values if your management station is a member of the same community as the SNMP agent.

Dell Networking supports RFC 4001, *Textual Conventions for Internet Work Addresses* that defines values representing a type of internet address. These values display for `ipAddressTable` objects using the `snmpwalk` command.

There are several UNIX SNMP commands that read data.

- Read the value of a single managed object.

```
snmpget -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```

- Read the value of the managed object directly below the specified object.

```
snmpgetnext -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```
- Read the value of many objects at once.

```
snmpwalk -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```

In the following example, the value “4” displays in the OID before the IP address for IPv4. For an IPv6 IP address, a value of “16” displays.

Example of Reading the Value of a Managed Object

```
> snmpget -v 2c -c mycommunity 10.11.131.161 sysUpTime.0
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32852616) 3 days, 19:15:26.16
> snmpget -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32856932) 3 days, 19:16:09.32
```

Example of Reading the Value of the Next Managed Object

```
> snmpgetnext -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0
SNMPv2-MIB::sysContact.0 = STRING:
> snmpgetnext -v 2c -c mycommunity 10.11.131.161 sysContact.0
SNMPv2-MIB::sysName.0 = STRING:
```

Example of Reading the Value of Many Managed Objects at Once

```
> snmpwalk -v 2c -c mycommunity 10.16.130.148 .1.3.6.1.2.1.1
SNMPv2-MIB::sysDescr.0 = STRING: Dell Networking OS
Operating System Version: 1.0
Application Software Version: E8-3-17-46
Series: I/O-Aggregator
Copyright (c) 1999-2012 by Dell Inc. All Rights Reserved.
Build Time: Sat Jul 28 03:20:24 PDT 2012
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.6027.1.4.2
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (77916) 0:12:59.16
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: FTOS
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysServices.0 = INTEGER: 4
```

Displaying the Ports in a VLAN using SNMP

Dell Networking OS identifies VLAN interfaces using an interface index number that is displayed in the output of the `show interface vlan` command.

Example of Identifying the VLAN Interface Index Number

```
Dell(conf)#do show interface vlan id 10
% Error: No such interface name.
R5(conf)#do show interface vlan 10
Vlan 10 is down, line protocol is down
Address is 00:01:e8:cc:cc:ce, Current address is 00:01:e8:cc:cc:ce
Interface index is 1107787786
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
ARP type: ARPA, ARP Timeout 04:00:00
```


To display the ports in a VLAN, send an `snmpget` request for the object `dot1qStaticEgressPorts` using the interface index as the instance number, as shown in the following example.

Example of Viewing the Ports in a VLAN in SNMP

```
snmpget -v2c -c mycommunity 10.11.131.185 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

The table that the Dell Networking system sends in response to the `snmpget` request is a table that contains hexadecimal (hex) pairs, each pair representing a group of eight ports.

- Seven hex pairs represent a stack unit. Seven pairs accommodate the greatest number of ports available on an Aggregator, 12 ports. The last stack unit is assigned eight pairs, the eighth pair is unused.

The first hex pair, 00 in the previous example, represents ports 1 to 7 in Stack Unit 0. The next pair to the right represents ports 8 to 15. To resolve the hex pair into a representation of the individual ports, convert the hex pair to binary. Consider the first hex pair 00, which resolves to 0000 0000 in binary:

- Each position in the 8-character string is for one port, starting with Port 1 at the left end of the string, and ending with Port 8 at the right end. A 0 indicates that the port is not a member of the VLAN; a 1 indicates VLAN membership.

All hex pairs are 00, indicating that no ports are assigned to VLAN 10. In the following example, Port 0/2 is added to VLAN 10 as untagged; the first hex pair changes from 00 to 04.

Example of Viewing VLAN Ports Using SNMP (Port Assigned)

```
R5(conf)#do show vlan id 10
```

```
Codes: * - Default VLAN, G - GVRP VLANs
Q: U - Untagged, T - Tagged
    x - Dot1x untagged, X - Dot1x tagged
    G - GVRP tagged, M - Vlan-stack
```

```
NUM Status Description Q Ports
10 Inactive          U TenGi 0/2
```


```
[Unix system output]
```

```
> snmpget -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING: 40 00 00 00 00 00 00 00
00 00 00 00
```

The value 40 is in the first set of 7 hex pairs, indicating that these ports are in Stack Unit 0. The hex value 40 is 0100 0000 in binary. As described, the left-most position in the string represents Port 1. The next position from the left represents Port 2 and has a value of 1, indicating that Port 0/2 is in VLAN 10. The remaining positions are 0, so those ports are not in the VLAN.

Fetching Dynamic MAC Entries using SNMP

The Aggregator supports the RFC 1493 dot1d table for the default VLAN and the dot1q table for all other VLANs.

 **NOTE:** The table contains none of the other information provided by the show vlan command, such as port speed or whether the ports are tagged or untagged.


 **NOTE:** The 802.1q Q-BRIDGE MIB defines VLANs regarding 802.1d, as 802.1d itself does not define them. As a switchport must belong a VLAN (the default VLAN or a configured VLAN), all MAC address learned on a switchport are associated with a VLAN. For this reason, the Q-Bridge MIB is used for MAC address query. Moreover, specific to MAC address query, the MAC address indexes dot1dTpFdbTable only for a single forwarding database, while dot1qTpFdbTable has two indices — VLAN ID and MAC address — to allow for multiple forwarding databases and considering that the same MAC address is learned on multiple VLANs. The VLAN ID is added as the first index so that MAC addresses are read by the VLAN, sorted lexicographically. The MAC address is part of the OID instance, so in this case, lexicographic order is according to the most significant octet.

Table 18. MIB Objects for Fetching Dynamic MAC Entries in the Forwarding Database

MIB Object	OID	MIB	Description
dot1dTpFdbTable	.1.3.6.1.2.1.17.4.3	Q-BRIDGE MIB	List the learned unicast MAC addresses on the default VLAN.
dot1qTpFdbTable	.1.3.6.1.2.1.17.7.1.2. 2	Q-BRIDGE MIB	List the learned unicast MAC addresses on non-default VLANs.
dot3aCurAggFdb Table	.1.3.6.1.4.1.6027.3.2. 1.1.5	F10-LINK-AGGREGATION -MIB	List the learned MAC addresses of aggregated links (LAG).

In the following example, R1 has one dynamic MAC address, learned off of port TenGigabitEthernet 0/7, which is a member of the default VLAN, VLAN 1. The SNMP walk returns the values for dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus.

Each object is comprised of an OID concatenated with an instance number. In the case of these objects, the instance number is the decimal equivalent of the MAC address; derive the instance number by converting each hex pair to its decimal equivalent. For example, the decimal equivalent of E8 is 232, and so the instance number for MAC address 00:01:e8:06:95:ac is 0.1.232.6.149.172.

The value of dot1dTpFdbPort is the port number of the port off which the system learns the MAC address. In this case, of TenGigabitEthernet 0/7, the manager returns the integer 118.

Example of Fetching Dynamic MAC Addresses on the Default VLAN

```
-----MAC Addresses on Dell Networking
System-----
Dell#show mac-address-table
VlanId  Mac Address      Type      Interface      State
1       00:01:e8:06:95:ac Dynamic  Tengig 0/7     Active
-----Query from Management Station-----
```

```
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.4.3.1
SNMPv2-SMI::mib-2.17.4.3.1.1.0.1.232.6.149.172 = Hex-STRING: 00 01 E8 06 95 AC
```

Example of Fetching Dynamic MAC Addresses on a Non-default VLANs

In the following example, TenGigabitEthernet 0/7 is moved to VLAN 1000, a non-default VLAN. To fetch the MAC addresses learned on non-default VLANs, use the object dot1qTpFdbTable. The instance number is the VLAN number concatenated with the decimal conversion of the MAC address.

```
-----MAC Addresses on Dell Networking
System-----
Dell#show mac-address-table
VlanId  Mac Address      Type      Interface      State
1000    00:01:e8:06:95:ac  Dynamic  Tengig 0/7     Active
-----Query from Management Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.7.1.2.2.1
```

Example of Fetching MAC Addresses Learned on a Port-Channel

Use dot3aCurAggFdbTable to fetch the learned MAC address of a port-channel. The instance number is the decimal conversion of the MAC address concatenated with the port-channel number.

```
-----MAC Addresses on Dell Networking
System-----
Dell(conf)#do show mac-address-table
VlanId  Mac Address      Type      Interface      State
1000    00:01:e8:06:95:ac  Dynamic  Po 1           Active
-----Query from Management Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.4.1.6027.3.2.1.1.5
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.1.1000.0.1.232.6.149.172.1 = INTEGER: 1000
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.2.1000.0.1.232.6.149.172.1 = Hex-STRING: 00 01 E8 06 95 AC
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.3.1000.0.1.232.6.149.172.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.4.1000.0.1.232.6.149.172.1 = INTEGER: 1
```

Deriving Interface Indices

The Dell Networking OS assigns an interface number to each (configured or unconfigured) physical and logical interface. Display the interface index number using the `show interfacecommand` from EXEC Privilege mode, as shown in the following example.

The interface index is a binary number with bits that indicate the slot number, port number, interface type, and card type of the interface. The Dell Networking OS converts this binary index number to decimal, and displays it in the output of the `show interface` command.

Starting from the least significant bit (LSB):

- the first 14 bits represent the card type
- the next 4 bits represent the interface type
- the next 7 bits represent the port number
- the next 5 bits represent the slot number
- the next 1 bit is 0 for a physical interface and 1 for a logical interface
- the next 1 bit is unused

For example, the index 44634369 is 10101010010001000100000001 in binary. The binary interface index for TenGigabitEthernet 0/4 of an Aggregator. Notice that the physical/logical bit and the final, unused bit

are not given. The interface is physical, so this must be represented by a 0 bit, and the unused bit is always 0. These two bits are not given because they are the most significant bits, and leading zeros are often omitted.


For interface indexing, slot and port numbering begins with binary one. If the Dell Networking system begins slot and port numbering from 0, binary 1 represents slot and port 0. In the Aggregator the first interface is 0/1 and 0/0s lindex is unused and lindex creation logic is not changed. Because Zero is reserved for logical interfaces, it starts from 1. For the first interface, port number is set to 1. Adding it causes an increment by 1 for the next interfaces, so it only starts from 2. Therefore, the port number is set to 4 for 0/3.

Example of Deriving the Interface Index Number

```
Dell#show interface tengig 0/2
TenGigabitEthernet 0/2 is up, line protocol is up
Hardware is Dell Force10Eth, address is 00:01:e8:0d:b7:4e
Current address is 00:01:e8:0d:b7:4e
Interface index is 72925242
[output omitted]
```

Monitor Port-Channels

To check the status of a Layer 2 port-channel, use f10LinkAggMib (1.3.6.1.4.1.6027.3.2). In the following example, Po 1 is a switchport and Po 2 is in Layer 3 mode.

 **NOTE:** The interface index does not change if the interface reloads or fails over. If the unit is renumbered (for any reason) the interface index changes during a reload.

Example of SNMP Trap for Monitored Port-Channels

```
[senthilnathan@lithium ~]$ snmpwalk -v 2c -c public 10.11.1.1 .
1.3.6.1.4.1.6027.3.2.1.1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.2 = INTEGER: 2
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.2.1 = Hex-STRING: 00 01 E8 13 A5 C7
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.2.2 = Hex-STRING: 00 01 E8 13 A5 C8
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.3.1 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.3.2 = INTEGER: 1107755010
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.4.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.4.2 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.5.1 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.5.2 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.6.1 = STRING: "Tengig 0/4 " << Channel
member for Po1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.6.2 = STRING: "Tengig 0/5 " << Channel
member for Po2
dot3aCommonAggFdbIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.1.1107755009.1 = INTEGER: 1107755009
dot3aCommonAggFdbVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.2.1107755009.1 = INTEGER: 1
dot3aCommonAggFdbTagConfig
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.3.1107755009.1 = INTEGER: 2 (Tagged 1
or Untagged 2)
dot3aCommonAggFdbStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.4.1107755009.1 = INTEGER: 1 << Status
active, 2 - status inactive
```

If you learn the MAC address for the LAG, the LAG status also displays.

```

dot3aCurAggVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.1.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggMacAddr
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.2.1.0.0.0.0.0.1.1 = Hex-STRING: 00 00
00 00 00 01
dot3aCurAggIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.3.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.4.1.0.0.0.0.0.1.1 = INTEGER: 1 <<
Status active, 2 - status
inactive

```

For L3 LAG, you do not have this support.

```

SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.33865785 = INTEGER: 33865785
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed
interface state to down: Tengig 0/1"
2010-02-10 14:22:39 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed
interface state to down: Po 1"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500932) 23:36:49.32 SNMPv2-
MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.33865785 = INTEGER: 33865785 SNMPv2-
SMI::enterprises.6027.3.1.1.4.1.2 =
STRING: "OSTATE_UP: Changed interface state to up: Tengig 0/1"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500934) 23:36:49.34 SNMPv2-
MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_UP: Changed
interface state to up: Po 1"

```

Entity MIBS

The Entity MIB provides a mechanism for presenting hierarchies of physical entities using SNMP tables. The Entity MIB contains the following groups, which describe the physical elements and logical elements of a managed system. The following tables are implemented for the Aggregator.

- **Physical Entity:** A physical entity or physical component represents an identifiable physical resource within a managed system. Zero or more logical entities may utilize a physical resource at any given time. Determining which physical components are represented by an agent in the EntPhysicalTable is an implementation-specific matter. Typically, physical resources (for example, communications ports, backplanes, sensors, daughter-cards, power supplies, and the overall chassis), which you can manage via functions associated with one or more logical entities, are included in the MIB.
- **Containment Tree:** Each physical component may be modeled as contained within another physical component. A containment-tree is the conceptual sequence of *entPhysicalIndex* values that uniquely specifies the exact physical location of a physical component within the managed system. It is generated by following and recording each *entPhysicalContainedIn* instance up the tree towards the root, until a value of zero indicating no further containment is found.

Example of Sample Entity MIBS outputs

```
Dell#show inventory optional-module
```

```
Unit Slot Expected Inserted Next Boot Status/Power(On/Off)
```

```
-----  
1 0 SFP+ SFP+ AUTO Good/On  
1 1 QSFP+ QSFP+ AUTO Good/On  
* - Mismatch
```

```
Dell#
```

```
The status of the MIBS is as follows:
```

```
$ snmpwalk -c public -v 2c 10.16.150.162 .1.3.6.1.2.1.47.1.1.1.1.2  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.1 = ""  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.2 = STRING: "PowerEdge-FN-410S-IOA"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.3 = STRING: "Chassis 0 container"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.4 = STRING: "Module 0"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.5 = STRING: "Unit: 0 Port 1 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.6 = STRING: "Unit: 0 Port 2 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.7 = STRING: "Unit: 0 Port 3 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.8 = STRING: "Unit: 0 Port 4 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.9 = STRING: "Unit: 0 Port 5 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.10 = STRING: "Unit: 0 Port 6 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.11 = STRING: "Unit: 0 Port 7 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.12 = STRING: "Unit: 0 Port 8 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.13 = STRING: "Unit: 0 Port 9 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.14 = STRING: "Unit: 0 Port 10 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.15 = STRING: "Unit: 0 Port 11 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.16 = STRING: "Unit: 0 Port 12 10G Level"  
shathishmuthu@login-maa-107 ~ $
```

SNMP Traps for Link Status

To enable SNMP traps for link status changes, use the `snmp-server enable traps snmp linkdown linkup` command.

Standard VLAN MIB

When the Aggregator is in Standalone mode, where all the 4000 VLANs are part of all server side interfaces as well as the single uplink LAG, it takes a long time (30 seconds or more) for external management entities to discover the entire VLAN membership table of all the ports. Support for current status OID in the standard VLAN MIB is expected to simplify and speed up this process. This OID provides 4000 VLAN entries with port membership bit map for each VLAN and reduces the scan for (4000 X Number of ports) to 4000.

Enhancements

1. The `dot1qVlanCurrentEgressPorts` MIB attribute has been enhanced to support logical LAG interfaces.
2. Current status OID in standard VLAN MIB is accessible over SNMP.
3. The bitmap supports 42 bytes for physical ports and 16 bytes for the LAG interfaces (up to a maximum of 128 LAG interfaces).
4. A 59 byte buffer bitmap is supported and in that bitmap:
 - First 42 bytes represent the physical ports.
 - Next 16 bytes represent logical ports 1-128.
 - An additional 1 byte is reserved for future.

Fetching the Switchport Configuration and the Logical Interface Configuration

Important Points to Remember

- The SNMP should be configured in the chassis and the chassis management interface should be up with the IP address.
- If a port is configured in a VLAN, the respective bit for that port will be set to 1 in the specific VLAN.
- In the aggregator, all the server ports and uplink LAG 128 will be in switchport. Hence, the respective bits are set to 1.

The following output is for the default VLAN.

Example of dot1qVlanCurrentUntaggedPorts output

```
snmpwalk -Os -c public -v 1 10.16.151.151 1.3.6.1.2.1.17.7.1.4.2.1.5
mib-2.17.7.1.4.2.1.5.0.1107525633 = Hex-STRING: FF FF FF FF 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 01 00
```

The last byte is free byte . The bit for LAGs starts from 43 byte. If server LAG 1 is created with server ports Te 0/6 and Te 0/7, the respective bit for the ports are unset and the bit for LAG 1 is set in default VLAN. The corresponding output will be as follows:

```
snmpwalk -Os -c public -v 1 10.16.151.151 1.3.6.1.2.1.17.7.1.4.2.1.5
mib-2.17.7.1.4.2.1.5.0.1107525633 = Hex-STRING: F9 FF FF FF 00 00 00 00 00 00
00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 80 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 01 00
```

In the above example, the 43rd byte is set to 80. The 43rd byte is for LAG IDs from 1 to 8. But, only one LAG po 1 is set as switch port. Hence, the binary bits will be 10000000. If this converted to Hexadecimal, the value will be 80. Similarly, the first byte for Te 0/1 to Te 0/8 server ports, as the 6th and 7th byte are removed from switch port, the respective bits are set to 0. In binary, the value is 11111001 and the corresponding hex decimal value is F9.

In standalone mode, there are 4000 VLANs, by default. The SNMP output will display for all 4000 VLANs. To view a particular VLAN, issue the `snmp query` with VLAN interface ID.

```
Dell#show interface vlan 1010 | grep "Interface index"
      Interface index is 1107526642
```

Use the output of the above command in the snmp query.

```
snmpwalk -Os -c public -v 1 10.16.151.151
1.3.6.1.2.1.17.7.1.4.2.1.4.0.1107526642
mib-2.17.7.1.4.2.1.4.0.1107526642 = Hex-STRING: F9 FF FF FF 00 00 00 00 00 00
00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 80 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 01 00
```

MIB Support to Display the Available Memory Size on Flash

Dell Networking provides more MIB objects to display the available memory size on flash memory. The following table lists the MIB object that contains the available memory size on flash memory.

Table 19. MIB Objects for Displaying the Available Memory Size on Flash via SNMP

MIB Object	OID	Description
chStackUnitFlashUsageUtil	1.3.6.1.4.1.6027.3.19.1.2.8.1.6	Contains flash memory usage in percentage.

The chStackUnitUtilTable MIB table contains the chStackUnitFlashUsageUtil MIB object which contains the flash memory usage percent.

Viewing the Available Flash Memory Size

- To view the available flash memory using SNMP, use the following command.

```
snmpget -v2c -c public 192.168.60.120 .1.3.6.1.4.1.6027.3.10.1.2.9.1.6.1  
enterprises.6027.3.10.1.2.9.1.5.1 = Gauge32: 24
```

The output above displays that 24% of the flash memory is used.

MIB Support to Display the Software Core Files Generated by the System

Dell Networking provides MIB objects to display the software core files generated by the system. The chSysSwCoresTable contains the list of software core files generated by the system. The following table lists the related MIB objects.

Table 20. MIB Objects for Displaying the Software Core Files Generated by the System

MIB Object	OID	Description
chSysSwCoresTable	1.3.6.1.4.1.6027.3.19.1.2.9	This is the table that contains the list of software core files generated by the system.
chSysCoresEntry	1.3.6.1.4.1.6027.3.19.1.2.9.1	Entry number.
chSysCoresInstance	1.3.6.1.4.1.6027.3.19.1.2.9.1.1	Stores the indexed information about the available software core files.
chSysCoresFileName	1.3.6.1.4.1.6027.3.19.1.2.9.1.2	Contains the core file names and the file paths.
chSysCoresTimeCreated	1.3.6.1.4.1.6027.3.19.1.2.9.1.3	Contains the time at which core files are created.

MIB Object	OID	Description
chSysCoresStackUnitNumber	1.3.6.1.4.1.6027.3.19.1.2.9.1.4	Contains information that includes which stack unit or processor the core file was originated from.
chSysCoresProcess	1.3.6.1.4.1.6027.3.19.1.2.9.1.5	Contains information that includes the process names that generated each core file.

Viewing the Software Core Files Generated by the System

- To view the viewing the software core files generated by the system, use the following command.

```
snmpwalk -v2c -c public 192.168.60.120 .1.3.6.1.4.1.6027.3.10.1.2.10

enterprises.6027.3.10.1.2.10.1.1.1.1 = 1
enterprises.6027.3.10.1.2.10.1.1.1.2 = 2
enterprises.6027.3.10.1.2.10.1.1.1.3 = 3
enterprises.6027.3.10.1.2.10.1.1.2.1 = 1
enterprises.6027.3.10.1.2.10.1.2.1.1 = "/CORE_DUMP_DIR/flashmnr.core.gz"
enterprises.6027.3.10.1.2.10.1.2.1.2 = "/CORE_DUMP_DIR/FTP_STK_MEMBER/
f10cp_l2mgr_131108080758_Stk1.acore.gz"
enterprises.6027.3.10.1.2.10.1.2.1.3 = "/CORE_DUMP_DIR/FTP_STK_MEMBER/
f10cp_vrrp_140522124357_Stk1.acore.gz"
enterprises.6027.3.10.1.2.10.1.2.2.1 =
"/CORE_DUMP_DIR/FTP_STK_MEMBER/f10cp_sysd_140617134445_Stk0.acore.gz"
enterprises.6027.3.10.1.2.10.1.3.1.1 = "Fri Mar 14 11:51:46 2014"
enterprises.6027.3.10.1.2.10.1.3.1.2 = "Fri Nov 8 08:11:16 2013"
enterprises.6027.3.10.1.2.10.1.3.1.3 = "Fri May 23 05:05:16 2014"
enterprises.6027.3.10.1.2.10.1.3.2.1 = "Tue Jun 17 14:19:26 2014"
enterprises.6027.3.10.1.2.10.1.4.1.1 = 0
enterprises.6027.3.10.1.2.10.1.4.1.2 = 1
enterprises.6027.3.10.1.2.10.1.4.1.3 = 1
enterprises.6027.3.10.1.2.10.1.4.2.1 = 0
enterprises.6027.3.10.1.2.10.1.5.1.1 = "flashmnr"
enterprises.6027.3.10.1.2.10.1.5.1.2 = "l2mgr"
enterprises.6027.3.10.1.2.10.1.5.1.3 = "vrrp" Hex: 76 72 72 70
enterprises.6027.3.10.1.2.10.1.5.2.1 = "sysd" Hex: 73 79 73 64
```

The output above displays that the software core files generated by the system.

Stacking

An Aggregator auto-configures to operate in standalone mode. To use an Aggregator in a stack, you must manually configure it using the CLI to operate in stacking mode.

Stacking is supported on the FN410S and FN410T Aggregators with ports 9 and 10 as the stack ports. The Aggregator supports both ring and daisy-chain topology and stacking of the same type. FN 410S and FN 410T Aggregators support two-unit in-chassis stacking and up to six units stacking across the chassis.

In Stack mode, the lower two external Ethernet ports (ports 9 and 10) operate as stack links. In Programmable MUX (PMUX) mode, you can configure any of the external Ethernet ports to operate as stack links.

Stacking provides a single point of management for high availability and higher throughput. To configure a stack, you must use the CLI.

Supported Modes

Stacking, PMUX

Configuring a Switch Stack

To configure and bring up a switch stack, follow these steps:

1. Connect the ports on the base module of two Aggregators using 10G direct attach or SFP+ fibre cables.
2. Configure each Aggregator to operate in stacking mode.
3. Reload each Aggregator, one after the other in quick succession.

Stacking Prerequisites

Before you cable and configure a stack of the Aggregators, review the following prerequisites.


- All Aggregators in the stack must be powered up with the initial or startup configuration before you attach the cables.
- All stacked Aggregators must run the same Dell Networking OS version. To check the version that a switch is running, use the `show version` command. To download a Dell Networking OS version, go to <http://support.dell.com>.
- Stacking is supported only with other Aggregators. A maximum of six Aggregators are supported in a single stack.
- A maximum of four stack groups is supported on a stacked Aggregator.
- Interconnect the stack units by following the instructions in [Cabling the Switch Stack](#).
- You cannot stack a Standalone IOA and a PMUX.

Master Selection Criteria

A Master is elected or re-elected based on the following considerations, in order:

1. The switch with the highest priority at boot time.
2. The switch with the highest MAC address at boot time.
3. A unit is selected as Standby by the administrator, and a fail over action is manually initiated or occurs due to a Master unit failure.

No record of previous stack mastership is kept when a stack loses power. As it reboots, the election process will once again determine the Master and Standby switches. As long as the priority has not changed on any members, the stack will retain the same Master and Standby.

 **NOTE:** Each stack members' role (including the Master and Standby) can be defined by the user at any time by setting the priority.

If the entire stack is powered OFF and ON again, the unit that was the Master before the reboot will remain the Master after the stack resumes operation. However, when a stack is powered on, all members are in **sleep** mode for 5 seconds while waiting on the previous Master to join the stack. If the previous Master fails to join within 5 seconds, the remaining members (including the Standby) elect a new Master.

Configuring Priority and stack-group

Perform the following steps to configure the priorities and stack-groups for each of the switches.


1. Set the priorities for the stack-unit.

CONFIGURATION mode

```
stack-unit unit-number priority 1-14
```

```
Dell(conf)# stack-unit 0 priority 12
```

Setting the priority will determine which switch will become the management (Master) switch. The switch with the highest priority number is elected Master. The default priority is 0.

 **NOTE:** It is best practice to assign priority values to all switches before stacking them in order to acquire and retain complete control over each units role in the stack.

2. Configure the stack-group for each stack-unit.

CONFIGURATION mode

```
stack-unit unit-no stack-group stack-group-id
```

```
Dell(conf)# stack-unit 0 stack-group 0  
Dell (conf)#02:26:05: %STKUNIT0-M:CP %IFMGR-6-STACK_PORTS_ADDED:  
Ports Fo 0/33 have been configured as stacking ports.  
Please save and reload for config to take effect
```

Stack-groups are easier to think of simply as stack ports. For example, using the `stack-group 0` command simply turns the lower port (port 9) into a stacking port. Similarly, `stack-group 1`, `stack-group 2` and `stack-group 3` commands correspond to ports 10, 11 and 12 respectively.

 **NOTE:** Stack-group is supported only in PMUX mode.

3. Continue to run the `stack-unit 0 stack-group <0-3>` command to add additional stack ports to the switch, using the stack-group mapping.


Cabling the Switch Stack

Dell PowerEdge FN I/O Aggregators are connected to operate as a single stack in a ring topology using the SFP+ or Base-T ports on the front end ports 9 and 10. To create a stack in either a ring or daisy-chain topology, you can use two units on the same chassis or up to six units across multiple chassis.

Prerequisite: Before you attach the stacking cables, all Aggregators in the stack must be powered up with the default or reconfigured settings.

To connect stacking ports, use only SFP+ transceivers, SFP+ cables and BaseT and its supported cables (separately purchased). For example:

1. Insert a cable in port 9 on the first aggregator.
2. Connect the cable to port 10 on the next aggregator.
3. Continue this pattern on up to 6 aggregators.
4. Connect a cable from port 9 on the last aggregator to port 10 on the first aggregator. This creates a ring topology.

 **NOTE:** The resulting topology allows the stack to function as a single switch with resilient failover capabilities.

Accessing the CLI

To configure a stack, you must access the stack master in one of the following ways.

- For remote out-of-band management (OOB), enter the OOB management interface IP address into a Telnet or secure shell (SSH) client and log in to the switch using the user ID and password to access the CLI.
- For local management, use the attached console connection to the master switch to log in to the CLI. Console access to the stack CLI is available on the master only.
- For remote in-band management from a network management station, enter the virtual local area network (VLAN) IP address of the management port and log in to the switch to access the CLI.

Configuring and Bringing Up a Stack

After you attach the cables in a stack of Aggregators, perform the following steps to configure and bring up the stack:

1. Set up a connection to the CLI on an Aggregator as described in [Accessing the CLI](#).
2. Connect the terminal to the console port on an Aggregator. Enter the following commands to access the CLI and configure the base module ports for stacking mode:

```
Login: username
Password: *****
Dell> enable
Dell# configure
Dell(conf)# stack-unit 0 iom-mode stack
```

Where `stack-unit 0` defines the default stack-unit number in the initial configuration of a switch.


3. Save the stacking configuration on the SFP+ or 10G Base-T ports:
Dell# write memory

4. Reboot the Aggregator by entering the `reload` command in EXEC Privilege mode:

```
Dell# reload
```

Repeat the above steps on each Aggregator in the stack by entering the `stackunit 0 iom-mode stack` command and saving the configuration.

If the stacked switches all reboot at approximately the same time, the Aggregator with the highest MAC address is automatically elected as the master switch. The Aggregator with the next highest MAC address is elected as the standby master.

 **NOTE:** You can ensure that a stacked switch becomes the master by rebooting the switch first, and waiting for it to come up before rebooting the second switch in the stack. It is recommended to have the boot image for the stack units via boot from flash.

The Status LED on the Aggregator that is elected as the master switch is blue. Perform VLAN and other software configuration for the stack by connecting to the console port on the master unit. To determine which switch is the stack master, enter the `show system` command at the terminal.

To remove an Aggregator from the stack, enter the `stack-unit 0 iom-mode standalone` command in Global configuration mode, save the configuration and reload the Aggregator for the change to take effect.

In stack mode, all VLAN membership are removed and the port is assigned only to the default VLAN1. You must configure additional VLAN membership, as required.

Adding a Stack Unit

You can add a new unit to an existing stack both when the unit has no stacking ports (stack groups) configured and when the unit already has stacking ports configured. If the units to be added to the stack have been previously used, they are assigned the smallest available unit ID in the stack.

To add a standalone Aggregator to a stack, follow these steps:

1. Power on the switch.
2. Attach SFP+ or direct attach cables to connect 10G ports on the switch to one or more switches in the stack.
3. Log on to the CLI and enter global configuration mode.
Login: username

Password: *****

Dell> enable

Dell# configure
4. Configure the Aggregator to operate in stacking mode.
CONFIGURATION mode

`stack-unit 0 iom-mode`

`stack`
5. Reload the switch. Dell Operating System automatically assigns a number to the new unit and adds it as member switch in the stack. The new unit synchronizes its running and startup configurations with the stack.

```
EXEC Privilege mode
```

```
reload
```

If an Aggregator is already configured to operate in stacking mode, simply attach SFP+ or direct attach cables to connect 10G ports on the base module of each stacked Aggregator. The new unit synchronizes its running and startup configurations with the stack.

Dell Networking OS Behavior: When you add a new Aggregator to a stack:

- If the new unit has been configured with a stack number that is already assigned to a stack member, the stack avoids a numbering conflict by assigning the new switch the first available stack number.
- If the stack has been provisioned for the stack number that is assigned to the new unit, the pre-configured provisioning must match the switch type. If there is a conflict between the provisioned switch type and the new unit, a mismatch error message is displayed.

Resetting a Unit on a Stack

Use the following `reset` command to reload any of the member units or the standby in a stack. If you try to reset the stack master, the following error message is displayed:

```
% Error: Reset of master unit is not allowed.
```

To reset a unit on a stack, use the following command:

- Reset any designated stack member, except the management unit.

```
EXEC Privilege mode
```

```
reset stack-unit unit-number {hard}
```

- Hard reset any stack unit including master unit.

```
EXEC Privilege mode
```

```
power-cycle stack-unit [unit-number]
```

Removing an Aggregator from a Stack

To remove an Aggregator from a stack, follow the below steps:

1. Disconnect the stacking cables from the unit. The unit can be powered on or off and can be online or offline.
2. Log on to the CLI and enter Global Configuration mode.

```
Login: username
```

```
Password: *****
```

```
Dell> enable
```

```
Dell# configure
```

3. Configure the Aggregator to operate in standalone mode.

```
CONFIGURATION
```

```
stack-unit 0 iom-mode standalone
```

4. Log on to the CLI and reboot each switch, one after another, in as short a time as possible.

```
EXEC PRIVILEGE
```

```
reload
```


The switch functions in standalone mode but retains the running and startup configuration that was last synchronized by the master switch while it operated as a stack unit.

Merging Two Operational Stacks

The recommended procedure for merging two operational stacks is as follows:

1. Always power off all units in one stack before connecting to another stack.
2. Add the units as a group by unplugging one stacking cable in the operational stack and physically connecting all unpowered units.
3. Completely cable the stacking connections, making sure the redundant link is also in place.

Two operational stacks can also be merged by reconnecting stack cables without powering down units in either stack. Connecting a powered-up standalone unit to an existing stack leads to same behavior as when merging two operational stacks. In such cases, Manager re-election is done and the Manager with the higher MAC address wins the election. The losing stack manager resets itself and all its member units. After the reset, all the losing stack members join the winning stack to form a single stack. The winning stack remains functional through the merge process. If the stack merge is performed in this way, then it is strongly recommended that the user set the admin management preference of the desired winner stack manager to a higher value than the stack manager that should lose the election.

 **NOTE:** In case of a stack, when one unit member resets and joins the stack, momentary drops will be observed, which is recovered after few seconds.

Verifying a Stack Configuration

The following lists the status of a stacked switch according to the color of the System Status light emitting diodes (LEDs) on its front panel.

- Blue indicates the switch is operating as the stack master or as a standalone unit.
- Off indicates the switch is a member or standby unit.
- Amber indicates the switch is booting or a failure condition has occurred.

Using Show Commands

To display information on the stack configuration, use the `show` commands on the master switch.

- Displays stacking roles (master, standby, and member units) and the stack MAC address.
`show system [brief]`
- Displays the stack groups allocated on a stacked switch. The range is from 0 to 5.
`show system stack-unit unit-number stack-group configured`
- Displays the port numbers that correspond to the stack groups on a switch. The valid stack-unit numbers are from 0 to 5.
`show system stack-unit unit-number stack-group`
- Displays the type of stack topology (ring or daisy chain) with a list of all stacked ports, port status, link speed, and peer stack-unit connection.
`show system stack-ports [status | topology]`

Troubleshooting a Switch Stack

To perform troubleshooting operations on a switch stack, use the following commands on the master switch.

1. Displays the status of stacked ports on stack units.
`show system stack-ports`
2. Displays the master standby unit status, failover configuration, and result of the last master-standby synchronization; allows you to verify the readiness for a stack failover.
`show redundancy`
3. Displays input and output flow statistics on a stacked port.
`show hardware stack-unit unit-number stack-port port-number`
4. Clears statistics on the specified stack unit. The valid stack-unit numbers are from 0 to 5.
`clear hardware stack-unit unit-number counters`
5. Displays the current operational mode of the Aggregator (standalone or stacking) and the mode in which the Aggregator will operate at the next reload.
`show system stack-unit unit-number iom-mode`

Failure Scenarios

The following sections describe some of the common fault conditions that can happen in a switch stack and how they are resolved.

Stack Member Fails

- **Problem:** A unit that is not the stack master fails in an operational stack.
- **Resolution:** If a stack member fails in a daisy chain topology, a split stack occurs. If a member unit fails in a ring topology, traffic is re-routed over existing stack links.

The following syslog messages are generated when a member unit fails:

```
Dell#May 31 01:46:17: %STKUNIT3-M:CP %IPC-2-STATUS: target stack unit 4 not responding
```

```
May 31 01:46:17: %STKUNIT3-M:CP %CHMGR-2-STACKUNIT_DOWN: Major alarm: Stack unit 4 down - IPC timeout
```

```
Dell#May 31 01:46:17: %STKUNIT3-M:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/49,53
```

```
Dell#May 31 01:46:18: %STKUNIT5-S:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/49,53
```

Unplugged Stacking Cable

- **Problem:** A stacking cable is unplugged from a member switch. The stack loses half of its bandwidth from the disconnected switch.
- **Resolution:** Intra-stack traffic is re-routed on another link using the redundant stacking port on the switch. A recalculation of control plane and data plane connections is performed.

Master Switch Fails

- **Problem:** The master switch fails due to a hardware fault, software crash, or power loss.
- **Resolution:** A failover procedure begins:
 1. Keep-alive messages from the Aggregator master switch time out after 60 seconds and the switch is removed from the stack.
 2. The standby switch takes the master role. Data traffic on the new master switch is uninterrupted. Protocol traffic is managed by the control plane.
 3. A member switch is elected as the new standby. Data traffic on the new standby is uninterrupted. The control plane prepares for operation in Warm Standby mode.

Stack-Link Flapping Error

Problem/Resolution: Stacked Aggregators monitor their own stack ports and disable any stack port that flaps five times within 10 seconds. If the stacking ports that flap are on the master or standby, KERN-2-INT error messages note the units.

To re-enable a downed stacking port, power cycle the stacked switch on which the port is installed.

The following is an example of the stack-link flapping error message.

```
-----MANAGEMENT
UNIT-----
Error: Stack Port 49 has flapped 5 times within 10 seconds. Shutting down this
stack port now.
Error: Please check the stack cable/module and power-cycle the stack.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times
within 10
seconds. Shutting down this stack port now.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/
module and
power-cycle the stack.
-----STANDBY
UNIT-----
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times
within 10
seconds. Shutting down this stack port now.
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/module
and power-cycle the stack.
```

Master Switch Recovers from Failure

- **Problem:** The master switch recovers from a failure after a reboot and rejoins the stack as the standby unit or member unit. Protocol and control plane recovery requires time before the switch is fully online.
- **Resolution:** When the entire stack is reloaded, the recovered master switch becomes the master unit of the stack.

Stack Unit in Card-Problem State Due to Incorrect Dell Networking OS Version

- **Problem:** A stack unit enters a Card-Problem state because the switch has a different Dell Networking OS version than the master unit. The switch does not come online as a stack unit.
- **Resolution:** To restore a stack unit with an incorrect Dell Networking OS version as a member unit, disconnect the stacking cables on the switch and install the correct Dell Networking OS version. Then

add the switch to the stack as described in [Adding a Stack Unit](#). To verify that the problem has been resolved and the stacked switch is back online, use the `show system brief` command.

```
Dell#show system brief
Stack MAC : 00:1e:c9:f1:00:9b
-- Stack Info --
Unit UnitType  Status          ReqTyp          CurTyp          Version        Ports
-----
0  Management online          PE-FN-410S-IOA PE-FN-410S-IOA 1-0(0-1864) 12
1  Standby   card problem PE-FN-410S-IOA unknown          12
2  Member   not present
3  Member   not present
4  Member   not present
5  Member   not present
```

Card Problem — Resolved

```
Dell#show system brief
Stack MAC : 00:1e:c9:f1:04:82
-- Stack Info --
Unit UnitType  Status          ReqTyp          CurTyp          Version        Ports
-----
0  Management online          PE-FN-410S-IOA PE-FN-410S-IOA 1-0(0-1864) 12
1  Standby   online PE-FN-410S-IOA PE-FN-410S-IOA 1-0(0-1864) 12
2  Member   not present
3  Member   not present
4  Member   not present
5  Member   not present
```

Stack Unit in Card-Problem State Due to Configuration Mismatch

- **Problem:** A stack unit enters a Card-Problem state because there is a configuration mismatch between the logical provisioning stored for the stack-unit number on the master switch and the newly added unit with the same number.
- **Resolution:** From the master switch, reload the stack by entering the `reload` command in EXEC Privilege mode. When the stack comes up, the card problem will be resolved.

Upgrading a Switch Stack

To upgrade all switches in a stack with the same Dell Networking OS version, follow these steps.

1. Copy the new Dell Networking OS image to a network server.
2. Download the Dell Networking OS image by accessing an interactive CLI that requests the server IP address and image filename, and prompts you to upgrade all member stack units.

EXEC Privilege mode

```
upgrade system { flash: | ftp: | scp: | tftp: | usbflash: } partition
```

Specify the system partition on the master switch into which you want to copy the Dell Networking OS image. The system then prompts you to upgrade all member units with the new Dell Networking OS version.

The valid values are a: and b:.

3. Reboot all stack units to load the Dell Networking OS image from the same partition on all switches in the stack.

CONFIGURATION mode

```
boot system stack-unit all primary system partition
```

4. Save the configuration.

```
EXEC Privilege
```

```
write memory
```

5. Reload the stack unit to activate the new Dell Networking OS version.

```
CONFIGURATION mode
```

```
reload
```

Example of Upgrading all Stacked Switches

The following example shows how to upgrade all switches in a stack, including the master switch.

```
Dell# upgrade system ftp: A:
Address or name of remote host []: 10.11.200.241
Source file name []: //FTOS-XL-8.3.17.0.bin
User name to login remote host: ftp
Password to login remote host:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Erasing IOM Primary Image, please wait
.!.....
...
.....Writing.....
...
...
...
...
31972272 bytes successfully copied
System image upgrade completed successfully.
Upgrade system image for all stack-units [yes/no]: yes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
Image upgraded to all
Dell# configure
Dell(conf)# boot system stack-unit all primary system: A:
Dell(conf)# end
Dell# write memory
Jan 3 14:01:48: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config to
startup-config
in flash by default
Synchronizing data to peer Stack-unit
!!!!
Dell# reload
Proceed with reload [confirm yes/no]: yes
```

Upgrading a Single Stack Unit

Upgrading a single stacked switch is necessary when the unit was disabled due to an incorrect Dell Networking OS version. This procedure upgrades the image in the boot partition of the member unit from the corresponding partition in the master unit.

To upgrade an individual stack unit with a new Dell Networking OS version, follow the below steps:

1. Download the Dell Networking OS image from the master's boot partition to the member unit, and upgrade the relevant boot partition in the single stack-member unit.
EXEC Privilege mode

```
upgrade system stack-unit unit-number partition
```

2. Reboot the stack unit from the master switch to load the Dell Networking OS image from the same partition.

```
CONFIGURATION mode
```

```
boot system stack-unit unit-number primary system partition
```

3. Save the configuration.

```
EXEC Privilege mode
```

```
write memory
```

4. Reset the stack unit to activate the new Dell Networking OS version.

```
EXEC Privilege mode
```

```
power-cycle stack-unit unit-number
```

Example of Upgrading a Single Stack Unit

The following example shows how to upgrade an individual stack unit.

```
Dell# upgrade system stack-unit 1 A:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Image upgraded to Stack unit 1
Dell# configure
Dell(conf)# boot system stack-unit 1 primary system: A:
Dell(conf)# end
Dell#Jan 3 14:27:00: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from console
Dell# write memory
Jan 3 14:27:10: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config to
startup-config in flash
by default
Synchronizing data to peer Stack-unit
!!!!
....
Dell# power-cycle stack-unit 1
Proceed with power-cycle? Confirm [yes/no]:yes
```

Broadcast Storm Control

On the Aggregator, the broadcast storm control feature is enabled by default on all ports, and disabled on a port when an iSCSI storage device is detected. Broadcast storm control is re-enabled as soon as the connection with an iSCSI device ends.

Broadcast traffic on Layer 2 interfaces is limited or suppressed during a broadcast storm. You can view the status of a broadcast-storm control operation by using the `show io-aggregator broadcast storm-control status` command. You can disable broadcast storm control by using the `no io-aggregator broadcast storm-control` command.

Dell Networking OS Behavior: If broadcast traffic exceeds 1000 Mbps, the Aggregator limits it to 1000 Mbps per port-pipe.

Supported Modes

Standalone, PMUX, VLT, Stacking

Disabling Broadcast Storm Control

To disable broadcast storm control on an Aggregator, use the `no io-aggregator broadcast storm-control` command from CONFIGURATION mode.

To re-enable broadcast storm control, enter the `io-aggregator broadcast storm-control` command.

Displaying Broadcast-Storm Control Status

To display the status of a current storm control operation, use the `show io-aggregator broadcast storm-control status` command from EXEC Privilege mode.

Configuring Storm Control

The following configurations are available only in PMUX mode.

1. To configure the percentage of broadcast traffic allowed on an interface, use the `storm-control broadcast [packets_per_second in]` command from INTERFACE mode.
2. To configure the percentage of multicast traffic allowed on an interface, use the `storm-control multicast [packets_per_second in]` command from INTERFACE mode.
3. To configure the percentage of unknown-unicast traffic allowed on an interface, use the `storm-control unknown-unicast [packets_per_second in]` command from INTERFACE mode.

System Time and Date

The Aggregator auto-configures the hardware and software clocks with the current time and date. If necessary, you can manually set and maintain the system time and date using the CLI commands described in this chapter.

- Setting the Time for the Software Clock
- Setting the Time Zone
- Setting Daylight Savings Time

Supported Modes

Standalone, PMUX, VLT, Stacking

Setting the Time for the Software Clock

You can change the order of the `month` and `day` parameters to enter the time and date as *time day month year*. You cannot delete the software clock.

The software clock runs only when the software is up. The clock restarts, based on the hardware clock, when the switch reboots.

To set the software clock, use the following command.

- Set the system software clock to the current time and date.
EXEC Privilege mode

```
clock set time month day year
```

- *time*: Enter the time in hours:minutes:seconds. For the hour variable, use the 24-hour format; for example, 17:15:00 is 5:15 pm.
- *month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *day*: Enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *year*: Enter a four-digit number as the year. The range is from 1993 to 2035.

Example of the `clock set` Command

```
Dell#clock set 12:11:00 21 may 2012  
Dell#
```

Setting the Timezone

Universal time coordinated (UTC) is the time standard based on the International Atomic Time standard, commonly known as Greenwich Mean time. When determining system time, you must include the

differentiator between the UTC and your local timezone. For example, San Jose, CA is the Pacific Timezone with a UTC offset of -8.

To set the clock timezone, use the following command.

- Set the clock to the appropriate timezone.

CONFIGURATION mode

```
clock timezone timezone-name offset
```

– *timezone-name*: Enter the name of the timezone. Do not use spaces.

– *offset*: Enter one of the following:

* a number from 1 to 23 as the number of hours in addition to UTC for the timezone.

* a minus sign (-) then a number from 1 to 23 as the number of hours.

Example of the `clock timezone` Command

```
Dell#conf
Dell(conf)#clock timezone Pacific -8
Dell#
```

Setting Daylight Savings Time

Dell Networking OS supports setting the system to daylight savings time once or on a recurring basis every year.

Setting Daylight Saving Time Once

Set a date (and time zone) on which to convert the switch to daylight saving time on a one-time basis. To set the clock for daylight savings time once, use the following command.

- Set the clock to the appropriate timezone and daylight saving time.

CONFIGURATION mode

```
clock summer-time time-zone date start-month start-day start-year start-time
end-month end-day end-year end-time [offset]
```

– *time-zone*: Enter the three-letter name for the time zone. This name displays in the show clock output.

– *start-month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.

– *start-day*: enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.

– *start-year*: enter a four-digit number as the year. The range is from 1993 to 2035.

– *start-time*: enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.

– *end-month*: enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.

– *end-day*: enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.

– *end-year*: enter a four-digit number as the year. The range is from 1993 to 2035.

– *end-time*: enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.

- *offset*: (OPTIONAL) enter the number of minutes to add during the summer-time period. The range is from 1 to 1440. The default is **60 minutes**.

Example of the `clock summer-time` Command

```
Dell(conf)#clock summer-time pacific date Mar 14 2012 00:00 Nov 7 2012 00:00
Dell(conf)#
```

Setting Recurring Daylight Saving Time

Set a date (and time zone) on which to convert the switch to daylight saving time on a specific day every year.

If you have already set daylight saving for a one-time setting, you can set that date and time as the recurring setting with the `clock summer-time time-zone recurring` command.

To set a recurring daylight saving time, use the following command.

- Set the clock to the appropriate timezone and adjust to daylight saving time every year.
CONFIGURATION mode


```
clock summer-time time-zone recurring start-week start-day start-month start-
time end-week end-day end-month end-time [offset]
```

- *time-zone*: Enter the three-letter name for the time zone. This name displays in the show clock output.
- *start-week*: (OPTIONAL) Enter one of the following as the week that daylight saving begins and then enter values for *start-day* through *end-time*:
 - * *week-number*: Enter a number from 1 to 4 as the number of the week in the month to start daylight saving time.
 - * *first*: Enter the keyword *first* to start daylight savings time in the first week of the month.
 - * *last*: Enter the keyword *last* to start daylight saving time in the last week of the month.
- *start-month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *start-day*: Enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *start-year*: Enter a four-digit number as the year. The range is from 1993 to 2035.
- *start-time*: Enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.
- *end-week*: If you entered a start-week, enter the one of the following as the week that daylight saving ends:
 - * *week-number*: Enter a number from 1 to 4 as the number of the week in the month to end daylight saving time.
 - * *first*: Enter the keyword *first* to end daylight saving time in the first week of the month.
 - * *last*: Enter the keyword *last* to end daylight saving time in the last week of the month.
- *end-month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *end-day*: Enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *end-year*: Enter a four-digit number as the year. The range is from 1993 to 2035.
- *end-time*: Enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.

- *offset*: (OPTIONAL) Enter the number of minutes to add during the summer-time period. The range is from 1 to 1440. The default is **60 minutes**.

Example of the `clock summer-time recurring` Command

```
Dell(conf)#clock summer-time pacific recurring Mar 14 2012 00:00 Nov 7 2012
00:00
Dell(conf)#
```

 **NOTE:** If you enter <CR> after entering the `recurring` command parameter, and you have already set a one-time daylight saving time/date, the system uses that time and date as the recurring setting.

Example of Clock Summer-Time Recurring Parameters

```
Dell(conf)#clock summer-time pacific recurring ?
<1-4>      Week number to start
first      Week number to start
last       Week number to start
<cr>
Dell(conf)#clock summer-time pacific recurring
Dell(conf)#
```

Uplink Failure Detection (UFD)

Supported Modes

Standalone, PMUX, VLT, Stacking

Feature Description


UFD provides detection of the loss of upstream connectivity and, if used with network interface controller (NIC) teaming, automatic recovery from a failed link.


A switch provides upstream connectivity for devices, such as servers. If a switch loses its upstream connectivity, downstream devices also lose their connectivity. However, the devices do not receive a direct indication that upstream connectivity is lost because connectivity to the switch is still operational.

UFD allows a switch to associate downstream interfaces with upstream interfaces. When upstream connectivity fails, the switch disables the downstream links. Failures on the downstream links allow downstream devices to recognize the loss of upstream connectivity.

For example, as shown in the following illustration, Switches S1 and S2 both have upstream connectivity to Router R1 and downstream connectivity to the server. UFD operation is shown in Steps A through C:

- In Step A, the server configuration uses the connection to S1 as the primary path. Network traffic flows from the server to S1 and then upstream to R1.
- In Step B, the upstream link between S1 and R1 fails. The server continues to use the link to S1 for its network traffic, but the traffic is not successfully switched through S1 because the upstream link is down.
- In Step C, UFD on S1 disables the link to the server. The server then stops using the link to S1 and switches to using its link to S2 to send traffic upstream to R1.

 **NOTE:** In Standalone and VLT modes, the UFD group number is 1 by default and cannot be changed.



- A. Switches 1 and 2 have upstream and downstream connections to Router1 and Server via primary Links.
- B. Upstream link between Switch1 and Router1 fails. Downstream link to Server stays up temporarily.
- C. Switch1 disables downstream link to Server. Server starts to connect with Router1 using backup link to Switch2; Switch2 starts to use the backup link to Router1.

Figure 28. Uplink Failure Detection

How Uplink Failure Detection Works

UFD creates an association between upstream and downstream interfaces. The association of uplink and downlink interfaces is called an *uplink-state group*.

An interface in an uplink-state group can be a physical interface or a port-channel (LAG) aggregation of physical interfaces.

An enabled uplink-state group tracks the state of all assigned upstream interfaces. Failure on an upstream interface results in the automatic disabling of downstream interfaces in the uplink-state group. As a result, downstream devices can execute the protection or recovery procedures they have in place to establish alternate connectivity paths, as shown in the following illustration.




Figure 29. Uplink Failure Detection Example

If only one of the upstream interfaces in an uplink-state group goes down, a specified number of downstream ports associated with the upstream interface are put into a Link-Down state. You can configure this number and is calculated by the ratio of the upstream port bandwidth to the downstream port bandwidth in the same uplink-state group. This calculation ensures that there is no traffic drops due to insufficient bandwidth on the upstream links to the routers/switches.

By default, if all upstream interfaces in an uplink-state group go down, all downstream interfaces in the same uplink-state group are put into a Link-Down state.

Using UFD, you can configure the automatic recovery of downstream ports in an uplink-state group when the link status of an upstream port changes. The tracking of upstream link status does not have a major impact on central processing unit (CPU) usage.

UFD and NIC Teaming

To implement a rapid failover solution, you can use uplink failure detection on a switch with network adapter teaming on a server.

For more information, refer to [Network Interface Controller \(NIC\) Teaming](#).

For example, as shown previously, the switch/ router with UFD detects the uplink failure and automatically disables the associated downstream link port to the server. To continue to transmit traffic upstream, the server with NIC teaming detects the disabled link and automatically switches over to the backup link in order to continue to transmit traffic upstream.

Important Points to Remember

When you configure UFD, the following conditions apply.

- You can configure up to 16 uplink-state groups. By default, no uplink state groups are created in PMUX mode and uplink state group 1 is created in Standalone and VLT modes.
 - An uplink-state group is considered to be operationally *up* if it has at least one upstream interface in the Link-Up state.
 - An uplink-state group is considered to be operationally *down* if it has no upstream interfaces in the Link-Up state. No uplink-state tracking is performed when a group is disabled or in an Operationally Down state.
- You can assign physical port or port-channel interfaces to an uplink-state group in PMUX mode.
 - You can assign an interface to only one uplink-state group. Configure each interface assigned to an uplink-state group as either an upstream or downstream interface, but not both.
 - You can assign individual member ports of a port channel to the group. An uplink-state group can contain either the member ports of a port channel or the port channel itself, but not both.
 - If you assign a port channel as an upstream interface, the port channel interface enters a Link-Down state when the number of port-channel member interfaces in a Link-Up state drops below the configured `minimum number of members` parameter.
- If one of the upstream interfaces in an uplink-state group goes down, either a user-configurable set of downstream ports or all the downstream ports in the group are put in an Operationally Down state with an UFD Disabled error. The order in which downstream ports are disabled is from the lowest numbered port to the highest.
 - If one of the upstream interfaces in an uplink-state group that was down comes up, the set of UFD-disabled downstream ports (which were previously disabled due to this upstream port going down) is brought up and the UFD Disabled error is cleared.
- If you disable an uplink-state group, the downstream interfaces are not disabled regardless of the state of the upstream interfaces.
 - If an uplink-state group has no upstream interfaces assigned, you cannot disable downstream interfaces when an upstream link goes down.
- To enable the debug messages for events related to a specified uplink-state group or all groups, use the `debug uplink-state-group [group-id]` command, where the `group-id` is from 1 to 16.
 - To turn off debugging event messages, use the `no debug uplink-state-group [group-id]` command.

- For an example of debug log message, refer to [Clearing a UFD-Disabled Interface](#).

Uplink Failure Detection (SMUX mode)

In Standalone or VLT modes, by default, all the server-facing ports are tracked by the operational status of the uplink LAG. If the uplink LAG goes down, the aggregator loses its connectivity and is no longer operational. All the server-facing ports are brought down after the specified defer-timer interval, which is 10 seconds by default. If you have configured VLAN, you can reduce the defer time by changing the defer-timer value or remove it by using the `no defer-timer` command.

1. View the Uplink status group.

EXEC Privilege mode

```
show uplink-state-group
```

```
Dell#show uplink-state-group
Uplink State Group: 1   Status: Enabled, Down
```

2. Enable the uplink group tracking.

UPLINK-STATE-GROUP mode

```
enable
```

```
Dell(conf)#uplink-state-group 1
Dell(conf-uplink-state-group-1)#enable
```

To disable the uplink group tracking, use the `no enable` command.

3. Change the default timer.

UPLINK-STATE-GROUP mode

```
defer-timer seconds
```

```
Dell(conf)#uplink-state-group 1
Dell(conf-uplink-state-group-1)#defer-timer 20
Dell(conf-uplink-state-group-1)#show config
!
uplink-state-group 1
  downstream TenGigabitEthernet 0/1-12
  upstream Port-channel 128
  defer-timer 20
```

Configuring Uplink Failure Detection (PMUX mode)

To configure UFD, use the following commands.

1. Create an uplink-state group and enable the tracking of upstream links on the switch/router.

CONFIGURATION mode

```
uplink-state-group group-id
```

- *group-id*: values are from 1 to 16.

To delete an uplink-state group, use the `no uplink-state-group group-id` command.

2. Assign a port or port-channel to the uplink-state group as an upstream or downstream interface.

UPLINK-STATE-GROUP mode

```
{upstream | downstream} interface
```

For interface, enter one of the following interface types:

- TenGigabit Ethernet: enter `tengigabitethernet {slot/port | slot/port-range}`
- Port channel: enter `port-channel {1-128 | port-channel-range}`

Where *port-range* and *port-channel-range* specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example:

```
upstream tengigabitethernet 0/1-2,5,9,11-12
downstream port-channel 1-3,5
```

- A comma is required to separate each port and port-range entry.

To delete an interface from the group, use the `no {upstream | downstream} interface` command.

3. Assign a port or port-channel to the uplink-state group as an upstream or downstream interface.
UPLINK-STATE-GROUP mode

```
{upstream | downstream} interface
```

For interface, enter one of the following interface types:

- 10 Gigabit Ethernet: enter `tengigabitethernet {slot/port | slot/port-range}`
- Port channel: enter `port-channel {1-128 | port-channel-range}`

Where *port-range* and *port-channel-range* specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example:

```
upstream gigabitethernet 0/1-2,5,9,11-12
downstream port-channel 1-3,5
```

- A comma is required to separate each port and port-range entry.

To delete an interface from the group, use the `no {upstream | downstream} interface` command.

4. (Optional) Configure the number of downstream links in the uplink-state group that will be disabled (Oper Down state) if one upstream link in the group goes down.

UPLINK-STATE-GROUP mode

```
downstream disable links {number | all}
```

- *number*: specifies the number of downstream links to be brought down. The range is from 1 to 1024.
- *all*: brings down all downstream links in the group.

The default is no downstream links are disabled when an upstream link goes down.

To revert to the default setting, use the `no downstream disable links` command.

5. (Optional) Enable auto-recovery so that UFD-disabled downstream ports in the uplink-state group come up when a disabled upstream port in the group comes back up.

UPLINK-STATE-GROUP mode

```
downstream auto-recover
```

The default is auto-recovery of UFD-disabled downstream ports is enabled.

To disable auto-recovery, use the `no downstream auto-recover` command.

6. Specify the time (in seconds) to wait for the upstream port channel (LAG 128) to come back up before server ports are brought down.

UPLINK-STATE-GROUP mode

```
defer-timer seconds
```



NOTE: This command is available in Standalone and VLT modes only.

The range is from 1 to 120.

7. (Optional) Enter a text description of the uplink-state group.

UPLINK-STATE-GROUP mode

```
description text
```

The maximum length is 80 alphanumeric characters.

8. (Optional) Disable upstream-link tracking without deleting the uplink-state group.

UPLINK-STATE-GROUP mode

```
no enable
```

The default is upstream-link tracking is automatically enabled in an uplink-state group.

To re-enable upstream-link tracking, use the `enable` command.

Clearing a UFD-Disabled Interface (in PMUX mode)

You can manually bring up a downstream interface in an uplink-state group that UFD disabled and is in a UFD-Disabled Error state.

To re-enable one or more disabled downstream interfaces and clear the UFD-Disabled Error state, use the following command.

- Re-enable a downstream interface on the switch/router that is in a UFD-Disabled Error State so that it can send and receive traffic.

EXEC mode

```
clear ufd-disable {interface interface | uplink-state-group group-id}
```

For *interface*, enter one of the following interface types:

- 10 Gigabit Ethernet: `enter tengigabitethernet {slot/port | slot/port-range}`
- Port channel: `enter port-channel {1-128 | port-channel-range}`

- * Where *port-range* and *port-channel-range* specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example:

```
tengigabitethernet 0/1-2,5,9,11-12  
port-channel 1-3,5
```

- * A comma is required to separate each port and port-range entry.

```
clear ufd-disable {interface interface | uplink-state-group group-id}: re-  
enables all UFD-disabled downstream interfaces in the group. The range is from 1 to 16.
```


Example of Syslog Messages Before and After Entering the `clear ufd-disable uplink-state-group` Command

The following example message shows the Syslog messages that display when you clear the UFD-Disabled state from all disabled downstream interfaces in an uplink-state group by using the `clear ufd-disable uplink-state-group group-id` command. All downstream interfaces return to an operationally up state.

```
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to
down: Te 0/1
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to
down: Te 0/2
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to
down: Te 0/3
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/1
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/2
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/3
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed uplink state group state
to down: Group
3
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/6
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/6

Dell(conf-if-range-te-0/1-3)#do clear ufd-disable uplink-state-group 3

00:11:50: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/6
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/6
```

Displaying Uplink Failure Detection

To display information on the UFD feature, use any of the following commands.

- Display status information on a specified uplink-state group or all groups.
EXEC mode

```
show uplink-state-group [group-id] [detail]
```

- *group-id*: The values are 1 to 16.
- *detail*: displays additional status information on the upstream and downstream interfaces in each group.

- Display the current status of a port or port-channel interface assigned to an uplink-state group.
EXEC mode

```
show interfaces interface
```

interface specifies one of the following interface types:

- 10 Gigabit Ethernet: enter `tengigabitethernet slot/port`.
- Port channel: enter `port-channel {1-128}`.

If a downstream interface in an uplink-state group is disabled (Oper Down state) by uplink-state tracking because an upstream port is down, the message `error-disabled[UFD]` displays in the output.

- Display the current configuration of all uplink-state groups or a specified group.
EXEC mode or UPLINK-STATE-GROUP mode

(For EXEC mode) `show running-config uplink-state-group [group-id]`

(For UPLINK-STATE-GROUP mode) `show configuration`

- *group-id*: The values are from 1 to 16.

Example of Viewing Uplink State Group Status

```
Dell# show uplink-state-group
```

```
Uplink State Group: 1 Status: Enabled, Up
Uplink State Group: 3 Status: Enabled, Up
Uplink State Group: 5 Status: Enabled, Down
Uplink State Group: 6 Status: Enabled, Up
Uplink State Group: 7 Status: Enabled, Up
Uplink State Group: 16 Status: Disabled, Up
```

```
Dell# show uplink-state-group 16
```

```
Uplink State Group: 16 Status: Disabled, Up
```

```
Dell#show uplink-state-group detail
```

```
(Up): Interface up      (Dwn): Interface down  (Dis): Interface disabled
```

```
Uplink State Group      : 1          Status: Enabled, Up
Defer Timer              : 10 sec
Upstream Interfaces     : Po 128 (Up)
Downstream Interfaces   : Te 0/1 (Dwn) Te 0/2 (Dwn) Te 0/3 (Up) Te 0/4 (Dwn) Te
0/5 (Up)
                        Te 0/6 (Dwn) Te 0/7 (Up) Te 0/8 (Up)
```

Dell#

Example of Viewing Interface Status with UFD Information

```
Dell#show interfaces tengigabitethernet 0/7
TenGigabitEthernet 0/7 is up, line protocol is down (error-disabled[UFD])
Hardware is Force10Eth, address is 00:01:e8:32:7a:47
  Current address is 00:01:e8:32:7a:47
Interface index is 280544512
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:25:46
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:01:23
```

Examples of Viewing UFD Output

```
Dell#show running-config uplink-state-group
!
uplink-state-group 1
no enable
downstream TenGigabitEthernet 0/3
upstream TenGigabitEthernet 0/1
Dell#
```

```
Dell(conf-uplink-state-group-16)# show configuration
!
uplink-state-group 16
no enable
description test
downstream disable links all
downstream TengigabitEthernet 0/4
upstream TengigabitEthernet 0/5
upstream Port-channel 8
```

Sample Configuration: Uplink Failure Detection

The following example shows a sample configuration of UFD on a switch/router in which you configure as follows.

- Configure uplink-state group 3.
- Add downstream links GigabitEthernet 0/1, 0/2, 0/5, 0/9, 0/11, and 0/12.
- Configure two downstream links to be disabled if an upstream link fails.
- Add upstream links GigabitEthernet 0/3 and 0/4.
- Add a text description for the group.
- Verify the configuration with various show commands.

Example of Configuring UFD

```
Dell(conf)#uplink-state-group 3
Dell(conf-uplink-state-group-3)#
```

```
00:23:52: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP: Changed uplink state group Admin
state to up:
Group 3
```

```
Dell(conf-uplink-state-group-3)#downstream tengigabitethernet 0/1-2,5,9,11-12
Dell(conf-uplink-state-group-3)#downstream disable links 2
Dell(conf-uplink-state-group-3)#upstream tengigabitethernet 0/3-4
Dell(conf-uplink-state-group-3)#description Testing UFD feature
Dell(conf-uplink-state-group-3)#show config
```

```
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4
```

```
Dell#show running-config uplink-state-group
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4
```

```
Dell#show uplink-state-group 3
```

```
Uplink State Group: 3 Status: Enabled, Up
```

```
Dell#show uplink-state-group detail
```

```
(Up): Interface up (Dwn): Interface down (Dis): Interface disabled
```

```
Uplink State Group      : 3 Status: Enabled, Up
Upstream Interfaces     : Te 0/3(Up) Te 0/4(Up)
Downstream Interfaces   : Te 0/1(Up) Te 0/2(Up) Te 0/5(Up) Te 0/9(Up) Te 0/11(Up)
                        : Te 0/12(Up)
```

```
< After a single uplink port fails >
```

```
Dell#show uplink-state-group detail
```

```
(Up): Interface up (Dwn): Interface down (Dis): Interface disabled
```

Uplink State Group : 3 Status: Enabled, Up
Upstream Interfaces : Te 0/3 (Dwn) Te 0/4 (Up)
Downstream Interfaces : Te 0/1 (Dis) Te 0/2 (Dis) Te 0/5 (Up) Te 0/9 (Up) Te
0/11 (Up)
Te 0/12 (Up)

PMUX Mode of the IO Aggregator

This chapter provides an overview of the PMUX mode.

I/O Aggregator (IOA) Programmable MUX (PMUX) Mode

IOA PMUX is a mode that provides flexibility of operation with added configurability. This involves creating multiple LAGs, configuring VLANs on uplinks and the server side, configuring data center bridging (DCB) parameters, and so forth.

By default, IOA starts up in IOA Standalone mode. You can change to PMUX mode by executing the following commands and then reloading the IOA. After the IOA reboots, the IOA operates in PMUX mode. PMUX mode supports both stacking and VLT operations.

Configuring and Changing to PMUX Mode

After the IOA is operational in the default Standalone mode:

1. Connect the terminal to the console port on the IOA to access the CLI and enter the following commands:

```
Login: username
Password: *****
Dell> enable
Dell#
Dell#show system stack-unit 0 iom-mode
Unit Boot-Mode Next-Boot
-----
0 standalone standalone
Dell#
```

2. Change IOA mode to PMUX mode.

```
Dell(conf)# stack-unit 0 iom-mode programmable-mux
```

Where `stack-unit 0` defines the default stack-unit number.

3. Delete the startup configuration file.

```
Dell# delete startup-config
```

4. Reboot the IOA by entering the `reload` command.

```
Dell# reload
```

5. Repeat the above steps for each member of the IOA in PMUX mode.

After system is up, you can see the PMUX mode status:

```
Dell#sh system stack-unit 0 iom-mode
Unit          Boot-Mode          Next-Boot
```

```
-----  
0          programmable-mux          programmable-mux  
Dell#
```

The IOA is now ready for PMUX operations.


Configuring the Commands without a Separate User Account

Starting with Dell Networking OS version 9.3(0.0), you can configure the PMUX mode CLI commands without having to configure a new, separate user profile. The user profile you defined to access and log in to the switch is sufficient to configure the PMUX mode commands.

The IOA PMUX Mode CLI Commands section lists the PMUX mode CLI commands that you can now configure without a separate user account.

Virtual Link Trunking (VLT)

VLT allows physical links between two chassis to appear as a single virtual link to the network core. VLT eliminates the requirement for Spanning Tree protocols by allowing link aggregation group (LAG) terminations on two separate distribution or core switches, and by supporting a loop-free topology. VLT provides Layer 2 multipathing, creating redundancy through increased bandwidth and enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

 **NOTE:** When you launch the VLT link, the VLT peer-ship is not established if any of the following is **TRUE**:

- The VLT System-MAC configured on both the VLT peers do not match.
- The VLT Unit-Id configured on both the VLT peers are identical.
- The VLT System-MAC or Unit-Id is configured only on one of the VLT peers.
- The VLT domain ID is not the same on both peers.

If the VLT peer-ship is already established, changing the System-MAC or Unit-Id does not cause VLT peer-ship to go down.

Also, if the VLT peer-ship is already established and the VLT Unit-Id or System-MAC are configured on both peers, then changing the CLI configurations on the VLT Unit-Id or System-MAC is rejected if any of the following become **TRUE**:

- After making the CLI configuration change, the VLT Unit-Id becomes identical on both peers.
- After making the CLI configuration change, the VLT System-MAC do not match on both peers.

When the VLT peer-ship is already established, you can remove the VLT Unit-Id or System-MAC configuration from either or both peers. However, removing configuration settings can cause the VLT ports to go down if you configure the Unit-Id or System-MAC on only one of the VLT peers.

Overview

VLT allows physical links between two chassis to appear as a single virtual link to the network core or other switches such as Edge, Access, or top-of-rack (ToR).

VLT reduces the role of spanning tree protocols (STPs) by allowing link aggregation group (LAG) terminations on two separate distribution or core switches, and by supporting a loop-free topology. (To prevent the initial loop that may occur prior to VLT being established, use a spanning tree protocol.

VLT provides Layer 2 multipathing, creating redundancy through increased bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

Virtual link trunking offers the following benefits:

- Allows a single device to use a LAG across two upstream devices.
- Eliminates STP-blocked ports.
- Provides a loop-free topology.
- Uses all available uplink bandwidth.
- Provides fast convergence if either the link or a device fails.
- Optimized forwarding with virtual router redundancy protocol (VRRP).
- Provides link-level resiliency.
- Assures high availability.

As shown in the following example, VLT presents a single logical Layer 2 domain from the perspective of attached devices that have a virtual link trunk terminating on separate chassis in the VLT domain. However, the two VLT chassis are independent Layer2/Layer3 (L2/L3) switches for devices in the upstream network. L2/L3 control plane protocols and system management features function normally in VLT mode. Features such as VRRP and internet group management protocol (IGMP) snooping require state information coordinating between the two VLT chassis. IGMP and VLT configurations must be identical on both sides of the trunk to ensure the same behavior on both sides.

Setting up VLT

The following figure shows the sample VLT topology.

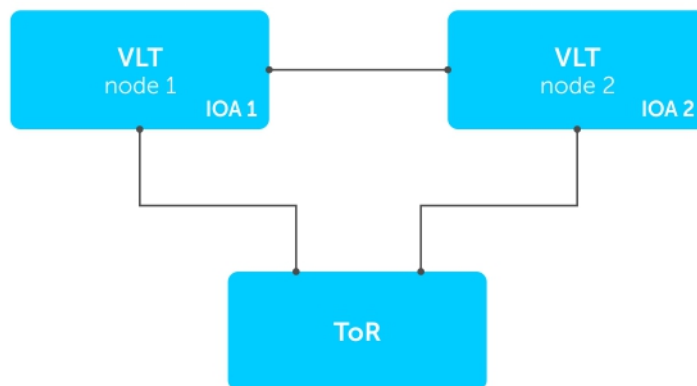


Figure 30. Sample VLT Topology

Port 9 will be used as ICL link and this one 10G port is connected between the two Aggregators.

In PMUX VLT, you can choose any uplink ports for configuring VLT.


Important Points to Remember

- VLT port channel interfaces must be switch ports.
- Dell Networking strongly recommends that the VLTi (VLT interconnect) be a static LAG and that you disable LACP on the VLTi.
- If the `lacp-ungroup` feature is not supported on the ToR, reboot the VLT peers one at a time. After rebooting, verify that VLTi (ICL) is active before attempting DHCP connectivity.

Configuration Notes

When you configure VLT, the following conditions apply.

- VLT domain
 - A VLT domain supports two chassis members, which appear as a single logical device to network access devices connected to VLT ports through a port channel.
 - A VLT domain consists of the two core chassis, the interconnect trunk, backup link, and the LAG members connected to attached devices.
 - Each VLT domain has a unique MAC address that you create or VLT creates automatically.
 - ARP tables are synchronized between the VLT peer nodes.
 - VLT peer switches operate as separate chassis with independent control and data planes for devices attached on non-VLT ports.
 - One chassis in the VLT domain is assigned a primary role; the other chassis takes the secondary role. The primary and secondary roles are required for scenarios when connectivity between the chassis is lost. VLT assigns the primary chassis role according to the lowest MAC address. You can configure the primary role.
 - In a VLT domain, the peer switches must run the same Dell Networking OS software version.
 - Separately configure each VLT peer switch with the same VLT domain ID and the VLT version. If the system detects mismatches between VLT peer switches in the VLT domain ID or VLT version, the VLT Interconnect (VLTi) does not activate. To find the reason for the VLTi being down, use the `show vlt statistics` command to verify that there are mismatch errors, then use the `show vlt brief` command on each VLT peer to view the VLT version on the peer switch. If the VLT version is more than one release different from the current version in use, the VLTi does not activate.
 - The chassis members in a VLT domain support connection to orphan hosts and switches that are not connected to both switches in the VLT core.
- VLT interconnect (VLTi)
 - The VLT interconnect supports a maximum of two 10G ports.
 - A VLT interconnect over 1G ports is *not* supported.
 - The port channel must be in Default mode (not Switchport mode) to have VLTi recognize it.
 - The system automatically includes the required VLANs in VLTi. You do not need to manually select VLANs.
 - VLT peer switches operate as separate chassis with independent control and data planes for devices attached to non-VLT ports.
 - Port-channel link aggregation (LAG) across the ports in the VLT interconnect is required; individual ports are not supported. Dell Networking strongly recommends configuring a static LAG for VLTi.
 - The VLT interconnect synchronizes L2 and L3 control-plane information across the two chassis.
 - The VLT interconnect is used for data traffic only when there is a link failure that requires using VLTi in order for data packets to reach their final destination.
 - Unknown, multicast, and broadcast traffic can be flooded across the VLT interconnect.

- MAC addresses for VLANs configured across VLT peer chassis are synchronized over the VLT interconnect on an egress port such as a VLT LAG. MAC addresses are the same on both VLT peer nodes.
 - ARP entries configured across the VLTi are the same on both VLT peer nodes.
 - If you shut down the port channel used in the VLT interconnect on a peer switch in a VLT domain in which you did not configure a backup link, the switch's role displays in the `show vlt brief` command output as Primary instead of Standalone.
 - When you change the default VLAN ID on a VLT peer switch, the VLT interconnect may flap.
 - In a VLT domain, the following software features are supported on VLTi: link layer discovery protocol (LLDP), flow control, port monitoring, jumbo frames, and data center bridging (DCB).
 - When you enable the VLTi link, the link between the VLT peer switches is established if the following configured information is true on both peer switches:
 - * the VLT system MAC address matches.
 - * the VLT unit-id is not identical.
-  **NOTE:** If you configure the VLT system MAC address or VLT unit-id on only one of the VLT peer switches, the link between the VLT peer switches is not established. Each VLT peer switch must be correctly configured to establish the link between the peers.
- If the link between the VLT peer switches is established, changing the VLT system MAC address or the VLT unit-id causes the link between the VLT peer switches to become disabled. However, removing the VLT system MAC address or the VLT unit-id may disable the VLT ports if you happen to configure the unit ID or system MAC address on only one VLT peer at any time.
 - If the link between VLT peer switches is established, any change to the VLT system MAC address or unit-id fails if the changes made create a mismatch by causing the VLT unit-ID to be the same on both peers and/or the VLT system MAC address does not match on both peers.
 - If you replace a VLT peer node, preconfigure the switch with the VLT system MAC address, unit-id, and other VLT parameters before connecting it to the existing VLT peer switch using the VLTi connection.
- VLT backup link
 - In the backup link between peer switches, heartbeat messages are exchanged between the two chassis for health checks. The default time interval between heartbeat messages over the backup link is 1 second. You can configure this interval. The range is from 1 to 5 seconds. DSCP marking on heartbeat messages is CS6.
 - In order that the chassis backup link does not share the same physical path as the interconnect trunk, Dell Networking recommends using the management ports on the chassis and traverse an out-of-band management network. The backup link can use user ports, but not the same ports the interconnect trunk uses.
 - The chassis backup link does not carry control plane information or data traffic. Its use is restricted to health checks only.
 - Virtual link trunks (VLTs) between access devices and VLT peer switches
 - To connect servers and access switches with VLT peer switches, you use a VLT port channel, as shown in [Overview](#).
 - The discovery protocol running between VLT peers automatically generates the ID number of the port channel that connects an access device and a VLT switch. The discovery protocol uses LACP properties to identify connectivity to a common client device and automatically generates a VLT number for port channels on VLT peers that connects to the device. The discovery protocol requires that an attached device always runs LACP over the port-channel interface.
 - VLT provides a loop-free topology for port channels with endpoints on different chassis in the VLT domain.
 - VLT uses shortest path routing so that traffic destined to hosts via directly attached links on a chassis does not traverse the chassis-interconnect link.

- VLT allows multiple active parallel paths from access switches to VLT chassis.
- VLT supports port-channel links with LACP between access switches and VLT peer switches. Dell Networking recommends using static port channels on VLTi.
- If VLTi connectivity with a peer is lost but the VLT backup connectivity indicates that the peer is still alive, the VLT ports on the Secondary peer are orphaned and are shut down.
- Software features supported on VLT port-channels
 - For information about configuring IGMP Snooping in a VLT domain, refer to [VLT and IGMP Snooping](#).
 - All system management protocols are supported on VLT ports, including SNMP, RMON, AAA, ACL, DNS, FTP, SSH, Syslog, NTP, RADIUS, SCP, TACACS+, Telnet, and LLDP.
 - Enable Layer 3 VLAN connectivity VLT peers by configuring a VLAN network interface for the same VLAN on both switches.
 - Dell Networking does not recommend enabling peer-routing if the CAM is full. To enable peer-routing, a minimum of two local DA spaces for wild card functionality are required.
- Software features supported on VLT physical ports
 - In a VLT domain, the following software features are supported on VLT physical ports: 802.1p, LLDP, flow control, port monitoring, and jumbo frames.
- Software features not supported with VLT
 - In a VLT domain, the following software features are supported on non-VLT ports: 802.1x, , DHCP snooping, FRRP, IPv6 dynamic routing, ingress and egress QOS.
- Failure scenarios
 - On a link failover, when a VLT port channel fails, the traffic destined for that VLT port channel is redirected to the VLTi to avoid flooding.
 - When a VLT switch determines that a VLT port channel has failed (and that no other local port channels are available), the peer with the failed port channel notifies the remote peer that it no longer has an active port channel for a link. The remote peer then enables data forwarding across the interconnect trunk for packets that would otherwise have been forwarded over the failed port channel. This mechanism ensures reachability and provides loop management. If the VLT interconnect fails, the VLT software on the primary switch checks the status of the remote peer using the backup link. If the remote peer is up, the secondary switch disables all VLT ports on its device to prevent loops.
 - If all ports in the VLT interconnect fail, or if the messaging infrastructure fails to communicate across the interconnect trunk, the VLT management system uses the backup link interface to determine whether the failure is a link-level failure or whether the remote peer has failed entirely. If the remote peer is still alive (heartbeat messages are still being received), the VLT secondary switch disables its VLT port channels. If keepalive messages from the peer are not being received, the peer continues to forward traffic, assuming that it is the last device available in the network. In either case, after recovery of the peer link or reestablishment of message forwarding across the interconnect trunk, the two VLT peers resynchronize any MAC addresses learned while communication was interrupted and the VLT system continues normal data forwarding.
 - If the primary chassis fails, the secondary chassis takes on the operational role of the primary.
- The SNMP MIB reports VLT statistics.

Primary and Secondary VLT Peers

Primary and Secondary VLT Peers are supported on the Aggregator.

To prevent issues when connectivity between peers is lost, you can designate Primary and Secondary roles for VLT peers . You can elect or configure the Primary Peer. By default, the peer with the lowest MAC address is selected as the Primary Peer.

If the VLTi link fails, the status of the remote VLT Primary Peer is checked using the backup link. If the remote VLT Primary Peer is available, the Secondary Peer disables all VLT ports to prevent loops.

If all ports in the VLTi link fail or if the communication between VLTi links fails, VLT checks the backup link to determine the cause of the failure. If the failed peer can still transmit heartbeat messages, the Secondary Peer disables all VLT member ports and any Layer 3 interfaces attached to the VLAN associated with the VLT domain. If heartbeat messages are not received, the Secondary Peer forwards traffic and assumes the role of the Primary Peer. If the original Primary Peer is restored, the VLT peer is reassigned as the Primary Peer and the other peer must be reassigned as a Secondary Peer. Peer role changes are reported as SNMP traps.

VLT Bandwidth Monitoring

When bandwidth usage of the VLTi (ICL) exceeds 80%, a syslog error message (shown in the following message) and an SNMP trap are generated.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-LAG (port-channel 25) crosses threshold. Bandwidth usage (80 )
```

When the bandwidth usage drops below the 80% threshold, the system generates another syslog message (shown in the following message) and an SNMP trap.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-LAG (port-channel 25) reaches below threshold. Bandwidth usage (74 )VLT show remote port channel status
```

VLT and IGMP Snooping

When configuring IGMP Snooping with VLT, ensure the configurations on both sides of the VLT trunk are identical to get the same behavior on both sides of the trunk.

When you configure IGMP snooping on a VLT node, the dynamically learned groups and multicast router ports are automatically learned on the VLT peer node.

VLT Port Delayed Restoration

When a VLT node boots up, if the VLT ports have been previously saved in the start-up configuration, they are not immediately enabled.

To ensure MAC and ARP entries from the VLT per node are downloaded to the newly enabled VLT node, the system allows time for the VLT ports on the new node to be enabled and begin receiving traffic.

The `delay-restore` feature waits for all saved configurations to be applied, then starts a configurable timer. After the timer expires, the VLT ports are enabled one-by-one in a controlled manner. The delay between bringing up each VLT port-channel is proportional to the number of physical members in the port-channel. The default is 90 seconds.

If you enable IGMP snooping, IGMP queries are also sent out on the VLT ports at this time allowing any receivers to respond to the queries and update the multicast table on the new node.


This delay in bringing up the VLT ports also applies when the VLTi link recovers from a failure that caused the VLT ports on the secondary VLT peer node to be disabled.

Non-VLT ARP Sync

In the Dell Networking OS version 9.2(0.0), ARP entries (including ND entries) learned on other ports are synced with the VLT peer to support station move scenarios.

Prior to Dell Networking OS version 9.2.(0.0), only ARP entries learned on VLT ports were synced between peers.

Additionally, ARP entries resulting from station movements from VLT to non-VLT ports or to different non-VLT ports are learned on the non-VLT port and synced with the peer node. The peer node is updated to use the new non-VLT port.

 **NOTE:** ARP entries learned on non-VLT, non-spanned VLANs are not synced with VLT peers.

Verifying a VLT Configuration

To monitor the operation or verify the configuration of a VLT domain, use any of the following `show` commands on the primary and secondary VLT switches.

- Display information on backup link operation.

EXEC mode

```
show vlt backup-link
```

- Display general status information about VLT domains currently configured on the switch.

EXEC mode

```
show vlt brief
```

- Display detailed information about the VLT-domain configuration, including local and peer port-channel IDs, local VLT switch status, and number of active VLANs on each port channel.

EXEC mode

```
show vlt detail
```

- Display the VLT peer status, role of the local VLT switch, VLT system MAC address and system priority, and the MAC address and priority of the locally-attached VLT device.

EXEC mode

```
show vlt role
```

- Display the current configuration of all VLT domains or a specified group on the switch.

EXEC mode

```
show running-config vlt
```

- Display statistics on VLT operation.

EXEC mode

```
show vlt statistics
```

- Display the current status of a port or port-channel interface used in the VLT domain.

EXEC mode

```
show interfaces interface
```

– *interface*: specify one of the following interface types:

* 10-Gigabit Ethernet: enter `tengigabitethernet slot/port`.

* Port channel: enter port-channel {1-128}.

Example of the show vlt backup-link Command

```
Dell_VLTpeer1# show vlt backup-link
```

```
VLT Backup Link
-----
Destination:          10.11.200.18
Peer HeartBeat status: Up
HeartBeat Timer Interval: 1
HeartBeat Timeout:    3
UDP Port:             34998
HeartBeat Messages Sent: 1026
HeartBeat Messages Received: 1025
```

```
Dell_VLTpeer2# show vlt backup-link
```

```
VLT Backup Link
-----
Destination:          10.11.200.20
Peer HeartBeat status: Up
HeartBeat Timer Interval: 1
HeartBeat Timeout:    3
UDP Port:             34998
HeartBeat Messages Sent: 1030
HeartBeat Messages Received: 1014
```

Example of the show vlt brief Command

```
Dell_VLTpeer1# show vlt brief
```

```
VLT Domain Brief
-----
Domain ID:            1000
Role:                 Secondary
Role Priority:        32768
ICL Link Status:     Up
HeartBeat Status:    Up
VLT Peer Status:     Up
Local Unit Id:       0
Version:              5(1)
Local System MAC address: 00:01:e8:8a:e9:70
Remote System MAC address: 00:01:e8:8a:e7:e7
Configured System MAC address: 00:0a:0a:01:01:0a
Remote system version: 5(1)
Delay-Restore timer:  90 seconds
```

```
Dell_VLTpeer2# show vlt brief
```

```
VLT Domain Brief
-----
Domain ID:            1000
Role:                 Primary
Role Priority:        32768
ICL Link Status:     Up
HeartBeat Status:    Up
VLT Peer Status:     Up
Local Unit Id:       1
Version:              5(1)
Local System MAC address: 00:01:e8:8a:e7:e7
Remote System MAC address: 00:01:e8:8a:e9:70
Configured System MAC address: 00:0a:0a:01:01:0a
```

Remote system version: 5(1)
Delay-Restore timer: 90 seconds

Example of the show vlt detail Command

Dell_VLTpeer1# show vlt detail

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
100	100	UP	UP	10, 20, 30
127	2	UP	UP	20, 30

Dell_VLTpeer2# show vlt detail

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
2	127	UP	UP	20, 30
100	100	UP	UP	10, 20, 30

Example of the show vlt role Command

Dell_VLTpeer1# show vlt role

VLT Role

VLT Role: Primary
System MAC address: 00:01:e8:8a:df:bc
System Role Priority: 32768
Local System MAC address: 00:01:e8:8a:df:bc
Local System Role Priority: 32768

Dell_VLTpeer2# show vlt role

VLT Role

VLT Role: Secondary
System MAC address: 00:01:e8:8a:df:bc
System Role Priority: 32768
Local System MAC address: 00:01:e8:8a:df:e6
Local System Role Priority: 32768

Example of the show running-config vlt Command

Dell_VLTpeer1# show running-config vlt
!
vlt domain 30
peer-link port-channel 60
back-up destination 10.11.200.18

Dell_VLTpeer2# show running-config vlt
!
vlt domain 30
peer-link port-channel 60
back-up destination 10.11.200.20

Example of the show vlt statistics Command

```
Dell_VLTpeer1# show vlt statistics
```

```
VLT Statistics
-----
HeartBeat Messages Sent:      987
HeartBeat Messages Received: 986
ICL Hello's Sent:            148
ICL Hello's Received:        98
```

```
Dell_VLTpeer2# show vlt statistics
```

```
VLT Statistics
-----
HeartBeat Messages Sent:      994
HeartBeat Messages Received: 978
ICL Hello's Sent:            89
ICL Hello's Received:        89
```

VLT Sample Configurations

To configure VLT, configure a backup link and interconnect trunk, create a VLT domain, configure a backup link and interconnect trunk, and connect the peer switches in a VLT domain to an attached access device (switch or server).

Review the following examples of VLT configurations.

Configuring Virtual Link Trunking (VLT Peer 1)

Configure the backup link.

```
Dell_VLTpeer1(conf)#interface ManagementEthernet 0/0
Dell_VLTpeer1(conf-if-ma-0/0)#ip address 10.11.206.23/
Dell_VLTpeer1(conf-if-ma-0/0)#no shutdown
Dell_VLTpeer1(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
Dell_VLTpeer1(conf)#interface port-channel 100
Dell_VLTpeer1(conf-if-po-100)#channel-member TenGigE 0/6,7
Dell_VLTpeer1(conf-if-po-100)#no shutdown
Dell_VLTpeer1(conf-if-po-100)#exit
```

Enable VLT and create a VLT domain with a backup-link VLT interconnect trunk (VLTi).

```
Dell_VLTpeer1(conf)#vlt domain 999
Dell_VLTpeer1(conf-vlt-domain)#peer-link port-channel 100
Dell_VLTpeer1(conf-vlt-domain)#back-up destination 10.11.206.35
Dell_VLTpeer1(conf-vlt-domain)#exit
```

Configure the port channel to an attached device.

```
Dell_VLTpeer1(conf)#interface port-channel 110
Dell_VLTpeer1(conf-if-po-110)#switchport
Dell_VLTpeer1(conf-if-po-110)#channel-member TenGigE 0/5
Dell_VLTpeer1(conf-if-po-110)#no shutdown
Dell_VLTpeer1(conf-if-po-110)#vlt-peer-lag port-channel 110
Dell_VLTpeer1(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
Dell_VLTpeer1# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

NUM Status Description Q Ports
10 Active                U Po110(Te 0/5)
                       T Po100(Te 0/6,7)
```

Configuring Virtual Link Trunking (VLT Peer 2)

Configure the backup link.

```
Dell_VLTpeer2(conf)#interface ManagementEthernet 0/0
Dell_VLTpeer2(conf-if-ma-0/0)#ip address 10.11.206.35/
Dell_VLTpeer2(conf-if-ma-0/0)#no shutdown
Dell_VLTpeer2(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
Dell_VLTpeer2(conf)#interface port-channel 100
Dell_VLTpeer2(conf-if-po-100)#channel-member TenGigE 0/3,4
Dell_VLTpeer2(conf-if-po-100)#no shutdown
Dell_VLTpeer2(conf-if-po-100)#exit
```

Enable VLT and create a VLT domain with a backup-link VLT interconnect (VLTi).

```
Dell_VLTpeer2(conf)#vlt domain 999
Dell_VLTpeer2(conf-vlt-domain)#peer-link port-channel 100
Dell_VLTpeer2(conf-vlt-domain)#back-up destination 10.11.206.23
Dell_VLTpeer2(conf-vlt-domain)#exit
```

Configure the port channel to an attached device.

```
Dell_VLTpeer2(conf)#interface port-channel 110
Dell_VLTpeer2(conf-if-po-110)#switchport
Dell_VLTpeer2(conf-if-po-110)#channel-member TenGigE 0/8
Dell_VLTpeer2(conf-if-po-110)#no shutdown
Dell_VLTpeer2(conf-if-po-110)#vlt-peer-lag port-channel 110
Dell_VLTpeer2(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
Dell_VLTpeer2# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

NUM Status Description Q Ports
10 Active                U Po110(Te 0/8)
                       T Po100(Te 0/3,4)
```

Verifying a Port-Channel Connection to a VLT Domain (From an Attached Access Switch)

On an access device, verify the port-channel connection to a VLT domain.

```
Dell_TORswitch(conf)# show running-config interface port-channel 11
!  
interface Port-channel 11  
switchport  
channel-member TenGigE 0/1,2  
no shutdown
```

Troubleshooting VLT

To help troubleshoot different VLT issues that may occur, use the following information.


 **NOTE:** For information on VLT Failure mode timing and its impact, contact your Dell Networking representative.

Table 21. Troubleshooting VLT

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
Bandwidth monitoring	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above the 80% threshold and when it drops below 80%.	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above its threshold.	Depending on the traffic that is received, the traffic can be offloaded inVLTi.
Domain ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	Verify the domain ID matches on both VLT peers.
Dell Networking OS Version mismatch	A syslog error message is generated.	A syslog error message is generated.	Follow the correct upgrade procedure for the unit with the mismatched Dell Networking OS version.
Remote VLT port channel status	N/A	N/A	Use the <code>show vlt detail</code> and <code>show vlt brief</code> commands to view the VLT port channel status information.
System MAC mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify that the unit ID of VLT peers is not the same on both units and

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
Unit ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message is generated.	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message is generated.	that the MAC address is the same on both units. Verify the unit ID is correct on both VLT peers. Unit ID numbers must be sequential on peer units; for example, if Peer 1 is unit ID "0", Peer 2 unit ID must be "1".
Version ID mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify the Dell Networking OS software versions on the VLT peers is compatible. For more information, refer to the <i>Release Notes</i> for this release.
VLT LAG ID is not configured on one VLT peer	A syslog error message is generated. The peer with the VLT configured remains active.	A syslog error message is generated. The peer with the VLT configured remains active.	Verify the VLT LAG ID is configured correctly on both VLT peers.
VLT LAG ID mismatch	The VLT port channel is brought down. A syslog error message is generated.	The VLT port channel is brought down. A syslog error message is generated.	Perform a mismatch check after the VLT peer is established.

NPIV Proxy Gateway

The N-port identifier virtualization (NPIV) Proxy Gateway (NPG) feature provides FCoE-FC bridging capability on the FN 2210S Aggregator, allowing server CNAs to communicate with SAN fabrics over the FN 2210S Aggregator.

NPIV Proxy Gateway Configuration

The Aggregator switches function as a top-of-rack edge switch that supports Converged Enhanced Ethernet (CEE) traffic — Fibre Channel over Ethernet (FCoE) for storage, Interprocess Communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

The NPG provides FCoE-FC bridging capability on the Aggregator.

This chapter describes how to configure and use an NPIV proxy gateway on the Aggregator in a SAN.

NPIV Proxy Gateway Operations and Capabilities

Benefits of an NPIV Proxy Gateway

The Aggregator functions as a top-of-rack edge switch that supports CEE traffic — FCoE for storage, IPC for servers, and Ethernet LAN (IP cloud) for data — as well as FC links to one or more SAN fabrics.

Using an NPG helps resolve the following problems in a storage area network:

- Fibre Channel storage networks typically consist of servers connected to edge switches, which are connected to SAN core switches. As the SAN grows, it is necessary to add more ports and SAN switches. This results in an increase in the required domain IDs, which may surpass the upper limit of 239 domain IDs supported in the SAN network. An NPG avoids the need for additional domain IDs because it is deployed outside the SAN and uses the domain IDs of core switches in its FCoE links.
- With the introduction of 10GbE links, FCoE is being implemented for server connections to optimize performance. However, a SAN traditionally uses Fibre Channel to transmit storage traffic. FCoE servers require an efficient and scalable bridging feature to access FC storage arrays, which an NPG provides.

NPIV Proxy Gateway Operation

Consider a sample scenario of NPG operation. An FX2 server chassis configured as an NPG does not join a SAN fabric, but functions as an FCoE-FC bridge that forwards storage traffic between servers and core SAN switches. The core switches forward SAN traffic to and from FC storage arrays.

An FX2 chassis FC port is configured as an N (node) port that logs in to an F (fabric) port on the upstream FC core switch and creates a channel for N-port identifier virtualization. NPIV allows multiple N-port fabric logins at the same time on a single, physical Fibre Channel link.

Converged Network Adapter (CNA) ports on servers connect to the FX2 chassis Ten-Gigabit Ethernet ports and log in to an upstream FC core switch through the N port. Server fabric login (FLOGI) requests are converted into fabric discovery (FDISC) requests before being forwarded to the FC core switch.

Servers use CNA ports to connect over FCoE to an Ethernet port in ENode mode on the NPIV proxy gateway. FCoE transit with FIP snooping is automatically enabled and configured on the FX2 gateway to prevent unauthorized access and data transmission to the SAN network. FIP is used by server CNAs to discover an FCoE switch operating as an FCoE forwarder (FCF).

The NPIV proxy gateway aggregates multiple locally connected server CNA ports into one or more upstream N port links, conserving the number of ports required on an upstream FC core switch while providing an FCoE-to-FC bridging functionality. The upstream N ports on an FX2 can connect to the same or multiple fabrics.

Using an FCoE map applied to downstream (server-facing) Ethernet ports and upstream (fabric-facing) FC ports, you can configure the association between a SAN fabric and the FCoE VLAN that connects servers over the NPIV proxy gateway to FC switches in the fabric. An FCoE map virtualizes the upstream SAN fabric as an FCF to downstream CNA ports on FCoE-enabled servers as follows:

- As soon as an FC N port comes online (`no shutdown` command), the NPG starts sending FIP multicast advertisements, which contain the fabric name derived from the 64-bit worldwide name (WWN) of the principal SAN switch. (The principal switch in a fabric is the FC switch with the lowest domain ID.)
- When you apply the FCoE map to a server-facing Ethernet port in ENode mode, ACLs are automatically configured to allow only FCoE traffic from servers that perform a successful FLOGI on the FC switch. All other traffic on the VLAN is denied.

You can specify one or more upstream N ports in an FCoE map. The FCoE map also contains the VLAN ID of the dedicated VLAN used to transmit FCoE traffic between the SAN fabric and servers.

NPIV Proxy Gateway: Protocol Services

The Aggregator with the NPG provides the following protocol services:

- Fibre Channel service to create N ports and log in to an upstream FC switch.
- FCoE service to perform:
 - Virtualization of FC N ports on an NPG so that they appear as FCoE FCFs to downstream servers.
 - NPIV service to perform the association and aggregation of FCoE servers to upstream F ports on core switches (through N ports on the NPG). Conversion of server FLOGIs and FDISCs, which are received over the Aggregator with the ENode ports, are converted into FDISCs addressed to the upstream F ports on core switches.

NPIV Proxy Gateway Functionality

The Aggregator with the NPG provides the following functionality in a storage area network:

- FIP Snooping bridge that provides security for FCoE traffic using ACLs.
- FCoE gateway that provides FCoE-to-FC bridging. N-port virtualization using FCoE maps exposes upstream F ports as FCF ports to downstream server-facing ENode ports on the NPG.

NPIV Proxy Gateway: Terms and Definitions

The following table describes the terms used in an NPG configuration on the Aggregator.

Table 22. Aggregator with the NPIV Proxy Gateway: Terms and Definitions

Term	Description
FC port	Fibre Channel port on the Aggregator that operates in autosensing, 2, 4, or 8-Gigabit mode. On an NPIV proxy gateway, an FC port can be used as a downlink for a server connection and an uplink for a fabric connection.
F port	Port mode of an FC port connected to an end node (N) port on an Aggregator with the NPIV proxy gateway.
N port	Port mode of an Aggregator with the FC port that connects to an F port on an FC switch in a SAN fabric. On an Aggregator with the NPIV proxy gateway, an N port also functions as a proxy for multiple server CNA-port connections.
ENode port	Port mode of a server-facing Aggregator with the Ethernet port that provides access to FCF functionality on a fabric.
CNA port	N-port functionality on an FCoE-enabled server port. A converged network adapter (CNA) can use one or more Ethernet ports. CNAs can encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.
DCB map	Template used to configure DCB parameters, including priority-based flow control (PFC) and enhanced transmission selection (ETS), on CEE ports.
Fibre Channel fabric	Network of Fibre Channel devices and storage arrays that inter-operate and communicate.
FCF	Fibre Channel forwarder: FCoE-enabled switch that can forward FC traffic to both downstream FCoE and upstream FC devices. An NPIV proxy gateway functions as an FCF to export upstream F port configurations to downstream server CNA ports.
FC-MAP	FCoE MAC-address prefix — The unique 24-bit MAC address prefix in FCoE packets used to generate a fabric-provided MAC address (FPMA). The FPMA is required to send FCoE packets from a server to a SAN fabric.
FCoE map	Template used to configure FCoE and FC parameters on Ethernet and FC ports in a converged fabric.
FCoE VLAN	VLAN dedicated to carrying only FCoE traffic between server CNA ports and a SAN fabric. (FCoE traffic must travel in a VLAN.) When you apply an FCoE map on a port, FCoE is enabled on the port. All non-FCoE traffic is dropped on an FCoE VLAN.
FIP	FCoE Initialization Protocol: Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP is used by server CNAs to discover

Term	Description
	an upstream FCoE switch operating as an FCF. FIP keepalive messages maintain the connection between an FCoE initiator and an FCF.
NPIV	N-port identifier virtualization: The capability to map multiple FCoE links from downstream ports to a single upstream FC link.
principal switch	The switch in a fabric with the lowest domain number. The principal switch accesses the master name database and the zone/zone set database.

DCB Maps

A Data Center Bridging (DCB) map is used to configure DCB functionality, such as PFC and ETS, on the Aggregator with the Ethernet ports that support CEE traffic and are DCBx-enabled, by default.

By default, no PFC and ETS settings in a DCB map are applied to the Aggregator with the Ethernet ports when they are enabled. On an Aggregator with the NPG, you must configure PFC and ETS parameters in a DCB map and then apply the map to server-facing Ethernet ports.

FCoE Maps

An FCoE map is used to identify the SAN fabric to which FCoE storage traffic is sent. Using an FCoE map, an Aggregator with the NPG operates as an FCoE-FC bridge between an FC SAN and FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric.

An FCoE map applies the following parameters on server-facing Ethernet and fabric-facing FC ports on the Aggregator:

- The dedicated FCoE VLAN used to transport FCoE storage traffic.
- The FC-MAP value used to generate a fabric-provided MAC address.
- The association between the FCoE VLAN ID and FC fabric ID where the desired storage arrays are installed. Each Fibre Channel fabric serves as an isolated SAN topology within the same physical network.
- The priority used by a server to select an upstream FCoE forwarder (FCF priority).
- FIP keepalive (FKA) advertisement timeout.

NOTE:

In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN must be unique. Use one FCoE map to access one SAN fabric. You cannot use the same FCoE map to access different fabrics.

When you configure an Aggregator with the NPG, FCoE transit with FIP snooping is automatically enabled and configured using the parameters in the FCoE map applied to server-facing Ethernet and fabric-facing FC interfaces.

After you apply an FCoE map on an FC port, when you enable the port (`no shutdown`), the NPG starts sending FIP multicast advertisements on behalf of the FC port to downstream servers in order to advertise the availability of a new FCF port on the FCoE VLAN. The FIP advertisement also contains a keepalive message to maintain connectivity between a SAN fabric and downstream servers.


Configuring an NPIV Proxy Gateway

Prerequisite: Before you configure an NPIV proxy gateway (NPG) on an Aggregator, ensure that the following features are enabled.

- DCB is enabled by default on the Aggregator.
- Autonegotiated DCBx is enabled for converged traffic by default with the Ethernet ports on all Aggregators.
- FCoE transit with FIP snooping is automatically enabled when you configure Fibre Channel on the Aggregator.

To configure an NPG operation on an Aggregator, follow these general configuration steps:

1. Enabling Fibre Channel Capability on the Switch
2. Creating a DCB map
3. Applying a DCB map on server-facing Ethernet ports
4. Creating an FCoE VLAN
5. Creating an FCoE map
6. Applying an FCoE map on server-facing Ethernet ports
7. Applying an FCoE Map on fabric-facing FC ports

 **NOTE:** All these configurations are available only in PMUX mode and you cannot perform these configurations in Standalone mode.

Default Configurations in Standalone mode

By default, the following configurations are set in Standalone mode:

1. All the FC port are applied with the default FCoE map.
2. All the 10G server facing ports with the DCBX frames will have the default FCoE map and default DCB map assigned, and without the DCBX frames, DCB_MAP_PFC_OFF will be applied.

Default DCB map

```
Dell(conf)#do show qos dcb-map SAN_DCB_MAP
```

```
-----
```

```
State      :Complete
```

```
PfcMode:ON
```

```
-----
```

```
PG:0 TSA:ETS BW:30 PFC:OFF
```

```
Priorities:0 1 2 5 6 7
```

```
PG:1 TSA:ETS BW:30 PFC:OFF
```

```
Priorities:4
```

```
PG:2 TSA:ETS BW:40 PFC:ON
```

```
Priorities:3
```

Default FCoE map

```
Dell(conf)#do show fcoe-map
```

```

Fabric Name      SAN_FABRIC
Fabric Id       1002
Vlan Id        1002
Vlan priority   3
FC-MAP         0efc00
FKA-ADV-Period 8
Fcf Priority    128
Config-State   ACTIVE
Oper-State     UP
Members
Fc 0/9
Te 0/4

```

DCB_MAP_PFC_OFF

```
Dell(conf)#do show qos dcb-map DCB_MAP_PFC_OFF
```

```

-----
State      :In-Progress
PfcMode:OFF
-----

```

```
Dell(conf)#
```

Enabling Fibre Channel Capability on the Switch

Enable the Fibre Channel capability on an Aggregator that you want to configure as an NPG for the Fibre Channel protocol. When you enable Fibre Channel capability, FCoE transit with FIP snooping is automatically enabled on all VLANs on the switch, using the default FCoE transit settings.

Task	Command	Command Mode
Enable the Fibre Channel capability on an Aggregator for the Fibre Channel protocol.	<code>feature fc</code>	CONFIGURATION

Creating a DCB Map

Configure the priority-based flow control (PFC) and enhanced traffic selection (ETS) settings in a DCB map before you apply them on downstream server-facing ports on an Aggregator.

Step	Task	Command	Command Mode
1	Create a DCB map to specify PFC and ETS settings for groups of dot1p priorities.	<code>dcb-map name</code>	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group. Configure whether the priority group traffic should be handled with strict-priority scheduling. The sum of all allocated bandwidth percentages must be 100 percent. Strict-priority traffic is serviced first. Afterward, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.	<code>priority-group group_num {bandwidth percentage strict-priority} pfc {on off}</code>	DCB MAP

Step	Task	Command	Command Mode
	<p>Restriction: You can enable PFC on a maximum of two priority queues.</p> <p>Repeat this step to configure PFC and ETS traffic handling for each priority group, for example: <code>priority-group 0 bandwidth 60 pfc off</code> <code>priority-group 1 bandwidth 20 pfc on</code> <code>priority-group 2 bandwidth 20 pfc on</code> <code>priority-group 4 strict-priority pfc off</code></p>		
3	<p>Specify the priority group ID number to handle VLAN traffic for each dot1p class-of-service: 0 through 7. Leave a space between each priority group number. For example, <code>priority-pgid 0 0 0 1 2 4 4 4</code> where dot1p priorities 0, 1, and 2 are mapped to priority group 0; dot1p priority 3 is mapped to priority group 1; dot1p priority 4 is mapped to priority group 2; dot1p priorities 5, 6, and 7 are mapped to priority group 4.</p> <p>All priorities that map to the same egress queue must be in the same priority group.</p>	<pre>priority-pgid dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num</pre>	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (no `priority pgid` command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share port bandwidth.
- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, create a new DCB map with the desired PFC and ETS settings, and apply the new map to the interfaces to override the previous DCB map settings. Then, delete the original dot1p priority-to-priority group mapping.
- If you delete the dot1p priority-to-priority group mapping (no `priority pgid` command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt the network operation.

Applying a DCB Map on Server-facing Ethernet Ports

You can apply a DCB map only on a physical Ethernet interface and can apply only one DCB map per interface.

Step	Task	Command	Command Mode
1	Enter CONFIGURATION mode on a server-facing port or port channel to apply a DCB map.	<pre>interface {tengigabitEthernet slot/port }</pre>	CONFIGURATION

Step	Task	Command	Command Mode
	You cannot apply a DCB map on a port channel. However, you can apply a DCB map on the ports that are members of the port channel.		
2	Apply the DCB map on an Ethernet port or port channel. The port is configured with the PFC and ETS settings in the DCB map, for example:	<code>dcb-map name</code>	INTERFACE
	<pre>Dell# interface tengigabitEthernet 0/0 Dell(config-if-te-0/0)# dcb-map SAN_DCBI</pre>		
	Repeat this step to apply a DCB map to more than one port or port channel.		

Creating an FCoE VLAN

Create a dedicated VLAN to send and receive Fibre Channel traffic over FCoE links between servers and a fabric over an NPG. The NPG receives FCoE traffic and forwards decapsulated FC frames over FC links to SAN switches in a specified fabric.

Step	Task	Command	Command Mode
1	Create the dedicated VLAN for FCoE traffic. Range: 2–4094. VLAN 1002 is commonly used to transmit FCoE traffic.	<code>interface vlan vlan- id</code>	CONFIGURATION

When you apply an FCoE map to an Ethernet port, the port is automatically configured as a tagged member of the FCoE VLAN.

Creating an FCoE Map

An FCoE map consists of:

- An association between the dedicated VLAN, used to carry FCoE traffic, and the SAN fabric where the storage arrays are installed. Use a separate FCoE VLAN for each fabric to which the FCoE traffic is forwarded. Any non-FCoE traffic sent on a dedicated FCoE VLAN is dropped.
- The FC-MAP value, used to generate the fabric-provided MAC address (FPMA). The FPMA is used by servers to transmit FCoE traffic to the fabric. You can associate an FC-MAP with only one FCoE VLAN and conversely, associate an FCoE VLAN with only one FC-MAP.
- FCF priority, the priority used by a server CNA to select an upstream FCoE forwarder (FCF).
- FIP keepalive (FKA) advertisement timeout.

The values for the FCoE VLAN, fabric ID and FC-MAP must be unique. Apply an FCoE map on downstream server-facing Ethernet ports and upstream fabric-facing Fibre Channel ports.

Step	Task	Command	Command Mode
1	Create an FCoE map that contains parameters used in the communication between servers and a SAN fabric.	<code>fcoe-map map-name</code>	CONFIGURATION
2	Configure the association between the dedicated VLAN and the fabric where the desired storage arrays are installed. The fabric and VLAN ID numbers must be the same. Fabric and VLAN ID range: 2–4094. For example: <code>fabric id 10 vlan 10</code>	<code>fabric-id fabric-num</code> <code>vlan vlan-id</code>	FCoE MAP
3	Add a text description of the settings in the FCoE map. Maximum: 32 characters.	<code>description text</code>	FCoE MAP
4	Specify the FC-MAP value used to generate a fabric-provided MAC address, which is required to send FCoE traffic from a server on the FCoE VLAN to the FC fabric specified in Step 2. Enter a unique MAC address prefix as the FC-MAP value for each fabric. Range: 0EFC00–0EFCFF. Default: None.	<code>fc-map fc-map-value</code>	FCoE MAP
5	Configure the priority used by a server CNA to select the FCF for a fabric login (FLOGI). Range: 1–255. Default: 128.	<code>fcf-priority</code> <code>priority</code>	FCoE MAP
6	Enable the monitoring FIP keepalive messages (if it is disabled) to detect if other FCoE devices are reachable. Default: FIP keepalive monitoring is enabled.	<code>keepalive</code>	FCoE MAP
7	Configure the time interval (in seconds) used to transmit FIP keepalive advertisements. Range: 8–90 seconds. Default: 8 seconds.	<code>fka-adv-period</code> <code>seconds</code>	FCoE MAP

Applying an FCoE Map on Server-facing Ethernet Ports

You can apply multiple FCoE maps on an Ethernet port or port channel. When you apply an FCoE map on a server-facing port or port channel:

- The port is configured to operate in hybrid mode (accept both tagged and untagged VLAN frames).
- The associated FCoE VLAN is enabled on the port or port channel.

When you enable a server-facing Ethernet port, the servers respond to the FIP advertisements by performing FLOGIs on upstream virtualized FCF ports. The NPG forwards the FLOGIs as FDISC messages to a SAN switch.

Step	Task	Command	Command Mode
1	Configure a server-facing Ethernet port or port channel with an FCoE map.	<code>interface {tengigabitEthernet slot/port port- channel num}</code>	CONFIGURATION
2	Apply the FCoE/FC configuration in an FCoE map on the Ethernet port. Repeat this step to apply an FCoE map to more than one port, for example: Dell# interface tengigabitEthernet 0/0 Dell(config-if-te-0/0)# fcoe-map SAN_FABRIC_A Dell# interface port-channel 3 Dell(config-if-te-0/0)# dcb-map SAN_DCB1 Dell(config-if-po-3)# fcoe-map SAN_FABRIC_A Dell# interface fortygigabitEthernet 0/48 Dell(config-if-fo-0/0)# fcoe-map SAN_FABRIC_A	<code>fcoe-map map-name</code>	INTERFACE or INTERFACE PORT_CHANNEL
3	Enable the port for FCoE transmission using the map settings.	<code>no shutdown</code>	INTERFACE

Applying an FCoE Map on Fabric-facing FC Ports

The Aggregator, with the FC ports, are configured by default to operate in N port mode to connect to an F port on an FC switch in a fabric. You can apply only one FCoE map on an FC port.

When you apply an FCoE map on a fabric-facing FC port, the FC port becomes part of the FCoE fabric, whose settings in the FCoE map are configured on the port and exported to downstream server CNA ports.

Each Aggregator, with the FC port, is associated with an Ethernet MAC address (FCF MAC address). When you enable a fabric-facing FC port, the FCoE map applied to the port starts sending FIP multicast advertisements using the parameters in the FCoE map over server-facing Ethernet ports. A server sees the FC port, with its applied FCoE map, as an FCF port.

Step	Task	Command	Command Mode
1	Configure a fabric-facing FC port.	<code>interface fibrechannel slot/ port</code>	CONFIGURATION
2	Apply the FCoE and FC fabric configurations in an FCoE map to the port. Repeat this step to apply an FCoE map to more than one FC port, for example: Dell# interface fi 0/9 Dell(config-if-fc-0/9)# fabric SAN_FABRIC_A	<code>fabric map-name</code>	INTERFACE FIBRE_CHANNEL

Step	Task	Command	Command Mode
3	Enable the port for FC transmission.	no shutdown	INTERFACE FIBRE_CHANNEL

Important Points to Remember

You can apply a DCB or FCoE map to a range of Ethernet or Fibre Channel interfaces by using the `interface range` command; for example:

```
Dell(config)# interface range tengigabitEthernet 1/12 - 23 , tengigabitEthernet 2/24 - 35
```

```
Dell(config)# interface range fibrechannel 0/0 - 3 , fibrechannel 0/8 - 11
```

Enter the keywords `interface range` followed by an interface type and port range. A port range must contain spaces before and after the dash. Separate each interface type and port range with a space, comma, and space as shown in the preceding examples.

Sample Configuration

1. Configure a DCB map with PFC and ETS settings:

```
Dell(config)# dcb-map SAN_DCB_MAP
```

```
Dell(config-dcbx-name)# priority-group 0 bandwidth 60 pfc off
```

```
Dell(config-dcbx-name)# priority-group 1 bandwidth 20 pfc on
```

```
Dell(config-dcbx-name)# priority-group 2 bandwidth 20 pfc on
```

```
Dell(config-dcbx-name)# priority-group 4 strict-priority pfc off
```

```
Dell(conf-dcbx-name)# priority-pgid 0 0 0 1 2 4 4 4
```

2. Apply the DCB map on a downstream (server-facing) Ethernet port:

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

```
Dell(config-if-te-0/0)#dcb-map SAN_DCB_MAP
```

3. Create the dedicated VLAN to be used for FCoE traffic:

```
Dell(conf)#interface vlan 1002
```

4. Configure an FCoE map to be applied on downstream (server-facing) Ethernet and upstream (core-facing) FC ports:

```
Dell(config)# fcoe-map SAN_FABRIC_A
```

```
Dell(config-fcoe-name)# fabric-id 1002 vlan 1002
```

```
Dell(config-fcoe-name)# description "SAN_FABRIC_A"
```


```
Dell(config-fcoe-name)# fc-map 0efc00
```

- ```
Dell(config-fcoe-name)# keepalive
```
- ```
Dell(config-fcoe-name)# fcf-priority 128
```
- ```
Dell(config-fcoe-name)# fka-adv-period 8
```
5. Enable an upstream FC port:
- ```
Dell(config)# interface fibrechannel 0/0
```
- ```
Dell(config-if-fc-0)# no shutdown
```
6. Enable a downstream Ethernet port:
- ```
Dell(config)#interface tengigabitEthernet 0/0
```
- ```
Dell(conf-if-te-0)# no shutdown
```

## Displaying NPIV Proxy Gateway Information

To display information on the NPG operation, use the show commands in the following table:

**Table 23. Displaying NPIV Proxy Gateway Information**

| Command                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>show interfaces status</code>           | Displays the operational status of Ethernet and Fibre Channel interfaces on the Aggregator with the NPG.<br><br> <b>NOTE:</b> Although the show interface status command displays the Fiber Channel (FC) interfaces with the abbreviated label of 'Fc' in the output, if you attempt to specify a FC interface by using the <code>interface fc</code> command in the CLI interface, an error message is displayed. You must configure FC interfaces by using the <code>interface fi</code> command in CONFIGURATION mode. |
| <code>show fcoe-map [brief   map-name]</code> | Displays the Fibre Channel and FCoE configuration parameters in FCoE maps. Enter the <b>brief</b> keyword to display an overview of currently configured FCoE maps.<br><br>Enter the name of an FCoE map to display the FC and FCoE parameters configured in the map to be applied on the Aggregator with the FC ports.                                                                                                                                                                                                                                                                                     |
| <code>show qos dcb-map map-name</code>        | Displays configuration parameters in a specified DCB map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>show npiv devices [brief]</code>        | Displays information on FCoE and FC devices currently logged in to the NPG.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <code>show fc switch</code>                   | Displays the FC mode of operation and worldwide node (WWN) name of an Aggregator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

### show interfaces status Command Example

```
Dell# show interfaces status
Port Description Status Speed Duplex Vlan
```



```

Te 0/1 Up 10000 Mbit Full 1-4094
Te 0/2 Down Auto Auto 1-1001,1003-4094
Te 0/3 Up 10000 Mbit Full 1-1001,1003-4094
Te 0/4 Down Auto Auto 1-1001,1003-4094
Te 0/5 Up 10000 Mbit Full 1-4094
Te 0/6 Up 10000 Mbit Full 1-4094
Te 0/7 Up 10000 Mbit Full 1-4094
Te 0/8 toB300 Down Auto Auto 1-1001,1003-4094
Fc 0/9 Up 8000 Mbit Full --
Fc 0/10 Up 8000 Mbit Full --
Te 0/11 Down Auto Auto --
Te 0/12 Down Auto Auto --

```

**Table 24. show interfaces status Field Descriptions**

| Field       | Description                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port        | Server-facing 10GbE Ethernet (Te), or fabric-facing Fibre Channel (FC) port with <i>slot/port</i> information.                                                                                                                                                                                                                                                                          |
| Description | Text description of port.                                                                                                                                                                                                                                                                                                                                                               |
| Status      | Operational status of port:<br><br>Ethernet ports - up (transmitting FCoE and LAN storage traffic) or down (not transmitting traffic).<br><br>Fibre Channel ports - up (link is up and transmitting FC traffic) or down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer SW port), or removed (port has been shut down). |
| Speed       | Transmission speed (in Megabits per second) of Ethernet and FC ports, including auto-negotiated speed (Auto).                                                                                                                                                                                                                                                                           |
| Duplex      | Data transmission mode: Full (allows communication in both directions at the same time), Half (allows communication in both directions but not at the same time), Auto (auto-negotiated transmission).                                                                                                                                                                                  |
| VLAN        | VLAN IDs of the VLANs in which the port is a member.                                                                                                                                                                                                                                                                                                                                    |

## show fcoe-map Command Examples

```

Dell# show fcoe-map brief
Fabric-Name Fabric-Id Vlan-Id FC-MAP FCF-Priority Config-
State Oper-State
fid_1003 1003 1003 0efc03 128
ACTIVE UP
fid_1004 1004 1004 0efc04 128
ACTIVE DOWN

```

```
Dell# show fcoe-map fid_1003
```

```

Fabric Name fid_1003
Fabric Id 1003
Vlan Id 1003
Vlan priority 3
FC-MAP 0efc03
FKA-ADV-Period 8

```

```

Fcf Priority 128
Config-State ACTIVE
Oper-State UP
Members
Fc 0/9
Te 0/11 Te 0/12

```

**Table 25. show fcoe-map Field Descriptions**

| Field          | Description                                                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric-Name    | Name of a SAN fabric.                                                                                                                                                                                                                                           |
| Fabric ID      | The ID number of the SAN fabric to which FC traffic is forwarded.                                                                                                                                                                                               |
| VLAN ID        | The dedicated VLAN used to transport FCoE storage traffic between servers and a fabric over the NPG. The configured VLAN ID must be the same as the fabric ID.                                                                                                  |
| VLAN priority  | FCoE traffic uses VLAN priority 3. This setting is not user-configurable.                                                                                                                                                                                       |
| FC-MAP         | FCoE MAC-address prefix value - The unique 24-bit MAC address prefix that identifies a fabric.                                                                                                                                                                  |
| FKA-ADV-period | Time interval (in seconds) used to transmit FIP keepalive advertisements.                                                                                                                                                                                       |
| FCF Priority   | The priority used by a server to select an upstream FCoE forwarder.                                                                                                                                                                                             |
| Config-State   | Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).          |
| Oper-State     | Operational status of the link to the fabric: up (link is up and transmitting FC traffic), down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer FC port), or removed (port has been shut down). |
| Members        | Aggregator with the FC ports, which are members of the dedicated FCoE VLAN that carries storage traffic to the specified fabric.                                                                                                                                |

## show qos dcb-map Command Examples

```

Dell# show qos dcb-map dcbmap2

State :Complete
PfcMode:ON

PG:0 TSA:ETS BW:50 PFC:OFF
Priorities:0 1 2 4 5 6 7

PG:1 TSA:ETS BW:50 PFC:ON
Priorities:3

```

**Table 26. show qos dcb-map Field Descriptions**

| Field      | Description                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State      | Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured. |
| PFC Mode   | PFC configuration in the DCB map: On (enabled) or Off.                                                                                                                 |
| PG         | Priority group configured in the DCB map.                                                                                                                              |
| TSA        | Transmission scheduling algorithm used in the DCB map: Enhanced Transmission Selection (ETS).                                                                          |
| BW         | Percentage of bandwidth allocated to the priority group.                                                                                                               |
| PFC        | PFC setting for the priority group: On (enabled) or Off.                                                                                                               |
| Priorities | 802.1p priorities configured in the priority group.                                                                                                                    |

**show npiv devices brief Command Example**

```
Dell# show npiv devices brief
```

```
Total NPIV Devices = 2
```

```

ENode-Intf ENode-WWPN FCoE-Vlan Fabric-Intf Fabric-Map
LoginMethod Status

Te 0/11 20:01:00:10:18:f1:94:20 1003 Fc 0/9
fid_1003 FLOGI LOGGED_IN
Te 0/12 10:00:00:00:c9:d9:9c:c5 1003 Fc 0/10
fid_1003 FDISC LOGGED_IN
```

**Table 27. show npiv devices brief Field Descriptions**

| Field              | Description                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------|
| Total NPIV Devices | Number of downstream ENodes connected to a fabric over the Aggregator with the NPG.                               |
| ENode-Intf         | Aggregator with the Ethernet interface ( <i>slot/port</i> ) to which a server CNA is connected.                   |
| ENode-WWPN         | Worldwide port name (WWPN) of a server CNA port.                                                                  |
| FCoE-Vlan          | VLAN ID of the dedicated VLAN used to transmit FCoE traffic to and from the fabric.                               |
| Fabric-Intf        | Fabric-facing Fibre Channel port ( <i>slot/port</i> ) on which FC traffic is transmitted to the specified fabric. |

| Field        | Description                                                                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric-Map   | Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.                                                                          |
| Login Method | Method used by the server CNA to log in to the fabric; for example: FLOGI - ENode logged in using a fabric login (FLOGI). FDISC - ENode logged in using a fabric discovery (FDISC). |
| Status       | Operational status of the link between a server CNA port and a SAN fabric: Logged In - Server has logged in to the fabric and is able to transmit FCoE traffic.                     |

## show npiv devices Command Example

```
Dell# show npiv devices
ENode[0]:
ENode MAC : 00:10:18:f1:94:21
ENode Intf : Te 0/11
FCF MAC : 5c:f9:dd:ef:10:c8
Fabric Intf : Fc 0/9
FCoE Vlan : 1003
Fabric Map : fid_1003
ENode WWPN : 20:01:00:10:18:f1:94:20
ENode WWNN : 20:00:00:10:18:f1:94:21
FCoE MAC : 0e:fc:03:01:02:01
FC-ID : 01:02:01
LoginMethod : FLOGI
Secs : 5593
Status : LOGGED_IN

ENode[1]:
ENode MAC : 00:10:18:f1:94:22
ENode Intf : Te 0/12
FCF MAC : 5c:f9:dd:ef:10:c9
Fabric Intf : Fc 0/10
FCoE Vlan : 1003
Fabric Map : fid_1003
ENode WWPN : 10:00:00:00:c9:d9:9c:cb
ENode WWNN : 10:00:00:00:c9:d9:9c:cd
FCoE MAC : 0e:fc:03:01:02:02
FC-ID : 01:02:01
LoginMethod : FDISC
Secs : 5593
Status : LOGGED_IN
```

**Table 28. show npiv devices Field Descriptions**

| Field                 | Description                                                                                                     |
|-----------------------|-----------------------------------------------------------------------------------------------------------------|
| ENode <i>[number]</i> | Server CNA that has successfully logged in to a fabric over an Aggregator with the Ethernet port in ENode mode. |
| Enode MAC             | MAC address of a server CNA port.                                                                               |
| Enode Intf            | Port number of a server-facing Ethernet port operating in ENode mode.                                           |

| Field       | Description                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FCF MAC     | Fibre Channel forwarder MAC: MAC address of Aggregator with the FCF interface.                                                                                                                                                                              |
| Fabric Intf | Fabric-facing Aggregator with the Fibre Channel port ( <i>slot/port</i> ) on which FCoE traffic is transmitted to the specified fabric.                                                                                                                     |
| FCoE VLAN   | ID of the dedicated VLAN used to transmit FCoE traffic from a server CNA to a fabric and configured on both the server-facing Aggregator with the server CNA port.                                                                                          |
| Fabric Map  | Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.                                                                                                                                                  |
| Enode WWPN  | Worldwide port name of the server CNA port.                                                                                                                                                                                                                 |
| Enode WWNN  | Worldwide node name of the server CNA.                                                                                                                                                                                                                      |
| FCoE MAC    | Fabric-provided MAC address (FPMA). The FPMA consists of the FC-MAP value in the FCoE map and the FC-ID provided by the fabric after a successful FLOGI. In the FPMA, the most significant bytes are the FC-MAP; the least significant bytes are the FC-ID. |
| FC-ID       | FC port ID provided by the fabric.                                                                                                                                                                                                                          |
| LoginMethod | Method used by the server CNA to log in to the fabric; for example, FLOGI or FDISC.                                                                                                                                                                         |
| Secs        | Number of seconds that the fabric connection is up.                                                                                                                                                                                                         |
| State       | Status of the fabric connection: logged in.                                                                                                                                                                                                                 |

## show fc switch Command Example

```
Dell# show fc switch
Switch Mode : NPG
Switch WWN : 10:00:5c:f9:dd:ef:10:c0
Dell#
```


**Table 29. show fc switch Command Description**

| Field       | Description                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------|
| Switch Mode | Fibre Channel mode of operation of an Aggregator. Default: NPG (configured as an NPIV proxy gateway).           |
| Switch WWN  | Factory-assigned worldwide node (WWN) name of the Aggregator. The Aggregator WWN name is not user-configurable. |

## Displaying NPIV Proxy Gateway Information

To display information on the NPG operation, use the show commands in the following table:

**Table 30. Displaying NPIV Proxy Gateway Information**

| Command                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>show interfaces status</code>           | Displays the operational status of Ethernet and Fibre Channel interfaces on the Aggregator with the NPG.<br><br> <b>NOTE:</b> Although the <code>show interface status</code> command displays the Fiber Channel (FC) interfaces with the abbreviated label of 'Fc' in the output, if you attempt to specify a FC interface by using the <code>interface fc</code> command in the CLI interface, an error message is displayed. You must configure FC interfaces by using the <code>interface fi</code> command in CONFIGURATION mode. |
| <code>show fcoe-map [brief   map-name]</code> | Displays the Fibre Channel and FCoE configuration parameters in FCoE maps. Enter the <b>brief</b> keyword to display an overview of currently configured FCoE maps.<br><br>Enter the name of an FCoE map to display the FC and FCoE parameters configured in the map to be applied on the Aggregator with the FC ports.                                                                                                                                                                                                                                                                                                 |
| <code>show qos dcb-map map-name</code>        | Displays configuration parameters in a specified DCB map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <code>show npiv devices [brief]</code>        | Displays information on FCoE and FC devices currently logged in to the NPG.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <code>show fc switch</code>                   | Displays the FC mode of operation and worldwide node (WWN) name of an Aggregator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

### show interfaces status Command Example

```
Dell# show interfaces status
Port Description Status Speed Duplex Vlan
Te 0/1 Up 10000 Mbit Full 1-4094
Te 0/2 Down Auto Auto 1-1001,1003-4094
Te 0/3 Up 10000 Mbit Full 1-1001,1003-4094
Te 0/4 Down Auto Auto 1-1001,1003-4094
Te 0/5 Up 10000 Mbit Full 1-4094
Te 0/6 Up 10000 Mbit Full 1-4094
Te 0/7 Up 10000 Mbit Full 1-4094
Te 0/8 toB300 Down Auto Auto 1-1001,1003-4094
Fc 0/9 Up 8000 Mbit Full --
Fc 0/10 Up 8000 Mbit Full --
Te 0/11 Down Auto Auto --
Te 0/12 Down Auto Auto --
```

**Table 31. show interfaces status Field Descriptions**

| Field       | Description                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------|
| Port        | Server-facing 10GbE Ethernet (Te), or fabric-facing Fibre Channel (FC) port with <i>slot/port</i> information. |
| Description | Text description of port.                                                                                      |
| Status      | Operational status of port:                                                                                    |

Ethernet ports - up (transmitting FCoE and LAN storage traffic) or down (not transmitting traffic).

Fibre Channel ports - up (link is up and transmitting FC traffic) or down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer SW port), or removed (port has been shut down).

|        |                                                                                                                                                                                                        |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Speed  | Transmission speed (in Megabits per second) of Ethernet and FC ports, including auto-negotiated speed (Auto).                                                                                          |
| Duplex | Data transmission mode: Full (allows communication in both directions at the same time), Half (allows communication in both directions but not at the same time), Auto (auto-negotiated transmission). |
| VLAN   | VLAN IDs of the VLANs in which the port is a member.                                                                                                                                                   |

## show fcoe-map Command Examples

```
Dell# show fcoe-map brief
Fabric-Name Fabric-Id Vlan-Id FC-MAP FCF-Priority Config-
State Oper-State
fid_1003 1003 1003 0efc03 128
ACTIVE UP
fid_1004 1004 1004 0efc04 128
ACTIVE DOWN
```

```
Dell# show fcoe-map fid_1003
```

```
Fabric Name fid_1003
Fabric Id 1003
Vlan Id 1003
Vlan priority 3
FC-MAP 0efc03
FKA-ADV-Period 8
Fcf Priority 128
Config-State ACTIVE
Oper-State UP
Members
Fc 0/9
Te 0/11 Te 0/12
```

**Table 32. show fcoe-map Field Descriptions**

| Field         | Description                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric-Name   | Name of a SAN fabric.                                                                                                                                          |
| Fabric ID     | The ID number of the SAN fabric to which FC traffic is forwarded.                                                                                              |
| VLAN ID       | The dedicated VLAN used to transport FCoE storage traffic between servers and a fabric over the NPG. The configured VLAN ID must be the same as the fabric ID. |
| VLAN priority | FCoE traffic uses VLAN priority 3. This setting is not user-configurable.                                                                                      |

|                |                                                                                                                                                                                                                                                                 |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FC-MAP         | FCoE MAC-address prefix value - The unique 24-bit MAC address prefix that identifies a fabric.                                                                                                                                                                  |
| FKA-ADV-period | Time interval (in seconds) used to transmit FIP keepalive advertisements.                                                                                                                                                                                       |
| FCF Priority   | The priority used by a server to select an upstream FCoE forwarder.                                                                                                                                                                                             |
| Config-State   | Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).          |
| Oper-State     | Operational status of the link to the fabric: up (link is up and transmitting FC traffic), down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer FC port), or removed (port has been shut down). |
| Members        | Aggregator with the FC ports, which are members of the dedicated FCoE VLAN that carries storage traffic to the specified fabric.                                                                                                                                |

## show qos dcb-map Command Examples

```
Dell# show qos dcb-map dcbmap2
```

```
State :Complete
PfcMode:ON

PG:0 TSA:ETS BW:50 PFC:OFF
Priorities:0 1 2 4 5 6 7

PG:1 TSA:ETS BW:50 PFC:ON
Priorities:3
```

**Table 33. show qos dcb-map Field Descriptions**

| Field      | Description                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State      | Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured. |
| PFC Mode   | PFC configuration in the DCB map: On (enabled) or Off.                                                                                                                 |
| PG         | Priority group configured in the DCB map.                                                                                                                              |
| TSA        | Transmission scheduling algorithm used in the DCB map: Enhanced Transmission Selection (ETS).                                                                          |
| BW         | Percentage of bandwidth allocated to the priority group.                                                                                                               |
| PFC        | PFC setting for the priority group: On (enabled) or Off.                                                                                                               |
| Priorities | 802.1p priorities configured in the priority group.                                                                                                                    |



## show npiv devices brief Command Example

```
Dell# show npiv devices brief
```

```
Total NPIV Devices = 2
```

```


ENode-Intf ENode-WWPN FCoE-Vlan Fabric-Intf Fabric-Map
LoginMethod Status

Te 0/11 20:01:00:10:18:f1:94:20 1003 Fc 0/9
fid_1003 FLOGI LOGGED_IN
Te 0/12 10:00:00:00:c9:d9:9c:cb 1003 Fc 0/10
fid_1003 FDISC LOGGED_IN
```

**Table 34. show npiv devices brief Field Descriptions**

| Field              | Description                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total NPIV Devices | Number of downstream ENodes connected to a fabric over the Aggregator with the NPG.                                                                                                 |
| ENode-Intf         | Aggregator with the Ethernet interface ( <i>slot/port</i> ) to which a server CNA is connected.                                                                                     |
| ENode-WWPN         | Worldwide port name (WWPN) of a server CNA port.                                                                                                                                    |
| FCoE-Vlan          | VLAN ID of the dedicated VLAN used to transmit FCoE traffic to and from the fabric.                                                                                                 |
| Fabric-Intf        | Fabric-facing Fibre Channel port ( <i>slot/port</i> ) on which FC traffic is transmitted to the specified fabric.                                                                   |
| Fabric-Map         | Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.                                                                          |
| Login Method       | Method used by the server CNA to log in to the fabric; for example: FLOGI - ENode logged in using a fabric login (FLOGI). FDISC - ENode logged in using a fabric discovery (FDISC). |
| Status             | Operational status of the link between a server CNA port and a SAN fabric: Logged In - Server has logged in to the fabric and is able to transmit FCoE traffic.                     |

## show npiv devices Command Example

```
Dell# show npiv devices
ENode[0]:
ENode MAC : 00:10:18:f1:94:21
ENode Intf : Te 0/11
FCF MAC : 5c:f9:dd:ef:10:c8
Fabric Intf : Fc 0/9
FCoE Vlan : 1003
Fabric Map : fid_1003
ENode WWPN : 20:01:00:10:18:f1:94:20
```

```

ENode WWNN : 20:00:00:10:18:f1:94:21
FCoE MAC : 0e:fc:03:01:02:01
FC-ID : 01:02:01
LoginMethod : FLOGI
Secs : 5593
Status : LOGGED_IN

```

```

ENode[1]:
ENode MAC : 00:10:18:f1:94:22
ENode Intf : Te 0/12
FCF MAC : 5c:f9:dd:ef:10:c9
Fabric Intf : Fc 0/10
FCoE Vlan : 1003
Fabric Map : fid_1003
ENode WWPN : 10:00:00:00:c9:d9:9c:cb
ENode WWNN : 10:00:00:00:c9:d9:9c:cd
FCoE MAC : 0e:fc:03:01:02:02
FC-ID : 01:02:01
LoginMethod : FDISC
Secs : 5593
Status : LOGGED_IN

```

**Table 35. show npiv devices Field Descriptions**

| Field                 | Description                                                                                                                                                                                                                                                 |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENode <i>[number]</i> | Server CNA that has successfully logged in to a fabric over an Aggregator with the Ethernet port in ENode mode.                                                                                                                                             |
| Enode MAC             | MAC address of a server CNA port.                                                                                                                                                                                                                           |
| Enode Intf            | Port number of a server-facing Ethernet port operating in ENode mode.                                                                                                                                                                                       |
| FCF MAC               | Fibre Channel forwarder MAC: MAC address of Aggregator with the FCF interface.                                                                                                                                                                              |
| Fabric Intf           | Fabric-facing Aggregator with the Fibre Channel port ( <i>slot/port</i> ) on which FCoE traffic is transmitted to the specified fabric.                                                                                                                     |
| FCoE VLAN             | ID of the dedicated VLAN used to transmit FCoE traffic from a server CNA to a fabric and configured on both the server-facing Aggregator with the server CNA port.                                                                                          |
| Fabric Map            | Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.                                                                                                                                                  |
| Enode WWPN            | Worldwide port name of the server CNA port.                                                                                                                                                                                                                 |
| Enode WWNN            | Worldwide node name of the server CNA.                                                                                                                                                                                                                      |
| FCoE MAC              | Fabric-provided MAC address (FPMA). The FPMA consists of the FC-MAP value in the FCoE map and the FC-ID provided by the fabric after a successful FLOGI. In the FPMA, the most significant bytes are the FC-MAP; the least significant bytes are the FC-ID. |
| FC-ID                 | FC port ID provided by the fabric.                                                                                                                                                                                                                          |

| Field       | Description                                                                         |
|-------------|-------------------------------------------------------------------------------------|
| LoginMethod | Method used by the server CNA to log in to the fabric; for example, FLOGI or FDISC. |
| Secs        | Number of seconds that the fabric connection is up.                                 |
| State       | Status of the fabric connection: logged in.                                         |

## show fc switch Command Example

```
Dell# show fc switch
Switch Mode : NPG
Switch WWN : 10:00:5c:f9:dd:ef:10:c0
Dell#
```

**Table 36. show fc switch Command Description**

| Field       | Description                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------|
| Switch Mode | Fibre Channel mode of operation of an Aggregator. Default: NPG (configured as an NPIV proxy gateway).           |
| Switch WWN  | Factory-assigned worldwide node (WWN) name of the Aggregator. The Aggregator WWN name is not user-configurable. |

## Upgrade Procedures

To find the upgrade procedures, go to the *Dell Networking OS Release Notes* for your system type to see all the requirements needed to upgrade to the desired Dell Networking OS version. To upgrade your system type, follow the procedures in the *Dell Networking OS Release Notes*.

### Get Help with Upgrades

Direct any questions or concerns about the Dell Networking OS upgrade procedures to the Dell Technical Support Center. You can reach Technical Support:

- On the web: <http://support.dell.com/>
- By email: Dell-Force10\_Technical\_Support@Dell.com
- By phone: US and Canada: 866.965.5800, International: 408.965.5800.

# Debugging and Diagnostics

This chapter contains the following sections:

- Debugging Aggregator Operation
- Software Show Commands
- Offline Diagnostics
- Trace Logs
- Show Hardware Commands

## Supported Modes

Standalone, PMUX, VLT

## Debugging Aggregator Operation

This section describes common troubleshooting procedures to use for error conditions that may arise during Aggregator operation.

### All interfaces on the Aggregator are operationally down

This section describes how you can troubleshoot the scenario in which all the interfaces are down.

**Symptom:** All Aggregator interfaces are down.

**Resolution:** Ensure the port channel 128 is up and that the Aggregator-facing port channel on the top-of-rack switch is correctly configured.

#### Steps to Take:

1. Verify that uplink port-channel 128 is up (show interfaces port-channel 128 brief command) and display the status of member ports (show uplink-state-group 1 detail command).

```
Dell#show interfaces port-channel 128 brief
Codes: L - LACP Port-channel

 LAG Mode Status Uptime Ports
L 128 L2L3 up 17:36:24 Te 0/1 (Up)
 Te 0/2 (Up)
 Te 0/3 (Dwn)
 Te 0/4 (Dwn)
 Te 0/5 (Up)
 Te 0/6 (Dwn)
 Te 0/7 (Dwn)
 Te 0/8 (Up)
```

```
Dell#show uplink-state-group 1 detail
```

```

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled

Uplink State Group : 1 Status: Enabled, Up
Defer Timer : 10 sec
Upstream Interfaces : Po 128 (Up)
Downstream Interfaces : Te 0/1 (Up) Te 0/2 (Up) Te 0/3 (Dwn) Te 0/4 (Dwn) Te
0/5 (Up)
 Te 0/6 (Dwn) Te 0/7 (Dwn) Te 0/8 (Up)

```

2. Verify that the downstream port channel in the top-of-rack switch that connect to the Aggregator is configured correctly.

## Broadcast, unknown multicast, and DLF packets switched at a very low rate

**Symptom:** Broadcast, unknown multicast, and DLF packets are switched at a very low rate.

By default, broadcast storm control is enabled on an Aggregator and rate limits the transmission of broadcast, unknown multicast, and DLF packets to 1Gbps. This default behavior is designed to avoid unnecessarily flooding these packets on all (4094) VLANs on all Aggregator interfaces (default configuration).

**Resolution:** Disable broadcast storm control globally on the Aggregator.

### Steps to Take:

1. Display the current status of broadcast storm control on the Aggregator (show io-aggregator broadcast storm-control status command).
 

```

Dell#show io-aggregator broadcast storm-control status

Storm-Control Enabled

Broadcast Traffic limited to 1000 Mbps

```
2. Disable broadcast storm control (no io-aggregator broadcast storm-control command) and re-display its status.
 

```

Dell#config terminal
Dell(conf)#no io-aggregator broadcast storm-control
Dell(conf)#end
Dell#show io-aggregator broadcast storm-control status

Storm-Control Disabled

```

## Flooded packets on all VLANs are received on a server

**Symptom:** All packets flooded on all VLANs on an Aggregator are received on a server, even if the server is configured as a member of only a subset of VLANs. This behavior happens because all Aggregator ports are, by default, members of all (4094) VLANs.

**Resolution:** Configure a port that is connected to the server with restricted VLAN membership.

### Steps to Take:

1. Display the current port mode for Aggregator L2 interfaces (show interfaces switchport interface command)..
 

```

Dell#show interfaces switchport tengigabitethernet 0/1

Codes: U - Untagged, T - Tagged
 x - Dot1x untagged, X - Dot1x tagged
 G - GVRP tagged, M - Trunk, H - VSN tagged
 i - Internal untagged, I - Internal tagged, v - VLT untagged, V -
VLT tagged

```

```
Name: TenGigabitEthernet 0/1
802.1QTagged: Hybrid
SMUX port mode: Auto VLANs enabled
Vlan membership:
Q Vlans
U 1
T 2-4094
```

```
Native VlanId: 1
```

2. Assign the port to a specified group of VLANs (vlan tagged command) and re-display the port mode status..

```
Dell(conf)#interface tengigabitethernet 0/1
Dell(conf-if-te-0/1)#vlan tagged 2-5,100,4010
Dell(conf-if-te-0/1)#
```

```
Dell#show interfaces switchport tengigabitethernet 0/1
```

```
Codes: U - Untagged, T - Tagged
 x - Dot1x untagged, X - Dot1x tagged
 G - GVRP tagged, M - Trunk, H - VSN tagged
 i - Internal untagged, I - Internal tagged, v - VLT untagged, V -
VLT tagged
```

```
Name: TenGigabitEthernet 0/1
802.1QTagged: Hybrid
SMUX port mode: Admin VLANs enabled
Vlan membership:
Q Vlans
U 1
T 2-5,100,4010
```

```
Native VlanId: 1
```

## Software show Commands

Use the show version and show system stack-unit 0 commands as a part of troubleshooting an Aggregator's software configuration.

**Table 37. Software show Commands**

| Command      | Description                                                                          |
|--------------|--------------------------------------------------------------------------------------|
| show version | Display the current version of Dell Networking OS software running on an Aggregator. |

### show version Command Example

```
Dell#show version
Dell Real Time Operating System Software
Dell Operating System Version: 2.0
Dell Application Software Version: 9-4(0-180)
Copyright (c) 1999-2014 by Dell Inc. All Rights Reserved.
Build Time: Sun Mar 30 20:15:19 PDT 2014
Build Path: /work.local/build/toolSpaces/tools05/E9-4-0/SW/SRC
Dell Networking OS uptime is 9 hour(s), 21 minute(s)

System image file is "dv-ci-stomp-tc-1-a1"
```

```
System Type: PE-FN-410S-IOA
Control Processor: MIPS RMI XLP with 2147483648 bytes of memory, core(s) 1.

128M bytes of boot flash memory.


 1 12-port GE/TE (FN)
 12 Ten GigabitEthernet/IEEE 802.3 interface(s)
Dell#
```

## Offline Diagnostics

The offline diagnostics test suite is useful for isolating faults and debugging hardware.

The diagnostics tests are grouped into three levels:

- **Level 0** — Level 0 diagnostics check for the presence of various components and perform essential path verifications. In addition, Level 0 diagnostics verify the identification registers of the components on the board.
- **Level 1** — A smaller set of diagnostic tests. Level 1 diagnostics perform status, self-test, access, and read-write tests for all the components on the board and test their registers for appropriate values. In addition, Level 1 diagnostics perform extensive tests on memory devices (for example, SDRAM, flash, NVRAM, EEPROM) wherever possible.
- **Level 2** — The full set of diagnostic tests. Level 2 diagnostics are used primarily for on-board MAC level, Physical level, external Loopback tests, and more extensive component diagnostics. Various components on the board are put into Loopback mode and test packets are transmitted through those components. These diagnostics also perform snake tests using virtual local area network (VLAN) configurations.

 **NOTE:** Diagnostic is not allowed in Stacking mode, including member stacking. Avoid stacking before executing the diagnostic tests in the chassis.

## Important Points to Remember

- You can only perform offline diagnostics on an offline standalone unit. You cannot perform diagnostics if the ports are configured in a stacking group. Remove the port(s) from the stacking group before executing the diagnostic test.
- Diagnostics only test connectivity, not the entire data path.
- Diagnostic results are stored on the flash of the unit on which you performed the diagnostics.
- When offline diagnostics are complete, the unit or stack member reboots automatically.

## Running Offline Diagnostics


To run offline diagnostics, use the following commands.  
For more information, refer to the examples following the steps.

1. Place the unit in the offline state.  
EXEC Privilege mode

```
offline stack-unit
```

You cannot enter this command on a MASTER or Standby stack unit.



 **NOTE:** The system reboots when the offline diagnostics complete. This is an automatic process. The following warning message appears when you implement the `offline stack-unit` command:

```
Warning - offline of unit will bring down all the protocols and
the unit will be operationally down, except for running Diagnostics.
Please make sure that stacking/fanout not configured for Diagnostics
execution.
Also reboot/online command is necessary for normal operation after the
offline command is issued.
Proceed with Offline [confirm yes/no]:
```

```
Dell#offline stack-unit 0
Warning - offline of unit will bring down all the protocols and
the unit will be operationally down, except for running Diagnostics.
Please make sure that stacking/fanout not configured for Diagnostics
execution.
Also reboot/online command is necessary for normal operation after the
offline command is issued.
Proceed with Offline [confirm yes/no]:yes
Dell#
```

**2. Confirm the offline status.**

EXEC Privilege mode

```
show system brief
```

```
Dell#show system brief
```

```
Stack MAC : 00:1e:c9:de:03:7b
```

```
-- Stack Info --
Unit UnitType Status ReqTyp CurTyp
Version Ports

 0 Management offline PE-FN-410S-IOA PE-FN-410S-IOA
1-0(0-1862) 12
 1 Member not present
 2 Member not present
 3 Member not present
 4 Member not present
 5 Member not present
```

```
Dell#
```

## Trace Logs

In addition to the syslog buffer, the Dell Networking OS buffers trace messages which are continuously written by various software tasks to report hardware and software events and status information.

Each trace message provides the date, time, and name of the Dell Networking OS process. All messages are stored in a ring buffer. You can save the messages to a file either manually or automatically after failover.

## Auto Save on Crash or Rollover

Exception information for MASTER or standby units is stored in the `flash://TRACE_LOG_DIR` directory. This directory contains files that save trace information when there has been a task crash or timeout.

- On a MASTER unit, you can reach the `TRACE_LOG_DIR` files by FTP or by using the `show file` command from the `flash://TRACE_LOG_DIR` directory.
- On a Standby unit, you can reach the `TRACE_LOG_DIR` files only by using the `show file` command from the `flash://TRACE_LOG_DIR` directory.

 **NOTE:** Non-management member units do not support this functionality.


### Example of the `dir flash:` Command

```
Dell#dir flash://TRACE_LOG_DIR
Directory of flash://TRACE_LOG_DIR
 1 drwx 4096 Jan 17 2011 15:02:16 +00:00 .
 2 drwx 4096 Jan 01 1980 00:00:00 +00:00 ..
 3 -rwx 100583 Feb 11 2011 20:41:36 +00:00 failure_trace0_RPM0_CP

flash: 2143281152 bytes total (2069291008 bytes free)
```

## Using the Show Hardware Commands

The `show hardware` command tree consists of commands used with the Aggregator switch. These commands display information from a hardware sub-component and from hardware-based feature tables.

 **NOTE:** Use the `show hardware` commands only under the guidance of the Dell Technical Assistance Center.

- View internal interface status of the stack-unit CPU port which connects to the external management interface.

EXEC Privilege mode

```
show hardware stack-unit {0-5} cpu management statistics
```

- View driver-level statistics for the data-plane port on the CPU for the specified stack-unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} cpu data-plane statistics
```

This view provides insight into the packet types entering the CPU to see whether CPU-bound traffic is internal (IPC traffic) or network control traffic, which the CPU must process.

- View the modular packet buffers details per stack unit and the mode of allocation.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer total-buffer
```

- View the modular packet buffers details per unit and the mode of allocation.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} total-buffer
```

- View the forwarding plane statistics containing the packet buffer usage per port per stack unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64 | all} buffer-
info
```

- View the forwarding plane statistics containing the packet buffer statistics per COS per port.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64} queue {0-14 |
all} buffer-info
```

- View input and output statistics on the party bus, which carries inter-process communication traffic between CPUs.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} cpu party-bus statistics
```

- View the ingress and egress internal packet-drop counters, MAC counters drop, and FP packet drops for the stack unit on per port basis.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} drops unit {0-0} port {33-56}
```

This view helps identifying the stack unit/port pipe/port that may experience internal drops.

- View the input and output statistics for a stack-port interface.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} stack-port {33-56}
```

- View the counters in the field processors of the stack unit.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} counters
```

- View the details of the FP Devices and Hi gig ports on the stack-unit.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} details
```

- Execute a specified bShell command from the CLI without going into the bShell.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} execute-shell-cmd {command}
```

- View the Multicast IPMC replication table from the bShell.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} ipmc-replication
```

- View the internal statistics for each port-pipe (unit) on per port basis.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} port-stats [detail]
```

- View the stack-unit internal registers for each port-pipe.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} register
```

- View the tables from the bShell through the CLI without going into the bShell.  
EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} table-dump {table name}
```

## Environmental Monitoring

Aggregator components use environmental monitoring hardware to detect transmit power readings, receive power readings, and temperature updates.

To receive periodic power updates, you must enable the following command.

- Enable environmental monitoring.  
`enable optic-info-update interval`

### Example of the `show interfaces transceiver` Command

```
Dell#show interfaces tengigabitethernet 0/9 transceiver
SFP is present
SFP+ 9 Serial Base ID fields
SFP+ 9 Id = 0x03
SFP+ 9 Ext Id = 0x04
SFP+ 9 Connector = 0x21
SFP+ 9 Transceiver Code = 0x00 0x00 0x00 0x00 0x00 0x04 0x00 0x00
SFP+ 9 Encoding = 0x00
SFP+ 9 BR Nominal = 0x67
SFP+ 9 Length(SFM) Km = 0x00
SFP+ 9 Length(OM3) 2m = 0x00
SFP+ 9 Length(OM2) 1m = 0x00
SFP+ 9 Length(OM1) 1m = 0x00
SFP+ 9 Length(Copper) 1m = 0x01
SFP+ 9 Vendor Rev = A
SFP+ 9 Laser Wavelength = 256 nm
SFP+ 9 CheckCodeBase = 0xf2
SFP+ 9 Serial Extended ID fields
SFP+ 9 Options = 0x00 0x00
SFP+ 9 BR max = 0
SFP+ 9 BR min = 0
SFP+ 9 Vendor SN = APF11040012888
SFP+ 9 Datecode = 110207
SFP+ 9 CheckCodeExt = 0xb3
SFP+ 9 DOM is not supported

Dell#
```

## Recognize an Over-Temperature Condition

An overtemperature condition occurs, for one of two reasons: the card genuinely is too hot or a sensor has malfunctioned.


Inspect cards adjacent to the one reporting the condition to discover the cause.

- If directly adjacent cards are not normal temperature, suspect a genuine overheating condition.
- If directly adjacent cards are normal temperature, suspect a faulty sensor.

When the system detects a genuine over-temperature condition, it powers off the card. To recognize this condition, look for the following system messages:

```
CHMGR-2-MAJOR_TEMP: Major alarm: chassis temperature high (temperature reaches
or exceeds threshold of
[value]C)
CHMGR-2-TEMP_SHUTDOWN_WARN: WARNING! temperature is [value]C; approaching
shutdown threshold of [value]C
```

To view the programmed alarm thresholds levels, including the shutdown value, use the `show alarms threshold` command.

 **NOTE:** When the ingress air temperature exceeds 61°C, the Status LED turns Amber and a major alarm is triggered.

#### Example of the `show alarms threshold` Command

```
Dell#show alarms threshold
-- Temperature Limits (deg C) --


Unit0 Ingress-Air Off Ingress-Air Major Off Major Shutdown
 58 61 84 86 90
Dell#
```

## Troubleshoot an Over-Temperature Condition

To troubleshoot an over-temperature condition, use the following information.

1. Use the `show environment` commands to monitor the temperature levels.
2. Check air flow through the system. Ensure that the air ducts are clean and that all fans are working correctly.
3. After the software has determined that the temperature levels are within normal limits, you can re-power the card safely. To bring back the line card online, use the `power-on` command in EXEC mode.

In addition, Dell Networking requires that you install blanks in all slots without a line card to control airflow for adequate system cooling.

 **NOTE:** Exercise care when removing a card; if it has exceeded the major or shutdown thresholds, the card could be hot to the touch.

#### Example of the `show environment` Command

```
Dell#show environment
-- Unit Environment Status --
Unit Status Temp Voltage TempStatus

* 0 online 59C ok 2

* Management Unit

-- Thermal Sensor Readings (deg C) --
Unit Sensor0 Sensor1 Sensor2

0 55 45 58

Dell#
```

## Recognize an Under-Voltage Condition

If the system detects an under-voltage condition, it sends an alarm.

To recognize this condition, look for the following system message: `%CHMGR-1-CARD_SHUTDOWN:`  
Major alarm: Line card 2 down - auto-shutdown due to under voltage.


This message indicates that the specified card is not receiving enough power. In response, the system first shuts down Power over Ethernet (PoE).

## Troubleshoot an Under-Voltage Condition

To troubleshoot an under-voltage condition, check that the correct number of power supplies are installed and their Status light emitting diodes (LEDs) are lit.

The following table lists information for SNMP traps and OIDs, which provide information about environmental monitoring hardware and hardware components.

**Table 38. SNMP Traps and OIDs**

| OID String                            | OID Name              | Description                                                                                                                                                                                                    |
|---------------------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Receiving Power</b>                |                       |                                                                                                                                                                                                                |
| .1.3.6.1.4.1.6027.3.10.1.2.5.1.6      | chSysPortXfpRecvPower | OID displays the receiving power of the connected optics.                                                                                                                                                      |
| <b>Transmitting power</b>             |                       |                                                                                                                                                                                                                |
| .1.3.6.1.4.1.6027.3.10.1.2.5.1.8      | chSysPortXfpTxPower   | OID displays the transmitting power of the connected optics.                                                                                                                                                   |
| <b>Temperature</b>                    |                       |                                                                                                                                                                                                                |
| .1.3.6.1.4.1.6027.3.10.1.2.5.1.7      | chSysPortXfpRecvTemp  | OID displays the temperature of the connected optics.                                                                                                                                                          |
|                                       |                       |  <b>NOTE:</b> These OIDs only generate if you enable the <code>enable optic-info-update-interval is enabled</code> command. |
| <b>Hardware MIB Buffer Statistics</b> |                       |                                                                                                                                                                                                                |
| .1.3.6.1.4.1.6027.3.16.1.1.4          | fpPacketBufferTable   | View the modular packet buffers details per stack unit and the mode of allocation.                                                                                                                             |
| .1.3.6.1.4.1.6027.3.16.1.1.5          | fpStatsPerPortTable   | View the forwarding plane statistics containing the packet buffer usage per port per stack unit.                                                                                                               |
| .1.3.6.1.4.1.6027.3.16.1.1.6          | fpStatsPerCOSTable    | View the forwarding plane statistics containing the packet buffer statistics per COS per port.                                                                                                                 |

## Buffer Tuning

Buffer tuning allows you to modify the way your switch allocates buffers from its available memory and helps prevent packet drops during a temporary burst of traffic.

The application-specific integrated circuit (ASICs) implement the key functions of queuing, feature lookups, and forwarding lookups in hardware.

Forwarding processor (FP) ASICs provide Ethernet MAC functions, queueing, and buffering, as well as store feature and forwarding tables for hardware-based lookup and forwarding decisions. 1G and 10G interfaces use different FPs.

You can tune buffers at three locations

1. CSF — Output queues going from the CSF.
2. FP Uplink — Output queues going from the FP to the CSF IDP links.
3. Front-End Link — Output queues going from the FP to the front-end PHY.


All ports support eight queues, four for data traffic and four for control traffic. All eight queues are tunable.

Physical memory is organized into cells of 128 bytes. The cells are organized into two buffer pools — the dedicated buffer and the dynamic buffer.

- **Dedicated buffer** — this pool is reserved memory that other interfaces cannot use on the same ASIC or by other queues on the same interface. This buffer is always allocated, and no dynamic re-carving takes place based on changes in interface status. Dedicated buffers introduce a trade-off. They provide each interface with a guaranteed minimum buffer to prevent an overused and congested interface from starving all other interfaces. However, this minimum guarantee means that the buffer manager does not reallocate the buffer to an adjacent congested interface, which means that in some cases, memory is under-used.
- **Dynamic buffer** — this pool is shared memory that is allocated as needed, up to a configured limit. Using dynamic buffers provides the benefit of statistical buffer sharing. An interface requests dynamic buffers when its dedicated buffer pool is exhausted. The buffer manager grants the request based on three conditions:
  - The number of used and available dynamic buffers.
  - The maximum number of cells that an interface can occupy.
  - Available packet pointers (2k per interface). Each packet is managed in the buffer using a unique packet pointer. Thus, each interface can manage up to 2k packets.

You can configure dynamic buffers per port on both 1G and 10G FPs and per queue on CSFs. By default, the FP dynamic buffer allocation is 10 times oversubscribed. For the 48-port 1G card:

- Dynamic Pool = Total Available Pool(16384 cells) – Total Dedicated Pool = 5904 cells
- Oversubscription ratio = 10
- Dynamic Cell Limit Per port =  $59040/29 = 2036$  cells



**Figure 31. Buffer Tuning Points**

## Deciding to Tune Buffers

Dell Networking recommends exercising caution when configuring any non-default buffer settings, as tuning can significantly affect system performance. The default values work for most cases. As a guideline, consider tuning buffers if traffic is bursty (and coming from several interfaces). In this case:

- Reduce the dedicated buffer on all queues/interfaces.
- Increase the dynamic buffer on all interfaces.
- Increase the cell pointers on a queue that you are expecting will receive the largest number of packets.

To define, change, and apply buffers, use the following commands.

- Define a buffer profile for the FP queues.  
CONFIGURATION mode  
  
`buffer-profile fp fsqueue`
- Define a buffer profile for the CSF queues.  
CONFIGURATION mode  
  
`buffer-profile csf csqueue`
- Change the dedicated buffers on a physical interface.



BUFFER PROFILE mode

`buffer dedicated`

- Change the maximum number of dynamic buffers an interface can request.

BUFFER PROFILE mode

`buffer dynamic`

- Change the number of packet-pointers per queue.

BUFFER PROFILE mode

`buffer packet-pointers`

- Apply the buffer profile to a CSF to FP link.

CONFIGURATION mode

`buffer csf linecard`

**Dell Networking OS Behavior:** If you attempt to apply a buffer profile to a non-existent port-pipe, the system displays the following message: `%DIFFSERV-2-DSA_BUFF_CARVING_INVALID_PORT_SET: Invalid FP port-set 2 for linecard 2. Valid range of port-set is <0-1>.` However, the configuration still appears in the running-config.

Configuration changes take effect immediately and appear in the running configuration. Because under normal conditions all ports do not require the maximum allocation, the configured dynamic allocations can exceed the actual amount of available memory; this allocation is called oversubscription. If you choose to oversubscribe the dynamic allocation, a burst of traffic on one interface might prevent other interfaces from receiving the configured dynamic allocation, which causes packet loss.

You cannot allocate more than the available memory for the dedicated buffers. If the system determines that the sum of the configured dedicated buffers allocated to the queues is more than the total available memory, the configuration is rejected, returning a syslog message similar to the following: `00:04:20: %S50N:0 %DIFFSERV-2-DSA_DEVICE_BUFFER_UNAVAILABLE: Unable to allocate dedicated buffers for stack-unit 0, port pipe 0, egress port 25 due to unavailability of cells.`

**Dell Networking OS Behavior:** When you remove a buffer-profile using the `no buffer-profile [fp | csf]` command from CONFIGURATION mode, the buffer-profile name still appears in the output of the `show buffer-profile [detail | summary]` command. After a stack unit reset, the buffer profile correctly returns to the default values, but the profile name remains. Remove it from the `show buffer-profile [detail | summary]` command output by entering `no buffer [fp-uplink | csf] stack-unit port-set buffer-policy` from CONFIGURATION mode and `no buffer-policy` from INTERFACE mode.

To display the allocations for any buffer profile, use the `show` commands.

To display the default buffer profile, use the `show buffer-profile {summary | detail}` command from EXEC Privilege mode.

### Example of Viewing the Default Buffer Profile

```
Dell#show buffer-profile detail interface tengigabitethernet 0/1
Interface tengig 0/1
Buffer-profile -
Dynamic buffer 194.88 (Kilobytes)
```

| Queue# | Dedicated Buffer<br>(Kilobytes) | Buffer Packets |
|--------|---------------------------------|----------------|
| 0      | 2.50                            | 256            |
| 1      | 2.50                            | 256            |
| 2      | 2.50                            | 256            |
| 3      | 2.50                            | 256            |
| 4      | 9.38                            | 256            |
| 5      | 9.38                            | 256            |
| 6      | 9.38                            | 256            |
| 7      | 9.38                            | 256            |

### Example of Viewing the Buffer Profile Allocations

```
Dell#show running-config interface tengigabitethernet 0/6 !
interface TenGigabitEthernet 0/6
mtu 9252
switchport
no shutdown
buffer-policy myfsbufferprofile
```

### Example of Viewing the Buffer Profile (Interface)

```
Dell#show buffer-profile detail int te 0/2
Interface Te 0/2
Buffer-profile fsqueue-fp
Dynamic buffer 1256.00 (Kilobytes)
Queue# Dedicated Buffer Buffer Packets
 (Kilobytes)
0 3.00 256
1 3.00 256
2 3.00 256
3 3.00 256
4 3.00 256
5 3.00 256
6 3.00 256
7 3.00 256
```

### Example of Viewing the Buffer Profile (Linecard)

```
Dell#show buffer-profile detail fp-uplink stack-unit 0 port-set 0
Linecard 0 Port-set 0
Buffer-profile fsqueue-hig
Dynamic Buffer 1256.00 (Kilobytes)
Queue# Dedicated Buffer Buffer Packets
 (Kilobytes)
0 3.00 256
1 3.00 256
2 3.00 256
3 3.00 256
4 3.00 256
5 3.00 256
6 3.00 256
7 3.00 256
```

## Using a Pre-Defined Buffer Profile

The Dell Networking OS provides two pre-defined buffer profiles, one for single-queue (for example, non-quality-of-service [QoS]) applications, and one for four-queue (for example, QoS) applications. You must reload the system for the global buffer profile to take effect, a message similar to the following displays: % Info: For the global pre-defined buffer profile to take effect, please save the config and reload the system..

**Dell Networking OS Behavior:** After you configure buffer-profile global 1Q, the message displays during every bootup. Only one reboot is required for the configuration to take effect; afterward you may ignore this bootup message.

**Dell Networking OS Behavior:** The buffer profile does not returned to the default, 4Q. If you configure 1Q, save the running-config to the startup-config, and then delete the startup-config and reload the chassis. The only way to return to the default buffer profile is to remove the 1Q profile configured and then reload the chassis.

If you have already applied a custom buffer profile on an interface, the `buffer-profile global` command fails and a message similar to the following displays: % Error: User-defined buffer profile already applied. Failed to apply global pre-defined buffer profile. Please remove all user-defined buffer profiles.

Similarly, when you configure `buffer-profile global`, you cannot not apply a buffer profile on any single interface. A message similar to the following displays: % Error: Global pre-defined buffer profile already applied. Failed to apply user-defined buffer profile on interface Gi 0/1. Please remove global pre-defined buffer profile.

If the default buffer profile (4Q) is active, the system displays an error message instructing you to remove the default configuration using the `no buffer-profile global` command.

To apply a predefined buffer profile, use the following command.

- Apply one of the pre-defined buffer profiles for all port pipes in the system.  
CONFIGURATION mode

```
buffer-profile global [1Q|4Q]
```

## Sample Buffer Profile Configuration

The two general types of network environments are sustained data transfers and voice/data.

Dell Networking recommends a single-queue approach for data transfers.

### Example of a Single Queue Application with Default Packet Pointers

```
!
buffer-profile fp fsqueue-fp
buffer dedicated queue0 3 queue1 3 queue2 3 queue3 3 queue4 3 queue5 3 queue6 3
queue7 3
buffer dynamic 1256
!
buffer-profile fp fsqueue-hig
buffer dedicated queue0 3 queue1 3 queue2 3 queue3 3 queue4 3 queue5 3 queue6
3 queue7 3
buffer dynamic 1256
```

```

!
buffer fp-uplink stack-unit 0 port-set 0 buffer-policy fsqueue-hig
buffer fp-uplink stack-unit 0 port-set 1 buffer-policy fsqueue-hig
!
Interface range gi 0/1 - 48
buffer-policy fsqueue-fp

Dell#show run int Tengig 0/10
!
interface TenGigabitEthernet 0/10

```

## Troubleshooting Packet Loss

The `show hardware stack-unit` command is intended primarily to troubleshoot packet loss. To troubleshoot packet loss, use the following commands.

- `show hardware stack-unit 0-5 cpu data-plane statistics`
- `show hardware stack-unit 0-5 cpu party-bus statistics`
- `show hardware stack-unit 0-5 drops unit 0-0 port 1-56`
- `show hardware stack-unit 0-5 stack-port 33-56`
- `show hardware stack-unit 0-5 unit 0-0 {counters | details | port-stats [detail] | register | ipmc-replication | table-dump}:`
- `show hardware {layer2| layer3} {eg acl |in acl} stack-unit 0-5 port-set 0-0`
- `show hardware layer3 qos stack-unit 0-5 port-set 0-0`
- `show hardware system-flow layer2 stack-unit 0-5 port-set 0-1 [counters]`
- `clear hardware stack-unit 0-5 counters`
- `clear hardware stack-unit 0-5 unit 0-0 counters`
- `clear hardware stack-unit 0-5 cpu data-plane statistics`
- `clear hardware stack-unit 0-5 cpu party-bus statistics`
- `clear hardware stack-unit 0-5 stack-port 33-56`

### Displaying Drop Counters

To display drop counters, use the following commands.

- Identify which stack unit, port pipe, and port is experiencing internal drops.
 

```
show hardware stack-unit 0-11 drops [unit 0 [port 0-63]]
```
- Display drop counters.
 

```
show hardware stack-unit drops unit port
```

### Example of the `show hardware stack-unit` Command to View Drop Counters Statistics

```

Dell#show hardware stack-unit 0 drops
UNIT No: 0
Total Ingress Drops :0
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0
UNIT No: 1
Total Ingress Drops :0
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0

Dell#show hardware stack-unit 0 drops unit 0

```

```

Port# :Ingress Drops :IngMac Drops :Total Mmu Drops :EgMac Drops :Egress
Drops
1 0 0 0 0 0
2 0 0 0 0 0
3 0 0 0 0 0
4 0 0 0 0 0
5 0 0 0 0 0
6 0 0 0 0 0
7 0 0 0 0 0
8 0 0 0 0 0

```

```
Dell#show hardware stack-unit 0 drops unit 0 port 1
```

```

--- Ingress Drops ---
Ingress Drops : 30
IBP CBP Full Drops : 0
PortSTPnotFwd Drops : 0
IPv4 L3 Discards : 0
Policy Discards : 0
Packets dropped by FP : 14
(L2+L3) Drops : 0
Port bitmap zero Drops : 16
Rx VLAN Drops : 0

--- Ingress MAC counters---
Ingress FCS Drops : 0
Ingress MTUExceeds : 0

--- MMU Drops ---
HOL DROPS : 0
TxPurge CellErr : 0
Aged Drops : 0

--- Egress MAC counters---
Egress FCS Drops : 0

--- Egress FORWARD PROCESSOR Drops ---
IPv4 L3UC Aged & Drops : 0
TTL Threshold Drops : 0
INVALID VLAN CNTR Drops : 0
L2MC Drops : 0
PKT Drops of ANY Conditions : 0
Hg MacUnderflow : 0
TX Err PKT Counter : 0

```

## Dataplane Statistics

The `show hardware stack-unit cpu data-plane statistics` command provides insight into the packet types coming to the CPU.

The command output in the following example has been augmented, providing detailed RX/ TX packet statistics on a per-queue basis. The objective is to see whether CPU-bound traffic is internal (so-called *party bus* or IPC traffic) or network control traffic, which the CPU must process.

### Example of Viewing Dataplane Statistics

```
Dell#show hardware stack-unit 2 cpu data-plane statistics
```

```

bc pci driver statistics for device:
 rxHandle :0
 noMhdr :0
 noMbuf :0

```

```

noClus :0
recvd :0
dropped :0
recvToNet :0
rxError :0
rxDatapathErr :0
rxPkt (COS0) :0
rxPkt (COS1) :0
rxPkt (COS2) :0
rxPkt (COS3) :0
rxPkt (COS4) :0
rxPkt (COS5) :0
rxPkt (COS6) :0
rxPkt (COS7) :0
rxPkt (UNIT0) :0
rxPkt (UNIT1) :0
rxPkt (UNIT2) :0
rxPkt (UNIT3) :0
transmitted :0
txRequested :0
noTxDesc :0
txError :0
txReqTooLarge :0
txInternalError :0
txDatapathErr :0
txPkt (COS0) :0
txPkt (COS1) :0
txPkt (COS2) :0
txPkt (COS3) :0
txPkt (COS4) :0
txPkt (COS5) :0
txPkt (COS6) :0
txPkt (COS7) :0
txPkt (UNIT0) :0

```

The `show hardware stack-unit cpu party-bus statistics` command displays input and output statistics on the party bus, which carries inter-process communication traffic between CPUs

### Example of Viewing Party Bus Statistics

```

Dell#show hardware stack-unit 2 cpu party-bus statistics
Input Statistics:
 27550 packets, 2559298 bytes
 0 dropped, 0 errors
Output Statistics:
 1649566 packets, 1935316203 bytes
 0 errors

```

### Displaying Stack Port Statistics

The `show hardware stack-unit stack-port` command displays input and output statistics for a stack-port interface.

#### Example of Viewing Stack Unit Statistics

```

Dell#show hardware stack-unit 2 stack-port 49
Input Statistics:
 27629 packets, 3411731 bytes
 0 64-byte pkts, 27271 over 64-byte pkts, 207 over 127-byte pkts
 17 over 255-byte pkts, 56 over 511-byte pkts, 78 over 1023-byte pkts
 0 Multicasts, 5 Broadcasts
 0 runts, 0 giants, 0 throttles

```

```

0 CRC, 0 overrun, 0 discarded
Output Statistics:
1649714 packets, 1948622676 bytes, 0 underruns
0 64-byte pkts, 27234 over 64-byte pkts, 107970 over 127-byte pkts
34 over 255-byte pkts, 504838 over 511-byte pkts, 1009638 over 1023-byte pkts
0 Multicasts, 0 Broadcasts, 1649714 Unicasts
0 throttles, 0 discarded, 0 collisions
Rate info (interval 45 seconds):
Input 00.00 Mbits/sec, 2 packets/sec, 0.00% of line-rate
Output 00.06 Mbits/sec, 8 packets/sec, 0.00% of line-rate
Dell#

```

## Restoring the Factory Default Settings

Restoring factory defaults deletes the existing NVRAM settings, startup configuration and all configured settings such as stacking or fanout.

To restore the factory default settings, use the `restore factory-defaults stack-unit {0-5 | all} {clear-all | nvram}` command in EXEC Privilege mode.

 **CAUTION: There is no undo for this command.**

### Important Points to Remember

- When you restore all the units in a stack, all units in the stack are placed into stand-alone mode.
- When you restore a single unit in a stack, only that unit is placed in stand-alone mode. No other units in the stack are affected.
- When you restore the units in stand-alone mode, the units remain in stand-alone mode after the restoration.
- After the restore is complete, the units power cycle immediately.

The following example shows the using the `restore factory-defaults` command to restore the Factory Default Settings.

### Restoring the Factory Default Settings

```

Dell#restore factory-defaults stack-unit 0 nvram

* Warning - Restoring factory defaults will delete the existing *
* persistent settings (stacking, fanout, etc.) *
* After restoration the unit(s) will be powercycled immediately. *
* Proceed with caution ! *
***** Proceed
with factory
settings? Confirm [yes/no]:yes
-- Restore status -- Unit Nvram Config


0 Success

Power-cycling the unit(s).
....

```

## Standards Compliance

This chapter describes standards compliance for Dell Networking products.

 **NOTE:** Unless noted, when a standard cited here is listed as supported by the Dell Networking Operating System (OS), the system also supports predecessor standards. One way to search for predecessor standards is to use the <http://tools.ietf.org/> website. Click "Browse and search IETF documents," enter an RFC number, and inspect the top of the resulting document for obsolescence citations to related RFCs.

### IEEE Compliance

The following is a list of IEEE compliance.

|               |                                            |
|---------------|--------------------------------------------|
| 802.1AB       | LLDP                                       |
| 802.1D        | Bridging                                   |
| 802.1p        | L2 Prioritization                          |
| 802.1Q        | VLAN Tagging, Double VLAN Tagging, GVRP    |
| 802.3ac       | Frame Extensions for VLAN Tagging          |
| 802.3ad       | Link Aggregation with LACP                 |
| 802.3ae       | 10 Gigabit Ethernet (10GBASE-W, 10GBASE-X) |
| 802.3ak       | 10 Gigabit Ethernet (10GBASE-CX4)          |
| 802.3i        | Ethernet (10BASE-T)                        |
| 802.3u        | Fast Ethernet (100BASE-FX, 100BASE-TX)     |
| 802.3x        | Flow Control                               |
| 802.1Qaz      | Enhanced Transmission Selection            |
| 802.1Qbb      | Priority-based Flow Control                |
| ANSI/TIA-1057 | LLDP-MED                                   |
| MTU           | 12,000 bytes                               |

### RFC and I-D Compliance

The Dell Networking OS supports the following standards. The standards are grouped by related protocol. The columns showing support by platform indicate which version of Dell Networking OS first supports the standard.



## General Internet Protocols

The following table lists the Dell Networking OS support per platform for general internet protocols.

**Table 39. General Internet Protocols**

| <b>RFC#</b> | <b>Full Name</b>                                                                        |
|-------------|-----------------------------------------------------------------------------------------|
| 768         | User Datagram Protocol                                                                  |
| 793         | Transmission Control Protocol                                                           |
| 854         | Telnet Protocol Specification                                                           |
| 959         | File Transfer Protocol (FTP)                                                            |
| 1321        | The MD5 Message-Digest Algorithm                                                        |
| 1350        | The TFTP Protocol (Revision 2)                                                          |
| 1661        | The Point-to-Point Protocol (PPP)                                                       |
| 1989        | PPP Link Quality Monitoring                                                             |
| 1990        | The PPP Multilink Protocol (MP)                                                         |
| 1994        | PPP Challenge Handshake Authentication Protocol (CHAP)                                  |
| 2474        | Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers |
| 2698        | A Two Rate Three Color Marker                                                           |
| 3164        | The BSD syslog Protocol                                                                 |

## General IPv4 Protocols

The following table lists the Dell Networking OS support per platform for general IPv4 protocols.

**Table 40. General IPv4 Protocols**

| <b>RFC#</b> | <b>Full Name</b>                                                             |
|-------------|------------------------------------------------------------------------------|
| 791         | Internet Protocol                                                            |
| 792         | Internet Control Message Protocol                                            |
| 826         | An Ethernet Address Resolution Protocol                                      |
| 1191        | Path MTU Discovery                                                           |
| 1305        | Network Time Protocol (Version 3) Specification, Implementation and Analysis |
| 1542        | Clarifications and Extensions for the Bootstrap Protocol                     |
| 1812        | Requirements for IP Version 4 Routers                                        |
| 2131        | Dynamic Host Configuration Protocol                                          |

## Network Management

The following table lists the Dell Networking OS support per platform for network management protocol.

**Table 41. Network Management**

| RFC# | Full Name                                                                                                     |
|------|---------------------------------------------------------------------------------------------------------------|
| 1155 | Structure and Identification of Management Information for TCP/IP-based Internets                             |
| 1156 | Management Information Base for Network Management of TCP/IP-based internets                                  |
| 1157 | A Simple Network Management Protocol (SNMP)                                                                   |
| 1212 | Concise MIB Definitions                                                                                       |
| 1215 | A Convention for Defining Traps for use with the SNMP                                                         |
| 1493 | Definitions of Managed Objects for Bridges [except for the dot1dTpLearnedEntryDiscards object]                |
| 1901 | Introduction to Community-based SNMPv2                                                                        |
| 2011 | SNMPv2 Management Information Base for the Internet Protocol using SMIv2                                      |
| 2012 | SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2                          |
| 2013 | SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2                                 |
| 2024 | Definitions of Managed Objects for Data Link Switching using SMIv2                                            |
| 2096 | IP Forwarding Table MIB                                                                                       |
| 2570 | Introduction and Applicability Statements for Internet Standard Management Framework                          |
| 2571 | An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks                |
| 2572 | Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)                          |
| 2574 | User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)              |
| 2575 | View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)                      |
| 2576 | Coexistence Between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework |
| 2578 | Structure of Management Information Version 2 (SMIv2)                                                         |

| <b>RFC#</b>            | <b>Full Name</b>                                                                                                                                                                                                              |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2579                   | Textual Conventions for SMIv2                                                                                                                                                                                                 |
| 2580                   | Conformance Statements for SMIv2                                                                                                                                                                                              |
| 2618                   | RADIUS Authentication Client MIB, except the following four counters:<br>radiusAuthClientInvalidServerAddresses<br>radiusAuthClientMalformedAccessResponses<br>radiusAuthClientUnknownTypes<br>radiusAuthClientPacketsDropped |
| 3635                   | Definitions of Managed Objects for the Ethernet-like Interface Types                                                                                                                                                          |
| 2674                   | Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions                                                                                                               |
| 2787                   | Definitions of Managed Objects for the Virtual Router Redundancy Protocol                                                                                                                                                     |
| 2819                   | Remote Network Monitoring Management Information Base: Ethernet Statistics Table, Ethernet History Control Table, Ethernet History Table, Alarm Table, Event Table, Log Table                                                 |
| 2863                   | The Interfaces Group MIB                                                                                                                                                                                                      |
| 2865                   | Remote Authentication Dial In User Service (RADIUS)                                                                                                                                                                           |
| 3273                   | Remote Network Monitoring Management Information Base for High Capacity Networks (64 bits): Ethernet Statistics High-Capacity Table, Ethernet History High-Capacity Table                                                     |
| 3416                   | Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)                                                                                                                                        |
| 3418                   | Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)                                                                                                                                           |
| 3434                   | Remote Monitoring MIB Extensions for High Capacity Alarms, High-Capacity Alarm Table (64 bits)                                                                                                                                |
| ANSI/TIA-1057          | The LLDP Management Information Base extension module for TIA-TR41.4 Media Endpoint Discovery information                                                                                                                     |
| draft-grant-tacacs -02 | The TACACS+ Protocol                                                                                                                                                                                                          |
| IEEE 802.1AB           | Management Information Base module for LLDP configuration, statistics, local system data and remote systems data components.                                                                                                  |

| RFC#                         | Full Name                                                                                                                                                                                    |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IEEE 802.1AB                 | The LLDP Management Information Base extension module for IEEE 802.1 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)                                       |
| IEEE 802.1AB                 | The LLDP Management Information Base extension module for IEEE 802.3 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)                                       |
| sFlow.org                    | sFlow Version 5                                                                                                                                                                              |
| sFlow.org                    | sFlow Version 5 MIB                                                                                                                                                                          |
| FORCE10-IF-EXTENSION-MIB     | Force10 Enterprise IF Extension MIB (extends the Interfaces portion of the MIB-2 (RFC 1213) by providing proprietary SNMP OIDs for other counters displayed in the "show interfaces" output) |
| FORCE10-LINKAGG-MIB          | Force10 Enterprise Link Aggregation MIB                                                                                                                                                      |
| FORCE10-COPY-CONFIG-MIB      | Force10 File Copy MIB (supporting SNMP SET operation)                                                                                                                                        |
| FORCE10-MONMIB               | Force10 Monitoring MIB                                                                                                                                                                       |
| FORCE10-PRODUCTS-MIB         | Force10 Product Object Identifier MIB                                                                                                                                                        |
| FORCE10-SS-CHASSIS-MIB       | Force10 S-Series Enterprise Chassis MIB                                                                                                                                                      |
| FORCE10-SMI                  | Force10 Structure of Management Information                                                                                                                                                  |
| FORCE10-SYSTEM-COMPONENT-MIB | Force10 System Component MIB (enables the user to view CAM usage information)                                                                                                                |
| FORCE10-TC-MIB               | Force10 Textual Convention                                                                                                                                                                   |
| FORCE10-TRAP-ALARM-MIB       | Force10 Trap Alarm MIB                                                                                                                                                                       |
| FORCE10-FIPS NOOPING-MI B    | Force10 FIP Snooping MIB (Based on T11-FCoE-MIB mentioned in FC-BB-5)                                                                                                                        |
| FORCE10-DCB -MIB             | Force10 DCB MIB                                                                                                                                                                              |
| IEEE 802.1Qaz                | Management Information Base extension module for IEEE 802.1 organizationally defined discovery information (LDP-EXT-DOT1-DCBX-MIB)                                                           |
| IEEE 802.1Qbb                | Priority-based Flow Control module for managing IEEE 802.1Qbb                                                                                                                                |

## MIB Location

You can find Force10 MIBs under the Force10 MIBs subhead on the Documentation page of iSupport:

<https://www.force10networks.com/csportal20/KnowledgeBase/Documentation.aspx>

You also can obtain a list of selected MIBs and their OIDs at the following URL:

[https://www.force10networks.com/csportal20/MIBs/MIB\\_OIDs.aspx](https://www.force10networks.com/csportal20/MIBs/MIB_OIDs.aspx)

Some pages of iSupport require a login. To request an iSupport account, go to:

<https://www.force10networks.com/CSPortal20/Support/AccountRequest.aspx>

If you have forgotten or lost your account information, contact Dell TAC for assistance.