Cisco MDS 9000 Series Interfaces Configuration Guide, Release 8.x

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

Preface

Preface

This preface describes the audience, organization of, and conventions used in the Cisco MDS 9000 Series Configuration Guides. It also provides information on how to obtain related documentation, and contains the following chapters:

Audience

To use this installation guide, you need to be familiar with electronic circuitry and wiring practices, and preferably be an electronic or electromechanical technician.

Document Conventions

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Warnings use the following conventions:

Warning This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device. Statement 1071.

Related Documentation

The documentation set for the Cisco MDS 9000 Series Switches includes the following documents.

Release Notes

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-release-notes-list.html

Regulatory Compliance and Safety Information

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/hw/regulatory/compliance/RCSI.html

Compatibility Information

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-device-support-tables-list.html

Installation and Upgrade

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-installation-guides-list.html

Configuration

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-installation-and-configuration-guides-list.html

CLI

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-command-reference-list.html

Troubleshooting and Reference

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/tsd-products-support-troubleshoot-and-alerts.html

To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/doclocater.html

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the . RSS feeds are a free service.

New and Changed Information

Table 1: New and Changed Interfaces Features, on page 1 summarizes the new and changed information in this document, and shows the releases in which each feature is supported. Your software release might not support all the features in this document. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release.

Table 1: New and Changed Interfaces Features

Feature Name	Description	Release	Where Documented
Port Beaconing	This feature can be used to identify individual switch and directly attached peer ports in a data center environment.	8.3(1)	Configuring Interfaces, on page 40
Buffer-to-Buffer Credit Recovery	This feature is supported for F ports.	8.2(1)	Configuring Interface Buffers, on page 105
Fibre Channel over Ethernet (FCoE)	New FCoE commands were introduced and some FCoE commands were modified to align with the commands used in Fibre Channel.	8.2(1)	Congestion Detection, Avoidance, and Isolation, on page 125
Port Monitor	The link connecting a core switch to a Cisco NPV switch should be treated as an Inter-Switch Link (ISL) (core port) in the port monitor. Previously, core ports were included as access ports and were subject to any portguard actions configured. This allows portguard actions on true access (edge) ports, while ports connecting to Cisco NPV switches remain unaffected.	8.1(1)	Configuring Interfaces, on page 40

Congestion Drop Timeout and No-Credit Frame Timeout Values for Fibre Channel	The link connecting a core switch to a Cisco NPV switch should be treated as an ISL (core port) for the purposes of congestion-drop, no-credit-drop, and slowport-monitor thresholds for Fibre Channel. Previously, core ports were subject to any change in the congestion-drop or no-credit-drop mode F value.	8.1(1)	Congestion Detection, Avoidance, and Isolation, on page 125
Slow Drain Detection and Congestion Isolation	The new Congestion Isolation feature can detect a slow-drain device via port monitor or manual configuration and isolate it from other normally performing devices on an ISL. Once the traffic to the slow-drain device is isolated, the traffic to the rest of the normally behaving devices remain unaffected. Traffic isolation is accomplished via the following three features: **1.** Extended Receiver Ready—This feature allows each ISL between supporting switches to be split into four separate virtual links, with each virtual link assigned its own buffer-to-buffer credits. One virtual link is for control traffic, one is for high-priority traffic, one is for slow devices, and the remaining one is for normal traffic. **2.** Congestion Isolation—This feature allows devices to be categorized as slow by either configuration command or by the port monitor. **3.** Port monitor portguard action for Congestion Isolation—Port monitor has a new portguard option to allow the categorization of a device as slow, so that it can have all the traffic flowing to the device routed to the slow virtual link.	8.1(1)	Congestion Detection, Avoidance, and Isolation, on page 125

Interface Overview

This chapter provides an overview of the interfaces and its features.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Trunks and PortChannels

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Series. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports.

PortChannels aggregate multiple physical ISLs into one logical link with higher bandwidth and port resiliency for both Fibre Channel and FICON traffic. With this feature, up to 16 expansion ports (E-ports) or trunking E-ports (TE-ports) can be bundled into a PortChannel. ISL ports can reside on any switching module, and they do not need a designated master port. If a port or a switching module fails, the PortChannel continues to function properly without requiring fabric reconfiguration.

Cisco NX-OS software uses a protocol to exchange PortChannel configuration information between adjacent switches to simplify PortChannel management, including misconfiguration detection and autocreation of PortChannels among compatible ISLs. In the autoconfigure mode, ISLs with compatible parameters automatically form channel groups; no manual intervention is required.

PortChannels load balance Fibre Channel traffic using a hash of source FC-ID and destination FC-ID, and optionally the exchange ID. Load balancing using PortChannels is performed over both Fibre Channel and FCIP links. Cisco NX-OS software also can be configured to load balance across multiple same-cost FSPF routes.

Fibre Channel Port Rate Limiting

The Fibre Channel port rate-limiting feature for the Cisco MDS 9100 Series controls the amount of bandwidth available to individual Fibre Channel ports within groups of four host-optimized ports. Limiting bandwidth on one or more Fibre Channel ports allows the other ports in the group to receive a greater share of the available bandwidth under high-utilization conditions. Port rate limiting is also beneficial for throttling WAN traffic at the source to help eliminate excessive buffering in Fibre Channel and IP data network devices.

Maximum NPIV Limit

The maximum number of NPIV logins is not configurable at the port level on edge switches operating in NPV mode. Starting with Cisco MDS 9000 Release 6.2(7), the maximum NPIV limit feature is supported on core NPIV switches, which include Cisco MDS 9513, MDS 9710, and MDS 9250i switches. The maximum NPIV limit per-port feature allows you to configure a per-port limit. If a maximum limit is configured, whenever an FDISC is received, it checks if the maximum NPIV limit is exceeded, then it will reject the FLOGI. If the maximum NPIV limit is not exceeded, if the limit is exceeded, then it will process the FLOGI. The **trunk-max-npiv-limit** command is used for F ports in trunking mode with multiple VSANs. If a port's operational mode goes into trunking mode, this parameter is used.

Extended Credits

Full line-rate Fibre Channel ports provide at least 255 standard buffer credits . Adding credits lengthens distances for the Fibre Channel SAN extension. Using extended credits, up to 4095 buffer credits from a pool of more than 6000 buffer credits for a module can be allocated to ports as needed to greatly extend the distance for Fibre Channel SANs.

Note This feature is supported on all Cisco MDS Director Class Fabric Switches and it is not supported on any Cisco MDS Fabric switches.

N Port Virtualization

Cisco NX-OS software supports industry-standard N port identifier virtualization (NPIV), which allows multiple N port fabric logins concurrently on a single physical Fibre Channel link. HBAs that support NPIV can help improve SAN security by enabling zoning and port security to be configured independently for each virtual machine (OS partition) on a host. In addition to being useful for server connections, NPIV is beneficial for connectivity between core and edge SAN switches.

N port virtualizer (NPV) is a complementary feature that reduces the number of Fibre Channel domain IDs in core-edge SANs. Cisco MDS 9000 Series Multilayer switches operating in the NPV mode do not join a fabric; they only pass traffic between core switch links and end devices, which eliminates the domain IDs for these switches. NPIV is used by edge switches in the NPV mode to log in to multiple end devices that share a link to the core switch. This feature is available only for Cisco MDS Blade Switch Series, the Cisco MDS 9124 Multilayer Fabric Switch, Cisco MDS 9134 Multilayer Fabric Switch, Cisco MDS 9148 Multilayer Fabric Switch, Cisco MDS 9148S Multilayer Fabric Switch, and Cisco MDS 9396S Multilayer Fabric Switch.

FlexAttach

One of the main problems in a SAN environment is the time and effort required to install and replace servers. The process involves both SAN and server administrators, and the interaction and coordination between them can make the process time consuming. To alleviate the need for interaction between SAN and server administrators, the SAN configuration should not be changed when a new server is installed or an existing server is replaced. FlexAttach addresses these problems by reducing configuration changes and the time and coordination required by SAN and server administrators when installing and replacing servers. This feature is available only for Cisco MDS 9000 Blade Switch Series, the Cisco MDS 9124, Cisco MDS 9134, Cisco MDS 9148 Multilayer Fabric Switch, Cisco MDS 9148S Multilayer Fabric Switch, and Cisco MDS 9396S switches when NPV mode is enabled.

Configuring Interfaces

This chapter provides information about interfaces and how to configure interfaces.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Feature History for Interfaces

Table 2: New and Changed Features, on page 13 lists the New and Changed features.

Table 2: New and Changed Features

Feature Name	Release	Feature Information
Interfaces and Port Channels		
Port Beaconing	8.3(1)	This feature can be used to identify individual switch and directly attached peer ports in a data center environment. The following command was introduced: **beacon interface fc** *slot/port* {**both** | **local** | **peer**} [**status** {**normal** | **warning** | **critical**}] [**duration** *seconds*] [**frequency** *number*]
Interface Modes	8.1(1)	The link connecting from a core switch to a Cisco N-Port Virtualizer (NPV) switch must be treated as an ISL (core port) in interfaces and port channels. Port monitor may take portguard action on the link if it is treated as an edge port, which will result in the loss of connectivity to the devices that are connected to the Cisco NPV switch. The following command was introduced: **switchport logical-type** {**auto** | **core** | **edge**}
Port Monitor		
Port Monitor	8.1(1)	The **port-type** {**access-port** | **trunks** | **all**} command was replaced with the **logical-type** {**core** | **edge** | **all**} command, where **port-type** was replaced with **logical-type**, **access-port** was replaced with **edge**, and **trunks** was replaced with **core**. The following command was modified: **logical-type** {**core** | **edge** | **all**}

Feature Name	Release	Feature Information
Port Monitor Policy	8.1(1)	A new port monitor portguard action (cong-isolate) was introduced for the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters. The *cong-isolate* portguard action was added to the following commands: • **counter credit-loss-reco** • **counter tx-credit-not-available** • **counter tx-slowport-oper-delay** • **counter tx-wait**

Information About Interfaces

The main function of a switch is to relay frames from one data link to another. To relay the frames, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces.

Interface Description

For Fibre Channel interfaces, you can configure the description parameter to provide a recognizable name for an interface. Using a unique name for each interface allows you to quickly identify an interface when you are looking at a listing of multiple interfaces. You can also use the description to identify the traffic or the use for a specific interface.

Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, and ST port (see Figure 1: Cisco MDS 9000 Series Switch Port Modes, on page 15). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Figure 1: Cisco MDS 9000 Series Switch Port Modes

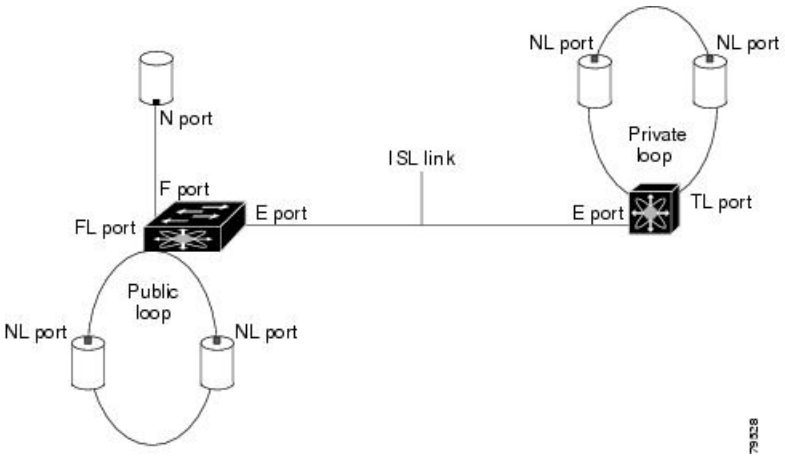

Note Interfaces are created in VSAN 1 by default. For more information about VSAN, see the Cisco MDS 9000 Series NX-OS Fabric Configuration Guide.

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.

- The operational status represents the current status of a specified attribute, such as the interface speed. This status cannot be changed and is read-only. Some values, for example, operational speed, may not be valid when the interface is down.

Note When a module is removed and replaced with the same type of module, the original configuration is retained. If a different type of module is inserted, the original configuration is no longer retained.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port can be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined for remote N ports and NL ports. E ports support Class 2, Class 3, and Class F services.

An E port connected to another switch can also be configured to form a port channel. For more details about configuring a port channel, see Configuring PortChannels, on page 201.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port can be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support Class 2 and Class 3 services.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port can be connected to one or more NL ports (including FL ports in other switches) to form a public, arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support Class 2 and Class 3 services.

NP Ports

An NP port is a port on a device that is in NPV mode and connected to the core switch via an F port. NP ports function like N ports, except that in addition to providing N port operations, they also function as proxies for multiple physical N ports.

For more details about NP ports and NPV, see Configuring N Port Virtualization, on page 233 .

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It can be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Series Multilayer Switches. These switches expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all the frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Series Multilayer Switches. For more details about trunking, see Configuring Trunking, on page 181. TE ports support Class 2, Class 3, and Class F services.

TF Port

In trunking F port (TF port) mode, an interface functions as a trunking expansion port. It can be connected to another trunked N port (TN port) or trunked NP port (TNP port) to create a link between a core switch and an NPV switch or an host bus adapter (HBA) in order to carry tagged frames. TF ports are specific to Cisco MDS 9000 Series Multilayer Switches. They expand the functionality of F ports to support VSAN trunking.

In TF port mode, all the frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Series Multilayer Switches. For more details about trunking, see Configuring Trunking, on page 181. TF ports support Class 2, Class 3, and Class F services.

TNP Port

In trunking NP port (TNP port) mode, an interface functions as a trunking expansion port. It can be connected to a trunked F port (TF port) to create a link to a core NPIV switch from an NPV switch in order to carry tagged frames.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Series. It monitors network traffic that passes though a Fibre Channel interface. This is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames; they only transmit a copy of the source traffic. The SPAN feature is non-intrusive and does not affect switching of network traffic in SPAN source ports. For more details about SPAN, see the Cisco MDS 9000 Series NX-OS System Management Configuration Guide.

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Series Multilayer Switches. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic. For more details about SPAN, see the Cisco MDS 9000 Series NX-OS System Management Configuration Guide.

Fx Port

Interfaces configured as Fx ports can operate in either F port mode or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode, for example, preventing an interface to connect to another switch.

Auto Mode

Interfaces configured in auto mode can operate in F port, FL port, E port, TE port, or TF port mode. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port mode or FL port mode depending on the N port mode or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Series Multilayer Switches, it may become operational in TE port mode. For more details about trunking, see Configuring Trunking, on page 181.

TL ports and SD ports are not determined during initialization and are administratively configured.

Interface States

An interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface, as described in Table 3: Administrative States , on page 18.

Table 3: Administrative States

Administrative State	Description
Up	Interface is enabled.
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

Operational state indicates the current operational state of an interface, as described in Table 4: Operational States , on page 18.

Table 4: Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic, as required. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed.
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE mode or TF mode.

Reason Codes

Reason codes are dependent on the operational state of an interface, as described in Table 5: Reason Codes for Interface States , on page 19.

Table 5: Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.
Down	Down	Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	See Table 6: Reason Codes for Nonoperational States , on page 20. Note that only some of the reason codes are listed in Table 6: Reason Codes for Nonoperational States , on page 20.

Note Only some of the reason are listed in the table.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code, as described in Table 6: Reason Codes for Nonoperational States , on page 20.

Table 6: Reason Codes for Nonoperational States

Reason Code (Long Version)	Description	Applicable Modes
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The Cisco NX-OS software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons: • Configuration failure • Incompatible buffer-to-buffer credit configuration To make the interface operational, you must first fix the error conditions causing this state, and administratively shut down or enable the interface.	
Fibre Channel redirect failure	A port is isolated because a Fibre Channel redirect is unable to program routes.	
No port activation license available	A port is not active because it does not have a port license.	
SDM failure	A port is isolated because SDM is unable to program routes.	

Reason Code (Long Version)	Description	Applicable Modes
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to the other side of the link E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. This might occur if more than one FL port exists in the same loop, in which case, all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
Port Channel administratively down	The interfaces belonging to a port channel are down.	Only port channel interfaces
Suspended due to incompatible speed	The interfaces belonging to a port channel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to a port channel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a port channel must be connected to the same pair of switches.	

Graceful Shutdown

Interfaces on a port are shut down by default (unless you modified the initial configuration).

The Cisco NX-OS software implicitly performs a graceful shutdown in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface.
- If a Cisco NX-OS software application executes a port shutdown as part of its function.

A graceful shutdown ensures that no frames are lost when the interface is shutting down. When a shutdown is triggered either by you or the Cisco NX-OS software, the switches connected to the shutdown link coordinate with each other to ensure that all the frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shutdown is not possible in the following situations:

- If you physically remove the port from the switch.
- If In-Order Delivery (IOD) is enabled. For more details about IOD, see Cisco MDS 9000 Series NX-OS Fabric Configuration Guide.
- If the Min_LS_interval interval is higher than 10 seconds. For information about Fabric Shortest Path First (FSPF) global configuration, see Cisco MDS 9000 Series NX-OS Fabric Configuration Guide

  .

Note This feature is triggered only if both the switches at either end of the E port interface are Cisco MDS switches and are running Cisco SAN-OS Release 2.0(1b) or later, or Cisco MDS NX-OS Release 4.1(1a) or later.

Port Administrative Speeds

By default, the port administrative speed for an interface is automatically calculated by the switch.

Autosensing

Auto sensing speed is enabled on all 4-Gbps and 8-Gbps switching module interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on 4-Gbps switching modules, and 8 Gbps on 8-Gbps switching modules. When auto sensing is enabled for an interface operating in dedicated rate mode, 4 Gbps of bandwidth is reserved even if the port negotiates at an operating speed of 1 Gbps or 2 Gbps.

To avoid wasting unused bandwidth on 48-port and 24-port 4-Gbps and 8-Gbps Fibre Channel switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 Gbps or 8 Gbps. This feature shares the unused bandwidth within the port group, provided the bandwidth does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports that are configured for auto sensing.

Tip When migrating a host that supports up to 2-Gbps traffic (that is, not 4 Gbps with auto-sensing capabilities) to the 4-Gbps switching modules, use auto sensing with a maximum bandwidth of 2 Gbps. When migrating a host that supports up to 4-Gbps traffic (that is, not 8 Gbps with auto-sensing capabilities) to the 8-Gbps switching modules, use auto sensing with a maximum bandwidth of 4 Gbps.

Frame Encapsulation

The **switchport encap eisl** command applies only to SD port interfaces. This command determines the frame format for all the frames transmitted by the interface in SD port mode. If the encapsulation is set to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources. For information about encapsulation, see the Cisco MDS 9000 Series NX-OS System Management Configuration Guide.

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface** *SD_port_interface* command output. For information about encapsulation, see the Cisco MDS 9000 Series NX-OS System Management Configuration Guide.

Port Beaconing

The Port Beaconing feature can be used to identify individual switch and directly attached peer ports in a data center environment. This feature may be used by a switch administrator to help a data center operations personnel to identify ports that need to be serviced by replacing cables or small form-factor pluggable transceivers (SFPs).

The switch administrator can specify a status, duration, and blink rate for switch port beacon LEDs. Port Beacon LEDs of any directly attached peer port may also be controlled if the peer supports the Link Cable Beaconing (LCB) Fibre Channel protocol. Port beacon LEDs on either end or both ends of a link may be controlled using a single command.

Bit Error Rate Thresholds

The bit error rate (BER) threshold is used by a switch to detect an increased error rate before performance degradation seriously affects traffic.

Bit errors occur because of the following reasons:

- Faulty or bad cable
- Faulty or bad Gigabit Interface Converter (GBIC) or Small Form-Factor Pluggable (SFP)
- GBIC or SFP is specified to operate at 1 Gbps, but is used at 2 Gbps
- GBIC or SFP is specified to operate at 2 Gbps, but is used at 4 Gbps
- Short-haul cable is used for long haul or long-haul cable is used for short haul
- Momentary synchronization loss
- Loose cable connection at one end or both ends
- Improper GBIC or SFP connection at one end or both ends

A BER threshold is detected when 15 error bursts occur in an interval of minimum 45 seconds and a maximum of 5-minute period with a sampling interval of 3 seconds. By default, the switch disables the interface when the threshold is reached. Use the **shutdown** and **no shutdown** command sequence to re-enable the interface.

You can configure the switch to not disable an interface when the threshold is crossed. By default, the threshold disables the interface.

Disabling the Bit Error Rate Threshold

By default, the threshold disables the interface. However, you can configure the switch to not disable an interface when the threshold is crossed.

To disable the BER threshold for an interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Prevent the detection of BER events from disabling the interface:

switch(config-if)# **switchport ignore bit-errors**

(Optional) Prevent the detection of BER events from enabling the interface:

switch(config-if)# **no switchport ignore bit-errors**

Tip Regardless of the setting of the **switchport ignore bit-errors** command, a switch generates a syslog message when the BER threshold is exceeded.

SFP Transmitter Types

The SFP hardware transmitters are identified by their acronyms when displayed using the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** and **show interface fc** *slot/port* **transceiver** commands display both values (ID and transmitter type) for Cisco-supported SFPs. Table 7: SFP Transmitter Acronym Definitions , on page 24 defines the acronyms used in the command output. For information about how to display interface information, see the Displaying Interface Information, on page 59.

Table 7: SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
Short wave laser	swl
Long wave laser	lwl
Long wave laser cost reduced	lwcr
Electrical	elec
Extended transmitters assigned to Cisco-supported SFPs	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530

CWDM-1550	c1550
CWDM-1570	c1570
CWDM-1590	c1590
CWDM-1610	c1610

Portguard

The Portguard feature is intended for use in environments where systems do not adapt quickly to a port going down and up (single or multiple times). For example, if a large fabric takes 5 seconds to stabilize after a port goes down, but the port actually goes up and down once per second, a severe failure might occur in the fabric, including devices becoming permanently unsynchronized.

The Portguard feature provides the SAN administrator with the ability to prevent this issue from occurring. A port can be configured to stay down after a specified number of failures in a specified time period. This allows the SAN administrator to automate fabric stabilization, thereby avoiding problems caused by the up-down cycle.

Using the Portguard feature, the SAN administrator can restrict the number of error events and bring a malfunctioning port to down state dynamically once the error events exceed the event threshold. A port can be configured such that it shuts down when specific failures occur.

There are two types of portguard, *Port Level* type and *Port Monitor* type. While the former is a basic type where event thresholds are configurable on a per port basis, the latter allows the configuration of policies that are applied to all the ports of the same type, for example, all E ports or all F ports.

Note We recommend against the simultaneous use of both types of portguard for a given port.

Port Level Portguard

The following is the list of events that can be used to trigger port-level portguard actions:

- TrustSec violation—Link fails because of excessive TrsustSec violation events.
- Bit errors—Link fails because of excessive bit error events.
- Signal loss—Link fails because of excessive signal loss events.
- Signal synchronization loss—Link fails because of excessive signal synchronization events.
- Link reset—Link fails because of excessive link reset events.
- Link down—Link fails because of excessive link down events.
- Credit loss (Loop F ports only)—Link fails because of excessive credit loss events.

A link failure occurs when it receives two bad frames in an interval of 10 seconds and the respective interface will be error disabled. A general link failure caused by link down is the superset of all other causes. The sum of the number of all other causes equals the number of link down failures. This means that a port is brought to down state when it reaches the maximum number of allowed link failures or the maximum number of specified causes.

Port level portguard can be used to shut down misbehaving ports based on certain link event types. Event thresholds are configurable for each event type per port which makes them customizable between host, array, and tape F ports, or between intra- and inter-data center E ports, for example.

The events listed above might get triggered by certain events on a port, such as:

- Receipt of Not Operational Signal (NOS)
- Too many hardware interrupts
- The cable is disconnected
- The detection of hardware faults
- The connected device is rebooted (F ports only)
- The connected modules are rebooted (E ports only)

Port Monitor Portguard

The Port Monitor Portguard feature allows a port to be automatically error disabled or flapped when a given event threshold is reached.

Note Absolute counters do not support portguard action. However, TX Slowport Oper Delay counter supports Congestion Isolation portguard action.

The following is the list of events that can be used to trigger the Port Monitor portguard actions:

- err-pkt-from-xbar
- err-pkt-to-xbar
- credit-loss-reco
- link-loss
- signal-loss
- sync-loss
- rx-datarate
- invalid-crc
- invalid-words
- link-loss
- tx-credit-not-available
- tx-datarate
- tx-discards
- tx-slowport-oper-delay
- txwait

- tx-discards

Port Monitor

The Port Monitor feature can be used to monitor the performance and status of ports and generate alerts when problems occur. You can configure thresholds for various counters and enable event triggers when the values cross the threshold.

For rising and falling thresholds, a syslog is generated only when the error count crosses these threshold values.

Table 8: Default Port Monitor Policy with Threshold Values, on page 27 displays the default port monitor policy with threshold values. The unit for threshold values (rising and falling) differs across different counters.

Note The link connecting a core switch to a Cisco NPV switch should be treated as an Inter-Switch Link (ISL) (core port) in the port monitor. Previously, core ports were included as access ports and were subject to any portguard actions configured. This allows portguard actions on true access (edge) ports, while ports connecting to Cisco NPV switches remain unaffected. Use the interface level **switchport logical-type** command to change the logical type for the links between an NPIV switch and a Cisco NPV switch.

Note NP ports are not monitored in port monitor.

Table 8: Default Port Monitor Policy with Threshold Values

Counter	Threshold Type	Interval (Seconds)	Rising Threshold	Event	Falling Threshold	Event	Warning Threshold	Port Monitor Portguard
link-loss	Delta	60	5	4	1	4	Not enabled	Not enabled
sync-loss	Delta	60	5	4	1	4	Not enabled	Not enabled
signal-loss	Delta	60	5	4	1	4	Not enabled	Not enabled
state-change	Delta	60	5	4	0	4	Not enabled	Not enabled
invalid-words	Delta	60	5	4	0	4	Not enabled	Not enabled
invalid-crc	Delta	60	5	4	1	4	Not enabled	Not enabled
tx-discards	Delta	60	200	4	10	4	Not enabled	Not enabled

lr-rx	Delta	60	5	4	1	4	Not enabled	Not enabled
lr-tx	Delta	60	5	4	1	4	Not enabled	Not enabled
timeout-discards	Delta	60	200	4	10	4	Not enabled	Not enabled
credit-loss-reco	Delta	60	1	4	0	4	Not enabled	Not enabled
tx-credit-not-available	Delta	1	10% [1]	4	0%	4	Not enabled	Not enabled
rx-datarate	Delta	60	80%	4	20%	4	Not enabled	Not enabled
tx-datarate	Delta	60	80%	4	20%	4	Not enabled	Not enabled
tx-slowport-oper-delay [2]	Absolute	60	50 ms	4	0 ms	4	Not enabled	Not enabled
txwait [3]	Delta	60	40%	4	0%	4	Not enabled	Not enabled
err-pkt-from-port_ASIC Error Pkt from Port	—	—	—	—	—	—	—	—
err-pkt-to-xbar_ASIC Error Pkt to xbar	—	—	—	—	—	—	—	—
err-pkt-from-xbar_ASIC Error Pkt from xbar	—	—	—	—	—	—	—	—

[1] tx-credit-not-available and TXWait are configured as a percentage of the polling interval. So, if 10% is configured with a 1 second polling interval, the tx-credit-not-available will alert when the port does not have tx credits available for 100 ms.

[2]
- For all platforms, if the default value for tx-slowport-oper-delay is modified, ISSD to a version lower than Cisco MDS NX-OS Release 6.2(13) will be restricted. To proceed with ISSD, use the **no** form of the **counter tx-slowport-oper-delay** command to roll back to the default value.
- This counter was introduced in Cisco NX-OS Release 6.2(13).

[3]
- For all platforms, if the default value for txwait is modified, ISSD to a version lower than Cisco MDS NX-OS Release 6.2(13) will be restricted. To proceed with ISSD, use the **no** form of the **counter txwait** command to roll back to the default value.

- This counter was introduced in Cisco NX-OS Release 6.2(13).

Table 9: Recommended Units for Port Monitor Policy

Counter	Threshold Type	Interval (Seconds)	Rising Threshold	Event	Falling Threshold	Event	Warning Threshold
link-loss	Delta	Seconds	Number	Event ID	Number	Event ID	Number
sync-loss	Delta	Seconds	Number	Event ID	Number	Event ID	Number
signal-loss	Delta	Seconds	Number	Event ID	Number	Event ID	Number
state-change	Delta	Seconds	Number	Event ID	Number	Event ID	Number
invalid-words	Delta	Seconds	Number	Event ID	Number	Event ID	Number
invalid-crc	Delta	Seconds	Number	Event ID	Number	Event ID	Number
tx-discards	Delta	Seconds	Number	Event ID	Number	Event ID	Number
lr-rx	Delta	Seconds	Number	Event ID	Number	Event ID	Number
lr-tx	Delta	Seconds	Number	Event ID	Number	Event ID	Number
timeout-discards	Delta	Seconds	Number	Event ID	Number	Event ID	Number
credit-loss-reco	Delta	Seconds	Number	Event ID	Number	Event ID	Number
tx-credit-not-available	Delta	Seconds	Percentage	Event ID	Percentage	Event ID	Percentage
rx-datarate	Delta	Seconds	Percentage	Event ID	Percentage	Event ID	Percentage
tx-datarate	Delta	Seconds	Percentage	Event ID	Percentage	Event ID	Percentage
tx-slowport-oper-delay	Absolute	Seconds	Milliseconds	Event ID	Milliseconds	Event ID	Milliseconds
txwait	Delta	Seconds	Percentage	Event ID	Percentage	Event ID	Percentage

Note

- The err-pkt-from-port_ASIC Error Pkt from Port, err-pkt-to-xbar_ ASIC Error Pkt to xbar, and err-pkt-from-xbar_ASIC Error Pkt from xbar counters were introduced in Cisco NX-OS Release 5.2(2a) and are not supported on one rack unit and two rack unit switches.
- We recommend that you use the delta threshold type for all the counters except the tx-slowport-oper-delay counter which uses absolute threshold type.
- The rx-datarate and tx-datarate are calculated using the inoctets and outoctets on an interface.
- The unit for threshold values (rising and falling) differs across different counters.
- tx-slowport-oper-delay wait is applicable only for advanced 16-Gbps and 32-Gbps modules and switches.
- You must configure slow-port monitoring using the **system timeout slowport-monitor** command in order to get alerts for tx-slowport-count and tx-slowport-oper-delay for a particular port type. (See the **system timeout slowport-monitor** command in the Cisco MDS 9000 Series Command Reference.)
- Absolute counters do not support port-guard action. However, tx-slowport-oper-delay counter supports Congestion Isolation port-guard action.
- txwait is applicable only for advanced 16-Gbps and 32-Gbps modules and switches. In the default configuration, the port monitor sends an alert if the transmit credit is not available for 400 ms (40%) in 1 second.

  txwait sends alerts when there are multiple slow-port events that have not hit the slow-port monitor threshold, but have together hit the txwait threshold configured. For example, if there are 40 discrete 10-ms intervals of 0 TX credits in 1 second, tx-slowport-oper-delay does not find these credits; txwait finds the credits and sends an alert.
- The state-change counter records the port down-to-port up action as one state change that is similar to *flap*. This is the reason the state-change counter does not have the portguard action set as *flap*.
- When the portguard action is set as *flap*, you will get alerts only through syslog.
- Only the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters use the **cong-isolate** keywords to detect slow flow on a device. For more information, see Configuring a Port Monitor Policy, on page 51.

The following counters were added from Cisco MDS NX-OS Release 5.2(2a) that are not included in the default policy:

Note

- Crossbar (Xbar) counters are supported only on the Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module (DS-X9448-768K9), Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9), and Cisco MDS 9000 24/10-Port SAN Extension Module (DS-X9334-K9).
- Crossbar counters do not work as expected when check interval is configured.
- Crossbar counters work only when the poll-interval is set to 300 seconds.

- err-pkt-from-port_ASIC Error Pkt from port

- err-pkt-to-xbar_ ASIC Error Pkt to xbar—This counter provides information about the number of error packets that were sent from the crossbar on a module to the crossbar on a supervisor.
- err-pkt-from-xbar_ ASIC Error Pkt from xbar—This counter provides information about the number of error packets that were sent to the crossbar on a module from the crossbar on a supervisor.

Table 10: Slowdrain Port-Monitor Policy Threshold Value, on page 31 displays the threshold value of the slow-drain port-monitor policy:

Table 10: Slowdrain Port-Monitor Policy Threshold Value

Counter	Threshold Type	Interval (Seconds)	Rising Threshold	Event	Falling Threshold	Event	Port Monitor Portguard
Credit Loss Reco	Delta	1	1	4	0	4	Not enabled
TX Credit Not Available	Delta	1	10	4	0	4	Not enabled

Note If no other port monitor policy is explicitly activated, the slowdrain policy is activated. The default policy shows only the default counter monitor values.

Warning Threshold

Port Monitor warning thresholds can be used to generate syslog messages before rising and falling thresholds are reached. A single threshold is configurable per Port Monitor counter. A syslog is generated whenever the counter crosses the configured warning threshold in either the rising or falling direction. This allows the user to track counters that are not severe enough to hit the rising threshold, but where nonzero events are of interest.

The warning threshold must be equal or less than the rising threshold and equal or greater than the falling threshold.

The warning threshold is optional; warning syslogs are only generated when it is specified in a counter configuration.

Use Case—Warning Threshold

Let us consider two scenarios with the following configurations:

- Rising threshold is 30
- Warning threshold is 10
- Falling threshold is 0

This example displays the syslog generated when the error count is less than the rising threshold value, but has reached the warning threshold value:

Syslog Generated When the Error Count is Less Than the Rising Threshold Value

```
%PMON-SLOT2-4-WARNING_THRESHOLD_REACHED_UPWARD: Invalid Words has reached warning threshold
 in the upward direction (port fc2/18 [0x1091000], value = 10).
```

```
%PMON-SLOT2-5-WARNING_THRESHOLD_REACHED_DOWNWARD: Invalid Words has reached warning threshold
 in the downward direction (port fc2/18 [0x1091000], value = 5).
```

In the first polling interval, the errors triggered for the counter (Invalid Words) are 10, and have reached the warning threshold value. A syslog is generated, indicating that the error count is increasing (moving in the upward direction).

In the next polling interval, the error count decreases (moves in the downward direction), and a syslog is generated, indicating that the error count has decreased (moving in the downward direction).

This example displays the syslog that is generated when the error count crosses the rising threshold value:

Syslog Generated When the Error Count Crosses the Rising Threshold Value

```
%PMON-SLOT2-4-WARNING_THRESHOLD_REACHED_UPWARD: Invalid Words has reached warning threshold
 in the upward direction (port fc2/18 [0x1091000], value = 30).

%PMON-SLOT2-3-RISING_THRESHOLD_REACHED: Invalid Words has reached the rising threshold
(port=fc2/18 [0x1091000], value=30).

%SNMPD-3-ERROR: PMON: Rising Alarm Req for Invalid Words counter for port fc2/18(1091000),
 value is 30 [event id 1 threshold 30 sample 2 object 4 fcIfInvalidTxWords]

%PMON-SLOT2-5-WARNING_THRESHOLD_REACHED_DOWNWARD: Invalid Words has reached warning threshold
 in the downward direction (port fc2/18 [0x1091000], value = 3).

%PMON-SLOT2-5-FALLING_THRESHOLD_REACHED: Invalid Words has reached the falling threshold
(port=fc2/18 [0x1091000], value=0).

%SNMPD-3-ERROR: PMON: Falling Alarm Req for Invalid Words counter for port fc2/18(1091000),
 value is 0 [event id 2 threshold 0 sample 2 object 4 fcIfInvalidTxWords]
```

This example displays the syslog generated when the error count is more than the warning threshold value and less than the rising threshold value:

Syslog Generated When the Error Count is More than the Warning Threshold Value and Less than the Rising Threshold Value

```
%PMON-SLOT2-4-WARNING_THRESHOLD_REACHED_UPWARD: Invalid Words has reached warning threshold
 in the upward direction (port fc2/18 [0x1091000], value = 15).

%PMON-SLOT2-5-WARNING_THRESHOLD_REACHED_DOWNWARD: Invalid Words has reached warning threshold
 in the downward direction (port fc2/18 [0x1091000], value = 3).
```

The errors generated for the counter (Invalid Words) are 30 when the counter has crossed both the warning and rising threshold values. A syslog is generated when no further errors are triggered.

As there are no further errors in this poll interval, the consecutive polling interval will have no errors, and the error count decreases (moves in downward direction) and reaches the falling threshold value, which is zero. A syslog is generated for the falling threshold.

Port Monitor Check Interval

Check interval polls for values more frequently within a poll interval so that the errors are detected much earlier and appropriate action can be taken.

With the existing poll interval, it is not possible to detect errors at an early stage. Users have to wait till the completion of the poll interval to detect the errors.

By default, the check interval functionality is not enabled.

Note

- The port monitor check interval feature is supported only on the Cisco MDS 9710 Multilayer Director, Cisco MDS 9718 Multilayer Directors, and Cisco MDS 9706 Multilayer Directors.
- Check interval is supported on both counters, absolute and delta.
- We recommend that you configure the poll interval as a multiple of the check interval.
- When a port comes up, the check interval will not provide an alert regarding invalid words for the port until the poll interval expires. We recommend that you bring up a set of ports at a given time in the module instead of all the ports.

Port Group Monitor

Note Port Group Monitor functionality only applies to modules that support oversubscription.

The ports on a line card are divided into fixed groups called port groups that share a link of fixed bandwidth to the backplane. Since the total port bandwidth can exceed the backplane link bandwidth, frames will be queued, introducing traffic delays. The Port Group Monitor functionality can be used to monitor this oversubscription in both the transmit and receive directions to allow ports to be rebalanced between port groups before the delays become unacceptable.

When the Port Group Monitor feature is enabled and when a policy consisting of polling interval in seconds and the rising and falling thresholds in percentage are specified, the port group monitor generates a syslog if port group traffic goes above the specified percentage of the maximum supported bandwidth for that port group (for receive and for transmit). Another syslog is generated if the value falls below the specified threshold.

Table shows the threshold values for the default Port Group Monitor policy:

Table 11: Default Port Group Monitor Policy Threshold Values

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	% Falling Threshold
RX Datarate	Delta	60	80	20
TX Datarate	Delta	60	80	20

Note When a port group monitor is enabled in a 1-rack box, and if any of the thresholds is met for the receive performance and transmit performance counters, the port group monitor is not supported.

Interface Types

Management Interfaces

You can remotely configure a switch through the management interface (mgmt0). To configure a connection on the mgmt0 interface, configure either the IPv4 parameters (IP address, subnet mask, and default gateway), or the IPv6 parameters (IP address, subnet mask, and default gateway) so that the switch is reachable.

Before you configure the management interface manually, obtain the switch's IPv4 address, subnet mask, and default gateway, or the IPv6 address, depending on which IP version you are configuring.

The management port (mgmt0) auto senses and operates in full-duplex mode at a speed of 10, 100, or 1000 Mbps. Auto sensing supports both the speed mode and the duplex mode. On a Supervisor-1 module, the default speed is 100 Mbps and the default duplex mode is auto. On a Supervisor-2 module, the default speed and the default duplex mode are set to auto.

Note Explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

VSAN Interfaces

VSANs are applicable to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN, and then use this interface to send frames to the corresponding VSAN. To use this feature, configure the IP address for this VSAN.

Note VSAN interfaces cannot be created for non existing VSANs.

Prerequisites for Interfaces

Before you begin configuring the interfaces, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, enter the **show module** command in EXEC mode. For information about verifying the module status, refer to the Cisco MDS 9000 Series NX-OS Fundamentals Configuration Guide.

Guidelines and Limitations

From Cisco MDS NX-OS Release 7.3(x) or earlier, ports were classified as port type access ports, trunks, or all in the port monitor. Access ports were mode (T)F ports and trunks were mode (T)E ports (ISLs). Since ports connecting to Cisco NPV switches are mode (T)F, they were included under the port type access ports. These Cisco NPV ports behave like ISLs, but they are a multi-user connection to a switch and not an end device. Because of this, it is not preferred to take portguard actions on the access ports for port-monitor counters pertaining to slow-drain conditions.

From Cisco MDS NX-OS Release 8.1(1), the port monitor has implemented a different classification mechanism. Instead of port type access ports, trunks, or all, a logical type core, edge, or all value can be configured. Core ports are mode T(E) ports and ports connecting core switches to Cisco NPV switches. Edge ports are mode F ports connecting to end devices. With this new classification, portguard actions can safely be configured especially pertaining to slow drain type conditions such that when the problem is detected and the action is taken, it is only on the ports connected to end devices. It is still valid to configure portguard actions for logical type core ports, but this should only be done for counters pertaining to physical errors on the port (such as link loss, invalid words, invalid CRC, and so on).

The MDS NX-OS will automatically classify all F port-channels and trunking F ports as logical-type core. It will classify all non-trunking F ports, including those to both Cisco and non-Cisco NPV switches, as logical-type edge.

If a Cisco NPV switch or non-Cisco NPV switch cannot take portguard types of actions then classifying the ports connected to it as logical-type edge is appropriate.

The logical type of a port is displayed using the **show interface** and **show interface brief** commands.

Note When you use the **logical-type** command to define a port type, the command overrides the default port type.

In the port monitor, you can configure the policies per port type (core and edge) so that portguard action can be taken on the ports when certain criteria are met. Generally, edge policies are configured to take portguard action on ports and the core policies will not be configured with portguard action. If the link between a core switch and a Cisco NPV switch is treated as an edge port, portguard action is taken on such ports which will result in the loss of connectivity to all the devices connected to the Cisco NPV switch.

For any Cisco NPV switch that supports its own Port Monitor policies, it is best to implement these portguard actions on the Cisco NPV switch itself. Hence, we recommend that all non-trunking F ports connected to Cisco NPV switches be manually configured to a logical type of core, using the **switchport logical-type core** command. This will ensure that port monitor core policy is applied to the port connected to a Cisco NPV switch. We also recommend that Port Monitor be implemented on the Cisco NPV switch, if supported.

For more information, see Interface Modes, on page 15.

Guidelines for Configuring Port Monitor Check Interval

- Check interval should be configured before activating any port monitor policies.

Note The value of the check interval is common across counters and policies.

- Check interval should be less than the poll interval.
- Check interval is applicable to all the active port monitor policies configured.
- Users should deactivate all the active port monitor policies before enabling, modifying, or disabling the check interval functionality.
- Check interval cannot be enabled when an active policy is configured.
- Software downgrade to a version that does not support the check interval functionality is restricted when the check interval functionality is enabled.
- We recommend that you do not have a portguard action set to the state-change counter when an interface state is changed from down state to up state.
- We recommend that you do not use the default policy when the check interval is configured.

Check Interval

Let us consider a scenario where the poll interval, rising threshold and check interval are configured with the following values:

- Poll interval is 100 seconds
- Rising threshold is 30
- Check interval is 20 seconds

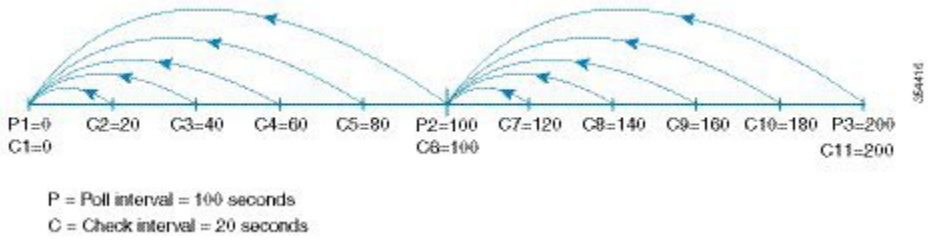

The check interval starts its interval, C1, along with the poll interval at P1. If an error occurs between the check intervals C2 and C3, the check intervals C2 and C3 are higher than the configured rising threshold value of 30, an alert (syslog or trap or both) is generated at C3, alerting the user that an error has occurred at that particular port.

Note You can configure longer poll intervals to capture events across poll intervals. For example, configure a poll interval of 24 hours with a check interval of 30 seconds, with the rising threshold value being checked cumulatively every 30 seconds.

Guidelines for VSAN Interface Configuration

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN; it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.

Tip After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature. See the Cisco MDS 9000 Series NX-OS IP Services Configuration Guide.

Guidelines and Limitations for Port Beaconing

- The port beacon LED on directly attached peers can only be controlled when the link to the peer is up and operational.
- If you enable port beacon mode on a port using the **beacon interface** command and then enable beacon mode using the **switchport beacon** command, the beacon mode takes precedence and the port beacon mode will be disabled. If you disable the beacon mode, the port beacon mode will continue to be disabled until you enable the port beacon mode again.
- If you send a port beaconing request from Switch A to Switch B using the **beacon interface** command and then if you enable **switchport beacon** locally on Switch B, the **switchport beacon** command takes precedence over the port beaconing request and stops the LED activity on Switch B. However, if you run the **show interface** command on Switch A, the output will continue to show the port beaconing status for the port on Switch B until the specified duration is reached.
- If you enable port beacon mode on a port using the **beacon interface** command and then perform a system switchover using the **system switchover** command, the **show interface** command on the switch does not show the port beaconing status as on. However, the port LED to which the port beaconing request was sent continues to beacon with the specified parameters until the specified duration is reached or when you run the **switchport beacon** command to override the port beaconing request for the port.
- If you send a port beaconing request with the duration set to 0 from Switch A that is running Cisco MDS NX-OS Release 8.3(1) or later releases to Switch B and then downgrade Switch A to Cisco MDS NX-OS Release 8.2(2) or earlier releases, the port LED on Switch B to which the port beaconing request was sent continues to beacon with the specified parameters until you run the **switchport beacon** command to override the port beaconing request for the port on Switch B.
- This feature is not supported on Cisco MDS switches that are operating in the Cisco NPV mode.
- This feature is not supported on port-channel interfaces. It is supported only on individual Fibre Channel interfaces or port-channel members.

Default Settings

Table 12: Default Interface Parameters , on page 39 lists the default settings for interface parameters.

Table 12: Default Interface Parameters

Parameters	Default
Interface mode	Auto
Interface speed	Auto
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	On (unless changed during initial setup) on non-NPV and NPIV core switches. Off on NPV switches.
Trunk-allowed VSANs or VF-IDs	1 to 4093
Interface VSAN	Default VSAN (1)
Beacon mode	Off (disabled)
EISL encapsulation	Disabled
Data field size	2112 bytes

Configuring Interfaces

For more information on configuring mgmt0 interfaces, refer to the Cisco MDS 9000 Series NX-OS Fundamentals Configuration Guide and Cisco MDS 9000 Series NX-OS IP Services Configuration Guide.

For more information on configuring Gigabit Ethernet interfaces, see the Cisco MDS 9000 Series NX-OS IP Services Configuration Guide.

Configuring a Fibre Channel Interface

To configure a Fibre Channel interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc 1/1**

When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

Configuring a Range of Fibre Channel Interfaces

To configure a range of interfaces, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the range of Fibre Channel interfaces and enter interface configuration submode3:

switch(config)# **interface fc1/1 - 4 , fc2/1 - 3**

Note When using this command, provide a space before and after the comma.

Setting the Interface Administrative State

To set the interface administrative state, you must first gracefully shut down the interface and enable traffic flow.

Shutting Down an Interface

To gracefully shut down an interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Gracefully shut down the interface and administratively disable the traffic flow; this is the default state

switch(config-if)# **shutdown**

Enabling Traffic Flow

To enable traffic flow, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Enable traffic flow to administratively allow traffic when the no prefix is used (provided the operational state is up):

switch(config-if)# **no shutdown**

Configuring an Interface Mode

To configure the interface mode, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Configure the administrative mode of the port. You can set the operational state to auto, E, F, FL, Fx, TL, NP, or SD port mode:

switch(config-if)# **switchport mode F**

Note Fx ports refer to an F port or an FL port (host connection only), but not E ports.

Step 4 Configure interface mode to auto negotiate an E, F, FL, or TE port mode (not TL or SD port modes) of operation:

switch(config-if)# **switchport mode auto**

Note
- TL ports and SD ports cannot be configured automatically. They must be administratively configured.
- You cannot configure Fibre Channel interfaces on Storage Services Modules (SSM) in auto mode.

Configuring the MAX NPIV Limit

Note Both the **max-npiv-limit** and **trunk-max-npiv-limit** can be configured on a port or port channel. If the port or port channel becomes a trunking port, **trunk-max-npiv-limit** is used for limit checks.

To configure the maximum NPIV limit, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc 3/29**

Step 3 Configure switch port mode F on the Fibre Channel interface:

switch(config-if)# **switchport mode F**

Step 4 Specify the maximum login value for this port:

switch(config-if)# **switchport max-npiv-limit 100**

The valid range is from 1 to 256.

Configuring the System Default F Port Mode

The **system default switchport mode F** command sets the administrative mode of all Fibre Channel ports to mode F, while avoiding traffic disruption caused by the formation of unwanted ISLs. This command is part of the setup utility that runs during bootup after a **write erase** or **reload** command is issued. It can also be executed from the command line in configuration mode. This command changes the configuration of the following ports to administrative mode F:

- All ports that are down and that are not out of service.
- All F ports that are up, whose operational mode is F, and whose administrative mode is not F.

The **system default switchport mode F** command does not affect the configuration of the following ports:

- All user-configured ports, even if they are down.
- All non-F ports that are up. However, if non-F ports are down, this command changes the administrative mode of those ports.

Note

- To ensure that ports that are a part of ISLs do not get changed to port mode F, configure the ports in port mode E, rather than in auto mode.
- When the command is executed from the command line, the switch operation remains graceful. No ports are flapped.

To set the administrative mode of Fibre Channel ports to mode F in the CLI, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Sets administrative mode of Fibre Channel ports to mode F (if applicable):

switch(config)# **system default switchport mode F**

(Optional) Set the administrative mode of Fibre Channel ports to the default (unless user configured), use the following command:

switch(config)# **no system default switchport mode F**

Note For detailed information about the switch setup utility, see the Cisco MDS 9000 Series NX-OS Fundamentals Configuration Guide.

Setup Utility

Setup Utility , on page 43 shows the command in the setup utility and the command from the command line.

```
Configure default switchport mode F (yes/no) [n]: y

switch(config)# system default switchport mode F
```

Configuring ISL Between Two Switches

Note Ensure that the Fibre Channel cable is connected between the ports and perform a no-shut operation on each port.

E-port mode is used when a port functions as one end of an ISL setting. When you set the port mode to E, you restrict the port coming up as an E port (trunking or nontrunking, depending on the trunking port mode).

To configure the port mode to E:

Step 1 Enter configuration mode:

switch#**configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc 3/29**

Step 3 Configure switch port mode E on the Fibre Channel interface:

switch(config)# **switchport mode E**

Note Ensure that you perform the task of setting the port mode to E on both the switches between which you are attempting to bring up the ISL link.

Configuring the Port Administrative Speeds

Note Changing the port administrative speed is a disruptive operation.

To configure the port speed of the interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the Fibre Channel interface and enter interface configuration mode:

switch(config)# **interface fc 1/1**

Step 3 Configure the port speed of the interface to 1000 Mbps:

switch(config-if)# **switchport speed 1000**

All the 10-Gbps capable interfaces, except the interface that is being configured, must be in the out-of-service state. At least one other 10-Gbps capable interface must be in the in-service state.

(Optional) Revert to the factory default (auto) administrative speed of the interface:

switch(config-if)# **no switchport speed**

Configuring Port Speed Group

To configure the port speed group of the interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the Fibre Channel interface and enter interface configuration mode:

switch(config)# **interface fc 1/1**

Step 3 Configure the port speed group to 10 Gbps:

switch(config-if)# **speed group 10g**

The preferred way of changing the speed group is the **10g-speed-mode** command.

(Optional) Unset the port speed group and revert to the factory default (auto) administrative speed group of the interface:

switch(config-if)# **no speed group 10g**

Configuring the Interface Description

The interface description can be any alphanumeric string that is up to 80 characters long.

To configure a description for an interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Configure the description of the interface:

switch(config-if)# **switchport description cisco-HBA2**

The string can be up to 80 characters long.

(Optional) Clear the description of the interface:

switch(config-if)# **no switchport description**

Configuring a Port Logical Type

The logical port type can be used to override the default type assigned by the Cisco NX-OS to a port. Previously, point to point F and TF ports were used by a single edge device with a single login to the switch. With the adoption of the Cisco NPV technology, these types of switch ports can now have multiple logins from multiple edge devices on a single port. In such cases, the ports are no longer dedicated to a single edge device, but are shared by multiple devices similar to Inter-Switch Links (ISLs). The **switchport logical-type** command allows the you to change the port type so that port monitor and congestion timeout features apply core type policies and not the more aggressive edge type policies to such links.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Configure a logical type for an interface:

switch(config-if)# **switchport logical-type** {**auto** | **core** | **edge**}

(Optional) Remove the logical type from an interface:

switch(config-if)# **no switchport logical-type** {**auto** | **core** | **edge**}

Specifying a Port Owner

Using the Port Owner feature, you can specify the owner of a port and the purpose for which a port is used so that the other administrators are informed.

Note The Portguard and Port Owner features are available for all ports regardless of the operational mode.

To specify or remove a port owner, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the port interface:

switch(config)# **interface fc1/1**

Step 3 Specify the owner of the switch port:

switch(config)# **switchport owner** *description*

The description can include the name of the owner and the purpose for which the port is used, and can be up to 80 characters long.

(Optional) Remove the port owner description:

switch(config)# **no switchport owner**

(Optional) Display the owner description specified for a port, use one of the following commands:

- switch# **show running interface fc** *module-number/interface-number*
- switch# **show port internal info interface fc** *module-number/interface-number*

Configuring Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface. Note that configuring the beacon mode has no effect on the operation of the interface.

To configure a beacon mode for a specified interface or range of interfaces, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc1/1**

Step 3 Enable the beacon mode for the interface:

switch(config-if)# **switchport beacon**

(Optional) Disable the beacon mode for the interface:

switch(config-if)# **no switchport beacon**

Tip The flashing green light turns on automatically when an external loopback that causes the interfaces to be isolated is detected. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

Configuring the Port Beacon LED

To configure the port beacon LEDs on one or both ends of a link, perform this step:

switch# **beacon interface fc** *slot/port* {**both** | **local** | **peer**} [**status** {**normal** | **warning** | **critical**}] [**duration** *seconds*] [**frequency** *number*]

Configuring a Switch Port Attribute Default Value

You can configure default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

To configure a default value for a switch port attribute, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Configure the default setting for the administrative state of an interface as up (the factory default setting is down):

switch(config)# **no system default switchport shutdown**

Note This command is applicable only to interfaces for which no user configuration exists for the administrative state.

(Optional) Configure the default setting for the administrative state of an interface as down:

switch(config)# **system default switchport shutdown**

Note This command is applicable only to interfaces for which no user configuration exists for the administrative state.

(Optional) Configure the default setting for the administrative trunk mode state of an interface as Auto:

switch(config)# **system default switchport trunk mode auto**

Note The default setting is On.

Configuring the Port-Level Portguard

All portguard causes are monitored over a common time interval with the same start and stop times. The *link down* counter is not a specific event, but the aggregation of all other cause counters in the same time interval.

To configure a port-level portguard for a interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the interface:

switch(config)# **interface fc1/1**

Step 3 Enable portguard error disabling of the interface if the link goes down once:

switch(config-if)# **errdisable detect cause link-down**

(Optional) Enable portguard error disabling of the interface if the link flaps a certain number of times within the specified time, in *seconds*:

switch(config-if)# **errdisable detect cause link-down** [**num-times** *number* **duration** *seconds*]

(Optional) Remove the portguard configuration for the interface:

switch(config-if)# **no errdisable detect cause link-down**

The link resumes flapping and sending error reports normally.

Step 4 Enable portguard error disabling of the interface if the specified error occurs once:

switch(config-if)# **errdisable detect cause** {**trustsec-violation** | **bit-errors** | **credit-loss** | **link-reset** | **signal-loss** | **sync-loss**}

(Optional) Enable portguard error disabling of the interface if the specified error occurs a certain number times within the specified time, in *seconds*:

switch(config-if)# **errdisable detect cause** {**trustsec-violation** | **bit-errors** | **credit-loss** | **link-reset** | **signal-loss** | **sync-loss**} [**num-times** *number* **duration** *seconds*]

(Optional) Remove the portguard configuration for the interface:

switch(config-if)# **no errdisable detect cause** {**trustsec-violation** | **bit-errors** | **credit-loss** | **link-reset** | **signal-loss** | **sync-loss**}

The link resumes flapping and sending error reports normally.

Note The portguard credit loss event is triggered only on loop interfaces; it is not triggered on point-to-point interfaces.

This example shows how to configure portguard to set an interface to Error Disabled state if the link flaps five times within 120 seconds due to multiple causes. The portguard controls the interface in the following manner:

- The interface will be error disabled due to link down if there are link failures due to bit errors 2 times and link failures due to credit loss 3 times in 120 seconds.
- The interface will be error disabled due to bit errors if there are link failures due to bit errors 5 times in 120 seconds.
- The interface will be error disabled due to credit loss if there are link failures due to credit loss 5 times in 120 seconds.

Example

This example shows how to configure portguard to bring a port to down state if the link flaps 5 times within 120 seconds based on multiple causes:

```
switch# configure terminal

switch (config)# interface fc1/1

switch (config-if)# errdisable detect cause link-down num-times 5 duration 120

switch (config-if)# errdisable detect cause bit-errors num-times 5 duration 120

switch (config-if)# errdisable detect cause credit-loss num-times 5 duration 120
```

The above example sets the configuration to the following status:

- The port will be error disabled due to bit errors if the port suffers link failure due to bit errors 5 times in 120 seconds.
- The port will be error-disabled due to credit loss if the port suffers link failure due to credit loss 5 times in 120 seconds.
- The port will be error-disabled due to link down if the port suffers link failure due to bit errors 2 times and link-failure due to credit loss 3 times in 120 seconds.

This example shows the internal information about a port in down state because of TrustSec violation:

```
switch# show interface fc1/9
fc1/9 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:09:54:7f:ee:eb:dc:00
    Peer port WWN is 20:49:8c:60:4f:53:bb:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is auto max 16 Gbps
    Operating Speed is 4 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 500
    Receive B2B Credit is 500
    B2B State Change Number is 14
```

```
Receive data field Size is 2112
Beacon is turned off
Logical type is core
Belongs to port-channel2
Trunk vsans (admin allowed and active) (1-2,5)
Trunk vsans (up)                       (1-2)
Trunk vsans (isolated)                 (5)
Trunk vsans (initializing)             ()
5 minutes input rate 448 bits/sec,56 bytes/sec, 0 frames/sec
5 minutes output rate 384 bits/sec,48 bytes/sec, 0 frames/sec
  783328 frames input,58490580 bytes
    0 discards,0 errors
    0 invalid CRC/FCS,0 unknown class
    0 too long,0 too short
  783799 frames output,51234876 bytes
    0 discards,0 errors
  56 input OLS,63  LRR,8 NOS,277 loop inits
  49 output OLS,27 LRR, 49 NOS, 43 loop inits
  500 receive B2B credit remaining
  500 transmit B2B credit remaining
  500 low priority transmit B2B credit remaining
Last clearing of "show interface" counters :  never
```

Tip

- Link down is the superset of all other causes. A port is brought to down state if the total number of other causes equals to the number of allowed link-down failures.
- Even if the link does not flap due to failure of the link, and portguard is not enabled, the port goes into a down state if too many invalid FLOGI requests are received from the same host. Use the **shut** and the **no shut** commands consecutively to bring up the link.

Configuring a Port Monitor

Configuring a portguard action is optional for each counter in a port monitor policy, and is disabled by default.

Enabling a Port Monitor

To enable or disable a port monitor, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable port monitoring:

switch(config)# **port-monitor enable**

(Optional) Disable port monitoring:

switch(config)# **no port-monitor enable**

Configuring the Check Interval

To configure the check interval, perform these steps:

Step 1 Enter the configuration mode:

switch# **configure terminal**

Step 2 Configure the check interval time to 30 seconds

switch# **port-monitor check-interval** *30*

To disable check interval use the following command:

switch# **no port-monitor check-interval**

Configuring a Port Monitor Policy

To configure a port monitor policy, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the policy name and enter port monitoring policy configuration mode:

switch(config)# **port-monitor name** *policyname*

(Optional) Remove the policy name:

switch(config)# **no port-monitor name** *policyname*

Step 3 Apply policy type:

switch(config-port-monitor)# **logical-type** {**core** | **edge** | **all**}

Step 4 Specify the counter parameters:

switch(config-port-monitor)# **counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-discards** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *RMON-ID* **warning-threshold** *count2* **falling-threshold** *count3* **event** *RMON-ID* **portguard** {**errordisable** | **flap** | **cong-isolate**}

Note

- We recommend that you use the delta threshold type for all the counters except the tx-slowport-oper-delay counter which uses absolute threshold type.
- The rx-datarate and tx-datarate are calculated using the inoctets and outoctets on an interface.
- You must activate the **err-pkt-from-port**, **err-pkt-from-xbar**, and **err-pkt-to-xbar** counters using the **monitor counter** *name* command, before specifying the counter parameters.
- Counters **err-pkt-from-xbar**, **err-pkt-from-port**, and **err-pkt-to-xbar** support delta threshold type only.
- Counter **tx-slowport-oper-delay** supports **absolute** threshold type only.
- Counter **tx-slowport-oper-delay** does not support portguard action.
- You must first enable ER_RDY flow-control mode using the **system fc flow-control er_rdy** command and then enable congestion isolation using the **feature congestion-isolation** command before setting the portguard action as congestion isolate (cong-isolate). For more information, see Configuring the Port-Monitor Portguard Action for Congestion Isolation, on page 163.
- The **cong-isolate** port monitor portguard action is used only for configuring the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters.

(Optional) Revert to the default values for a counter:

switch(config-port-monitor)# **no counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-discards** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *RMON-ID* **warning-threshold** *count2* **falling-threshold** *count3* **event** *RMON-ID* **portguard** {**errordisable** | **flap** | **cong-isolate**}

(Optional) Monitor a counter:

switch(config-port-monitor)# **monitor counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-discards** | **tx-slowport-oper-delay** | **txwait**}

A port monitor currently recognizes two kinds of ports:

- Logical-type edge ports are normally F ports that are connected to end devices.
- Logical-type core ports are E ports (ISLs) or (T)F ports connected to Cisco NPV switches. Some of the edge port counter thresholds and port-guard actions might not be appropriate on the TF ports in the port-monitor configurations. Specifically, portguard *disable*, *flap*, and *isolate* actions can affect multiple end devices on the F ports. Therefore, performing disable, flap, or isolate actions should be avoided on an N-Port Identifier Virtualization (NPIV) system.

Activating a Port Monitor Policy

To activate a port monitor policy, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Activate the specified port monitor policy:

switch(config)# **port-monitor activate** *policyname*

(Optional) Activate the default port monitor policy:

switch(config)# **port-monitor activate**

(Optional) Deactivate the specified port monitoring policy:

switch(config)# **no port-monitor activate** *policyname*

Configuring Port Monitor Portguard

To configure a port monitor portguard action, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the policy name and enter port monitoring policy configuration mode:

switch(config)# **port-monitor name** *policyname*

(Optional) Remove the policy:

switch(config)# **no port-monitor name** *policyname*

Step 3 Specify a counter, its parameters, and a portguard action for a counter:

switch(config-port-monitor)# **counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-discards** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *RMON-ID* **warning-threshold** *count2* **falling-threshold** *count3* **event** *RMON-ID* **portguard** {**errordisable** | **flap** | **cong-isolate**}

Note

- We recommend that you use the delta threshold type for all the counters except the tx-slowport-oper-delay counter which uses absolute threshold type.
- The rx-datarate and tx-datarate are calculated using the inoctets and outoctets on an interface.
- You must activate the **err-pkt-from-port**, **err-pkt-from-xbar**, and **err-pkt-to-xbar** counters using the **monitor counter** *name* command, before specifying the counter parameters.
- Counters **err-pkt-from-xbar**, **err-pkt-from-port**, and **err-pkt-to-xbar** support delta threshold type only.
- Counter **tx-slowport-oper-delay** supports **absolute** threshold type only.
- Counter **tx-slowport-oper-delay** does not support portguard action.
- You must first enable ER_RDY flow-control mode using the **system fc flow-control er_rdy** command and then enable congestion isolation using the **feature congestion-isolation** command before setting the portguard action as congestion isolate (cong-isolate). For more information, see Configuring the Port-Monitor Portguard Action for Congestion Isolation, on page 163.
- The **cong-isolate** port monitor portguard action is used only for configuring the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters.

Configuring Port Group Monitor

Enabling a Port Group Monitor

To enable a port group monitor, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable port monitoring:

switch(config)# **port-group-monitor enable**

(Optional) Disable port monitoring:

switch(config)# **no port-group-monitor enable**

Configuring a Port Group Monitor Policy

To configure a port group monitor policy, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the policy name and enter port group monitoring policy configuration mode:

switch(config)# **port-group-monitor name** *policyname*

(Optional) Remove the policy:

switch(config)# **no port-group-monitor name** *policyname*

Step 3 Specify the delta receive or transmit counter poll interval (in seconds) and thresholds (in percentage):

switch(config-port-group-monitor)# **counter** {**rx-datarate** | **tx-datarate**} **poll-interval** *seconds* **delta rising-threshold** *percentage1* **falling-threshold** *percentage2*

(Optional) Revert to the default policy:

switch(config-port-group-monitor)# **no counter tx-datarate**

For more information on reverting to the default policy, see Reverting to the Default Policy for a Specific Counter and Port Group Monitor.

Step 4 Turn on datarate monitoring:

switch(config-port-group-monitor)# **monitor counter** {**rx-datarate** | **tx-datarate**}

(Optional) Turn off datarate monitoring:

switch(config-port-group-monitor)# **no monitor counter** {**rx-datarate** | **tx-datarate**}

For more information on turning off transmit datarate monitoring, see Turning Off Specific Counter Monitoring.

Note On 8-Gbps and higher speed modules, port errors are monitored using the **invalid-crc** and **invalid-words** counters. The **err-pkt-from-port** counter is supported only on 4-Gbps modules.

Reverting to the Default Policy for a Specific Counter

The following examples display the default values for counters:

```
switch(config)# port-group-monitor name PGMON_policy
switch(config-port-group-monitor)# counter tx-datarate poll-interval 200 delta
rising-threshold 75 falling-threshold 0
 switch(config)# show port-group-monitor PGMON_policy
Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
-------------------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ---------------------
RX Datarate   Delta      200       75                    0
TX Datarate   Delta      60        80                    20

switch(config-port-group-monitor)# no counter tx-datarate
switch(config)# show port-group-monitor PGMON_policy
Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
-------------------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ---------------------
RX Datarate   Delta      60       80                   10
```

```
TX Datarate   Delta      60       80                  10
-------------------------------------------------------------------------------------------
```

Turning Off Specific Counter Monitoring

The following examples display turning off counter monitoring:

```
switch(config)# port-group-monitor name PGMON_policy
switch(config-port-group-monitor)# no monitor counter rx-datarate
switch(config)# show port-group-monitor PGMON_policy
Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
-------------------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
TX Datarate   Delta      60       100                 80
-------------------------------------------------------------------------------------------
```

Activating a Port Group Monitor Policy

To activate a port group monitor policy, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Activate the specified port group monitor policy:

switch(config)# **port-group-monitor activate** *policyname*

(Optional) Activate the default port group monitor policy:

switch(config)# **port-group-monitor activate**

(Optional) Deactivate the specified port group monitor policy:

switch(config)# **no port-group-monitor activate** *policyname*

Configuring Management Interfaces

Configuring the Management Interface Over IPv4

To configure the mgmt0 Ethernet interface to connect over IPv4, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the management Ethernet interface on the switch and enter interface configuration submode:

switch(config)# **interface mgmt0**

Step 3 Configure the IPv4 address and IPv4 subnet mask:
switch(config-if)# **ip address 10.16.1.2 255.255.255.0**

Step 4 Enable the interface:
switch(config-if)# **no shutdown**

Step 5 Return to configuration mode:
switch(config-if)# **exit**

Step 6 Configure the default gateway IPv4 address:
switch(config)# **ip default-gateway 1.1.1.4**

Step 7 Return to user EXEC mode:
switch(config)# **exit**
(Optional) Save your configuration changes to the file system:
switch# **copy running-config startup-config**

Configuring the Management Interface Over IPv6

To configure the mgmt0 Ethernet interface to connect over IPv6, perform these steps:

Step 1 Enter configuration mode:
switch# **configure terminal**

Step 2 Select the management Ethernet interface on the switch and enter interface configuration submode:
switch(config)# **interface mgmt0**

Step 3 Enable IPv6 and assign a link-local address on the interface:
switch(config-if)# **ipv6 enable**

Step 4 Specify an IPv6 unicast address and prefix length on the interface:
switch(config-if)# **ipv6 address 2001:0db8:800:200c::417a/64**

Step 5 Enable the interface:
switch(config-if)# **no shutdown**

Step 6 Return to user EXEC mode:
switch(config)# **exit**
(Optional) Save your configuration changes to the file system:
switch# **copy running-config startup-config**

Creating VSAN Interfaces

To create a VSAN interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Configure a VSAN with the ID 2:

switch(config)# **interface vsan 2**

Step 3 Enable the VSAN interface:

switch(config-if)# **no shutdown**

Verifying Interfaces Configuration

Displaying Interface Information

Run the **show interface** command from user EXEC mode. This command displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch.

The following example displays the status of interfaces:

Displays All Interfaces

```
switch# show interface
fc1/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:01:54:7f:ee:de:c5:00
    Admin port mode is SD
    snmp link state traps are enabled
    Port mode is SD
    Port vsan is 1
    Admin Speed is 8 Gbps
    Operating Speed is 8 Gbps
    Rate mode is dedicated
    Beacon is turned off
    Logical type is Unknown(0)
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      4 frames input,304 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      4 frames output,304 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      0 output OLS,0 LRR, 0 NOS, 0 loop inits
      1 receive B2B credit remaining
      0 transmit B2B credit remaining
      0 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:10:49 2017

    Last clearing of "show interface" counters :  never
.
.
.
fc3/8 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:88:54:7f:ee:de:c5:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Admin Speed is auto max 32 Gbps
    Operating Speed is 16 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 64
    Receive B2B Credit is 32
```

```
    Receive data field Size is 2112
    Beacon is turned off
    Logical type is core
    Trunk vsans (admin allowed and active) (1-7,200,400)
    Trunk vsans (up)                       (1-2)
    Trunk vsans (isolated)                 (6-7,200,400)
    Trunk vsans (initializing)             (3-5)
    5 minutes input rate 13438472736 bits/sec,1679809092 bytes/sec, 779072 frames/sec
    5 minutes output rate 13438477920 bits/sec,1679809740 bytes/sec, 779073 frames/sec
      99483764407 frames input,213691124011124 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      99485576094 frames output,213695013798564 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,1 LRR, 0 NOS, 0 loop inits
      32 receive B2B credit remaining
      62 transmit B2B credit remaining
      62 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:11:47 2017

    Last clearing of "show interface" counters :  never
.
.
.
fc3/15 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:8f:54:7f:ee:de:c5:00
    Admin port mode is F, trunk mode is off
    snmp link state traps are enabled
    Port mode is F, FCID is 0xe003c0
    Port vsan is 1
    Admin Speed is auto max 32 Gbps
    Operating Speed is 16 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 80
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Logical type is edge
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      29 frames input,2600 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      36 frames output,2948 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,1 LRR, 0 NOS, 0 loop inits
      32 receive B2B credit remaining
      80 transmit B2B credit remaining
      80 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:11:50 2017

    Last clearing of "show interface" counters :  never
```

You can also specify arguments (a range of interfaces or multiple specified interfaces) to display interface information. You can specify a range of interfaces by issuing a command in the following format:

interface fc1/1 - 5 , fc2/5 - 7

Note The spaces are required before and after the dash (-) and before and after the comma (,).

The following example displays the status of a range of interfaces:

Displays Multiple, Specified Interfaces

```
switch# show interface fc3/9 , fc3/12
fc3/9 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:89:54:7f:ee:de:c5:00
    Peer port WWN is 20:09:00:2a:6a:a4:0b:00
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is auto
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is ER_RDY

    Transmit B2B Credit for vl0 is 15
    Transmit B2B Credit for vl1 is 15
    Transmit B2B Credit for vl2 is 40
    Transmit B2B Credit for vl3 is 430
    Receive B2B Credit for vl0 is 15
    Receive B2B Credit for vl1 is 15
    Receive B2B Credit for vl2 is 40
    Receive B2B Credit for vl3 is 430
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    FCSP Status: Successfully authenticated
    Trunk vsans (admin allowed and active) (1-7,200,400)
    Trunk vsans (up)                       (1-7)
    Trunk vsans (isolated)                 (200,400)
    Trunk vsans (initializing)             ()
    5 minutes input rate 1175267552 bits/sec,146908444 bytes/sec, 67007 frames/sec
    5 minutes output rate 1175268256 bits/sec,146908532 bytes/sec, 67005 frames/sec
      8563890817 frames input,18703349820904 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      8563735031 frames output,18703009725636 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,3 LRR, 0 NOS, 0 loop inits
      70 receive B2B credit remaining
      500 transmit B2B credit remaining
      485 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:11:49 2017

    Last clearing of "show interface" counters :  never

fc3/12 is trunking
```

```
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:8c:54:7f:ee:de:c5:00
    Peer port WWN is 20:0c:00:2a:6a:a4:0b:00
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is auto
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is ER_RDY

    Transmit B2B Credit for vl0 is 15
    Transmit B2B Credit for vl1 is 15
    Transmit B2B Credit for vl2 is 40
    Transmit B2B Credit for vl3 is 430
    Receive B2B Credit for vl0 is 15
    Receive B2B Credit for vl1 is 15
    Receive B2B Credit for vl2 is 40
    Receive B2B Credit for vl3 is 430
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    FCSP Status: Successfully authenticated
    Trunk vsans (admin allowed and active) (1-7,200,400)
    Trunk vsans (up)                       (1-7)
    Trunk vsans (isolated)                 (200,400)
    Trunk vsans (initializing)             ()
    5 minutes input rate 1175267840 bits/sec,146908480 bytes/sec, 67008 frames/sec
    5 minutes output rate 1175265056 bits/sec,146908132 bytes/sec, 67007 frames/sec
      8564034952 frames input,18703367929364 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      8563736100 frames output,18703012026724 bytes
        0 discards,0 errors
      1 input OLS,1  LRR,1 NOS,0 loop inits
      1 output OLS,2 LRR, 0 NOS, 0 loop inits
      70 receive B2B credit remaining
      500 transmit B2B credit remaining
      485 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:11:50 2017

    Last clearing of "show interface" counters :  never
```

The following example displays the status of a specified interface:

Displays a Specific Interface

```
switch# show interface fc3/9
fc3/9 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:89:54:7f:ee:de:c5:00
    Peer port WWN is 20:09:00:2a:6a:a4:0b:00
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
```

```
    Port vsan is 1
    Admin Speed is auto
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is ER_RDY

    Transmit B2B Credit for vl0 is 15
    Transmit B2B Credit for vl1 is 15
    Transmit B2B Credit for vl2 is 40
    Transmit B2B Credit for vl3 is 430
    Receive B2B Credit for vl0 is 15
    Receive B2B Credit for vl1 is 15
    Receive B2B Credit for vl2 is 40
    Receive B2B Credit for vl3 is 430
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    FCSP Status: Successfully authenticated
    Trunk vsans (admin allowed and active) (1-7,200,400)
    Trunk vsans (up)                       (1-7)
    Trunk vsans (isolated)                 (200,400)
    Trunk vsans (initializing)             ()
    5 minutes input rate 1175263296 bits/sec,146907912 bytes/sec, 67007 frames/sec
    5 minutes output rate 1175266272 bits/sec,146908284 bytes/sec, 67007 frames/sec
      8570830922 frames input,18718506849280 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      8570675128 frames output,18718166747180 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,3 LRR, 0 NOS, 0 loop inits
      70 receive B2B credit remaining
      500 transmit B2B credit remaining
      485 low priority transmit B2B credit remaining
    Interface last changed at Mon Apr 24 23:11:49 2017

    Last clearing of "show interface" counters :  never
```

The following example displays the description of interfaces:

Displays Port Description

```
switch# show interface description
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
fc3/1              test intest
fc3/2              --
fc3/3              --
fc3/4              TE port
fc3/5              --
fc3/6              --
fc3/10             Next hop switch 5
fc3/11             --
fc3/12             --
fc3/16             --
```

```
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
port-channel 1     --
port-channel 5     --
port-channel 6     --
```

The following example displays a summary of information:

Displays Interface Information in a Brief Format

```
switch# show interface brief
-------------------------------------------------------------------------------------------

Interface  Vsan   Admin  Admin   Status       SFP    Oper  Oper   Port     Logical

                  Mode   Trunk                       Mode  Speed  Channel   Type
                         Mode                               (Gbps)
-------------------------------------------------------------------------------------------

fc1/1       1      E      on      up           swl    E     8      --       core

fc1/2       1      auto   on      sfpAbsent    --     --    --     --     --

fc1/3       1      F      on      up           swl    F     8      --       core
```

The following example displays a summary of information:

Displays Interface Counters

```
switch# show interface counters
fc3/1
    5 minutes input rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
    5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
    3502 frames input, 268400 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    3505 frames output, 198888 bytes
      0 discards
    1 input OLS, 1 LRR, 1 NOS, 0 loop inits
    2 output OLS, 1 LRR, 1 NOS, 0 loop inits
    1 link failures, 1 sync losses, 1 signal losses
.
.
.
fc9/8
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
```

```
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
sup-fc0
    114000 packets input, 11585632 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    113997 packets output, 10969672 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
mgmt0
    31557 packets input, 2230860 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    26618 packets output, 16824342 bytes, 0 underruns
      0 output errors, 0 collisions, 7 fifo
      0 carrier errors
vsan1
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
```

Note Interfaces 9/8 and 9/9 are not trunking ports and display Class 2, 3, and F information as well.

The following example displays the brief counter information of interfaces:

Displays Interface Counters in Brief Format

```
switch# show interface counters brief
-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------ ------------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
-------------------------------------------------------------------------------
```

```
fc3/1              0        3871                  0        3874
fc3/2              0        3902                  0        4232
fc3/3              0        3901                  0        4138
fc3/4              0        3895                  0        3894
fc3/5              0        3890                  0        3897
fc9/8              0        0                     0        0
fc9/9              0        5                     0        4
fc9/10             0        4186                  0        4182
fc9/11             0        4331                  0        4315
-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------ ------------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
-------------------------------------------------------------------------------
port-channel 1     0        0                     0        0
port-channel 2     0        3946                  0        3946
```

You can run the **show interface transceiver** command only on a switch in the Cisco MDS 9100 Series if the SFP is present, as show in the following example:

Displays Transceiver Information

```
switch# show interface transceiver

fc1/1 SFP is present
    name is CISCO-AGILENT
    part number is QFBR-5796L
    revision is
    serial number is A00162193
    fc-transmitter type is short wave laser
    cisco extended id is unknown (0x0)
...
fc1/9 SFP is present
    name is FINISAR CORP.
    part number is FTRJ-1319-7D-CSC
    revision is
    serial number is H11A6ER
    fc-transmitter type is long wave laser cost reduced
    cisco extended id is unknown (0x0)
...
```

The following example displays the entire running configuration, with information about all the interfaces. The interfaces have multiple entries in the configuration files to ensure that the interface configuration commands execute in the correct order when the switch reloads.

Displays the Running Configuration for All Interfaces

```
switch# show running-config
...
interface fc9/1
  switchport speed 2000
...
interface fc9/1
  switchport mode E
...
interface fc9/1
```

```
  channel-group 11 force
  no shutdown
```

The following example displays the running configuration information for a specified interface. The interface configuration commands are grouped together:

Displays the Running Configuration for a Specified Interface

```
switch# show running-config interface fc1/1
interface fc9/1
  switchport speed 2000
  switchport mode E
  channel-group 11 force
  no shutdown
```

Displays the Running Configuration after the System Default Switchport Mode F Command is Executed, on page 67 displays the running configuration after the **system default switchport mode F** command is executed.

The following example displays the running configuration after the **system default switchport mode F** command is executed:

Displays the Running Configuration after the System Default Switchport Mode F Command is Executed

```
switch# show running-config
version 3.1(3)
system default switchport mode F
interface fc4/1
interface fc4/2
interface fc4/3
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
interface fc4/8
interface fc4/9
interface fc4/10
```

The following example displays the running configuration after two interfaces are individually configured for FL mode:

Displays the Running Configuration after Two Interfaces are Individually Configured for Mode FL

```
switch# show running-config
version 3.1(3)
system default switchport mode F
interface fc4/1
  switchport mode FL
interface fc4/2
interface fc4/3
  switchport mode FL
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
```

```
interface fc4/8
interface fc4/9
interface fc4/1
```

The following example displays interface information in a brief format after the **system default switchport mode F** command is executed:

Displays Interface Information in a Brief Format after the System Default Switchport Mode F Command is Executed

```
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port      Logical
                  Mode   Trunk                         Mode  Speed  Channel   Type
                         Mode                          (Gbps)
--------------------------------------------------------------------------------
fc4/1      1      F      --      notConnected     swl    --           --         --
fc4/2      1      F      --      notConnected     swl    --           --         --
fc4/3      1      F      --      notConnected     swl    --           --         --
fc4/4      1      F      --      notConnected     swl    --           --         --
fc4/5      1      F      --      sfpAbsent        --     --           --         --
fc4/6      1      F      --      sfpAbsent        --     --           --         --
fc4/7      1      F      --      sfpAbsent        --     --           --         --
fc4/8      1      F      --      sfpAbsent        --     --           --         --
fc4/9      1      F      --      sfpAbsent        --     --           --         --
```

The following example displays interface information in a brief format after two interfaces are individually configured for FL mode:

Displays Interface Information in a Brief Format after Two Interfaces Are Individually Configured for Mode FL

```
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port        Logical
                  Mode   Trunk                         Mode  Speed  Channel     Type
                         Mode                          (Gbps)
--------------------------------------------------------------------------------
fc4/1      1      FL     --      notConnected     swl    --           --          --
fc4/2      1      F      --      notConnected     swl    --           --          --
fc4/3      1      FL     --      notConnected     swl    --           --          --
fc4/4      1      F      --      notConnected     swl    --           --          --
fc4/5      1      F      --      sfpAbsent        --     --           --          --
fc4/6      1      F      --      sfpAbsent        --     --           --          --
fc4/7      1      F      --      sfpAbsent        --     --           --          --
fc4/8      1      F      --      sfpAbsent        --     --           --          --
fc4/9      1      F      --      sfpAbsent        --     --           --          --
fc4/10     1      F      --      sfpAbsent        --     --           --          --
```

Displaying the Port-Level Portguard

The following command displays information about an interface that is set to error-disabled state by the portguard because of a TrustSec violation:

```
switch# show interface fc8/3

fc8/3 is down (Error disabled - port down due to trustsec violation) Hardware is Fibre
Channel, SFP is short wave laser w/o OFC (SN) Port WWN is 21:c3:00:0d:ec:10:57:80
Admin port mode is E, trunk mode is on snmp link state traps are enabled
Port vsan is 1
Receive data field Size is 2112 Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 11274 frames input, 1050732 bytes
  0 discards, 0 errors
  0 CRC, 0 unknown class
  0 too long, 0 too short
 11242 frames output, 971900 bytes
  0 discards, 0 errors
 11 input OLS, 34 LRR, 10 NOS, 0 loop inits
 72 output OLS, 37 LRR, 2 NOS, 0 loop inits
Interface last changed at Sun Nov 27 07:34:05 1988
```

An interface may be error disabled for several reasons. To recover an error-disabled interface, use the **shutdown** and **no shutdown** commands in interface configuration mode to re-enable the link.

Displaying Port Monitor Status and Policies

The following commands display information about the Port Monitor feature:

Note The port *Logical type* is displayed as the *Port type*.

```
switch# show port-monitor
--------------------------------------------------------------------------------------------
Port Monitor : enabled
--------------------------------------------------------------------------------------------
Congestion-Isolation : enabled
--------------------------------------------------------------------------------------------
Policy Name  : default
Admin status : Not Active
Oper status  : Not Active
Port type    : All Ports
------------------------------------------------------
------------------------------------------------------------------------
Counter                Threshold  Interval Rising Threshold event Falling Threshold  event
 Warning Threshold     PMON Portguard
-------                ---------  -------- ---------------- ----- ------------------ -----
 --------------------- --------------
Link Loss                Delta      60       5                4     1                  4
   Not enabled           Not enabled
Sync Loss                Delta      60       5                4     1                  4
   Not enabled           Not enabled
Signal Loss              Delta      60       5                4     1                  4
   Not enabled           Not enabled
Invalid Words            Delta      60       1                4     0                  4
   Not enabled           Not enabled
Invalid CRC's            Delta      60       5                4     1                  4
   Not enabled           Not enabled
State Change             Delta      60       5                4     0                  4
   Not enabled           Not enabled
TX Discards              Delta      60       200              4     10                 4
```

```
   Not enabled           Not enabled
LR RX                    Delta      60       5               4     1                 4
   Not enabled           Not enabled
LR TX                    Delta      60       5               4     1                 4
   Not enabled           Not enabled
Timeout Discards         Delta      60       200             4     10                4
   Not enabled           Not enabled
Credit Loss Reco         Delta      60       1               4     0                 4
   Not enabled           Not enabled
TX Credit Not Available  Delta      60       10%             4     0%                4
   Not enabled           Not enabled
RX Datarate              Delta      60       80%             4     20%               4
   Not enabled           Not enabled
TX Datarate              Delta      60       80%             4     20%               4
   Not enabled           Not enabled
TX-Slowport-Oper-Delay   Absolute   60       50ms            4     0ms               4
   Not enabled           Not enabled
TXWait                   Delta      60       40%             4     0%                4
   Not enabled           Not enabled
------------------------------------------------------------
--------------------------------------------------------------------------------
```

```
switch# show port-monitor active
Policy Name  : sample
Admin status : Active
Oper status  : Active
Port type    : All Ports
------------------------------------------------------------
--------------------------------------------------------------------------------
Counter                 Threshold  Interval Rising Threshold event Falling Threshold  event
 Warning Threshold     PMON Portguard
-------                 ---------  -------- ---------------- ----- ------------------ -----
 --------------------- --------------
Link Loss                Delta      60       5               4     1                 4
   Not enabled           Not enabled
Sync Loss                Delta      60       5               4     1                 4
   Not enabled           Not enabled
Signal Loss              Delta      60       5               4     1                 4
   Not enabled           Not enabled
Invalid Words            Delta      60       5               4     1                 4
   Not enabled           Not enabled
Invalid CRC's            Delta      60       5               4     1                 4
   Not enabled           Not enabled
State Change             Delta      60       5               4     0                 4
   Not enabled           Not enabled
TX Discards              Delta      60       50              4     0                 4
   Not enabled           Not enabled
LR RX                    Delta      60       5               4     1                 4
   Not enabled           Not enabled
LR TX                    Delta      60       5               4     1                 4
   Not enabled           Not enabled
Timeout Discards         Delta      60       200             4     10                4
   Not enabled           Not enabled
Credit Loss Reco         Delta      1        1               4     0                 4
   Not enabled           Cong-isolate
TX Credit Not Available  Delta      1        10%             4     0%                4
   Not enabled           Cong-isolate
RX Datarate              Delta      60       80%             4     70%               4
   Not enabled           Not enabled
TX Datarate              Delta      60       80%             4     70%               4
   Not enabled           Not enabled
ASIC Error Pkt from Port Delta      60       50              4     10                4
   Not enabled           Not enabled
```

```
ASIC Error Pkt to xbar   Delta      60       50              4     10                4
   Not enabled           Not enabled
ASIC Error Pkt from xbar Delta      60       50              4     10                4
   Not enabled           Not enabled
TX-Slowport-Oper-Delay   Absolute   1        50ms            4     0ms               4
   Not enabled           Cong-isolate
TXWait                   Delta      1        40%             4     0%                4
   Not enabled           Cong-isolate
--------------------------------------------------------
--------------------------------------------------------------------
```

```
switch# show port-monitor sample
Policy Name  : sample
Admin status : Active
Oper status  : Active
Port type    : All Edge Ports
-----------------------------------------------------------------------------------------------------
Counter          Threshold  Interval Rising Threshold  event Falling Threshold  event
portgurard
-----------------------------------------------------------------------------------------------------
Link Loss         Delta        60            5           4            1           4
Not enabled
Sync Loss         Delta        60            5           4            1           4
Not enabled
Signal Loss       Delta        60            5           4            1           4
Not enabled
Invalid Words     Delta        60            1           4            0           4
Not enabled
Invalid CRC's     Delta        60            5           4            1           4
Not enabled
TX Discards       Delta        60            200         4            10          4
Not enabled
LR RX             Delta        60            5           4            1           4
Not enabled
LR TX             Delta        60            5           4            1           4
Not enabled
Timeout Discards  Delta        60            200         4            10          4
Not enabled
Credit Loss Reco  Delta        1             1           4            0           4
Not enabled
TX Credit Not
Available         Delta        1             10%         4            0%          4
Not enabled
RX Datarate       Delta        60            80%         4            20%         4
Not enabled
TX Datarate       Delta        60            80%         4            20%         4
Not enabled
TX-Slowport-Count Delta        1             5           4            0           4
Not enabled
TX-Slowport-Oper
-Delay            Absolute     1             50ms        4            0ms         4
Not enabled
TXWait            Delta        1             40%         4            0%          4
Not enabled
-----------------------------------------------------------------------------------------------------
```

```
switch# show port-monitor default
Policy Name  : default
Admin status : Not Active
Oper status  : Not Active
```

```
Port type    : All Ports
------------------------------------------------------------------------------------------------
Counter            Threshold  Interval Rising    event  Falling     event  Warning
PMON
                                       Threshold        Threshold          Threshold
Portguard
------------------------------------------------------------------------------------------------
Link Loss          Delta      60       5          4     1            4     Not enabled
Not enabled
Sync Loss          Delta      60       5          4     1            4     Not enabled
Not enabled
Signal Loss        Delta      60       5          4     1            4     Not enabled
Not enabled
Invalid Words      Delta      60       1          4     0            4     Not enabled
Not enabled
Invalid CRC's      Delta      60       5          4     1            4     Not enabled
Not enabled
State Change       Delta      60       5          4     0            4     Not enabled
Not enabled
TX Discards        Delta      60       200        4     10           4     Not enabled
Not enabled
LR RX              Delta      60       5          4     1            4     Not enabled
Not enabled
LR TX              Delta      60       5          4     1            4     Not enabled
Not enabled
Timeout Discards   Delta      60       200        4     10           4     Not enabled
Not enabled
Credit Loss Reco   Delta      60       1          4     0            4     Not enabled
Not enabled
TX Credit Not      Delta      60       10%        4     0%           4     Not enabled
Not enabled
Available
RX Datarate        Delta      60       80%        4     20%          4     Not enabled
Not enabled
TX Datarate        Delta      60       80%        4     20%          4     Not enabled
Not enabled
TX-Slowport-       Absolute   60       50ms       4     0ms          4     Not enabled
Not enabled
Oper-Delay
TXWait             Delta      60       40%        4     0%           4     Not enabled
Not enabled
------------------------------------------------------------------------------------------------
```

```
switch# show port-monitor slowdrain
Policy Name  : slowdrain
Admin status : Not Active
Oper status  : Not Active
Port type    : All Edge Ports
------------------------------------------------------------------------------------------------
Counter            Threshold    Interval    Rising Threshold    event    Falling Threshold
    event    PMON Portguard
------------------------------------------------------------------------------------------------
Credit Loss Reco     Delta           1            1               4           0
       4         Not enabled
TX Credit Not
Available            Delta           1            10%             4           0%
       4         Not enabled
------------------------------------------------------------------------------------------------
```

```
switch# show port-monitor slowportdetect
Policy Name  : slowportdetect
```

```
Admin status : Not Active
Oper status  : Not Active
Port type    : All Ports
------------------------------------------------------------------------------------------------------
Counter                 Threshold  Interval Rising     event  Falling    event  Warning
   PMON
                                            Threshold          Threshold          Threshold
   Portguard
------------------------------------------------------------------------------------------------------
Credit Loss Reco          Delta        1        2        2        0         2    Not enabled
  Cong-isolate
TX Credit Not Available   Delta        1        2%       2        0%        2    Not enabled
  Cong-isolate
TX-Slowport-Oper-Delay    Absolute     1        2ms      2        0ms       2    Not enabled
  Cong-isolate
TXWait                    Delta        1        2%       2        0%        2    Not enabled
  Cong-isolate
------------------------------------------------------------------------------------------------------
```

Displaying Port Group Monitor Status and Policies

The following examples display information about the port group monitor:

```
switch# show port-group-monitor status
Port Group Monitor : Enabled
Active Policies : pgm2
Last 100 logs :
switch#
switch# show port-group-monitor
------------------------------------------------------------------------------------------------
Port Group Monitor : enabled
------------------------------------------------------------------------------------------------
Policy Name  : pgm1
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
------------------------------------------------------------------------------------------------
Counter      Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------      ---------  -------- -------------------- ----------------------
RX Datarate   Delta      60       50                   10
TX Datarate   Delta      60       50                   10
------------------------------------------------------------------------------------------------
Policy Name  : pgm2
Admin status : Active
Oper status  : Active
Port type    : All Port Groups
------------------------------------------------------------------------------------------------
Counter       Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------      ---------  -------- -------------------- ----------------------
RX Datarate   Delta      60       80                   10
TX Datarate   Delta      60       80                   10
------------------------------------------------------------------------------------------------
Policy Name  : default
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
------------------------------------------------------------------------------------------------
Counter       Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------      ---------  -------- -------------------- ----------------------
RX Datarate   Delta      60       80                   20
```

```
TX Datarate   Delta      60       80                  20
-------------------------------------------------------------------------------------------
switch# show port-group-monitor active
Policy Name  : pgm2
Admin status : Active
Oper status  : Active
Port type    : All Port Groups
-------------------------------------------------------------------------------------------
Counter       Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------      ---------  -------- -------------------- ----------------------
RX Datarate   Delta      60       80                  10
TX Datarate   Delta      60       80                  10
-------------------------------------------------------------------------------------------
switch# show port-group-monitor PGMON_policy
PPolicy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
-------------------------------------------------------------------------------------------
Counter      Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------      ---------  -------- -------------------- ----------------------
RX Datarate   Delta      26       450                 250
TX Datarate   Delta      60       100                 80
-------------------------------------------------------------------------------------------
```

Displaying the Management Interface Configuration

The following command displays the management interface configuration:

```
switch# show interface mgmt 0
mgmt0 is up
    Hardware is FastEthernet
    Address is 000c.30d9.fdbc
    Internet address is 10.16.1.2/24
    MTU 1500 bytes, BW 100 Mbps full Duplex
    26388 packets input, 6101647 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    10247 packets output, 2389196 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Displaying VSAN Interface Information

To following example displays the VSAN interface information:

```
switch# show interface vsan 2
vsan2 is up, line protocol is up
    WWPN is 10:00:00:05:30:00:59:1f, FCID is 0xb90100
    Internet address is 10.1.1.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
```

Transmit-Wait History Graph

The transmit-wait history for the slow ports on 16-Gbps and 32-Gbps modules and switches can be displayed in the form of a graph over a period of time. The total transmit-wait time for each time period is displayed as a column of #. The actual value appears above each column as a vertically printed number. The following graphs can be displayed:

- Seconds scale—The transmit-wait history for the port over the last 60 seconds. The Y-axis value is the total transmit-wait time for each second, in milliseconds.
- Minutes scale—The transmit-wait history for the port over the last 60 seconds. The Y-axis value is the total transmit-wait time for each minute, in seconds, to one decimal place.
- Hours scale—The transmit-wait history for the port over the last 60 seconds. The Y-axis value is the total transmit-wait time for each hour, in minutes.

To display the transmit-wait history for a given interval of time, use the following commands:

Display the transmit-wait history graph for the period when transmit credit is not available for a given interval of time (seconds, minutes, or hours):

switch# **show process creditmon txwait-history [module x [port y]]**

Display the transmit-wait time in 2.5 microsecond units, as well as in seconds:

switch# **show logging onboard txwait**

Note The transmit-wait delta values are logged periodically (every 20 seconds) into the OBFL when transmit wait increases by at least 100 ms in the 20-second interval.

Display the total transmit-wait value for a particular interface in 2.5-microsecond units:

switch# **show interface fc*x/y* counters**

The following example displays the transmit-wait history graph, in seconds, for 16-Gbps modules:

```
switch(config)# show process creditmon txwait-history module 1 port 81

TxWait history for port fc1/81:
===============================

   455555555455555455555455559999999999999999999999999999999999
     900000000800009000008100011111231922322211321121112112113111
     433799991899990359909838608935137962088988254848894870461938
1000                              #
900                          ####################################
800                          ####################################
700                          ####################################
600                          ####################################
500 ############################################################
400 ############################################################
300 ############################################################
200 ############################################################
100 ############################################################
    0....5....1....1....2....2....3....3....4....4....5....5....6
              0    5    0    5    0    5    0    5    0    5    0
```

```
  Tx Credit Not Available per second (last 60 seconds)
                # = TxWait (ms)
```

The following example displays the transmit-wait history graph, in minutes, for 16-Gbps modules:

```
   5555555555555555555555555555555555555555555555555555555555555555
   0555555555555555555555555555555555555555555555555555555555555555
   ............................................................
   0354444454444454454544444444456353635343754344533435545453 44
60
54  ############################################################
48 ############################################################
42 ############################################################
36 ############################################################
30 ############################################################
24 ############################################################
18 ############################################################
12 ############################################################
6 ############################################################
  0....5....1....1....2....2....3....3....4....4....5....5....6
            0    5    0    5    0    5    0    5    0    5    0

          Tx Credit Not Available per minute (last 60 minutes)
                # = TxWait (secs)
```

The following example displays the transmit wait history graph, in hours, for 16-Gbps modules:

```
  3333333333333333333333333333333333333333
  3333333333333333333333333333333333333332
     2222222222222222222222222222222222222229
     7777887777778778777787778767777777787679000000000000000000000000000000000
3600
3240 ########################################
2880 ########################################
2520 ########################################
2160 ########################################
1800 ########################################
1440 ########################################
1080 ########################################
 720 ########################################
 360 ########################################
    0....5....1....1....2....2....3....3....4....4....5....5....6....6....7.7
              0    5    0    5    0    5    0    5    0    5    0    5    0 2

          Tx Credit Not Available per hour (last 72 hours)
                # = TxWait (secs)
```

The following example displays the transmit-wait OBFL logging for 16-Gbps modules:

```
switch# show logging onboard txwait
Notes:
    - sampling period is 20 seconds
    - only txwait delta value >= 100 ms are logged
---------------------------------
 Module: 1 txwait count
---------------------------------

-------------------------------------------------------------------------------
| Interface | Delta TxWait Time     | Congestion | Timestamp                |
|           | 2.5us ticks | seconds |            |                          |
-------------------------------------------------------------------------------
```

```
|   fc1/11  |  3435973    | 08      |      42%   | Sun Sep 30 05:23:05 2001 |
|   fc1/11  |  6871947    | 17      |      85%   | Sun Sep 30 05:22:25 2001 |
```

Configuring Fibre Channel Interfaces

This chapter provides information about Fibre Channel interfaces, its features, and how to configure the Fibre Channel interfaces.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About Fibre Channel Interfaces

Forward Error Correction

The Transmitter Training Signal (TTS) is defined in FC-FS-4(Clause 5.5). It provides the capability for FC ports to negotiate the following two capabilities:

1. Transmitter training, which enables a receiver to send feedback to a transmitter to assist the transmitter in adapting to the characteristics of the link that connects them.
2. FEC usage.

The TTS is not used by 4 and 8-gigabit FC ports. From 32-gigabit speed and higher, its use is mandatory. For 16-gigabit FC ports, EA variants must transmit the TTS during the link speed negotiation, but the use of it by the receiver is optional, and EL variants must not use TTS.

Forward Error Correction (FEC) is defined in IEEE 802.3TM clause 74 and is implemented in FC without modification. FEC is not supported on 4, 8 and 16-gigabit EL ports. Its use is optional on 16-gigabit EA ports. The TTS is the mechanism that allows FEC negotiation on such ports.

For more information on configuring FEC and TTS, see the Configuring FEC, on page 91 section.

Dynamic Bandwidth Management

On port switching modules where bandwidth is shared, the bandwidth available to each port within a port group can be configured based on the port rate mode and speed configurations. Within a port group, some ports can be configured in dedicated rate mode while others operate in shared mode.

Ports configured in dedicated rate mode are allocated the required bandwidth to sustain a line rate of traffic at the maximum configured operating speed, and ports configured in shared mode share the available remaining bandwidth within the port group. Bandwidth allocation among the shared mode ports is based on the operational speed of the ports.

Unutilized bandwidth from the dedicated ports is shared among only the shared ports in a port group as per the ratio of the configured operating speed. A port cannot be brought up unless the reserved bandwidth is quarantined for the shared ports. For dedicated ports, configured bandwidth is taken into consideration while calculating available bandwidth for the port group. This behavior can be changed using bandwidth fairness by using the **rate-mode bandwidth-fairness module** *number* command.

Out-of-Service Interfaces

On supported modules and fabric switches, you might need to allocate all the shared resources for one or more interfaces to another interface in the port group or module. You can take interfaces out of service to release shared resources that are needed for dedicated bandwidth. When an interface is taken out of service, all shared resources are released and made available to the other interface in the port group or module. These shared resources include bandwidth for the shared mode port, rate mode, BB_credits, and extended BB_credits. All shared resource configurations are returned to their default values when the interface is brought back into service. Corresponding resources must be made available in order for the port to be successfully returned to service.

Caution If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces in the same port group.

Bandwidth Fairness

This feature improves fairness of bandwidth allocation among all ports and provides better throughput average to individual data streams. Bandwidth fairness can be configured per module.

Caution When you disable or enable bandwidth fairness, the change does not take effect until you reload the module.

Use the show module bandwidth-fairness command to check whether ports in a module are operating with bandwidth fairness enabled or disabled.

```
switch# show module 2 bandwidth-fairness
Module 2 bandwidth-fairness is enabled
```

Upgrade or Downgrade Scenario

When you are upgrading from a release earlier than Cisco SAN-OS Release 3.1(2), all modules operate with bandwidth fairness disabled until the next module reload. After the upgrade, any new module that is inserted has bandwidth fairness enabled.

When you are downgrading to a release earlier than Cisco SAN-OS Release 3.1(2), all modules keep operating in the same bandwidth fairness configuration prior to the downgrade. After the downgrade, any new module that is inserted has bandwidth fairness disabled.

Note After the downgrade, any insertion of a module or module reload will have bandwidth fairness disabled.

Guidelines and Limitations

Port Index Limitations

Note Generation 1 and Generation 2 modules are not supported in Cisco MDS NX-OS Release 8.1(x).

Cisco MDS 9000 switches allocate index identifiers for the ports on the modules. These port indexes cannot be configured. You can combine Generation 1, Generation 2, Generation 3, and Generation 4 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2 or Generation 3.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1, Generation 2, Generation 3, or Generation 4 switching modules are installed in the chassis.

Note On a switch with the maximum limit of 252 as port index, any new module that exceeds the limit does not power up when installed.

You can use the **show port index-allocation** command to display the allocation of port indexes on the switch.

```
switch# show port index-allocation
Module index distribution:
-----------------------------------------------------+
Slot | Allowed |      Allotted indices info    |
     |  range  | Total |          Index values |
-----------|---------|-------|----------------------|
1    |  -----  |   -   | (None)                |
2    |  -----  |   -   | (None)                |
3    |  -----  |   -   | (None)                |
4    |  -----  |   -   | (None)                |
5    |   0-1023|   32  | 0-31                  |
6    |  -----  |   -   | (None)                |
9    |  -----  |   -   | (None)                |
10   |  -----  |   -   | (None)                |
11   |  -----  |   -   | (None)                |
12   |  -----  |   -   | (None)                |
13   |   0-1023|   48  | 32-79                 |
SUP  | 253-255 |   3   | 253-255               |
```

When a module does not power up because of a resource limitation, you can display the reason by using the **show module** command.

```
switch# show module
Mod  Ports  Module-Type                          Model              Status
---  -----  ---------------------------------- ------------------ ----------
5    32     1/2/4/8/10 Gbps Advanced FC Module  DS-X9232-256K9     ok
7    0      Supervisor/Fabric-2                 DS-X9530-SF2-K9    active *
13   48     1/2/4/8/10 Gbps Advanced FC Module  DS-X9248-256K9     ok
Mod  Sw              Hw      World-Wide-Name(s) (WWN)
---  --------------  ------  --------------------------------------------------
5    5.2(2)          0.207   21:01:00:0d:ec:b7:28:c0 to 21:20:00:0d:ec:b7:28:c0
7    5.2(2)          1.9     --
13   5.2(2)          0.212   23:01:00:0d:ec:b7:28:c0 to 23:30:00:0d:ec:b7:28:c0
Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
5    68-ef-bd-a8-45-cc to 68-ef-bd-a8-45-d0  JAF1450CHQT
7    00-24-c4-60-00-f8 to 00-24-c4-60-00-fc  JAE141502L2
13   68-ef-bd-a8-40-00 to 68-ef-bd-a8-40-04  JAF1450BMBP
Xbar Ports  Module-Type                          Model              Status
---  -----  ---------------------------------- ------------------ ----------
1    0      Fabric Module 3                     DS-13SLT-FAB3      ok
2    0      Fabric Module 3                     DS-13SLT-FAB3      ok
Xbar Sw              Hw      World-Wide-Name(s) (WWN)
---  --------------  ------  --------------------------------------------------
1    NA              0.4     --
2    NA              0.4     --
Xbar MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    NA                                      JAF1451AMHG
2    NA                                      JAF1451AMHN
* this terminal session
```

The running configuration is updated when modules are installed. If you save the running configuration to the startup configuration (using the copy running-config startup-config command), during reboot the switch powers up the same set of modules as before the reboot regardless of the sequence in which the modules initialize. You can use the **show port index-allocation startup** command to display the index allocation the switch uses at startup.

```
switch# show port index-allocation startup
Startup module index distribution:
-----------------------------------------------------+
Slot | Allowed |      Alloted indices info            |
     |  range  | Total |          Index values        |
-----|---------|-------|------------------------------|
 1   |  -----  |   34  | 0-31,80-81                   |
 2   |  -----  |   32  | 32-63                        |
 3   |  -----  |   16  | 64-79                        |(Slot 1 shares 80-81)
 4   |  -----  |   48  | 96-127,224-239               |
SUP  | 253-255 |   3   | 253-255                      |
```

Note The output of the **show port index-allocation startup** command does not display anything in the Allowed range column because the command extracts the indices from the persistent storage service (PSS) and displaying an allowed range for startup indices is meaningless.

If a module fails to power up, you can use the **show module** *slot* **recovery-steps** command to display the reason.

PortChannel Limitations

Note Generation 1 and Generation 2 modules are not supported in Cisco MDS NX-OS Release 8.1(x).

PortChannels have the following restrictions:

- The maximum number of PortChannels allowed is 256 if all switching modules are Generation 2 or Generation 3, or both.
- The maximum number of PortChannels allowed is 128 whenever there is a Generation 1 switching module in use with a Generation 2 or Generation 3 switching module.
- Ports need to be configured in dedicated rate mode on the Generation 2 and Generation 3 switching module interfaces to be used in the PortChannel.
- Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9) supports 32 Gbps, 16 Gbps, 8 Gbps, and 4 Gbps speed. However, a single 32-Gbps SFP supports only 32 Gbps, 16 Gbps, and 8 Gbps speed and a single 16-Gbps SFP supports only 16 Gbps, 8 Gbps, and 4 Gbps speed. You must not configure speed values other than the values recommended for these SFPs.

The Generation1, Generation 2, and Generation 3 modules have the following restrictions for PortChannel configuration:

- Generation 1 switching module interfaces do not support auto speed with a maximum of 2 Gbps.
- Generation 1 and Generation 2 module interfaces do not support auto speed with maximum of 4 Gbps.
- Generation 2 and Generation 3 switching module interfaces cannot be forcefully added to a PortChannel if sufficient resources are not available.

Note Before adding a Generation 2 or Generation 3 interface to a PortChannel, use the **show port-resources module** command to check for resource availability.

When configuring PortChannels on switches with Generation 1, Generation 2, and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 2 and Generation 3 interfaces to auto with a maximum of 2 Gbps.
- Configure the Generation 1 switching modules followed by the Generation 2 switching modules, and then the Generation 3 switching modules, and then configure the PortChannel.

When configuring PortChannels on switches with only Generation 2 and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 3 interfaces to auto with a maximum of 4 Gbps.
- Configure the Generation 2 switching modules, followed by the Generation 3 switching modules, and then configure the PortChannel.

Table 13: PortChannel Configuration and Addition Results , on page 86 describes the results of adding a member to a PortChannel for various configurations.

Table 13: PortChannel Configuration and Addition Results

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
No members	Any	Any	Generation 1 or Generation 2 or Generation 3 or Generation 4	Force	Pass
	Auto	Auto	Generation 1 or Generation 2 or Generation 3 or Generation 4	Normal or force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass or fail[4]
	Auto	Auto max 4000	Generation 3 or Generation 4		
	Auto max 2000	Auto	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto max 2000	Auto max 4000	Generation 3 or or Generation 4		
	Auto max 4000	Auto	Generation 2 or Generation 3 or or Generation 4		
	Auto max 4000	Auto max 2000	Generation 2 or Generation 3 or or Generation 4		
Generation 1 interfaces	Auto	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass or fail[5]
	Auto max 4000	Auto	Generation 1 or Generation 2		
	Auto max 4000	Auto	Generation 3		

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
Generation 2 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
Generation 3 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3	Normal	Fail
				Force	Pass
Generation 4 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3 or Generation 4	Normal	Fail
				Force	Pass

[4] If resources are not available.

[5] If resources are not available.

Use the **show port-channel compatibility parameters** command to obtain information about PortChannel addition errors.

Configuring Fibre Channel Interfaces

Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode

To configure the Fibre Channel switching modules when starting with the default configuration or when migrating from shared rate mode to dedicated rate mode, perform these steps:

Step 1 Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the Taking Interfaces out of Service, on page 98.

Step 2 Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, 8 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

See the Dynamic Bandwidth Management, on page 81.

Step 3 Configure the rate mode (dedicated or shared).

See the Configuring FEC, on page 91.

Step 4 Configure the port mode.

See the Configuring an Interface Mode , on page 41.

Note ISL ports cannot operate in shared rate mode.

Step 5 Configure the BB_credits and extended BB_credits, as necessary.

See the Configuring Buffer-to-Buffer Credits, on page 115 and Extended Buffer-to-Buffer Credits, on page 112 .

Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode

To configure the Fibre Channel switching modules migrating from dedicated rate mode to shared rate mode, perform these steps:

Step 1 Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the Taking Interfaces out of Service, on page 98.

Step 2 Configure the BB_credits and extended BB_credits, as necessary.

See the Configuring Buffer-to-Buffer Credits, on page 115 and Extended Buffer-to-Buffer Credits, on page 112

Step 3 Configure the port mode.

See the Configuring an Interface Mode , on page 41.

Note ISL ports cannot operate in shared rate mode.

Step 4 Configure the rate mode (dedicated or shared) to use.

See the Configuring FEC, on page 91.

Step 5 Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

See the Dynamic Bandwidth Management, on page 81.

Configuring Port Speed

Note Changing port speed and rate mode disrupts traffic on the port. Traffic on other ports in the port group is not affected.

To configure the port speed on an interface, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc 1/1**

Selects the interface and enters interface configuration submode.

Step 3 switch(config-if)# **switchport speed** {**1000** | **2000** | **4000** | **8000** | **10000** | **16000** | **32000**}

Configures the port speed in megabits per second. The auto parameter enables autosensing on the interface.

Step 4 switch(config-if)# **switchport speed auto**

Configures autosensing for an interface.

Note The auto speed configurations are available only for the specific modules.

Step 5 switch(config-if)# **no switchport speed**

Reverts to the default speed for the interface (auto).

Use the **show interface** command to verify the port speed configuration for an interface.

```
switch# show interface fc 9/1
fc9/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 22:01:00:05:30:01:9f:02
    Admin port mode is F
    snmp traps are enabled
    Port mode is F, FCID is 0xeb0002
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is shared
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      226 frames input, 18276 bytes
        0 discards, 0 errors
```

```
  0 CRC, 0 unknown class
  0 too long, 0 too short
326 frames output, 21364 bytes
  0 discards, 0 errors
0 input OLS, 0 LRR, 1 NOS, 0 loop inits
3 output OLS, 2 LRR, 0 NOS, 0 loop inits
16 receive B2B credit remaining
64 transmit B2B credit remaining
```

Configuring FEC

FEC has the following restrictions:

- FEC is supported on the DS-X9334-K9, DS-X9648-1536K9, and DS-X9448-768K9 modules in the Cisco MDS 9700 Series switch. FEC is also supported on the Cisco MDS 9132T Fibre Channel Switch and Cisco MDS 9396S Multilayer Fabric Switch.
- FEC fallback is not supported on the Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9) when its interfaces are configured at 16-Gbps Fibre Channel fixed speed. However, FEC fallback is supported on the Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module (DS-X9448-768K9) when its interfaces are configured at 16-Gbps Fibre Channel fixed speed.
- Modifying FEC configuration briefly disrupts traffic on the port.
- FEC can be configured only on fixed speed 16-Gbps Fibre Channel ports. FEC is not supported for ports configured with 2000/4000/8000/auto/auto-max maximum speed.
- 32-Gbps Fibre Channel ports come up automatically in FEC. You do not have to configure the ports using the **switchport fec** and **switchport fec tts** commands, as they are meant only for 16-Gbps Fibre Channel ports and have no effect on 32-Gbps Fibre Channel ports.
- From Cisco MDS NX-OS Release 6.2(11c), FEC with Transmitter Training Signal (TTS) is supported on the Cisco MDS 9396S 16-Gbps Multilayer Fabric Switch and Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module (DS-X9448-768K9), except in Cisco MDS NX-OS Release 6.2(13).

  From Cisco MDS NX-OS Release 8.2(1), FEC with Transmitter Training Signal (TTS) is supported by default and cannot be disabled on the Cisco MDS 9132T Fibre Channel Switch and Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9).
- From Cisco MDS NX-OS Release 8.2(1), FEC with TTS feature is supported in Simple Network Management Protocol (SNMP) and Device Manager (DM). This feature is not supported in Cisco MDS NX-OS Release 8.1(1) or earlier.

  Note FEC is supported in DM.
- DM uses SNMP to get switch updates.

To configure FEC on an interface on the 48-port 16-Gbps Fibre Channel switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc 1/1**

Selects the interface and enters interface configuration submode.

Step 3 switch(config-if)# **switchport speed 16000**

Sets the port speed to 16 Gbps.

Step 4 switch(config-if)# **switchport fec**

Note The **switchport fec** command works only in fixed 16-Gbps speed and a message stating the same appears when you execute this command.

Enables FEC for the 16-Gbps interface.

- FEC is active if it is configured on both local and peer switches.
- FEC is not active if it is configured only on the local switch, but not on the peer switch.

Step 5 switch(config-if)# **switchport fec tts**

(Optional) Enables TTS, that allows negotiation of FEC. This command is only accepted on ports with fixed 16-Gbps speed and FEC enabled.

Note The **switchport fec tts** command can be used only after configuring FEC using the **switchport fec** command.

Use the **show interface** command to verify the port speed configuration for an interface on the 48-port, 16-Gbps Fibre Channel switching module.

```
switch# show interface fc3/15
fc3/15 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:8f:54:7f:ee:ea:3a:00
    Admin port mode is auto, trunk mode is off
    snmp link state traps are enabled
    Port mode is F, FCID is 0xdf0000
    Port vsan is 100
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 128
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    admin fec state is down
    oper fec state is down
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      16072969258 frames input,34396153854332 bytes
        23 discards,45 errors
        22 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      8040504998 frames output,17206679580576 bytes
        1344 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      306 output OLS,304 LRR, 1 NOS, 4 loop inits
      32 receive B2B credit remaining
      128 transmit B2B credit remaining
      128 low priority transmit B2B credit remaining
```

```
Interface last changed at Wed Mar 12 21:23:36 2014
Last clearing of "show interface" counters  :never
```

Configuring Rate Mode

Note Changing port speed and rate mode disrupts traffic on the port.

To configure the rate mode (dedicated or shared) on an interface on a Fibre Channel switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc 1/1**

Selects the interface and enters interface configuration submode.

Step 3 switch(config-if)# **switchport rate-mode dedicated**

Reserves dedicated bandwidth for the interface.

Note If you cannot reserve dedicated bandwidth on an interface, you might have exceeded the port group maximum bandwidth. Use the **show port-resources** command to determine what resources are already allocated.

Step 4 switch(config-if)# **switchport rate-mode shared**

Reserves shared (default) bandwidth for the interface.

Step 5 switch(config-if)# **no switchport rate-mode**

Reverts to the default state (shared).

Disabling Restrictions on Oversubscription Ratios

Note Before disabling restrictions on oversubscription ratios, ensure that you have explicitly shut down shared ports.

To disable restrictions on oversubscription ratios on a Fibre Channel switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no rate-mode oversubscription-limit module 1**

Disables restrictions on oversubscription ratios for a module.

Note You must enter this command separately for each module for which you want to remove the restrictions.

Step 3 switch(config)# **exit**

Exits configuration mode.

Step 4 switch# **copy running-config startup-config**

Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Use the **show running-config** command to view oversubscription ratios for a module. If oversubscription ratios are enabled, then no restriction appears in the output.

```
switch# show running-config
version 3.1(1)
...
no rate-mode oversubscription-limit module 2
interface fc2/1
  switchport speed 2000
interface fc2/1
...
```

Examples

To disable restrictions on oversubscription ratios for ports on a 48-port Generation 2 switch that is configured with both shared and dedicated ports, perform these steps:

Step 1 To disable restrictions on oversubscription ratios, you must shut down any shared ports. Use the show port-resources command to view the configuration on a module and to identify shared ports.

Example:

```
switch# show port-resources module 2
Module 2
Available dedicated buffers are 4656
 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc2/1                                    16        4.0  shared
  fc2/2                                    16        4.0  shared
  fc2/3                                    16        4.0  dedicated
  fc2/4                                    16        4.0  shared
  fc2/5                                    16        4.0  shared
  fc2/6                                    16        4.0  dedicated
  fc2/7                                    16        4.0  dedicated
  fc2/8                                    16        4.0  shared
  fc2/9                                    16        4.0  shared
```

```
  fc2/10                                        16        4.0  shared
  fc2/11                                        16        4.0  shared
  fc2/12                                        16        4.0  shared
  ...
 Port-Group 4
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc2/37                                        16        4.0  shared
  fc2/38                                        16        4.0  shared
  fc2/39                                        16        4.0  dedicated
  fc2/40                                        16        4.0  dedicated
  fc2/41                                        16        4.0  dedicated
  fc2/42                                        16        4.0  shared
  fc2/43                                        16        4.0  shared
  fc2/44                                        16        4.0  shared
  fc2/45                                        16        4.0  shared
  fc2/46                                        16        4.0  shared
  fc2/47                                        16        4.0  shared
  fc2/48                                        16        4.0  shared
```

Step 2 Shut down all shared ports for which you want to remove restrictions on oversubscription ratios.

Example:

```
switch (config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch (config-if)# shutdown
```

Step 3 Display the interface status to confirm the shutdown of all shared ports.

Example:

```
switch(config-if)# end
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                                (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      down             swl    --           --
fc2/2      1      FX     --      down             swl    --           --
fc2/3      1      T      --      up               swl    --           --
fc2/4      1      FX     --      down             swl    --           --
fc2/5      1      FX     --      down             swl    --           --
fc2/6      1      TE     --      up               swl    --           --
fc2/7      1      TE     --      up               swl    --           --
fc2/8      1      FX     --      down             swl    --           --
...
fc2/48     1      FX     --      down             swl    --           --
```

Step 4 Disable restrictions on oversubscription ratios for the ports.

Example:

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
```

```
switch(config)# no rate-mode oversubscription-limit module 2
```

Step 5 Bring up the ports that you shut down in step 2, and display their status to confirm that they are no longer shut down.

Example:

```
switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch(config-if)# no shutdown
switch(config-if)# end
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                                (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      up              swl    --           --
fc2/2      1      FX     --      up              swl    --           --
fc2/3      1      T      --      up              swl    --           --
fc2/4      1      FX     --      up              swl    --           --
fc2/5      1      FX     --      up              swl    --           --
fc2/6      1      TE     --      up              swl    --           --
fc2/7      1      TE     --      up              swl    --           --
fc2/8      1      FX     --      up              swl    --           --
...
fc2/48     1      FX     --      up              swl    --           --
```

Step 6 Confirm that the ports are now operating with no restrictions on oversubscription ratios.

Example:

```
switch# show running-config | include oversubscription-limit
no rate-mode oversubscription-limit module 2 <---indicates no restrictions on oversubscription ratios
```

Step 7 Save the new oversubscription ratio configuration to the startup configuration.

Example:

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

Enabling Restrictions on Oversubscription Ratios

Note

- You must enable restrictions on oversubscription ratios before you can downgrade modules to a previous release.
- Before enabling restrictions on oversubscription ratios, ensure that you have explicitly configured shared ports to out-of-service mode.

To enable restrictions on oversubscription ratios on a Fibre Channel switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48**

Specifies the port interfaces for which you want to enable restrictions on oversubscription ratios.

Step 3 switch(config-if)# **shutdown**

Shuts down shared ports.

Step 4 switch(config-if)# **out-of-service**

Takes shared ports out of service.

Step 5 switch# **rate-mode oversubscription-limit module 1**

Enables restrictions on oversubscription ratios for the module.

Note You must enter this command separately for each module for which you want to add the restriction.

Step 6 switch# **configure terminal**

Enters configuration mode.

Step 7 switch(config)# **interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48**

Specifies the port interfaces for which you want to enable restrictions on oversubscription ratios.

Step 8 switch(config)# **no out-of-service**

Returns all shared ports to service.

Step 9 switch(config-if)# **no shutdown**

Bring up shared ports.

Step 10 switch(config)# **exit**

Exits configuration mode.

Step 11 switch# **copy running-config startup-config**

Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Enabling Bandwidth Fairness

To enable bandwidth fairness on a switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **rate-mode bandwidth-fairness module 1**

Enables bandwidth fairness for a module.

Note You must enter this command separately for each module for which you want to enable bandwidth fairness. You must reload the module for the command to take effect.

Step 3 switch(config)# **exit**

Exits configuration mode.

Disabling Bandwidth Fairness

Note If you disable bandwidth fairness, up to a 20 percent increase in internal bandwidth allocation is possible for each port group; however, bandwidth fairness is not guaranteed when there is a mix of shared and full-rate ports in the same port group.

To disable bandwidth fairness on a switching module, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no rate-mode bandwidth-fairness module 1**

Disables bandwidth fairness for a module.

Note You must enter this command separately for each module for which you want to disable bandwidth fairness. You must reload the module for the command to take effect.

Step 3 switch(config)# **exit**

Exits configuration mode.

Taking Interfaces out of Service

You can take interfaces out of service on Generation 2 and Generation 3 switching modules. When an interface is out of service, all the shared resources for the interface are released as well as the configuration associated with those resources.

Note

- The interface must be disabled using a **shutdown** command before it can be taken out of service.
- The interface cannot be a member of a PortChannel.
- Taking interfaces out of service releases all the shared resources to ensure that they are available to other interfaces. This causes the configuration in the shared resources to revert to default when the interface is brought back into service. Also, an interface cannot come back into service unless the default shared resources for the port are available. The operation to free up shared resources from another port is disruptive.

To take an interface out of service, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc 1/1**

Selects the interface and enters interface configuration submode.

Step 3 switch(config-if)# **no channel-group**

Removes the interface from a PortChannel.

Step 4 switch(config-if)# **shutdown**

Disables the interface.

Step 5 switch(config-if)# **out-of-service**

Takes the interface out of service.

Use the **show port-resources module** command to verify the out-of-service configuration for interfaces on a Generation 2 and Generation 3 switching module.

This example shows a 24-port 4-Gbps module:

```
switch# show port-resources module 9
Module 9
Available dedicated buffers are 5429
 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/1                                     16        4.0  shared
  fc9/2 (out-of-service)
  fc9/3                                     16        4.0  shared
  fc9/4                                     16        4.0  shared
  fc9/5                                     16        4.0  shared
```

```
  fc9/6                                   16        4.0  shared
...
```

Releasing Shared Resources in a Port Group

When you want to reconfigure the interfaces in a port group on a Generation 2 or Generation 3 module, you can return the port group to the default configuration to avoid problems with allocating shared resources.

Note

- The interface cannot be a member of a PortChannel.
- Releasing shared resources disrupts traffic on the port. Traffic on other ports in the port group is not affected.

To release the shared resources for a port group, perform these steps:

Step 1 switch# **configure t**

Enters configuration mode.

Step 2 switch(config)# **interface fc 1/1**

Selects the interface and enters interface configuration submode.

Tip You can use an interface range to release the resources for all interfaces in a port group.

Step 3 switch(config-if)# **no channel-group**

Removes the interface from a PortChannel.

Step 4 switch(config-if)# **shutdown**

Disables the interface.

Step 5 switch(config-if)# **out-of-service**

Takes the interface out of service.

Step 6 switch(config-if)# **no out-of-service**

Makes the interface available for service. Repeat Step 2 through Step 6 for all the interfaces in the port group.

Disabling ACL Adjacency Sharing for System Image Downgrade

Fibre Channel ACL adjacency sharing is enabled by default on the switches with an active Generation 2 switching module as of Cisco MDS SAN-OS Release 3.0(3), and with an active Generation 3 module as of MDS NX-OS Release 4.1(1). Fibre Channel ACL adjacency sharing improves the performance for zoning and inter-VSAN routing (IVR) network address translation (NAT). To prevent disruptions when downgrading

the system image on your switch to a release prior to Cisco SAN-OS Release 3.0(3), enter the following command in EXEC mode:

```
switch# system no acl-adjacency-sharing
```

To reenable Fibre Channel ACL adjacency sharing on your switch, enter the following command in EXEC mode:

```
switch# system acl-adjacency-sharing
```

Verifying Fibre Channel Interfaces Configuration

To display Fibre Channel interface configuration information, perform one of the following tasks:

Command	Purpose
show module	Displays the module.
show module *slot* **recovery-steps**	Displays the slot for the module.
show port-resources module *slot*	Displays the port resources for the slot.
show interface fc *slot/port*	Displays the slot or port information. FEC admin and operational states are displayed.
show interface brief	Displays the interface.
show port index-allocation	Displays the port in the index allocation.
show port index-allocation startup	Displays the startup port in the index allocation.
show port-channel compatibility parameters	Displays the PortChannel compatibility parameters.
show module *slot* **bandwidth-fairness**	Displays the module slot bandwidth fairness information.

For detailed information about the fields in the output from these commands, refer to the Cisco MDS 9000 Series Command Reference.

Displaying FEC Module Interfaces

This example shows a 32-Gbps FC interface status:

Note 36-gigabit FC ports comes up automatically in FEC and need not be configured.

```
switch# show interface fc 10/21 brief
 --------------------------------------------------------------------------------
---------
Interface  Vsan   Admin  Admin   Status       SFP    Oper  Oper   Port     Logical
                  Mode   Trunk                       Mode  Speed  Channel   Type
                         Mode                               (Gbps)
--------------------------------------------------------------------------------
fc10/21     1      auto   on      trunking     swl    TE     32     --       core


switch# show interface fc10/21
fc10/21 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 22:55:54:7f:ee:ea:1f:00
    Peer port WWN is 22:24:54:7f:ee:ea:1d:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
```

```
    Port mode is TE
    Port vsan is 1
    Admin Speed is auto max 32 Gbps
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY

    Transmit B2B Credit is 500
    Receive B2B Credit is 500
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    fec is enabled by default
    Logical type is core
    Trunk vsans (admin allowed and active)  (1)
    Trunk vsans (up)                        (1)
    Trunk vsans (isolated)                  ()
```

Displaying SFP Diagnostic Information

You can use the **show interface transceiver** command to display small form-factor pluggable (SFP) diagnostic information.

```
switch# show interface transceiver
...
fc2/9 sfp is present
    Name is CISCO-Switch
    Manufacturer's part number is XXXX-XXXX-XXX
    Revision is V01
    Serial number is XXXXXX
    Cisco part number is XXXXX
    Cisco pid is XXXXXX
    FC Transmitter type is short wave laser w/o OFC (SN)
    FC Transmitter supports short distance link length
    Transmission medium is multimode laser with 50 um aperture (M5)
    Supported speeds are - Min speed: 8000 Mb/s, Max speed: 32000 Mb/s
    Nominal bit rate is 28000 Mb/s
    Link length supported for 50/125um OM3 fiber is 70 m
    Cisco extended id is unknown (0x0)

    No tx fault, rx loss, no sync exists, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
       Temperature           :  30.61 C
        Voltage              :   3.35 V
        Current              :   4.10 mA
        Optical Tx Power     :  -2.44 dBm
        Optical Rx Power     :  N/A   dBm      --
        Tx Fault count  : 0
    Note: ++  high-alarm; +  high-warning; --  low-alarm; -low-warning
...
```

Configuration Examples for Fibre Channel Interfaces

Configuration Example for FEC Module Interfaces

These steps describe how to configure FEC module interfaces:

Step 1 Select the interfaces fc 4/1 through fc 4/2.

Example:

```
switch# configure terminal
switch(config)# interface fc 4/1 - 2
```

Step 2 Configure the FEC on the interfaces.

Example:

```
switch(config-if)# switchport speed 16000
switch(config-if)# switchport fec
```

Step 3 Enable the interfaces and return to configuration mode.

Example:

```
switch(config-if)# no shutdown
switch(config-if)# exit
```

Step 4 Select the interfaces fc 4/3 through fc 4/4.

Example:

```
switch# configure terminal
switch(config)# interface fc 4/3 - 4
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

Example:

```
switch(config-if)# switchport speed 16000
switch(config-if)# switchport fec
```

Configuring Interface Buffers

This chapter provides information about interfaces buffers, its features, and how to configure the interface buffers.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About Interface Buffers

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination.

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that Fibre Channel switches do not run out of buffers, so that switches do not drop frames. Buffer-to-buffer credits are negotiated on a per-hop basis.

The receive buffer-to-buffer credit (fcrxbbcredit) value may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.

Performance Buffers

Regardless of the configured receive buffer-to-buffer credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Series, you can specify the amount of performance buffers allocated in addition to the configured receive buffer-to-buffer credit value.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

Buffer Pools

In the architecture of 4-Gbps, 8-Gbps, and 16-Gbps modules, receive buffers shared by a set of ports are called *buffer groups*. The receive buffer groups are organized into *global* and *local* buffer pools.

The receive buffers allocated from the global buffer pool to be shared by a port group are called a global receive buffer pool. Global receive buffer pools include the following buffer groups:

- Reserved internal buffers
- Allocated buffer-to-buffer credit buffers for each Fibre Channel interface (user configured or assigned by default)
- Common unallocated buffer pool for buffer-to-buffer credits, if any, to be used for additional buffer-to-buffer credits as needed

Buffer-to-Buffer Credit Buffers for Switching Modules

This section describes how buffer credits are allocated to Cisco MDS 9000 Series Multilayer switches.

48-Port 32-Gbps Fibre Channel Module Buffer-to-Buffer Credit Buffers

Note The 48-Port 32-Gbps Fibre Channel module buffer-to-buffer credit buffer allocation is also applicable to Cisco MDS 9132T, MDS 9148T, and MDS 9396T switches.

Table lists the buffer-to-buffer credit buffer allocation for the 48-port 32-Gbps Fibre Channel switching module (DS-X9648-1536K9).

Table 14: 48-Port 32-Gbps Switching Module Buffer-to-Buffer Credit Buffer Allocation

Buffer-to-Buffer Credit Buffer Allocation	**Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 4-Gbps to 32-Gbps Speed	
	ISL	**Fx Port**
Default buffer-to-buffer credit buffers	500	32
Maximum buffer-to-buffer credit buffers	500	500
Extended Buffer-to-Buffer Credit Buffer Allocation	**Extended Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 4-Gbps to 32-Gbps Speed	
	ISL	**Fx Port**
Maximum extended buffer-to-buffer credit buffers	300	300

Note The DS-X9648-1536K9 module is a full rate card.

The following guidelines apply to buffer-to-buffer credit buffers on the 48-port 32-Gbps Fibre Channel switching modules:

- Buffer-to-buffer credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers when all other ports in a port group are moved to out of service.
- Buffer-to-buffer credit buffers for Fx port mode connections can be configured from a minimum of 1 buffers to a maximum of 500 buffers when all other ports in a port group are moved to out of service.
- If the user has installed an enterprise license, per port credits in a port group can be increased up to 800 using extended buffer to buffer credits.
- If the user has installed an enterprise license, per port credits in a port group can be increased up to 8191 using extended buffer to buffer credits when ports are moved to out of service.

- The extended buffer feature can be activated by using the following commands:

```
switch(config)# interface fc1/5
switch(config-if)# switchport fcrxbbcredit extended 4095
```

Note In Cisco MDS 9700 Series Switches module, each port group comprises of 16 ports, and there are 3 port groups per ASIC. Port group buffers can be allocated to any combination of ports in that port group using extended buffer configuration. Refer to the **show port-resource module** *module_number* command for details about buffers supported by port-groups.

48-Port 16-Gbps Fibre Channel Module Buffer-to-Buffer Credit Buffers

Table 15: 48-Port 16-Gbps Switching Module Buffer-to-Buffer Credit Buffer Allocation, on page 109 lists the buffer-to-buffer credit buffer allocation for the 48-port 16-Gbps Fibre Channel switching module (DS-X9448-768K9).

Table 15: 48-Port 16-Gbps Switching Module Buffer-to-Buffer Credit Buffer Allocation

Buffer-to-Buffer Credit Buffer Allocation	**Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 8-Gbps to 16-Gbps Speed	
	ISL	**Fx Port**
Default buffer-to-buffer credit buffers	500	32
Maximum buffer-to-buffer credit buffers	500	500
Extended Buffer-to-Buffer Credit Buffer Allocation	**Extended Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 8-Gbps to 16-Gbps Speed	
	ISL	**Fx Port**
Maximum extended buffer-to-buffer credit buffers	2150	2150

Note The DS-X9448-768K9 module is a 16 Gbps line-rate module.

The following guidelines apply to buffer-to-buffer credit buffers on the 48-port 16-Gbps Fibre Channel switching modules:

- Buffer-to-buffer credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers.
- Buffer-to-buffer credit buffers for Fx port mode connections can be configured from a minimum of 1 buffers to a maximum of 500 buffers.
- If the user has installed an enterprise license, per port credits in a port group can be increased up to 2650 using extended buffer to buffer credits.
- If the user has installed an enterprise license, per port credits in a port group can be increased up to 4144 using extended buffer to buffer credits when ports are moved to out of service. However, the cli restricts the per port credits to be increased to only 4095.

Note In Cisco MDS 9700 Series Switches module, each port group comprises of 4 ports, and there are 12 port groups per ASIC. Port group buffers can be allocated to any combination of ports in that port group using extended buffer configuration. Refer to the **show port-resource module** *module_number* command for details about buffers supported by port-groups.

Buffer-to-Buffer Credit Buffers for Fabric Switches

This section describes how buffer credits are allocated to Cisco MDS 9000 Fabric switches.

Cisco MDS 9396S Fabric Switch Buffer-to-Buffer Credit Buffers

Table 16: 96-Port 16-Gbps Switch Buffer-to-Buffer Credit Buffer Allocation, on page 110 lists the buffer-to-buffer credit buffer allocation for the 96-port 16-Gbps Fibre Channel switch.

Table 16: 96-Port 16-Gbps Switch Buffer-to-Buffer Credit Buffer Allocation

Buffer-to-Buffer Credit Buffer Allocation	**Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 16-Gbps Speed	
	ISL	**Fx Port**
Default buffer-to-buffer credit buffers	500	32
Maximum buffer-to-buffer credit buffers	500	500

Note Cisco MDS 9396S is a 16 Gbps line-rate switch.

The following guidelines apply to buffer-to-buffer credit buffers on the 96-port 16-Gbps Fibre Channel switch:

- Buffer-to-buffer credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers.
- Buffer-to-buffer credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers.
- Per port credits can be increased up to 4095 using extended buffer to buffer credits if the user has installed an enterprise license.

Note In MDS 9396S Fabric switch, total buffer available are 99600 for 24 port groups. One port group comprises of 4 ports, and there are 2 port groups per ASIC. Each port-group consists of total 4150 buffers. These buffers can be allocated to any combination of port(s) using extended buffer configuration. Please refer **show port-resource module** *module_number* command for details about buffers supported by port-groups.

Cisco MDS 9250i and Cisco MDS 9148S Fabric Switch Buffer-to-Buffer Credit Buffers

Table 17: 40/48-Port 16-Gbps Switch Buffer-to-Buffer Credit Buffer Allocation, on page 111 lists the buffer-to-buffer credit buffer allocation for 40/48-port 16-Gbps Cisco MDS 9250i and 9148S Fabric switches.

Table 17: 40/48-Port 16-Gbps Switch Buffer-to-Buffer Credit Buffer Allocation

Buffer-to-Buffer Credit Buffer Allocation	**Buffer-to-Buffer Credit Buffers Per Port**	
	Dedicated Rate Mode 16-Gbps Speed	
	ISL	**Fx Port**
Default buffer-to-buffer credit buffers	64	64
Maximum buffer-to-buffer credit buffers	253	253

Note Cisco MDS 9148S and Cisco MDS 9250i are 16 Gbps line-rate switches.

The following guidelines apply to buffer-to-buffer credit buffers on the 40/48-port 9250i/9148S Fabric switches:

- Buffer-to-buffer credit buffers can be configured from a minimum of 1 buffer to a maximum of 64 buffers per port when the ports are in F or FL mode.
- Buffer-to-buffer credit buffers can be configured from a minimum of 2 buffers to a maximum of 64 buffers per port when the ports are in E or TE mode.
- Buffer-to-buffer credit buffers for F or FL port can be configured for a single port in a port group from a minimum of 1 buffer to a maximum of 253 buffers when all other ports in a port group are moved to out of service.

- Buffer-to-buffer credit buffers for E or TE port can be configured for a single port in a port group from a minimum of 2 buffer to a maximum of 253 buffers when all other ports in a port group are moved to out of service.

Note The ports that are moved to out-of-service need not be licensed.

Extended Buffer-to-Buffer Credits

To facilitate buffer-to-buffer credits for long-haul links, the extended buffer-to-buffer credits feature allows you to configure the receive buffers above the maximum value on all 4-Gbps, 8-Gbps, advanced 8-Gbps, 16-Gbps, and 32-Gbps switching modules. When necessary, you can reduce the buffers on one port and assign them to another port, exceeding the default maximum. The minimum extended buffer-to-buffer credits per port is 256 and the maximum is 4095.

In general, you can configure any port in a port group to dedicated rate mode. To do this, you must first release the buffers from the other ports before configuring larger extended buffer-to-buffer credits for a port.

Note The ENTERPRISE_PKG license is required to use extended buffer-to-buffer credits on 4-Gbps, 8-Gbps, advanced 8-Gbps, 16-Gbps, and 32-Gbps switching modules. Also, extended buffer-to-buffer credits are not supported by ports in shared rate mode.

All ports on the 4-Gbps, 8-Gbps, 16-Gbps, and 32-Gbps switching modules support extended buffer-to-buffer credits. There are no limitations for how many extended buffer-to-buffer credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can take interfaces out of service to make more extended buffer-to-buffer credits available to other ports.

Buffer-to-Buffer Credit Recovery

Although Fibre Channel standards require low bit and frame error rates, there is a likelihood of errors occuring. When these errors affect certain Fibre Channel primitives, credit loss might occur. When credits are lost, performance degradation might occur. When all credits are lost, transmission of frames in that direction stops. The Fibre Channel standards introduces a feature for two attached ports to detect and correct such scenarios nondisruptively. This feature is called *buffer-to-buffer credit recovery*.

A credit can be lost in either of these scenarios:

- An error corrupts the start-of-frame (SoF) delimiter of a frame. The receiving port fails to recognize the frame and subsequently does not send a corresponding receiver ready (R_RDY) primitive to the sender. The sending port does not replenish the credit to the receiving port.
- An error corrupts an R_RDY primitive. The receiving port fails to recognize the R_RDY and does not replenish the corresponding credit to the sending port.

The Buffer-to-Buffer Credit Recovery feature can help recover from the two specified scenarios. It is a per-hop feature and is negotiated between two directly attached peer ports when the link comes up, by exchanging parameters. Buffer-to-buffer credit recovery is enabled when a receiver acknowledges a nonzero buffer-to-buffer state change number (BB_SC_N).

Buffer-to-buffer credit recovery functions as follows:

1. The local port and peer port agree to send checkpoint primitives to each other for frames and R_RDYs, starting from the time the link comes up.
2. If a port detects frame loss, it sends the corresponding number of R_RDYs to replenish the lost credits at the peer port.
3. If a port detects R_RDY loss, the port internally replenishes the lost credits to the interface buffer pool.

Buffer-to-buffer credit recovery implementation is as follows:

1. Buffer-to-buffer state change SOF (BB_SCs) primitives are transmitted every $2^{BB\_SC\_N}$ number of frames sent. This enables an attached port to determine if any frames are lost. If frames loss is detected, the receiver of the BB_SCs transmits the appropriate number of R_RDYs to compensate for the lost frames.
2. Buffer-to-buffer state change R_RDY (BB_SCr) primitives are transmitted every $2^{BB\_SC\_N}$ number of R_RDY primitives sent. This enables an attached port to determine if any R_RDY primitives are lost. If R_RDY primitive loss is detected, the receiver of the BB_SCr increments the number of transmit credits by the appropriate number to compensate for the lost R_RDYs.

The Buffer-to-Buffer Credit Recovery feature can be used on any nonarbitrated loop link. This feature is most useful on unreliable links, such as Metropolitan Area Networks (MANs) or WANs, but can also help on shorter, high-loss links, such as a link with a faulty fiber connection.

Note The Buffer-to-Buffer Credit Recovery feature is not compatible with the distance extension (DE) feature, also known as buffer-to-buffer credit spoofing. If you use intermediate optical equipment, such as dense wavelength-division multiplexing (DWDM) or Fibre Channel bridges that use DE on Inter-Switch Links (ISLs) between switches, then buffer-to-buffer credit recovery on both sides of an ISL must be disabled.

The following are the guidelines and restrictions for the Buffer-to-Buffer Credit Recovery feature:

- E ports
  - This feature is enabled by default on ISLs (E ports).
  - This feature works on an ISL between a Cisco switch and a peer switch from any vendor, provided this feature is supported on the peer switch.
  - This feature is supported only on links that are in R_RDY flow control mode. It is not supported on links that are in ER_RDY flow control mode.
- F ports
  - This feature is enabled by default on F ports.
  - This feature works on an F port between a Cisco switch and a peer device from any vendor, provided this feature is supported on the peer device.

Note Some host bus adapters (HBAs) do not support the Buffer-to-Buffer Credit Recovery feature. Others support this feature at only certain speeds. Check with your HBA vendor about the exact configurations supported.

- N-Port ID Virtualization (NPIV) ports do not support buffer-to-buffer credit recovery for Cisco N-Port Virtualizer (Cisco NPV) switch logins.

Receive Data Field Size

You can configure the receive data field size for Fibre Channel interfaces. The default data field size is 2112 bytes, which supports frame lengths up to 2148 byte, the maximum size of Fibre Channel frames.

Configuring Interface Buffers

Configuring Buffer-to-Buffer Credits

Note When you configure port mode to auto or E, and rate mode to dedicated for all the ports in the global buffer pool, you must reconfigure buffer credits on one or more ports (other than the default mode).

To configure a single pool of buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be in R_RDY flow-control mode.

Before you begin

Enable the Receiver Ready (R_RDY) mode on ISLs before configuring the shared buffer-to-buffer credit pool. For more information, see Disabling Extended Receiver Ready, on page 161.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc** *slot/port*

Step 3 Set the buffer-to-buffer credits as a single pool on an interface:

switch(config-if)# **switchport fcrxbbcredit** *credits* **mode** {**E** | **Fx**}

(Optional) Reset the buffer-to-buffer credits on the interface to the default value:

switch(config-if)# **switchport fcrxbbcredit default**

Configuring Buffer-to-Buffer Credits for Virtual Links

Note When you configure port mode to auto or E, and rate mode to dedicated for all the ports in the global buffer pool, you must reconfigure buffer credits on one or more ports (other than the default mode).

To configure per-virtual-link buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be an ISL in ER_RDY flow-control mode.

Before you begin

Enable the Extended Receiver Ready (ER_RDY) mode on ISLs before configuring the virtual-link credits. For more information, see Enabling Extended Receiver Ready, on page 160.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc** *slot/port*

Step 3 Set the buffer-to-buffer credits per virtual-link on an ISL:

switch(config-if)# **switchport vl-credit vl0** *credits* **vl1** *credits* **vl2** *credits* **vl3** *credits*

Step 4 (Optional) Reset the buffer-to-buffer credits on the ISL to the default value:

switch(config-if)# **switchport vl-credit default**

Configuring Performance Buffers

To configure performance buffers for a Fibre Channel interface, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enters interface configuration submode:

switch(config)# **interface fc***slot*/*port*

Step 3 Set the number of performance buffers on an interface:

switch(config-if)# **switchport fcrxbbcredit performance-buffers** *perf_bufs*

(Optional) Reset the number of performance buffers on an interface to the default value:

switch(config-if)# **switchport fcrxbbcredit performance-buffers default**

Configuring Extended Buffer-to-Buffer Credits

Note You cannot configure regular buffer-to-buffer credits after configuring the extended buffer-to-buffer credits.

To configure a single pool of extended buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be in R_RDY flow-control mode.

Before you begin

Enable the Receiver Ready (R_RDY) mode on ISLs before configuring the shared buffer-to-buffer credit pool. For more information, see Disabling Extended Receiver Ready, on page 161.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable the extended Buffer-to-Buffer Credits feature:

switch(config)# **fcrxbbcredit extended enable**

Step 3 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc** *slot/port*

Step 4 Set the extended buffer-to-buffer credits as a single pool on an interface:

switch(config-if)# **switchport fcrxbbcredit extended** *extend_bufs* **mode** {**E** | **Fx**}

Step 5 (Optional) Reset the extended buffer-to-buffer credits on the interface to the default value:

switch(config-if)# **switchport fcrxbbcredit extended default**

Configuring Extended Buffer-to-Buffer Credits for Virtual Links

Note You cannot configure regular buffer-to-buffer credits after configuring the extended buffer-to-buffer credits.

To configure per-virtual-link extended buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be an ISL in ER_RDY flow-control mode.

Before you begin

Enable the Extended Receiver Ready (ER_RDY) mode on ISLs before configuring the virtual link credits. For more information, see Enabling Extended Receiver Ready, on page 160.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable the extended Buffer-to-Buffer Credits feature:

switch(config)# **fcrxbbcredit extended enable**

Step 3 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc** *slot/port*

Step 4 Set the extended buffer-to-buffer credits per virtual link on an ISL:

switch(config-if)# **switchport vl-credit extended vl0** *credits* **vl1** *credits* **vl2** *credits* **vl3** *credits*

Step 5 (Optional) Reset the extended buffer-to-buffer credits on the ISL to the default value:

switch(config-if)# **switchport vl-credit extended default**

Configuring Buffer-to-Buffer Credit Recovery

Buffer-to-buffer credit recovery is enabled by default on all Fibre Channel ports.

To disable the buffer-to-buffer credit recovery on a port, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select the interface and enter interface configuration submode:

switch(config)# **interface fc** *slot*/*port*

Step 3 Disable buffer-to-buffer credit recovery on the interface:

switch(config-if)# **no switchport fcbbscn**

Step 4 (Optional) To enable buffer-to-buffer credit recovery on an interface if it was disabled:

switch(config-if)# **switchport fcbbscn**

Configuring Receive Data Field Size

Note From Cisco MDS NX-OS 8.2(1), the **switchport fcrxbufsize** command is obsolete on the Cisco MDS 9700 48-port 16-Gbps Fibre Channel Switching Module and the Cisco MDS 9700 48-port 32-Gbps Fibre Channel Switching Module. The receive data field size is permanently set to 2112 bytes. Any receive data field size configuration from earlier Cisco MDS NX-OS versions is ignored.

To configure the receive data field size, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

switch(config)# **interface fc** *slot*/*port*

Step 3 Set the data field size for the selected interface:

switch(config-if)# **switchport fcrxbufsize** *bytes*

Step 4 (Optional) Reset the receive data field size on the interface to the default value:

switch(config-if)# **no switchport fcrxbufsize**

Configuration Examples for Interface Buffers

This example shows how to enable buffer-to-buffer credit recovery on an interface if it is disabled:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcbbscn
```

This example shows how to configure default credits on an interface:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit default
```

This example shows how to configure 50 receive buffer credits on an interface:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit 50
```

This example shows how to configure 4095 extended buffer credits to an interface:

```
switch# configure terminal
switch(config)# fcrxbbcredit extended enable
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit extended 4095
```

This examples shows how to assign 45 performance buffers to a selected interface:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit performance-buffers 45
```

This example shows how to set the received frame data field size for an interface to 2000 bytes:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbufsize 2000
```

This example shows how to assign buffer-to-buffer credits per virtual link on an ISL:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport vl-credit  vl0 12 vl1 10 vl2 29 vl3 349
```

This example shows how to assign extended buffer-to-buffer credits per virtual link on an ISL:

```
switch# configure terminal
switch(config)# fcrxbbcredit extended enable
switch(config)# interface fc 1/1
```

```
switch(config-if)# switchport vl-credit extended vl0 20 vl1 25 vl2 40 vl3 349
```

Verifying Interface Buffer Configuration

This example shows which of the interfaces on a specified module are in R_RDY flow-control mode:

```
switch# show flow-control r_rdy module 3
fc3/17
fc3/18
```

This example shows how to verify the buffer-to-buffer credit information for all interfaces:

```
sswitch# show interface bbcredit
fc2/1 is down (SFP not present)
.
.
.
fc2/17 is trunking
Transmit B2B Credit is 255
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 375
12 receive B2B credit remaining
255 transmit B2B credit remaining
fc2/21 is down (Link failure or not-connected)
.
.
.
fc2/31 is up
Transmit B2B Credit is 0
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 48
12 receive B2B credit remaining
0 transmit B2B credit remaining
```

This example shows how to verify buffer-to-buffer credit information for a specific Fibre Channel interface:

```
switch# show interface fc2/31 bbcredit
fc2/31 is up
Transmit B2B Credit is 0
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 48
12 receive B2B credit remaining
0 transmit B2B credit remaining
```

This example shows how to verify the type of buffers and data field size a port supports:

```
switch# show interface fc1/1 capabilities
fc1/1
Min Speed is 2 Gbps
Max Speed is 16 Gbps
FC-PH Version (high, low) (0,6)
Receive data field size (max/min) (2112/256) bytes
Transmit data field size (max/min) (2112/128) bytes
Classes of Service supported are Class 2, Class 3, Class F
Class 2 sequential delivery supported
Class 3 sequential delivery supported
Hold time (max/min) (100000/1) micro sec
BB state change notification supported
```

```
Maximum BB state change notifications 14
Rate Mode change not supported

Rate Mode Capabilities Dedicated
Receive BB Credit modification supported yes
FX mode Receive BB Credit (min/max/default) (1/500/32)
ISL mode Receive BB Credit (min/max/default) (2/500/500)
Performance buffer modification supported yes
FX mode Performance buffers (min/max/default) (1/0/0)
ISL mode Performance buffers (min/max/default) (1/0/0)

Out of Service capable yes
Beacon mode configurable yes
Extended B2B credit capable yes
On demand port activation license supported no
```

This example shows how to verify the operational receive data field size for a port:

```
switch# show interface fc 4/1
fc4/1 is down (SFP not present)
Hardware is Fibre Channel
Port WWN is 20:c1:8c:60:4f:c9:53:00
Admin port mode is auto, trunk mode is on
snmp link state traps are enabled
Port vsan is 1
Receive data field Size is 2112
Beacon is turned off
Logical type is Unknown(0)
5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
4 frames input,304 bytes
0 discards,0 errors
0 invalid CRC/FCS,0 unknown class
0 too long,0 too short
4 frames output,304 bytes
0 discards,0 errors
0 input OLS,0 LRR,0 NOS,0 loop inits
0 output OLS,0 LRR, 0 NOS, 0 loop inits
Last clearing of "show interface" counters : never
```

This example shows how to verify credit mode and credit allocation for an ISL:

```
switch# show interface fc9/1
.
.
.
Port flow-control is ER_RDY

Transmit B2B Credit for vl0 is 15
Transmit B2B Credit for vl1 is 15
Transmit B2B Credit for vl2 is 40
Transmit B2B Credit for vl3 is 430
Receive B2B Credit for vl0 is 15
Receive B2B Credit for vl1 is 15
Receive B2B Credit for vl2 is 40
Receive B2B Credit for vl3 is 430
.
.
.
```

Troubleshooting Interface Buffer Credits

Use the **show logging onboard interrupt-stats** command to view the number of times a port sent extra R_RDYs or incremented transmit buffer to buffer credits to restore credit counts:

```
switch# show logging onboard interrupt-stats
...
--------------------------------------------------------------------------------
INTERRUPT COUNTS INFORMATION FOR DEVICE: FCMAC
--------------------------------------------------------------------------------
 Interface|                                      |         |    Time Stamp
 Range    |            Interrupt Counter Name    |  Count  |MM/DD/YY HH:MM:SS
          |                                      |         |
--------------------------------------------------------------------------------
fc1/1     |IP_FCMAC_INTR_ERR_BB_SCR_INCREMENT    |1        |01/01/17 20:00:00
fc1/1     |IP_FCMAC_INTR_ERR_BB_SCS_RESEND       |1        |01/01/17 10:00:00
...
```

Congestion Detection, Avoidance, and Isolation

This chapter provides information about devices that cause congestion in a Fibre Channel or Fibre Channel over Ethernet (FCoE) network and includes information about how to identify and avoid or isolate such devices. These devices can be both slow devices as well as devices that are attempting to over utilize the bandwidth of their links or interfaces.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Feature History for Congestion Detection, Avoidance, and Isolation

Feature Name	Release	Feature Information
Fibre Channel over Ethernet (FCoE)	8.2(1)	New FCoE commands were introduced and some FCoE commands were modified to align with the commands used in Fibre Channel. The following commands were modified: • The congestion drop timeout command has changed from **system default interface congestion timeout** *milliseconds* **mode** {**core** | **edge**} to **system timeout fcoe congestion-drop** {*milliseconds* | **default**} **mode** {**core** | **edge**} • The pause drop timeout command has changed from **system default interface pause timeout** *milliseconds* **mode** {**core** | **edge**} to **system timeout fcoe pause-drop** {*milliseconds* | **default**} **mode** {**core** | **edge**} • The output for the **show interface vfc** *slot/port* **counter details** and **show interface priority-flow-control** commands were modified to add the receive and transmit pause frame information in the output. • The **show logging onboard** command was modified to add the **txwait**, **rxwait**, and **error-stats** keywords. The following commands were introduced: • **show hardware internal txwait-history** [**module** *number* | **port** *number*] • **show hardware internal rxwait-history** [**module** *number* | **port** *number*]

Feature Name	Release	Feature Information
Extended Receiver Ready	8.1(1)	This feature allows each Inter-Switch Link (ISL) between supporting switches to be split into four separate virtual links, with each virtual link assigned its own buffer-to-buffer credits. The following commands were introduced: • **show flow-control** {**er_rdy** | **r_rdy**} [**module** *number*] • **switchport vl-credit** {**default** | **vl0** *value* **vl1** *value* **vl2** *value* **vl3** *value*} • **system fc flow-control** {**default** | **er_rdy** | **r_rdy**}
Congestion Isolation	8.1(1)	This feature allows devices to be categorized as slow by either configuration command or by the port monitor. The following commands were introduced: • **congestion-isolation** {**include** | **exclude**} **pwwn** *pwwn* **vsan** *vsan-id* • **feature congestion-isolation** • **show congestion-isolation** {**exclude-list** | **global-list** | **ifindex-list** | **include-list** | **pmon-list** | **remote–list** | **status**} The *cong-isolate* portguard action was added to the following commands: • **counter credit-loss-reco** • **counter tx-credit-not-available** • **counter tx-slowport-oper-delay** • **counter tx-wait**

Feature Name	Release	Feature Information
Congestion Drop Timeout, No-Credit Frame Timeout, and Slow-Port Monitor Timeout Values for Fibre Channel	8.1(1)	The following features were modified: • The link connecting a core switch to a Cisco NPV switch should be treated as an ISL (core port) for the purpose of congestion-drop, no-credit-drop, slow-port monitor, and port monitor. To accomplish this, **logical-type** {**all** | **core** | **edge**} feature was introduced. • The Fibre Channel congestion drop timeout value's range was changed from 100-500 ms to 200-500 ms. The following commands were modified: **Note** From Cisco MDS NX-OS Release 8.1(1), E ports are treated as core and F ports are treated as edge in the **system timeout congestion-drop**, **system timeout no-credit-drop**, and **system timeout slowport-monitor** commands. • **system timeout congestion-drop** *milliseconds* **logical-type** {**core** | **edge**} • **system timeout no-credit-drop** *milliseconds* **logical-type edge** • **system timeout slowport-monitor** *milliseconds* **logical-type** {**core** | **edge**} • **switchport logical-type** {**auto** | **core** | **edge**}

Information About SAN Congestion

Information About SAN Congestion Caused by Slow-Drain Devices

Most SAN edge devices use Class 2 or Class 3 Fibre Channel services that have link-level flow control. The Flow Control feature allows a receiving port to back-pressure the upstream-sending port when the receiving port reaches its capacity to accept frames. When an edge device does not accept frames from the fabric for an extended period of time, it creates a condition in the fabric known as slow drain. If the upstream source of a slow-edge device is an ISL, it results in credit starvation, or slow drain in that ISL. This credit starvation then affects the unrelated flows that use the same ISL link. Although, the flow control mechanisms are different between Fibre Channel and FCoE, similar congestion can occur. Regardless of the protocol of the device causing the congestion, the congestion can propagate back to the source of the frames via both FC and FCoE links.

Fibre Channel buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that each side of the Fibre Channel link is able to control the rate of incoming frames. BB_credits are set on a per-hop basis. Each side of a Fibre Channel connection informs the other connection of the number of buffers that are available for it to receive frames. The sender can only send frames if the receiver has buffers. For each frame received, the receiver transmits a R_RDY (also known as BB_credit) to the sender of that frame. If there is some processing delay in the receiver, it can withhold the BB_credits, thereby limiting the rate it is receiving frames. If the receiver withholds the BB_credits to a significant amount, it cause congestion in the SAN. This BB_credit mechanism works independently in each direction.

In FCoE, the flow control mechanism is called priority flow control (PFC). PFC consists of a receiver sending class-based pause frames to a sender. PFC pause frames contain a value called a quanta. The quanta determines how long the class of traffic is paused. There are two types of PFC pause frames—non-zero quanta and zero quanta. A PFC pause frame with a non-zero quanta signals the receiver to stop sending frames immediately for a specified amount of time. A PFC pause frame with a zero quanta signals the receiver that it can resume sending frames immediately. As the receiver experiences some processing delay or its buffers reach a defined threshold, it can transmit a PFC pause frame with a non-zero quanta. After the buffers are sufficiently freed, the receiver can transmit another PFC pause frame containing a zero quanta which in turn signals the sender to resume traffic. This PFC pause mechanism works in each direction independently of the other.

Devices that do not accept frames at the rate generated by the sender can be both Fibre Channel and FCoE. The underlying flow control mechanism is different between the Fibre Channel and FCoE. But, Fibre Channel and FCoE can equally cause congestion in the SAN. These devices are referred to as slow drain devices.

Slow-drain devices can be detected and actions can be taken to drop all or old frames that exceed the configured threshold and queued to the slow drain devices, reset credits on the affected ports, flap the affected ports, error-disable the affected ports, or isolate the traffic flows to the slow-drain devices. This is achieved using the Congestion Detection, Congestion Avoidance, and Congestion Isolation features.

Information About SAN Congestion Caused by Over Utilization

Small Computer Systems Interface (SCSI) host devices request data via various SCSI *read* commands. These SCSI *read* commands contain a data length field which is the amount of data requested in the specific *read* request. Likewise, SCSI targets request data via the SCSI Xfr_rdy command and the amount of data requested is contained in the burst size. The rate of these *read* or Xfr_rdy requests coupled with the amount of data requested can result in more data flowing to the specific end device than its link can support at the given time. This is compounded by speed mismatches, hosts zoned to multiple targets, and targets zoned to multiple hosts.

The switch infrastructure (SAN) can buffer some of this excess but if the rate of requests is continuous then the switch's queues can fill and Fibre Channel or FCoE back pressure can result. This back pressure is done by withholding BB_credits on Fibre Channel and by sending PFC pauses with non-zero quantas on FCoE. The resulting effects to the SAN can look identical to slow drain, but the root cause is much different. The main mechanism for detecting this is monitoring the Tx datarate of the end device ports. This can be done via port monitor.

Information About Congestion Detection, Avoidance, and Isolation

Information About Congestion Detection

The following features are used to detect congestion on Cisco MDS switches:

- Display of credits agreed to on the port along with remaining credits (Fibre Channel only): The credits agreed to in both directions in the FLOGI (F ports) and Exchange Link Parameters (ELP for ISLs) are displayed via the **show interface** command. Additionally, the instantaneous value of the remaining credits is also displayed in output of the **show interface** command. The credits agreed to is static and unchaging information (at least when the link is up). However, the remaining credits values are constantly changing. If the remaining credits approach or reach zero, it indicates congestion on that port.

  The following example displays the transmit and receive credits information for an F port:

```
switch# show interface fc9/16
fc9/16 is up
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port mode is F, FCID is 0x0c0100
Transmit B2B Credit is 16
Receive B2B Credit is 32
…
32 receive B2B credit remaining
16 transmit B2B credit remaining
```

  The following example displays the transmit and receive credits information for an E port that is in R_RDY mode:

```
switch# show interface fc1/5
fc1/5 is trunking (Not all VSANs UP on the trunk)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Transmit B2B Credit is 64
Receive B2B Credit is 500
…
500 receive B2B credit remaining
64 transmit B2B credit remaining
```

  The following example displays the transmit and receive credits information for an F port that is in ER_RDY mode:

```
switch# show interface fc9/1 | i i fc | credit
fc9/1 is trunking
Transmit B2B Credit for vl0:15 vl1:15 vl2:40 vl3:430
Receive B2B Credit for vl0:15 vl1:15 vl2:40 vl3:430
0 invalid CRC/FCS,0 unknown class
Transmit B2B credit remaining for virtual link 0-3: 15,15,40,428
Receive B2B credit remaining for virtual link 0-3: 15,15,40,430
```

- Tx and Rx transitions to zero (Fibre Channel only): When a port approaches or reaches zero remaining credits in either direction, the *transitions to zero* counter is incremented. This indicates that the port is running out of credits, but does not indicate the duration the port is at zero credits. The port could have been at zero credits momentarily or for a longer time. TxWait gives a better view of the impact of running out of credits, because it gives the actual duration the port was at zero Tx credits remaining. Transitions to zero are shown in the **show interface counters** and **show interface counters detailed** commands.

  The following example displays the *transition to zero* status of the transmit and receive credits:

```
switch# show interface fc1/1 counters | i fc | transitions
fc1/1
0 Transmit B2B credit transitions to zero
0 Receive B2B credit transitions to zero
```

- Priority-flow-control pauses (FCoE only): This gives a count of PFC pause frames sent and received on an interface. Priority-flow-control pauses is just a count and includes both PFC pauses with a non-zero quanta (actual pause frames) and PFC pauses with a zero quanta (resume frames). This does not give any indication of the amount of time the port is actually paused. The port could have been at zero credits momentarily or for a longer time. TxWait and RxWait give a better view of the impact of these pause frames because they give the actual amount of time the port was paused in each direction. PFC pauses can be seen via the **show interface** and **show interface priority-flow-control** commands.

  The following example displays the *transition to zero* status of the transmit and receive credits:

```
switch# show interface eth3/1
Ethernet3/1 is up
admin state is up, Dedicated Interface
Belongs to Epo540
...snip
RX
555195 unicast packets 105457 multicast packets 0 broadcast packets
...snip
230870335 Rx pause
TX
326283313 unicast packets 105258 multicast packets 0 broadcast packets
...snip
0 Tx pause
```

  The following example displays the *transition to zero* status of the transmit and receive credits:

```
switch# show interface priority-flow-control
RxPause: No. of pause frames received
TxPause: No. of pause frames transmitted
TxWait: Time in 2.5uSec a link is not transmitting data[received pause]
RxWait: Time in 2.5uSec a link is not receiving data[transmitted pause]
=============================================================================
Interface        Admin Oper  (VL bmap) VL  RxPause    TxPause RxWait-2.5us(sec)
TxWait-2.5us(sec))
=============================================================================
ethernet-
port-channel540 Auto    NA     (8)     3    456200000  0         0(0)
152866694355(382166)
Ethernet2/1     Auto    On     (8)     3    4481929    0         0(0)
5930346153(14825)
...snip
Ethernet2/48    Auto    Off
Ethernet3/1     Auto    On     (8)     3    0          0         0(0)             0(0)
...snip
```

```
Ethernet3/6     Auto    Off
Ethernet3/7     Auto    On     (8)     3    0          0         0(0)             0(0)
```

- Slowport monitor (Fibre Channel only): A slowport monitor threshold duration is specified to detect ports that are at zero transmit credits for the specified duration. When a port is at zero Tx credits for the specified threshold value, the switch records an entry in the slowport monitor log and in logging onboard. It is shown in the **show process creditmon slowport-monitor-events** and **show logging onboard slowport-monitor-events** commands. The information shown in outputs of these commands is identical, except that the slowport-monitor log only holds the last 10 events per port whereas the logging onboard holds the events in chronological order and can hold more events compared to the slowport monitor.

  Events are recorded at a maximum frequency of 100ms. When the count goes up the operational delay is shown. That indicates the length of time where the port was at 0 Tx credits. If the count goes up by more than one from the previous entry then the operational delay is the average operational delay for the multiple events.

  In the following example, at 02/02/18 18:12:37.308 the slowport detection count was 276 and the previous value was 273. That indicates there were three intervals of time in the previous 100ms where the port was at 0 Tx credits for 1ms or more. The average time the port was at 0 credits is shown in the oper delay column (4ms). That indicates there was a total of 12ms of time the port was at 0 Tx credits in the previous 100ms. That 12ms was in three separate intervals.

  Also, port monitor can generate a slowport monitor alert. By default this is off. See port-monitor below:

  The **show process creditmon slowport-monitor-events** [**module** *number*] [**port** *number*] command shows the last 10 events per port.

```
switch# show process creditmon slowport-monitor-events

        Module: 01      Slowport Detected: NO


        Module: 09      Slowport Detected: YES
================================================================================
 Interface = fc9/2
--------------------------------------------------------------------
| admin  | slowport  | oper  |            Timestamp            |
| delay  | detection | delay |                                 |
| (ms)   | count     | (ms)  |                                 |
--------------------------------------------------------------------
|   1    |       289 |    2  | 1. 02/02/18 21:33:20.853        |
|   1    |       279 |   10  | 2. 02/02/18 21:33:20.749        |
|   1    |       279 |   19  | 3. 02/02/18 21:33:20.645        |
|   1    |       276 |    4  | 4. 02/02/18 18:12:37.308        |
|   1    |       273 |    3  | 5. 02/02/18 17:07:44.395        |
|   1    |       258 |    2  | 6. 02/02/18 13:33:08.451        |
|   1    |       254 |    1  | 7. 02/02/18 12:49:01.899        |
|   1    |       253 |   14  | 8. 02/02/18 12:49:01.794        |
|   1    |       242 |    1  | 9. 02/02/18 10:07:33.594        |
|   1    |       242 |    3  |10. 02/02/18 10:07:32.865        |
```

  The **show logging onboard slowport-monitor-events** command shows all slowport-monitor-events on a module by module basis.

```
switch# show logging onboard slowport-monitor-events module 9

---------------------------------
```

```
 Module: 9 slowport-monitor-events
---------------------------------

----------------------------
 Show Clock
----------------------------
2018-02-03 12:27:45

---------------------------------
 Module: 9 slowport-monitor-events
---------------------------------

-------------------------------------------------------------------------------
| admin   | slowport   | oper   |            Timestamp          | Interface
| delay   | detection | delay |                                |
| (ms)    | count      | (ms)   |                               |
-------------------------------------------------------------------------------
|   1     |       289 |     2   | 02/02/18 21:33:20.853         | fc9/2
|   1     |       279 |    10   | 02/02/18 21:33:20.749         | fc9/2
|   1     |       277 |    19   | 02/02/18 21:33:20.645         | fc9/2
|   1     |       276 |     4   | 02/02/18 18:12:37.308         | fc9/2
…snip
```

- Txwait (Fibre Channel and FCoE): Txwait is a measure of the time a port is at zero transmit credits in Fibre Channel or in a received PFC pause state in FCoE. It increments by one every 2.5 microseconds that a port is unable to transmit.

  Txwait is shown in the following ways:

  - Cumulative since interface counters were last cleared in the **show interface counters** command.
  - Percent unable to transmit for the last 1 second, 1 minute, 1 hour, and 72 hours in the **show interface counters** command.
  - A graphical representation of txwait for the last 60 seconds, 60 minutes, and 72 hours. In Fibre Channel, it is shown in the **show process creditmon txwait-history** command and in FCoE it is shown in the **show hardware internal txwait-history** command.
  - An entry in the On-Board Failure Log (OBFL) when a port accumulates 100 millisecond or more txwait in a 20-second interval. This is shown in the **show logging onboard txwait** command.

  The following example show interface counters command shows “18271182848 2.5us TxWait due to lack of transmit credits”. These are cumulative since the counters were last cleared or since the module first came up. In the example TxWait incremented 18271182848 times. To convert to seconds, multiply by 2.5 and divide by 1,000,000. 18271182848 * 2.5 / 1000000 = 45677.95712 seconds.

  The example also shows the percentage credits were not available over four time periods. 93% credit unavailability over the last 60 seconds would be approximately 55.8 seconds (0.93 * 60) of time the interface was at 0 Tx credits(TxWait). This would be a TxWait value of 55.8 * 1000000 / 2.5 = 22320000.

```
switch#  show interface fc1/1 counters
fc1/1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 143200 bits/sec, 17900 bytes/sec, 9 frames/sec
    38221762 frames input, 81793981988 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 81793981988 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
```

```
    1667797665 frames output, 3569086410860 bytes
      0 class-2 frames, 0 bytes
      1667797665 class-3 frames, 3569086410860 bytes
      0 class-f frames, 0 bytes
      30446972 discards, 0 errors
    30446972 timeout discards, 25 credit loss
    0 input OLS, 25 LRR, 0 NOS, 0 loop inits
    43 output OLS, 34 LRR, 24 NOS, 0 loop inits
    24 link failures, 8 sync losses, 5 signal losses
     463928 Transmit B2B credit transitions to zero
     16099 Receive B2B credit transitions to zero
      18271182848 2.5us TxWait due to lack of transmit credits
      Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/93%/49%/16%
      32 receive B2B credit remaining
      58 transmit B2B credit remaining
    Last clearing of "show interface" counters : never
```

The following example show process creditmon txwait-history command gives a histogram view of credit unavailability(TxWait). There are three graphs:

- The last 60 seconds—Each column is a second.
- The last 60 minutes—Each column is a minute
- The last 72 hours—Each column is an hour

“Time 0” on the very left is the most current time on each graph(the time when command was issued). Above each column is the amount of time Tx credits were not available in that time period in a vertical numerical format. The value in the following 60 second graph 9 seconds prior to when the command was issued, of 334ms, indicates that in that one second interval nine seconds prior to when the command was entered, there was TxWait of 334ms.

Every 60 seconds the prior 60 seconds of values are totaled and averaged and rolled into the most current minute in the “per minute” graph. Every 60 minutes the values in the “per minute” graph are totaled and averaged into the “per hour” graph.

```
switch# show process creditmon txwait-history

TxWait history for port fc1/1:
==============================
             399993 699996                          99998      24
             399990 299995                         899990      42
     000000004866970499661000000000000000000000000299996000000051
1000          ####   ####                           ####
 900          ####   ####                           ####
 800          ####   ####                           #####
 700          ####   #####                          #####
 600          ####  ######                          #####
 500          ####  ######                          #####
 400          ####  ######                          #####       #
 300         ###### ######                          #####       #
 200         ###### ######                          #####      ##
 100         ###### ######                         ######      ##
    0....5....1....1....2....2....3....3....4....4....5....5....6
              0    5    0    5    0    5    0    5    0    5    0

          Tx Credit Not Available per second (last 60 seconds)
                 # = TxWait (ms)

   3614662  2545516555 45544 55355555  5551 2555  654  554  244
   10180070029314403740584841943527260098480078803087006752063 6
```

```
   ............................................................
   80310030098893202440334416996826130048720435606082003916023 1
60  #  ##    #    # #   #    #    #    ##    ###  ##    #
54  #  ##    # ## ####  ##   ## #####  ###   ###  ##   ##
48  # ###    # ## #### ####  ## #####  ###   ###  ###  ###    #
42  # ###    #### #### ##### ## #####  ###   ###  ###  ###   ##
36  # ###    #### #### ##### ########  ###   ###  ###  ###   ##
30 ## ####   #### #### ##### ########  ###   ###  ###  ###   ##
24 ## ####  ##### #### ##### ########  ###   ###  ###  ###  ###
18 ## ####  ##### #### ##### ########  #### ####  ###  ###  ###
12 #######  ########## ##### ########  #### ####  ###  ###  ###
 6 #######  ########## ##### ########  #### #### ####  ###  ###
  0....5....1....1....2....2....3....3....4....4....5....5....6
            0    5    0    5    0    5    0    5    0    5    0

          Tx Credit Not Available per minute (last 60 minutes)
                 # = TxWait (secs)

    2122222211222222221
    27234106299161216437
    96976917464816256020
    960359822158274304920000000000000000000000000000000000000000000000000000000000000000000
3600
3240
2880
2520    ##  #    #   ##
2160 # #######  ########
1800 ####################
1440 ####################
1080 ####################
 720 ####################
 360 ####################
   0....5....1....1....2....2....3....3....4....4....5....5....6....6....7.7
            0    5    0    5    0    5    0    5    0    5    0    5    0 2

          Tx Credit Not Available per hour (last 72 hours)
                 # = TxWait (secs)
```

Also, port monitor can generate a txwait alert (Fibre Channel only).

- TxWait or RxWait (FCoE only): TxWait is a measure of time the VFC is in a receive PFC pause state preventing it from transmitting. Rxwait is a measure of time a port is in a transmit PFC pause state preventing the adjacent device from transmitted to it. It increments by one every 2.5 microseconds that a port is unable to receive.

  TxWait or RxWait is shown in various ways.

  - Cumulative since interface counters were last cleared in the **show interface counters** command.
  - Percent unable to receive for the last 1 second, 1 minute, 1 hour, and 72 hours in the **show interface counters** command.
  - A graphical representation of txwait or rxwait for the last 60 seconds, 60 minutes, and 72 hours. In FCoE, it is shown in the **show hardware internal txwait-history** command for txwait or the **show hardware internal rxwait-history** command for rxwait.
  - An entry in the OBFL when a port accumulates 100 millisecond or more rxwait in a 20-second interval. This is shown in the **show logging onboard txwait** command for txwait or the **show logging onboard rxwait** command for rxwait.

The following example show interface counters command shows “1104349910 2.5 us TxWait due to pause frames (VL3)”. These are cumulative since the counters were last cleared or since the module first came up. In the example TxWait incremented 1104349910 times. To convert to seconds, multiply by 2.5 and divide by 1,000,000. 1104349910 * 2.5 / 1000000 = 2760.874 seconds of time the VFC was unable to transmit.

The following example show interface counters command shows “205484298144 2.5us RxWait due to PFC Pause frames (VL3)”. These are cumulative since the counters were last cleared or since the module first came up. In the example RxWait incremented 205484298144 times. To convert to seconds, multiply by 2.5 and divide by 1,000,000. 205484298144 * 2.5 / 1000000 = 513710.745 seconds of time the VFC was unable to receive.

The example also shows the percentage of time the VFC was paused in each direction over the last 1 second, 1 minute, 1 hour and 72 hours. For TxWait this is the percentage of time the VFC received PFC pauses. For RxWait this is the percentage of time the VFC was sending pause frames preventing the other side from transmitting. In the example, in the last one minute the VFC was prevented from transmitting (TxWait) 33% of the time(20 seconds).

```
switch# show interface vfc-po540 counters

vfc-po540
    1571394073 fcoe in packets
    3322884900540 fcoe in octets
    79445277 fcoe out packets
    69006091691 fcoe out octets
    1104349910 2.5 us TxWait due to pause frames (VL3)
    205484298144 2.5 us RxWait due to pause frames (VL3)
    0 Tx frames with pause opcode (VL3)
    3302000 Rx frames with pause opcode (VL3)
Percentage pause in TxWait per VL3  for last 1s/1m/1h/72h: 0%/33%/0%/0%
Percentage pause in RxWait per VL3  for last  1s/1m/1h/72h: 0%/0%/0%/30%
```

- Tx-credit-not-available or rx-credit-not-available (Fibre Channel only): Tx-credit-not-available or rx-credit-not-available is a counter that increments by 1 when a port is at zero transmit credits for 100 consecutive milliseconds. This is done by a software process. Tx-credit-not-available is available in the **show logging onboard error-stats** command, and also available using the **show system internal snmp credit-not-available** command. Also, port monitor can generate a tx-credit-not-available or rx-credit-not-available alert.

  Show logging onboard error-stats records various counters when they change value. When a counter’s value changes the current value is recorded along with the date and time. These are only updated periodically at regular intervals like 20 seconds. So if a counter’s value changed in the 20 second interval several times, then only the value that existed at the end of the 20 second interval is recorded. To determine the amount the counter incremented the previous value must be located and subtracted from the current value. In the following example the FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO counter incremented to 7377473 for port fc1/13 on 12/06/17 06:15:20. The previous value was 7377461 at 12/06/17 06:15:00. So in the time interval between 06:15:00 and 06:15:20 the FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO incremented by 7377473 – 7377461 = 12.

```
switch# show logging onboard error-stats | i WT
fc1/13             |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO                        |7377473
|12/06/17 06:15:20
fc1/6              |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO                        |337
|12/06/17 06:15:20
fc1/4              |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO                        |7428998
|12/06/17 06:15:20
```

```
fc1/13              |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO                      |7377461
|12/06/17 06:15:00
fc1/6               |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO                      |324
|12/06/17 06:15:00
fc1/4               |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO                      |7428994
|12/06/17 06:15:00
```

Show logging onboard error-stats has several different counters that pertain to SAN congestion. Most of these counters are module/switch dependent. For Tx-credit-not-available or Rx-credit-not-available they are the following:

FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO[5, 50i,48S,96S]

F32_MAC_KLM_CNTR_TX_WT_AVG_B2B_ZERO[6]

Count of the number of times that an interface was at zero Tx B2B credits for 100 ms. This status typically indicates congestion at the device attached on that interface.

FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO[5,50i,48S,96S]

F32_MAC_KLM_CNTR_RX_WT_AVG_B2B_ZERO[6]

Count of the number of times an interface was at zero Rx B2B credits for 100 ms; this status typically indicates that the switch is withholding R_Rdy primitive to the device attached on that interface due to congestion in the path to devices with which it is communicating.

Also, port monitor can generate tx-credit-not-available or rx-credit-not-available alerts (Fibre Channel only). See the port monitor section.

- Interface priority flow control (FCoE only): Cisco MDS switches track PFC pause frames (transmitted and received). This is available in the **show interface vfc**, **show interface priority-flow-control**, and **show logging onboard error-stats** commands.
- Port monitor tx datarate counter (Fibre Channel only): Port-monitor tx-datarate counter can be configured to provide information when a port becomes highly utilized in the transmit (Tx) direction. Port monitor has two thresholds called a rising-threshold and a falling threshold. The rising-threshold is when the port's Tx datarate hits or exceeds the percentage of the operational link speed. The falling-threshold is when the Tx datarate falls to that percentage or below. For each event an alert is generated. The time the port was highly utilized is the time between the rising-threshold and falling-threshold alerts. These alerts are recorded in the RMON log in all releases. In NX-OS 8.2(1) and later these are logged in OBFL and are shown via the **show logging onboard datarate** command.
- Port monitor (Fibre Channel only): Port monitor can generate alerts for various congestion-related counters. Port monitor has two thresholds called a rising-threshold and a falling threshold. The rising-threshold is when the port's counter hits or exceeds the configured value. The falling-threshold is when the port's counter falls to that configured value or below. For each event an alert is generated. The time the port was between the rising threshold and the falling threshold is when the event was occurring. These alerts are recorded in the RMON log in all releases.

Table 18: Features to Detect Slow Drain, on page 138 describes the features that help detect slow drain:

Table 18: Features to Detect Slow Drain

Feature Name	Description
Port monitor's credit-loss-reco counter	credit-loss-reco counter resets a link when there is not enough transmit credits available for 1 second for edge ports and 1.5 second for core ports.

Feature Name	Description
Port monitor's invalid-crc counter	invalid-crc counter represents the total number of frames with CRC errors received by a port.
Port monitor's invalid-words counter	invalid-words counter represents the total number of invalid words received by a port.
Port monitor's link-loss counter	link-loss counter represents the total number of link failures encountered by a port.
Port monitor's lr-rx counter	lr-rx counter represents the total number link reset primitive sequence received by a port.
Port monitor's lr-tx counter	lr-tx counter represents the total number link reset primitive sequence transmitted by a port.
Port monitor's rx-datarate counter	rx-datarate counter receives frame rate in bytes per seconds.
Port monitor's signal-loss counter	signal-loss counter represents the number of times a port encountered laser or signal loss.
Port monitor's state-change counter	state-change counter represents the number of times a port has transitioned to an operational up state.
Port monitor's sync-loss counter	sync-loss counter represents the number of times a port experienced loss of synchronization in RX.
Port monitor's tx-credit-not-available counter	tx-credit-not-available counter increments by one if there is no transmit buffer-to-buffer credits available for a duration of 100 ms.
Port monitor's timeout-discards counter	timeout-discards counter represents the total number of frames dropped at egress due to congestion timeout or no-credit-drop timeout.
Port monitor's tx-datarate counter	tx-datarate counter represents the transmit frame rate in bytes per seconds.
Port monitor's tx-discards counter	tx-discards counter represents the total number of frames dropped at egress due to timeout, abort, offline, and so on.
Port monitor's tx-slowport-count counter	tx-slowport-count counter represents the number of times slow port events were detected by a port for the configured slowport-monitor timeout. This is applicable only for Generation 3 modules.
Port monitor's tx-slowport-oper-delay counter	tx-slowport-oper-delay counter captures average credit delay (or R_RDY delay) experienced by a port. The value is in milliseconds.
Port monitor's txwait counter	txwait counter is an aggregate time counter that counts transmit wait time of a port. Transmit wait is a condition when port experiences no transmit credit available (tx b2b = 0) and frames are waiting for transmission.

Information About Congestion Avoidance

Congestion avoidance focuses on minimizing or completely avoiding the congestion that results from frames being queued to congested ports.

Cisco MDS switches have multiple features designed to void congestion in SAN:

- Congestion-drop timeout threshold (Fibre Channel and FCoE): The congestion-drop timeout threshold determines the amount of time a queued Fibre Channel or FCoE frame will stay in the switch awaiting transmission. Once the threshold is reached the frame is discarded as a *timeout drop*. The lower the value the quicker these queued frames are dropped and the result buffer freed. This can relieve some back pressure in the switch, especially on ISLs. By default it is 500 ms but can be configured as low as 200 ms in 1 ms increments. It is configured using the **system timeout congestion-drop** (Fibre Channel) and **system timeout fcoe congestion-drop** (FCoE) commands.
- No-credit-drop timeout threshold (Fibre Channel only): No-credit-drop timeout threshold is used to time when a Fibre Channel port is at zero Tx credits. Once a Fibre Channel port hits zero Tx credits the timer is started. If the configured threshold is reached then all frames queued to that port will be dropped regardless of their actual age in the switch. Furthermore, as long as the port remains at zero Tx credits, all newly arriving frames are immediately dropped. This can have a dramatic effect on reliving congestion especially on upstream ISLs. This allows unrelated flows to move continuously. This is off by default. If configured, it should be set to a value that is lower than the configured (or defaulted) Fibre Channel congestion-drop timeout. It is configured via the system timeout no-credit-drop command. The no-credit timeout functionality is only used for edge ports because these ports are directly connected to the slow-drain devices.
- Pause-drop timeout threshold (FCoE only): Pause-drop timeout threshold is used to time when a FCoE port is in a continuous state of Rx pause (unable to transmit). After an FCoE port receives a PFC pause with a non-zero quanta, the timer is started. If the port continues to receive PFC pauses with a non-zero quanta such that it remains in the Rx pause state continuously for the pause-drop threshold, then all frames queued to that port will be dropped regardless of their actual age in the switch. Furthermore, as long as the port remains in a Rx pause state, all newly arriving frames are immediately dropped. This can have a dramatic effect on relieving congestion especially on the upstream ISLs. This allows unrelated flows to move continuously. This is on by default with a value of 500 ms. If configured, it should be set to a value that is lower than the configured (or defaulted) FCoE congestion-drop timeout. It is configured via the **system timeout fcoe pause-drop** commands (available from Cisco MDS NX-OS Release 8.2(1) onwards). The FCoE pause-drop timeout functionality is only used for edge ports, because these ports are directly connected to the slow-drain devices.
- Port monitor with portguard actions of flap and error disable: For more information, see the Port Monitor, on page 27 section.

Note The *no-credit timeout* functionality is only used for edge ports because these ports are directly connected to the slow-drain devices. The no-credit timeout functionality is not supported on Generation 1 modules.

Information About Congestion Isolation

The Congestion Isolation feature can detect a slow-drain device via port monitor or manual configuration and isolate the slow-drain device from other normally performing devices on an ISL. After the traffic to the slow-drain device is isolated, the traffic to the rest of the normally behaving devices will be unaffected. Traffic isolation is accomplished using the following three features:

- Extended Receiver Ready—This feature allows each ISL between supporting switches to be split into four separate virtual links, with each virtual link assigned its own buffer-to-buffer credits. Virtual link 0 used to carry control traffic, virtual link 1 is used to carry high-priority traffic, virtual link 2 is used to carry slow devices, and virtual link 3 is used to carry normal traffic.
- Congestion Isolation—This feature allows devices to be categorized as slow by either configuration command or by port monitor.
- Port monitor portguard action for Congestion Isolation—Port monitor has a new portguard option to allow the categorization of a device as slow so that it can have all traffic flowing to the device routed to the slow virtual link.

Extended Receiver Ready

Note Extended Receiver Ready (ER_RDY) feature functions only on Fibre Channel Inter-Switch Links (ISL) and only between switches that support this feature.

ER_RDY primitives are used as an alternative to Receiver Ready (R_RDY). ER_RDY primitives virtualize a physical link into multiple virtual links (VLs) that are assigned individual buffer-to-buffer credits, thereby controlling the flow to the physical link. The ER_RDY feature is used by Congestion Isolation to route slow flows to a specific VL, called a low-priority VL (VL2), so that all the normal flows are unaffected. ER_RDY supports up to four VLs.

Figure 2: Traffic Flow Using Virtual Links, on page 141 shows VLs managing the good flow and slow flow. VL0 (red link) is used for control traffic, VL1 (orange link) is used for high-priority traffic, VL2 (blue link) is used for slow traffic, and VL3 (green link) is used for normal-data traffic. Slow flow detected at Host H2 is automatically assigned to VL2, which prevents the congestion of the link and allows the good flow from Host H1 to use either the VL1 or VL3 depending on the flow priority.

Figure 2: Traffic Flow Using Virtual Links

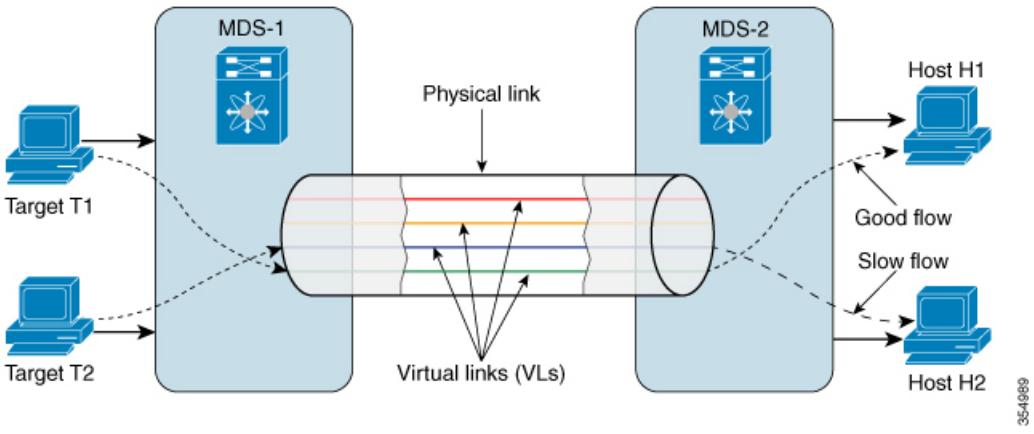

Table 19: Virtual Link-to-QoS Priority Mapping, on page 142 provides VL-to-QoS priority mapping information. Use this information while setting a zone QoS priority in a zone where Congestion Isolation is enabled in order to avoid QoS priority flow from being treated as slow flow.

Table 19: Virtual Link-to-QoS Priority Mapping

Virtual Link	QoS Priority
VL0 (control traffic)	7
VL1 (not used for any traffic)	5, 6
VL2 (slow traffic)	2, 3, 4
VL3 (normal traffic)	0, 1

Congestion Isolation

The Congestion Isolation feature uses VL capabilities to isolate the flows to the slow devices on an ISL to a low-priority VL that has less buffer-to-buffer credits than the buffer-to-buffer credits used for the normal traffic VL. Traffic in the direction of the slow device is routed to a low-priority VL. Normal devices continue to use the normal VL that has more buffer-to-buffer credits. Slow devices can be marked as slow either via the port monitor or manually, using the **congestion-isolation include pwwn** *pwwn* **vsan** *vsan-id* command.

Note When a device is manually marked as slow using the **congestion-isolation include pwwn** *pwwn* **vsan** *vsan-id* command or automatically detected as slow via the port monitor, the Fibre Channel Name Server (FCNS) database registers the slow-device attribute (slow-dev) for the device and distributes the information to the entire fabric.

You must ensure that the following requirements are met before enabling the Congestion Isolation feature:

- Flows must traverse ISLs because Congestion Isolation functions only across Fibre Channel ISLs.
- ISLs or port channels must be in ER_RDY flow-control mode.
- If you want the port monitor to automatically detect the slow devices, the port-monitor policies must be configured to use the congestion isolation port-guard action (cong-isolate).

  Optionally, devices can be configured manually as slow using the **congestion-isolation include pwwn** *pwwn* **vsan** *vsan-id* command.

Port-Monitor Portguard Action for Congestion Isolation

The cong-isolate port-monitor portguard action automatically isolates a port after a given event rising-threshold is reached.

Note Absolute counters do not support portguard actions. However, the tx-slowport-oper-delay absolute counter supports Congestion Isolation portguard action (cong-isolate).

The following is the list of counters that you can use to trigger the Congestion Isolation port-monitor portguard action (cong-isolate):

- credit-loss-reco

- tx-credit-not-available
- tx-slowport-oper-delay
- txwait

Guidelines and Limitations for Congestion Detection, Avoidance, and Isolation

Guidelines and Limitations for Congestion Detection

The **show tech-support slowdrain** command contains all the congestion detection indications, counters, and log messages as well as other commands that allow an understanding of the switches, MDS NX-OS versions, and topology. Since, congestion can propagate from one switch to another, the **show tech-support slowdrain** command should be gathered from all the switches at approximately the same time to have the best view of where the congestion started and how it spread. This can be easily done via the DCNM SAN client using the **Tools**-> **Run CLI** feature. This feature will issue a command or commands to all the switches in the fabric and consolidates the individual switch outputs into a single fabric zip file.

Some commands display simple counters such as the **show interface counters** command, whereas some commands display counter information with accompanying date and time stamps. The commands that display counters with accompanying date and time stamps are mostly the **show logging onboard** commands.

There are various “sections” of show logging onboard that contain information pertaining to slow drain and over utilization. Most “sections” will update periodically and include counters only when they actually change in the prior interval. Different sections have different update periods. They are:

- Error-stats—Includes many error counters accompanying date and time stamps
- Txwait—Includes interfaces that record 100ms or more of TxWait in a 20 second interval. The values displayed are not the current value of TxWait, but only deltas from the previous 20 second interval. If TxWait incremented by the equivalent of less than 100ms there is no entry.
- Rxwait—Includes interfaces that record 100ms or more of RxWait in a 20 second interval. The values displayed are not the current value of RxWait, but only deltas from the previous 20 second interval. If RxWait incremented by the equivalent of less than 100ms there is no entry.

When a counter increments in the interval the current value of the counter is displayed along with the date and time when the counter was checked. To determine the amount the counter incremented, the delta value, in the interval the current value must be subtracted from the previously recorded value.

For example, the following show logging onboard error-stats output shows that when the counter was checked at 01/12/18 11:37:55 the timeout-drop counter, F16_TMM_TOLB_TIMEOUT_DROP_CNT, for port fc1/ was a value of 743. The previous time it incremented was 12/20/17 06:31:47 and it was a value of 626. This means that since error-stats interval is 20 seconds, between at 01/12/18 11:37:35 and at 01/12/18 11:37:55 the counter incremented by 743 – 626 = 117 frames. There were 117 frames discarded at timeout-drops during that interval.

```
switch# show logging onboard error-stats

----------------------------
 Show Clock
----------------------------
2018-01-24 15:01:35

---------------------------------
 Module: 1 error-stats
---------------------------------
```

```
-------------------------------------------------------------------------------------------------

 ERROR STATISTICS INFORMATION FOR DEVICE DEVICE: FCMAC
-------------------------------------------------------------------------------------------------

    Interface      |                                          |         |    Time Stamp
      Range        |                Error Stat Counter Name   |  Count  |MM/DD/YY HH:MM:SS
                   |                                          |         |
-------------------------------------------------------------------------------------------------
fc1/8              |F16_TMM_TOLB_TIMEOUT_DROP_CNT             |743      |01/12/18 11:37:55
fc1/8              |F16_TMM_TOLB_TIMEOUT_DROP_CNT             |626      |12/20/17 06:31:47
fc1/5              |F16_TMM_TOLB_TIMEOUT_DROP_CNT             |627      |12/20/17 06:31:47
fc1/3              |F16_TMM_TOLB_TIMEOUT_DROP_CNT             |556      |12/20/17 06:31:47
fc1/8              |F16_TMM_TOLB_TIMEOUT_DROP_CNT             |623      |12/20/17 04:05:05
```

Guidelines and Limitations for Congestion Avoidance

The default value for system timeout congestion-drop is 500 ms.

System timeout no-credit-drop is disabled by default. When configured, this feature reduces the effects of slow drain in the fabric. However, if it is configured to a value that is too low, it can cause disruption. The disruption is caused because many frames are discarded when a device withholds credits for even a short duration. The lower the value, the quicker it can start discarding frames that are queued from an upstream ISL to this (slow) port. This will relieve the back pressure or congestion on that ISL and allow other normally performing devices to continue their operation. The actual value chosen is fabric and implementation dependent.

Following are some guidelines for choosing the system timeout congestion-drop value:

- 200 ms—Safe value for most fabric
- 100 ms—Aggressive value
- 50 ms—Very aggressive value

Generally, before configuring the no-credit-drop value, the switches should be checked for the presence of large amounts of continuous time at zero Tx credits. The show logging onboard start time mm/dd/yy-hh:mm:ss error-stats command can be run looking for instances of the FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO counter indicating 100ms intervals at zero credits. Additionally, the port-monitor tx-credit-not-available and the show system internal snmp credit-not-available command will show similar information. Only when the fabric only shows very limited amounts of 100ms at zero Tx credits should no-credit-drop be considered. If there are large amounts of ports with 100ms at zero Tx credits, then the problems with those end devices should be investigated and resolved prior to configuring no-credit-drop.

Note No-credit-drop can only be configured on ports that are classified *logical-type edge*. These are typically F ports.

Slowport-monitor, if configured, must have a value lower than no-credit-drop since it will only indicate a slow port if the port has no credits for at least the amount of time configured and there are frames queued for transmit. Since no-credit-drop will drops any frames queued for transmit, if no-credit-drop is configured for a value equal to or less than slowport-monitor, there will be no frames queued for transmit and slowport-monitor will not detect the slow port.

Guidelines and Limitations for Congestion Isolation

Extended Receiver Ready

- ER_RDY is supported only on Fibre Channel ports on Cisco MDS 9700 Series with Cisco MDS 9700 16-Gbps Fibre Channel Switching Module (DS-X9448-768K9), Cisco MDS 9000 Series 24/10 SAN Extension Module (DS-X9334-K9) (Fibre Channel ports only), Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9), and MDS 9396S switches. In a fabric consisting of supported and unsupported switches (mixed fabric), this feature may not work effectively. In a mixed fabric, ER_RDY flow-control mode is used only between supported switches and R_RDY flow-control mode is used between unsupported switches.
- Trunking must be enabled on all ISLs in the topology for ER_RDY flow-control mode to work.
- After the **system fc fow-control er_rdy** command is configured on both the local switch and its adjacent switch, the ISLs connecting the switches should be flapped to put the ISLs in ER_RDY flow-control mode. In port channels, these links can be flapped one at a time, preventing loss of connectivity.
- For migration purposes, port channels can have their member links in both R_RDY and ER_RDY flow-control modes. This is to facilitate nondisruptive conversion from R_RDY to ER_RDY flow-control mode. Do not allow this inconsistent state to persist longer that it takes to perform the conversion from R_RDY to ER_RDY flow-control mode.
- VL1 is reserved for host bus adapter (HBA) and zone quality of service (QoS).
- Inter VSAN Routing (IVR), Fibre Channel Redirect (FCR), Fibre Channel Over TCP/IP (FCIP), and Fibre Channel over Ethernet (FCoE) are not supported in ER_RDY flow-control mode.
- In-Order Delivery (IOD) may get affected when the flow-control mode is initially set to ER_RDY and when the device's flows are moved from one VL to another VL.
- Switches running releases prior to Cisco MDS NX-OS Release 8.1(1) in a fabric are unaware of slow devices. Upon upgrading to Cisco MDS NX-OS Release 8.1(x) or later, these switches become aware of the slow devices.
- If you have configured the buffer-to-buffer credits using the **switchport fcrxbbcredit** *value* command in the Cisco MDS NX-OS Release 7.3(x) or earlier, upgraded to Cisco MDS NX-OS Release 8.1(1), and set flow-control mode to ER_RDY, the buffer-to-buffer credits that are already configured get distributed to the VLs in the following manner:
  - If the buffer-to-buffer credits value that is configured is 50, the default buffer-to-buffer credit values 5, 1, 4, and 40 are allocated to VL0, VL1, VL2, and VL3 respectively.
  - If the buffer-to-buffer credits value that is configured is more than 34 and less than 50, the buffer-to-buffer credits get distributed in the ratio 5:1:4:40.
  - If the buffer-to-buffer credits value that is configured is more than 50, the default values 5, 1, 4, and 40 are allocated to VL0, VL1, VL2, and VL3 respectively. The remaining buffer-to-buffer credits get distributed in the ratio 15:15:40:430 (VL0:VL1:VL2:VL3).
  - If you are upgrading or if you are in the Cisco MDS NX-OS Release 8.1(1), if ER_RDY was enabled, and if the buffer-to-buffer credits value that is configured is less than 34, the VLs are stuck in the initialization state because the control lane (VL0) is allocated 0 credits. To recover from this situation, shutdown the link and allocate more than 34 buffer-to-buffer credits using the **switchport**

fcrxbbcredit *value* or allocate at least one buffer-to-buffer credit to VL0, using the **switchport vl-credit vl0** *value* **vl1** *value* **vl2** *value* **vl3** *value* command.

Note The sum of the buffer-to-buffer credits configured for VLs cannot exceed 500.

- If you had configured the buffer-to-buffer credits using the **switchport fcrxbbcredit** *value* **mode E** command, and used the **switchport vl-credit vl0** *value* **vl1** *value* **vl2** *value* **vl3** *value* command to set the new buffer-to-buffer credits values for the VLs, the sum of the configured buffer-to-buffer credits for VLs are pushed to the **switchport fcrxbbcredit** *value* **mode E** command.
- Use the **no switchport fcrxbbcredit** *value* or **switchport vl-credit default** command to set the default buffer-to-buffer credits value for the VLs.
- If you have configured the extended buffer-to-buffer credits using the **switchport fcrxbbcredit extended** *value* in the Cisco MDS NX-OS Release 7.3(x) or earlier, upgraded to Cisco MDS NX-OS Release 8.1(1), and set the flow-control mode to ER_RDY, the extended buffer-to-buffer credits that are already configured are distributed to the VLs in the following manner:
  - If the buffer-to-buffer credits value that is configured is less than 50, the minimum values 5, 1, 4, and 40 are allocated to VL0, VL1, VL2, and VL3 respectively.
  - If the buffer-to-buffer credits value that is configured is more than 34 and less than 50, the buffer-to-buffer credits get distributed in the ratio 5:1:4:40.
  - If the buffer-to-buffer credits value that is configured is more than 50, the minimum values 15, 15, 4, and 430 are allocated to VL0, VL1, VL2, and VL3 respectively. The remaining buffer-to-buffer credits are distributed in the ratio 30:30:100:3935 (VL0:VL1:VL2:VL3).
  - If you are upgrading to or if you are in the Cisco MDS NX-OS Release 8.1(1), ER_RDY is enabled, and the buffer-to-buffer credits value configured is less than 34, the VLs are stuck in the initialization state because the control lane (VL0) is allocated 0 credits. To recover from this situation, shutdown the link and allocate more than 34 buffer-to-buffer credits using the **switchport fcrxbbcredit** *value* or allocate at least one buffer-to-buffer credit to VL0, using the **switchport vl-credit vl0** *value* **vl1** *value* **vl2** *value* **vl3** *value* command.

Note The sum of the extended buffer-to-buffer credits configured for VLs cannot exceed 4095 on a Cisco MDS 9700 16-Gbps Fibre Channel Switching Module, and 8270 on a Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module.

- You cannot configure regular buffer-to-buffer credits after you configure the extended buffer-to-buffer credits. You must first disable the extended buffer-to-buffer credits using the **no fcrxbbcredit extended enable** command and then configure the regular buffer-to-buffer credits.
- You cannot disable the extended buffer-to-buffer credits configuration even if one link is running in the extended buffer-to-buffer credits mode.
- ER_RDY is not supported on interfaces whose speed is set to 10-Gbps.

Congestion Isolation

- Congestion Isolation is disabled by default.
- Congestion Isolation and its configurations are applicable only to the switch being configured, and not to the entire fabric.
- If you enable the ER_RDY and Congestion Isolation features on a supported switch before adding it to a fabric that is using ER_RDY flow-control mode, the ISLs that are connected between the supported switch and its adjacent switch are automatically in the ER_RDY flow-control mode and you need not flap the links on the switch for the links to use the ER_RDY flow-control mode.
- In a fabric consisting of supported and unsupported switches, Congestion Isolation functions as desired only between supported switches. Congestion Isolation functionality between unsupported devices is unpredictable.
- After a device is detected as slow, only the traffic moving in the direction of the slow device is routed to a low-priority VL (VL2). Traffic in the reverse direction is not classified as slow, and is unaffected.
- When a slow device is detected or a device is configured as slow, the switch sends an FCNS notification to all the other switches that are capable of supporting the Congestion Isolation feature and also to the switches that may not have this feature enabled. If the switch is capable of supporting this feature but does not have it enabled, then the FCNS notification is rejected and the following messages are displayed at the originating switch:
  - %FCNS-2-CONGESTION_ISOLATION_FAILURE: %$VSAN vsan-id%$ SWILS-RJT received from domain domain-id for congestion-isolation. Issue includes CLI/FCNS DB refresh on the remote domain.
  - %FCNS-2-CONGESTION_ISOLATION_INT_ERROR: %$VSAN 237%$ Error reason: Congestion-Isolation disabled on the remote domain. Please enable the feature on the remote domain.

  If the Congestion Isolation feature is configured on all the intended switches, these messages do not have any negative effect and can be ignored. For example, if a Cisco MDS switch is connected via FCoE ISLs then the Congestion Isolation feature does not apply to this switch and these messages can be ignored. However, ER_RDY and Congestion Isolation features can be configured on an FCoE connected switch preventing the messages from being displayed.
- Figure 3: Traffic Flow When Multiple Targets are Connected shows a fabric that has multiple targets connected to switch SW1 and two hosts (Host H1 and Host H2) connected to switch SW2. Both hosts H1 and H2 are zoned with all four targets T1 to T4. Host H2 is detected as a slow device. The traffic from the targets to host H2 is marked as slow and is routed to VL2. Since VL2 has fewer buffer-to-buffer credits and because host H2 is itself withholding buffer-to-buffer credits from SW2, traffic on VL2 from SW1 to SW2 will be constrained by what host H2 can receive. This results in switch SW1 withholding buffer-to-buffer credits from all four targets T1 to T4. This will affect all traffic being sent by the targets to any destination. Consequently, other hosts zoned with the targets, like host H1, will also see their traffic affected. This is an expected behavior. In such a situation, resolve the slow-drain condition for the traffic to flow normally.

Figure 3: Traffic Flow When Multiple Targets are Connected

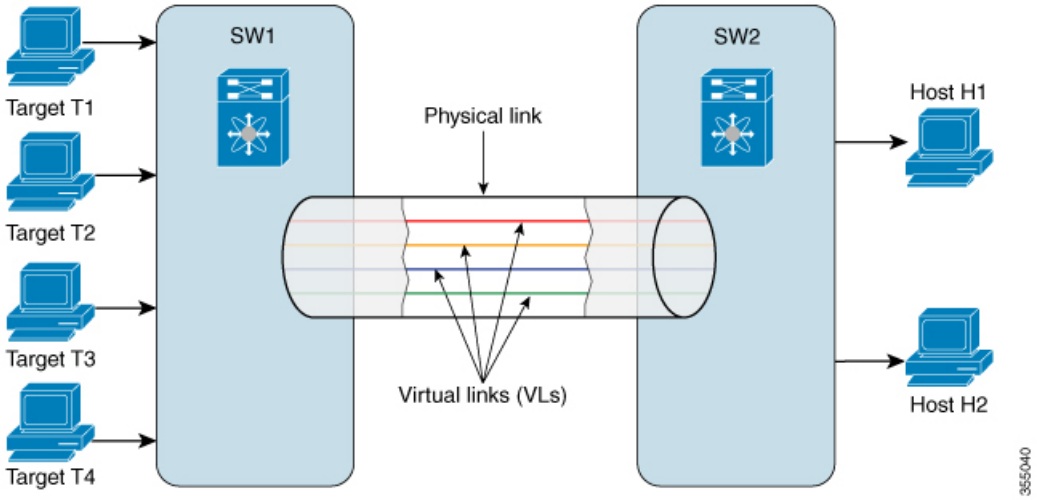

- If in a zone, the zone QoS priority is set to medium and Congestion Isolation is enabled on the switches in the zone, the traffic with zone QoS priority medium are treated as slow, and Congestion Isolation routes the traffic to the low-priority VL (VL2). To avoid this situation, set the zone QoS priority to low or high.
- When a link to a Cisco NPV switch carrying multiple fabric logins (FLOGIs) is detected as a slow device, all the devices connected to the Cisco NPV switch are marked as slow devices.
- Downgrading from a supported release to an unsupported release is disabled after the Congestion Isolation feature is enabled. To downgrade to an unsupported release:

1. If cong-isolate port monitor portguard action is configured in a port monitor policy, use the **no counter** {**credit-loss-reco** | **tx-credit-not-available** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *event-id* **warning-threshold** *count2* **falling-threshold** *count3* **event** *event-id* **portguard cong-isolate** command to remove the action from the policy.
2. Use the **no congestion-isolation** {**include** | **exclude**} **pwwn** *pwwn* **vsan** *vsan-id* command to remove any devices that are manually included or excluded as slow-drain devices.
3. Use the **no feature congestion-isolation** command to disable the Congestion Isolation feature.
4. Use the **no system fc flow-control er_rdy** command to reset the flow-control mode to R_RDY.
5. Flap all the ISLs using the **shutdown** and **no shutdown** commands.
6. Use the **show flow-control r_rdy** command to display the ISLs currently functioning in R_RDY mode.
7. Use the **show flow-control er_rdy** command to display the ISLs currently functioning in ER_RDY mode.

Note The port monitor detects slow devices when a given rising-threshold is reached and triggers the congestion isolation feature in the switch to move traffic to that slow device into the slow Virtual Link (VL2). The switch does not automatically remove any devices from congestion isolation. This must be done manually once the problem with the slow device is identified and resolved.

Configuring Congestion Avoidance

The following features can be configured for congestion avoidance:

- Congestion-drop
- No-credit-drop
- Pause-drop
- Port-monitor portguard action for congestion avoidance (errdisable and flap)

Configuring Congestion Detection

Most of the features used for congestion detection are enabled by default and do not require any additional configuration. These features include txwait, rxwait, interface priority flow control, OBFL error stats, and tx-credit-not-available. The following congestion detection features are configurable.

Modules and switches included in “Module and Switch Support” section of Table 20.

- 16-Gbps modules or switches:
  - Cisco MDS 9700 Series 16-Gbps Fibre Channel Module (DS-X9448-768K9)
  - Cisco MDS 9000 Series 24/10 SAN Extension Module (DS-X9334-K9)
  - Cisco MDS 9250i Fabric Switch
  - Cisco MDS 9148S Fabric Switch
  - Cisco MDS 9396S Fabric Switch
  - 32-Gbps modules or switches:
    - Cisco MDS 9000 Series 32-Gbps Fibre Channel Module (DS-X9648-1536K9)
    - Cisco MDS 9132T Fibre Channel Switch
- 10-Gbps FCoE module:
  - Cisco MDS 9700 48-Port 10-Gbps Fibre Channel over Ethernet (DS-X9848-480K9)
- 40-Gbps FCoE module:
  - Cisco MDS 9700 40-Gbps 24-Port Fibre Channel over Ethernet Module (DS-X9824-960K9)

Table 20: Slow Port Monitor Support on Fibre Channel and FCoE Switching Modules, on page 152 displays the congestion detection features supported on different Fibre Channel and FCoE switching modules for the Cisco MDS NX-OS Release 8.x.

Table 20: Slow Port Monitor Support on Fibre Channel and FCoE Switching Modules

Function	Module and Switch Support	
	16 Gbps and 32 Gbps Fibre Channel	10 Gbps and 40 Gbps FCoE
Txwait OBFL logging	Yes	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Txwait port monitor counter	Yes	No
Txwait interface counter	Yes	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Txwait interface unable to transmit for the last 1 second, 1 minute, 1 hour, and 72 hours	Yes	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
A graphical representation of txwait for the last 60 seconds, 60 minutes, and 72 hours	Yes	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Rxwait OBFL logging	No	Yes
Rxwait interface counter	No	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Rxwait interface unable to receive for the last 1 second, 1 minute, 1 hour, and 72 hours	No	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
A graphical representation of rxwait for the last 60 seconds, 60 minutes, and 72 hours	No	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Port monitor slow-port counter	Yes	No
OBFL error stats	Yes	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.
Interface priority flow control	No	Yes, from Cisco MDS NX-OS Release 8.2(1) onwards.

Configuring the Slow-Port Monitor Timeout Value for Fibre Channel

The slow-port monitor functionality is similar to the no-credit frame timeout and drop functionality, except that it does not drop frames; it only logs qualifying events. When a Fibre Channel egress port has no transmit credits continuously for the slow-port monitor timeout period, the event is logged. No frames are dropped unless the no-credit frame timeout period is reached and no-credit frame timeout drop is enabled. If the no-credit frame timeout drop is not enabled, no frames are dropped until the congestion frame timeout period is reached.

Slow-port monitoring is implemented in the hardware, with the slow-port monitor functionality being slightly different in each generation of hardware. The 16-Gbps and 32-Gbps modules and switches can detect each instance of the slow-port monitor threshold being crossed. The slow-port monitor log is updated at 100-ms intervals. A log for a slow-port event on a 16-Gbps and 32-Gbps module or system increments the exact number of times the threshold is reached.

Slow port monitor can also generate an alert and syslog message via port monitor.

To configure the slow-port monitor timeout value, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the slow-port monitor timeout value:

switch(config)# **system timeout slowport-monitor** *milliseconds* **logical-type** {**core | edge**}

Valid values for the slow-port monitor timeout are:

- 32-Gbps and 16-Gbps modules or switches—1 to 500 ms in 1-ms increments.

Note For 32-Gbps modules, trunking (TE, TF, and TNP) ports will use the core timeout value and non-trunking ports (F, E, and NP) or edge ports will use the edge timeout value.

(Optional) Revert to the default slow-port monitor timeout value (50 ms) for the specified port type:

switch(config)# **system timeout slowport-monitor default logical-type** {**core | edge**}

(Optional) Disable the slow-port monitor:

switch(config)# **no system timeout slowport-monitor default logical-type** {**core | edge**}

Configuring Slow Port Monitor for Port Monitor

Slow port monitor can be configured in port monitor via the tx-slowport-oper-delay counter. The **system timeout slowport-monitor** command also must be configured with a value that is less than or equal to the tx-slowport-oper-delay rising threshold. The port monitor logical type must also match the **system timeout slowport-monitor logical-type** command. Failure to do this results in no port monitor alerts being generated for tx-slowport-oper-delay.

Configuring the Transmit Average Credit-Not-Available Duration Threshold and Action in Port Monitor

Cisco MDS monitors its ports that are at zero transmit credits for 100 ms or more. This is called transmit average credit-not-available duration. The Port Monitor feature can monitor this using the TX Credit Not Available counter. When the transmit average credit-not-available duration exceeds the threshold set in the port monitor policy, an SNMP trap with interface details is sent, indicating the transmit average credit not available duration event along with a syslog message. Additionally, the following events may be configured:

- A warning message is displayed.
- The port can be error disabled.

- The port can be flapped.

The Port Monitor feature provides the CLI to configure the thresholds and actions. The threshold configuration is configured as a percentage of the interval. The thresholds can be 0 to 100 percent in multiples of 10, and the interval can be 1 second to 1 hour. The default is 10 percent of a 1-second interval and generates a SNMP trap and syslog message when the transmit-average-credit-not-available duration hits 100 ms.

The following edge port monitor policy is active by default. No port monitor policy is enabled for core ports by default.

```
switch# show port-monitor slowdrain

Policy Name  : slowdrain
Admin status : Not Active
Oper status  : Not Active
Port type    : All Access Ports
--------------------------------------------------------------------------------
Counter          Threshold Interval Rising    event Falling   event PMON Portguard
                                    Threshold       Threshold
--------------------------------------------------------------------------------
Credit Loss Reco  Delta      1          1        4       0       4    Not enabled
TX Credit Not
Available         Delta      1          10%      4       0%      4    Not enabled
```

The following example shows how to configure a new policy similar to the slowdrain policy with the tx-credit not available threshold set to 200 ms:

Note The default *slowdrain* port monitor policy cannot be modified and a new policy needs to be configured.

```
switch# configure
switch(config)# port-monitor name slowdrain_tx200ms
switch(config-port-monitor)# logical-type edge
switch(config-port-monitor)# no monitor counter all
switch(config-port-monitor)# monitor counter credit-loss-reco
switch(config-port-monitor)# monitor counter tx-credit-not-available
switch(config-port-monitor)# counter tx-credit-not-available poll-interval 1 delta
rising-threshold 20 event 4 falling-threshold 0 event 4
switch(config-port-monitor)# no port-monitor activate slowdrain
switch(config)# port-monitor activate slowdrain_tx200ms
switch(config)# end

switch# show port-monitor active
Policy Name  : slowdrain_tx200ms
Admin status : Active
Oper status  : Active
Port type    : All Edge Ports
--------------------------------------------------------------------------------------------------------
Counter Threshold  Interval     Rising event     Falling   event  Warning PMON
       Threshold  Threshold   Threshold     Portguard
--------------------------------------------------------------------------------------------------------
Credit Loss
Reco    Delta  1  1    4    0   4    Not enabled   Not enabled
TX Credit
Not Available Delta        1       20%   4    0%  4    Not enabled   Not enabled
```

Configuring Other Congestion Related Port Monitor Counters

The following port-monitor counters related to SAN congestion can be configured:

Table 21: Port-Monitor Counters

Counter Name	Description
invalid-words	Represents the total number of invalid words received by a port.
link-loss	Represents the total number of link failures encountered by a port.
lr-rx	Represents the total number link reset primitive sequence received by a port.
lr-tx	Represents the total number of link reset primitive sequence transmitted by a port.
rx-datarate	Receives frame rate in bytes per seconds.
signal-loss	Represents the number of times a port encountered laser or signal loss.
state-change	Represents the number of times a port has transitioned to an operational up state.
sync-loss	Represents the number of times a port experienced loss of synchronization in RX.
tx-credit-not-available	Increments by one if there is no transmit buffer-to-buffer credits available for a duration of 100 ms.
timeout-discards	Represents the total number of frames dropped at egress due to congestion timeout or no-credit-drop timeout.
tx-datarate	Represents the transmit frame rate in bytes per seconds.
tx-discards	Represents the total number of frames dropped at egress due to timeout, abort, offline, and so on.
tx-slowport-count	Represents the number of times slow port events were detected by a port for the configured slowport-monitor timeout. This is applicable only for generation 3 modules.
tx-slowport-oper-delay	Captures average credit delay (or R_RDY delay) experienced by a port. The value is in milliseconds.

Configuring Congestion Avoidance

The following features can be configured for congestion avoidance:

- Congestion-drop
- No-credit-drop
- Pause-drop
- Port-monitor portguard action for congestion avoidance (errdisable and flap)

Configuring the Congestion Drop Timeout Value for FCoE

When an FCoE frame takes longer than the congestion drop timeout period to be transmitted by the egress port, the frame is dropped. This dropping of frames is useful in controlling the effect of slow egress ports that are paused almost continuously (long enough to cause congestion), but not long enough to trigger the pause timeout drop. Frames dropped due to the congestion drop threshold are counted as egress discards against the egress port. Egress discards release buffers in the upstream ingress ports of a switch, allowing the unrelated flows to move continuously through the ports.

The congestion drop timeout value is 500 ms by default for all port types. We recommend that you retain the default timeout for core ports, and consider configuring a lower value for edge ports. The congestion drop timeout value should be equal to or greater than the pause drop timeout value for that port type.

To configure the congestion drop timeout value for FCoE, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Depending on the Cisco MDS NX-OS release version you are using, use one of the following commands to configure the system-wide FCoE congestion drop timeout, in milliseconds, for core or edge ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

  switch(config)# **system default interface congestion timeout** *milliseconds* **mode {core | edge}**

  The FCoE congestion drop timeout range is from 100 to 1000 ms.

  Note To prevent premature packet drops, the minimum value recommended for FCoE congestion drop timeout is 200 ms.

- Cisco MDS NX-OS Release 8.2(1) and later releases

  switch(config)# **system timeout fcoe congestion-drop** {*milliseconds* | **default**} **mode {core | edge}**

  The FCoE congestion drop timeout range is from 200 to 500 ms.

  Note In Cisco MDS NX-OS Release 8.1(1) and earlier releases, the FCoE congestion drop timeout value could be configured to as low as 100 ms. However, under certain circumstances configuring a congestion drop timeout value of 100 ms led to premature packet drops. In Cisco MDS NX-OS 8.2(1) and later releases, the minimum congestion drop timeout value was set to 200 ms to prevent premature packet drops. Therefore, we do not recommended that you specify a congestion drop timeout value of less than 200 ms in Cisco MDS NX-OS Release 8.1(1) and earlier releases.

(Optional) Depending on the Cisco MDS NX-OS release version you are using, use one of the following commands to revert to the default FCoE congestion drop timeout value of 500 milliseconds:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

  switch(config)# **no system default interface congestion timeout** *milliseconds* **mode {core | edge}**

- Cisco MDS NX-OS Release 8.2(1) and later releases

  switch(config)# **no system timeout fcoe congestion-drop** {*milliseconds* | **default**} **mode {core | edge}**

Configuring Pause Drop Timeout for FCoE

When an FCoE port is in a state of continuous pause during the FCoE pause drop timeout period, all the frames that are queued to that port are dropped immediately. As long as the port continues to remain in the pause state, the newly arriving frames destined for the port are dropped immediately. These drops are counted as egress discards on the egress port, and free up buffers in the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

To reduce the effect of a slow-drain device on unrelated traffic flows, configure a lower-pause drop timeout value than the congestion frame timeout value, for edge ports. This causes the frames that are destined for a slow port to be dropped immediately after the FCoE pause drop timeout period has occurred, rather than waiting for the congestion timeout period to drop them.

By default, the FCoE pause drop timeout is enabled on all ports and the value is set to 500 ms. We recommend that you retain the default timeout core ports and consider configuring a lower value for edge ports.

To configure the FCoE pause drop timeout value, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to configure the system-wide FCoE pause drop timeout value, in milliseconds, for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

  switch(config)# **system default interface pause timeout** *milliseconds* **mode {core | edge}**

- Cisco MDS NX-OS Release 8.2(1) and later releases

  switch(config)# **system timeout fcoe pause-drop** {*milliseconds* | **default**} **mode {core | edge}**

The range is from 100 to 500 milliseconds.

(Optional) Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to enable the FCoE pause drop timeout to the default value of 500 milliseconds for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

  switch(config)# **system default interface pause mode {core | edge}**

- Cisco MDS NX-OS Release 8.2(1) and later releases

  switch(config)# **system timeout fcoe pause-drop default mode {core | edge}**

(Optional) Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to disable the FCoE pause drop timeout for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

  switch(config)# **no system default interface pause mode {core | edge}**

- Cisco MDS NX-OS Release 8.2(1) and later releases

  switch(config)# **no system timeout fcoe pause-drop default mode {core | edge}**

Configuring the Congestion Drop Timeout Value for Fibre Channel

When a Fibre Channel frame takes longer than the congestion timeout period to be transmitted by the egress port, the frame is dropped. This option of the frames being dropped is useful for controlling the effect of slow egress ports that lack transmit credits almost continuously; long enough to cause congestion, but not long enough to trigger the no-credit timeout drop. These drops are counted as egress discards on the egress port, and release buffers into the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

By default, the congestion timeout value is 500 ms for all port types. We recommend that you retain the default timeout for core ports and configure a lower value (not less than 200 ms) for edge ports. The congestion timeout value should be equal to or greater than the no-credit frame timeout value for that port type.

To configure the congestion frame timeout value for the Fibre Channel, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Configure the Fibre Channel congestion drop timeout value, in milliseconds, for the specified port type:

switch(config)# **system timeout congestion-drop** *milliseconds* **logical-type** {**core | edge**}

The range is 200-500 ms in multiples of 10.

Step 3 (Optional) Revert to the default value for the congestion timeout for the specified port type:

switch(config)# **no system timeout congestion-drop default logical-type** {**core | edge**}

Configuring the No-Credit Frame Timeout Value for Fibre Channel

When a Fibre Channel egress port has no transmit credits continuously for the no-credit timeout period, all the frames that are already queued to that port are dropped immediately. As long as the port remains in this condition, newly arriving frames destined for that port are dropped immediately. These drops are counted as egress discards on the egress port, and release buffers in the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

No-credit dropping can be enabled or disabled. By default, frame dropping is disabled and the frame timeout value is 500 ms for all port types. We recommend that you retain the default frame timeout for core ports and configure a lower value (300 ms) for edge ports. If the slow-drain events continue to affect unrelated traffic flows, the frame timeout value for the edge ports can be lowered to drop the previous slow-drain frames. This

frees the ingress buffers for frames of unrelated flows, thus reducing the latency of the frames through the switch.

Note

- The no-credit frame timeout value should always be less than the congestion frame timeout for the same port type, and the edge port frame timeout values should always be lower than the core port frame timeout values.
- The slow-port monitor delay value should always be less than the no-credit frame timeout value for the same port type.

On 16-Gbps and later modules and systems, the no-credit timeout value can be 1 to 500 ms in multiples of 1 ms. Dropping starts immediately after the no-credit condition comes into existence for the configured timeout value.

To configure the no-credit timeout value, perform these steps:

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the no-credit timeout value:

switch(config)# **system timeout no-credit-drop** *milliseconds* **logical-type edge**

(Optional) Revert to the default no-credit timeout value (500 ms):

switch(config)# **system timeout no-credit-drop default logical-type edge**

(Optional) Disable the no-credit drop timeout value:

switch(config)# **no system timeout no-credit-drop logical-type edge**

Displaying Credit Loss Recovery Actions

When a port is at zero transmit credits for 1 full second (F ports) and 1.5 seconds (E ports), it is called a credit loss condition. Cisco MDS initiates credit loss recovery by transmitting a Link Credit Reset (LCR). If the end device responds with a Link Credit Reset Response (LCRR), the port is back at its fully agreed number of B2B credits in both directions. If an LRR is not received, the port is shut down.

When the port detects the credit loss condition and recovers, some of the following actions might occur:

1. An SNMP trap with interface details can be sent, indicating the credit loss event.
2. The port can be error disabled.
3. The port can be flapped.

When the configured threshold is exceeded, one or more of these actions can be combined together. These actions can be turned on or off depending on the situation. The Port Monitor feature provides the CLI to configure the thresholds and action.

The 1 second (F ports) and 1.5 seconds (E ports) timers that are set for the switch to initiate CLR are fixed and cannot be changed.

To verify a port monitor policy to generate SNMP alerts and take other actions in the quantity and timing of these events, perform these steps:

Step 1 Display the last 10 credit loss events per interface per module:

switch# **show process creditmon credit-loss-events** [**module x**]

Step 2 Display a chronological log of credit loss events per module:

switch# **show process creditmon credit-loss-event-history** [**module x**]

Note When a port sees the credit loss condition and fails to recover, the port flaps. This function is already a part of the portguard, and you can configure the supported actions using the Portguard feature.

Configuring Congestion Isolation

The Congestion Isolation feature allows slow devices to be put into their own virtual link automatically as the port monitor detects the slow-drain condition.

The following port-monitor counters are used to detect slow drain and isolate the devices on an interface.

- credit-loss-reco
- tx-credit-not-available
- tx-slowport-oper-delay
- txwait

Configure the Slow-Drain Device Detection and Congestion Isolation feature in the following sequence:

1. Configure the Extended Receiver Ready feature. For more information, see Enabling Extended Receiver Ready, on page 160.
2. Configure the Congestion Isolation feature. For more information, see Enabling Congestion Isolation, on page 162.
3. Configure a port-monitor policy with one or more counters containing the portguard action *cong-isolate*. For more information, see Configuring the Port-Monitor Portguard Action for Congestion Isolation, on page 163.

Configuring Extended Receiver Ready

Enabling Extended Receiver Ready

To enable Extended Receiver Ready (ER_RDY) on a switch, perform these steps:

Before you begin

You must enable ER_RDY flow-control mode using the **system fc flow-control er_rdy** command on the local and adjacent switches, and then flap the ISLs connecting the local and adjacent switches to enable ER_RDY flow-control mode on the ISLs.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable ER_RDY flow-control mode:

switch(config)# **system fc flow-control er_rdy**

Note Enable the ER_RDY flow-control mode on both the connected switches for an existing Inter-Switch Link (ISL) before proceeding to step 3.

Step 3 Select a Fibre Channel interface and enter interface configuration submode:

switch(config-if)# **interface fc** *slot/port*

Step 4 Gracefully shut down the interface and administratively disable traffic flow:

switch(config-if)# **shutdown**

Step 5 Enable traffic flow on the interface:

switch(config-if)# **no shutdown**

Step 6 Return to privileged executive mode:

switch(config-if)# **end**

Step 7 Verify if the link is in ER_RDY flow-control mode:

switch# **show flow-control er_rdy**

Disabling Extended Receiver Ready

To disable Extended Receiver Ready (ER_RDY) on a switch, perform these steps:

Before you begin

1. Remove the congestion-isolation portguard action for the links in the port-monitor policy. For more information, see Configuring the Port-Monitor Portguard Action for Congestion Isolation, on page 163.
2. Disable the Congestion Isolation feature. For more information, see Enabling Congestion Isolation, on page 162.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Disable ER_RDY flow-control mode:

switch(config)# **no system fc flow-control**

Step 3 Select a Fibre Channel interface and enter interface configuration submode:

switch(config-if)# **interface fc** *slot/port*

Step 4 Gracefully shut down the interface and administratively disable traffic flow:

switch(config-if)# **shutdown**

Step 5 Enable traffic flow on the interface:

switch(config-if)# **no shutdown**

Step 6 Return to privileged executive mode:

switch(config-if)# **end**

Step 7 Verify if the link is in R_RDY flow-control mode:

switch# **show flow-control r_rdy**

Configuring Congestion Isolation

Enabling Congestion Isolation

To enable Congestion Isolation, perform these steps:

Before you begin

Configure Extended Receiver Ready. For more information, see Enabling Extended Receiver Ready, on page 160.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable Congestion Isolation:

switch(config)# **feature congestion-isolation**

(Optional) Manually include or exclude a device to be detected as a slow-drain device:

switch(config)# **congestion-isolation** {**include** | **exclude**} **pwwn** *pwwn* **vsan** *vsan-id*

Disabling Congestion Isolation

To disable Congestion Isolation, perform these steps:

Before you begin

Remove the congestion-isolation portguard action for links in the port-monitor policy. For more information, see Configuring the Port-Monitor Portguard Action for Congestion Isolation, on page 163.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Clear the slow-drain devices that were manually added to be detected as slow devices or excluded from being detected as slow devices:

switch(config)# **no congestion-isolation** {**include** | **exclude**} **pwwn** *pwwn* **vsan** *vsan-id*

Step 3 Disable Congestion Isolation:

switch(config)# **no feature congestion-isolation**

Removing an Interface

Port monitor detects slow devices when a given threshold is reached and triggers the congestion isolation feature in the switch to move traffic to that slow device into the slow Virtual Link (VL2). The switch does not automatically remove any devices from congestion isolation. This must be done manually once the problem with the slow device is identified and resolved.

To manually remove an interface from being detected as slow, perform these steps:

Remove an interface from being detected as slow by the port monitor:

switch#: **congestion-isolation remove interface** *slot*/*port*

Configuring the Port-Monitor Portguard Action for Congestion Isolation

To configure a port-monitor portguard action for Congestion Isolation, perform these steps:

Before you begin

1. Configure Extended Receiver Ready. For more information, see Enabling Extended Receiver Ready, on page 160.
2. Configure Congestion Isolation. For more information, see Enabling Congestion Isolation, on page 162.

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Specify the policy name and enter port monitoring policy configuration mode:

switch(config)# **port-monitor name** *policyname*

Step 3 Specify the counter parameters for the portguard to take congestion isolation action on a port:

switch(config-port-monitor)# **counter** {**credit-loss-reco** | **tx-credit-not-available** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *event-id* **warning-threshold** *count2* **falling-threshold** *count3* **event** *event-id* **portguard cong-isolate**

Note Absolute counters do not support portguard actions. However, the tx-slowport-oper-delay absolute counter supports Congestion Isolation portguard action.

(Optional) Revert to the default values for a counter:

switch(config-port-monitor)# **no counter** {**credit-loss-reco** | **tx-credit-not-available** | **tx-slowport-oper-delay** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *event-id* **warning-threshold** *count2* **falling-threshold** *count3* **event** *event-id* **portguard cong-isolate**

Step 4 Return to configuration mode:

switch(config-port-monitor)# **exit**

Step 5 Activate the specified port-monitor policy:

switch(config)# **port-monitor activate** *policyname*

(Optional) Deactivate the specified port-monitoring policy:

switch(config)# **no port-monitor activate** *policyname*

Verifying Slow-Drain Device Detection and Congestion Isolation

Table 22: Verifying Slow-Drain Device Detection and Congestion Isolation

Command	Purpose
show congestion-isolation {**exclude-list** | **global-list** | **ifindex-list** | **include-list** | **pmon-list** | **remote–list** | **status**}	Displays information about the devices configured for Congestion Isolation.
show flow-control {**er_rdy** | **r_rdy**} **module***number*	Displays the interfaces that are in ER_RDY flow-control mode. **Note** If the module number is not specified, the command will display the output for all the modules.
show system fc flow-control	Displays Fibre Channel flow-control mode of a switch.

Configuration Examples for Congestion Detection, Avoidance, and Isolation

Configuration Examples for Congestion Detection

This example shows how to configure the FCoE congestion drop timeout to the default value of 500 milliseconds for a core port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface congestion timeout 500 mode core
```

This example shows how to configure the FCoE congestion drop timeout to the default value of 500 milliseconds for a core port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe congestion-drop default mode core
```

This example shows how to configure the FCoE congestion drop timeout to the default value of 500 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface congestion timeout 500 mode edge
```

This example shows how to configure the FCoE congestion drop timeout to the default value of 500 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe congestion-drop default mode edge
```

This example shows how to configure the FCoE congestion drop timeout to the value of 200 milliseconds for a core port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface congestion timeout 200 mode core
```

This example shows how to configure the FCoE congestion drop timeout to the value of 200 milliseconds for a core port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe congestion-drop 200 mode core
```

This example shows how to configure the FCoE congestion drop timeout to the value of 200 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
```

```
switch(config)# system default interface congestion timeout 200 mode edge
```

This example shows how to configure the FCoE congestion drop timeout to the value of 200 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe congestion-drop 200 mode edge
```

This example shows how to configure the FCoE pause drop timeout value of 100 milliseconds for a core port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface pause timeout 100 mode core
```

This example shows how to configure the FCoE pause drop timeout value of 200 milliseconds for a core port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe pause-drop 200 mode core
```

This example shows how to configure the FCoE pause drop timeout value of 100 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface pause timeout 100 mode edge
```

This example shows how to configure the FCoE pause drop timeout value of 200 milliseconds for a edge port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe pause-drop 200 mode edge
```

This example shows how to configure the FCoE pause drop timeout to the default of 500 milliseconds for the core port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface pause mode core
```

This example shows how to configure the FCoE pause drop timeout to the default of 500 milliseconds for the core port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe pause-drop default mode core
```

This example shows how to configure the FCoE pause drop timeout to the default of 500 milliseconds for the edge port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# system default interface pause mode edge
```

This example shows how to configure the FCoE pause drop timeout to the default value of 500 milliseconds for an edge port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# system timeout fcoe pause-drop default mode edge
```

This example shows how to disable the FCoE pause drop timeout for a core port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# no system default interface pause mode core
```

This example shows how to disable the FCoE pause drop timeout for a core port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# no system timeout fcoe pause-drop default mode core
```

This example shows how to disable the FCoE pause drop timeout for an edge port type in Cisco MDS NX-OS Release 8.1(1) and earlier releases:

```
switch# configure terminal
switch(config)# no system default interface pause mode edge
```

This example shows how to disable the FCoE pause drop timeout for an edge port type in Cisco MDS NX-OS Release 8.2(1) and later releases:

```
switch# configure terminal
switch(config)# no system timeout fcoe pause-drop default mode edge
```

Configuration Examples for Congestion Avoidance

Note

- From Cisco MDS NX-OS Release 8.1(1), mode E is treated as logical-type core and mode F is treated as logical-type edge.
- The port *Logical type* is displayed as the *Port type*.

This example shows hot to check the currently active port-monitor policy:

```
switch# show port-monitor active
Policy Name  : sample
Admin status : Active
Oper status  : Active
Port type    : All Ports
--------------------------------------------------------------------------------
Counter     Threshold  Interval Rising     event Falling   event Warning      PMON
                                Threshold        Threshold       Threshold    Portguard
--------------------------------------------------------------------------------
```

```
Link
Loss          Delta      10        6          4     5          4   Not enabled  Flap
Sync
Loss          Delta      60        5          4     1          4   Not enabled  Not enabled
Signal
Loss          Delta      60        5          4     1          4   Not enabled   Not enabled
Invalid
Words         Delta      60        1          4     0          4   Not enabled   Not enabled
Invalid
CRC's         Delta      30        20         2     10         2   Not enabled   Not enabled
State
Change        Delta      60        5          4     0          4   Not enabled   Not enabled
TX
Discards      Delta      60        200        4     10         4   Not enabled   Not enabled
LR RX         Delta      60        5          4     1          4   Not enabled   Not enabled
LR TX         Delta      60        5          4     1          4   Not enabled   Not enabled
Timeout
Discards      Delta      60        200        4     10         4   Not enabled   Not enabled
Credit
Loss Reco     Delta      1         1          4     0          4   Not enabled   Not enabled
TX Credit
Not Available Delta      3         40%        4     2%         4   Not enabled   Not enabled
RX Datarate   Delta      60        80%        4     20%        4   Not enabled   Not enabled
TX Datarate   Delta      60        80%        4     20%        4   Not enabled   Not enabled
ASIC Error
Pkt to xbar   Delta      300       5          4     0          4   Not enabled   Not enabled
-------------------------------------------------------------------------------------------
```

This example shows how to configure the Fibre Channel congestion drop timeout value of 210 milliseconds for logical type core:

```
switch# configure terminal
switch(config)# system timeout congestion-drop 210 logical-type core
```

This example shows how to configure the Fibre Channel congestion drop timeout to the default value of 200 milliseconds for logical type core:

```
switch# configure terminal
switch(config)# system timeout congestion-drop default logical-type core
```

This example shows how to configure the Fibre Channel no-credit drop timeout value of 100 milliseconds for logical type edge:

```
switch# configure terminal
switch(config)# system timeout no-credit-drop 100 logical-type edge
```

This example shows how to configure the Fibre Channel no-credit drop timeout to the default value of 500 milliseconds for logical type edge:

Note The no-credit drop timeout value is disabled by default.

```
switch# configure terminal
switch(config)# system timeout no-credit-drop default logical-type edge
```

This example shows how to disable the Fibre Channel no-credit drop timeout for logical type edge when it is enabled:

```
switch# configure terminal
switch(config)# no system timeout no-credit-drop logical-type edge
```

This example shows how to configure the Fibre Channel hardware slowport monitoring value of 10 milliseconds for logical type edge:

```
switch# configure terminal
switch(config)# system timeout slowport-monitor 10 logical-type edge
```

This example shows how to configure the Fibre Channel hardware slowport monitoring to the default value of 50 milliseconds for logical type edge:

Note The slowport monitoring value is disabled by default.

```
switch# configure terminal
switch(config)# system timeout slowport-monitor default logical-type edge
```

This example shows how to disable the Fibre Channel hardware slowport monitoring for logical type edge when it is enabled:

```
switch# configure terminal
switch(config)# no system timeout slowport-monitor logical-type edge
```

Configuration Examples for Congestion Isolation

This example shows how to enable ER_RDY flow-control mode:

```
switch# configure terminal
switch(config)# system fc flow-control er_rdy
Flap the ISLs to activate ER_RDY mode on E ports.Use the CLI show flow-control r_rdy to
list the ports that are still in R_RDY mode
```

This example shows how to disable ER_RDY flow-control mode:

Note You need to disable the Congestion Isolation feature before disabling the ER_RDY flow-control mode.

```
switch# configure terminal
switch(config)# no feature congestion-isolation
switch(config)# no system fc flow-control
```

This example shows how to enable Congestion Isolation:

```
switch# configure terminal
switch(config)# feature congestion-isolation
Flap the ISLs to activate ER_RDY mode on E ports. Use the CLI show flow-control r_rdy to
list the ports that are still in R_RDY mode
```

This example shows how to disable Congestion Isolation:

```
switch# configure terminal
switch(config)# no feature congestion-isolation
Flap the ISLs to activate ER_RDY mode on E ports.Use the CLI show flow-control r_rdy to
list the ports that are still in R_RDY mode
```

This example shows how to manually configure a device as a slow device. The configured device will be permanently treated as a slow device until it is removed from congestion isolation using the **no congestion-isolation include pwwn** *pwwn* **vsan vsan-id** command. All traffic to this device traversing the device's ISLs that are in ER_RDY flow-control mode will be routed to the low-priority VL (VL2).

```
switch# configure terminal
switch(config)# congestion-isolation include pwwn 10:00:00:00:c9:f9:16:8d vsan 4
```

This example shows how to configure a device that is to be excluded from automatic congestion isolation by the port monitor. Even when the rising threshold of a port-monitor counter is reached and the portguard action is set to cong-isolate, this device will not be isolated as a slow device, and traffic to this device traversing the device's ISLs that are in ER_RDY flow-control mode will not be routed to the low-priority VL (VL2).

```
switch# configure terminal
switch(config)# congestion-isolation exclude pwwn 10:00:00:00:c9:f9:16:8d vsan 4
```

Slow devices can be identified either via the port monitor, or manually included or excluded using the **congestion-isolation** {**include** | **exclude**} **pwwn** *pwwn* **vsan** *vsan-id* command. On removing the exclude configuration, if the device is detected as slow by the port monitor, the device will again be marked as slow. Also, if the exclude configuration is already used for a device marked as slow by the port monitor, the device will no longer behave as a slow device.

This example shows how to manually remove an interface from being detected as slow in the port monitor:

1. Identifying the interface that you want to remove from being detected as slow.

```
switch# show congestion-isolation ifindex-list
======================
Ifindex: 1088000(fc2/9)                <<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<< interface
 fc2/9 marked slow
```

2. Identifying the host that is using the interface.

```
switch# show congestion-isolation pmon-list

PMON detected list for vsan 1    : PWWN(FCID)
===============================================

PMON detected list for vsan 2    : PWWN(FCID)
===============================================
```

```
PMON detected list for vsan 3    : PWWN(FCID)
================================================
21:00:00:24:ff:4f:70:46(0x040020)          <<<<<<<<<<<<<<<<<<<<<<host behind interface
 fc2/9 marked slow

PMON detected list for vsan 4    : PWWN(FCID)
================================================

PMON detected list for vsan 5    : PWWN(FCID)
================================================
```

3. Remove the interface from being marked as slow.

```
switch# congestion-isolation remove interface fc2/9    <<<<<<<<<<< CLI to remove an
interface from being marked as slow by PMON
```

4. Verifying if the interface is removed from being detected as slow.

```
switch# show congestion-isolation pmon-list

PMON detected list for vsan 1    : PWWN(FCID)
================================================

PMON detected list for vsan 2    : PWWN(FCID)
================================================

PMON detected list for vsan 3    : PWWN(FCID)
================================================

<<<<<<<<<<<<<< host behind interface fc2/9 removed from isolation
PMON detected list for vsan 4    : PWWN(FCID)
================================================

PMON detected list for vsan 5    : PWWN(FCID)
================================================
```

Verifying Congestion Detection, Avoidance, and Isolation

Verifying Congestion Detection and Avoidance

The following commands display slow-port monitor events:

Note These commands are applicable for both supervisor and module prompts.

Display slow-port monitor events per module:

switch# **show process creditmon slowport-monitor-events [module x [port y]]**

Display the slow-port monitor events on the Onboard Failure Logging (OBFL):

switch# **show logging onboard slowport-monitor-events**

Note The slow-port monitor events are logged periodically into the OBFL.

The following example displays the credit monitor or output of the **creditmon slow-port monitor-events** command for the 16-Gbps and 32-Gbps modules and switches:

```
switch# show process creditmon slowport-monitor-events

        Module: 06      Slowport Detected: YES
================================================================================
Interface = fc6/3
--------------------------------------------------------------------
| admin  | slowport  | oper  |            Timestamp            |
| delay  | detection | delay |                                 |
| (ms)   | count     | (ms)  |                                 |
--------------------------------------------------------------------
|   1    |     46195 |    1  | 1. 10/14/12 21:46:51.615        |
|   1    |     46193 |   50  | 2. 10/14/12 21:46:51.515        |
|   1    |     46191 |   50  | 3. 10/14/12 21:46:51.415        |
|   1    |     46189 |   50  | 4. 10/14/12 21:46:51.315        |
|   1    |     46187 |   50  | 5. 10/14/12 21:46:51.215        |
|   1    |     46185 |   50  | 6. 10/14/12 21:46:51.115        |
|   1    |     46183 |   50  | 7. 10/14/12 21:46:51.015        |
|   1    |     46181 |   50  | 8. 10/14/12 21:46:50.915        |
|   1    |     46179 |   50  | 9. 10/14/12 21:46:50.815        |
|   1    |     46178 |   50  |10. 10/14/12 21:46:50.715        |
--------------------------------------------------------------------
```

Note For 16-Gbps modules, 32-Gbps modules, and Cisco MDS 9700, 9148S, 9250i, and 9396S switches, if **no-credit-drop** timeout is configured, the maximum value of **tx-slowport-oper-delay** as shown in slow-port monitor events is limited by the **no-credit-drop timeout**. So, the maximum value for **tx-slowport-oper-delay** can reach the level of the **no-credit-drop** timeout even if the actual slow-port delay from the device is higher because the frames are forcefully dropped by the hardware when **tx-slowport-oper-delay** reaches the level of the **no-credit-drop** timeout.

Verifying Congestion Isolation

Note TxWait on FCoE ethernet or Virtual Fibre Channels (VFC) interfaces is the amount of time a port cannot transmit because of the received Priority Flow Control (PFC) pause frames.

RxWait on FCoE ethernet or VFCs is the amount of time a port cannot receive because of the port transmitting PFC pause frames.

Both TxWait and RxWait are in units of 2.5us and are converted to seconds in some command outputs. To convert to seconds multiple the TxWait or RxWait value by 2.5 and divide by 1,000,000.

This example displays the status and statistics of priority-flow-control on all interfaces:

```
switch# show interface priority-flow-control
RxPause: No. of pause frames received
TxPause: No. of pause frames transmitted
TxWait: Time in 2.5uSec a link is not transmitting data[received pause]
RxWait: Time in 2.5uSec a link is not receiving data[transmitted pause]
================================================================================
Interface          Admin Oper (VL bmap) VL  RxPause  TxPause RxWait-     TxWait-
                                                             2.5us(sec)  2.5us(sec)
================================================================================

port-channel1      Auto NA       (8)     3    0        0        0(0)      0(0)
port-channel350    Auto NA       (8)     3    0        0        0(0)      0(0)
port-channel351    Auto NA       (8)     3    0        0        0(0)      0(0)
port-channel552    Auto NA       (8)     3    111506   0        0(0)      5014944(12)
port-channel700    Auto NA       (8)     3    0        0        0(0)      0(0)
Ethernet2/17       Auto Off
Ethernet2/18       Auto Off
Ethernet2/19       Auto Off
Ethernet2/20       Auto Off
Ethernet2/25       Auto On       (8)     3    0        0        0(0)      0(0)
Ethernet2/26       Auto On       (8)     3    0        0        0(0)      0(0)
```

This example displays the detailed configuration and statistics of a specified virtual Fibre Channel interface:

```
switch# show interface vfc 9/11 counters details
vfc9/11
    3108091433 fcoe in packets
    6564116595616 fcoe in octets
    30676987 fcoe out packets
    2553913687 fcoe out octets
    0 2.5us TxWait due to pause frames (VL3)
```

```
    134795 2.5us RxWait due to pause frames (VL3)
    0 Tx frames with pause opcode (VL3)
    0 Rx frames with pause opcode (VL3)
Percentage pause in TxWait per VL3  for last 1s/1m/1h/72h: 0%/0%/0%/0%
Percentage pause in RxWait per VL3  for last  1s/1m/1h/72h: 0%/0%/0%/0%
```

This example displays the transmit wait history information for Ethernet2/16:

```
switch# show hardware internal txwait-history module 2 port 16
TxWait history for port Eth2/16:
  ==============================
       6666666666666666666667666666666666666666666666666666666666666666
       7667777777777777776671777676677776767777777776777676768778676
       2853241601145111129728012626411014272052114264119551532106 17
  1000
   900
   800
   700 ############################################################
   600 ############################################################
   500 ############################################################
   400 ############################################################
   300 ############################################################
   200 ############################################################
   100 ############################################################
      0....5....1....1....2....2....3....3....4....4....5....5....6
                0    5    0    5    0    5    0    5    0    5    0

            TxWait per second (last 60 seconds)
                   # = TxWait (ms)

     44444444444
     00000000006000000000000000000000000000000000000000000000000000
     ............................................................
     33433332333000000000000000000000000000000000000000000000000000
  60
  54
  48
  42 ###########
  36 ###########
  30 ###########
  24 ###########
  18 ###########
  12 ###########
   6 ############
    0....5....1....1....2....2....3....3....4....4....5....5....6
              0    5    0    5    0    5    0    5    0    5    0

            TxWait per minute (last 60 minutes)
                   # = TxWait (secs)



     000000000000000000000000000000000000000000000000000000000000000000000000000000000000
  3600
  3240
  2880
  2520
  2160
  1800
  1440
  1080
   720
```

```
 360
    0....5....1....1....2....2....3....3....4....4....5....5....6....6....7.7
              0    5    0    5    0    5    0    5    0    5    0    5    0 2

    TxWait per hour (last 72 hours)
                # = TxWait (secs)
```

This example displays the receive wait history information for Ethernet10/41:

```
switch# show hardware internal rxwait-history module 10 port 41
RxWait history for port Eth10/41:
  ==============================
       000000000000000000000000000000000000000000000000000000000000
  1000
   900
   800
   700 ############################################################
   600 ############################################################
   500 ############################################################
   400 ############################################################
   300 ############################################################
   200 ############################################################
   100 ############################################################
      0....5....1....1....2....2....3....3....4....4....5....5....6
                0    5    0    5    0    5    0    5    0    5    0

             RxWait per second (last 60 seconds)
                    # = RxWait (ms)

     33333333331
     000000000000000079999999999999999200000000000000000000000000
     ............................................................
     000000000000000096666666666666663000000000000000000000000000
  60
  54
  48
  42 ############
  36 ############
  30 ############
  24 ############
  18 ############
  12 ############
   6 ############
    0....5....1....1....2....2....3....3....4....4....5....5....6
              0    5    0    5    0    5    0    5    0    5    0

             RxWait per minute (last 60 minutes)
                    # = RxWait (secs)



       000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
  3600
  3240
  2880
  2520
  2160
  1800
  1440
  1080
   720
   360
```

```
      0....5....1....1....2....2....3....3....4....4....5....5....6....6....7.7
                0    5    0    5    0    5    0    5    0    5    0    5    0 2

      RxWait per hour (last 72 hours)
                  # = RxWait (secs)
```

This example displays the onboard failure log(OBFL) for TxWait caused by receiving PFC pause frames:

```
module# show logging onboard txwait
---------------------------------
 Module: 2 txwait count
---------------------------------
----------------------------
 Show Clock
----------------------------
2017-09-22 06:22:17
Notes:
     - Sampling period is 20 seconds
     - Only txwait delta >= 100 ms are logged
-------------------------------------------------------------------------------------
|   Interface     | Delta TxWait Time     | Congestion | Timestamp                |
|                 | 2.5us ticks | seconds |            |                          |
-------------------------------------------------------------------------------------
|   Eth2/1(VL3)   |  2508936    |    6    |      31%   | Fri Sep 22 05:29:21 2017 |
|   Eth2/1(VL3)   |  3355580    |    8    |      41%   | Mon Sep 11 17:55:52 2017 |
|   Eth2/1(VL3)   |  8000000    |   20    |     100%   | Mon Sep 11 17:55:31 2017 |
|   Eth2/1(VL3)   |  8000000    |   20    |     100%   | Mon Sep 11 17:55:11 2017 |
|   Eth2/1(VL3)   |  8000000    |   20    |     100%   | Mon Sep 11 17:54:50 2017 |
```

This example displays the onboard failure log(OBFL) for RxWait caused by transmitting PFC pause frames:

```
module# show logging onboard rxwait
---------------------------------
 Module: 14 rxwait count
---------------------------------
----------------------------
 Show Clock
----------------------------
2017-09-22 11:53:53
Notes:
     - Sampling period is 20 seconds
     - Only rxwait delta >= 100 ms are logged
-------------------------------------------------------------------------------------
|   Interface       | Delta RxWait Time     | Congestion | Timestamp                |
|                   | 2.5us ticks | seconds |            |                          |
-------------------------------------------------------------------------------------
|   Eth14/21(VL3)   |  2860225    |    7    |      35%   | Thu Sep 21 23:59:46 2017 |
|   Eth14/30(VL3)   |    42989    |    0    |       0%   | Thu Sep 14 14:53:57 2017 |
|   Eth14/29(VL3)   |    45477    |    0    |       0%   | Thu Sep 14 14:47:56 2017 |
|   Eth14/30(VL3)   |    61216    |    0    |       0%   | Thu Sep 14 14:47:56 2017 |
|   Eth14/29(VL3)   |    43241    |    0    |       0%   | Thu Sep 14 14:47:36 2017 |
|   Eth14/30(VL3)   |    43845    |    0    |       0%   | Thu Sep 14 14:47:36 2017 |
|   Eth14/29(VL3)   |    79512    |    0    |       0%   | Thu Sep 14 14:47:16 2017 |
|   Eth14/30(VL3)   |    62529    |    0    |       0%   | Thu Sep 14 14:47:16 2017 |
|   Eth14/29(VL3)   |    50699    |    0    |       0%   | Thu Sep 14 14:45:56 2017 |
|   Eth14/30(VL3)   |    47839    |    0    |       0%   | Thu Sep 14 14:45:56 2017 |
```

This example displays the error statistics onboard failure log (OBFL) for a switch:

```
switch# show logging onboard error-stats
```

```
---------------------------
 Show Clock
---------------------------
2017-09-22 15:35:31


-------------------------------------------------------------------------------------------
 STATISTICS INFORMATION FOR DEVICE ID 166 DEVICE Clipper MAC
-------------------------------------------------------------------------------------------
 Port |                                        |               |    Time Stamp   |In
Range |            Error Stat Counter Name     |    Count      |MM/DD/YY HH:MM:SS|st
      |                                        |               |                 |Id
-------------------------------------------------------------------------------------------
11    |GD rx pause transitions of XOFF-XON VL3 |2147           |09/22/17 00:11:24|02
11    |GD uSecs VL3 is in internal pause rx state |7205308     |09/22/17 00:11:24|02
11    |GD rx frames with pause opcode for VL3  |6439           |09/22/17 00:11:24|02
11    |PL SW pause event (vl3)                 |113            |09/22/17 00:11:24|02
```

This example show how to verify system flow-control mode:

```
switch# show system fc flow-control
System flow control is ER_RDY
```

This example shows how to verify the Congestion Isolation status:

```
switch# show congestion-isolation status
Flow Control Mode     : ER_RDY
Congestion Isolation  : Enabled
Sampling Interval     : 1
Timeout               : 0
ESS Cap Details
-----------------------------------------
VSAN: 0x1(1)
Enabled domain-list: 0x4(4 - local)
Disabled domain-list: None
Unsupported domain-list: 0x61(97)
VSAN: 0x2(2)
Enabled domain-list: 0x4(4 - local)
Disabled domain-list: None
Unsupported domain-list: 0xb8(184)
VSAN: 0x3(3)
Enabled domain-list: 0x4(4 - local)
Disabled domain-list: None
Unsupported domain-list: None
VSAN: 0x4(4)
Enabled domain-list: 0x4(4 - local) 0xbb(187)
Disabled domain-list: None
Unsupported domain-list: None
```

This example shows how to verify the list of devices that were detected as slow on a local switch:

```
switch# show congestion-isolation pmon-list vsan 4
PMON detected list for vsan 4    : PWWN(FCID)
================================================
10:00:00:00:c9:f9:16:8d(0xbe0000)
```

This example shows how to verify the global list of devices that were detected as slow in a fabric when the Congestion Isolation feature was enabled. The global list should be the same on all switches in the fabric where the Congestion Isolation feature is enabled.

```
switch# show congestion-isolation global-list vsan 4
Global list for vsan 4 PWWN(FCID)
=========================================
10:00:00:00:c9:f9:16:8d(0xbe0000)
```

This example shows the list of devices that were detected as slow on remote switches (not locally detected slow devices):

```
switch# show congestion-isolation remote-list vsan 4
Remote list for vsan 4    : PWWN(FCID)
=================================================
10:00:00:00:c9:f9:16:8d(0xbe0000)
```

This example shows a single device that is marked as slow (feature slow-dev) either via the port monitor or the **congestion isolation include** command:

```
switch# show congestion-isolation include-list vsan 4
Include list for vsan 4   : PWWN(FCID) (online/offline)
=========================================================
10:00:00:00:c9:f9:16:8d(0xbe0000) - (Online)

switch# show fcns database vsan 4
VSAN 4:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x040000    N     10:00:40:55:39:0c:80:85 (Cisco)         ipfc
0x040020    N     21:00:00:24:ff:4f:70:47 (Qlogic)        scsi-fcp:target
0xbe0000    N     10:00:00:00:c9:f9:16:8d (Emulex)        scsi-fcp:init slow-dev <<<slow
device
[testing]Total number of entries = 3
```

This example shows the list of devices that were manually configured using Congestion Isolation exclude list command on a local switch:

```
switch# show congestion-isolation exclude-list vsan 4
Exclude list for vsan 4   : PWWN(FCID) (online/offline)
=========================================================
10:00:00:00:c9:f9:16:8d(0xbe0000) - (Online)
```

Configuring Trunking

This chapter provides information about trunking and how to configure the trunking.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Series Multilayer Switches. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports (see Figure 4: Trunking E Ports, on page 183 and Figure 5: Trunking F Ports, on page 184).

Trunking E Ports

Trunking the E ports enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

Figure 4: Trunking E Ports

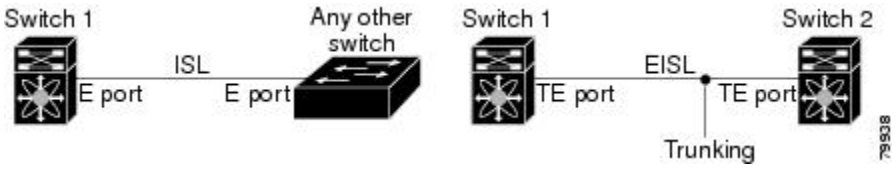

Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c_Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Trunking F Ports

Trunking F ports allows interconnected ports to transmit and receive tagged frames in more than one VSAN, over the same physical link.

Figure 5: Trunking F Ports, on page 184 represents the possible trunking scenarios in a SAN with MDS core switches, NPV switches, third-party core switches, and HBAs.

Figure 5: Trunking F Ports

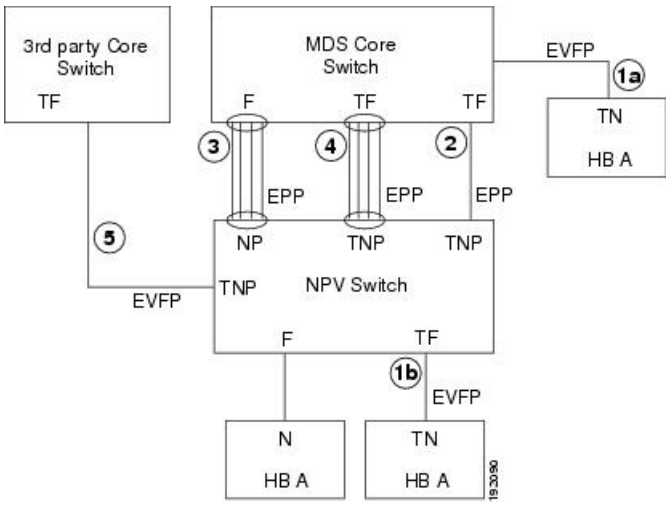

Link Number	Link Description
1a and 1b	F port trunk with N port.[6]
2	F port trunk with NP port.
3	F PortChannnel with NP port.
4	Trunked F PortChannel with NP port.
5	Trunking NP port with third-party core switch F port

[6] These features are not supported currently.

Key Concepts

The trunking feature includes the following key concepts:

- TE port—If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- TF port—If trunk mode is enabled in an F port (see the link 2 in Figure 5: Trunking F Ports, on page 184) and that port becomes operational as a trunking F port, it is referred to as a TF port.
- TN port—If trunk mode is enabled (not currently supported) in an N port (see the link 1b in Figure 5: Trunking F Ports, on page 184) and that port becomes operational as a trunking N port, it is referred to as a TN port.
- TNP port—If trunk mode is enabled in an NP port (see the link 2 in Figure 5: Trunking F Ports, on page 184) and that port becomes operational as a trunking NP port, it is referred to as a TNP port.
- TF PortChannel—If trunk mode is enabled in an F PortChannel> (see the link 4 in Figure 5: Trunking F Ports, on page 184) and that PortChannel becomes operational as a trunking F PortChannel, it is referred to as TF PortChannel. Cisco Port Trunking Protocol (PTP) is used to carry tagged frames.
- TF-TN port link—A single link can be established to connect an F port to an HBA to carry tagged frames (see the link 1a and 1b in Figure 5: Trunking F Ports, on page 184) using Exchange Virtual Fabrics

Protocol (EVFP). A server can reach multiple VSANs through a TF port without inter-VSAN routing (IVR).

- TF-TNP port link—A single link can be established to connect an TF port to an TNP port using the PTP protocol to carry tagged frames (see the link 2 in Figure 5: Trunking F Ports, on page 184). PTP is used because PTP also supports trunking PortChannels.

Note The TF-TNP port link between a third-party NPV core and a Cisco NPV switch is established using the EVFP protocol.

- A Fibre Channel VSAN is called Virtual Fabric and uses a VF_ID in place of the VSAN ID. By default, the VF_ID is 1 for all ports. When an N port supports trunking, a pWWN is defined for each VSAN and called a logical pWWN. In the case of MDS core switches, the pWWNs for which the N port requests additional FC_IDs are called virtual pWWNs.

Trunking Protocols

The trunking protocol is important for trunking operations on the ports. The protocols enable the following activities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

Table 23: Supported Trunking Protocols, on page 185 specifies the protocols used for trunking and channeling.

Table 23: Supported Trunking Protocols

Trunk Link	Default
TE-TE port link	Cisco EPP (PTP)
TF-TN port link[7]	FC-LS Rev 1.62 EVFP
TF-TNP port link	Cisco EPP (PTP)
E or F PortChannel	Cisco EPP (PCP)
TF Port Channel	Cisco EPP (PTP and PCP)
Third-party TF-TNP port link [8]	FC-LS Rev 1.62 EVFP

[7] These features are not currently supported.
[8] These features are not currently supported.

By default, the trunking protocol is enabled on E ports and disabled on F ports. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected. The TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

Note We recommend that both ends of a trunking link belong to the same port VSAN. On certain switches or fabric switches where the port VSANs are different, one end returns an error and the other end is not connected.

Trunk Modes

By default, trunk mode is enabled on all Fibre Channel interfaces (Mode: E, F, FL, Fx, ST, and SD) on non-NPV switches. On NPV switches, by default, trunk mode is disabled. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The trunk mode configuration at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends (see Table 24: Trunk Mode Status Between Switches , on page 186).

Table 24: Trunk Mode Status Between Switches

Your Trunk Mode Configuration			Resulting State and Port Mode	
Port Type	**Switch 1**	**Switch 2**	**Trunking State**	**Port Mode**
E ports	On	Auto or on	Trunking (EISL)	TE port
	Off	Auto, on, or off	No trunking (ISL)	E port
	Auto	Auto	No trunking (ISL)	E port
Port Type	**Core Switch**	**NPV Switch**	**Trunking State**	**Link Mode**
F and NP ports	On	Auto or on	Trunking	TF-TNP link
	Auto	On	Trunking	TF-TNP link
	Off	Auto, on, or off	No trunking	F-NP link

Tip The preferred configuration on the Cisco MDS 9000 Series Multilayer Switches is one side of the trunk set to auto and the other side set to on.

Note When connected to a third-party switch, the trunk mode configuration on E ports has no effect. The ISL is always in a trunking disabled state. In the case of F ports, if the third-party core switch ACC's physical FLOGI with the EVFP bit is configured, then EVFP protocol enables trunking on the link.

Trunk-Allowed VSAN Lists and VF_IDs

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

Switch 1 (see Figure 6: Default Allowed-Active VSAN Configuration, on page 188) has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational (see Figure 6: Default Allowed-Active VSAN Configuration, on page 188).

For all F, N, and NP ports, the default VF_ID is 1 when there is no VF_ID configured. The trunk-allowed VF_ID list on a port is same as the list of trunk-allowed VSANs. VF_ID 4094 is called the control VF_ID and it is used to define the list of trunk-allowed VF-IDs when trunking is enabled on the link.

If F port trunking and channeling is enabled, or if **switchport trunk mode on** is configured in NPV mode for any interface, or if NP PortChannel is configured, the VSAN and VF-ID ranges available for the configuration are as described in Table 25: VSAN and VF-ID Reservations, on page 187.

Table 25: VSAN and VF-ID Reservations

VSAN or VF-ID	Description
000h	Cannot be used as virtual fabric identifier.
001h(1) to EFFh(3839)	This VSAN range is available for user configuration.
F00h(3840) to FEEh(4078)	Reserved VSANs and they are not available for user configuration.
FEFh(4079)	EVFP isolated VSAN.
FF0h(4080) to FFEh(4094)	Used for vendor-specific VSANs.
FFFh	Cannot be used as virtual fabric identifier.

Note If the VF_ID of the F port and the N port do not match, then no tagged frames can be exchanged.

Figure 6: Default Allowed-Active VSAN Configuration

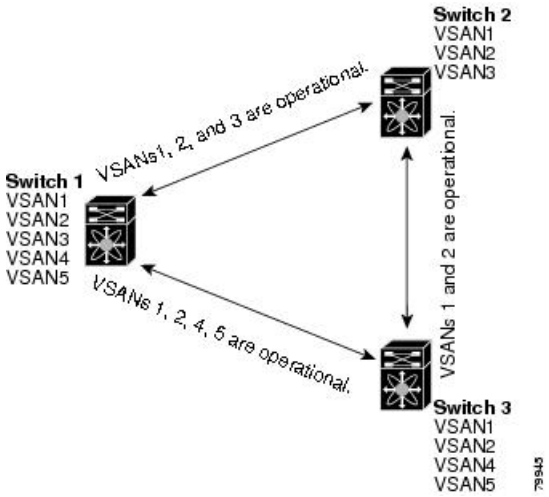

You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using Figure 6: Default Allowed-Active VSAN Configuration, on page 188 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 7: Operational and Allowed VSAN Configuration, on page 188). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Figure 7: Operational and Allowed VSAN Configuration

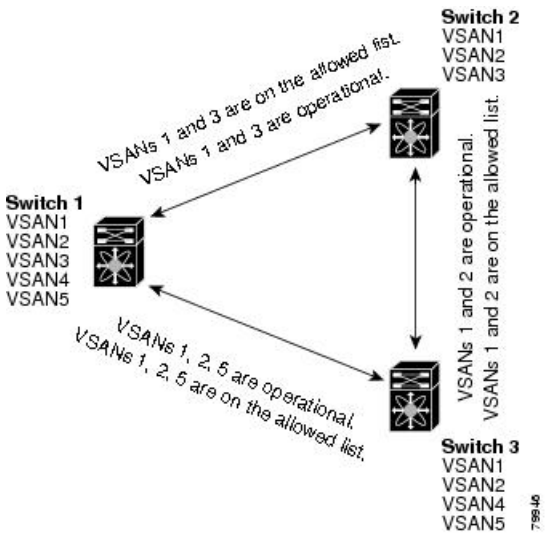

Guidelines and Limitations

General Guidelines and Limitations

The trunking feature has the following general configuration guidelines and limitations:

- You will see the **switchport trunk mode off** command added to F ports after upgrading from Cisco MDS NX-OS Release 8.1(1) to Cisco MDS NX-OS Release 8.2(1).
- F ports support trunking in Fx mode.
- The trunk-allowed VSANs configured for TE, TF, and TNP links are used by the trunking protocol to determine the allowed active VSANs in which frames can be received or transmitted.
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.
- Trunking F ports and trunking F PortChannels are not supported on the following hardware:
  - 91x4 switches, if NPIV is enabled and used as the NPIV core switch.
  - Generation 1 2-Gbps Fibre Channel switching modules.
- On core switches, the FC-SP authentication will be supported only for the physical FLOGI from the physical pWWN.
- No FC-SP authentication is supported by the NPV switch on the server F ports.
- MDS does not enforce the uniqueness of logical pWWNs across VSANs.
- DPVM is not supported on trunked F port logins.
- The DPVM feature is limited to the control of the port VSAN, since the EVFP protocol does not allow changing the VSAN on which a logical pWWN has done FLOGI.
- The port security configuration will be applied to both the first physical FLOGI and the per VSAN FLOGIs.
- Trunking is not supported on F ports that have FlexAttach enabled.
- On MDS 91x4 core switches, hard zoning can be done only on F ports that are doing either NPIV or trunking. However, in NPV mode, this restriction does not apply since zoning is enforced on the core F port.

Note Fibre Channel Security Protocol (FC-SP) is not supported for 6.2(1) release on MDS 9710, but targeted for a future release.

Upgrade and Downgrade Limitations

The trunking and channeling feature includes the following upgrade and downgrade limitations:

- When F port trunking or channeling is configured on a link, the switch cannot be downgraded to Cisco MDS SAN-OS Release 3.x and NX-OS Release 4.1(1b), or earlier.
- If you are upgrading from a SAN-OS Release 3.x to NX-OS Release 5.0(1), and you have not created VSAN 4079, the NX-OS software will automatically create VSAN 4079 and reserve it for EVFP use.

If VSAN 4079 is reserved for EVFP use, the **switchport trunk allowed vsan** command will filter out VSAN 4079 from the allowed list, as shown in the following example:

```
switch(config-if)# switchport trunk allowed vsan 1-4080
1-4078,4080
```

- If you have created VSAN 4079, the upgrade to NX-OS Release 5.0(1) will have no affect on VSAN 4079.
- If you downgrade after NX-OS Release 5.0(1), the VSAN will no longer be reserved for EVFP use.

Difference Between TE Ports and TF-TNP Ports

In case of TE ports, the VSAN will in be initializing state when VSAN is coming up on that interface and when peers are in negotiating phase. Once the handshake is done, VSAN will be moved to up state in the successful case, and isolated state in the case of failure. Device Manager will show the port status as amber during initializing state and it will be green once VSANs are up.

This example shows the trunk VSAN states of a TE port:

```
switch# show interface fc2/15
fc2/15 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:4f:00:0d:ec:6d:2b:40
    Peer port WWN is 20:0a:00:0d:ec:3f:ab:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 250
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (isolated)                 (100-101)
    Trunk vsans (initializing)             ()
```

In case of TF ports, after the handshake, one of the allowed VSANs will be moved to the up state. All other VSANs will be in initializing state even though the handshake with the peer is completed and successful. Each VSAN will be moved from initializing state to up state when a server or target logs in through the trunked F or NP ports in the corresponding VSAN.

Note In case of TF or TNP ports, the Device Manager will show the port status as amber even after port is up and there is no failure. It will be changed to green once all the VSAN has successful logins.

This example shows a TF port information after the port is in the up state:

```
sw7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

This example shows the TF port information when a server logs in on noninternal FLOGI VSAN. VSAN 2183 is moved to the up state when the server logs in to VSAN 2183.

```
w7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2
182-2183)
    Trunk vsans (up)                       (1,2183)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

Trunking Misconfiguration Examples

If you do not configure the VSANs correctly, issues with the connection may occur. For example, if you merge the traffic in two VSANs, both VSANs will be mismatched. The trunking protocol validates the VSAN interfaces at both ends of a link to avoid merging VSANs (see Figure 8: VSAN Mismatch, on page 192).

Figure 8: VSAN Mismatch

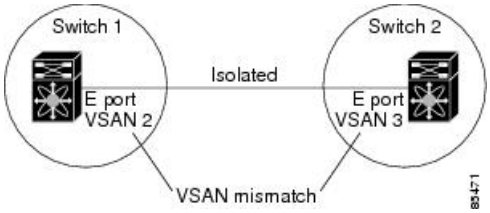

The trunking protocol detects potential VSAN merging and isolates the ports involved (see Figure 8: VSAN Mismatch, on page 192).

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Series Multilayer Switches (see Figure 9: Third-Party Switch VSAN Mismatch, on page 192).

Figure 9: Third-Party Switch VSAN Mismatch

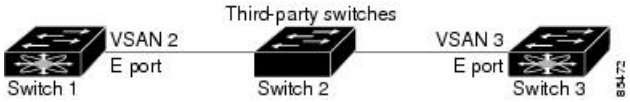

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. Cisco DCNM-SAN helps detect such topologies.

Default Settings

Table 26: Default Trunk Configuration Parameters , on page 193 lists the default settings for trunking parameters.

Table 26: Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	ON on non-NPV and MDS core switches. OFF on NPV switches.
Allowed VSAN list	1 to 4093 user-defined VSAN IDs.
Allowed VF-ID list	1 to 4093 user-defined VF-IDs.
Trunking protocol on E ports	Enabled.
Trunking protocol on F ports	Disabled.

Configuring Trunking

Enabling the Cisco Trunking and Channeling Protocols

To enable or disable the Cisco trunking and channeling protocol, perform these steps:

Before you begin

To avoid inconsistent configurations, disable all ports with a **shutdown** command before enabling or disabling the trunking protocols.

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **trunk protocol enable**

Enables the Cisco PTP trunking protocol (default).

Step 3 switch(config)# **no trunk protocol enable**

Disables the Cisco PTP trunking protocol.

Enabling the F Port Trunking and Channeling Protocol

To enable or disable the F port trunking and channeling protocol, perform these steps:

Before you begin

To avoid inconsistent configurations, shut all ports before enabling or disabling the trunking protocols.

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature fport-channel-trunk**

Enables the F port trunking and channeling protocol (default).

Step 3 switch(config)# **no feature fport-channel-trunk**

Disables the F port trunking and channeling protocol.

Configuring Trunk Mode

To configure trunk mode, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1**

Configures the specified interface.

Step 3 switch(config-if)# **switchport trunk mode on**

Enables (default) the trunk mode for the specified interface.

switch(config-if)# **switchport trunk mode off**

(Optional) Disables the trunk mode for the specified interface.

switch(config-if)# **switchport trunk mode auto**

(Optional) Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1**

Configures the specified interface.

Step 3 switch(config-if)# **switchport trunk allowed vsan 2-4**

Changes the allowed list for the specified VSANs.

Step 4 switch(config-if)# **switchport trunk allowed vsan add 5**

Expands the specified VSAN (5) to the new allowed list.

switch(config-if)# **no switchport trunk allowed vsan 2-4**

(Optional) Deletes VSANs 2, 3, and 4.

switch(config-if)# **no switchport trunk allowed vsan add 5**

(Optional) Deletes the expanded allowed list.

Verifying Trunking Configuration

To display trunking configuration information, perform one of the following tasks:

Command	Purpose
show interface fc *slot/port*	Displays the interface configuration information that includes trunking, trunk mode, allowed VSANs, and status.
show trunk protocol	Displays whether the trunk protocol is enabled.
show interface trunk vsan *numbers*	Displays whether the interface is trunking, and the allowed VSAN list for each trunking interface.

For detailed information about the fields in the output from these commands, refer to the Cisco MDS NX-OS Command Reference.

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples Displays a Trunked Fibre Channel Interface, on page 196 to Displays Per VSAN Information on Trunk Ports, on page 197.

Displays a Trunked Fibre Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
    Hardware is Fibre Channel
    Port WWN is 20:0d:00:05:30:00:58:1e
    Peer port WWN is 20:0d:00:05:30:00:59:1e
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Receive B2B Credit is 255
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      233996 frames input, 14154208 bytes, 0 discards
        0 CRC, 0 unknown class
        0 too long, 0 too short
      236 frames output, 13818044 bytes, 0 discards
      11 input OLS, 12 LRR, 10 NOS, 28 loop inits
      34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

Displays the Trunking Protocol

```
switch# show trunk protocol
Trunk protocol is enabled
```

Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
 Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
    Vsan 1 is up, FCID is 0x760001
    Vsan 2 is up, FCID is 0x6f0001
fc3/11 is trunking
    Belongs to port-channel 6
 Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
 Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
```

Configuration Example for F Port Trunking

This example shows how to configure trunking and bring up the TF-TNP link between an F port in the NPIV core switch and an NP port in the NPV switch:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch:

Example:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

Example:

```
switch(config)# feature npiv
```

Step 3 Configure the port mode to auto, F, or Fx on the MDS core switch:

Example:

```
switch(config)# interface fc1/2
switch(config-if)# switchport mode F
```

Step 4 Configure the trunk mode to ON on the MDS core switch:

Example:

```
switch(config-if)# switchport trunk mode on
```

Step 5 Configure the port mode to NP on the NPV switch:

Example:

```
switch(config)# interface fc1/2
switch(config-if)# switchport mode NP
```

Step 6 Configure the trunk mode to ON on the NPV switch:

Example:

```
switch(config-if)# switchport trunk mode on
```

Step 7 Set the port administrative state on NPIV and NPV switches to ON:

Example:

```
switch(config)# interface fc1/2
switch(config-if)# shut
```

```
switch(config-if)# no shut
```

Step 8 Save the configuration.

Example:

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

Configuring PortChannels

This chapter provides information about PortChannels and how to configure the PortChannels.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About PortChannels

PortChannels Overview

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy (see Figure 10: PortChannel Flexibility, on page 203). PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

Figure 10: PortChannel Flexibility

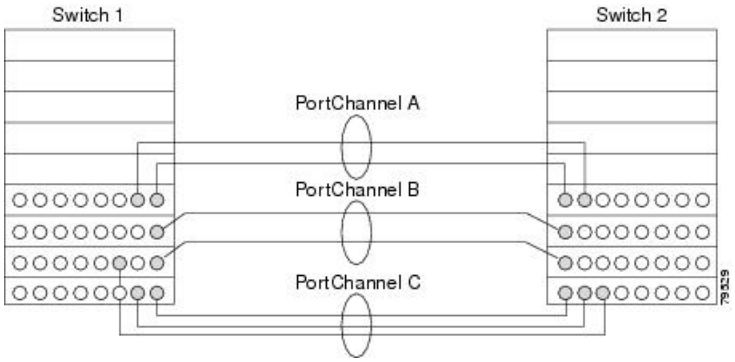

PortChannels on Cisco MDS 9000 Series Multilayer Switches allow flexibility in configuration. This illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

E PortChannels

An E PortChannel refers to the aggregation of multiple E ports into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

A PortChannel has the following features and restrictions:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).

- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

Note See the Cisco MDS 9000 Series NX-OS Fabric Configuration Guide for information about failover scenarios for PortChannels and FSPF links.

F and TF PortChannels

An F PortChannel is also a logical interface that combines a set of F ports connected to the same Fibre Channel node and operates as one link between the F ports and the NP ports. The F PortChannels support bandwidth utilization and availability like the E PortChannels. F PortChannels are mainly used to connect MDS core and NPV switches to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN.

An F PortChannel trunk combines the functionality and advantages of a TF port and an F PortChannel. This logical link uses the Cisco PTP and PCP protocols over Cisco EPP (ELS).

Note If a Cisco MDS 9124 or 9134 switch is used as a core switch, only a nontrunking F PortChannel is supported. Trunking is not supported on this platform when NPIV enabled.

PortChanneling and Trunking

Trunking is a commonly used storage industry term. However, the Cisco NX-OS software and switches in the Cisco MDS 9000 Series Multilayer Switches implement trunking and PortChanneling as follows:

- PortChanneling enables several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. For example, when trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Series Multilayer Switches. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see Figure 11: Trunking Only, on page 204 and Figure 12: PortChanneling and Trunking, on page 205). See Configuring Trunking, on page 181 for information on trunked interfaces.

Figure 11: Trunking Only

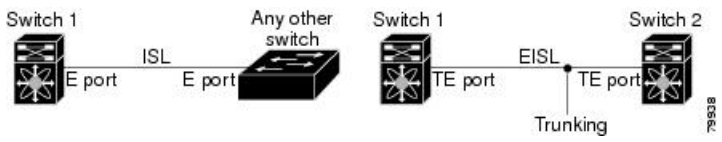

Figure 12: PortChanneling and Trunking

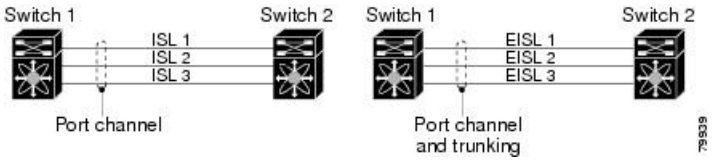

PortChanneling and trunking are used separately across an ISL.

- PortChanneling—Interfaces can be channeled between the following sets of ports:
  - E ports and TE ports
  - F ports and NP ports
  - TF ports and TNP ports
- Trunking—Trunking permits carrying traffic on multiple VSANs between switches.

See the Cisco MDS 9000 Series NX-OS Fabric Configuration Guide.

- Both PortChanneling and trunking can be used between TE ports over EISLs.

Load Balancing

Two methods support the load-balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

Figure 13: SID1 and DID1-Based Load Balancing, on page 206 illustrates how source ID 1 (SID1) and destination ID1 (DID1) based load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 13: SID1 and DID1-Based Load Balancing

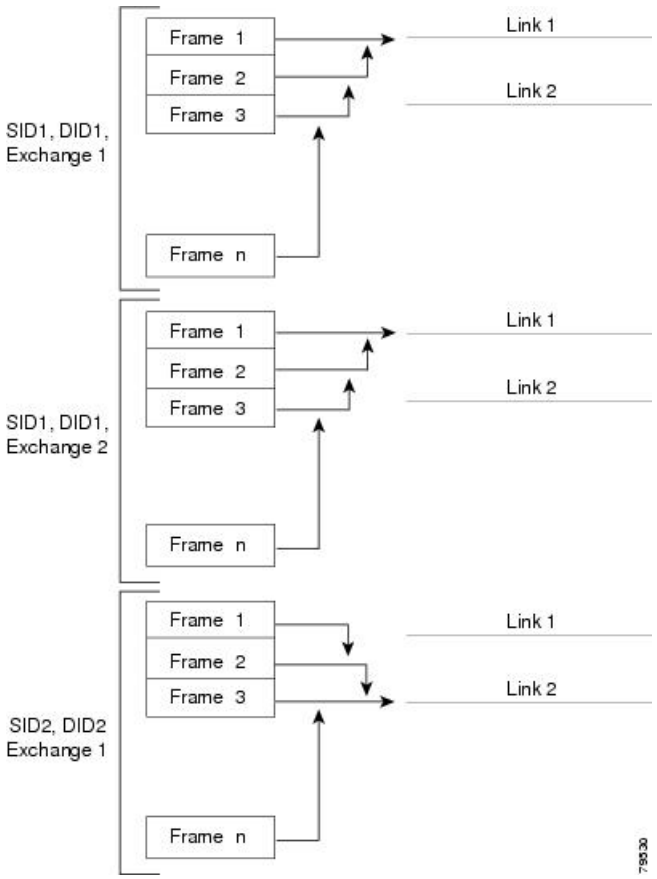

Figure 14: SID1, DID1, and Exchange-Based Load Balancing, on page 207 illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 14: SID1, DID1, and Exchange-Based Load Balancing

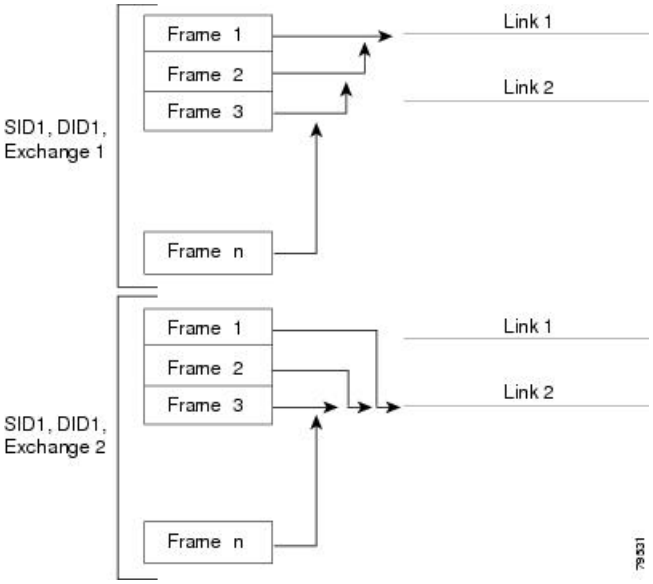

For more information on configuring load balancing and in-order delivery features, see the Cisco MDS 9000 Series NX-OS Fabric Configuration Guide.

PortChannel Modes

You can configure each PortChannel with a channel group mode parameter to determine the PortChannel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- ON (default)—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON mode require you to explicitly enable and disable the PortChannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- ACTIVE—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. The ACTIVE PortChannel mode allows automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 27: Channel Group Configuration Differences , on page 208 compares ON and ACTIVE modes.

Table 27: Channel Group Configuration Differences

ON Mode	ACTIVE Mode
No protocol is exchanged.	A PortChannel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel.	Moves interfaces to the isolated state if its operational values are incompatible with the PortChannel.
When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end.	When you add or modify a PortChannel interface, the PortChannel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a PortChannel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

PortChannel Deletion

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode used (ACTIVE and ON), the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the Graceful Shutdown, on page 21).

If you delete the PortChannel for one port, then the individual ports within the deleted PortChannel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

Interfaces in a PortChannel

You can add or remove a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel. Removing an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

Note For information about PortChannel support on Generation 2 switching modules, see the PortChannel Limitations, on page 85.

Interface Addition to a PortChannel

You can add a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel.

A port can be configured as a member of a static PortChannel only if the following configurations are the same in the port and the PortChannel:

- Speed
- Mode
- Rate mode
- Port VSAN
- Trunking mode
- Allowed VSAN list or VF-ID list

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the Generation 1 PortChannel Limitations, on page 216 and Graceful Shutdown, on page 21).

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a PortChannel. The compatibility check is performed before a port is added to the PortChannel.

The check ensures that the following parameters and settings match at both ends of a PortChannel:

- Capability parameters (type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, rate mode, port VSAN, allowed VSAN list, and port security).

Note Ports in shared rate mode cannot form a PortChannel or a trunking PortChannel.

- Operational parameters (remote switch WWN and trunking mode).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the ON mode.

- An interface enters the isolated state if the interface is configured in the ACTIVE mode.

Forcing an Interface Addition

You can force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the addition.

Note When PortChannels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the Graceful Shutdown, on page 21) sections.

Interface Deletion from a PortChannel

When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the Generation 1 PortChannel Limitations, on page 216 and Graceful Shutdown, on page 21 sections).

PortChannel Protocols

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco NX-OS software provides robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

A protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default.

The PortChannel protocol expands the PortChannel functional model in Cisco MDS switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a PortChannel. The protocol ensures that a set of ports are eligible to be part of the same PortChannel. They are only eligible to be part of the same PortChannel if all the ports have a compatible partner.

The PortChannel protocol uses two subprotocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications such as write acceleration, work for PortChannels over FCIP links.
- Autocreation protocol—Automatically aggregates compatible ports into a PortChannel.

Channel Group Creation

Note Channel groups are not supported on internal ports in the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

Assuming link A1-B1 comes up first (see Figure 15: Autocreating Channel Groups, on page 211), that link is operational as an individual link. When the next link comes up, for example, A2-B2, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is because of the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 15: Autocreating Channel Groups

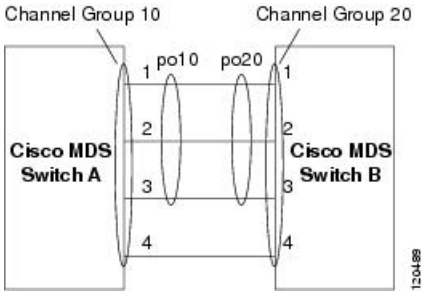

The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

Table 28: Channel Group Configuration Differences , on page 211 identifies the differences between user-configured and auto-configured channel groups.

Table 28: Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.

User-Configured Channel Group	Autocreated Channel Group
You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration.	All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.
Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface.	Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist.

Note Autocreation is not supported as of MDS NX-OS Release 4.1(1b) and later.

Autocreation

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated PortChannel.
  - A port is aggregated with another compatible port to form a new PortChannel.
- Newly created PortChannels are allocated from the maximum possible PortChannel (128 for Generation 1 or a combination of Generation 1 and Generation 2 switches, or 256 for Generation 2 switches) in a decreasing order based on availability. If all 128 (or 256) numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel.
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.
- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.

- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.

Tip When enabling autocreation in any switch in the Cisco MDS 9000 Series Multilayer Switches, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to an autocreated PortChannel.

Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. Once performed, this task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and the autocreation of channel group is implicitly disabled for all member ports.

Tip If you enable persistence, be sure to enable it at both ends of the PortChannel.

Prerequisites for PortChannels

Before configuring a PortChannel, consider the following guidelines:

- Configure the PortChannel across switching modules to implement redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.

Note On switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, you can configure a maximum of 128 PortChannels. On switches with only Generation 2 switching modules, or Generation 2 and Generation 3 switching modules, you can configure a maximum of 256 PortChannels.

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled because an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see Figure 17: Misconfigured Configurations, on page 218 for an example of an invalid configuration).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and reenable the links.

If all three conditions are not met, the faulty link is disabled.

Enter the **show interface** command for that interface to verify that the PortChannel is functioning as required.

Default Settings

Table 29: Default PortChannel Parameters , on page 215 lists the default settings for PortChannels.

Table 29: Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default PortChannel mode	ON mode on non-NPV and NPIV core switches. ACTIVE mode on NPV switches.
Autocreation	Disabled.

Guidelines and Limitations

General Guidelines and Limitations

Cisco MDS 9000 Series Multilayer switches support the following number of PortChannels per switch:

- Switches with only Generation 1 switching modules do not support F and TF PortChannels.
- Switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, support a maximum of 128 PortChannels. Only Generation 2 ports can be included in the PortChannels.
- Switches with only Generation 2 switching modules or Generation 2 and Generation 3 modules support a maximum of 256 PortChannels with 16 interfaces per PortChannel.
- A PortChannel number refers to the unique identifier for each channel group. This number ranges from of 1 to 256.

Generation 1 PortChannel Limitations

This section includes the restrictions on creation and addition of PortChannel members to a PortChannel on Generation 1 hardware:

- The 32-port 2-Gbps or 1-Gbps switching module.
- The MDS 9140 and 9120 switches.

When configuring the host-optimized ports on Generation 1 hardware, the following PortChannel guidelines apply:

- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules; only the first port of each group of 4 ports is included in a PortChannel.
  - You can configure only the first port in each 4-port group as an E port (for example, the first port in ports 1–4, the fifth port in ports 5–8, and so on). If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2–4, 6–8, and so on) are not usable and remain in the shutdown state.
  - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.

F and TF PortChannel Limitations

The following guidelines and restrictions are applicable for F and TF PortChannels:

- The ports must be in F mode.
- Automatic creation is not supported.

- The PortChannel interface must be in ACTIVE mode when multiple FCIP interfaces are grouped with WA.
- ON mode is not supported. Only ACTIVE-ACTIVE mode is supported. By default, the mode is ACTIVE on the NPV switches.
- Devices logged in through F PortChannel on an MDS switch are not supported in IVR non-NAT configuration. The devices are supported only in IVR NAT configuration.
- Port security rules are enforced only on physical pWWNs at the single link level.
- FC-SP authenticates only the first physical FLOGI of every PortChannel member.
- Since the FLOGI payload carries only the VF bits to trigger the use of a protocol after the FLOGI exchange, those bits will be overridden. In the case of the NPV switches, the core has a Cisco WWN and will try to initiate the PCP protocol.
- The name server registration of the N ports logging in through an F PortChannel will use the fWWN of the PortChannel interface.
- DPVM configuration is not supported.
- The PortChannel port VSAN cannot be configured using DPVM.
- The Dynamic Port VSAN Management (DPVM) database will be queried only for the first physical FLOGI of each member, so that the port VSAN can be configured automatically.
- DPVM does not bind FC_IDs to VSANs, but pWWNs to VSANs. It will be queried only for the physical FLOGI.

Valid and Invalid PortChannel Examples

PortChannels are created with default values. You can change the default configuration just like any other physical interface.

Figure 16: Valid PortChannel Configurations, on page 217 provides examples of valid PortChannel configurations.

Figure 16: Valid PortChannel Configurations

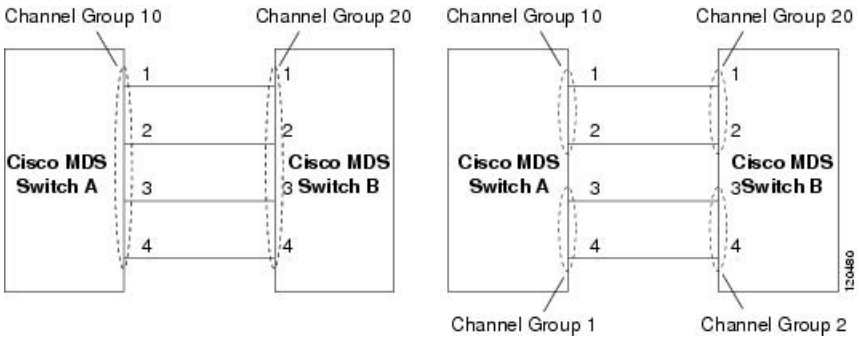

Figure 17: Misconfigured Configurations, on page 218 provides examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 17: Misconfigured Configurations

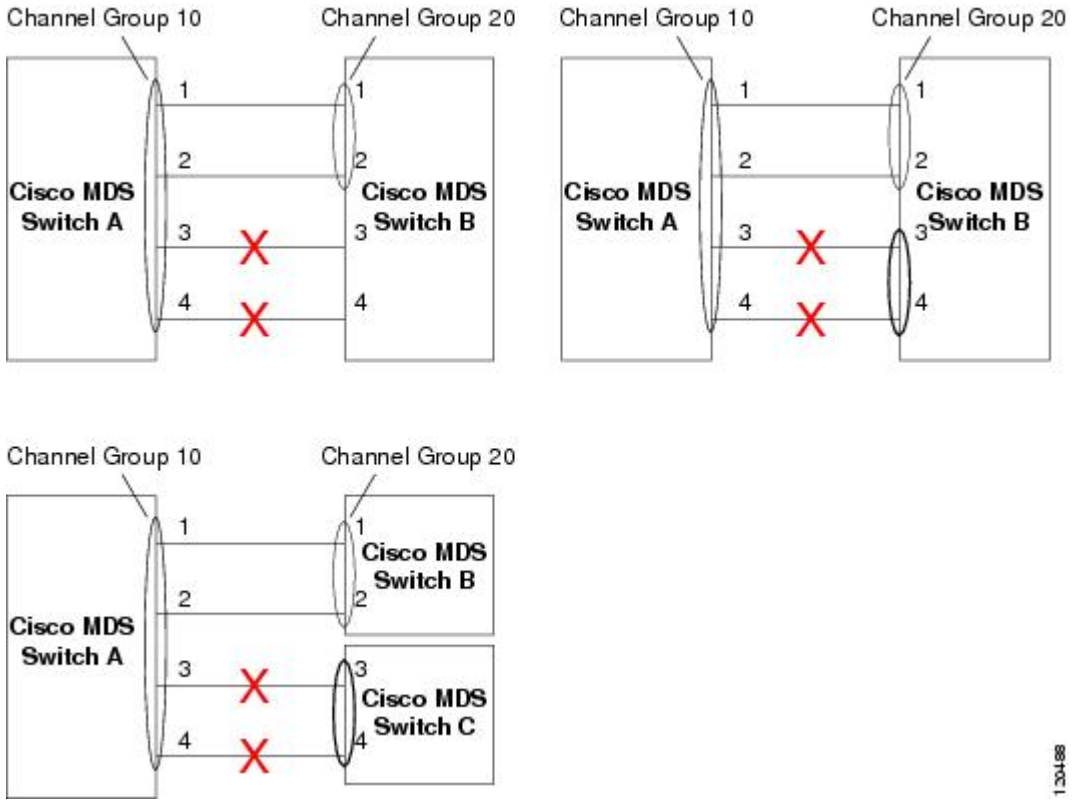

Configuring PortChannels

Configuring PortChannels Using the Wizard Creating a PortChannel

To create a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

Configures the specified PortChannel (1) using the default ON mode.

Configuring the PortChannel Mode

By default, the CLI and the Device Manager create the PortChannel in ON mode in the NPIV core switches and ACTIVE mode on the NPV switches. DCNM-SAN creates all PortChannels in ACTIVE mode. We recommend that you create PortChannels in ACTIVE mode.

Note An F PortChannel is supported only on ACTIVE mode.

To configure ACTIVE mode, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

Configures the specified PortChannel (1) using the default ON mode.

Step 3 switch(config-if)# **channel mode active**

Configures the ACTIVE mode.

switch(config-if)# **no channel mode active**

(Optional) Reverts to the default ON mode.

Deleting PortChannels

To delete a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no interface port-channel 1**

Deletes the specified PortChannel (1), its associated interface mappings, and the hardware associations for this PortChannel.

Adding an Interface to a PortChannel

To add an interface to a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/15**

Configures the specified port interface (fc1/15).

Step 3 switch(config-if)# **channel-group 15**

Adds physical Fibre Channel port 1/15 to channel group 15. If channel group 15 does not exist, it is created. The port is shut down.

Adding a Range of Ports to a PortChannel

To add a range of ports to a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1 - 5**

Configures the specified range of interfaces. In this example, interfaces from 1/1 to 1/5 are configured.

Step 3 switch(config-if)# **channel-group 2**

Adds physical interfaces 1/1, 1/2, 1/3, 1/4, and 1/5 to channel group 2. If channel group 2 does not exist, it is created.

If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces.

What to do next

Note By default, the CLI adds a interface normally to a PortChannel, while DCNM-SAN adds the interface by force, unless specified explicitly.

Forcing an Interface Addition

To force the addition of a port to a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1**

Specifies the interface fc1/1.

Step 3 switch(config-if)# **channel-group 1 force**

Forces the addition of the physical port for interface fc1/1 to channel group 1. The port is shut down.

Deleting an Interface From a PortChannel

To delete a physical interface (or a range of physical interfaces) from a PortChannel, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1**

Enters the selected physical interface level.

Step 3 switch(config)# **interface fc1/1 - 5**

Enters the selected range of physical interfaces.

Step 4 switch(config-if)# **no channel-group 2**

Deletes the physical Fibre Channel interfaces in channel group 2.

Enabling and Configuring Autocreation

To configure automatic channel groups, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc8/13**

Enters the configuration mode for the selected interface(s).

Step 3 switch(config- if)# **channel-group auto**

Automatically creates the channel group for the selected interface(s).

switch(config- if)# **no channel-group auto**

(Optional) Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled.

Converting to Manually Configured Channel Groups

You can convert autocreated channel group to a user-configured channel group using the **port-channel** *channel-group-number* **persistent** EXEC command. If the PortChannel does not exist, this command is not executed.

Verifying PortChannel Configuration

To display PortChannel configuration information, perform one of the following tasks:

Command	Purpose
show port-channel summary	Displays a summary of PortChannels within the switch. A one-line summary of each PortChannel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the PortChannel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a PortChannel and can change if the port goes down. The FOP is also identified by an asterisk (*).
show port-channel database	Displays the PortChannel configured in the default ON mode and ACTIVE mode.
show port-channel consistency	Displays the consistency status without details.
show port-channel consistency detail	Displays the consistency status with details.
show port-channel usage	Displays the PortChannel usage.
show port-channel compatibility-parameters	Displays the PortChannel compatibility.
show interface fc *slot/port*	Displays autocreated PortChannels.
show port-channel database interface port-channel *number*	Displays the specified PortChannel interface.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Series NX-OS Command Reference*.

You can view specific information about existing PortChannels at any time from EXEC mode. The following **show** commands provide further details on existing PortChannels. You can force all screen output to go to a printer or save it to a file. See Examples Displays the PortChannel Summary, on page 223 to Displays the PortChannel Summary, on page 223.

Displays the PortChannel Summary

```
switch# show port-channel summary
 ----------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
----------------------------------------------------------------------------
port-channel 77                2                 0                 --
port-channel 78                2                 0                 --
port-channel 79                2                 2                 fcip200
```

Displays the PortChannel Configured in the Default ON Mode

```
switch# show port-channel database

port-channel 77
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
             fcip2    [down]
port-channel 78
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

Displays the PortChannel Configured in the ACTIVE Mode

```
switch# show port-channel database

port-channel 77
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
             fcip2    [down]
port-channel 78
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

The **show port-channel consistency** command has two options: without details and with details.

Displays the Consistency Status without Details

```
switch# show port-channel consistency
Database is consistent
```

Displays the Consistency Status with Details

```
switch# show port-channel consistency detail
Authoritative port-channel database:
============================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
============================================
database 1: from module 5
============================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
============================================
database 2: from module 4
============================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
...
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

Displays the PortChannel Usage

```
switch# show port-channel usage
Totally 3 port-channel numbers used
===================================
Used  :   77 - 79
Unused:   1 - 76 , 80 - 256
```

Use the existing **show** commands to obtain further details on autocreated channel group attributes. Autocreated PortChannels are indicated explicitly to help differentiate them from the manually created PortChannels.

Displays the PortChannel Compatibility

```
switch# show port-channel compatibility-parameters
physical port layer                 fibre channel or ethernet
    port mode                           E/AUTO only
    trunk mode
    speed
    port VSAN
    port allowed VSAN list
```

Displays Autocreated PortChannels

```
switch# show interface fc1/1
fc1/1 is trunking
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:0a:00:0b:5f:3b:fe:80
    ...
    Receive data field Size is 2112
    Beacon is turned off
    Port-channel auto creation is enabled
 Belongs to port-channel 123
...
```

Displays the Specified PortChannel Interface

```
switch# show port-channel database interface port-channel 128
port-channel 128
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    Channel is auto created
    First operational port is fc1/1
    1 ports in total, 1 ports up
    Ports:   fc1/1   [up] *
```

Displays the PortChannel Summary

```
switch# show port-channel summary
--------------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
port-channel 1                 1                 0                 --
port-channel 2                 1                 1                 fc8/13
port-channel 3                 0                 0                 --
port-channel 4                 0                 0                 --
port-channel 5                 1                 1                 fc8/3
port-channel 6                 0                 0                 --
```

Configuration Examples for F and TF PortChannels

This example shows how to configure F PortChannel in shared mode and bring up the link (not supported on the MDS 91x4 switches) between F ports on the NPIV core switches and NP ports on the NPV switches:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch.

Example:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

Example:

```
switch(config)# feature npiv
```

Step 3 Create the PortChannel on the MDS core switch:

Example:

```
switch(config)# interface port-channel 1
switch(config-if)# switchport mode F
switch(config-if)# channel mode active
switch(config-if)# switchport trunk mode off
switch(config-if)# switchport rate-mode shared
switch(config-if)# exit
```

Step 4 Configure the PortChannel member interfaces on the core switch:

Example:

```
switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport trunk mode off
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# channel-group 1
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the PortChannel on the NPV switch:

Example:

```
switch(config)# interface port-channel 1
switch(config-if)# switchport mode NP
switch(config-if)# switchport rate-mode shared
switch(config-if)# exit
```

Step 6 Configure the PortChannel member interfaces on the NPV switch:

Example:

```
switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# switchport mode NP
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport trunk mode off
switch(config-if)# channel-group 1
switch(config-if)# no shut
switch(config-if)# exit
```

Step 7 Set the administrative state of all the PortChannel member interfaces in both NPIV core switch and the NPV switch to ON:

Example:

```
switch(config)# interface fc1/1-3
switch(config-if)# shut
switch(config-if)# >no shut
switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# >no shut
```

Configuration Examples for F and TF PortChannels (Dedicated Mode)

Note The speed configuration must be the same for all member interfaces in a PortChannel. While configuring the channel in dedicated mode, ensure that required bandwidth is available to the ports.

This example shows how to configure channeling in dedicated mode and bring up the TF-TNP PortChannel link between TF ports in the NPIV core switch, and TNP ports in the NPV switch:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch:

Example:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

Example:

```
switch(config)# feature npiv
```

Step 3 Create the PortChannel on the MDS core switch:

Example:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport mode F
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# channel mode active
switch(config-if)# exit
```

Step 4 Configure the PortChannel member interfaces on the MDS core switch in dedicated mode:

Example:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the PortChannel in dedicated mode on the NPV switch:

Example:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode NP
switch(config-if)# no shut
switch(config-if)# exit
```

Step 6 Configure the PortChannel member interfaces on the NPV switch in dedicated mode:

Example:

```
switch(config)# interface fc3/1-3
switch(config-if)# shut
switch(config-if)# switchport mode NP
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Step 7 Set the administrative state of all the PortChannel member interfaces in both NPIV core switch and the NPV switch to ON:

Example:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# no shut
switch(config)# interface fc3/1-3
switch(config-if)# shut
switch(config-if)# no shut
```

Configuring N Port Virtualization

This chapter provides information about N port virtualization and how to configure N port virtualization.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About N Port Virtualization

NPV Overview

N port virtualization (NPV) reduces the number of Fibre Channel domain IDs in SANs. Switches operating in the NPV mode do not join a fabric. They pass traffic between NPV core switch links and end devices, which eliminates the domain IDs for these edge switches.

NPV is supported by the following Cisco MDS 9000 switches and Cisco Nexus 5000 Series switches only:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco MDS 9148 Multilayer Fabric Switch
- Cisco MDS 9148S Multilayer Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter
- Cisco MDS 9396S Multilayer Fabric Switch
- Cisco Nexus 5000 Series switches

Note NPV is available on these switches only while in NPV mode; if in switch mode, NPV is not available.

N Port Identifier Virtualization

N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port. This feature allows multiple applications on the N port to use different identifiers and allows access control, zoning, and port security to be implemented at the application level Figure 18: NPIV Example, on page 235 shows an example application using NPIV.

Figure 18: NPIV Example

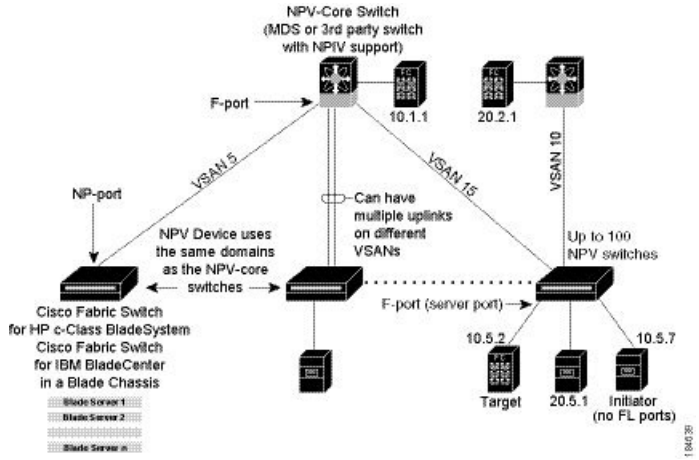

You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

Note All of the N port identifiers are allocated in the same VSAN.

N Port Virtualization

Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to edge devices. Such a model is cost-effective because the per port cost for director class switches is much higher than that of fabric switches. However, as the number of ports in the fabric increases, the number of switches deployed also increases, and you can end up with a significant increase in the number of domain IDs. This challenge becomes even more difficult when additional blade chassis are deployed in Fibre Channel networks.

NPV addresses the increase in the number of domain IDs needed to deploy a large number of the ports by making a fabric or blade switch appear as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric or blade switch. NPV aggregates multiple locally connected N ports into one or more external NP links, which shares the domain ID of the NPV core switch among multiple NPV switches. NPV also allows multiple devices to attach to same port on the NPV core switch, which reduces the need for more ports on the core

For more information on scalability limits, see the Cisco MDS NX-OS Configuration Limits guide.

Figure 19: Cisco NPV Fabric Configuration

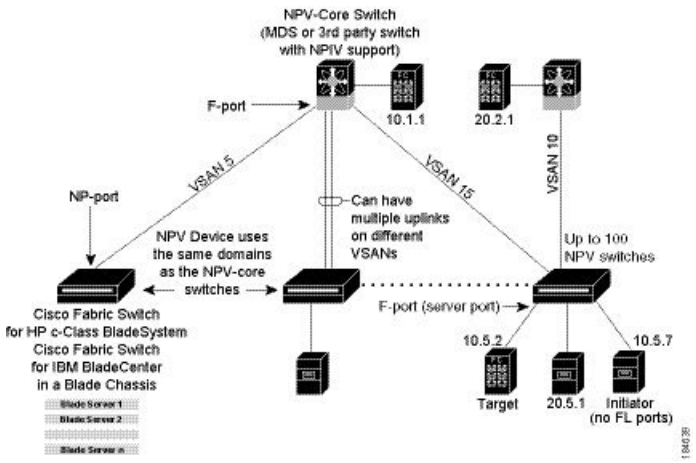

While NPV is similar to N port identifier virtualization (NPIV), it does not offer exactly the same functionality. NPIV provides a means to assign multiple FC IDs to a single N port, and allows multiple applications on the N port to use different identifiers. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of NPIV to get multiple FCIDs allocated from the core switch on the NP port.

Note For the Cisco MDS 9124, Cisco MDS 9134, and Cisco MDS 9148 legacy switches and the Cisco MDS 9148S and Cisco MDS 9396S switches supporting NPV mode, the nested NPV switches are not supported in the topology.

Figure 20: Cisco NPV Configuration-Interface View, on page 237 shows a more granular view of an NPV configuration at the interface level.

Figure 20: Cisco NPV Configuration-Interface View

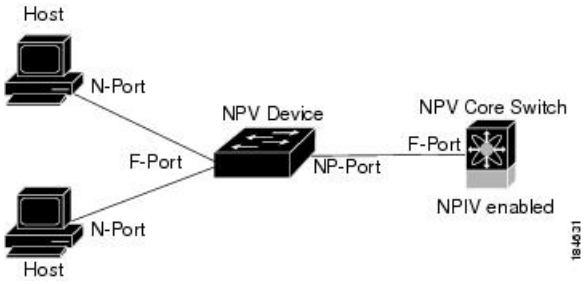

NPV Mode

A switch is in NPV mode after a user has enabled NPV and the switch has successfully rebooted. NPV mode applies to an entire switch. All end devices connected to a switch that is in NPV mode must log in as an N port to use this feature (loop-attached devices are not supported). All links from the edge switches (in NPV mode) to the NPV core switches are established as NP ports (not E ports), which are used for typical interswitch links. NPIV is used by the switches in NPV mode to log in to multiple end devices that share a link to the NPV core switch.

Note In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink to the core from the NPV device. For traffic beyond the NPV device, core switches will enforce in-order delivery if needed and/or configured.

After entering NPV mode, only the following commands are available:

Command	Description
aaa	Configure aaa functions.
banner	Configure banner message.
boot	Configure boot variables.
callhome	Enter the callhome configuration mode.
cfs	CFS configuration commands.
cli	Configure CLI commands.
clock	Configure time-of-day clock.

Command	Description
crypto	Set crypto settings.
event	Event Manager commands.
fcanalyzer	Configure cisco fabric analyzer.
feature	Command to enable/disable features.
fips	Enable/Disable FIPS mode.
flex-attach	Configure Flex Attach.
hardware	Hardware Internal Information.
hw-module	Enable/Disable OBFL information.
interface	Configure interfaces.
ip	Configure IP features.
ipv6	Configure IPv6 features.
license	Modify license features.
line	Configure a terminal line.
logging	Modify message logging facilities.
module	Configure for module.
no	Negate a command or set its defaults.
npv	Config commands for FC N_port Virtualizer.
ntp	NTP Configuration.
password	Password for the user
port-group-monitor	Configure port group monitor.
port-monitor	Configure port monitor.
power	Configure power supply.
poweroff	Power off a module in the switch.
radius	Configure RADIUS configuration.
radius-server	Configure RADIUS related parameters.
rate-mode	Configure rate mode oversubscription limit.
rmon	Remote Monitoring.
role	Configure roles.

Command	Description
snmp	Configure snmp.
snmp-server	Configure snmp server.
span	Enter SPAN configuration mode.
ssh	SSH to another system.
switchname	Configure system's network name.
system	System management commands.
terminal	Configure terminal settings.
this	Shows info about current object (mode's instance).
username	Configure user information.
vsan	Enter the vsan configuration mode.
wwn	Set secondary base MAC addr and range for additional WWNs.

NP Ports

An NP port (proxy N port) is a port on a device that is in NPV mode and connected to the NPV core switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

NP Links

An NP link is basically an NPIV uplink to a specific end device. NP links are established when the uplink to the NPV core switch comes up; the links are terminated when the uplink goes down. Once the uplink is established, the NPV switch performs an internal FLOGI to the NPV core switch, and then (if the FLOGI is successful) registers itself with the NPV core switch's name server. Subsequent FLOGIs from end devices in this NP link are converted to FDISCs. For more details refer to the Internal FLOGI Parameters, on page 239 section.

Server links are uniformly distributed across the NP links. All the end devices behind a server link will be mapped to only one NP link.

Internal FLOGI Parameters

When an NP port comes up, the NPV device first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fWWN (fabric port WWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based sWWN (switch WWN) of the NPV device used as nWWN (node WWN) in the internal FLOGI.

After completing its FLOGI request, the NPV device registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc1/4) of the NP port is embedded in the symbolic port name in the name server registration of the NPV device itself.
- The IP address of the NPV device is registered as the IP address in the name server registration of the NPV device.

Note The BB_SCN of internal FLOGIs on NP ports is always set to zero. The BB_SCN is supported at the F-port of the NPV device.

Figure 21: Internal FLOGI Flows, on page 240 shows the internal FLOGI flows between an NPV core switch and an NPV device.

Figure 21: Internal FLOGI Flows

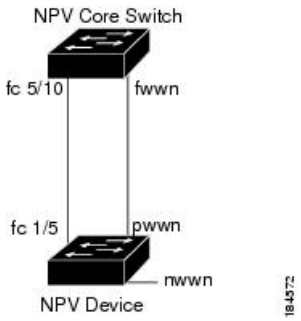

Table 30: Internal FLOGI Parameters , on page 240 identifies the internal FLOGI parameters that appear in .

Table 30: Internal FLOGI Parameters

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the NPV device.
fWWN	The fWWN of the F port on the NPV core switch.
symbolic port name	The switch name and NP port interface string. **Note** If there is no switch name available, then the output will display “switch.” For example, switch: fc1/5.
IP address	The IP address of the NPV device.
symbolic node name	The NPV switch name.

Although fWWN-based zoning is supported for NPV devices, it is not recommended because:

- Zoning is not enforced at the NPV device (rather, it is enforced on the NPV core switch).
- Multiple devices behind an NPV device log in via the same F port on the core (they use same fWWN and cannot be separated into different zones).
- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

Default Port Numbers

Port numbers on NPV-enabled switches will vary depending on the switch model. For details about port numbers for NPV-eligible switches, see the Cisco NX-OS Series Licensing Guide.

NPV CFS Distribution over IP

NPV devices use only IP as the transport medium. CFS uses multicast forwarding for CFS distribution. NPV devices do not have ISL connectivity and FC domain. To use CFS over IP, multicast forwarding has to be enabled on the Ethernet IP switches all along the network that physically connects the NPV switch. You can also manually configure the static IP peers for CFS distribution over IP on NPV-enabled switches. For more information, see the Cisco MDS 9000 Series NX-OS System Management Configuration Guide.

NPV Traffic Management

Auto

Before Cisco MDS SAN-OS Release 3.3(1a), NPV supported automatic selection of external links. When a server interface is brought up, an external interface with the minimum load is selected from the available links. There is no manual selection on the server interfaces using the external links. Also, when a new external interface was brought up, the existing load was not distributed automatically to the newly available external interface. This newly brought up interface is used only by the server interfaces that come up after this interface.

Traffic Map

As in Cisco MDS SAN-OS Release 3.3(1a) and NX-OS Release 4.1(1a), NPV supports traffic management by allowing you to select and configure the external interfaces that the server uses to connect to the core switches.

Note When the NPV traffic management is configured, the server uses only the configured external interfaces. Any other available external interface will not be used.

The NPV traffic management feature provides the following benefits:

- Facilitates traffic engineering by providing dedicated external interfaces for the servers connected to NPV.
- Uses the shortest path by selecting external interfaces per server interface.
- Uses the persistent FC ID feature by providing the same traffic path after a link break, or reboot of the NPV or core switch.
- Balances the load by allowing the user to evenly distribute the load across external interfaces.

Disruptive

Disruptive load balance works independent of automatic selection of interfaces and a configured traffic map of external interfaces. This feature forces reinitialization of the server interfaces to achieve load balance when this feature is enabled and whenever a new external interface comes up. To avoid flapping the server interfaces too often, enable this feature once and then disable it whenever the needed load balance is achieved.

If disruptive load balance is not enabled, you need to manually flap the server interface to move some of the load to a new external interface.

Multiple VSAN Support

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs on the NPV-enabled switch. The correct uplink must be selected based on the VSAN that the uplink is carrying.

Guidelines and Limitations

NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- NPV core switches must support NPIV.
- You can have up to 105 NPV devices.
- Nondisruptive upgrades are supported. See the Cisco MDS 9000 Series NX-OS Fundamentals Configuration Guide.
- Port tracking is supported. See the Cisco MDS 9000 Series NX-OS Security Configuration Guide.
- You can configure zoning for end devices that are connected to edge switches using all available member types on a core switch. However, the preferred way of zoning servers connected to any switch in NPV mode is via pWWN, device-alias, and fc alias. Multiple servers should be configured in the same zone only when using smart zoning. The smart zoning feature is available on all MDS switches. For more information, see the Smart Zoning section in the Configuring and Managing Zones chapter of the *Cisco MDS 9000 Series Fabric Configuration Guide*.
- Port security is supported on the NPV core switch for devices logged in via NPV.
- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NPV core switch links (in the same VSAN) upon initial login. If there are multiple NPV core switch links in the same VSAN, then you cannot assign a specific one to an end device.
- Both servers and targets can be connected to an NPV device.
- Remote SPAN is not supported.
- Local switching is not supported; all traffic is switched using the NPV core switch.
- NPV devices can connect to multiple NPV core switches. In other words, different NP ports can be connected to different NPV core switches.
- NPV supports NPIV-capable servers. This capability is called nested NPIV.
- Connecting two Cisco NPV switches together is not supported.
- Only F, NP, and SD ports are supported in NPV mode.
- In the case of servers that are booted over the SAN with NPV, if an NPV link failover occurs, servers will lose access to their boot LUN temporarily.
- NPV switches do not recognize the BB_SCN configuration on the xNP ports because of interoperability issues with the third-party core switches.

NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when the automatic traffic engineering by the NPV device is not sufficient for the network requirements.
- Do not configure traffic maps for all the servers. For non-configured servers, NPV will use automatic traffic engineering.
- Configure the Persistent FC ID on the core switch. Traffic engineering directs the associated server interface to external interfaces that lead to the same core switch. The server will be assigned the same FC ID for every log in. This guideline is not applicable if a 91x4 switch is used as the core switch.
- Server interfaces configured to a set of external interfaces cannot use any other available external interfaces, even if the configured interfaces are not available.
- Do not configure disruptive load balancing because this involves moving a device from one external interface to another interface. Moving the device between external interfaces requires NPV relogin to the core switch through F port leading to traffic disruption.
- Link a set of servers to a core switch by configuring the server to a set of external interfaces that are linked to the core switch.

DPVM Configuration Guidelines

When NPV is enabled, the following requirements must be met before you configure DPVM on the NPV core switch:

- You must explicitly configure the WWN of the internal FLOGI in DPVM. If DPVM is configured on the NPV core switch for an end device that is connected to the NPV device, then that end device must be configured to be in the same VSAN. Logins from a device connected to an NPV device will fail if the device is configured to be in a different VSAN. To avoid VSAN mismatches, ensure that the internal FLOGI VSAN matches the port VSAN of the NP port.
- The first login from an NP port determines the VSAN of that port. If DPVM is configured for this first login, which is the internal login of the NPV device, then the NPV core switch's VSAN F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see the Cisco MDS 9000 Series NX-OS Fabric Configuration Guide.

NPV and Port Security Configuration Guidelines

Port security is enabled on the NPV core switch on a per interface basis. To enable port security on the NPV core switch for devices logging in via NPV, you must adhere to the following requirements:

- The internal FLOGI must be in the port security database so that, the port on the NPV core switch will allow communications and links.
- All of the end device pWWNs must also be in the port security database.

Once these requirements are met, you can enable port security as you would in any other context. For details about enabling port security, see the Cisco MDS 9000 Series NX-OS Security Configuration Guide.

Connecting an NPIV-Enabled Cisco MDS Fabric Switch

This topic provides information about connecting an NPIV-enabled Cisco MDS 9396S Multilayer Fabric Switch to an NPV switch running Cisco MDS NX-OS Release 6.2(13) and earlier.

When trunking is enabled on the NPV ports of any MDS switch (released before the Cisco MDS 9396S Multilayer Fabric Switch) that runs on an Cisco MDS NX-OS Release 6.2(13) and earlier, and you connect an NPIV enabled Cisco MDS 9396S Multilayer Fabric Switch, use ports fc1/1 through fc1/63.

Note Trunking failure can occur in both non-portChannel (individual physical NP uplinks) and portChannel NP uplinks. To avoid trunking failure, ensure that you upgrade the NPV switch to Cisco MDS NX-OS Release 6.2(13) or later.

Configuring N Port Virtualization

Enabling N Port Identifier Virtualization

You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

Note All of the N port identifiers are allocated in the same VSAN.

To enable or disable NPIV on the switch, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature npiv**

Enables NPIV for all VSANs on the switch.

switch(config)# **no feature npiv**

(Optional) Disables (default) NPIV on the switch.

Configuring NPV

When you enable NPV, the system configuration is erased and the system reboots with the NPV mode enabled.

Note We recommend that you save the current configuration either on bootflash or a TFTP server before NPV (if the configuration is required for later use). Use the following commands to save either your non-NPV or NPV configuration:

switch# **copy running bootflash:filename**

The configuration can be reapplied later using the following command:

switch# **copy bootflash:filename running-config**

Note NPV cannot be enabled or disabled from the ASCII configuration file. You can enable or disable only from the command line.

To configure NPV using the CLI, perform the following steps:

Step 1 switch# **configure terminal**

On the NPV core switch, enters configuration mode.

Step 2 switch(config)# **feature npiv**

Enables NPIV mode on the NPV core switch.

switch(config)# **no feature npiv**

(Optional) Disables NPIV mode on the NPV core switch.

Step 3 switch(config)# **interface fc 2/1**

Configures the NPIV core switch port as an F port.

Changes Admin status to bring up the interfaces.

Step 4 switch(config)# **vsan database**

Configures the port VSANs for the F port on the NPIV core switch.

Step 5 switch(config)# **npv enable**

Enables NPV mode on a NPV device (module, Cisco MDS 9124, Cisco MDS 9134, Cisco MDS 9148 Fabric Switch, Cisco MDS 9148S Multilayer Fabric Switch Cisco MDS 9250i Multilayer Fabric Switch, and Cisco MDS 9396S Multilayer Fabric Switch). The module or switch is rebooted, and when it comes back up, is in NPV mode.

Note A write-erase is performed during the reboot.

Step 6 switch(config)# **interface fc 1/1**

On the NPV device, selects the interfaces that will be connected to the aggregator switch and configure them as NP ports.

Changes Admin status to bring up the interfaces.

Step 7 switch(config)# **vsan database**

Configures the port VSANs for the NP port on the NPV device.

Step 8 switch(config-if)# **exit**

Exits interface mode for the port.

Step 9 switch(config)# **interface fc 1/2 - 6**

Selects the remaining interfaces (2 through 6) on the NPV-enabled device and configures them as F ports.

Changes Admin status to bring up the interfaces.

Step 10 switch(config)# **vsan database**

Configures the port VSANs for the F ports on the NPV device.

Step 11 switch(config-npv)# **no npv enable**

Terminates session and disables NPV mode, which results in a reload of the NPV device.

Configuring NPV Traffic Management

The NPV traffic management feature is enabled after configuring NPV. Configuring NPV traffic management involves configuring a list of external interfaces to the servers, and enabling or disabling disruptive load balancing.

Configuring List of External Interfaces per Server Interface

A list of external interfaces are linked to the server interfaces when the server interface is down, or if the specified external interface list includes the external interface already in use.

To configure the list of external interfaces per server interface, perform the following tasks:

Step 1 switch# **configure terminal**

Enters configuration mode on the NPV.

Step 2 switch(config)# **npv traffic-map server-interface** *svr-if-range* **external-interface fc** *ext-fc-if-range*

Allows you to configure a list of external FC interfaces per server interface by specifying the external interfaces in the svr-if-range. The server to be linked is specified in the ext-fc-if-range.

Step 3 switch(config)# **npv traffic-map server-interface** *svr-if-range* **external-interface port-channel** *ext-pc-if-range*

Allows you to configure a list of external PortChannel interfaces per server interface by specifying the external interfaces in the svr-if-range. The server to be linked is specified in the ext-pc-if-range.

Note While mapping non-PortChannel interfaces and PortChannel interfaces to the server interfaces, include them separately in two steps.

Step 4 switch(config)# **no npv traffic-map server-interface** *svr-if-range* **external-interface** *ext-if-range*

Disables the NPV traffic management feature on the NPV.

Enabling the Global Policy for Disruptive Load Balancing

Disruptive load balancing allows you to review the load on all the external interfaces and balance the load disruptively. Disruptive load balancing is done by moving the servers using heavily loaded external interfaces, to the external interfaces running with fewer loads.

To enable or disable the global policy for disruptive load balancing, perform the following tasks:

Step 1 switch# **configure terminal**

Enters configuration mode on the NPV.

Step 2 switch(config)# **npv auto-load-balance disruptive**

Enables disruptive load balancing on the NPV core switch.

Step 3 switch (config)# **no npv auto-load-balance disruptive**

Disables disruptive load balancing on the NPV core switch.

Verifying NPV Configuration

To display NPV configuration information, perform one of the following tasks:

Command	Purpose
show fcns database	Displays all the NPV devices in all the VSANs that the aggregator switch belongs to.
show fcns database detail	Displays additional details such as IP addresses, switch names, interface names about the NPV devices.
show npv flogi-table	Displays a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs.
show npv status	Displays the status of the different servers and external interfaces.
show npv traffic-map	Displays the NPV traffic map.
show npv internal info traffic-map	Displays the NPV internal traffic details.

For detailed information about the fields in the output from these commands, refer to the Cisco MDS 9000 Series NX-OS Command Reference.

Verifying NPV

To view all the NPV devices in all the VSANs that the aggregator switch belongs to, enter the **show fcns database** command.

```
switch# show fcns database

VSAN 1:
--------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init
Total number of entries = 4
```

For additional details (such as IP addresses, switch names, interface names) about the NPV devices you see in the **show fcns database** output, enter the **show fcns database detail** command.

```
switch# show fcns database detail

------------------------
VSAN:1 FCID:0x010000
------------------------
port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
```

```
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/1
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:01:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)
connected interface         :port-channel6
switch name (IP address)    :switch (192.0.2.1)
------------------------
VSAN:1 FCID:0x010001
------------------------
port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/2
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:02:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
connected interface         :port-channel6
switch name (IP address)    :switch (192.0.2.1)
```

If you need to contact support, enter the **show tech-support NPV** command and save the output so that support can use it to troubleshoot, if necessary.

To display a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, enter the **show npv flogi-table** command.

```
switch# show npv flogi-table
```

SERVER	VSAN	FCID
INTERFACE		
---		
fc1/19	1	0xee0
fc1/19	1	0xee0
fc1/19	1	0xee0
fc1/19	1	0xee0

```
Total number of flogi = 4.
```

To display the status of the different servers and external interfaces, enter the **show npv status** command.

```
switch# show npv status

npiv is enabled

External Interfaces:
====================
  Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
  Interface: fc1/2, VSAN: 3, FCID: 0x040000, State: Up

  Number of External Interfaces: 2

Server Interfaces:
==================
  Interface: fc1/7, VSAN: 2, NPIV: No, State: Up
  Interface: fc1/8, VSAN: 3, NPIV: No, State: Up

  Number of Server Interfaces: 2
```

Verifying NPV Traffic Management

To display the NPV traffic map, enter the **show npv traffic-map** command.

```
switch# show npv traffic-map

NPV Traffic Map Information:
----------------------------
Server-If    External-If(s)
----------------------------
fc1/1        fc1/5
----------------------------
```

To display the NPV internal traffic details, enter the **show npv internal info traffic-map** command.

```
switch# show npv internal info traffic-map

NPV Traffic Map Information:
------------------------------------------------------------
Server-If    Last Change Time               External-If(s)
------------------------------------------------------------
fc1/1        2015-01-15 03:24:16.247856       fc1/5
------------------------------------------------------------
```

Configuring FlexAttach Virtual pWWN

This chapter provides information about FlexAttach virtual pWWN and how to configure FlexAttach virtual pWWN.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About FlexAttach Virtual pWWN

FlexAttach Virtual pWWN

FlexAttach virtual pWWN feature facilitates server and configuration management. In a SAN environment, the server installation or replacement, requires interaction and coordination among the SAN and server administrators. For coordination, it is important that the SAN configuration does not change when a new server is installed, or when an existing server is replaced. FlexAttach virtual pWWN minimizes the interaction between the server administrator and the SAN administrator by abstracting the real pWWN using virtual pWWNs.

When FlexAttach virtual pWWN is enabled on an interface, a virtual pWWN is assigned to the server interface. The real pWWN is replaced by a virtual pWWN, which is used for a SAN configuration such as zoning.

Server administrators can benefit from FlexAttach in the following scenarios:

- Pre-configure—Pre-configure SAN for new servers that are not available physically yet. For example, they may be on order. FlexAttach can be enabled on the ports designated for the new servers and use the virtual WWNs assigned for configuring SAN. The new servers are then plugged into the fabric without any change needed in the SAN.
- Replacement to the same port—A failed server can be replaced onto the same port without changing the SAN. The new server gets a same pWWN as the failed server because the virtual pWWN is assigned to the port.
- Replacement to (spare)—A spare server, which is on the same NPV device or a different NPV device) can be brought online without changes to the SAN. This action is achieved by moving the virtual port WWN from the current server port to the spare port.
- Server Mobility—A server can be moved to another port on the same NPV device or another NPV device without changing the SAN. This is accomplished by moving the virtual pWWN to the new port. No change is needed if FlexAttach was configured using the physical port WWN of the server to the virtual port WWN mapping.

Difference Between San Device Virtualization and FlexAttach Port Virtualization

Table describes the difference between SAN device virtualization (SDV) and FlexAttach port virtualization.

Table 31: Difference Between SDV and FlexAttach Virtualization

SAN Device Virtualization (SDV)	FlexAttach Virtualization
Facilitates target and disk management, and only facilitates disk and data migration.	Facilitates server management and has no restriction on the end devices used.
WWN NAT and Fibre Channel ID (FC-ID) are allocated on the virtual device, both primary and secondary.	WWN and Network Address Transport (NAT) is allocated to host bus adapter (HBA).

SAN Device Virtualization (SDV)	FlexAttach Virtualization
FC-ID rewrite on the switch indicates a rewrite-capable switch on the path.	No rewrite requirements.
Configuration is distributed. This allows programming rewrites and connectivity anywhere.	Configuration distribution is not required for any of the interface-based configurations.
Configuration is secured to device alias.	Does not require device alias for virtual pWWN.
Does not allow automapping to the secondary device.	Allows automapping to the new HBA. Mapping process is manual for NPIV.

FlexAttach Virtual pWWN CFS Distribution

The FlexAttach virtual pWWN configuration is distributed for CFS through IPv4, and is enabled by default. The FlexAttach virtual pWWN distribution, by default, is on CFS region 201. The CFS region 201 links only to the NPV-enabled switches. Other CFS features such as syslog is on region 0. Region 0 will be linked through IPv4 for all NPV switches on the same physical fabric. If CFS has an option to link through IPv4 or ISL, then CFS will select the ISL path.

Note NPV switches do not have ISL (E or TE ports) and are linked through IPv4.

Security Settings for FlexAttach Virtual pWWN

Security settings for the FlexAttach virtual pWWN feature are done by port security at the NPV core. Node WWN of the end device is used to provide physical security.

For more details on enabling port security, refer to the Cisco MDS 9000 Series NX-OS Security Configuration Guide.

Guidelines and Limitations

Following are recommended guidelines and requirements when deploying FlexAttach virtual pWWN:

- FlexAttach configuration is supported only on NPV switches.
- Cisco Fabric Services (CFS) IP version 4 (IPv4) distribution should be enabled.
- Virtual WWNs should be unique across the fabric.

Configuring FlexAttach Virtual pWWN

Automatically Assigning FlexAttach Virtual pWWN

Automatic assignment of virtual pWWN can be configured on an NPV switch globally, per VSAN, or per port. When assigned automatically, a virtual WWN is generated from the device local switch WWN.

To assign a virtual pWWN automatically, perform this task:

Before you begin

The port must be in a shut state when the virtual pWWN is enabled.

Step 1 Enter configuration mode:

Example:

```
switch# configure terminal
```

Step 2 Assign FlexAttach virtual pWWN automatically for the interfaces:

Example:

```
switch(config)# flex-attach virtual-pwwn auto [interface interface-list]
```

To assign FlexAttach virtual pWWN automatically for the VSANs:

```
switch# (config)# flex-attach virtual-pwwn auto [vsan vsan-range]
```

Step 3 Commit the configuration:

Example:

```
switch(config)# flex-attach commit
```

Manually Assigning FlexAttach Virtual pWWN

Restrictions

The interface mentioned in the interface value must be in a shut state.

To assign virtual pWWN manually, perform this task:

Before you begin

- Some ports may be in automode, some in manual mode, and the virtual pWWNs need not be assigned.

- The port must be in a shut state when a virtual pWWN is assigned.

Step 1 Enter configuration mode:

Example:

```
switch# configure terminal
```

Step 2 Configure the FlexAttach virtual pWWN for the interface:

Example:

```
switch(config)# flex-attach virtual-pwwn vpwwn interface interface
```

(Optional) Configure the FlexAttach virtual pWWN for the interface in the VSAN:

```
switch(config)# flex-attach virtual-pwwn vpwwn interface interface [ vsan vsan]
```

Step 3 Commit the configuration:

```
switch(config)# flex-attach commit
```

Mapping pWWN to Virtual pWWN

You can configure virtual pWWNs through real pWWNs. This process is required for NPIV hosts containing multiple pWWNs, of which only FLOGI is mapped to the virtual pWWN. Subsequent FDSIDs will have different mappings.

Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch across the NPV switches. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.

Restrictions

- The specified virtual pWWN and the real pWWN must not be logged in.
- To map pWWN to virtual pWWN, perform this task:

Before you begin

The interface must be in a shut state and the specified virtual pWWN should not be logged in.

Step 1 Enter configuration mode:

Example:

```
switch# configure terminal
```

Step 2 Map the pWWN to the virtual pWWN:

Example:

```
switch(config)# flex-attach virtual-pwwn vpwwn pwwn pwwn
```

Step 3 Commit the configuration:

```
switch(config)# flex-attach commit
```

Verifying FlexAttach Virtual pWWN Configuration

To display FlexAttach configuration information, perform one of the following tasks:

Command	Purpose
show flex-attach virtual-pwwn	Displays the type and value of virtual pWWNs.
show fcns database	Displays if the end device is logged with the correct virtual WWNs.

For detailed information about the fields in the output from these commands, refer to the Cisco MDS 9000 Series NX-OS Command Reference.

To view and confirm that the type and value of virtual pWWNs are correct, enter the **show flex-attach virtual-pwwn** command.

Example: Displaying the Type and Value of Virtual pWWNs

```
switch# show flex-attach virtual-pwwn
VIRTUAL PORT WWNS ASSIGNED TO INTERFACES
------------------------------------------------------------------------
VSAN INTERFACE VIRTUAL-PWWN AUTO LAST-CHANGE
------------------------------------------------------------------------
1 fc1/1 00:00:00:00:00:00:00:00
1 fc1/2 22:73:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/3 22:5e:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/4 22:5f:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/5 22:74:00:05:30:01:6e:1e TRUE Thu Jan 31 01:26:24 2008
1 fc1/6 22:60:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/7 22:61:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/8 22:62:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/9 22:63:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/10 22:64:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/11 22:65:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
1 fc1/12 22:66:00:05:30:01:6e:1e TRUE Thu Jan 31 01:58:52 2008
```

Verifying the End Device

To verify that the end device is logged with the correct virtual WWNs, use the **show fcns database** command on the NPV core.

Example: Verifying the End Device

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init
Total number of entries = 4
```

Monitoring FlexAttach Virtual pWWN

Table lists the errors that might be displayed and provides the workarounds.

Table 32: FlexAttach Errors and Workarounds

Error	Description	Workaround
fc1/1 : interface is not down	FlexAttach configuration fails because the configuration is enabled for an active interface with the operation state as up.	To move the port to the shut state, enable the FlexAttach configuration, and then move the port to no shut state.
FlexAttach configuration is not distributed to the peers	The FlexAttach configuration on one peer NPV is not available to any other peer NPV.	FlexAttach configuration will not be distributed if **cfs ipv4 distribute**, or **cfs ipv6 distribute** is disabled. Enable **cfs ipv4 distribute**, or **cfs ipv6 distribute**.
Even with CFS distribution enabled Inagua does not become a peer with other NPVs	CFS over IP is enabled, and the Inagua in one blade center is not the peer NPV for other NPVs.	CFS over IP uses IP multicast to discover the NPV peers in the network. IBM MM does not support multicast and cannot act as a peer with NPV. This prevents the FlexAttach configuration from getting distributed to other peer NPVs in the network.
NP port uses physical pWWN instead of virtual pWWN configured through FlexAttach	This occurs when NP port uses physical pWWN instead of virtual pWWN, that is configured through FlexAttach.	FlexAttach is supported on server interfaces such as F ports, and not on external interfaces such as NP ports.
real port WWN and virtual WWN cannot be same	This occurs when you try to configure FlexAttach with a similar value for pWWN and virtual pWWN.	Use different values for pWWN and virtual pWWn, as similar values for pWWN and virtual pWWn are not allowed.
Virtual port WWN already exists	This occurs when you try to configure an already defined pWWN to a different interface.	Use an undefined virtual pWWN for a new interface.

Configuring Port Tracking

This chapter provides information about port tracking and how to configure port tracking.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Information About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information.

In Figure 22: Traffic Recovery Using Port Tracking, on page 265, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 22: Traffic Recovery Using Port Tracking

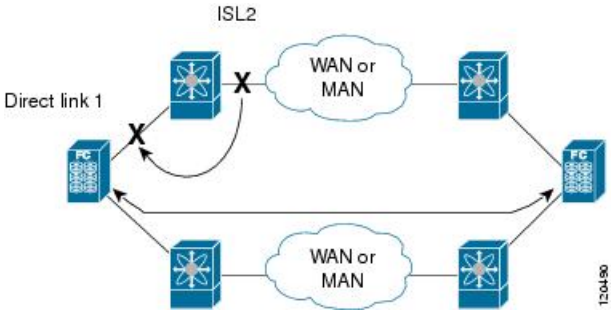

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco NX-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports.
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Guidelines and Limitations

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.
- Be aware that the linked port is automatically brought down when the tracked port goes down. Be aware that the linked port is automatically brought down when the tracked port goes down.

Default Settings

Table 33: Default Port Tracking Parameters , on page 267 lists the default settings for port tracking parameters.

Table 33: Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled.
Operational binding	Enabled along with port tracking.

Configuring Port Tracking

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco MDS 9000 Series Multilayer Switches. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature port-track**

Enables port tracking.

switch(config)# **no feature port-track**

(Optional) Removes the currently applied port tracking configuration and disables port tracking.

Information About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Binding a Tracked Port Operationally

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc8/6**

Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.

Note This link symbolizes the direct link (1) in .

Step 3 switch(config-if)# **port-track interface port-channel 1**

Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down.

Note This link symbolizes the ISL (2) in .

switch(config-if)# **no port-track interface port-channel 1**

(Optional) Removes the port tracking configuration that is currently applied to interface fc8/6.

Information About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In Figure 23: Traffic Recovery Using Port Tracking, on page 269, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 23: Traffic Recovery Using Port Tracking

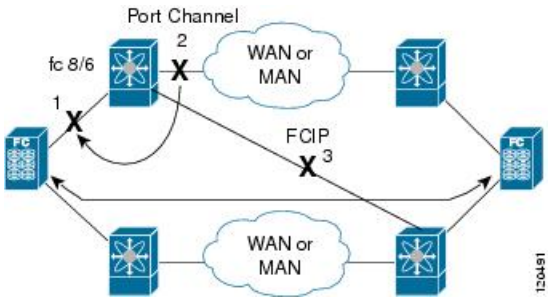

Tracking Multiple Ports

To track multiple ports, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc8/6**

Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.

Note This link symbolizes the direct link (1) in Figure 23: Traffic Recovery Using Port Tracking, on page 269.

Step 3 switch(config-if)# **port-track interface port-channel 1**

Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down.

Note This link symbolizes the ISL (2) in Figure 23: Traffic Recovery Using Port Tracking, on page 269.

Step 4 switch(config-if)# **port-track interface fcip 5**

Tracks interface fc8/6 with interface fcip 5. When FCIP 5 goes down, interface fc8/6 is also brought down.

Note This link symbolizes the ISL (3) in Figure 23: Traffic Recovery Using Port Tracking, on page 269.

Information About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

Tip The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc8/6**

Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.

Step 3 switch(config-if)# **port-track interface port-channel 1 vsan 2**

Enables tracking of the PortChannel in VSAN 2.

switch(config-if)# **no port-track interface port-channel 1 vsan 2**

(Optional) Removes the VSAN association for the linked port. The PortChannel link remains in effect.

Information AboutForceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

To forcefully shut down a tracked port, perform these steps:

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/5**

Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.

Step 3 switch(config-if)# **port-track force-shut**

Forcefully shuts down the tracked port.

switch(config-if)# **no port-track force-shut**

(Optional) Removes the port shutdown configuration for the tracked port.

Verifying Port Tracking Configuration

The **show** commands display the current port tracking settings for the Cisco MDS switch (see Examples Displays the Linked and Tracked Port Configuration, on page 272 to Displays a Forced Shutdown Configuration , on page 273).

Displays the Linked and Tracked Port Configuration

```
switch# show interface
...
fc8/6 is down (All tracked ports down
)     <-----------------------------------Linked port
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 21:c6:00:05:30:00:37:1e
    Admin port mode is auto, trunk mode is on
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface port-channel 1 vsan 2 (trunking) <-----Tracked port
Port tracked with interface fcip 5 <---------------------------------------Tracked port
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      269946 frames input, 22335204 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      205007 frames output, 10250904 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      2 output OLS, 2 LRR, 0 NOS, 1 loop inits
      0 receive B2B credit remaining
      0 transmit B2B credit remaining
...
```

Displays a Tracked Port Configuration for a Fibre Channel Interface

```
switch# show interface fc1/1
fc1/1 is down (Administratively down)
    Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
    Port WWN is 20:01:00:05:30:00:0d:de
    Admin port mode is FX
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface fc1/2 (down)
Port tracked with interface port-channel 1 vsan 2 (down)
Port tracked with interface fcip1 (down)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1 frames input, 128 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      1 frames output, 128 bytes
        0 discards, 0 errors
```

```
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 receive B2B credit remaining
    0 transmit B2B credit remaining
```

Displays a Tracked Port Configuration for a PortChannel Interface

```
switch# show interface port-channel 1
port-channel 1 is down (No operational members)
    Hardware is Fibre Channel
    Port WWN is 24:01:00:05:30:00:0d:de
    Admin port mode is auto, trunk mode is on
    Port vsan is 2
    Linked to 1 port(s)
      Port linked to interface fc1/1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    No members
```

Displays a Forced Shutdown Configuration

```
switch# show interface fc 1/5
fc1/5 is up
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:05:00:05:30:00:47:9e
    Admin port mode is F
    Port mode is F, FCID is 0x710005
    Port vsan is 1
    Speed is 1 Gbps
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Port track mode is force_shut <--this port remains shut even if the tracked port is
back up
```