

Cisco MDS 9000 Series Interfaces Configuration Guide, Release 9.x

First Published: 2022-09-02

Last Modified: 2026-03-05

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2026 Cisco Systems, Inc. All rights reserved.



Preface

- Preface, on page iii

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 sections:

Audience

This publication is for network administrators who install, configure, and maintain Cisco MDS 9000 Series Switches.

Document Conventions

Command descriptions use these conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.

Convention	Description
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use these conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

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.

Warning IMPORTANT SAFETY INSTRUCTIONS

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.

SAVE THESE INSTRUCTIONS

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to mds-docfeedback@cisco.com. We appreciate your feedback.

Related Documentation

The entire Cisco MDS 9000 Series switches documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/series.html

Documentation Roadmap

https://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/rel90.html

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



New and Changed Information

- Change Summary, on page 2

Change Summary

The following tables summarize 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://bst.cloudapps.cisco.com/bugsearch and the release notes for your software release.

Table 1: Feature History for Congestion Management

Feature Name	Release	Description	Where Documented
Transmit FC abort counter	9.4(5)	Support has been introduced for per-interface transmit FC Abort Sequence (ABTS) frame counters. This allows monitoring of health and performance of Cisco MDS FC fabrics.	Configuring Fibre Channel Interfaces
Link Negotiation Early Training Complete	9.4(5)	Support has been introduced for enhanced Fibre Channel link negotiation capability for simultaneous TC signaling and LSN completion during Fibre Channel link negotiation at 32 Gbps to interoperate with non-compliant devices.	Configuring Fibre Channel Interfaces
RDF indicates disabled FPM feature during device registration	9.4(5)	RDF indicates FPM feature not enabled while notifying registered devices.	Congestion Management
Test Smart Monitoring and Alerting	9.4(5)	Support has been introduced for generating test syslogs and SNMP traps for SMA. This allows testing of remote automation that processes these alerts.	Congestion Management
Congestion Isolation for over-utilization action	9.4(5)	Actions for congestion isolation for over-utilization are summarized.	Congestion Management
PMON Over-Util support for NPV mode	9.4(5)	Support has been introduced for tx-overutilization counter in NPV mode.	About PMON Transmit Overutilization Counter, on page 220

Feature Name	Release	Description	Where Documented
FPIN without PLOGI	9.4(5)	Support has been introduced to send FPIN notifications to the target port by the Fabric Controller even if PLOGI (Port Login) has not occurred.	Congestion Management
Smart Monitoring and Alerting	9.4(4)	Smart Monitoring and Alertng (SMA) feature is available in the production environment. This feature provides congestion alerting and real time monitoring.	Congestion Management
Fabric Congestion and Diagnostics	9.4(4)	Support to send on demand RDF requests directly to connected HBAs has been added.	Congestion Management
Over-Utilization Counter in PMON	9.4(4)	Support for the following are added: • FPIN peer congestion notification • FPIN warning level	About PMON Transmit Overutilization Counter, on page 220
Over-Utilization Counter in PMON	9.4(3)	Support for overutilization counter and its associated mechanisms are introduced to enhance network performance by providing better detection and management of congestion.	About PMON Transmit Overutilization Counter, on page 220
system transceiver power-control command to automatically power off the transceiver if interface is in admin shutdown.	9.4(2)	This command was introduced.	Configuring Transceiver (SFP) Power Saving , on page 133
RxWait feature for the 64G module and 64G Fabric switches	9.3(2)	RxWait is now displayed in the following commands for interfaces on 64G capable modules/switches: • show interface • show interface counters detailed • show interface rxwait-history • show logging onboard rxwait	Congestion Management

Feature Name	Release	Description	Where Documented
DIRL Supported Devices in Switch and NPV modes	9.3(2)	Updated list of devices that support DIRL in NPV mode for Release 9.3(2).	Congestion Management
HBA Extended Receiver Ready	9.3(1)	Added support on F and NP ports. HBA ER_RDY is in preview (beta) status and not to be used in the production environment.	Congestion Management, on page 177
DIRL NPV Support	9.3(1)	Enhanced to support the switches in NPV mode.	Congestion Management, on page 177
Fabric Notifications	9.2(1)	The Fabric Notification — FPIN and Congestion Signal feature are out of the preview (beta) status and is used in the production environment.	Congestion Management, on page 177
TxWait OBFL	9.2(1)	The TxWait OBFL file size was increased from 512 KB to 8 MB.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Congestion Isolation	8.5(1)	This feature is now handled by Fabric Performance Monitor (FPM). The following commands were introduced: • **feature fpm** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* The following commands were deprecated: • **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**} • **congestion-isolation remove interface** *slot*/*port*	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Congestion Isolation Recovery	8.5(1)	The Congestion Isolation Recovery feature automatically recovers the flow which was moved to low-priority VL after it was detected as slow back to normal VL; thereby, recovering the flow. The following commands were introduced: • **feature fpm** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* • **port-monitor cong-isolation-recover** {**recovery-interval** *seconds* | **isolate-duration** *hours* **num-occurrence** *number*} The **counter** port monitor command was modified to add the **cong-isolate-recover** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Fabric Notifications	8.5(1)	Fabric Notifications are used to notify end devices of performance impacting conditions and behaviors that affect the normal flow of IO such as link integrity degradation and congestion. The following commands were introduced: • **feature fpm** • **counter txwait warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* • **fpm fpin period** *seconds* • **fpm congestion-signal period** *seconds* • **show fpm** {**fpin** | **registration** {**congestion-signal** | **summary**} | **congested-device database** [**exclude** | **local** | **remote** | **static**]} **vsan** *id* • **port-monitor fpin** {**recovery-interval** *seconds* | **isolate-duration** *hours* **num-occurrence** *number*} The **counter** port monitor command was modified to add the **FPIN** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Dynamic Ingress Rate Limiting (DIRL)	8.5(1)	DIRL is used to automatically limit the amount of traffic that is flowing through a switch port that is congested. The following commands were introduced: • **feature fpm** • **fpm dirl** {**exclude list** | **reduction** *percentage* **recovery** *percentage* • **member** {**fc4-feature target** | **interface fc** *slot/port*} • **fpm dirl recover interface fc** *slot/port* • **show fpm** {**dirl exclude** | **fpin vsan** *id* | **ingress-rate-limit** {**events** | **status**} **interface fc***slot/port*} • **port-monitor dirl recovery-interval** *seconds* The **counter** port monitor command was modified to add the **DIRL** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Fibre Channel and Fibre Channel over Ethernet (FCoE)	8.4(1)	The following commands were modified: • The **show hardware internal rxwait-history** [**module** *number* | **port** *number*] command was changed to **show interface** [*interface-range*] **rxwait-history**. • The **show hardware internal txwait-history** [**module** *number* | **port** *number*] command was changed to **show interface** [*interface-range*] **txwait-history**. • The **show process creditmon txwait-history** [**module** *number* [**port** *number*]] command was changed to **show interface** [*interface-range*] **txwait-history**. The following command outputs were modified: • **show interface** *interface-range* **aggregate-counters** • **show interface** *interface-range* **counters** • **show interface** *interface-range* **counters detailed** • **show interface priority-flow-control** • **show interface vfc** *interface-range* **counters detailed**	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
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.	Congestion Management, on page 177
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 Management, on page 177

Feature Name	Release	Description	Where Documented
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**	Congestion Management, on page 177
Congestion Drop Timeout, No-Credit Frame Timeout, and Slow-Port Monitor Timeout Values for Fibre Channel	8.1(1)	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.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Slow Drain Detection and Congestion Isolation	8.1(1)	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 **2.** Congestion Isolation **3.** Port monitor portguard action for Congestion Isolation	Congestion Management, on page 177

Table 2: Feature History for Interfaces

Feature Name	Release	Description	Where Documented
Interfaces and Port Channels			
Port Beaconing	8.4(1)	This feature is supported on Cisco MDS switches that are operating in Cisco N-Port Virtualizer (Cisco NPV) mode.	Configuring Interfaces, on page 62
Port Beaconing	8.3(1)	This feature can be used to identify individual switch and directly attached peer ports in a data center environment.	Configuring Interfaces, on page 62
Port Beaconing	8.4(1)	This feature is supported on Cisco MDS switches that are operating in Cisco NPV mode.	Configuring Interfaces, on page 62
Port Monitor	8.4(1)	Added support to configure a logging severity level for port monitor syslog messages.	Configuring Interfaces, on page 62

Feature Name	Release	Description	Where Documented
Interfaces	8.4(1)	Fixed the output formatting of the **show logging onboard txwait** command.	Configuring Interfaces, on page 62
Port Beaconing	8.3(1)	This feature can be used to identify individual switch and directly attached peer ports in a data center environment.	Configuring Interfaces, on page 62
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.	Configuring Interfaces, on page 62
Port Monitor			
Port Monitor Policy	8.5(1)	A new port monitor portguard action (cong-isolate-recover) was introduced for the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters.	
Port Monitor	8.1(1)	The following commands were modified: • **port-type** {**access-port** | **trunks** | **all**} • **logical-type** {**core** | **edge** | **all**}	

Feature Name	Release	Description	Where Documented
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.	

Table 3: Feature History for Interface Buffers

Feature Name	Release	Description	Where Documented
Buffer-to-Buffer Credit Recovery	8.4(1)	Support for buffer-to-buffer credit recovery for NP ports.	
Buffer-to-Buffer Credit Recovery	8.2(1)	Support for buffer-to-buffer credit recovery for F ports.	
Enhanced Receiver Ready	8.1(1)	This feature was introduced. The following commands were introduced: • show flow-control er_rdy • switchport vl-credit • system fc flow-control er_rdy	

Table 4: Feature History for Port Channels

Feature Name	Release	Description	Where Documented
Port channels	8.4(1)	The default port channel mode is changed from On to Active mode.	

Table 5: Feature History for N Port Identifier Virtualization

Feature Name	Release	Description	Where Documented
N Port Virtualization (NPV) Load Balancing	8.5(1)	NPV load balancing scheme is enhanced to propose a mapping of server interfaces to external interfaces based on the throughput value so that the traffic can be evenly distributed on the external interfaces. The following commands were introduced: • **show npv traffic-map proposed** • **npv traffic-map analysis clear**	
N Port Identifier Virtualization	8.4(2)	The NPIV feature is enabled by default.	
NP Ports	8.4(1)	Buffer-to-Buffer State Change Notification (BBSCN) allowed on NP Ports	



Interface Overview

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

- Finding Feature Information, on page 18
- Trunks and Port Channels, on page 19
- Fibre Channel Port Rate Limiting, on page 20
- Maximum NPIV Limit, on page 21
- Extended Credits, on page 22
- N Port Virtualization, on page 23
- FlexAttach, on page 24

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://bst.cloudapps.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 Port Channels

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. E and F ports support trunking.

Port channels 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 port channel. ISL ports can reside on any switching module, and they do not need a designated primary port. If a port or a switching module fails, the port channel continues to function properly without requiring fabric reconfiguration.

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

Port channels load balance Fibre Channel traffic using a hash of source FC-ID and destination FC-ID, and optionally the exchange ID. Load balancing using port channels is performed over both Fibre Channel and FCIP links. Cisco NX-OS software can also 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, on page 26
- Feature History for Interfaces, on page 27
- Information About Interfaces, on page 29
- Prerequisites for Interfaces, on page 56
- Guidelines and Limitations, on page 57
- Default Settings, on page 61
- Configuring Interfaces, on page 62
- Verifying Interface Configuration, on page 85
- Transmit-Wait History Graph, on page 117

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://bst.cloudapps.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 6: New and Changed Features, on page 27 lists the New and Changed features.

Table 6: New and Changed Features

Feature Name	Release	Feature Information
Interfaces and Port Channels		
Display SFP power control and status	9.4(2)	Added FC SFP power control and status
Display SFP descriptions and Parameters	9.4(2)	Added support to display SFP descriptions and Parameters for Fibre Channel ports
Port Beaconing	8.4(1)	This feature is supported on Cisco MDS switches that are operating in Cisco NPV mode.
Port Monitor	8.4(1)	Added support to configure a logging severity level for port monitor syslog messages.
Interfaces	8.4(1)	Fixed the output formatting of the **show logging onboard txwait** command.
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		

Feature Name	Release	Feature Information
Port Monitor Policy	8.5(1)	A new port monitor portguard action (cong-isolate-recover) was introduced for the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters. The *cong-isolate-recover* 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**
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**}
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.

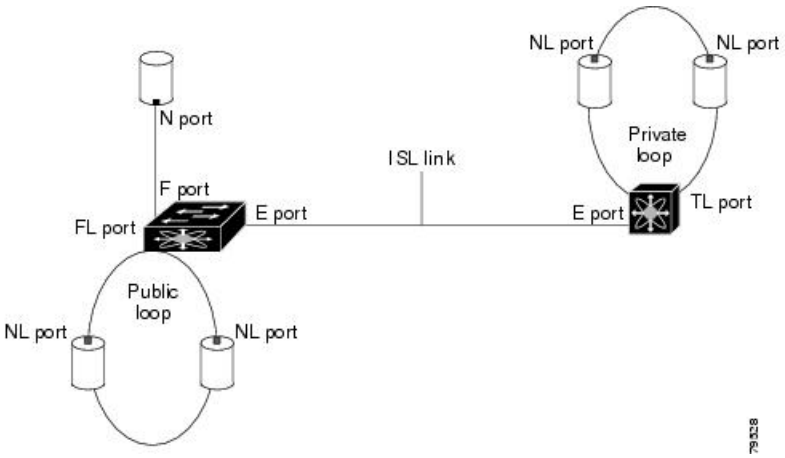
Interface Description

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 29). 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 Port Channels, on page 337.

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 373 .

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 317. 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 317. 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 317.

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 7: Administrative States , on page 32.

Table 7: 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 8: Operational States , on page 32.

Table 8: 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 9: Reason Codes for Interface States , on page 33.

Table 9: 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 10: Reason Codes for Nonoperational States , on page 34. Note that only some of the reason codes are listed in Table 10: Reason Codes for Nonoperational States , on page 34.

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 10: Reason Codes for Nonoperational States , on page 34.

Table 10: 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.

Debounce Timer

Debounce timers delay the notification of link changes that can decrease traffic loss due to a network reconfiguration.

There are two types of debounce timers:

- Sync Loss: This timer applies when a link is active. A link is active after the link initialization (LR-LRR-IDLE-IDLE) is successful. If there is synchronization loss for less than 100 ms when the Fibre Channel link is active, the interface does not bounce, but remains active. The value for debounce timer link down due to synchronization loss is 100 ms for Fibre Channel interfaces. This value cannot be configured. If there is synchronization loss for 100 ms or more when the Fibre Channel link is active, the interface goes down with the following message:

```
%PORT-5-IF_DOWN_LINK_FAILURE: %$VSAN vsan%$ Interface intf is down (Link failure loss
of sync)
```

- NOS/OLS: This timer applies when a Fibre Channel port is initializing prior to when it is active. A Fibre Channel port is initializing prior to FLOGI or ACC (FLOGI) for F ports and ELP or ACC (ELP) for E ports. During the port initialization if a Fibre Channel interface encounters multiple NOS/OLS sequences continuously for a threshold of 10 times in 2 seconds, the interface is going to be moved to the *errDisabled* state with the following message:

```
%PORT-5-IF_DOWN_LINK_FAILURE: %$VSAN vsan%$ Interface intf is down (Link failure due
to NOS/OLS debounce timeout)
```

  The value for NOS/OLS debounce timer is 2 seconds and not configurable.

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:

Procedure

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 11: SFP Transmitter Acronym Definitions , on page 39 defines the acronyms used in the command output. For information about how to display interface information, see the #unique_65.

Table 11: SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
Short wave laser	swl
Medium wave laser	mwl
Extended reach wave laser	erwl
Long wave laser	lwl
Long wave laser cost reduced	lwcr
Electrical	elec

Port Monitor

The Port Monitor feature can be used to monitor the performance and status of ports and generate alerts and syslog messages 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 counter value crosses these threshold values.

Table 12: Default Port Monitor Policy with Threshold Values for Releases Prior to Cisco MDS NX-OS Release 8.5(1), on page 40 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 From Cisco MDS NX-OS Release 8.3(1), NP ports are also monitored in port monitor.

Table 12: Default Port Monitor Policy with Threshold Values for Releases Prior to Cisco MDS NX-OS Release 8.5(1)

Counter	Threshold Type	Interval (Seconds)
link-loss	Delta	60
sync-loss	Delta	60
signal-loss	Delta	60
state-change	Delta	60
invalid-words	Delta	60
invalid-crc	Delta	60
tx-discards	Delta	60
lr-rx	Delta	60
lr-tx	Delta	60
timeout-discards	Delta	60
credit-loss-reco	Delta	60
tx-credit-not-available	Delta	1
rx-datarate	Delta	60
tx-datarate	Delta	60
tx-slowport-oper-delay [2]	Absolute	60
txwait[3]	Delta	60

[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.

If the tx-credit-not-available timer and the port monitor timer do not start at the same time or if the difference between the tx-credit-not-available timer and the port monitor timer is not zero, there will be a spike of rising and falling alarms from port monitor.

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 13: Default Port Monitor Policy with Threshold Values for Cisco MDS NX-OS Release 8.5(1) and Later Releases

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
link-loss	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
sync-loss	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
signal-loss	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
invalid-words	Delta	60	none	n/a	1	0	4	syslog, rmon	none	n/a	n/a
invalid-crcs	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
state-change	Delta	60	none	n/a	5	0	4	syslog, rmon	none	n/a	n/a
tx-discards	Delta	60	none	n/a	200	10	4	syslog, rmon	none	n/a	n/a
lr-rx	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
lr-tx	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a
timeout-discards	Delta	60	none	n/a	200	10	4	syslog, rmon	none	n/a	n/a
credit-loss-reco	Delta	60	none	n/a	1	0	4	syslog, rmon	none	n/a	n/a
tx-credit-not-available	Delta	60	none	n/a	10% [4]	0%	4	syslog, rmon	none	n/a	n/a

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
rx-datarate	Delta	10	none	n/a	80%	70%	4	syslog, rmon	none	n/a	n/a
tx-datarate	Delta	10	none	n/a	80%	70%	4	syslog, rmon	none	n/a	n/a
tx-slowport-oper-delay [5]	Absolute	60	none	n/a	50ms	0ms	4	syslog, rmon	none	n/a	n/a
txwait [6]	Delta	60	none	n/a	30%	10%	4	syslog, rmon	none	n/a	n/a
rx-datarate-burst	Delta	10	none	n/a	5@90%	1@90%	4	syslog, rmon, obfl	none	n/a	n/a
tx-datarate-burst	Delta	10	none	n/a	5@90%	1@90%	4	syslog, rmon, obfl	none	n/a	n/a
input-errors	Delta	60	none	n/a	5	1	4	syslog, rmon	none	n/a	n/a

[4] 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.

If the tx-credit-not-available timer and the port monitor timer do not start at the same time or if the difference between the tx-credit-not-available timer and the port monitor timer is not zero, there will be a spike of rising and falling alarms from port monitor.

[5]
- 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).

[6]
- 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 14: Recommended Units for Port Monitor Policy For Releases Prior to Cisco MDS NX-OS Release 8.5(1)

Counter	Threshold Type	Interval (Seconds)	Rising Threshold	Event	Falling Threshold	Event	Warning Threshold
link-loss	Delta	Seconds	Number	Event ID	Number	Event ID	Number

Counter	Threshold Type	Interval (Seconds)	Rising Threshold	Event	Falling Threshold	Event	Warning Threshold
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's	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
err-pkt-to-xbar	Delta	Seconds	Number	Event ID	Number	Event ID	Number
err-pkt-from-xbar	Delta	Seconds	Number	Event ID	Number	Event ID	Number

Table 15: Recommended Units for Port Monitor Policy For Cisco MDS NX-OS Release 8.5(1) and Later Releases

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
link-loss	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
sync-loss	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
signal-loss	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
invalid-words	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
invalid-crcs	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
state-change	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
tx-discards	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
lr-rx	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
lr-tx	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Delta	Seconds	Percentage	syslog, rmon	Percentage	Percentage	Event ID	syslog, rmon	none	n/a	n/a
rx-datarate	Delta	Seconds	Percentage	syslog, rmon	Percentage	Percentage	Event ID	syslog, rmon	none	n/a	n/a
tx-datarate	Delta	Seconds	Percentage	syslog, rmon	Percentage	Percentage	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Absolute	Seconds	Milliseconds	syslog, rmon	Milliseconds	Milliseconds	Event ID	syslog, rmon	none	n/a	n/a
txwait	Delta	Seconds	Percentage	syslog, rmon	Percentage	Percentage	Event ID	syslog, rmon	none	Percentage	Percentage
[illegible]	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a
[illegible]	Delta	Seconds	Number	syslog, rmon, obfl	Number	Number	Event ID	syslog, rmon, obfl	none	n/a	n/a
[illegible]	Delta	Seconds	Number	syslog, rmon, obfl	Number	Number	Event ID	syslog, rmon, obfl	none	n/a	n/a

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
input-errors	Delta	Seconds	Number	syslog, rmon	Number	Number	Event ID	syslog, rmon	none	n/a	n/a

Note

- From Cisco MDS NX-OS Release 8.1(1), the err-pkt-from-port—ASIC Error Pkt from Port counter is deprecated.
- 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.
- The tx-slowport-oper-delay wait counter 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.
- The txwait counter 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** and **cong-isolate-recover** keywords to detect slow flow on a device. For more information, see Configuring a Port Monitor Policy, on page 74.
- You can configure RMON alerts for rx-datarate-burst, tx-datarate-burst, sfp-rx-power-low-warn and sfp-tx-power-low-warn counters. However, RMON alerts will not be generated.

For more information on internal CRC errors and the various stages, see "Internal CRC Detection and Isolation" section in *Cisco MDS 9000 Series High Availability Configuration Guide*.

Table 16: Slowdrain Port-Monitor Policy Threshold Value For Releases Prior to Cisco MDS NX-OS Release 8.5(1), on page 46 displays the threshold value of the slow-drain port-monitor policy:

Table 16: Slowdrain Port-Monitor Policy Threshold Value For Releases Prior to Cisco MDS NX-OS Release 8.5(1)

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

Table 17: Slowdrain Port-Monitor Policy Threshold Value For Cisco MDS NX-OS Release 8.5(1) and Later Releases

Counter	Threshold Type	Interval (Secs)	Warning		Thresholds		Rising/Falling actions			Congestion-signal	
			Threshold	Alerts	Rising	Falling	Event	Alerts	PortGuard	Warning	Alarm
Credit Loss Reco	Delta	1	none	n/a	1	0	4	syslog, rmon	none	n/a	n/a
TX Credit Not Available	Delta	1	none	n/a	10	0	4	syslog, rmon	none	n/a	n/a
tx-datarate	Delta	10	none	n/a	80	70	4	syslog, obfl	none	n/a	n/a
TXWait[7]	Delta	1	none	n/a	30	10	4	syslog, rmon	none	n/a	n/a

[7] Supported with Release 9.3(1) and later

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.

Crossbar (Xbar) Counters

The Xbar counters monitor internal CRC errors. These are CRC errors that have been caused internally by one of the forwarding *stages* in the switch. These only apply to director class FC modules.

The following are the crossbar counters:

- err-pkt-from-port
- err-pkt-to-xbar

- err-pkt-from-xbar

The above crossbar (Xbar) counters 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).
- Check interval does not function or apply to the crossbar counters.

- err-pkt-from-port—ASIC Error Pkt from port



Note The err-pkt-from-port counter is deprecated from Cisco MDS NX-OS Release 8.1(1).

- err-pkt-to-xbar—ASIC Error Pkt to xbar: This counter provides information about the number of internal CRC errors detected at an FC ASIC on a module and sent to the crossbar ASIC in the same module (ingress direction). These are referred to as *stage 1* internal CRC errors.
- err-pkt-from-xbar—ASIC Error Pkt from xbar: This counter provides information about the number of internal CRC errors detected at an FC ASIC on a module that were received from the crossbar ASIC in the same module (egress direction). These are referred to as *stage 5* internal CRC errors.

These two err-pkt counters are handled differently than the normal port monitor counters. Every 10 seconds (nonconfigurable), the counters' values are obtained for each FC ASIC on each module (linecard). If the counter has increased by any value, then port monitor increments its internal err-pkt-to/from-xbar counter by 1 for that FC ASIC. 10 seconds later they are checked and incremented again in a similar manner. The port monitor internal err-pkt-to/from-xbar counter would have to increase for a specific FC ASIC to a value that equals or exceeds the configured rising threshold in the configured poll-interval time for it to trigger a rising threshold alert. For example, if the poll interval is 60 and the rising threshold for this counter is 3, then it indicates that the counter for a specific FC ASIC for a port range would have to increment in a minimum of 3 separate 10 second intervals within the poll interval of 60 seconds to generate a rising-threshold alert.



Note

- On the 2/4/8/10/16 Gbps Advanced FC module, DS-X9448-768K9, there are 6 FC ASICs each handling 8 ports.
- On the 1/10/40G IPS,2/4/8/10/16G FC module, DS-X9334-K9, there are 3 FC ASICs each handling 8 ports.
- On the 4/8/16/32 Gbps Advanced FC module, DS-X9648-1536K9, there are 3 FC ASICs each handling 16 ports.

SFP Counters

From Cisco MDS NX-OS Release 8.5(1), the SFP counters allow you to configure the low warning thresholds for *Tx Power* and *Rx Power* for SFPs so that you receive a syslog when these values drop below the configured values. SFPs are monitored once every 10 minutes (600 seconds). The rising threshold is the count of the times the Rx or Tx Power was less than or equal to the SFP's Rx or Tx Power low warning threshold multiplied by the percentage. Consequently, the rising threshold can at most increment by one, every 10 minutes. Configuring a rising threshold value that is more than the 600 multiple of the poll interval will display an error. For example, for a polling interval of 1200, the rising threshold will be 2 (1200/600) and cannot be more than 2. The SFP counters are not included in the default policy and the only alert action that is available is syslog. You can configure the polling interval using the port monitor **counter** command.

You can configure the SFP counters as below:

- Configuring a low warning threshold percentage of 100% allows this counter to trigger when the Rx Power is less than or equal to the SFP's Rx Power low warning threshold.
- Configuring a low warning threshold percentage less than 100% allows this counter to trigger when the Rx Power is above the SFP's Rx Power low warning threshold.
- Configuring a low warning threshold percentage of greater than 100% allows this counter to trigger when the Rx Power is less than the SFP's Rx Power low warning threshold (between low warning and low alarm).

Note

- The SFP counters are not part of the default port monitor policy. You must explicitly enable them using the **monitor counter** command.
- The minimum polling interval for SFP counters is 600 seconds. The polling interval must be in multiple of 600. You can configure the polling interval using the port monitor **counter** command.

For configuring the SFP counters, see Configuring a Port Monitor Policy, on page 74.

The following are the SFP counters:

- sfp-rx-power-low-warn: Specifies the number of times a port's SFP has reached a percentage of the SFP's Rx Power's low warning threshold. This threshold varies depending on the SFP type, speed, and manufacturer and can be displayed via the **show interface transceiver details** command. Hence, this threshold is not an absolute value but a percentage of each individual SFP's Rx Power low warning threshold. This percentage can be configured in the range of 50% to 150% to allow for alerting at values less than the Rx Power low warning threshold or greater than the Rx Power low warning threshold.. Hence, this is an absolute value and varies between 50% to 150%. The low warning threshold value is calculated as the actual low warning threshold value of the SFP times the specified percentage. If the Rx power is lesser than or equal to the low warning threshold value, then this counter is incremented.
- sfp-tx-power-low-warn: Specifies the number of times a port's SFP has reached a percentage of the SFP's Tx Power's low warning threshold. This threshold varies depending on the SFP type, speed, and manufacturer and can be displayed via the **show interface transceiver details** command. Hence, this threshold is not an absolute value but a percentage of each individual SFP's Tx Power low warning threshold. This percentage can be configured in the range of 50% to 150% to allow for alerting at values less than the Tx Power low warning threshold or greater than the Tx Power low warning threshold.. Hence, this is an absolute value and varies between 50% to 100%. The low warning threshold value is

calculated as the actual low warning threshold value of the SFP times the specified percentage. If the Tx power is lesser than or equal to the low warning threshold value, then this counter is incremented.

Datarate Burst Counters

From Cisco MDS NX-OS Release 8.5(1), the datarate burst counters monitor the number of times the datarate crosses the configured threshold datarate in 1 second intervals. If the number crosses the configured number for rising threshold, the configured alert actions are taken as the condition is met. Datarate burst counters are polled every second. The datarate burst counters are not included in the default policy. For configuring the datarate burst counters, see Configuring a Port Monitor Policy, on page 74.

The following are the datarate burst counters:

- rx-datarate-burst
- tx-datarate-burst

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

- From Cisco MDS NX-OS Release 8.5(1), port monitor does *early detection* and does not require the port monitor check interval feature to be configured, as it is redundant.
- The port monitor check interval feature is supported only on the Cisco MDS 9710 Multilayer Director, Cisco MDS 9718 Multilayer Directors, Cisco MDS 9706 Multilayer Directors, Cisco MDS 9250i, Cisco MDS 9148T, Cisco MDS 9396T, and Cisco MDS 9132T.
- 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 Monitor Early Detection

Prior to Cisco MDS NX-OS Release 8.5(1) and without check interval configured, port-monitor checked to determine if the warning or rising thresholds were reached only after the polling interval expired. Starting with Cisco MDS NX-OS Release 8.5(1), most port monitor counters are monitored every second so that port monitor can detect warning and rising thresholds and take alert actions as soon as the threshold is detected. There is no change in the falling threshold behavior.

Port Monitor Alerts

From Cisco MDS NX-OS Release 8.5(1), port monitor allows you to configure alerts for each counter so that you can tailor the alerts that port monitor generates with each counter. By default, all counters are configured for syslog and RMON alerts. Only the rx-datarate, tx-datarate, rx-datarate-burst, and tx-datarate-burst counters allow the configuration of the OBFL alert type. OBFL indicates that these counters record their events into Onboard Failure Logging. These are disposable via the **show logging onboard datarate** command.

The following alerts are supported:

- syslog: Generates a syslog when a configured threshold is reached. You can also configure an event ID (severity-level) for the syslogs that are generated when a rising or falling threshold is detected so that you can filter the logs using the severity level.

  The following severity levels are supported:

  - ALERT (1)
  - CRITICAL (2)
  - ERROR (3)
  - WARNING (4)
  - NOTICE (5)

- rmon: Generates an SNMP alert when a configured threshold is reached.
- obfl: Enables OBFL logging.

Note The OBFL alert is supported only for rx-datarate, tx-datarate, rx-datarate-burst, and tx-datarate-burst counters.

- none: Disables all alerts.

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 18: 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.

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, flapped, congestion-isolated, and so on 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.

Note From Cisco MDS NX-OS Release 8.5(1), the input errors, sfp-rx-power-low-warn, sfp-tx-power-low-warn, rx-datarate-burst, and tx-datarate-burst counters were added.

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

- credit-loss-reco
- link-loss
- signal-loss
- sync-loss
- rx-datarate
- invalid-crcs
- invalid-words
- timeout-discards
- tx-credit-not-available
- tx-datarate
- tx-discards
- tx-slowport-oper-delay
- txwait
- input-errors
- sfp-rx-power-low-warn
- sfp-tx-power-low-warn
- state-change
- rx-datarate-burst
- tx-datarate-burst

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. 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

- FEC on Cisco MDS 9000 Series switches should be disabled on 16 Gbps links.
- Ensure that enough buffer credits are configured on MDS ports to cover the distance for the specific frame size.
- Use logical link bundling via port channels across DWDM links even in case of different path lengths. The superior technology of port channels to any other logical bundling solution in the industry makes this possible and provides the benefit of higher reliability without impacting FSPF routing decisions.
- Check with your DWDM vendor for specific configurations and setting, specifically for NOS/LOS propagation.
- When a module initializes for the first time, the LCP timer starts, followed by the interfaces coming up. However, the PM obtains transceiver details from different MTS for all active ports except for the DOM value, which is the expected behavior. After the module reloads, the DOM value for interfaces during the initial LCP poll interval is set to 10 minutes.
- 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 29.

Guidelines for Configuring Port Monitor Check Interval

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

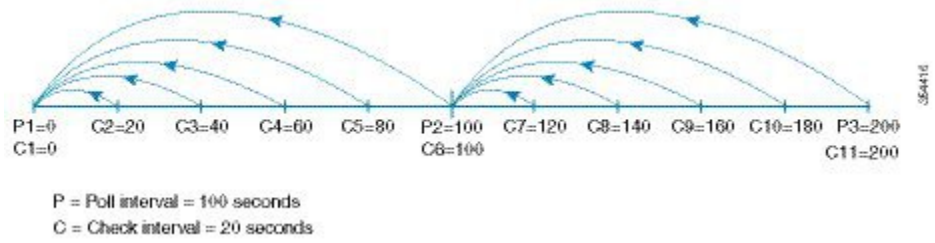
Note The value of the check interval is common across counters and policies.

- We recommend that you configure the check interval to be less than the poll interval. Also, configure the poll interval as a multiple of the check 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.
- From Cisco MDS NX-OS Release 8.4(1), this feature is supported on Cisco MDS switches that are operating in 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 19: Default Interface Parameters , on page 61 lists the default settings for interface parameters.

Table 19: 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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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 66 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:

Procedure

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:

Procedure

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 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:

Procedure

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.

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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*]

Note
The duration range is from 45 to 2000000 seconds. The duration must be equal to or greater than **num-times** multipled by 45.

(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 5 times within 225 seconds due to multiple causes. The portguard controls the interface in the following manner:

Example

This example shows how to configure portguard to bring a port to down state if the link flaps 5 times within 225 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 225
switch(config-if)# errdisable detect cause bit-errors num-times 5 duration 225
switch(config-if)# errdisable detect cause credit-loss num-times 5 duration 225
```

The above example sets the configuration to the following status:

- The port will be error disabled due to link down if the port suffers link failure due to link down 5 times in 225 seconds.
- The port will be error-disabled due to bit errors if the port suffers link failure due to bit errors 5 times in 225 seconds.
- The port will be error-disabled due to credit loss if the port suffers link failure due to credit loss 5 times in 225 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:

Procedure

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:

Procedure

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:

Procedure

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:

Releases prior to Cisco MDS NX-OS Release 8.5(1)

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** { **cong-isolate** | **errordisable** | **flap**}

Cisco MDS NX-OS Release 8.5(1) and later releases

switch(config-port-monitor)# **counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **input-errors** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **rx-datarate-burst** | **sfp-rx-power-low-warn** | **sfp-tx-power-low-warn** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-datarate-burst** | **tx-discards** | **tx-slowport-oper-delay** | **txwait** [**warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals**]} **poll-interval**

seconds {**absolute** | **delta**} **rising-threshold** *count3* **event** *RMON-ID* [**warning-threshold** *count4*] [**alerts** [**obfl** **rmon** **syslog** | **none**]] [**datarate** *count5*] [**falling-threshold** *count6*] [**portguard** {**DIRL** | **FPIN** | **cong-isolate** | **cong-isolate-recover** | **errordisable** | **flap**]}

Note

- A port monitor policy cannot be configured as a combination of cong-isolate, cong-isolate-recover, DIRL, and FPIN port guard actions. For example, if in a policy you configure the tx-datarate, tx-datarate-burst, and txwait with DIRL portguard action and then configure the credit-loss-reco counter with the cong-isolate portguard action, you will not be able to activate the policy.
- Port monitor polling interval must not be more than the configured recovery interval when the cong-isolate, cong-isolate-recover, DIRL, and FPIN port guard actions are configured.
- 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) and congestion isolation recovery (cong-isolate-recover).
- From Cisco MDS NX-OS Release 8.5(1), a new default *fabricmon_edge_policy* is introduced where FPIN is already configured for the supported counters.
- From Cisco MDS NX-OS Release 8.5(1), switches operating in the Cisco NPV mode do not support cong-isolate, cong-isolate-recover, DIRL, and FPIN portguard actions and the default *fabricmon_edge_policy*.
- When you configure a policy with the cong-isolate, cong-isolate-recover, DIRL, or FPIN portguard actions, you can expect multiple rising thresholds without waiting for a falling threshold.
- You must configure Exchange Diagnostic Capabilities (EDC) interval for congestion signal before configuring the TxWait **warning-signal-threshold** and **alarm-signal-threshold** values. For more information, see Configuring EDC Congestion Signal, on page 279.
- Ensure that you provide at least one minute delay between deactivation and activation of port monitor policy when configuring the portguard actions **cong-isolate**, **cong-isolate-recover**, and **FPIN**.
- The **cong-isolate**, **cong-isolate-recover**, **DIRL**, and **FPIN** portguard actions are applicable only for logical-type edge policies.
- The **cong-isolate** and **cong-isolate-recover** port monitor portguard actions are supported only for the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters.
- The **DIRL** port monitor portguard action is supported only for the tx-datarate, tx-datarate-burst, and txwait counters.
- The **FPIN** port monitor portguard action is supported only for the link-loss, sync-loss, signal-loss, invalid-words, invalid-crc, and txwait counters.

- For SFP counters, **sfp-rx-power-low-warn** and **sfp-tx-power-low-warn**, the polling interval must be configured in multiples of 600 (10 minutes) and the rising threshold value should not exceed the multiple value of the polling interval. For example, if the polling interval is configure as 1800, which is 3 times 600, then the rising threshold value should not be more than 3.
- The rx-datarate-burst and tx-datarate-burst counters are configured as the number of 1-second bursts above 90% (default) detected in a polling interval. You can change the default datarate burst threshold using the **counter tx-datarate-burst poll-interval** *seconds* **delta rising-threshold** *count* **event** *RMON-ID* **datarate** *percentage* command.

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

Releases prior to Cisco MDS NX-OS Release 8.5(1)

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** { **cong-isolate** | **errordisable** | **flap**}

Cisco MDS NX-OS Release 8.5(1) and later releases

switch(config-port-monitor)# **no counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **input-errors** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **rx-datarate-burst** | **sfp-rx-power-low-warn** | **sfp-tx-power-low-warn** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-datarate-burst** | **tx-discards** | **tx-slowport-oper-delay** | **txwait** [**warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals**]} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count3* **event** *RMON-ID* [**warning-threshold** *count4*] [**alerts** [**obfl rmon syslog** | **none**]] [**datarate** *count5*] [**falling-threshold** *count6*] [**portguard** {**DIRL** | **FPIN** | **cong-isolate** | **cong-isolate-recover** | **errordisable** | **flap**]}

(Optional) Monitor a counter:

Releases prior to Cisco MDS NX-OS Release 8.5(1)

switch(config-port-monitor)# **monitor counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **input-errors** | **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-count** | **tx-slowport-oper-delay** | **txwait**}

Cisco MDS NX-OS Release 8.5(1) and later releases

switch(config-port-monitor)# **monitor counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **input-errors** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **rx-datarate-burst** | **sfp-rx-power-low-warn** | **sfp-tx-power-low-warn** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-datarate-burst** | **tx-discards** | **tx-slowport-count** | **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:

Procedure

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 Logging Level for Port Monitor

To configure logging level for port monitor syslog messages, perform the steps below:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Configure a logging level for port monitor syslog messages:

switch(config)# **logging level pmon** *severity-level*

(Optional) Revert to the default logging level for the port monitor syslog messages:

switch(config)# **no logging level pmon**

Configuring Port Monitor Portguard

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

Procedure

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:

Releases prior to Cisco MDS NX-OS Release 8.5(1)

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** { **cong-isolate** | **errordisable** | **flap**}

Cisco MDS NX-OS Release 8.5(1) and later releases

switch(config-port-monitor)# **counter** {**credit-loss-reco** | **err-pkt-from-port** | **err-pkt-from-xbar** | **err-pkt-to-xbar** | **input-errors** | **invalid-crc** | **invalid-words** | **link-loss** | **lr-rx** | **lr-tx** | **rx-datarate** | **rx-datarate-burst** | **sfp-rx-power-low-warn** | **sfp-tx-power-low-warn** | **signal-loss** | **state-change** | **sync-loss** | **timeout-discards** | **tx-credit-not-available** | **tx-datarate** | **tx-datarate-burst** | **tx-discards** | **tx-slowport-oper-delay** | **txwait** [**warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals**]} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count3* **event** *RMON-ID* [**warning-threshold** *count4*] [**alerts** [**obfl rmon syslog** | **none**]] [**datarate** *count5*] [**falling-threshold** *count6*] [**portguard** {**DIRL** | **FPIN** | **cong-isolate** | **cong-isolate-recover** | **errordisable** | **flap**]}

Note

- A port monitor policy cannot be configured as a combination of cong-isolate, cong-isolate-recover, DIRL, and FPIN port guard actions. For example, if in a policy you configure the tx-datarate, tx-datarate-burst, and txwait with DIRL portguard action and then configure the credit-loss-reco counter with the cong-isolate portguard action, you will not be able to activate the policy.
- Port monitor polling interval must not be more than the configured recovery interval when the cong-isolate, cong-isolate-recover, DIRL, and FPIN port guard actions are configured.
- 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) and congestion isolation recovery (cong-isolate-recover).
- From Cisco MDS NX-OS Release 8.5(1), a new default *fabricmon_edge_policy* is introduced where FPIN is already configured for the supported counters.
- From Cisco MDS NX-OS Release 8.5(1), switches operating in the Cisco NPV mode do not support cong-isolate, cong-isolate-recover, DIRL, and FPIN portguard actions and the default *fabricmon_edge_policy*.
- When you configure a policy with the cong-isolate, cong-isolate-recover, DIRL, or FPIN portguard actions, you can expect multiple rising thresholds without waiting for a falling threshold.
- You must configure Exchange Diagnostic Capabilities (EDC) interval for congestion signal before configuring the TxWait **warning-signal-threshold** and **alarm-signal-threshold** values. For more information, see Configuring EDC Congestion Signal, on page 279.
- Ensure that you provide at least one minute delay between deactivation and activation of port monitor policy when configuring the portguard actions **cong-isolate**, **cong-isolate-recover**, and **FPIN**.
- The **cong-isolate**, **cong-isolate-recover**, **DIRL**, and **FPIN** portguard actions are applicable only for logical-type edge policies.
- The **cong-isolate** and **cong-isolate-recover** port monitor portguard actions are supported only for the credit-loss-reco, tx-credit-not-available, tx-slowport-oper-delay, and txwait counters.
- The **DIRL** port monitor portguard action is supported only for the tx-datarate, tx-datarate-burst, and txwait counters.
- The **FPIN** port monitor portguard action is supported only for the link-loss, sync-loss, signal-loss, invalid-words, invalid-crc, and txwait counters.
- For SFP counters, **sfp-rx-power-low-warn** and **sfp-tx-power-low-warn**, the polling interval must be configured in multiples of 600 (10 minutes) and the rising threshold value should not exceed the multiple value of the polling interval. For example, if the polling interval is configure as 1800, which is 3 times 600, then the rising threshold value should not be more than 3.
- The rx-datarate-burst and tx-datarate-burst counters are configured as the number of 1-second bursts above 90% (default) detected in a polling interval. You can change the default datarate burst threshold using the **counter tx-datarate-burst poll-interval** *seconds* **delta rising-threshold** *count* **event** *RMON-ID* **datarate** *percentage* command.

Configuring Port Group Monitor

Enabling a Port Group Monitor

To enable a port group monitor, perform these steps:

Procedure

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:

Procedure

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 the Monitoring of Specific Counter.

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:

Procedure

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:

Procedure

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:

Procedure

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:

Procedure

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 Interface Configuration

Display interface information

Run the **show interface** command to display physical interface related information. Without any arguments, this command displays the information for all the configured interfaces in the switch.

The command output format has evolved over time. Outputs may differ if not running the latest release of Cisco MDS NX-OS

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 20:01:00:2a:6a:26:a1:e0
    Peer port WWN is 21:10:00:24:ff:98:01:01
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is F, FCID is 0xe90c00
    Port vsan is 1
    Admin Speed is auto
    Operating Speed is 4 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY
    Transmit B2B Credit is 16
    Receive B2B Credit is 64
    B2B State Change: Admin(on), Oper(down)
    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
      16 frames input, 1312 bytes, 82 bytes/frame average frame size
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      17 frames output, 1232 bytes, 72 bytes/frame average frame size
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,1 LRR, 0 NOS, 0 loop inits
    64 receive B2B credit remaining
    16 transmit B2B credit remaining
    16 low priority transmit B2B credit remaining
    0 timeout discards, 0 credit loss
    0 link failures, 0 sync losses, 0 signal losses
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Interface last changed at Sat Apr 19 17:08:22 2025
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is FNS190413TV
      Cisco pid is DS-SFP-FC16G-SW
      Temperature 40.12 C, Voltage 3.34 V, Current 8.76 mA, TxPower -3.61 dBm, R
xPower -3.92 dBm
      Power control is not supported, Power is on
fc1/2 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:03:00:2a:6a:26:a1:e0
    Peer port WWN is 20:03:00:de:fb:74:e4:60
```

```
    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
    Operating Speed is 4 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY
    Transmit B2B Credit is 500
    Receive B2B Credit is 64
    B2B State Change: Admin(on), Oper(up), Negotiated Value(14)
    Receive data field Size is 2112
    Beacon is turned off
    Logical type is core
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 32 bits/sec,4 bytes/sec, 0 frames/sec
    5 minutes output rate 32 bits/sec,4 bytes/sec, 0 frames/sec
      493 frames input, 42776 bytes, 86 bytes/frame average frame size
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      493 frames output, 45400 bytes, 92 bytes/frame average frame size
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,2 LRR, 0 NOS, 1 loop inits
    64 receive B2B credit remaining
    500 transmit B2B credit remaining
    485 low priority transmit B2B credit remaining
    0 timeout discards, 0 credit loss
    0 link failures, 0 sync losses, 0 signal losses
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Interface last changed at Sat Apr 19 17:08:25 2025
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is FNS19130U1L
      Cisco pid is DS-SFP-FC16G-SW
      Temperature 43.41 C, Voltage 3.33 V, Current 8.62 mA, TxPower -2.50 dBm, R
xPower -3.96 dBm
      Power control is not supported, Power is on

fc1/3 is down (Link failure or not-connected)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:04:00:2a:6a:26:a1:e0
    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
      0 frames input, 0 bytes, 0 bytes/frame average frame size
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      0 frames output, 0 bytes, 0 bytes/frame average frame size
        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 timeout discards, 0 credit loss
    0 link failures, 0 sync losses, 0 signal losses
```

```
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is FNS18150JCU
      Cisco pid is DS-SFP-FC16G-SW
      Temperature 44.30 C, Voltage 3.32 V, Current 7.91 mA, TxPower -2.99 dBm, R
xPower -3.67 dBm
      Power control is not supported, Power is on

sup-fc0 is up
    Hardware is Fibre Channel
    Speed is 1 Gbps
    1978 packets input, 222120 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame
      0 overrun, 0 fifo
    1980 packets output, 252544 bytes
      0 underruns, 0 output errors
      0 collisions, 0 fifo
      0 carrier errors

IPStorage1/1 is down (SFP not present)
    Hardware is IPStorage, address is f40f.1b72.ca33
    Internet address(es):
    MTU 1500  bytes
    Port mode is IPS
    Beacon is turned off
    Auto-Negotiation is turned on
    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 packets input, 0 bytes, 0 bytes/packet average packet size
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    0 packets output, 0 bytes, 0 bytes/packet average packet size, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors

fcip1 is down (Tunnel port src interface unbound)
    Hardware is IPStorage
    Port WWN is 00:00:00:00:00:00:00:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    Belongs to port-channel50
    Peer Information
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    Ficon Load Balancer configured off for all vsans
    Ficon Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 2
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      0 Active TCP connections
        0 Attempts for active connections, 0 close of connections
    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 bytes/frame average frame size
         0 Class F frames input, 0 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
```

```
          0 Error frames timestamp error 0
        0 frames output, 0 bytes, 0 bytes/frame average frame size
          0 Class F frames output, 0 bytes
          0 Class 2/3 frames output, 0 bytes
          0 Error frames

mgmt0 is up
    Internet address is 10.197.111.181/24
    Hardware is GigabitEthernet
    Address is f40f.1b72.ca30
    MTU 1500 bytes, BW 1000 Mbps full Duplex
    42241 packets input, 3274206 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame
      0 overrun, 0 fifo
    110 packets output, 24487 bytes
      0 underruns, 0 output errors
      0 collisions, 0 fifo
      0 carrier errors

Ethernet1/1 is down (SFP not inserted)
admin state is down, Dedicated Interface
  Hardware: 10000 Ethernet, address: f40f.1b72.ca35 (bia f40f.1b72.ca35)
  MTU bytes (CoS values):  MTU  1500(0-2,4-7) bytes  MTU  2112(3) bytes
  BW 10000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, medium is broadcast
  Port mode is access
  auto-duplex, auto-speed
  Beacon is turned off
  Auto-Negotiation is turned on
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned off
  Switchport monitor is off
  Last link flapped never
  Last clearing of "show interface" counters never
  0 interface resets
  Load-Interval #1: 30 seconds
    30 seconds input rate 0 bits/sec, 0 packets/sec
    30 seconds output rate 0 bits/sec, 0 packets/sec
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  Load-Interval #2: 5 minute (300 seconds)
    300 seconds input rate 0 bits/sec, 0 packets/sec
    300 seconds output rate 0 bits/sec, 0 packets/sec
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  RX
    0 unicast packets  0 multicast packets  0 broadcast packets
    0 input packets  0 bytes
    0 jumbo packets  0 storm suppression packets
    0 runts  0 giants  0 CRC/FCS  0 no buffer
    0 input error  0 short frame  0 overrun   0 underrun  0 ignored
    0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
    0 input with dribble  0 input discard
    0 Rx pause
  TX
    0 unicast packets  0 multicast packets  0 broadcast packets
    0 output packets  0 bytes
    0 jumbo packets
    0 output error  0 collision  0 deferred  0 late collision
    0 lost carrier  0 no carrier  0 babble  0 output discard
    0 Tx pause

port-channel7 is down (No operational members)
    Hardware is Fibre Channel
    Port WWN is 24:07:00:2a:6a:26:a1:e0
```

```
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    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
      0 frames input, 0 bytes, 0 bytes/frame average frame size
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      0 frames output, 0 bytes, 0 bytes/frame average frame size
        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 timeout discards, 0 credit loss
    0 link failures, 0 sync losses, 0 signal losses
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    No members
```

Displays the counters for all interfaces

```
switch# show interface counters
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
    16 frames input, 1312 bytes,  82 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      16 class-3 frames, 1312 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    17 frames output, 1232 bytes,  72 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      17 class-3 frames, 1232 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors
    0 timeout discards, 0 credit loss
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    1 output OLS, 1 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
     3 Transmit B2B credit transitions to zero
     1 Receive B2B credit transitions to zero
      0 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    64 receive B2B credit remaining
    16 transmit B2B credit remaining
    16 low priority transmit B2B credit remaining
    Last clearing of "show interface" counters: never
fc1/2
    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
    14 frames input, 1124 bytes,  80 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      14 class-3 frames, 1124 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    15 frames output, 1036 bytes,  69 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      15 class-3 frames, 1036 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors
    0 timeout discards, 0 credit loss
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
```

```
    1 output OLS, 1 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
     3 Transmit B2B credit transitions to zero
     1 Receive B2B credit transitions to zero
```

Displays an particular interface

```
switch# show interface fc1/1
fc1/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:01:00:2a:6a:26:a1:e0
    Peer port WWN is 21:10:00:24:ff:98:01:01
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is F, FCID is 0xe90c00
    Port vsan is 1
    Admin Speed is auto
    Operating Speed is 4 Gbps
    Rate mode is dedicated
    Port flow-control is R_RDY
    Transmit B2B Credit is 16
    Receive B2B Credit is 64
    B2B State Change: Admin(on), Oper(down)
    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
      16 frames input, 1312 bytes, 82 bytes/frame average frame size
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      17 frames output, 1232 bytes, 72 bytes/frame average frame size
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      1 output OLS,1 LRR, 0 NOS, 0 loop inits
    64 receive B2B credit remaining
    16 transmit B2B credit remaining
    16 low priority transmit B2B credit remaining
    0 timeout discards, 0 credit loss
    0 link failures, 0 sync losses, 0 signal losses
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Interface last changed at Sat Apr 19 17:08:22 2025
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is FNS190413TV
      Cisco pid is DS-SFP-FC16G-SW
      Temperature 39.20 C, Voltage 3.34 V, Current 8.76 mA, TxPower -3.59 dBm, R
xPower -3.92 dBm
      Power control is not supported, Power is on
```

Displays the counters of an particular interface

```
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 0 bits/sec, 0 bytes/sec, 0 frames/sec
    16 frames input, 1312 bytes,  82 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      16 class-3 frames, 1312 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
```

```
    17 frames output, 1232 bytes,  72 bytes/frame average frame size
      0 class-2 frames, 0 bytes
      17 class-3 frames, 1232 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors
    0 timeout discards, 0 credit loss
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    1 output OLS, 1 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
     3 Transmit B2B credit transitions to zero
     1 Receive B2B credit transitions to zero
      0 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    64 receive B2B credit remaining
    16 transmit B2B credit remaining
    16 low priority transmit B2B credit remaining
    Last clearing of "show interface" counters: never
```

Displays the details of counters of an particular interface

```
switch# show interface fc1/1 counters detailed
fc1/1
 Rate Stats:
  Rx 5 min rate bit/sec:                                                0
  Tx 5 min rate bit/sec:                                                0
  Rx 5 min rate bytes/sec:                                              0
  Tx 5 min rate bytes/sec:                                              0
  Rx 5 min rate frames/sec:                                             0
  Tx 5 min rate frames/sec:                                             0

 Total Stats:
  Rx total frames:                                                     16
  Tx total frames:                                                     17
  Rx total bytes:                                                    1312
  Tx total bytes:                                                    1232
  Rx bytes/frame average frame size:                                   82
  Tx bytes/frame average frame size:                                   72
  Rx total multicast:                                                   0
  Tx total multicast:                                                   0
  Rx total broadcast:                                                   0
  Tx total broadcast:                                                   0
  Rx total unicast:                                                    16
  Tx total unicast:                                                    17
  Rx total discards:                                                    0
  Tx total discards:                                                    0
  Rx total errors:                                                      0
  Tx total errors:                                                      0
  Rx class-2 frames:                                                    0
  Tx class-2 frames:                                                    0
  Rx class-2 bytes:                                                     0
  Tx class-2 bytes:                                                     0
  Rx class-2 frames discards:                                           0
  Rx class-2 port reject frames:                                        0
  Rx class-3 frames:                                                   16
  Tx class-3 frames:                                                   17
  Rx class-3 bytes:                                                  1312
  Tx class-3 bytes:                                                  1232
  Rx class-3 frames discards:                                           0
  Rx class-f frames:                                                    0
  Tx class-f frames:                                                    0
  Rx class-f bytes:                                                     0
  Tx class-f bytes:                                                     0
  Rx class-f frames discards:                                           0
```

```
 Link Stats:
  Rx Link failures:                                                    0
  Rx Sync losses:                                                      0
  Rx Signal losses:                                                    0
  Rx Primitive sequence protocol errors:                               0
  Rx Invalid transmission words:                                       0
  Rx Invalid CRCs:                                                     0
  Rx Delimiter errors:                                                 0
  Rx fragmented frames:                                                0
  Rx frames with EOF aborts:                                           0
  Rx unknown class frames:                                             0
  Rx Runt frames:                                                      0
  Rx Jabber frames:                                                    0
  Rx too long:                                                         0
  Rx too short:                                                        0
  Rx Link Reset(LR) while link is active:                              0
  Tx Link Reset(LR) while link is active:                              0
  Rx Link Reset Responses(LRR):                                        0
  Tx Link Reset Responses(LRR):                                        1
  Rx Offline Sequences(OLS):                                           0
  Tx Offline Sequences(OLS):                                           1
  Rx Non-Operational Sequences(NOS):                                   0
  Tx Non-Operational Sequences(NOS):                                   0
  BB_SCs credit resend actions:                                        0
  BB_SCr Tx credit increment actions:                                  0

 Loop Stats:
  Rx F8 type LIP sequence errors:                                      0
  Tx F8 type LIP sequence errors:                                      0
  Rx Non F8 type LIP sequence errors:                                  0
  Tx Non F8 type LIP sequence errors:                                  0

 Congestion Stats:
  Tx Timeout discards:                                                 0
  Tx Credit loss:                                                      0
  TxWait 2.5us due to lack of transmit credits:                        0
  Percentage TxWait for last 1s/1m/1h/72h:                     0%/0%/0%/0%
  Rx B2B credit remaining:                                             0
  Tx B2B credit remaining:                                             0
  Tx Low Priority B2B credit remaining:                                0
  Rx B2B credit transitions to zero:                                   1
  Tx B2B credit transitions to zero:                                   3

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

Displays the FCIP interface

```
switch# show interface fcip 1
fcip1 is down (Tunnel port src interface unbound)
    Hardware is IPStorage
    Port WWN is 00:00:00:00:00:00:00:00
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    Belongs to port-channel50
    Peer Information
    Write acceleration mode is configured off
    Tape acceleration mode is configured off
    Tape Accelerator flow control buffer size is automatic
    FICON XRC Accelerator is configured off
    Ficon Load Balancer configured off for all vsans
    Ficon Tape acceleration configured off for all vsans
    IP Compression is disabled
    Maximum number of TCP connections is 2
```

```
    QOS control code point is 0
    QOS data code point is 0
    TCP Connection Information
      0 Active TCP connections
        0 Attempts for active connections, 0 close of connections
    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 bytes/frame average frame size
         0 Class F frames input, 0 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      0 frames output, 0 bytes, 0 bytes/frame average frame size
         0 Class F frames output, 0 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
```

Displays the counters of a FCIP interface

```
switch# show interface fcip 1 counters
fcip1
    TCP Connection Information
      0 Active TCP connections
        0 Attempts for active connections, 0 close of connections
    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 bytes/frame average frame size
         0 Class F frames input, 0 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      0 frames output, 0 bytes, 0 bytes/frame average frame size
         0 Class F frames output, 0 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
```

Displays the details of the counters for a FCIP interface

```
switch# show interface fcip 1 counters detailed
fcip1
    TCP Connection Information
      0 Active TCP connections
        0 Attempts for active connections, 0 close of connections
    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 bytes/frame average frame size
         0 Class F frames input, 0 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      0 frames output, 0 bytes, 0 bytes/frame average frame size
         0 Class F frames output, 0 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
```

Displays the IPS interface

```
switch# show interface ipstorage 1/1
IPStorage1/1 is down (SFP not present)
    Hardware is IPStorage, address is f40f.1b72.ca33
    Internet address(es):
    MTU 1500  bytes
    Port mode is IPS
```

```
    Beacon is turned off
    Auto-Negotiation is turned on
    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 packets input, 0 bytes, 0 bytes/packet average packet size
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    0 packets output, 0 bytes, 0 bytes/packet average packet size, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Displays the counters of an IPS interface

```
switch# show interface ipstorage 1/1 counters
IPStorage1/1
    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 packets input, 0 bytes, 0 bytes/packet average packet size
      0 multicast frames, 0 broadcast frames
      0 errors, 0 queue drops, 0 if-down drops, 0 RED drops
      0 bad ether type drop, 0 bad protocol drops
    0 packets output, 0 bytes, 0 bytes/packet average packet size, 0 underruns
      0 multicast, 0 broadcast
      0 errors, 0 collisions, 0 arp drops, 0 if-down drops
```

Displays the details of counters for an IPS interface

```
switch# show interface ipstorage 1/1 counters detailed
IPStorage1/1
 Rate Stats:
  Rx 5 minutes input rate bits/sec:                                    0
  Rx 5 minutes input rate bytes/sec:                                   0
  Rx 5 minutes input rate packets/sec:                                 0
  Tx 5 minutes input rate bits/sec:                                    0
  Tx 5 minutes input rate bytes/sec:                                   0
  Tx 5 minutes input rate packets/sec:                                 0

 Total Stats:
  Rx packets input:                                                    0
  Rx bytes:                                                            0
  Rx bytes/frame average frame size:                                   0
  Rx multicast frames:                                                 0
  Rx broadcast frames:                                                 0
  Tx packets output:                                                   0
  Tx bytes:                                                            0
  Tx bytes/packet average packet size:                                 0
  Tx multicast frames:                                                 0
  Tx broadcast frames:                                                 0
  Tx errors:                                                           0

 MAC Error Stats:
  Rx errors:                                                           0
  Rx queue drops:                                                      0
  Rx if-down drops:                                                    0
  Rx RED drops:                                                        0
  Rx bad ether type drops:                                             0
  Rx bad protocol drops:                                               0
  Rx packets with FCS or align errors:                                 0
  Rx packets with length < min:                                        0
  Rx packets with length < min & FCS error:                            0
  Rx packets with length > max:                                        0
  Rx packets with length > max & FCS error:                            0
  Rx packets with PCI error:                                           0
```

```
  Rx packets with GMX/SPX/PCI error:                                    0
  Rx packets dropped by RED buffer exhaustion:                          0
  Rx L2 multicast packets:                                              0
  Rx L2 broadcast packets:                                              0
  Tx underruns:                                                         0
  Tx collisions:                                                        0
  Tx ARP drops:                                                         0
  Tx if-down drops:                                                     0
```

Displays the port channel interface

```
switch# show interface port-channel50
port-channel50 is down (No operational members)
    Hardware is IPStorage
    Port WWN is 24:32:00:2a:6a:26:a1:e0
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    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
      0 frames input, 0 bytes, 0 bytes/frame average frame size
         0 Class F frames input, 0 bytes
         0 Class 2/3 frames input, 0 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      0 frames output, 0 bytes, 0 bytes/frame average frame size
         0 Class F frames output, 0 bytes
         0 Class 2/3 frames output, 0 bytes
         0 Error frames
    Member[1] : fcip1     [down]
    Member[2] : fcip2     [down]
```

Displaying interface information for other 9.x releases

The following examples are for releases prior to 9.4(4)

Displays the status of 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 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


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      auto   on      down         swl    --     --     --       --
fc1/2       1      auto   on      down         swl    --     --     --       --
fc1/3       1      auto   on      down         swl    --     --     --       --
fc1/4       1      auto   on      down         swl    --     --     --       --
fc1/5       1      auto   on      down         swl    --     --     --       --
fc1/6       1      auto   on      down         swl    --     --     --       --
fc1/7       1      auto   on      down         swl    --     --     --       --
fc1/8       1      auto   on      down         swl    --     --     --       --
fc1/9       1      auto   on      down         swl    --     --     --       --
fc1/10      1      auto   on      down         swl    --     --     --       --
fc1/11      1      auto   on      down         swl    --     --     --       --
fc1/12      1      auto   on      sfpAbsent    --     --     --     --       --

--------------------------------------------------------------------------------
Interface                Status                            Speed
                                                           (Gbps)
--------------------------------------------------------------------------------
sup-fc0                  up                                1

--------------------------------------------------------------------------------
Interface              Status         IP Address        Speed     MTU
--------------------------------------------------------------------------------
IPStorage1/1           outOfServc     10.1.1.32/24       1 Gbps    1500
IPStorage1/2           outOfServc     --                 1 Gbps    1500
IPStorage1/3           outOfServc     --                 1 Gbps    1500
IPStorage1/4           errDisabled    --                 1 Gbps    1500
IPStorage1/5           init           --                 25 Gbps   1500
IPStorage1/6           outOfServc     --                 40 Gbps   1500
```

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 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.

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 Interface Statistics

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

The following example displays the statistics of interfaces:

```
switch# ShowIntStats
2023/07/21 18:27:15 Link Stats:
```

Intf	LRR Rx	LRR Tx	FEC Corrected	Link FEC Failures	Sync FEC Loss Uncorrected	Signal Loss BB_SCs	Invalid Words	Invalid CRCs	NOS Rx	NOS Tx	OLS Rx	OLS Tx
fc1/1	0	4	0	3	0	0	0	0	0	3	3	4
fc1/2	0	2	0	0	0	0	0	0	0	0	0	1
fc1/3	2	2	0	1	0	0	0	0	0	1	1	2
fc1/4	2	2	0	1	0	0	3051	0	0	1	1	2
fc1/5	2	2	0	1	0	0	0	0	0	1	1	2
fc1/6	2	2	0	1	0	0	0	0	1	1	2	2
fc1/7	0	0	0	0	0	0	0	0	0	0	0	0
fc1/8	0	0	0	0	0	0	0	0	0	0	0	0
fc1/9	0	0	0	0	0	0	0	0	0	0	0	0
fc1/10	0	0	0	0	0	0	0	0	0	0	0	0
fc1/11	0	2	0	0	0	0	0	0	0	0	0	1
fc1/12	0	2	0	0	0	0	0	0	0	0	0	1
fc1/13	0	2	0	0	0	0	0	0	0	0	0	1
fc1/14	4	0	1	1	2	0	0	0	4	3	2	2
fc1/15	4	0	0	1	2	0	0	0	2	2	2	2
fc1/16	3	1	0	1	2	0	0	0	1	1	1	2
fc1/17	0	0	0	0	0	0	0	0	0	0	0	0
fc1/18	0	0	0	0	0	0	0	0	0	0	0	0
fc1/19	2	0	44222135	0	0	0	0	0	1	0	1	1
fc1/20	1	1	58311348	0	0	0	0	0	0	0	0	1
fc1/21	1	1	2776734	0	0	0	0	0	0	0	0	1
fc1/22	1	1	32918274	0	0	0	0	0	0	0	0	1
fc1/23	2	0	52394	0	0	0	0	0	0	0	1	1
fc1/24	0	4	0	3	0	0	0	0	0	3	3	4
fc1/25	0	4	0	3	0	0	0	0	0	3	3	4
fc1/26	0	4	0	3	0	0	0	0	0	3	3	4
fc1/27	0	4	0	3	0	0	0	0	0	3	3	4

```
| fc1/28          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |         0 |           0 |       0 |       0 |
| fc1/29          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |  48458710 |           0 |       0 |       0 |
| fc1/30          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |         0 |           0 |       0 |       0 |
| fc1/31          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   0
|   0 |   0 |         0 |           0 |       0 |       0 |
| fc1/32          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   1 |         0 |           0 |       0 |       0 |
| fc1/33          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |         0 |           0 |       0 |       0 |
| fc1/34          |        1 |      0 |      0 |        0 |        0 |   0 |   1 |   1 |   2
|   2 |   2 |         0 |           0 |       0 |       0 |
| fc1/35          |        1 |      0 |      0 |        0 |        0 |   0 |   1 |   1 |   2
|   2 |   2 |      1930 |           0 |       0 |       0 |
| fc1/36          |        1 |      0 |      0 |        0 |        0 |   0 |   1 |   1 |   2
|   2 |   2 |        53 |           0 |       0 |       0 |
| fc1/37          |        1 |      0 |      0 |        0 |        0 |   0 |   1 |   1 |   2
|   2 |   2 |         0 |           0 |       0 |       0 |
| fc1/38          |        1 |      0 |      0 |        0 |        0 |   0 |   1 |   1 |   2
|   2 |   2 |         0 |           0 |       0 |       0 |
| fc1/39          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |         0 |           0 |       0 |       0 |
| fc1/40          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |         0 |           0 |       0 |       0 |
| fc1/41          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |         0 |           0 |       0 |       0 |
| fc1/42          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |    460053 |           0 |       0 |       0 |
| fc1/43          |        0 |     12 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |     39898 |           0 |       0 |       0 |
| fc1/44          |       62 |  11386 |      1 |       62 |        6 |   0 |  62 |   0 |  67
|  58 |  72 |  45455329 |          23 |       0 |       0 |
| fc1/45          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |     56572 |           0 |       0 |       0 |
| fc1/46          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   0 |   2 |         0 |           0 |       0 |       0 |
| fc1/47          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |    154536 |           0 |       0 |       0 |
| fc1/48          |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   1
|   1 |   1 |         0 |           0 |       0 |       0 |
| port-channel21  |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   3
|   0 |   6 |         0 |           0 |       0 |       0 |
| port-channel22  |       62 |  11386 |      1 |       62 |        6 |   0 |  62 |   0 |  71
|  61 |  77 |  45600392 |          23 |       0 |       0 |
| port-channel31  |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   0
|   0 |   0 |         0 |           0 |       0 |       0 |
| port-channel32  |        0 |     12 |      0 |        0 |        0 |   0 |   0 |   0 |   5
|   3 |   7 |    499951 |           0 |       0 |       0 |
| port-channel41  |        4 |      0 |      0 |     3051 |        0 |   1 |   4 |   5 |   8
|   8 |   8 |         0 |           0 |       0 |       0 |
| port-channel42  |        8 |      1 |      2 |        0 |        0 |   7 |  11 |  10 |  16
|  21 |  11 |      1984 |           6 |       0 |       0 |
| port-channel131 |        0 |      0 |      0 |        0 |        0 |   1 |   0 |   2 |   5
|   7 |   3 | 104595439 |           0 |       0 |       0 |
| port-channel132 |        0 |      0 |      0 |        0 |        0 |   0 |   0 |   0 |   0
|   0 |   0 |         0 |           0 |       0 |       0
```

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

-stats 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 link statistics of a specified interface:

```
switch# ShowIntStats --link-stats fc1/1
2023/07/21 18:27:24 Link Stats:

———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
|       | Link     | Sync | Signal | Invalid | Invalid | NOS | NOS | OLS | OLS | LRR | LRR
 | FEC       | FEC         |        |         |
| Intf  | Failures | Loss | Loss   | Words   | CRCs    | Rx  | Tx  | Rx  | Tx  | Rx  | Tx
 | Corrected | Uncorrected | BB_SCs | BB_SCr |
+———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
| fc1/1 |        3 |    0 |      0 |       0 |       0 |   0 |   3 |   3 |   4 |   0 |   4
 |         0 |           0 |      0 |      0 |
———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
```

The following example displays the link statistics for a range of interfaces:

```
switch# ShowIntStats --link-stats fc1-10
2023/07/21 18:27:37 Link Stats:

———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
|       | Link     | Sync | Signal | Invalid | Invalid | NOS | NOS | OLS | OLS | LRR | LRR
 | FEC       | FEC         |        |         |
| Intf  | Failures | Loss | Loss   | Words   | CRCs    | Rx  | Tx  | Rx  | Tx  | Rx  | Tx
 | Corrected | Uncorrected | BB_SCs | BB_SCr |
+———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
| fc1/1  |        3 |    0 |      0 |       0 |       0 |   0 |   3 |   3 |   4 |   0 |
4 |         0 |           0 |      0 |      0 |
| fc1/2  |        0 |    0 |      0 |       0 |       0 |   0 |   0 |   1 |   0 |
2 |         0 |           0 |      0 |      0 |
| fc1/3  |        1 |    0 |      0 |       0 |       0 |   0 |   1 |   1 |   2 |   2 |
2 |         0 |           0 |      0 |      0 |
| fc1/4  |        1 |    0 |      0 |    3051 |       0 |   0 |   1 |   1 |   2 |   2 |
2 |         0 |           0 |      0 |      0 |
| fc1/5  |        1 |    0 |      0 |       0 |       0 |   0 |   1 |   1 |   2 |   2 |
2 |         0 |           0 |      0 |      0 |
| fc1/6  |        1 |    0 |      0 |       0 |       0 |   1 |   1 |   2 |   2 |   2 |
2 |         0 |           0 |      0 |      0 |
| fc1/7  |        0 |    0 |      0 |       0 |       0 |   0 |   0 |   0 |   0 |   0 |
0 |         0 |           0 |      0 |      0 |
| fc1/8  |        0 |    0 |      0 |       0 |       0 |   0 |   0 |   0 |   0 |   0 |
0 |         0 |           0 |      0 |      0 |
| fc1/9  |        0 |    0 |      0 |       0 |       0 |   0 |   0 |   0 |   0 |   0 |
0 |         0 |           0 |      0 |      0 |
| fc1/10 |        0 |    0 |      0 |       0 |       0 |   0 |   0 |   0 |   0 |   0 |
0 |         0 |           0 |      0 |      0 |

———————+————————+————+————————+————————+————+————+————+————+————+————+————+————+
```

The following example displays the general statistics for a range of interfaces:

```
switch# ShowIntStats --general-stats fc1/1-4
2023/07/21 18:27:57 General Stats:

-----+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------
|       | Frames   | Frames   | C3 Frames | C3 Frames | C2 Frames | C2 Frames | CF Frames
| CF Frames | Mcast     | Mcast     | Bcast     | Bcast     | Ucast     | Ucast     |
| Intf  | Rx        | Tx        | Rx        | Tx        | Rx        | Tx        | Rx
| Tx        | Frames Rx | Frames Tx | Frames Rx | Frames Tx | Frames Rx | Frames Tx | Frames Rx | Frames Tx |
+----+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+
| fc1/1 | 29819972 | 44716994 |  29819972 |  44716994 |         0 |         0 |         0
|         0 |         0 |         0 |         0 |         0 |  29819972 |  44716994 |
| fc1/2 |       71 |      112 |        71 |       112 |         0 |         0 |         0
|         0 |         0 |         0 |         0 |         0 |        71 |       112 |
| fc1/3 |   314596 |   309500 |         0 |         0 |         0 |         0 |    314596
|    309500 |         0 |         0 |         0 |         0 |    314596 |    309500 |
| fc1/4 |       45 |     3045 |         0 |         0 |         0 |         0 |        45
|      3045 |         0 |         0 |         0 |         0 |        45 |      3045 |

-----+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------+-------
```

The following example displays the congestion statistics for a range of interfaces:

```
switch# ShowIntStats --congestion-stats fc1/1-4
2023/07/21 18:28:32 Congestion Stats:

-------+------+------+--------+---------------+----------+--------+--------+--------+-----+-----
|       |      |      |        | TxWait % last | Timeout  | Credit | Active | Active | LRR
 | LRR |
| Intf  | TBBZ | RBBZ | TxWait | 1s/1m/1h/72h  | Discards | Loss   | LR Rx  | LR Tx  | Rx
 | Tx  |
+-------+------+------+--------+---------------+----------+--------+--------+--------+-----+-----+
| fc1/1 |    9 |    4 |      0 |   0%/0%/0%/0% |        0 |      0 |      0 |      0 |   0
 |   4 |
| fc1/2 |    5 |    2 |      0 |   0%/0%/0%/0% |        0 |      0 |      1 |      0 |   0
 |   2 |
| fc1/3 |    5 |    4 |      0 |   0%/0%/0%/0% |        0 |      0 |      0 |      2 |   2
 |   2 |
| fc1/4 |    5 |    4 |      0 |   0%/0%/0%/0% |        0 |      0 |      0 |      2 |   2
 |   2 |

-------+------+------+--------+---------------+----------+--------+--------+--------+-----+-----
```

The following example displays the details of the transceiver (SFP) statistics for a range of interfaces:

```
switch# ShowIntStats --transceiver-stats fc1/1-15
2023/07/21 18:29:06 Transceiver(SFP) Detail Stats:

-----+-------------+-------------+---------+----+--------+------+--------+--------+---------+-----
|        |              | Cisco           | Serial      |      | Nominal   |         |
     |         | Tx       | Rx          | Tx    |
| Intf   | Name         | PID             | Number      | Sync | Bit Rate  | Temp    |
Voltage | Current | Power    | Power       | Fault |
+------+-------------+-------------+---------+----+--------+------+--------+--------+---------+-----+
| fc1/1  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18150XH6 |   in | 14000Mb/s | 45.34C |
3.33V |  7.71mA | -2.81dBm |    -2.90dBm |     0 |
| fc1/2  |   CISCO-AVAGO | DS-SFP-FC32G-SW | AVD2101W00P |   in | 28000Mb/s | 43.39C |
3.30V |  7.50mA |  0.23dBm |    -1.00dBm |     0 |
```

```
| fc1/3  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18283Q6Z |   in | 14000Mb/s | 54.66C |
3.30V |  8.54mA | -3.51dBm |    -3.00dBm |     0 |
| fc1/4  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18150RLK |   in | 14000Mb/s | 53.76C |
3.32V |  7.97mA | -2.90dBm |   -11.78dBm |     0 |
| fc1/5  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18400JV0 |   in | 14000Mb/s | 49.35C |
3.32V |  8.10mA | -3.10dBm |    -2.72dBm |     0 |
| fc1/6  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18020AJZ |   in | 14000Mb/s | 52.11C |
3.31V |  7.85mA | -2.83dBm |    -4.45dBm |     0 |
| fc1/7  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18400DD0 |   no | 14000Mb/s | 54.80C |
3.33V |  8.01mA | -3.20dBm | -28.24dBm-- |     0 |
| fc1/8  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS180116DQ |   no | 14000Mb/s | 52.80C |
3.31V |  8.07mA | -2.98dBm | -27.45dBm-- |     0 |
| fc1/9  | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18400EC5 |   no | 14000Mb/s | 53.55C |
3.32V |  8.06mA | -3.32dBm | -25.09dBm-- |     0 |
| fc1/10 | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18021LXY |   no | 14000Mb/s | 54.33C |
3.31V |  8.04mA | -3.17dBm | -27.96dBm-- |     0 |
| fc1/11 | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18021KTD |   in | 14000Mb/s | 53.63C |
3.31V |  7.94mA | -3.24dBm |    -3.01dBm |     0 |
| fc1/12 | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS180116E6 |   in | 14000Mb/s | 54.29C |
3.31V |  8.02mA | -2.89dBm |    -2.62dBm |     0 |
| fc1/13 | CISCO-FINISAR | DS-SFP-FC16G-SW | FNS18400JVS |   in | 14000Mb/s | 52.54C |
3.31V |  7.85mA | -3.10dBm |    -4.02dBm |     0 |
| fc1/14 | CISCO-FINISAR | DS-SFP-FC32G-SW | FNS21240LWJ |   in | 28000Mb/s | 53.53C |
3.32V |  7.67mA | -1.38dBm |    -1.29dBm |     0 |
| fc1/15 |   CISCO-AVAGO | DS-SFP-FC32G-SW | AVD2101W02S |   in | 28000Mb/s | 48.12C |
3.29V |  7.50mA | -0.30dBm |    -1.41dBm |     0 |

-------+-------------+-----------------+-------------+------+-----------+--------+
Note: Only ports having transceiver are displayed
```

The following example displays the breif details such as device alias for a range of interfaces:

```
switch# ShowIntStats --brief fc1/1-4
2023/07/21 18:31:20 Interface Brief + Device-alias + Peer PWWN + Description:

-------+------+-------+-------+--------+-----+-------+-------+------+-----------------+-----------------+---------
|       |      |       | Admin |        |     |       | Oper  |      |         |         |
Device-alias             |                       |            |
|       |      | Admin | Trunk |        |     | Oper  | Speed |  Port    | logical | Name
 or                  | F Port                |             |
| Intf  | VSAN | Mode  | Mode  | Status | SFP | Mode  | (Gbps) | Channel | Type    |
Switchname                | Peer PWWN               | Description |
+-------+------+-------+-------+--------+-----+-------+-------+------+-----------------+-----------------+---------+
| fc1/1 | 1000 |     F |   off |       up | swl |    F |     16 |      -- |    edge |
                  NA |                      NA |          -- |
| fc1/2 | 1000 |     F |   off |       up | swl |    F |     32 |      -- |    edge |
M9148V-N24_init_qle2742_1 | 21:00:34:80:0d:6d:72:52 |          -- |
| fc1/3 |    1 |     E |    on | trunking | swl |   TE |     16 |      41 |    core |
                  NA |                      NA |          -- |
| fc1/4 |    1 |     E |    on | trunking | swl |   TE |     16 |      41 |    core |
                  NA |                      NA |          -- |

-------+------+-------+-------+--------+-----+-------+-------+------+-----------------+-----------------+---------
Note: Only upto 64 characters of discription are displayed
```

The following example displays the error statistics for a range of interfaces:

```
switch# ShowIntStats --errorsonly fc1/44-48
2023/07/21 18:34:37 Link Stats:
```

```
——————+————————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————

switch# ShowIntStats port-channel144
2023/07/21 18:38:19 Link Stats:

——————+————————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————+————

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
```

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:

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      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 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     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       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     event Falling     event  portgurard
                                     Threshold       Threshold
-------------------------------------------------------------------------------
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      event    Falling     event    PMON
                                            Threshold            Threshold            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     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
-------------------------------------------------------------------------------------------------
```

```
switch# show logging level pmon
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
PMon                  4                          4
```

Note The port monitor process does not display in the list of processes when you run the **show logging level** command. The **show logging level pmon** command must be issued to determine the logging level of port monitor.

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:
===============================

   455555555455554555554555599999999999999999999999999999999999
     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
     7777887777778778777787777877677777777876790000000000000000000000000000000000
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:

```
module-4# show logging onboard txwait

---------------------------------
 Module: 4 txwait count
---------------------------------

-----------------------------
 Show Clock
-----------------------------
2018-11-26 14:33:11

---------------------------------
```

```
 Module: 4 txwait
---------------------------------
 Notes:
     - Sampling period is 20 seconds
     - Only txwait delta >= 100 ms are logged

-------------------------------------------------------------------------------
|   Interface     | Delta TxWait Time     | Congestion | Timestamp               |
|                 | 2.5us ticks | seconds |            |                         |
-------------------------------------------------------------------------------
|   Eth4/1(VL3)   |  2758526    |    6    |      34%   | Mon Nov 26 14:32:28 2018 |
|   Eth4/1(VL3)   |  7982000    |   19    |      99%   | Mon Nov 26 14:32:08 2018 |
|   Eth4/1(VL3)   |  7976978    |   19    |      99%   | Mon Nov 26 14:31:48 2018 |
|   Eth4/1(VL3)   |  7974588    |   19    |      99%   | Mon Nov 26 14:31:28 2018 |
|   Eth4/1(VL3)   |  7970818    |   19    |      99%   | Mon Nov 26 14:31:08 2018 |
|   Eth4/1(VL3)   |  7965766    |   19    |      99%   | Mon Nov 26 14:30:48 2018 |
|   Eth4/1(VL3)   |  7976161    |   19    |      99%   | Mon Nov 26 14:30:28 2018 |
|   Eth4/1(VL3)   |  7538726    |   18    |      94%   | Mon Nov 26 14:30:08 2018 |
|   Eth4/1(VL3)   |  7968258    |   19    |      99%   | Mon Nov 26 14:29:48 2018 |
|      fc4/9      |  7987745    |   19    |      99%   | Mon Nov 26 14:33:08 2018 |
|      fc4/9      |  7991818    |   19    |      99%   | Mon Nov 26 14:32:48 2018 |
|      fc4/9      |  7992774    |   19    |      99%   | Mon Nov 26 14:32:28 2018 |
|      fc4/9      |  7992052    |   19    |      99%   | Mon Nov 26 14:32:08 2018 |
|      fc4/9      |  7991918    |   19    |      99%   | Mon Nov 26 14:31:48 2018 |
|      fc4/9      |  7991993    |   19    |      99%   | Mon Nov 26 14:31:28 2018 |
|      fc4/9      |  7987967    |   19    |      99%   | Mon Nov 26 14:31:08 2018 |
|      fc4/9      |  7992034    |   19    |      99%   | Mon Nov 26 14:30:48 2018 |
|      fc4/9      |  7991966    |   19    |      99%   | Mon Nov 26 14:30:28 2018 |
|      fc4/9      |  7990076    |   19    |      99%   | Mon Nov 26 14:30:08 2018 |
|      fc4/9      |  7991890    |   19    |      99%   | Mon Nov 26 14:29:48 2018 |
```



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, on page 122
- Information About Fibre Channel Interfaces, on page 123
- Guidelines and Limitations, on page 124
- Configuring Fibre Channel Interfaces, on page 127
- Verifying Fibre Channel Interfaces Configuration, on page 137
- Configuration Examples for Fibre Channel Interfaces, on page 143

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://bst.cloudapps.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

Forward Error Correction (FEC) allows you to send frames in a way that the receiver can detect and correct errors without the need of retransmitting the frames if there are any errors in the frames. Using FEC, you can transfer frames over impaired links because of an increased tolerance on the receiver side; in fact, in case of bit errors, FEC allows the receiver to correct them.

Transmitter Training Signal (TTS) provides the capability for FC ports to negotiate the following two capabilities:

1. 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. Allows to use FEC.

For more information on configuring FEC and TTS, see the Configuring FEC, on page 129 section.

Note Modifying the FEC configuration briefly disrupts traffic on the port.

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. 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 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.

Guidelines and Limitations

Port Channel Limitations

Port channels have the following restrictions:

Port Speed Information

- 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.
- Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module (DS-X9648-768K9) and the Cisco MDS 9000 24/10-Port SAN Extension Module (DS-X9334-K9) (Fibre Channel ports) supports 16 Gbps, 10 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps speeds. However, a single 16 Gbps SFP supports only 16 Gbps, 8 Gbps, and 4 Gbps speed and a single 8-Gbps SFP supports only 8 Gbps, 4 Gbps, and 2 Gbps speed. For 10 Gbps speeds, the 10 Gbps SFP supports only 10 Gbps. You must not configure speed values other than the values recommended for these SFPs.

The following table describes the results of adding a member to a port channel for various configurations.

Table 20: Port Channel Configuration and Addition Results

Port Channel Members	Configured Speed		New Member Type	Addition Type	Result
	Port Channel	New Member			
DS-X9448-768K9 and DS-X9334-K9	Auto	Auto max 4000	DS-X9448-768K9 and DS-X9334-K9	Normal	Fail
				Force	Pass
	Auto max 4000	Auto max 4000	DS-X9448-768K9 and DS-X9334-K9	Normal or Force	Pass
	Auto max 4000	Auto max 8000 or auto max 16000	DS-X9448-768K9 and DS-X9334-K9	Normal	Fail
				Force	Pass
	Auto max 8000 or auto max 16000	Auto max 4000	DS-X9448-768K9 and DS-X9334-K9	Normal	Fail
				Force	Pass

Port Channel Members	Configured Speed		New Member Type	Addition Type	Result
	Port Channel	New Member			
DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Auto	Auto	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal or Force	Pass
	Auto	Auto max 8000 or auto max 16000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal	Fail
				Force	Pass
	Auto	Auto max 32000	DS-X9648-1536K9	Normal	Fail
				Force	Pass
	Auto max 8000	Auto max 8000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal or Force	Pass
	Auto max 8000	Auto max 16000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal	Fail
				Force	Pass
	Auto max 8000	Auto max 32000	DS-X9648-1536K9	Normal	Fail
				Force	Pass
	Auto max 16000	Auto max 16000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal or Force	Pass
	Auto max 16000	Auto max 8000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal	Fail
				Force	Pass
	Auto max 16000	Auto max 32000	DS-X9648-1536K9	Normal	Fail
				Force	Pass
DS-X9648-1536K9	Auto max 32000	Auto max 32000	DS-X9648-1536K9	Normal or force	Pass
	Auto max 32000	Auto max 4000, auto max 8000, or auto max 16000	DS-X9448-768K9, DS-X9334-K9, and DS-X9648-1536K9	Normal	Fail
				Force	Pass

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

Configuring Fibre Channel Interfaces

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:

Procedure

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 Early TC During Link Negotiation

MDS 64 Gbps capable FC interfaces enforce strict compliance to the Fibre Channel link negotiation standards. This includes a requirement at 32 Gbps speed that Link Speed Negotiation (LSN) be completed before the Training Complete (TC) indication is sent. However, some devices from other vendors incorrectly require TC indication to be set simultaneously with LSN complete during link negotiation. This configuration option allows the MDS to interoperate with such non-compliant FC devices at 32 Gbps link speed.

Guidelines and Limitations:

- Early TC behavior is disabled by default. This option should be enabled only on interfaces connected to non-compliant devices.
- This option is supported only on Fibre Channel (FC) interfaces configured in F-mode. It can be configured regardless of interface speed but will only be applied during link negotiation at 32 Gbps speed.
- Before performing an In-Service Software Downgrade (ISSD) to a lower software version that does not support this feature, disable this configuration on all interfaces where it is enabled to ensure the downgrade is successful.
- This feature is supported only on the following 64 Gbps MDS Fabric Switches and 64 Gbps MDS 9700 linecards:
  - Cisco MDS 9124V 64-Gbps 24-Port Fabric Switch
  - Cisco MDS 9148V 64-Gbps 48-Port Fabric Switch
  - Cisco MDS 9396V 64-Gbps 96-Port Fabric Switch
  - Cisco MDS 9700 48-port 64-Gbps Switching Module (DS-X9748-3072K9)

Toconfigure early TC behavior for 64 Gbps capable F-Ports on MDS switches, follow these steps:

Before you begin

The switch must be running Cisco MDS NX-OS 9.4(5) or later version.

Procedure

Step 1 switch# **`configure terminal`**

Enters global configuration mode.

Step 2 switch(config)# **interface fc1/1**

Enters interface configuration submode.

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

This command disables the switch interface.

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

Sets the interface to F-port mode.

Step 5 switch(config-if)# **switchport link-negotiation early-tc**

This command enables early TC indication during link negotiation at 32 Gbps speed

Step 6 (Optional) switch(config-if)# **no switchport link-negotiation early-tc**

This command disables early TC indication during link negotiation and restores the standards-based link initialization behavior.

This is the default mode.

Step 7 switch(config-if)# **no shutdown**

This command enables the switch interface to start link negotiation.

Step 8 switch# **show interface fc1/1 | include negotiation**

This command verifies the Early TC During Link Negotiation

```
switch# show interface fc1/1 | include Negotiation
Link negotiation early TC is on
```

Configuring FEC

FEC has the following restrictions:

- FEC is supported on the DS-X9748-3072K9, DS-X9648-1536K9, DS-X9334-K9, and DS-X9448-768K9 modules in the Cisco MDS 9700 Series switch. FEC is also supported on the Cisco MDS 9132T, MDS 9220i, MDS 9396S, MDS 9148T, and MDS 9396T switches.
- FEC fallback[8] is not supported on the Cisco MDS 48-Port 64-Gbps Fibre Channel Switching Module (DS-X9748-3072K9) and Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9) when their 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.
- For Cisco MDS 48-Port 64-Gbps Fibre Channel Switching Module (DS-X9748-3072K9), ensure that both FEC and TTS are configured on local and peer switches to negotiate FEC at 16 Gbps. Also, you must configure the **switchport speed 16000** command first, then configure the **switchport fec** and **switchport fec tts** commands.
- Modifying the FEC configuration briefly disrupts traffic on the port.

[8] When the admin speed is auto and FEC is configured on either side of a link, but the link does not come up in FEC mode.

- FEC cannot be configured when auto speed is selected for operating speeds 2000/4000/8000/16000. However, FEC is always enabled on ports running at 32-Gbps and higher speeds but no configuration is required.
- Ports operating at 32 Gbps or higher speeds automatically negotiate FEC because FEC is required at those speeds. No FEC configuration is necessary as the **switchport fec** and **switchport fec tts** commands are meant only for 16-Gbps speeds where FEC is optional.
- 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 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.
- From Cisco MDS NX-OS Release 8.4(1), FEC admin state has changed from up or down to on or off respectively.

To configure FEC on an interface operating at 16-Gbps fixed speed, perform these steps:

Procedure

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.

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

Note
The **switchport fec** command works only on interfaces that support fixed 16 Gbps and higher speeds and a message stating the same appears when you execute this command.

Enables FEC for the 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 interfaces with fixed 16-Gbps speeds 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:

This example displays the FEC state when FEC is enabled:

```
switch# show interface fc3/15 | i fec
    admin fec state is on
    oper fec state is down
```

This example displays the FEC state when FEC is disabled:

```
switch# show interface fc3/15 | i fec
    admin fec state is off
    oper fec state is down
```

Configuring Rate Mode



Note

- Changing port speed and rate mode disrupts traffic on the port.
- Dedicated and shared rate modes are not supported on interfaces that support 16 Gbps or higher speeds.
- Interfaces that are on modules and switches that support 16 Gbps and higher speeds operate in dedicated mode.

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

Procedure

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).

Taking Interfaces out of Service

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 port channel.
- 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:

Procedure

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 port channel.

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

Disables the interface.

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

Takes the interface out of service.

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

```
switch# show port-resources module 1
Module 1
  Available dedicated buffers for global buffer #0 [port-group 1] are 2618
  Available dedicated buffers for global buffer #1 [port-group 2] are 2149
  Available dedicated buffers for global buffer #2 [port-group 3] are 2150
  Available dedicated buffers for global buffer #3 [port-group 4] are 1102
  Available dedicated buffers for global buffer #4 [port-group 5] are 2150
  Available dedicated buffers for global buffer #5 [port-group 6] are 2150
  Available dedicated buffers for global buffer #6 [port-group 7] are 2150
  Available dedicated buffers for global buffer #7 [port-group 8] are 2150
```

```
 Available dedicated buffers for global buffer #8 [port-group 9] are 2150
 Available dedicated buffers for global buffer #9 [port-group 10] are 2150
 Available dedicated buffers for global buffer #10 [port-group 11] are 2150
 Available dedicated buffers for global buffer #11 [port-group 12] are 2150

Port-Group 1
 Total bandwidth is 64.0 Gbps
 Allocated dedicated bandwidth is 64.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc1/1                                   500       16.0  dedicated
 fc1/2                                    32       16.0  dedicated
 fc1/3                                   500       16.0  dedicated
 fc1/4                                   500       16.0  dedicated

Port-Group 2
 Total bandwidth is 64.0 Gbps
 Allocated dedicated bandwidth is 52.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc1/5                                   500       16.0  dedicated
 fc1/6                                   500       16.0  dedicated
 fc1/7                                   500        4.0  dedicated
 fc1/8                                   500       16.0  dedicated

.
.
.

Port-Group 12
 Total bandwidth is 64.0 Gbps
 Allocated dedicated bandwidth is 64.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc1/45                                  500       16.0  dedicated
 fc1/46                                  500       16.0  dedicated
 fc1/47                                  500       16.0  dedicated
 fc1/48                                  500       16.0  dedicated
```

Configuring Transceiver (SFP) Power Saving

The transceiver power saving feature allows the switch to turn off the power to transceivers (SFPs) when the interface is put into an admin-down (shutdown) state. This reduces the switch's power comsuption resulting in a total power reduction.

You can configure the transceiver power saving feature to cut off the power to the transceiver regardless of the status of the interface.

Note

- When the feature is enabled, SFP is present and port is admin-down , then the power is off.
- When you upgrade to Cisco MDS NX-OS Release 9.4(2), by default the feature is enabled.
- When you downgrade from Cisco MDS NX-OS Release 9.4(2), ensure the feature is disabled.
- When you downgrade from Cisco MDS NX-OS Release 9.4(2) and upgrade back to Cisco MDS NX-OS Release 9.4(2), by default the feature is still disabled.
- Feature is available on these models:
  - Cisco MDS 9700 48-Port 64-Gbps Fibre Channel Switching Module
  - Cisco MDS 9124V
  - Cisco MDS 9148V
  - Cisco MDS 9396V

The feature applies to all types of transceivers in the supported switches and modules.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system transceiver power-control**

Enables the SFP power saving feature. The **no system transceiver power-control** disables the feature

Use the**show show systemtransceiverpower-control** command to display if the feature is enabled or disabled.

```
switch# show system transceiver power-control
Power control is on
```

Verifying Transceiver (SFP) Power Saving

The **show interface** command now displays the system transceiver power-control setting and if there is power to the transceiver.

```
switch# show interface fc1/25
fc1/25 is down (Link failure or not-connected)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:19:00:08:31:0a:b3:d0
    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
      384433 frames input,20817552 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      295817 frames output,27256872 bytes
        0 discards,0 errors
      0 input OLS,3  LRR,0 NOS,0 loop inits
      5 output OLS,3 LRR, 1 NOS, 0 loop inits
    0 Zone drops
    0 FIB drops for ports 25-48
    0 XBAR errors for ports 25-48
    0 Other drop count for ports 25-25
    0 timeout discards, 0 credit loss
    1 link failures, 1 sync losses, 0 signal losses
    6 Transmit B2B credit transitions to zero
    6 Receive B2B credit transitions to zero
    0 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    75 2.5us RxWait due to lack of receive credits
    Percentage RxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Interface last changed at Thu Jan 22 13:28:11 2026
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is ACW28240QUE
      Cisco pid is DS-SFP-FC64G-SW
      Temperature 36.20 C, Voltage 3.30 V, Current 6.30 mA, TxPower -0.89 dBm, RxPower
-40.00 dBm --
      Power control is on, Power is on
```

Configuring FC Abort Sequence Counter

The fc-abort sequence counter helps you monitor the health and performance of your Fiber Channel network on Cisco MDS switches. It does this by counting FC abort sequence (ABTS) frames at the switch egress interface which indicates when FC exchanges are unexpectedly stopped. This counter is displayed as part of the FC interface counters.

The FC abort sequence counter allows the network administrator to:

- Determine if any FC ABTS are being transmitted in the fabric. This indicates exchanges are not completing successfully.
- Identify potential network congestion or timeout problems since FC ABTS are used when an exchange does not complete within a period of time.
- Track trends in abort activity to troubleshoot and optimize performance without interrupting live traffic.

Guidelines and Limitations:

- Fc-abort sequence counter is supported starting from Cisco MDS NX-OS 9.4(5) on the following modules and switches:
  - Cisco MDS 9700 48-port 64-Gbps Switching Module (DS-X9748-3072K9)
  - Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9)
  - Cisco MDS 9396V 64-Gbps 96-Port Fibre Channel switch
  - Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel switch

- Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch
- Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switch
- Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switch
- Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switch

- Before downgrading the switch software (ISSD) to a version earlier than 9.4(5), disable this feature.
- This counter is disabled by default.

To configure fc-abort sequence counter on MDS switches, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters global configuration mode.

Step 2 switch(config)# **system counter fc-aborts**

This command enables the FC abort sequence counter

Step 3 switch(config)# **no system counter fc-aborts**

This command disables the FC abort sequence counter.

Step 4 switch(config)# **end**

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 port channel 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 Fibre Channel interface status:

Note 32-Gbps Fibre Channel 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** *interface-range* **transceiver details** command to display small form-factor pluggable (SFP) diagnostic information.

```
switch# show interface fc1/5 transceiver details
fc1/5 sfp is present
    Name is CISCO-AVAGO
    Manufacturer's part number is SFBR-5780APZ-CS2
    Revision is G2.3
    Serial number is AGD151785V6
    Cisco part number is 10-2418-01
    Cisco pid is DS-SFP-FC8G-SW
    FC Transmitter type is short wave laser w/o OFC (SN)
    FC Transmitter supports short distance link length
    Transmission medium is multimode laser with 62.5 um aperture (M6)
    Supported speeds are - Min speed: 2000 Mb/s, Max speed: 8000 Mb/s
    Nominal bit rate is 8500 Mb/s
    Link length supported for 50/125um OM2 fiber is 50 m
    Link length supported for 62.5/125um fiber is 21 m
    Link length supported for 50/125um OM3 fiber is 150 m
    Cisco extended id is unknown (0x0)

    No tx fault, no rx loss, in sync state, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
-------------------------------------------------------------------------------
                                    Alarms                  Warnings
                                 High        Low         High          Low
-------------------------------------------------------------------------------
  Temperature   50.26 C         75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.35 V          3.63 V      2.97 V      3.46 V        3.13 V
  Current        8.33 mA         8.50 mA     2.00 mA     8.50 mA       2.00 mA
  Tx Power      -2.45 dBm        1.70 dBm  -14.00 dBm   -1.30 dBm    -10.00 dBm
  Rx Power      -4.81 dBm        3.00 dBm  -17.30 dBm    0.00 dBm    -13.30 dBm
  Transmit Fault Count = 0
-------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
```

Beginning with Cisco MDS NX-OS Release 9.3(1) duplicate or Cisco non-Compatible SFPs are also detected. All duplicate SFPs will get error disabled and all duplicate SFP ports are reported in a syslog.

```
switch# show interface fc18/45

fc18/45 is down (Error disabled - Duplicate SFP serial number)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 23:28:8c:60:4f:32:30:80
    Admin port mode is auto, trunk mode is on
.
.
    Transceiver Information:
      Serial number is RS212900040004
      Cisco pid is DS-SFP-FC64G-SW
     Temperature 22.93 C, Voltage 3.34 V, Current 0.00 mA --, TxPower -40.00 dBm --, RxPower
 -40.00 dBm --
```

```
switch# show interface fc1/3-5 brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status       SFP    Oper  Oper   Port     Logical
                 Mode   Trunk                       Mode  Speed  Channel   Type
                        Mode                              (Gbps)
--------------------------------------------------------------------------------
fc1/3       1      E      on      trunking     swl    TE     32     12       core
fc1/4       1      auto   on      notConnected swl    --     --     --       --
fc1/5       101    auto   off     notConnected swl    --     --     --       --
```

Displaying SFP Power Saving

You can use the **show interface** *interface-range* **transceiver details** command to display small form-factor pluggable (SFP) power saving.

Port is Up with power saving feature is enabled

```
switch# show interface fc1/1 transceiver details
fc1/1 sfp is present
    Name is CISCO-INNOLIGHT
    Manufacturer's part number is TR-WR13L-NCI
    Hardware revision is 1B
    Serial number is INL280107P0
    Nominal bit rate is 57800 Mb/s
    Link length supported for 9/125um fiber is 10 km
    FC Transmitter type is long wave laser cost reduced
    FC Transmitter supports long distance link length
    Transmission medium is single mode (SM) laser
    Supported speeds are - Min speed: 16000 Mb/s, Max speed: 64000 Mb/s
    Cisco extended id is none (0x0)
    Cisco vendor id is (0x23)
    Cisco part number is 10-3589-01
    Cisco pid is DS-SFP-FC64G-LW
    Cisco version id is V01
    Firmware version is 0.24
    Power control is on, Power is off

    No tx fault, no rx loss, no sync exists, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
    Power control is on, Power is on

    No tx fault, no rx loss, in sync state, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
---------------------------------------------------------------------------
                                  Alarms                 Warnings
                              High        Low         High          Low
```

```
--------------------------------------------------------------------------------
  Temperature   32.16 C          75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.29 V           3.63 V      2.97 V      3.46 V        3.13 V
  Current        6.09 mA         11.00 mA     2.00 mA    10.00 mA       3.00 mA
  Tx Power      -2.62 dBm         1.70 dBm  -13.00 dBm   -1.30 dBm     -9.00 dBm
  Rx Power      -5.03 dBm         3.00 dBm  -15.90 dBm    0.00 dBm    -11.90 dBm
  Transmit Fault Count = 0
--------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
```

Port is down with power saving feature is enabled

```
switch# show interface fc10/1 transceiver details
fc10/1 sfp is present
    Name is CISCO-INNOLIGHT
    Manufacturer's part number is TR-WR13L-NCI
    Hardware revision is 1B
    Serial number is INL280107P0
    Nominal bit rate is 57800 Mb/s
    Link length supported for 9/125um fiber is 10 km
    FC Transmitter type is long wave laser cost reduced
    FC Transmitter supports long distance link length
    Transmission medium is single mode (SM) laser
    Supported speeds are - Min speed: 16000 Mb/s, Max speed: 64000 Mb/s
    Cisco extended id is none (0x0)
    Cisco vendor id is (0x23)
    Cisco part number is 10-3589-01
    Cisco pid is DS-SFP-FC64G-LW
    Cisco version id is V01
    Firmware version is 0.24
    Power control is on, Power is off

    No tx fault, no rx loss, no sync exists, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
--------------------------------------------------------------------------------
                                    Alarms                  Warnings
                                High        Low         High          Low
--------------------------------------------------------------------------------
  Temperature        NA         NA          NA          NA            NA
  Voltage            NA         NA          NA          NA            NA
  Current            NA         NA          NA          NA            NA
  Tx Power           NA         NA          NA          NA            NA
  Rx Power           NA         NA          NA          NA            NA
  Transmit Fault Count  : NA
--------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
        NA  Not available
```

Port is Up with power saving feature is disabled

```
switch# show interface fc1/17 transceiver details
fc1/17 sfp is present

    Power control is off, Power is on

    No tx fault, no rx loss, in sync state, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
--------------------------------------------------------------------------------
```

```
                                    Alarms                  Warnings
                                 High        Low         High          Low
-------------------------------------------------------------------------------
  Temperature   32.52 C          75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.30 V           3.63 V      2.97 V      3.46 V        3.13 V
  Current        7.07 mA          8.64 mA     2.38 mA     8.07 mA       2.86 mA
  Tx Power      -0.64 dBm         5.00 dBm  -12.20 dBm    2.00 dBm     -8.20 dBm
  Rx Power      -1.93 dBm         5.00 dBm  -15.20 dBm    2.00 dBm    -11.20 dBm
  Transmit Fault Count = 0
-------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
        NA  Not available
```

Port is down with power saving feature is disabled

```
switch# show interface fc1/18 transceiver details
fc1/18 sfp is present

    Power control is off, Power is on

    No tx fault, rx loss, no sync exists, diagnostic monitoring type is 0x68
    SFP Diagnostics Information:
-------------------------------------------------------------------------------
                                    Alarms                  Warnings
                                 High        Low         High          Low
-------------------------------------------------------------------------------
  Temperature   34.54 C          75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.32 V           3.63 V      2.97 V      3.46 V        3.13 V
  Current        0.00 mA   --     12.00 mA     3.00 mA    11.20 mA       3.60 mA
  Tx Power     -40.00 dBm  --      5.00 dBm  -12.20 dBm    4.00 dBm     -8.20 dBm
  Rx Power     -40.00 dBm  --      5.00 dBm  -15.20 dBm    2.00 dBm    -11.20 dBm
  Transmit Fault Count = 0
-------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
        NA  Not available
```

You can also use the **show interface** command to get similar information. Here is an example:

```
switch# show interface fc10/1
fc10/1 is down (Administratively down)
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 22:41:54:7f:ee:ea:6f: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
      1202512 frames input,105165544 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      573110 frames output,31269868 bytes
        0 discards,0 errors
      4 input OLS,1  LRR,0 NOS,0 loop inits
      5 output OLS,4 LRR, 3 NOS, 0 loop inits
    6 Zone drops
    0 FIB drops for ports 1-24
    0 XBAR errors for ports 1-24
```

```
    0 Other drop count for ports 1-1
    0 timeout discards, 0 credit loss
    3 link failures, 0 sync losses, 0 signal losses
    9 Transmit B2B credit transitions to zero
    5 Receive B2B credit transitions to zero
    0 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    25 2.5us RxWait due to lack of receive credits
    Percentage RxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    Interface last changed at Thu Oct 17 15:21:52 2024
    Last clearing of "show interface" counters: never
    Transceiver Information:
      Serial number is INL280107P0
      Cisco pid is DS-SFP-FC64G-LW
      Temperature NA, Voltage NA, Current NA, TxPower NA, RxPower NA
      Power control is on, Power is off
```

Displaying FC Abort Sequence Counter

To view the fc-abort sequence counter use the **show interface** command.

```
switch# show interface fc1/5
fc1/5 is up
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:05:00:2a:6a:64:af:80
    ...
    90 low priority transmit B2B credit remaining
    0 FC-Aborts
```

Configuration Examples for Fibre Channel Interfaces

Configuration Example for FEC Module Interfaces

These steps describe how to configure FEC module interfaces:

Procedure

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
```

Note
For port that is connected to DWDM devices, when the port speed is set to the default speed of **switchport speed auto**, the port may take some time to switch to the new port speed. Hence, set the port speed explicitly using the **switchport**

speed {**1000** | **2000** | **4000** | **8000** | **10000** | **16000** | **32000**} command for such ports to use the new port speed much faster.



Configuring Interface Buffers

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

- Finding Feature Information, on page 146
- Feature History for Interface Buffers, on page 147
- Information About Interface Buffers, on page 148
- Configuring Interface Buffers, on page 164
- Configuration Examples for Interface Buffers, on page 169
- Verifying Interface Buffer Configuration, on page 170
- Troubleshooting Interface Buffer Credits, on page 172

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://bst.cloudapps.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 Interface Buffers

Table 21: Feature History for Interface Buffers

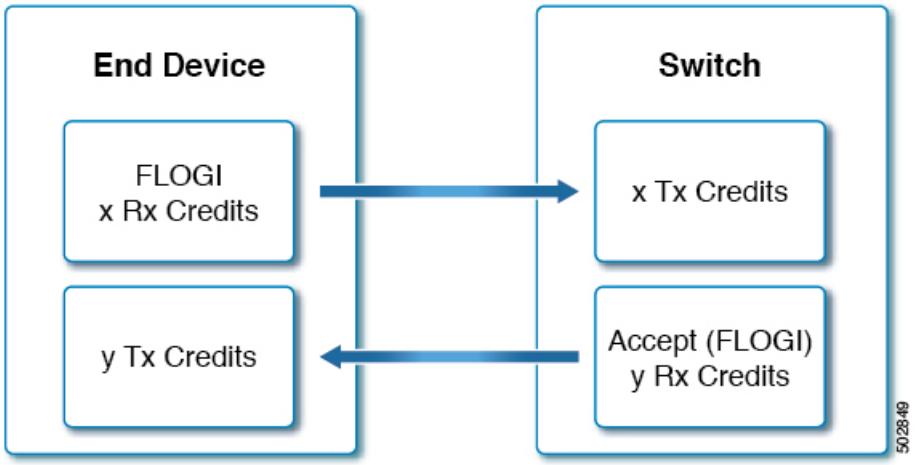
Feature Name	Release	Feature Information
Buffer-to-Buffer Credit Recovery	8.4(1)	Support for buffer-to-buffer credit recovery for NP ports.
Buffer-to-Buffer Credit Recovery	8.2(1)	Support for buffer-to-buffer credit recovery for F ports.
Enhanced Receiver Ready	8.1(1)	This feature was introduced. The following commands were introduced: • **show flow-control er_rdy** • **switchport vl-credit** • **system fc flow-control er_rdy**

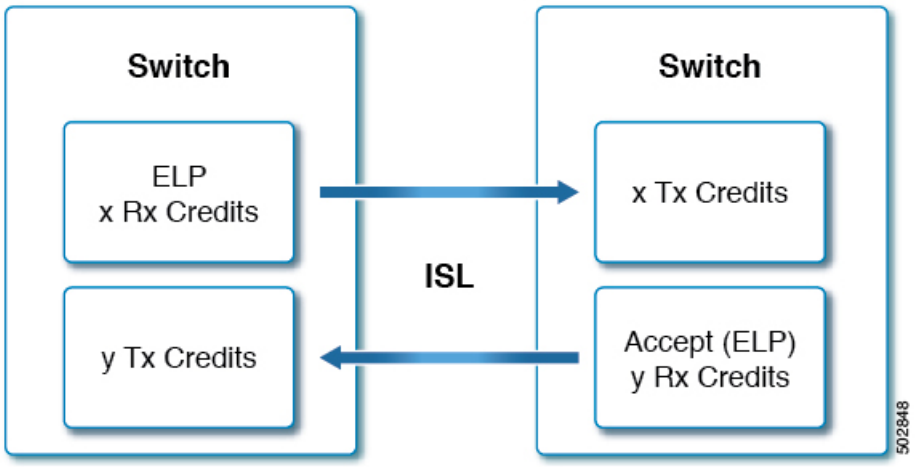
Information About Interface Buffers

Fibre Channel interfaces use buffer to buffer credits to ensure all packets are delivered to their destination without frame drops even if there is congestion in the network.

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a Fibre Channel link-level flow-control mechanism that ensures every frame (a Fibre Channel packet) that is sent has enough buffer space to be received. Each Fibre Channel link can be viewed as two unidirectional links each with their own set of BB_credits. During link initialization, each side informs the other side of the number of receive (Rx) BB_credits it has via the Exchange Link Parameters (ELP) and Accept (ELP) on an E port and FLOGI and Accept (FLOGI) on an F or NP port. When the Rx BB_credit number is received, it is stored as the transmit (Tx) BB_credit number. This way each side's Rx BB_credit number is the other side's Tx BB credit number for each direction on the link.





Each buffer location holds exactly one Fibre Channel frame regardless of size. As a frame is to be transmitted, the sender checks the remaining Tx BB_credit number. If it is greater than 0, a frame can be transmitted. The sender then decrements Tx BB_credit remaining number and transmits the frame. After the frame is received

and processed such that the receiver's buffer location is cleared, the receiver transmits a R_RDY primitive (BB_credit). When the BB credit is received, the sender increments its Tx BB_credit remaining number. This mechanism guarantees that a sender never transmits a frame that the receiver does not have a buffer to hold it in.

Note

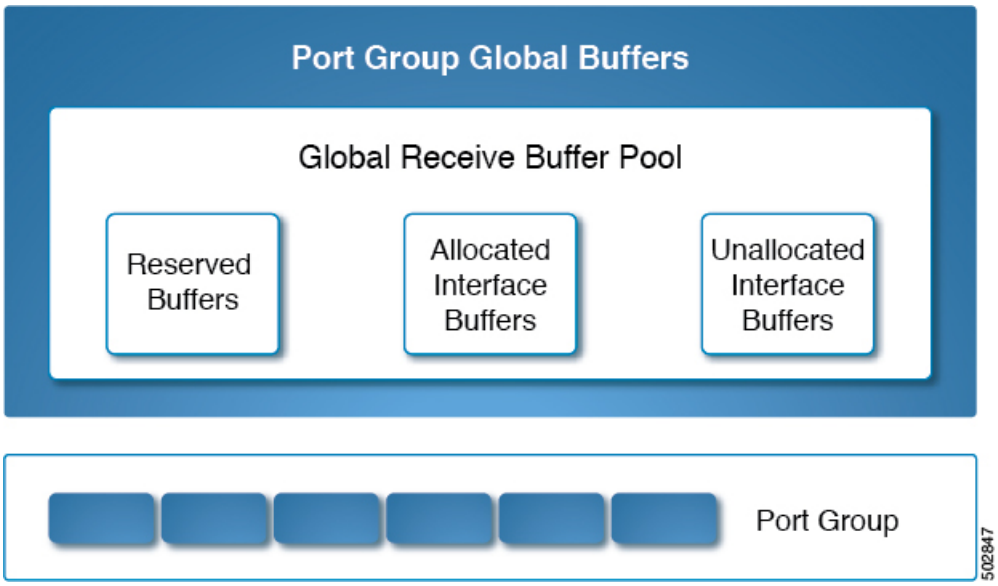
- Cisco MDS switches support mechanisms to avoid loss of R_RDYs which may cause issues on links. For more information, see the Buffer-to-Buffer Credit Recovery, on page 161 section.
- Reconfiguring BB_credits on an active link is a disruptive operation.
- The number of Rx BB_credits does not have to match on each side of the link.
- Only the Rx BB_credits can be configured on an individual interface because those are the only credits that an interface has control over.
- If the transmitter decrements the remaining Tx BB_credit remaining number and hits zero, the *Tx transition to zero counter* will be incremented by one. This typically indicates some level of congestion at the receiving device. Although it could also indicate that there are insufficient buffers for the speed and distance of the link.
- If the receiver does not transmit R_RDYs to the sender of frames, then once the number of frames equals to the Rx BB_credits are received the sender must stop sending since it has hit 0 Tx BB_credit remaining. The receiver will also hit 0 Rx BB_credits remaining and will increment its *Rx transition to zero* counter by one.
- Long-distance links may must have the number of BB_credits increased on both sides to ensure maximum performance.

Global Receive Buffer Pool

A port group is a set of contiguous ports that share common resources such as bandwidth and buffer credits from a global pool of buffers.

The global pool of buffers includes the global receive buffer pool. The global receive buffer pool includes the following buffer groups:

- Reserved internal buffers
- Allocated buffers for each Fibre Channel interface (user configured or assigned by default)
- Unallocated buffers, if any, to be used for additional buffers when required



Extended Buffer-to-Buffer Credits

Extended buffer-to-buffer credits are made possible by allocating extra buffers to specific interfaces. These extra buffers are taken from the unallocated buffer pool.

Note The ENTERPRISE_PKG license is required to use extended buffer-to-buffer credits on 16 Gbps, 32 Gbps, and 64 Gbps switching modules.

All ports on the 16 Gbps and 32 Gbps switching modules support extended buffer-to-buffer credits. There is a limitation on the maximum number of extended buffer-to-buffer credits you can assign to a port. If necessary, you can configure interfaces to use minimum credits to make more extended buffer-to-buffer credits available to other ports.

For long-distance ISLs, the extended buffer-to-buffer credits feature allows you to configure the receive buffers up to the level you need and within hardware resource limits. When necessary, you can reduce the buffers on one port and assign them to another port in the same port group. However, you must have first released the buffers from the other ports before configuring larger extended buffer-to-buffer credits for a port.

Extended BB_credits are typically used on long-distance ISL ports (E ports). If you require additional BB_credits on a port or a group of ports, buffers may need to be made available.

To allow most buffers to be available, perform these steps:

1. Configure ISL ports over different port-groups and modules.
2. Configure ports that are connected to end devices (F ports) from **mode auto** to **mode F** by using the **switchport mode f** command.

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. This will reduce the number of buffers allocated to those ports 500–32 (on most switch types) and add those saved buffers into the unallocated pool.

Without changing the default mode or speed, you can assign the remaining available BB_credits within the port group, using the **switchport fcrxbbcredit extended** *credits* command. Use the **show port-resources module** *slot* command to verify the updated BB_credits allocation among the interfaces.

For example, if there are 300 extended BB_credits available, we can assign these 300 BB_credits to a port that is having 500 BB_credits so that the port now has 800 BB_credits provided all other ports in the port group are configured with speed auto and mode auto.

```
switch# configure terminal
switch(config)# interface fc1/1
switch(config-if)# switchport fcrxbbcredit extended 800
```

When we configure ports in a port group to F Port mode, the reserved buffer credits for such ports are reduced from 500 BB_credits to 32 BB_credits and the remaining BB_credits are assigned to the unallocated interface buffer pool. If you need more BB_credits for that particular port group buffer pool, you can reduce the number of BB_credits being used by the F ports, using the **switchport fcrxbbcredit** *credits* command.

An alternate option is to configure the remaining ports in a port group to minimum credits to free up all BB_credits from these ports for extended BB_credits use.

Note The ENTERPRISE_PKG license is required to use extended buffer-to-buffer credits on 16 Gbps and 32 Gbps switching modules. All ports on the 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 configure interfaces to use minimum credits to make more extended buffer-to-buffer credits available to other ports.

Buffer-to-Buffer Credit Allocation

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

64 Gbps Fibre Channel Switching Module and Fabric Switches

Table 22: 64 Gbps Switching Modules/Fabric Switches Buffer-to-Buffer Credit Allocation

Number of port groups	• Cisco MDS 9700 48-Port 64-Gbps Fibre Channel Switching Module: 2 • Cisco MDS 9124V 24-Port 64-Gbps Fibre Channel Fabric Switch: 1 • Cisco MDS 9148V 48-Port 64-Gbps Fibre Channel Fabric Switch: 2 • Cisco MDS 9396V 96-Port 64-Gbps Fibre Channel Fabric Switch: 4
Default buffer-to-buffer credits	Auto/Mode E ports: 1000
	F ports: 100

Minimum configurable buffers per port	Auto/Mode E ports: 2
	Mode F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port group, when all other ports in the port group are **NOT** configured with minimum BB credit of 1 or 2	0 By default, all ports are in **switchport mode auto** with 1000 BB_credits reserved for each port. Hence, 0 BB_credits will be available in the global buffers pool. Configuring the **switchport fcrxbbcredit** command with a value less than 1000 BB_credits will release the credits less than 1000 BB_credits into the global buffers pool.. Another way of releasing BB_credits into global buffers pool is by configuring the **switchport mode f** command. This will release 900 BB_credits into the global buffer pool for each port that is configured as mode F. After the BB_credits are available in the global buffer pool, extended credits (credits greater than 1000) can be configured. However, the **feature fcrxbbcredit extended** command must be configured first.
Maximum configurable global buffers available per port, when all other ports in the port group are configured with minimum BB credit of 1 or 2.	16000
Default extended credits	3500

Note

- The maximum configurable buffer-to-buffer credits for an F port is 500 credits and for an E port is 1000 credits. However, when a port goes down and comes up in auto mode F, you are allowed to configure more than 500 credits because NX-OS cannot differentiate the port type when in auto mode.
- 16 Gbps and 32 Gbps switching modules or switches consume one BB_credit when ports are moved to 8 Gbps fixed speed. However, the 64 Gbps switching module does not consume BB_credits when ports are moved to 8 Gbps fixed speed.

32 Gbps Switching Modules or Switches

Table 23: 32 Gbps Switching Modules or Switches Buffer-to-Buffer Credit Allocation

Number of port groups	• Cisco MDS 9220i: 1 • Cisco MDS 9132T: 2 • Cisco MDS 9148T: 3 • Cisco MDS 9396T: 6 • Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module: 3

Default buffer-to-buffer credits	Auto/Mode E ports: 500
	F ports: 32
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port group, when all other ports in the port group are **NOT** configured with minimum BB credit of 1 or 2,	300
Maximum configurable global buffers available per port, when all other ports in the port group are configured with minimum BB credit of 1 or 2.	8170
Default extended BB credits	500

16 Gbps Switching Modules or Switches

Table 24: 16 Gbps Switching Modules or Switches Buffer-to-Buffer Credit Allocation

Number of port groups	• Cisco MDS 9396S: 24 • Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module: 12 • Cisco MDS 9700 24/10-Port SAN Extension Module: 6
Default buffer-to-buffer credits	Auto/Mode E ports: 500
	F ports: 32
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port group, when all other ports in the port group are **NOT** configured with minimum BB credit of 1 or 2.	2150

Maximum configurable global buffers available per port, when all other ports in the port group are configured with minimum BB credit of 1 or 2	4095
Default extended BB credits	500

Cisco MDS 9250i and Cisco MDS 9148S Fabric Switches

Table 25: Cisco MDS 9250i and Cisco MDS 9148S Fabric Switch Buffer-to-Buffer Credit Allocation

Number of port groups	• Cisco MDS 9250i: 10 • Cisco MDS 9148S: 12
Default buffer-to-buffer credits	Auto/Mode E ports: 64
	F ports: 64
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Maximum configurable global buffers available per port, when all other ports in the port group are configured with minimum BB credit of 1	253

Note The number of BB credits allocated for Cisco MDS 9250i and MDS 9148s is 64 BB credits per port but can be extended to 253 BB credits when other ports in the port group are moved to minimum credits without the need of an *Enterprise_PKG* license.

Examples: Buffer-to-Buffer Credit Allocation

Note In the command outputs, if the bandwidth is displayed as 32 Gbps, then the output is from either Cisco MDS 9700 switch with Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module, Cisco MDS 9220i, MDS 9132T, MDS 9148T, or MDS 9396T switches.

In the command outputs, if the bandwidth is displayed as 16 Gbps, then the output is from either Cisco MDS 9700 switch with Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module, Cisco MDS 9148S, or Cisco MDS 9250i switches.

The following example displays the default buffers when the switchport mode and speed are set to auto:

```
switch(config)# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 300 Available dedicated
 buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
------------------------------------------------------
Interfaces in the  B2B Credit Bandwidth    Rate Mode
Port-Group         Buffers    (Gbps)
------------------------------------------------------
fc1/1              500         32.0        dedicated
fc1/2              500         32.0        dedicated
fc1/3              500         32.0        dedicated
fc1/4              500         32.0        dedicated
fc1/5              500         32.0        dedicated
fc1/6              500         32.0        dedicated
fc1/7              500         32.0        dedicated
fc1/8              500         32.0        dedicated
fc1/9              500         32.0        dedicated
fc1/10             500         32.0        dedicated
fc1/11             500         32.0        dedicated
fc1/12             500         32.0        dedicated
fc1/13             500         32.0        dedicated
fc1/14             500         32.0        dedicated
fc1/15             500         32.0        dedicated
fc1/16             500         32.0        dedicated
```

The following example displays the buffer allocation when one port is set to E port mode, remaining ports are set to F Port mode, and all ports are set to speed auto:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 7320 Available dedicated
 buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
------------------------------------------------------------
Interfaces in the    B2B Credit  Bandwidth     Rate Mode
Port-Group           Buffers     (Gbps)
------------------------------------------------------------
fc1/1                  500         32.0         dedicated
fc1/2                  32          32.0         dedicated
fc1/3                  32          32.0         dedicated
fc1/4                  32          32.0         dedicated
fc1/5                  32          32.0         dedicated
fc1/6                  32          32.0         dedicated
fc1/7                  32          32.0         dedicated
fc1/8                  32          32.0         dedicated
fc1/9                  32          32.0         dedicated
fc1/10                 32          32.0         dedicated
fc1/11                 32          32.0         dedicated
fc1/12                 32          32.0         dedicated
fc1/13                 32          32.0         dedicated
fc1/14                 32          32.0         dedicated
```

```
fc1/15                  32          32.0          dedicated
fc1/16                  32          32.0          dedicated
```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set to F Port mode with speed auto, 16000, or 32000:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
-------------------------------------------------------------
Interfaces in the   B2B Credit   Bandwidth     Rate Mode
Port-Group          Buffers      (Gbps)
-------------------------------------------------------------
fc1/1               7820         32.0          dedicated
fc1/2               32           32.0          dedicated
fc1/3               32           32.0          dedicated
fc1/4               32           32.0          dedicated
fc1/5               32           32.0          dedicated
fc1/6               32           32.0          dedicated
fc1/7               32           32.0          dedicated
fc1/8               32           32.0          dedicated
fc1/9               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
fc1/1               32           32.0          dedicated
```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set to F Port mode with speed 8000:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 152.0 Gbps
-----------------------------------------------------
Interfaces in the   B2B Credit  Bandwidth   Rate Mode
Port-Group          Buffers     (Gbps)
-----------------------------------------------------
fc1/1               7580        32.0        dedicated
fc1/2               32          8.0         dedicated
fc1/3               32          8.0         dedicated
fc1/4               32          8.0         dedicated
fc1/5               32          8.0         dedicated
fc1/6               32          8.0         dedicated
fc1/7               32          8.0         dedicated
```

```
fc1/8               32          8.0         dedicated
fc1/9               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
fc1/1               32          8.0         dedicated
```

The following example displays the buffer allocation when two ports are set to E port mode with extended buffers, remaining ports are set to F, and all ports are set to speed auto:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
------------------------------------------------------
Interfaces in the   B2B Credit  Bandwidth   Rate Mode
Port-Group          Buffers     (Gbps)
------------------------------------------------------
fc1/1               3926        32.0        dedicated
fc1/2               3926        32.0        dedicated
fc1/3               32          32.0        dedicated
fc1/4               32          32.0        dedicated
fc1/5               32          32.0        dedicated
fc1/6               32          32.0        dedicated
fc1/7               32          32.0        dedicated
fc1/8               32          32.0        dedicated
fc1/9               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
fc1/1               32          32.0        dedicated
```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set out of service:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 94 Available dedicated
 buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
 buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 32.0 Gbps
------------------------------------------------------------------------------
Interfaces in the Port-Group B2B Credit Bandwidth Rate Mode
     Buffers    (Gbps)
------------------------------------------------------------------------------
fc1/1      8191  32.0      dedicated
```

```
fc1/2 (out-of-service)
fc1/3 (out-of-service)
fc1/4 (out-of-service)
fc1/5 (out-of-service)
fc1/6 (out-of-service)
fc1/7 (out-of-service)
fc1/8 (out-of-service)
fc1/9 (out-of-service)
fc1/10 (out-of-service)
fc1/11 (out-of-service)
fc1/12 (out-of-service)
fc1/13 (out-of-service)
fc1/14 (out-of-service)
fc1/15 (out-of-service)
fc1/16 (out-of-service)
```

The following example displays how to allocate maximum BB_credits on Cisco 9148S and 9250i switches:

The following example shows that the port-group 2 on the switch includes ports fc1/5-8 and each port has 64 credits:

```
switch# show port-resources module 1
.
.
.

Port-Group 2
Available dedicated buffers are 0
------------------------------------------------------------------------
Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                    Buffers     (Gbps)
------------------------------------------------------------------------
  fc1/5                              64          16.0    dedicated
  fc1/6                              64          16.0    dedicated
  fc1/7                              64          16.0    dedicated
  fc1/8                              64          16.0    dedicated
```

To allocate maximum BB_credits to port fc1/5, perform these steps:

1. Configure ports fc1/6-8 in the port-group to a minimum BB_Credit of 1:

```
switch# configure
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc1/6-8
switch(config-if)# switchport fcrxbbcredit 1
```

2. Configure port fc1/5 with the maximum BB_credits of 253:

```
switch# configure
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc1/5
switch(config-if)# switchport fcrxbbcredit 253
```

3. Verify the BB_credits allocation on port fc1/5:

```
switch# show port-resources module 1
```

```
.
.
.

Port-Group 2
Available dedicated buffers are 0
------------------------------------------------------------------
Interfaces in the Port-Group    B2B Credit  Bandwidth  Rate Mode
                                  Buffers     (Gbps)
------------------------------------------------------------------
fc1/5                               253       16.0      dedicated
fc1/6                                 1       16.0      dedicated
fc1/7                                 1       16.0      dedicated
fc1/8                                 1       16.0      dedicated
```

Long-Distance ISLs

When long-distance ISLs are required, you must have sufficient BB_credits configured to ensure that the ISL can run at maximum capacity. The simplest formula or the rule of thumb for computing BB_credits for long-distance ISL assumes a full sized Fibre Channel frame of approximately 2 KB and factors in the interface operating speed and one way distance of the ISL.

Note If the average frame size is less than 2 KB, the number of BB_credits must be increased.

Interface Speed	Minimum Number of BB_Credits Required Per Km (One Way)
1 Gbps	0.5 BB_Credit
2 Gbps	1 BB_Credit
4 Gbps	2 BB_credits
8 Gbps	4 BB_credits
16 Gbps	8 BB_credits
32 Gbps	16 BB_credits
64 Gbps	32 BB_credits

As per the table, to operate a 16-Gbps Fibre Channel ISL over 50 km, you would multiply the one way distance (50) times the minimum number of BB_credits per km (8). That is, a 50 km 16-Gbps ISL requires 400 BB_credits when the average frame size is approximately 2 KB. This is the minimum number of BB_credits that are required for the link to function at its best when utilized to its maximum. To accommodate a smaller average frame size than the maximum (full sized) value, more BB_credits would be required proportionally. Since each buffer is for a Fibre Channel frame irrespective of its size, when Fibre Channel frames are not full sized, more BB_credits are required to achieve full link utilization. In this case, an approximate yet simple formula for calculating BB_credits is the following:

BB_credits = (Minimum number of BB_credits required per km for interface speed x One-way distance (km)) / ((Average receive frame size (bytes) / 2150 bytes))

The following example displays the BB_credits calculation for a 16 Gbps link that is 50 km long with an average input frame size of approximately 1 KB (1075 bytes):

(8 BB_credits per km at 16 Gbps x 50 km) / (1075 / 2150) = 800 BB_credits

To take into consideration the actual average input frame size first determine the average frame size by dividing the total input bytes by the total frames input. The average frame size must be determined for the input direction (Rx side) on an interface since the receive BB_credits are being set. The total bytes and frames can be viewed in the **show interface counters** command output.

```
switch# show interface fc 2/7 counters
fc2/7
    5 minutes input rate 1048060640 bits/sec, 131007580 bytes/sec, 94786 frames/sec
    5 minutes output rate 253368512 bits/sec, 31671064 bytes/sec, 47717 frames/sec
    14079632456 frames input, 18624775031572 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    8089598629 frames output, 6040401816628 bytes
      0 discards, 0 errors
    0 timeout discards, 0 credit loss
    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
    15031 Transmit B2B credit transitions to zero
    0 Receive B2B credit transitions to zero
    11192 2.5us TxWait due to lack of transmit credits
    Percentage TxWait not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
    500 receive B2B credit remaining
    481 transmit B2B credit remaining
    481 low priority transmit B2B credit remaining
    Last clearing of "show interface" counters: 2d09h
```

In the above example, the calculation for the average frame size is:

18624775031572 (bytes) / 14079632456 (frames) = 1323 bytes/frame which is approximately 1.3 KB/frame

To complete the calculation:

(8 BB_credits per km at 16 Gbps x 50 km) / (1323 average bytes per frame / 2150 bytes) = 650 BB_credits

Thus, for an 8 Gbps link traversing 50 km carrying an average frame size of 1323 bytes would take a minimum of 650 BB_credits.

Note The required BB_credits for the other end of the same link may be different due to a different average frame size in the opposite direction. The average frame size must be calculated in a similar way from the adjacent interface.

For more information on how to change BB_credits, see the Extended Buffer-to-Buffer Credits, on page 150 section and see the **switchport fcrxbbcredit** *std_bufs* and **switchport fcrxbbcredit extended** *ext_bufs* command outputs.

Buffer-to-Buffer Credit Recovery

Although Fibre Channel standards require low bit and frame error rates, there is a likelihood of errors occurring. 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 using the **no switchport fcbbscn** command.

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 starting from Cisco MDS NX-OS Release 8.2(1).
  - 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.

- NP ports
  - The adjacent N-PortID Virtualization (NPIV) F port must also support this feature. Prior to Cisco MDS NX-OS Release 8.4(1), N-PortID Virtualization (NPIV) ports do not support buffer-to-buffer credit recovery for Cisco N-Port Virtualizer (Cisco NPV) switch logins.
  - This feature is enabled by default on NP ports starting from Cisco MDS NX-OS Release 8.4(1).

The count of times the buffer-to-buffer credits have been recovered for both types of recovery can be displayed using the **show interface counters detailed** command:

Prior to Cisco MDS NX-OS Release 8.4(1a) and earlier release:

```
switch# show interface fc1/1 counters detailed
fc1/1
...
0 BB_SCs credit resend actions, 0 BB_SCr Tx credit increment actions
```

From Cisco MDS NX-OS Release 8.4(2) and later release:

```
switch# show interface fc1/1 counters detailed
fc1/1
...

Congestion Stats:
Tx Timeout discards: 0
Tx Credit loss: 0
BB_SCs credit resend actions: 0
BB_SCr Tx credit increment actions: 0
```

Receive Data Field Size

By default, the maximum data field size is configured for Fibre Channel interfaces and cannot be reconfigured.

Configuring Interface Buffers

Configuring Buffer-to-Buffer Credits

Note When you configure port mode to auto or E 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 270.

Procedure

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 269.

Procedure

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 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 270.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable extended Rx B2B credit configuration:

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

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*

Note
If the ER_RDY flow-control mode is enabled using the **system fc flow-control er_rdy** command, then the configured credits are allocated to individual virtual lanes. For example, if the **switchport fcrxbbcredit extended 1000** command is configured for an interface, the extended buffers for the virtual lanes are configured as **switchport vl-credit extended vl0 16 vl1 16 vl2 47 vl3 921**.

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 269.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable extended Rx B2B credit configuration:

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

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 or enable the buffer-to-buffer credit recovery on a port, perform these steps:

Procedure

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:

- Cisco MDS NX-OS Release 8.4(1) and earlier releases

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

  Note
  The BB_SC_N value is set to the default value of 14.

- Cisco MDS NX-OS Release 8.4(2) and later releases

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

Caution
This command causes traffic disruption on the specified interface.

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:

Procedure

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 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 interface counters detailed** and the **show logging onboard interrupt-stats** commands 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
.
.
.
```

The BB_SCR credit recoveries use the underlying IP_FCMAC_INTR_ERR_BB_SCR_INCREMENT counter and the counter indicates the number of R_RDYs that were lost. The IP_FCMAC_INTR_ERR_BB_SCS_RESEND counter specifies the number of frames that were lost.

Use the **show interface** *port/slot* **counters** command to determine the interval the switch was unable to transmit frames since the counters were last cleared:

```
switch# show interface fc1/13 counters
.
.
.
     6252650 2.5us Txwaits due to lack of transmit credits
.
.
.
```

Txwait value can be converted to seconds using the following formula:

TxWait value in seconds = ((TxWait value in 2.5 μs ticks) x 2.5)/(1,000,000)

Using this formula, we can see that the switch was unable to transmit frames for more than 15 seconds.

Use the **show interface** *port/slot* **counters** command to determine the duration for which the Tx BB credits were zero for the last 1 second, 1 minute, 1 hour, and 72 hours:

```
switch# show interface fc1/13 counters
.
.
.
     Percentage Tx credits not available for last 1s/1m/1h/72h: 1%/5%/3%/2%
.
.
.
```

Use the **show logging onboard txwait module** *number* command to check duration for which the remaining Txwait BB credits were zero over the span of 20 seconds:

```
switch# show logging onboard txwait module 2

---------------------------------
  Module: 2 txwait count
---------------------------------

-----------------------------
  Show Clock
-----------------------------
2019-04-08 13:56:52
Notes:
      - Sampling period is 20 seconds
      - Only txwait delta >= 100 ms are logged

-------------------------------------------------------------------------------
| Interface | Delta TxWait Time     |Congestion| Timestamp                |
|           | 2.5us ticks | seconds |          |                          |
-------------------------------------------------------------------------------
|Eth2/2(VL3)|   882562    |    2    |   11%    | Tue Sep 11 08:52:34 2018|
|Eth2/1(VL3)|  4647274    |   11    |   58%    | Tue Sep 11 08:52:14 2018|
|Eth2/2(VL3)|  7529479    |   18    |   94%    | Tue Sep 11 08:52:14 2018|
|Eth2/1(VL3)|  7829159    |   19    |   97%    | Tue Sep 11 08:51:54 2018|
|Eth2/2(VL3)|  7923544    |   19    |   99%    | Tue Sep 11 08:51:54 2018|
|Eth2/1(VL3)|  5299754    |   13    |   66%    | Tue Sep 11 08:50:34 2018|
|Eth2/2(VL3)|   362484    |    0    |    4%    | Tue Sep 11 08:50:34 2018|
|Eth2/1(VL3)|  7924925    |   19    |   99%    | Tue Sep 11 08:50:14 2018|
|Eth2/2(VL3)|  2566450    |    6    |   32%    | Tue Sep 11 08:50:14 2018|
|Eth2/1(VL3)|  7935558    |   19    |   99%    | Tue Sep 11 08:49:54 2018|
|Eth2/2(VL3)|  6762560    |   16    |   84%    | Tue Sep 11 08:49:54 2018|
|Eth2/1(VL3)|  7908259    |   19    |   98%    | Tue Sep 11 08:49:34 2018|
|Eth2/2(VL3)|  5264976    |   13    |   65%    | Tue Sep 11 08:49:34 2018|
|Eth2/1(VL3)|  7925639    |   19    |   99%    | Tue Sep 11 08:49:14 2018|
```

Use the **show logging onboard error-stats** command o list the ports with zero remaining Tx BB credits for 100 ms:

```
switch# show logging onboard error-stats

-----------------------------
    Module:  1
-----------------------------

-----------------------------
 Show Clock
-----------------------------
2018-08-28 12:28:15

---------------------------------
 Module: 1 error-stats
---------------------------------


--------------------------------------------------------------------------------
 ERROR STATISTICS INFORMATION FOR DEVICE: FCMAC
--------------------------------------------------------------------------------
```

```
    Interface|                                    |         |    Time Stamp
      Range  |     Error Stat Counter Name        |  Count  |MM/DD/YY HH:MM:SS
             |                                    |         |
-------------------------------------------------------------------------------
fc7/2        |IP_FCMAC_CNT_STATS_ERRORS_RX_BAD_   |35806503 |03/17/19 11:32:44
              WORDS_FROM_DECODER
fc7/2        |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO      |2        |03/17/19 11:32:44
fc7/1        |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO      |1        |03/17/19 11:32:44
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |1        |03/15/19 22:10:25
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |16       |03/15/19 18:32:44
fc7/15       |F16_TMM_TOLB_TIMEOUT_DROP_CNT       |443      |03/15/19 15:39:42
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |12       |03/15/19 13:37:59
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |8        |03/15/19 13:29:59
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |4        |03/15/19 13:26:19
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |3        |01/01/17 13:12:14
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |25       |03/14/19 21:13:34
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |21       |03/14/19 21:06:34
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO      |17       |03/14/19 20:58:34
```

Use the **show interface** *port/slot* **bbcredit** command to verify BB credits information:

```
switch# show interface fc1/1 bbcredit
fc1/1 is up
    Transmit B2B Credit is 16
    Receive B2B Credit is 16
      17 receive B2B credit remaining
      16 transmit B2B credit remaining
```

Use the **show interface** *slot/port* **bbcredit** command to check for discrepency in the buffer-to-buffer credit values for a specific Fibre Channel interface:

```
switch# show interface fc2/1 bbcredit
fc2/1 is trunking
    Transmit B2B Credit is 500
    Receive B2B Credit is 500
    Receive B2B Credit performance buffers is 0
      500 receive B2B credit remaining
      500 transmit B2B credit remaining
      500 low priority transmit B2B credit remaining
      500 low priority transmit B2B credit remaining
```

Use the **show interface** *port/slot* **counters** command to display the Tx and Rx BB credit transitions to zero information:

```
switch# show interface fc1/13 counters
.
.
.
     33 Transmit B2B credit transitions to zero
     394351077 Receive B2B credit transitions to zero
.
.
.
```

Use the **show interface** *port/slot* **counters detailed** command to check for credit loss recovery:

Note

- In the **show interface** *port/slot* **counters** [**detailed**] command output, the *Transmit B2B credit transitions to zero* counter increments every time the transmit buffer-to-buffer credits goes to zero. When the ISLs are configured in the TX credit double-queue mode using the **system default tx-credit double-queue** command, some TX B2B credits are reserved for high-priority traffic and remaining credits are used for low-priority traffic from the total TX B2B credits configuration. Hence, when ISLs are in TX credit double-queue mode, this counter does not increment though the low-priority credits go to zero because the high-priority credits are still available.
- This command output is applicable for Cisco MDS NX-OS Release 8.4(2) and later releases. The command output varies if you are using Cisco MDS NX-OS Release 8.4(1a) or earlier releases.

```
switch# show interface fc1/4 counters detailed
fc1/4
  Rx 5 min rate bit/sec:                                                0
  Tx 5 min rate bit/sec:                                                0
  Rx 5 min rate bytes/sec:                                              0
  Tx 5 min rate bytes/sec:                                              0
  Rx 5 min rate frames/sec:                                             0
  Tx 5 min rate frames/sec:                                             0

Total Stats:
  Rx total frames:                                                      9
  Tx total frames:                                                     21
  Rx total bytes:                                                     716
  Tx total bytes:                                                    1436
  Rx total multicast:                                                   0
  Tx total multicast:                                                   0
  Rx total broadcast:                                                   0
  Tx total broadcast:                                                   0
  Rx total unicast:                                                     9
  Tx total unicast:                                                    21
  Rx total discards:                                                    0
  Tx total discards:                                                    0
  Rx total errors:                                                      0
  Tx total errors:                                                      0
 Rx class-2 frames:                                                    0
  Tx class-2 frames:                                                    0
  Rx class-2 bytes:                                                     0
  Tx class-2 bytes:                                                     0
  Rx class-2 frames discards:                                           0
  Rx class-2 port reject frames:                                        0
  Rx class-3 frames:                                                    9
  Tx class-3 frames:                                                   21
  Rx class-3 bytes:                                                   716
  Tx class-3 bytes:                                                  1436
  Rx class-3 frames discards:                                           0
  Rx class-f frames:                                                    0
  Tx class-f frames:                                                    0
  Rx class-f bytes:                                                     0
  Tx class-f bytes:                                                     0
  Rx class-f frames discards:                                           0

Link Stats:
  Rx Link failures:                                                     0
  Rx Sync losses:                                                       0
  Rx Signal losses:                                                     0
  Rx Primitive sequence protocol errors:                                0
```

```
  Rx Invalid transmission words:                                    0
  Rx Invalid CRCs:                                                  0
  Rx Delimiter errors:                                              0
  Rx fragmented frames:                                             0
  Rx frames with EOF aborts:                                        0
  Rx unknown class frames:                                          0
  Rx Runt frames:                                                   0
  Rx Jabber frames:                                                 0
  Rx too long:                                                      0
  Rx too short:                                                     0
  Rx FEC corrected blocks:                                          0
  Rx FEC uncorrected blocks:                                        0
  Rx Link Reset(LR) while link is active:                           0
  Tx Link Reset(LR) while link is active:                           0
  Rx Link Reset Responses(LRR):                                     0
  Tx Link Reset Responses(LRR):                                     1
  Rx Offline Sequences(OLS):                                        0
  Tx Offline Sequences(OLS):                                        1
  Rx Non-Operational Sequences(NOS):                                0
  Tx Non-Operational Sequences(NOS):                                0

Congestion Stats:
  Tx Timeout discards:                                              0
  Tx Credit loss:                                                   0
  BB_SCs credit resend actions:                                     0
  BB_SCr Tx credit increment actions:                               0
  TxWait 2.5us due to lack of transmit credits:                     0
  Percentage TxWait not available for last 1s/1m/1h/72h:    0%/0%/0%/0%
  Rx B2B credit remaining:                                         32
  Tx B2B credit remaining:                                         16
  Tx Low Priority B2B credit remaining:                            16
  Rx B2B credit transitions to zero:                                1
  Tx B2B credit transitions to zero:                                2

Other Stats:
  Zone drops:                                                       0
  FIB drops for ports  1-16:                                        0
  XBAR errors for ports  1-16:                                      0
  Other drop count:                                                 0

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



Congestion Management

This chapter provides information about congestion in a Fibre Channel or Fibre Channel over Ethernet (FCoE) network caused by end devices and provides information about how to identify and avoid or isolate such devices. These devices can be both slow devices and devices that are attempting to overutilize the bandwidth of their links.

- Finding Feature Information, on page 178
- Feature History for Congestion Management, on page 179
- SAN Congestion, on page 189
- About Congestion Management, on page 194
- Guidelines and Limitations for Congestion Management, on page 247
- Configuring Congestion Management, on page 260
- Configuration Examples for Congestion Management, on page 289
- Verifying Congestion Management, on page 301
- Action group example, on page 313
- SMA policy examples, on page 314
- Miscellaneous examples, on page 315

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://bst.cloudapps.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 Management

Table 26: Feature History for Congestion Management

Feature Name	Release	Description	Where Documented
Transmit FC abort counter	9.4(5)	Support has been introduced for per-interface transmit FC Abort Sequence (ABTS) frame counters. This allows monitoring of health and performance of Cisco MDS FC fabrics.	Configuring Fibre Channel Interfaces
Link Negotiation Early Training Complete	9.4(5)	Support has been introduced for enhanced Fibre Channel link negotiation capability for simultaneous TC signaling and LSN completion during Fibre Channel link negotiation at 32 Gbps to interoperate with non-compliant devices.	Configuring Fibre Channel Interfaces
RDF indicates disabled FPM feature during device registration	9.4(5)	RDF indicates FPM feature not enabled while notifying registered devices.	Congestion Management
Test Smart Monitoring and Alerting	9.4(5)	Support has been introduced for generating test syslogs and SNMP traps for SMA. This allows testing of remote automation that processes these alerts.	Congestion Management
Congestion Isolation for over-utilization action	9.4(5)	Actions for congestion isolation for over-utilization are summarized.	Congestion Management
PMON Over-Util support for NPV mode	9.4(5)	Support has been introduced for tx-overutilization counter in NPV mode.	About PMON Transmit Overutilization Counter, on page 220
FPIN without PLOGI	9.4(5)	Support has been introduced to send FPIN notifications to the target port by the Fabric Controller even if PLOGI (Port Login) has not occurred.	Congestion Management

Feature Name	Release	Description	Where Documented
Smart Monitoring and Alerting	9.4(4)	Smart Monitoring and Alertng (SMA) feature is available in the production environment. This feature provides congestion alerting and real time monitoring.	Congestion Management
Fabric Congestion and Diagnostics	9.4(4)	Support to send on demand RDF requests directly to connected HBAs has been added.	Congestion Management
Over-Utilization Counter in PMON	9.4(4)	Support for the following are added: • FPIN peer congestion notification • FPIN warning level	About PMON Transmit Overutilization Counter, on page 220
Over-Utilization Counter in PMON	9.4(3)	Support for overutilization counter and its associated mechanisms are introduced to enhance network performance by providing better detection and management of congestion.	About PMON Transmit Overutilization Counter, on page 220
system transceiver power-control command to automatically power off the transceiver if interface is in admin shutdown.	9.4(2)	This command was introduced.	Configuring Transceiver (SFP) Power Saving , on page 133
RxWait feature for the 64G module and 64G Fabric switches	9.3(2)	RxWait is now displayed in the following commands for interfaces on 64G capable modules/switches: • show interface • show interface counters detailed • show interface rxwait-history • show logging onboard rxwait	Congestion Management
DIRL Supported Devices in Switch and NPV modes	9.3(2)	Updated list of devices that support DIRL in NPV mode for Release 9.3(2).	Congestion Management

Feature Name	Release	Description	Where Documented
HBA Extended Receiver Ready	9.3(1)	Added support on F and NP ports. HBA ER_RDY is in preview (beta) status and not to be used in the production environment.	Congestion Management, on page 177
DIRL NPV Support	9.3(1)	Enhanced to support the switches in NPV mode.	Congestion Management, on page 177
Fabric Notifications	9.2(1)	The Fabric Notification — FPIN and Congestion Signal feature are out of the preview (beta) status and is used in the production environment.	Congestion Management, on page 177
TxWait OBFL	9.2(1)	The TxWait OBFL file size was increased from 512 KB to 8 MB.	Congestion Management, on page 177
Congestion Isolation	8.5(1)	This feature is now handled by Fabric Performance Monitor (FPM). The following commands were introduced: • **feature fpm** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* The following commands were deprecated: • **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**} • **congestion-isolation remove interface** *slot/port*	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Congestion Isolation Recovery	8.5(1)	The Congestion Isolation Recovery feature automatically recovers the flow which was moved to low-priority VL after it was detected as slow back to normal VL; thereby, recovering the flow. The following commands were introduced: • **feature fpm** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* • **port-monitor cong-isolation-recover** {**recovery-interval** *seconds* | **isolate-duration** *hours* **num-occurrence** *number*} The **counter** port monitor command was modified to add the **cong-isolate-recover** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Fabric Notifications	8.5(1)	Fabric Notifications are used to notify end devices of performance impacting conditions and behaviors that affect the normal flow of IO such as link integrity degradation and congestion. The following commands were introduced: • **feature fpm** • **counter txwait warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals** • **fpm congested-device** {**exclude** | **static**} **list** • **member pwwn** *pwwn* **vsan** *id* [**credit-stall**] • **fpm congested-device recover pwwn** *pwwn* **vsan** *id* • **fpm fpin period** *seconds* • **fpm congestion-signal period** *seconds* • **show fpm** {**fpin** | **registration** {**congestion-signal** | **summary**} | **congested-device database** [**exclude** | **local** | **remote** | **static**]} **vsan** *id* • **port-monitor fpin** {**recovery-interval** *seconds* | **isolate-duration** *hours* **num-occurrence** *number*} The **counter** port monitor command was modified to add the **FPIN** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Dynamic Ingress Rate Limiting (DIRL)	8.5(1)	DIRL is used to automatically limit the amount of traffic that is flowing through a switch port that is congested. The following commands were introduced: • **feature fpm** • **fpm dirl** {**exclude list** | **reduction** *percentage* **recovery** *percentage* • **member** {**fc4-feature target** | **interface fc** *slot/port*} • **fpm dirl recover interface fc** *slot/port* • **show fpm** {**dirl exclude** | **fpin vsan** *id* | **ingress-rate-limit** {**events** | **status**} **interface fc***slot/port*} • **port-monitor dirl recovery-interval** *seconds* The **counter** port monitor command was modified to add the **DIRL** port-guard action.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Fibre Channel and Fibre Channel over Ethernet (FCoE)	8.4(1)	The following commands were modified: • The **show hardware internal rxwait-history** [**module** *number* | **port** *number*] command was changed to **show interface** [*interface-range*] **rxwait-history**. • The **show hardware internal txwait-history** [**module** *number* | **port** *number*] command was changed to **show interface** [*interface-range*] **txwait-history**. • The **show process creditmon txwait-history** [**module** *number* [**port** *number*]] command was changed to **show interface** [*interface-range*] **txwait-history**. The following command outputs were modified: • **show interface** *interface-range* **aggregate-counters** • **show interface** *interface-range* **counters** • **show interface** *interface-range* **counters detailed** • **show interface priority-flow-control** • **show interface vfc** *interface-range* **counters detailed**	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
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.	Congestion Management, on page 177
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 Management, on page 177

Feature Name	Release	Description	Where Documented
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**	Congestion Management, on page 177
Congestion Drop Timeout, No-Credit Frame Timeout, and Slow-Port Monitor Timeout Values for Fibre Channel	8.1(1)	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.	Congestion Management, on page 177

Feature Name	Release	Description	Where Documented
Slow Drain Detection and Congestion Isolation	8.1(1)	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 2. Congestion Isolation 3. Port monitor portguard action for Congestion Isolation	Congestion Management, on page 177

SAN Congestion

SAN congestion occurs for one or more of the following three reasons:

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. This flow control feature allows a receiving port to back-pressure the upstream-sending port whenever the receiving port reaches its capacity to accept frames. When an edge device does not accept frames from the fabric for an extended time, it creates a congestion condition in the fabric that is 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 any unrelated flows that use the same shared ISL. This type of congestion can occur in both Fibre Channel and FCoE although the flow control mechanisms are different in each of them. Regardless of the protocol of the device causing the congestion, the congestion can propagate back to the source of the frames via both Fibre Channel and FCoE links.

Fibre Channel uses buffer-to-buffer credits (BB_credits). This is 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 side 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 an 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 from the sender, thereby limiting the rate at which it is receiving frames. If the receiver withholds the BB_credits for a significant amount, it causes congestion on that link. It may also cause congestion in the SAN as well. This BB_credit mechanism works independently in each direction of the traffic flow.

Frames and BB_credits are not sent reliably. If a frame is received that is so corrupt that it cannot be recognized, the receiver of that frame does not return a BB_credit. Or, if a frame is received intact and the BB_credit is returned but it is corrupted in transmission on the link, the receiver of that BB_credit does not recognize it as a BB_credit. In both cases, a transmit credit is lost. Credit Loss Recovery (LR or LRR) results when all the transmit credits are lost over time. The BB_SCN feature is used to recover such lost credits before completely running out of credits and causing congestion. Counts of frames and credits that are returned are periodically exchanged and if there is any discrepancy in the count then credits can be recovered. BB_SCN is available on all ISLs and is extended to F ports from Cisco MDS NX-OS Release 8.2(1). For F ports, the attached device must indicate support for BB_SCN in the FLOGI sent.

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 when it wants the sender to cease sending any frames of that class. PFC pause frames contain a value that is called a quanta. The quanta determines how long a class of traffic is paused. There are two types of PFC pause frames—nonzero quanta and zero quanta. A PFC pause frame with a nonzero 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 nonzero quanta. After the buffers are sufficiently available, 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 of the traffic flow independently of the other.

Devices that do not accept frames at the rate that is 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 mitigate the resulting congestion.

These actions include:

- Drop all or old frames that are queued to the slow drain interface that exceed the configured thresholds.
- Isolate the slow device to a separate logical virtual link on an ISL.
- Reset credits on the affected ports.
- Flap the affected ports.
- Error disable the affected ports.

These Congestion Detection, Congestion Avoidance, and Congestion Isolation features are used to detect slow-drain devices and take appropriate actions on them.

The slow drain condition can be classified in the following four levels:

- Level 3—Indicates severe congestion. Ports are without credits for a continuous amount of time and Credit Loss Recovery is initiated. For an F port, the duration when ports are without credits for a continuous amount of time is 1 second and for an E port it is 1.5 seconds. Credit Loss Recovery involves sending a Fibre Channel Link Credit Reset (LR) primitive to restore the BB_credits on the link in both directions. If the receiver responds with a Link Credit Reset Response (LRR), the credits are restored and the link resumes normal operation.

  If the congestion is severe, LRR may not be returned and the link fails with the *LR failed due to timeout* error. Credit Loss Recovery can be initiated from either side of the link. If MDS is the receiver of the LR (because the adjacent device initiated the Credit Loss Recovery), the only way MDS can return an LRR is when the input buffers of an interface are empty. If the interface still has frames that it had received but was unable to forward to the destination interface, the link fails with the *LR failed nonempty receive queue* error. If LR or LRR sequence is successful, the link returns to normal operation. Even if the link returns to its normal operation, the 1-second or 1.5-second time at zero Tx credit causes severe backwards congestion in the SAN. This backward congestion can work its way back all the way to the source of the frames. Severs or initiators typically see that a large amount of IO errors recorded due to many timeout drops that occur.

  When the link first initializes an LR and LRR, sequence occurs normally and does not indicate a level 3 slow drain condition.

  Although, severe congestion can occur in both Fibre Channel and FCoE the Link Credit Reset (LR or LRR) actions only apply to Fibre Channel.

- Level 2—Indicates moderate congestion that is causing frames to drop because the congestion drop timeout threshold has reached. Each frame that is received on an interface is timestamped. If the frame cannot be transmitted to the appropriate egress port within a congestion drop threshold of a switch, the frame is dropped to prevent excessive internal congestion in the switch. This is typically due to the adjacent device on the egress interface withholding credits (Fibre Channel) or sending PFC pauses. Each dropped frame is part of a SCSI (or other protocol) exchange and causes that exchange to fail. Servers or initiators record IO errors and terminate communication when SCSI exchanges fail. When the path between the initiator and target is over shared infrastructure, for example ISLs, other devices that are utilizing the shared infrastructure also sees timeout drops and large delays in their IO completion times. The congestion drop threshold is 500 ms by default and can be set to as low as 200 ms. The congestion drop threshold can be separately set for Fibre Channel and FCoE ports.
- Level 1 and Level 1.5—Indicates that delay occurs when frames cannot be transmitted immediately out of an egress port due to the port being without Tx buffer-to-buffer credits in Fibre Channel or in an Rx

Pause state for FCoE. The amount of delay is measured by the TxWait counter and can be calculated as a percentage of time. For example, if a port is unable to transmit for 200 ms (not necessarily continuous) in a 1-second interval then the TxWait congestion percentage for that 1-second interval is 20% for the specified interval. Level 1.5 indicates a more severe level of delay and is reserved for TxWait greater than or equal to 30%. Level 1 indicates instances when TxWait is less than 30%.

Almost always, higher levels of slow drain include the lower levels. For example, Level 3 slow drain includes level 2, level 1.5, and level 1 because the lack of ability to transmit causes delay and the delay causes timeout dropped frames. Longer delay causes Credit Loss Recovery to be initiated.

The following terms are used in the document:

- Buffer-to-Buffer (BB) credits (Fibre Channel only): BB_credits are a link flow control mechanism that is used in Fibre Channel. A Fibre Channel frame can only be transmitted if the *remaining Tx credit count* is greater than zero. When the frame is transmitted, the *remaining Tx credit count* is decremented by one. When the receiver of the frame processes the frame, it returns a credit that is called Receiver Ready (R_RDY). When an R_RDY is returned, the frame sender increments the *remaining Tx credit count* by one. If the *remaining Tx credit count* hits zero, no further frames can be transmitted until an R_RDY is received.
- R_RDY (Fibre Channel): A Fiber Channel primitive representing a Buffer-to-Buffer credit. For more information, see Buffer-to-Buffer (BB) credits (Fibre Channel only).
- ER_RDY (Extended R_RDY): A Fiber Channel primitive representing a Virtual Link based Buffer-to-Buffer credit. From Cisco MDS NX-OS 8.1(1), MDS introduced the Congestion-Isolation feature. This feature allows slow-drain devices to be isolated to a slow traffic virtual link (VL2) on an ISL (E port). The ISL must be in the Extended Receiver Ready (ER_RDY) mode for this feature to function. When an ISL is in ER_RDY mode, the link is logically partitioned into four separate virtual links. ER_RDY contains the VL number indicating which VL the BB credit is used for.
- PFC Pause (FCoE only): Priority Flow Control is a class-based flow control mechanism where class-based pause frames are sent to stop the flow of data in one direction for a specific class of service. PFC pause frames contain class bitmap and a value that is called a quanta. The class bitmap specifies which classes, or priorities, the pause frame applies to and the quanta determines how long a class of traffic is paused. There are two types of PFC pause frames: pause frames containing a nonzero quanta and pause frames containing a zero quanta. A PFC pause frame with a nonzero quanta signals the receiver to stop sending frames for the class immediately for a specified amount of time. A PFC pause frame with a zero quanta signals the receiver that it can resume sending frames for the class immediately. A PFC pause frame with a zero quanta can be called an *unpause* or *resume*.
- Transitions to zero (Fibre Channel only): When the *remaining Tx credit count* hits zero, the Tx transition to zero counter is incremented on the Tx side. On the Rx side (the side withholding the BB_credits), the Rx transition to zero counter is incremented. It is important to understand that the amount of time actually at zero *remaining Tx credits* is not represented by this counter. It could be for a short time that does not affect performance or it could be for a longer time that affects performance. Because of this, transitions to zero is not a good measure of congestion.
- TxWait (Fibre Channel and FCoE): TxWait is a measure of time when a port cannot transmit when it has frames queued in it. A port cannot transmit if it is at zero *remaining Tx credit count* (Fibre Channel) or if it has received a PFC pause frame. Each time TxWait increments, the port (or class) is unable to transmit for 2.5 microseconds. TxWait value can be converted to seconds by multiplying it by 2.5 and then dividing by 1,000,000.
- RxWait (FCoE only): RxWait is a measure of time where a port cannot receive frames. A port cannot receive frames if it has transmitted a PFC pause frame (FCoE). Each time RxWait increments, the

port (or class) is unable to receive for 2.5 microseconds. RxWait can be converted to seconds by multiplying it by 2.5 and then dividing by 1,000,000.

- Tx Credit not Available (Fibre Channel only): Tx Credit not Available is a software counter that increments by one when the port is at zero *remaining Tx credits* continuously for 100 ms.

  Timeout-drop (Fibre Channel and FCoE): A frame is dropped as a timeout drop when a received frame is unable to be transmitted out of the egress interface in the configured congestion-drop threshold time. This condition is typically due to congestion at the egress interface that is caused by a lack of Tx BB_credits (Fibre Channel) or in an Rx Pause state (FCoE). The default timeout drop value is 500 ms for both Fibre Channel and FCoE but can be configured to a value as low as 200 ms. Also, the frames that are dropped when the no-credit-drop (Fibre Channel) or pause-drop threshold is reached are also marked as timeout drops.

- Credit Loss Recovery (Fibre Channel only): Credit Loss Recovery occurs when a port is at zero *remaining Tx credits* continuously for 1 second (F or NP port) or 1.5 seconds (E port). When this condition occurs a Link Credit Reset (LR) Fibre Channel primitive is sent to reinitialize the credits (both directions) on the link. If a Link Credit Reset Response (LRR) is returned, all credits are restored and the link resumes to normal operation. If an LRR is not returned, the link fails and must completely reinitialize. For information about reasons for credit-loss-recovery, see Reasons for Credit Loss Recovery, on page 192.

- Link Credit Reset (LR) (Fibre Channel only): LR is a Fibre Channel primitive that is used at link initialization, as well as, to reinitialize BB_credits in both directions on an active link when credits are lost.

- Link Credit Reset Response (LRR) (Fibre Channel only): LRR is a Fibre Channel primitive that is a positive response to an LR.

About SAN Congestion Caused by Transmit Overutilization

Small Computer Systems Interface (SCSI) initiator 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 a 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 data, but if the rate of requests is continuous then the queues of a switch 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 on FCoE. The resulting effects to the SAN can look identical to slow drain, but the root cause is much different since the end device is not actually withholding buffer-to-buffer credits (or sending PFC Pauses). The main mechanism for detecting congestion caused by over utilization is by monitoring the Tx data rate of the end device ports. Port monitor can be used to detect congestion caused by over utilization.

Reasons for Credit Loss Recovery

Credit-loss-recovery can occur for the following distinct reasons:

- Frame or R_RDY corruption or loss: As discussed in the section on the BB_SCN feature, frames, and BB_credits (R_RDYs) can be corrupted and lost on the link. If the BB_SCN feature is negotiated between the end-point devices, then corruption or loss of frames can be detected and recovered as long as the number of lost or corrupted frames or BB_credits is less than the total number of credits over the detection window. If the interface completely runs out of transmit BB_credits either because BB_SCN was not negotiated on or the number of lost or corrupted frames or BB_credits was equal to the number of transmit BB_credits, then credit-loss-recovery is initiated. Frame and BB_credits that are lost or corrupted are due to some physical problem in the link. Check and replace SFPs, fiber cables, and patch panels first. Rarely the switch port or HBA could be a fault.
- Severe congestion: This is due to severe congestion in the end device. The reasons for this vary by end device type along with the OS and application so they cannot be described here.

To determine the reason for the credit-loss-recovery:

- Check for invalid CRCs, invalid transmission words, input errors, and any other signs of corrupted data on the interface with the credit-loss recovery. If there are any of these signs, then it is likely that the problem is due to corrupted or lost frames BB_credits. However, if there are no indications of invalid CRCs, invalid transmission words, or input errors, then the problem still could be due to corrupted or lost frames, or BB_credits. This is because a frame or a BB_credit could be corrupted and/or lost after it is transmitted by the MDS. If this is the case, then MDS would not know that has occurred and would not increment any counters indicating a problem. To check for these types of errors use the **show interface fc** *x/y* **counters detailed** command.
- Check for invalid CRCs, invalid transmission words, input errors, and any other signs of corrupted data on the adjacent device's interface or HBA. You can check for errors at the device itself (for example, at the host or target). Also, you can use the **show rdp fcid** *fcid_id* **vsan** *vsan_id* command to query the adjacent device's HBA for errors. Using this command it can be easily determined if there are invalid CRCs, invalid transmission words, or input errors on data being received from MDS. Note that not all HBAs support the **show rdp fcid** *fcid_id* **vsan** *vsan_id* command.
- Check for non-zero BB_SCN counts on the MDS interface. Non-zero BB_SCN counts indicate that BB_SCN is detecting a loss of some BB_credits or frames and is successfully recovering them. This is a good sign of some BB_credit and/or frames being lost or corrupted. To check for BB_SCN recovery occurrences, use the **show interface fc** *x/y* **counters detailed** command and look for the *BB_SCs credit resend actions* and *BB_SCr Tx credit increment actions* lines in the command output.
- Check if credit-loss-recovery is occurring for the same device on both A and B fabrics at the same or similar times. If that is the case, then it is unlikely that there is a similar physical problem with physical components on both links. The problem is most likely severe congestion being reflected back to the MDS switch port. To check for credit loss recovery occurrences use the **show interface fc** *x/y* **counters detailed** command and look for the *Tx Credit loss* line in the command output.
- Check for common or repetitive times of the day or week when this happens. Frames and BB_credits are not usually corrupted and/or lost only at certain times of the day or days of the week. This is a sign of severe congestion and not of BB_credit or frame loss or corruption.
- If the port experiencing Credit-Loss-Recovery is part of a port-channel (either F port-channel or E port-channel/ISL) and there are more than one port in the same port-channel experiencing Credit-Loss-Recovery, then most likely the problem is due to congestion. This is because the MDS load balances across all the members of a port-channel. Consequently, flows for one or more slow devices will be transmitted across all members in the port-channel and will affect all members. If only a single member of the port-channel is experiencing Credit-Loss-Recovery, then most likely the problem is due to physical components in the link.

About Congestion Management

About Congestion Detection

The following features are used to detect congestion on all slow-drain levels on Cisco MDS switches:

- **All Slow-Drain Levels**

  Display of credits agreed to along with the remaining credits on a port (Fibre Channel only)—The credits that are agreed to in both directions in FLOGI (F ports) and Exchange Link Parameters (ELP) for ISLs are displayed via the **show interface** command. Also, the instantaneous value of the remaining credits is also displayed in the output of the **show interface** command. The credits agreed to is static and unchanging information, at least when the link is up. However, the remaining credit values are constantly changing because each time a frame is transmitted, the Tx remaining count is decremented, and each time a credit is received, the Tx remaining count is incremented. When the remaining credits approach or reach zero, it indicates congestion on that port.

  The following example displays the transmitted and received credits information on 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 transmitted and received credits information on 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 transmitted and received credits information on an E 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
.
.
```

```
.
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
```

- **Level 3**

  Level 3 slow-drain condition is characterized by Fibre Channel BB_credits being unavailable continuously for 1 to 1.5 seconds. This condition causes the Credit Loss Recovery mechanism to be invoked to reinitialize both Tx and Rx credits on a link.

  For links in ER_RDY mode, Credit Loss Recovery link reset will still be initiated if Tx BB_credits are unavailable on virtual links 0, 1, and 3 for 1.5 seconds, and this duration cannot be changed or configured. For VL2, the slow VL, it will be initiated if Tx BB_credits are unavailable for 15 seconds, and this duration cannot be changed or configured.



Note In the ER_RDY mode, Credit Loss Recovery will reset the credits for all VLs.

Level 3 slow-drain condition is almost always accompanied by level 2 and level 1 or level 1.5 slow-drain condition.

Credit Loss Recovery that is initiated by either side of a link can be seen in the following ways:

The following example displays the count of Credit Loss Recovery being initiated by a switch on an interface for R_RDY port:



Note This command output is applicable for Cisco MDS NX-OS Release 8.4(2) and later releases. The command output varies if you are using Cisco MDS NX-OS Release 8.4(1a) or earlier releases.

```
switch# show interface fc1/4 counters detailed
fc1/4
  Rx 5 min rate bit/sec:                                                0
  Tx 5 min rate bit/sec:                                                0
  Rx 5 min rate bytes/sec:                                              0
  Tx 5 min rate bytes/sec:                                              0
  Rx 5 min rate frames/sec:                                             0
  Tx 5 min rate frames/sec:                                             0

Total Stats:
  Rx total frames:                                                      9
  Tx total frames:                                                     21
  Rx total bytes:                                                     716
  Tx total bytes:                                                    1436
  Rx total multicast:                                                   0
  Tx total multicast:                                                   0
  Rx total broadcast:                                                   0
  Tx total broadcast:                                                   0
  Rx total unicast:                                                     9
  Tx total unicast:                                                    21
  Rx total discards:                                                    0
  Tx total discards:                                                    0
  Rx total errors:                                                      0
  Tx total errors:                                                      0
```

```
 Rx class-2 frames:                                                 0
  Tx class-2 frames:                                                0
  Rx class-2 bytes:                                                 0
  Tx class-2 bytes:                                                 0
  Rx class-2 frames discards:                                       0
  Rx class-2 port reject frames:                                    0
  Rx class-3 frames:                                                9
  Tx class-3 frames:                                               21
  Rx class-3 bytes:                                               716
  Tx class-3 bytes:                                              1436
  Rx class-3 frames discards:                                       0
  Rx class-f frames:                                                0
  Tx class-f frames:                                                0
  Rx class-f bytes:                                                 0
  Tx class-f bytes:                                                 0
  Rx class-f frames discards:                                       0

Link Stats:
  Rx Link failures:                                                 0
  Rx Sync losses:                                                   0
  Rx Signal losses:                                                 0
  Rx Primitive sequence protocol errors:                            0
  Rx Invalid transmission words:                                    0
  Rx Invalid CRCs:                                                  0
  Rx Delimiter errors:                                              0
  Rx fragmented frames:                                             0
  Rx frames with EOF aborts:                                        0
  Rx unknown class frames:                                          0
  Rx Runt frames:                                                   0
  Rx Jabber frames:                                                 0
  Rx too long:                                                      0
  Rx too short:                                                     0
  Rx FEC corrected blocks:                                          0
  Rx FEC uncorrected blocks:                                        0
  Rx Link Reset(LR) while link is active:                           0
  Tx Link Reset(LR) while link is active:                           0
  Rx Link Reset Responses(LRR):                                     0
  Tx Link Reset Responses(LRR):                                     1
  Rx Offline Sequences(OLS):                                        0
  Tx Offline Sequences(OLS):                                        1
  Rx Non-Operational Sequences(NOS):                                0
  Tx Non-Operational Sequences(NOS):                                0

Congestion Stats:
  Tx Timeout discards:                                              0
  Tx Credit loss:                                                   0
  BB_SCs credit resend actions:                                     0
  BB_SCr Tx credit increment actions:                               0
  TxWait 2.5us due to lack of transmit credits:                     0
  Percentage TxWait not available for last 1s/1m/1h/72h:  0%/0%/0%/0%
  Rx B2B credit remaining:                                         32
  Tx B2B credit remaining:                                         16
  Tx Low Priority B2B credit remaining:                            16
  Rx B2B credit transitions to zero:                                1
  Tx B2B credit transitions to zero:                                2

Other Stats:
  Zone drops:                                                       0
  FIB drops for ports  1-16:                                        0
  XBAR errors for ports  1-16:                                      0
  Other drop count:                                                 0

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

The following example displays the interface counter information for HBA ER_RDY mode:



Note This command output is applicable for Cisco MDS NX-OS Release 9.3(1) and later releases.

```
switch# show interface fc1/19 counters detailed

fc1/19
  Rx 5 min rate bit/sec:                                    214440352
  Tx 5 min rate bit/sec:                                  13299539744
  Rx 5 min rate bytes/sec:                                   26805044
  Tx 5 min rate bytes/sec:                                 1662442468
  Rx 5 min rate frames/sec:                                    394096
  Tx 5 min rate frames/sec:                                   1182737

Total Stats:
  Rx total frames:                                      229691429454
  Tx total frames:                                      687972064890
  Rx total bytes:                                     14553243684900
  Tx total bytes:                                    961041345018896
  Rx total multicast:                                              0
  Tx total multicast:                                              0
  Rx total broadcast:                                              0
  Tx total broadcast:                                              0
  Rx total unicast:                                     229691429433
  Tx total unicast:                                     687972064797
  Rx total discards:                                               0
 Tx total discards:                                            11544
  Rx total errors:                                                 0
  Tx total errors:                                                 0
  Rx class-2 frames:                                               0
  Tx class-2 frames:                                               0
  Rx class-2 bytes:                                                0
  Tx class-2 bytes:                                                0
  Rx class-2 frames discards:                                      0
  Rx class-2 port reject frames:                                   0
  Rx class-3 frames:                                    229691429406
  Tx class-3 frames:                                    687972064710
  Rx class-3 bytes:                                   14553243684900
  Tx class-3 bytes:                                  961041345018896
  Rx class-3 frames discards:                                      0
  Rx class-f frames:                                               0
  Tx class-f frames:                                               0
  Rx class-f bytes:                                                0
  Tx class-f bytes:                                                0
  Rx class-f frames discards:                                      0

Link Stats:
  Rx Link failures:                                                0
  Rx Sync losses:                                                  0
  Rx Signal losses:                                                0
  Rx Primitive sequence protocol errors:                           0
  Rx Invalid transmission words:                                   0
  Rx Invalid CRCs:                                                 0
  Rx Delimiter errors:                                             0
  Rx fragmented frames:                                            0
  Rx frames with EOF aborts:                                       0
```

```
  Rx unknown class frames:                                                0
  Rx Runt frames:                                                         0
  Rx Jabber frames:                                                       0
  Rx too long:                                                            0
  Rx too short:                                                           0
  Rx FEC corrected blocks:                                                0
  Rx FEC uncorrected blocks:                                              0
  Rx Link Reset(LR) while link is active:                                11
  Tx Link Reset(LR) while link is active:                                 0
  Rx Link Reset Responses(LRR):                                           0
  Tx Link Reset Responses(LRR):                                          22
  Rx Offline Sequences(OLS):                                              0
  Tx Offline Sequences(OLS):                                             21
  Rx Non-Operational Sequences(NOS):                                     11
  Tx Non-Operational Sequences(NOS):                                      0
  BB_SCs credit resend actions:                                           0
  BB_SCr Tx credit increment actions:                                     0

Congestion Stats:
  Tx Timeout discards:                                                    0
  Tx Credit loss:                                                         0
  TxWait 2.5us due to lack of transmit credits for VL 0:                  0
  TxWait 2.5us due to lack of transmit credits for VL 1:                  0
  TxWait 2.5us due to lack of transmit credits for VL 2:                  0
  TxWait 2.5us due to lack of transmit credits for VL 3:           27223344
  Percentage VL3 TxWait for last 1s/1m/1h/72h:                 0%/0%/0%/0%
  Rx B2B credit remaining for VL 0:                                       0
  Rx B2B credit remaining for VL 1:                                      10
  Rx B2B credit remaining for VL 2:                                      10
  Rx B2B credit remaining for VL 3:                                      10
  Tx B2B credit remaining for VL 0:                                       0
  Tx B2B credit remaining for VL 1:                                      10
  Tx B2B credit remaining for VL 2:                                       3
  Tx B2B credit remaining for VL 3:                                       9
  Rx B2B credit transitions to zero for VL 0:                        505072
  Rx B2B credit transitions to zero for VL 1:                             7
  Rx B2B credit transitions to zero for VL 2:                           774
  Rx B2B credit transitions to zero for VL 3:                      32518514
  Tx B2B credit transitions to zero for VL 0:                         31356
  Tx B2B credit transitions to zero for VL 1:                             8
  Tx B2B credit transitions to zero for VL 2:                             8
  Tx B2B credit transitions to zero for VL 3:                      19932348

Other Stats:
  Zone drops:                                                             0
  FIB drops for ports 17-32:                                              0
  XBAR errors for ports 17-32:                                            0
  Other drop count:                                                       0

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

The following example displays instances of Credit Loss Recovery being initiated by a switch in OBFL error-stats:



Note The other slow drain indications displayed that accompany Credit Loss Recovery.

```
switch# show logging onboard error-stats

----------------------------
 Show Clock
```

```
----------------------------
2018-08-22 12:59:20

---------------------------------
 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
         |                                 |         |
-----------------------------------------------------------------------------
fc1/1    |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |14713116 |08/22/18 10:25:15
fc1/1    |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781669  |08/22/18 10:25:15
fc1/1    |FCP_SW_CNTR_CREDIT_LOSS          |18       |08/22/18 10:25:15
fc1/1    |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |13338566 |08/22/18 10:24:55
fc1/1    |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781544  |08/22/18 10:24:55
fc1/1    |FCP_SW_CNTR_CREDIT_LOSS          |10       |08/22/18 10:24:55
fc1/1    |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |11929676 |08/22/18 10:24:35
fc1/1    |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781418  |08/22/18 10:24:35
fc1/1    |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |11881213 |08/22/18 10:24:15
fc1/1    |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781307  |08/22/18 10:24:15
```

The following example displays instances of Credit Loss Recovery failing due to the adjacent device not returning an LR. This causes a link failure:

```
switch# show logging log | i i timeout
…
2018 Aug 17 12:54:59 MDS9710 %PORT-5-IF_DOWN_LINK_FAILURE: %$VSAN 1%$ Interface fc1/2
is down (Link failure Link reset failed due to timeout) port-channel228
2018 Aug 17 13:42:01 MDS9710 %PORT-5-IF_DOWN_LINK_FAILURE: %$VSAN 1%$ Interface fc1/2
is down (Link failure Link reset failed due to timeout)
```

The following example displays LRR received on a port:

```
switch# show interface fc1/1 counters detailed
fc1/1
      27651428465 frames, 59174056872960 bytes received
…
      0 link reset received while link is active             <<<<< Credit Loss Recovery
 initiated from the adjacent device
…
      18 link reset responses received                          <<<<< LRRs received
      0 link reset responses transmitted                        <<<<< LRRs transmitted
```

The following example displays a received LR failing due to severe ingress congestion on that interface:

```
switch# show log last 20
…
2018 Aug 22 10:21:44 MDS9710 %PORT-5-IF_DOWN_LINK_FAILURE: %$VSAN 237%$ Interface fc1/13
 is down (Link failure Link Reset failed nonempty recv queue)
```

- **Level 2**

  Level 2 slow-drain condition indicates that the links are so congested that the received frames that are destined for the congested links cannot be transmitted within the congestion-drop threshold. When this condition occurs, these frames are discarded or dropped as timeout-drops. These dropped frames cause SCSI exchanges to fail at the end hosts. Timeout discards would normally be accompanied by level 1 or level 1.5 congestion.

  Timeout-drops are displayed in the following ways:

  - Count of timeout-drops on an interface

```
switch# show interface fc1/1 counters | i fc | discard
fc1/13
      0 discards, 0 errors, 0 CRC/FCS
      14713116 discards, 0 errors    <<<<< total drops/discards
    14713116 timeout discards, 18 credit loss      <<<<< timeout drops/discards
```

    Discards—Specifies the total output discards or dropped frames. Discards are also known as frame drops.

    Timeout discards—Specifies the total output frames discarded due to congestion-drop or no-credit-drop threshold being reached.

  - Instances of timeout-drops in OBFL error-stats

```
switch# show logging onboard module 1 error-stats

----------------------------
 Show Clock
----------------------------
2018-08-22 17:15:32

---------------------------------
 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
           |                                    |         |
-------------------------------------------------------------------------------
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |14713116 |08/22/18 10:25:15
fc1/1      |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781669  |08/22/18 10:25:15
fc1/1      |FCP_SW_CNTR_CREDIT_LOSS          |18       |08/22/18 10:25:15
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT    |13338566 |08/22/18 10:24:55
fc1/1      |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO   |1781544  |08/22/18 10:24:55
fc1/1      |FCP_SW_CNTR_CREDIT_LOSS          |10       |08/22/18 10:24:55
```

  - Instances of timeout-drops in OBFL flow-control timeout-drops

```
switch# show logging onboard flow-control timeout-drops

---------------------------------
```

```
 Module: 1 flow-control timeout-drops
---------------------------------


----------------------------
 Show Clock
----------------------------
2018-08-22 17:16:57


-------------------------------------------------------------------------------
 ERROR STATISTICS INFORMATION FOR DEVICE DEVICE: FCMAC
-------------------------------------------------------------------------------
 Interface |                                |         |    Time Stamp
   Range   |    Error Stat Counter Name     |  Count  |MM/DD/YY HH:MM:SS
           |                                |         |
-------------------------------------------------------------------------------
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT   |14713116 |08/22/18 10:25:15
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT   |13338566 |08/22/18 10:24:55
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT   |11929676 |08/22/18 10:24:35
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT   |11881213 |08/22/18 10:24:15
fc1/1      |F16_TMM_TOLB_TIMEOUT_DROP_CNT   |11771790 |08/22/18 10:23:55
```

- **Level 1 or Level 1.5**

  Level 1 or level 1.5 slow-drain condition indicates that the interface is without transmit BB_credits at times. The interface can track the exact amount of time an interface is at zero transmit credits, in Fibre Channel and the exact amount of time FCoE class is paused in both directions. When an FCoE interface receives a PFC pause, it cannot transmit in a similar fashion to a Fibre Channel interface when the Fibre Channel interface is at zero transmit credits. This duration of time when an interface cannot transmit credits is called TxWait and is counted in 2.5 micro-second intervals. An FCoE interface transmitting a PFC pause (to prevent the other side from transmitting) is like a Fibre Channel interface not returning BB_credits. This duration of time when an interface cannot receive credits is called RxWait and is also counted in 2.5-micro intervals. Currently, RxWait is only measured for FCoE. In Fibre Channel, this duration of time an interface cannot receive credits is only measured by a software process. It is measured only when the interface is at zero Rx credits remaining for a continuous 100 ms amount of time.

  - Display of credit transitions to zero on a port (Fibre Channel only)—Whenever a port hits zero transmit or receive BB_credits, the transmit (Tx) or receive (Rx) BB_credits transitions to zero is incremented. When the transmit BB_credit transitions to zero is incremented, it indicates that the adjacent device has withheld BB_credits or BB_credits are lost. When the receive BB_credit transitions to zero is incremented, it indicates that the switchport is withholding BB_credit from an adjacent device. These interface counters can increment occasionally under normal conditions. These interface counters do not give any indication of the amount of time the interface was at zero credits. Therefore, these counters are not a preferred indication of congestion on a port. See the TxWait and RxWait counters for a better indication of Tx and Rx congestion on a port.

```
switch# show interface fc1/13 counters
fc1/13
    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 errors, 0 CRC/FCS
      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 errors
  0 timeout discards, 0 credit loss
  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
   0 Transmit B2B credit transitions to zero
   0 Receive B2B credit transitions to zero
    0 2.5us TxWait due to lack of transmit credits
    Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
    32 receive B2B credit remaining
    31 transmit B2B credit remaining
    31 low priority transmit B2B credit remaining
  Last clearing of "show interface" counters: 2d00h
```

Transmit B2B credit transitions to zero – Count of times the interface was at zero Tx B2B credits remaining and unable to transmit. This could be because the adjacent device withheld B2B credits from this interface, credits (or frames which should have generated credits) were lost, or because there were insufficient credits for the speed, average frame size, and distance of the link.
Receive B2B credit transitions to zero – Count of times the interface was at zero Rx B2B credits remaining. This is due to this interface withholding B2B credits.

- Display of the total amount of TxWait and RxWait on an interface. Each increment represents 2.5 microseconds of time an interface was at zero Tx or Rx credits. This can be displayed using the **show interface counters** and **show interface counters detailed** commands.

```
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 0 bits/sec, 0 bytes/sec, 0 frames/sec
    27651428465 frames input, 59174056872960 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 59174056872960 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    907817 frames output, 1942720200 bytes
      0 class-2 frames, 0 bytes
      907817 class-3 frames, 1942720200 bytes
      0 class-f frames, 0 bytes
      14713116 discards, 0 errors
    14713116 timeout discards, 18 credit loss
    0 input OLS, 18 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
     903218 Transmit B2B credit transitions to zero
     743093 Receive B2B credit transitions to zero
      108369199104 2.5us TxWait due to lack of transmit credits
      Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
      32 receive B2B credit remaining
      128 transmit B2B credit remaining
    Last clearing of "show interface" counters: 6w 4d
```

2.5us TxWait due to lack of transmit credits – Count of TxWait ticks in 2.5us since the interface counters have been cleared last. In this example, 108369199104 * 2.5 / 1000000 = 270922.99776 seconds of time the interface has not been able to transmit in the past 6 weeks and 4 days.
Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/0%/0%/0% – Percentage

```
 of TxWait as calculated in the last 1 second, 1 minute, 1 hour, and 72 hour
intervals.
```

- Display of TxWait, RxWait, and percentage Tx and Rx credits not available for the last 1 second, 1 minute, 1 hour, and 72 hour—This can be displayed using the **show interface counters detailed** command.

```
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 0 bits/sec, 0 bytes/sec, 0 frames/sec
    27651428465 frames input, 59174056872960 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 59174056872960 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    907817 frames output, 1942720200 bytes
      0 class-2 frames, 0 bytes
      907817 class-3 frames, 1942720200 bytes
      0 class-f frames, 0 bytes
      14713116 discards, 0 errors
    14713116 timeout discards, 18 credit loss
    0 input OLS, 18 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
     903218 Transmit B2B credit transitions to zero
     743093 Receive B2B credit transitions to zero
      108369199104 2.5us TxWait due to lack of transmit credits
      Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
      32 receive B2B credit remaining
      128 transmit B2B credit remaining
    Last clearing of "show interface" counters: 6w 4d
2.5us TxWait due to lack of transmit credits – Count of TxWait ticks in 2.5us since
 the interface counters have been cleared last. In this example, 108369199104 * 2.5
 / 1000000 = 270922.99776 seconds of time the interface has not been able to transmit
 in the past 6 weeks and 4 days.
Percentage Tx credits not available for last 1s/1m/1h/72h: 0%/0%/0%/0% - Percentage
 of TxWait as calculated in the last 1 second, 1 minute, 1 hour, and 72 hour
intervals.
```

```
switch# show interface vfc1/3 counters

vfc1/3
    3166 fcoe in packets
    460532 fcoe in octets
    3166 fcoe out packets
    1005564 fcoe out octets
    0 2.5 us TxWait due to pause frames for VL3
    0 2.5 us RxWait due to pause frames for VL3
    0 Tx frames with pause opcode for VL3
    0 Rx frames with pause opcode for 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%
```

- Display of histograms showing Tx credit unavailability TxWait (Fibre Channel) and PFC pause (TxWait and RxWait) for the last 60 seconds, 60 minutes, and 72 hours—You can display this information using the **show process creditmon txwait-history** (Fibre Channel) and **show system** {**txwait-history** | **rxwait-history**} (FCoE) commands.

Note From Cisco MDS NX-OS Release 8.4(1), the **show process creditmon txwait-history** and **show hardware internal** {**txwait-history** | **rxwait-history**} command has changed to the **show interface** [*interface-range*] {**txwait-history** | **rxwait-history**} command.

TxWait (or credit unavailability) increments because of lack of transmit BB_credits (Fibre Channel) or because of receiving PFC pause frames (FCoE).

RxWait (currently FCoE only) increments when the interface transmits PFC Pause frames.

There are three graphs for each command and each graph has the most recent second, minute, or hour unit on the X axis:

1. Seconds scale—Indicates the past 60 seconds, where each column represents a second of time. Above the histogram are the amounts of time, in milliseconds, that the ports were unable to transmit and is represented in a vertical format. In the first graph shown, 8 seconds before the command being run, there were 857 ms of TxWait (credit unavailability) in the 1-second interval. The most current second is displayed on the left.
2. Minutes scale—Indicates the past 60 minutes, where each column represents a minute of time. Above the histogram are the amounts of time, in seconds, that the ports were unable to transmit and is represented in a vertical format. In the second graph shown, a minute before the command being run, there was a 22.7 second of TxWait (credit unavailability) in the 1-minute interval. The most current minute is displayed on the left.
3. Hours scale—Indicates the past 72 hours, where each column represents an hour of time. Above the histogram are the amounts of time, in seconds, that the ports were unable to transmit and is represented in a vertical format. In the third graph shown, 24 hours before the command being run, there was a 342 seconds of TxWait (credit unavailability) in the 1-minute interval. And, 52 hours prior, there was a 220 seconds of TxWait in that hour. The most current hour is displayed on the left.

```
switch#  show interface fc1/1 txwait-history | no-more

TxWait history for port fc1/1:
==============================
            8999994              29999999999999999999997
            5888883              18887988888888889999998
     000000076362570000000000000664683546354643578887087000000000
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)
```

```
   245 45         1
   239121000000000800000000000000000000000000000000000000000000000000
   ............................................................
   772006000000000800000000000000000000000000000000000000000000000000
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)


                           3                           2
                           4                           2
     000000000000000000000002000000000000000000000000000000000000000000000000000
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 **show interface** [*interface-range*] **rxwait-history** (FCoE) command is similar to the TxWait history commands except that it shows the amount of time the port was receiving PFC pause; therefore, preventing the adjacent device from transmitting.

```
switch# show interface e1/47 rxwait-history

RxWait history for port Eth1/47:
  ==============================

      788777778777777787777777887787787777787778777777777777877777778
      900999990999999099999990099099099990999099999999999199999990
  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)

             1
   4444444445570000000000000000000000000000000000000000000000000
   ............................................................
   7777777775870000000000000000000000000000000000000000000000000
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)


     2                 1                          1
     7               2 5                          9
     000000000000000006060002000000000000000000000009000000000000000000000001
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)
```

- Display of delta TxWait and RxWait values in 20-second intervals where the delta TxWait is greater than or equal to 100 ms—You can use the **show logging onboard txwait** (Fibre Channel and FCoE) or **show logging onboard rxwait** (FCoE) commands to display the delta TxWait and RxWait values.

  TxWait and RxWait are logged to the persistent log (logging onboard or OBFL) whenever a port accumulates 100 ms or more of TxWait or RxWait in a 20-second interval. If a port accumulates less than 100 ms of TxWait or RxWait, nothing is logged for that 20-second interval.

Note From Cisco MDS NX-OS Release 9.2(1), the TxWait OBFL file size was increased from 512 KB to 8 MB. This requires a **clear logging onboard txwait** in certain situations. For more information, see Cisco MDS 9000 Series Command Reference.

The following information is displayed in the logging onboard TxWait and RxWait:

- Delta TxWait or RxWait ticks—Each tick represents 2.5 microseconds. Because the minimum value logged is the equivalent of 100 ms, the minimum value that is displayed in the output is 40,000.
- Delta TxWait or RxWait in seconds—TxWait value that is multiplied by 2.5 and then divided by 1,000,000 results in the TxWait value, in seconds. The TxWait value is displayed as an integer in the output. Therefore, TxWait value less than 1 second is displayed as 0.
- Congestion Percentage (%)—TxWait or RxWait value that is divided by 20 results in TxWait or RxWait, in seconds. This value gives a quick way of seeing how the congestion was in the 20-second interval.
- Timestamp—Indicates the date and time at the end of a 20-second interval when the delta TxWait was determined.

```
switch# show logging onboard txwait module 2

---------------------------------
  Module: 2 txwait count
---------------------------------

----------------------------
  Show Clock
----------------------------
2019-04-08 13:56:52
Notes:
      - Sampling period is 20 seconds
      - Only txwait delta >= 100 ms are logged

-------------------------------------------------------------------------
| Interface | Delta TxWait Time     |Congestion| Timestamp               |
|           | 2.5us ticks | seconds |          |                         |
-------------------------------------------------------------------------
|Eth2/2(VL3)|   882562    |    2    |   11%    | Tue Sep 11 08:52:34 2018|
|Eth2/1(VL3)|  4647274    |   11    |   58%    | Tue Sep 11 08:52:14 2018|
|Eth2/2(VL3)|  7529479    |   18    |   94%    | Tue Sep 11 08:52:14 2018|
|Eth2/1(VL3)|  7829159    |   19    |   97%    | Tue Sep 11 08:51:54 2018|
|Eth2/2(VL3)|  7923544    |   19    |   99%    | Tue Sep 11 08:51:54 2018|
|Eth2/1(VL3)|  5299754    |   13    |   66%    | Tue Sep 11 08:50:34 2018|
|Eth2/2(VL3)|   362484    |    0    |    4%    | Tue Sep 11 08:50:34 2018|
|Eth2/1(VL3)|  7924925    |   19    |   99%    | Tue Sep 11 08:50:14 2018|
|Eth2/2(VL3)|  2566450    |    6    |   32%    | Tue Sep 11 08:50:14 2018|
|Eth2/1(VL3)|  7935558    |   19    |   99%    | Tue Sep 11 08:49:54 2018|
|Eth2/2(VL3)|  6762560    |   16    |   84%    | Tue Sep 11 08:49:54 2018|
|Eth2/1(VL3)|  7908259    |   19    |   98%    | Tue Sep 11 08:49:34 2018|
|Eth2/2(VL3)|  5264976    |   13    |   65%    | Tue Sep 11 08:49:34 2018|
|Eth2/1(VL3)|  7925639    |   19    |   99%    | Tue Sep 11 08:49:14 2018|

switch# show logging onboard rxwait module 2
```

```
--------------------------------
  Module: 2 rxwait count
--------------------------------

----------------------------
  Show Clock
----------------------------
2019-04-08 13:58:03
Notes:
      - Sampling period is 20 seconds
      - Only rxwait delta >= 100 ms are logged

-------------------------------------------------------------------------------
| Interface | Delta RxWait Time     | Congestion | Timestamp                |
|           | 2.5us ticks | seconds |            |                          |
-------------------------------------------------------------------------------
|Eth2/1(VL7)|  6568902    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL6)|  6568927    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL5)|  6568951    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL4)|  6568975    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL3)|  6569000    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL2)|  6569024    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL1)|  6569050    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/1(VL0)|  6569075    |   16    |      82%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL7)|  7523430    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL6)|  7523455    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL5)|  7523479    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL4)|  7523504    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL3)|  7523528    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL2)|  7523552    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
|Eth2/2(VL1)|  7523578    |   18    |      94%   | Thu Aug  2 14:29:54 2018|
```

- Display of average Tx credit not available in 100-ms intervals—Cisco MDS switches have a software process that runs every 100 ms to check for ports that are in continuous state of 0 Tx credits remaining. The ports that are in the continuous state of 0 Tx credits are displayed in the output of the **show system internal snmp credit-not-available** [**module** *module*] and **show logging onboard error-stats** commands. These commands display 100 ms, 200 ms, or more of continuous state of 0 Tx credits remaining.

  The **show system internal snmp credit-not-available** [**module** *module*] command shows the Tx Credit Not Available alerts from port monitor. The alerts are in 100-ms intervals, as a percentage, of the configured port-monitor polling interval. If the Tx Credit Not Available (tx-credit-not-available) port-monitor counter is not configured in the active policy, no events are displayed.

  The *Duration of time not available* column is the percentage of polling interval where Tx credits were at zero and unavailable. In the command output, for the Event Time, Tue Aug 18 19:41:34 2018, the *Duration of time not available* is 10% and indicates 100 ms (10% of the polling interval of 1 second is 100 ms). At Tue Aug 18 19:52:52 2018, the port-monitor policy was changed so that the tx-credit-not-available counter's polling interval was 10 seconds and the rising-threshold was 20%. The *Duration of time not available* column shows 49% and indicates that almost 5 of the 10 seconds of Tx credits were at zero.

```
switch# show system internal snmp credit-not-available
Module: 1      Number of events logged: 20
-------------------------------------------------------------------------------
Port     Threshold Rising Interval(s) Event Time        Type      Duration of time

         /Falling                                                  not available
-------------------------------------------------------------------------------
```

```
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:41:34 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:42:14 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:42:15 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:42:55 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:42:56 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:44:34 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:44:35 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:48:50 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    20%
                                        19:48:51 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:49:31 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    20%
                                        19:49:32 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:51:42 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:51:43 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:52:51 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    10%
                                        19:52:52 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:53:14 2018
fc1/94   10/0(%)            1          Tue Aug 18       Rising    20%
                                        19:53:15 2018
fc1/94   10/0(%)            1          Tue Aug 18       Falling   0%
                                        19:58:36 2018
fc1/94   20/0(%)            10         Tue Aug 18       Rising    49%
                                        20:20:02 2018
fc1/94   20/0(%)            10         Tue Aug 18       Falling   0%
                                        20:21:45 2018
```

- Display of average Tx credit not available in logging onboard error-stats—The **show logging onboard error-stats** command displays the average Tx credit not available in 100-ms intervals as indicated by the FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO counter. This counter increments by 1 for every 100 ms that an interface is in a continuous state of 0 Tx credits remaining. The increments are recorded in the command output every 20 seconds. Information about other counters is also included in the command output.

```
switch# show logging onboard error-stats

----------------------------
    Module:  1
----------------------------

----------------------------
 Show Clock
----------------------------
2018-08-28 12:28:15

--------------------------------
```

```
 Module: 1 error-stats
---------------------------------


-------------------------------------------------------------------------------
 ERROR STATISTICS INFORMATION FOR DEVICE: FCMAC
-------------------------------------------------------------------------------
    Interface|                                  |         |    Time Stamp
      Range  |     Error Stat Counter Name      |  Count  |MM/DD/YY HH:MM:SS
             |                                  |         |
-------------------------------------------------------------------------------
fc7/2        |IP_FCMAC_CNT_STATS_ERRORS_RX_BAD_ |35806503 |03/17/19 11:32:44
              WORDS_FROM_DECODER
fc7/2        |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO    |2        |03/17/19 11:32:44
fc7/1        |FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO    |1        |03/17/19 11:32:44
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |1        |03/15/19 22:10:25
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |16       |03/15/19 18:32:44
fc7/15       |F16_TMM_TOLB_TIMEOUT_DROP_CNT     |443      |03/15/19 15:39:42
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |12       |03/15/19 13:37:59
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |8        |03/15/19 13:29:59
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |4        |03/15/19 13:26:19
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |3        |01/01/17 13:12:14
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |25       |03/14/19 21:13:34
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |21       |03/14/19 21:06:34
fc7/15       |FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO    |17       |03/14/19 20:58:34
```

- Display of Tx and Rx transitions to zero (Fibre Channel only)—When an interface reaches zero remaining credits in either direction, the *transitions to zero* counter is incremented. This incrementation of the counter indicates that a port is running out of credits, but does not indicate the duration that the port was at zero credits. The port could have been at zero credits momentarily or for a longer time. TxWait provides a better view of the impact of credits running out because it gives the actual duration that the port was at 0 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* counts 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)—Provides a count of PFC pause frames that are sent and received on an interface. PFC pause is a count and includes both PFC pauses with a nonzero quanta (actual pause frames) and PFC pauses with a zero quanta (unpause or resume frames). This count does not give any indication of the amount of time the port is paused. The port could have been paused momentarily or for a longer time. TxWait and RxWait give a better view of the impact of these pause frames because they provide the actual amount of time the port was paused in each direction. PFC pauses can be displayed via the **show interface** and **show interface priority-flow-control** commands.

  The following example displays the *pause* counts in the transmit and receive direction:

```
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 RxPause, TxPause counts and the corresponding RxWait, and TxWait for Ethernet ports used for FCoE:

```
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)
================================================================================

Epo540    Auto   NA    (8)     3   456200000  0      0(0)       152866694355(382166)
Eth2/1    Auto   On    (8)     3   4481929    0      0(0)       5930346153(14825)
...snip
Eth2/48   Auto   Off
Eth3/1    Auto   On    (8)     3   0          0      0(0)       0(0)
...snip
Eth3/6    Auto   Off
Eth3/7    Auto   On    (8)     3   0          0      0(0)       0(0)
```

The show interface command now displays the system transceiver power-control setting and if there is power to the transceiver.

```
switch# show interface fc1/25
fc1/25 is down (Link failure or not-connected)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:19:00:08:31:0a:b3:d0
    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
      384433 frames input,20817552 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      295817 frames output,27256872 bytes
        0 discards,0 errors
      0 input OLS,3  LRR,0 NOS,0 loop inits
      5 output OLS,3 LRR, 1 NOS, 0 loop inits
    0 Zone drops
    0 FIB drops for ports 25-48
    0 XBAR errors for ports 25-48
    0 Other drop count for ports 25-25
    0 timeout discards, 0 credit loss
    1 link failures, 1 sync losses, 0 signal losses
    6 Transmit B2B credit transitions to zero
    6 Receive B2B credit transitions to zero
    0 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
```

```
        75 2.5us RxWait due to lack of receive credits
        Percentage RxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
        Interface last changed at Thu Jan 22 13:28:11 2026
        Last clearing of "show interface" counters: never
        Transceiver Information:
          Serial number is ACW28240QUE
          Cisco pid is DS-SFP-FC64G-SW
          Temperature 36.20 C, Voltage 3.30 V, Current 6.30 mA, TxPower -0.89 dBm,
    RxPower -40.00 dBm --
          Power control is on, Power is on
```

- Slowport monitor (Fibre Channel only)—A threshold value of slowport monitor is specified to detect ports that are at zero transmit credits for a specified continuous duration. When a port is at zero Tx credits continuously for the specified threshold value, the switch records an entry in the slowport-monitor log and in logging onboard. This entry is shown in the **show process creditmon slowport-monitor-events** and **show logging onboard slowport-monitor-events** commands. The entry that is shown in the outputs of these commands is identical, except that the slowport-monitor log only holds the last ten events per port. However, the logging onboard holds the events in chronological order and can hold more events when compared to the slowport-monitor log.

  Events are recorded at a maximum frequency of 100 ms. When the count goes up, operational delay is displayed in the command output. Operational delay indicates the length of time when the port was at zero Tx credits. If the count goes up by more than one from the previous entry, then the operational delay is the average operational delay from multiple events in the 100 ms interval.

  In the following example, at 02/02/18 18:12:37.308 the slowport detection count was 276 and the previous value was 273. This example indicates that there were three intervals of time in the previous 100 ms where the port was at zero Tx credits for 1 ms or more. The average time the port was at zero credits is shown in the *oper delay* column (4 ms). Oper delay of 4 ms indicates that there was a total of 12 ms of time when the port was at zero Tx credits in the previous 100 ms. The 12-ms duration was in three separate intervals.

  Port monitor can also generate a slowport-monitor alert by using port monitor. By default, slowport-monitor alert is set to off. Slowport-monitor must be configured to get the port-monitor slowport-monitor alerts.

  The **show process creditmon slowport-monitor-events** [**module** *number*] [**port** *number*] command shows the last ten 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 per module.

```
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
```

- RxWait (FCoE only)—It is a measure of time that a port is in a transmit PFC pause state that is preventing the adjacent device from transmitting to the port. RxWait increments by 1 every 2.5 microseconds that a port is unable to receive.

  RxWait is shown in the following ways:

  - Cumulative count—Indicates the time the interface counters were last cleared, using the **show interface counters**, **show interface counters detailed** , and **show interface priority-flow-control** commands.
  - Count in percent—Indicates when the credits are unable to transmit for the last 1 second, 1 minute, 1 hour, and 72 hours, using the **show interface counters** and **show interface counters detailed** commands.
  - Graphical representation of the count for the last 60 seconds, 60 minutes, and 72 hours—In FCoE, the count is displayed using the **show interface** [*interface-range*] **rxwait-history** command.
  - On-Board Failure Log (OBFL)—An entry in the OBFL when a port accumulates 100 ms or more RxWait in a 20-second interval. This entry is displayed using the **show logging onboard rxwait** command.

  In the following example, the **show interface counters** command output displays the data “1104349910 2.5 us TxWait due to pause frames (VL3).” This data is cumulative from the time when the counters were last cleared or from the time when the module first came up. In this example, TxWait is incremented

1104349910 times. This data converted to seconds is (1104349910 * 2.5) / 1000000 = 2760.874 seconds. The VFC port channel was unable to transmit for 2760.874 seconds.

In the following example, the **show interface counters** command output shows the data "205484298144 2.5 us RxWait due to PFC Pause frames (VL3)." This data is cumulative from the time the counters were last cleared or from the time when the module first came up. In the example, RxWait is incremented 205484298144 times. This data converted to seconds is (205484298144 * 2.5) / 1000000 = 513710.745 seconds. The VFC port channel was unable to receive for 513710.745 seconds.

The following example also shows the percentage of time that 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 that the VFC received PFC pauses. For RxWait, this is the percentage of time that the VFC was sending pause frames preventing the other side from transmitting. In this example, in the last one minute the VFC was prevented from transmitting (TxWait) 33% of the time (20 seconds).

Note When the interface displayed is a VFC port channel or a VFC bound to an Ethernet port-channel, all values are cumulative for all members in the Ethernet port-channel.

```
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%
```

The **show logging onboard error-stats** command has several different counters that pertain to SAN congestion. Most of these counters are module or switch dependent. For information about tx-credit-not-available or rx-credit-not-available, the following counters are used:

- 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 BB_credits for 100 ms. This count typically indicates congestion at the device that is attached to an 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 BB_credits for 100 ms. This count typically indicates that the switch is withholding R_RDY primitive to a device attached to an interface of a switch due to congestion in the path to devices with which it is communicating.

  Also, port monitor can generate tx-credit-not-available alerts (Fibre Channel only). See the Port Monitor section.

- Overutilization—Configuring port monitor with the Tx datarate and Rx datarate counters allow the MDS to issue alerts, syslog entries, and record entries in the output of the **logging onboard datarate** command. In an all MDS environment, only Tx datarate is required to determine overutilization. In mixed environments, with other types of switches that do not support Tx datarate, configuring Rx datarate can help to determine the ingress rate from a non-MDS switch.

  Tx datarate and Rx datarate must be configured as follows and included in the active port-monitor policy:

```
counter tx-datarate poll-interval 10 delta rising-threshold 80 event 4 falling-threshold
 79 event 4
counter rx-datarate poll-interval 10 delta rising-threshold 80 event 4 falling-threshold
 79 event 4
```

  In the **show logging log** and **show logging onboard datarate** commands, the time an interface was running at a high Tx utilization is the time from rising threshold to falling threshold.

```
switch# show logging log
2018 Aug 24 13:09:07 %PMON-SLOT1-3-RISING_THRESHOLD_REACHED: TX Datarate has reached
the rising threshold (port=fc1/4 [0x1003000], value=820766704) .
2018 Aug 24 13:09:09 %PMON-SLOT12-5-FALLING_THRESHOLD_REACHED: TX Datarate has reached
 the falling threshold (port=fc12/11 [0x158a000], value=34050354) .
2018 Aug 24 13:09:18 %PMON-SLOT1-5-FALLING_THRESHOLD_REACHED: TX Datarate has reached
the falling threshold (port=fc1/4 [0x1003000], value=233513787) .
2018 Aug 24 13:09:42 %PMON-SLOT12-3-RISING_THRESHOLD_REACHED: TX Datarate has reached
the rising threshold (port=fc12/11 [0x158a000], value=878848923) .
2018 Aug 24 13:10:45 %PMON-SLOT12-5-FALLING_THRESHOLD_REACHED: TX Datarate has reached
 the falling threshold (port=fc12/11 [0x158a000], value=387111312) .
```

```
switch# show logging onboard datarate

----------------------------
    Module:  1
----------------------------

---------------------------------
 Module: 1 datarate
---------------------------------

----------------------------
 Show Clock
----------------------------
2018-08-28 15:43:33

---------------------------------
 Module: 1 datarate
---------------------------------
      - DATARATE INFORMATION FROM FCMAC

 ------------------------------------------------------------------------------------
| Interface |  Speed  |     Alarm-types      |   Rate  |         Timestamp         |
 ------------------------------------------------------------------------------------
|  fc1/94   |    4G   |  TX_DATARATE_FALLING |    57%  | Tue Aug 28 15:42:52 2018 |
|  fc1/94   |    4G   |  TX_DATARATE_RISING  |    86%  | Tue Aug 28 15:38:54 2018 |
|  fc1/94   |    4G   |  TX_DATARATE_FALLING |     8%  | Tue Aug 28 15:38:33 2018 |
|  fc1/94   |    4G   |  TX_DATARATE_RISING  |    85%  | Tue Aug 28 15:37:42 2018 |
```

Port Monitor

- Port monitor (Fibre Channel only)—Port monitor can generate alerts for various congestion-related counters. Port monitor has two thresholds that are called rising threshold and falling threshold. The rising threshold is when the counter of a port reaches or exceeds the configured threshold value. The falling threshold is when the counter of a port reaches or falls below a configured value. For each event, an alert is generated. The time that the port was between the rising threshold and falling threshold is when the event was occurring. These alerts are recorded in the RMON log in all releases.
- Port monitor does not have any effect on logging of the various congestion counters except in the case of tx-datarate and rx-datarate. From Cisco MDS NX-OS 8.2(1) and later releases, the alerts are logged in OBFL and are displayed in the **show logging onboard datarate** command. See the "Overutilization" section for the optimal tx-datarate and rx-datarate counter configuration to detect overutilization.

Table 27: Features to Detect Slow Drain, on page 216 describes the features that help detect the slow-drain condition:

Table 27: 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 seconds for core ports.
Port monitor's invalid-crc counter	Invalid-crc counter represents the total number of CRC errors that a port receives.
Port monitor's invalid-words counter	Invalid-words counter represents the total number of invalid words that a port receives.
Port monitor's link-loss counter	Link-loss counter represents the total number of link failures that a port encounters.
Port monitor's lr-rx counter	Lr-rx counter represents the total number of link reset primitive sequences that a port receives.
Port monitor's lr-tx counter	Lr-tx counter represents the total number of link reset primitive sequences that a port transmits.
Port monitor's rx-datarate counter	Rx-datarate counter receives frame rates 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 are no transmit buffer-to-buffer credits available for a duration of 100 ms.

Feature Name	Description
Port monitor's timeout-discards counter	Timeout-discards counter represents the total number of frames that are 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 that are 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 counter 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 a port experiences no transmit credit available (Tx B2B = 0) and frames are waiting for transmission.
Port monitor's tx-datarate-burst counter	Tx-datarate-burst counter monitors the number of times the datarate crosses the configured threshold datarate in 1 second intervals.
Port monitor's rx-datarate-burst counter	Rx-datarate-burst counter monitors the number of times the datarate crosses the configured threshold datarate in 1 second intervals.

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 avoid 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 39 section.

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 219 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

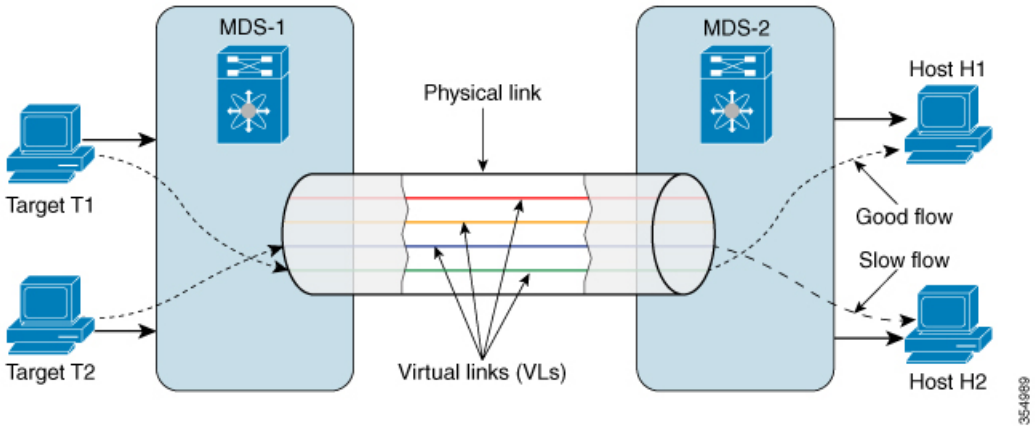


Table 28: Virtual Link-to-QoS Priority Mapping, on page 219 provides VL-to-QoS priority mapping information. Use this information while setting a zone QoS priority in a zone where Congestion Isolation is enabled in order to avoid QoS priority flow from being treated as slow flow.

Table 28: 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 congested 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 congested device is routed to a low-priority VL. Normal devices continue to use the normal VL that has more buffer-to-buffer credits. Congested devices can be marked as slow either via the port monitor or manually.

Note Prior to Cisco MDS NX-OS Release 8.5(1), when a device is manually marked as a congested device or automatically detected as a congested device via the port monitor, the Fibre Channel Name Server (FCNS) database registers the congested-device attribute (slow-dev) for the device and distributes the information to the entire fabric. For more information, see Configuring Congestion Isolation, on page 271.

From Cisco MDS NX-OS Release 8.5(1), when a device is manually marked as a congested device or automatically detected as a congested device via the port monitor, the information about the congested device will be displayed in the FPM database and FPM distributes this information to the entire fabric. For more information, see Configuring Congestion Isolation, on page 271.

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 congested devices.

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

About PMON Transmit Overutilization Counter

Benefits

The transmit over-utilization (tx-overutilization) port monitor counter allows detection and alerting when an FC interface is

- operating at or near its maximum transmit capacity, and

- receiving incoming packets more than its max capability of what it can actually process and potentially causing congestion in the network path.

Architecture

The tx-overutilization counter detects congestion occurring on an F port when the required transmit data rateexceeds the capacity of the interface at its current operational speed. This situation causes the network fabric to experience backpressure, leading to congestion.

The tx-overutilization counter operates by simultaneously monitoring Transmit Data Rate (tx-datarate) and Egress Buffer Occupancy values at the configured polling interval. Only when both values exceed their respective thresholds is a transmit over-utilization condition met and the configured actions are taken. The parameters with which tx-overutilization counter operates are as follows:

- Therising threshold, falling threshold and warning threshold monitor the transmit buffer occupancy. This is the percentage of the interface egress buffers that are currently occupied with frames awaiting transmission.
- tx-datarate:This is the average transmission rate percentage as calculated over the specified poll interval. This is configured as a percentage of the operational link speed.

Actions

When a tx-overutilization event occurs, there are several configurable actions that can be taken:

- Error Disable - The interface is brought down into error disable status.
- Flap - The interface is flapped (shutdown followed by a no shutdown).
- Dynamic Ingress Rate Limiting (DIRL): DIRL is used to automatically limit the amount of ingress traffic that is flowing through a switch port resulting in a corresponding reduction of egress traffic thus reducing or eliminating congestion. This action is performed by Fabric Performance Monitor (FPM).
- Fabric Performance Impact Notification (FPIN): FPIN frames are used to notify end devices of performance impacting conditions and behaviors that affect the normal flow of I/O such as link integrity degradation and congestion. The information provided may be used by the end devices to modify their behavior to address the reported conditions. FPINs are Extended Link Services (ELS) frames. In the case of the tx-overutilization counter, when the egress buffer occupancy percentage reaches or exceeds the rising-threshold and the tx-datarate threshold is breached, an FPIN with a descriptor type *Congestion* is sent with a event type of *Oversubscription* and a severity value of *warning*. When the tx-datarate falls below the configured value an FPIN with descriptor type *Congestion* and an event type of *Clear/None* is sent.



Note For warning, rising, congestion-signal-warning, and congestion-signal-alarm events, the transmission data rate and the egress buffer occupancy are monitored. However, for falling events, only the egress buffer occupancy is monitored.

- Congestion Signals: Congestion signals are link level Fibre Channel (FC) primitives sent to an end device by an attached an interface indicating transmission resources on that FC switch port are consumed above a warning or alarm threshold. This can be configured for both *TxWait* and *tx-overutilization* counters. End devices register with switches for which Congestion Signals are supported on transmit and receive and the specific intervals they support via the Exchange Diagnostic Capabilities (EDC) ELS. These

intervals are negotiated by the end device with the switch and can be configured (See the Configuring EDC Congestion Signal section). Depending on the threshold exceeded, port monitor sends *warning or alarm* signal primitives at the specified interval.

For more information on FPM, see: Congestion Isolation Recovery, on page 229 and Fabric Notifications—FPIN and Congestion Signal.

Note tx-overutilization counter is supported only for logical-type edge policies.

Supported Hardware and Software

Hardware:

- Cisco MDS 9700 48-Port 64-Gbps Fibre Channel Module
- Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Module
- Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switches
- Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switches
- Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switches
- Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switches
- Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel Switches
- Cisco MDS 9396V 64-Gbps 96-Port Fibre Channel Switches
- Cisco MDS 9220i 32-Gbps 20-Port Multiservice Fabric Switch

This feature is not supported by the following line cards and fabric switches:

- Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module
- Cisco MDS 9700 24/10-Port SAN Externsion Module
- Cisco MDS 9250i Multiservice Fabric Switch
- Cisco MDS 9396S 16-Gbps Multilayer Fabric Switch
- Cisco MDS 9148S 16-Gbps Multilayer Fabric Switch

16 Gbps fabric switches and line cards guidelines and limitations

These guidelines are applicable for 16 Gbps fabric switches and line cards

- No support for N Port Virtualization (NPV) with tx-overutilization counter configuration.

FPIN guidelines and limitations

These guidelines are applicable for FPIN.

- In CiscoMDS NX-OS 9.4(3) or earlier releases, peer FPIN congestion notification and peer FPIN clear congestion notification are not supported with tx-overutilization counter. Due to no support peer congestion

notification are not sent for transmit oversubscription events. Disable monitoring of tx-overutilization counter in the active Port Monitor policy to prevent notification failures.

- From Cisco MDS NX-OS Release 9.4(4) or later releases, peer FPIN congestion notification and peer FPIN clear congestion notification is supported with tx-overutilization counter.
- From Cisco MDS NX-OS Release 9.4(4) or earlier releases, FPIN sends notifications when any breach occurs as alarm severity only.
- From Cisco MDS NX-OS Release 9.4(4) or later releases, the following notifications are sent whenever breach occurs based on the warning threshold for oversubscription and txwait.
  - When only a warning breach occurs then the severity warning is displayed.
  - When warning then an alarm breach occurs then the severity warning is displayed.
  - When an alarm and then a warning breach occurs then the severity warning is displayed.
- For example, configure the warning threshold as 90, alarm threshold as 100 and falling threshold as 80, then
  - Warning condition is when threshold greater than or equal to 90 and lesser than 100 then warning threshold severity syslog is created.
  - Alarm condition is threshold greater than or equal to 100 alarm threshold then severity syslog is created.
  - Falling condition is when threshold lesser than or equal to 80 falling threshold severity syslog is created.

DIRL guidelines and limitations

These guidelines are applicable for DIRL.

- It is not recommend using the DIRL configuration with tx-overutilization. While configuring the DIRL port guard action command alongside tx-overutilization is permitted, it does not operate as intended.

NPV guidelines and limitations

These guidelines are applicable for NPV.

- From Cisco MDS NX-OS Release 9.4(5), the tx-overutilization counter is supported in NPV mode on 32 Gbps and 64 Gbps fabric switches and line cards.
- Congestion signals are not supported when NPV mode is off on the switch.

Configuring PMON Tx-Overutilization Counter

Use the following procedure to enable and configure Tx-overutilization Counter:

Before you begin

Create a port monitor policy and ensure it is not activated before configuring the tx-overutilization counter. For more information, see Enabling a Port Monitor.

Table 29: Table 1. The default settings of the tx-overutilization counters:

Parameter	Description	Default Value	Allowed Range
poll-interval	Frequency of counter checks	10 seconds	10 seconds
rising-threshold	The egress buffer occupancy percentage. This is the percentage of the interface's egress buffers occupied by frames queued for transmission. When the rising-threshold percentage is reached and the tx-datarate percentage is reached, port monitor takes the specified actions for example DIRL and FPIN.	80%	20% to 100%
falling-threshold	The egress buffer occupancy percentage. This is the percentage of egress buffers occupied by frames queued for transmission. When the egress buffer occupancy percentage is less than the falling-threshold Port Monitor stops taking the specified actions for example DIRL and FPIN and begins the recovery interval for DIRL.	30%	0% to 90%
warning-threshold [optional]	The egress buffer occupancy percentage. This is the percentage of egress buffers occupied by frames queued for transmission. When the egress buffer occupancy percentage is greater than or equal to the warning-threshold and less than the rising-threshold and tx-datarate > threshold, Port Monitor issues a syslog message only.	50%	10% to 99%

Parameter	Description	Default Value	Allowed Range
tx-datarate	The average Tx datarate percentage over the polling interval.	80%	70% to 100%
Warning Signal Threshold	The egress buffer occupancy percentage. This is the percentage of egress buffers occupied by frames queued for transmission. When the egress buffer occupancy percentage is greater than or equal to the warning-signal-threshold and less than the alarm-signal-threshold and tx-datarate > threshold, Port Monitor initiates the transmission of Warning level Congestion Signal primitives.	Disabled	10% to 95%
Alarm Signal Threshold	The egress buffer occupancy percentage. This is the percentage of egress buffers occupied by frames queued for transmission. When the egress buffer occupancy percentage is greater than or equal to the alarm-signal-threshold and tx-datarate > threshold, Port Monitor initiates the transmission of Alarm level Congestion Signal primitives.	Disabled	10% to 100%

Table 30: Port Monitor events and triggering conditions

Event	Triggering Conditions(s)
Rising threshold	Egress buffer occupancy >= rising threshold and tx-datrate >= threshold
Warning-threshold	Egress buffer occupancy > =warning-threshold and < rising-threshold and tx-datarate >= threshold
Falling-threshold	Egress buffer occupancy <= falling-threshold

Event	Triggering Conditions(s)
Congestion Signal Warning level	Egress buffer occupancy >= warning-signal-threshold and tx-datarate >= threshold
Congestion Signal Alarm level	Egress buffer occupancy >= alarm-signal-threshold and tx-datarate >= threshold
Stop sending Congestion Signals	Egress buffer occupancy < warning-signal-threshold

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the policy name:

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

Step 3 Enter the logical port type as edge.

switch(config-port-monitor)# **logical-type edge**

Step 4 Enable **TX-Overutilization** counter.

switch(config-port-monitor)# **monitor counter tx-overutilization**

Step 5 Set the counter monitoring thresholds and actions

switch (config-port-monitor)# **counter tx-overutilization poll-interval 10 delta rising-threshold 80 event 4 falling-threshold 70 alerts syslog rmon obfl portguard FPIN tx-datarate 90**

Step 6 Set the congestion signalling threshold

switch (config-port-monitor)# **counter tx-overutilization warning-signal-threshold 40 alarm-signal-threshold 80 portguard congestion-signals**

Verification of PMON Tx-Overutilization Counter Configuration

The following procedure show how to verify the **Tx-Overutilization** counter threshold and actions configuration.:

Procedure

Use the **show port-monitor** command to display counter thresholds and other parameters of port monitor policies:

The following example shows the updated thresholds for tx-overutilization counter after it is configured:

```
Policy Name  : testpolicy
Admin status : Not Active
Oper status  : Not Active
Logical type : All Edge Ports
```

Counter	Threshold	Interval	Warning	Thresholds	Rising/Falling actions	Congestion-signal						
	Type	(Secs)										
			Threshold		Alerts	Rising		Falling		Event		
Alerts	PortGuard		Warning	Alarm								
Link Loss	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
Sync Loss	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
Signal Loss	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
Invalid Words	Delta	60	none		n/a	1		0		4		
syslog,rmon	none		n/a		n/a							
Invalid CRC's	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
State Change	Delta	60	none		n/a	5		0		4		
syslog,rmon	none		n/a		n/a							
TX Discards	Delta	60	none		n/a	200		10		4		
syslog,rmon	none		n/a		n/a							
LR RX	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
LR TX	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							
Timeout Discards	Delta	60	none		n/a	200		10		4		
syslog,rmon	none		n/a		n/a							
Credit Loss Reco	Delta	1	none		n/a	1		0		4		
syslog,rmon	none		n/a		n/a							
TX Credit Not Available	Delta	1	none		n/a	10%		0%		4		
syslog,rmon	none		n/a		n/a							
RX Datarate	Delta	10	none		n/a	80%		70%		4		
syslog,rmon,obfl	none		n/a		n/a							
TX Datarate	Delta	10	none		n/a	80%		70%		4		
syslog,rmon,obfl	none		n/a		n/a							
TX-Slowport-Oper-Delay	Absolute	1	none		n/a	50ms		0ms		4		
syslog,rmon	none		n/a		n/a							
TXWait	Delta	1	none		n/a	30%		10%		4		
syslog,rmon	none		n/a		n/a							
tx-overutilization	Delta	10	none		n/a	80%@90%		70%@90%		4		
syslog,rmon,obfl	FPIN		40%@90%	80%@90%								
RX Datarate Burst	Delta	10	none		n/a	5@90%		1@90%		4		
syslog,rmon,obfl	none		n/a		n/a							
TX Datarate Burst	Delta	10	none		n/a	5@90%		1@90%		4		
syslog,rmon,obfl	none		n/a		n/a							
Input Errors	Delta	60	none		n/a	5		1		4		
syslog,rmon	none		n/a		n/a							

Verification of PMON Tx-Overutilization System Logging

TX-Overutilization related system log messages:

The following example shows Tx-Overutilization related system log messages:

2024 Nov 27 17:57:55 switch %PMON-SLOT1-4-WARNING_THRESHOLD_REACHED_UPWARD: Tx-Overutilization has reached warning threshold in the upward direction (port fc1/1 [0x1000000], value 97% tx-datarate 92%) .

The above system log shows that the tx-overutilization counter breached the configured warning threshold on upward direction.

2024 Nov 27 17:57:55 switch %PMON-SLOT1-4-CS_SENT: Counter Tx-Overutilization with a value 97% tx-datarate 92% : Sent Alarm Congestion Signal on port fc1/1[0x1000000].

The above system log shows that the alarm congestion signal was sent on port fc1/1[0x1000000], as Tx-Overutilization counter breached configured tx-data rate and alarm-signal thresholds.

2024 Nov 27 17:57:55 switch %PMON-SLOT1-4-RISING_THRESHOLD_REACHED_WARNING: Tx-Overutilization has reached the rising threshold (port=fc1/1 [0x1000000], value 97% tx-datarate 92%) .

The above system log shows that the Tx-Overutilization reached the rising threshold (port=fc1/1 [0x1000000], value 97% tx-datarate 92%)

2024 Nov 27 17:57:55 switch %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard action (FPIN) initiated for the Port fc1/1(1000000) Counter Tx-Overutilization .

The above system log shows that the Port guard action FPIN has been triggered by the policy setting for tx-overutilization counter.

2024 Nov 27 17:57:55 switch %FPM-5-OVERSUBSCRIPTION_EVENT: Congestion event with type Oversubscription triggered for device 50:08:01:60:00:8a:d3:10 (0xef0000) vsan 0x1

The above system log indicates that the FPIN has sent the congestion event to the connected device port whose pWWN and FC ID is 50:08:01:60:00:8a:d3:10 (0xef0000).

2024 Nov 27 17:57:55 switch %FPM-5-FPIN_SENT: FPIN 'Oversubscription' Congestion notification sent to device FCID 0xef0000 in VSAN 1.

The above system log indicates that the FPIN has sent the Over-subscription notification from the switch to the device port FCID 0xef0000

2024 Nov 27 17:58:00 switch %PMON-SLOT1-4-CS_STOP: Counter Tx-Overutilization with a value 0% tx-datarate 79% : Stop Alarm Congestion Signal on port fc1/1[0x1000000].

The above system log indicates that the system has stopped sending the alarm congestion signal.

2024 Nov 27 17:58:00 switch %PMON-SLOT1-4-WARNING_THRESHOLD_REACHED_DOWNWARD: Tx-Overutilization has reached warning threshold in the downward direction (port fc1/1 [0x1000000], value 0% tx-datarate 79%) .

The above system log shows that the egress buffer occupancy has fallen below the warning threshold (recovery and safe levels) hence the syslog was generated for the port with tx-datarate.

2024 Nov 27 17:58:00 switch %PMON-SLOT1-4-FALLING_THRESHOLD_REACHED_WARNING: Tx-Overutilization has reached the falling threshold (port=fc1/1 [0x1000000], value 0% tx-datarate 79%) .

The above system log shows that the egress buffer occupancy has fallen below falling threshold (recovery and safe levels).

2024 Nov 27 17:58:00 switch %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard action (FPIN) recovered for the Port fc1/1(1000000) Counter Tx-Overutilization .

The above system log shows that the Tx-Overutilizaton condition was recovered.

2024 Nov 27 17:58:00 switch %FPM-5-FPIN_SENT: FPIN 'Clear' Congestion notification sent to device FCID 0xef0000 in VSAN 1.

The above system log shows that the action was triggered as a part of recovery, **FPIN Clear/None** was sent to the device whose FCID is 0xef0000.

Congestion Isolation Recovery

Prior to Cisco MDS NX-OS Release 8.5(1) when a slow device was detected, the flows to the congested device were automatically moved to the low-priority VL using the Congestion Isolation feature. After the congested device recovered from congestion, the flows had to be manually moved flows from the low-priority VL to normal VL.

From Cisco MDS NX-OS Release 8.5(1), the Congestion Isolation Recovery feature automatically recovers the traffic to congested device from a low-priority VL to the normal VL. This recovery is done without any user intervention unlike the Congestion Isolation feature where user had to manually recover the traffic flowing to the congested device from low-priority VL to normal VL after the device had recovered from congestion.

The *cong-isolate-recover* portguard action is available in the port monitor policy for the supported slowdrain counters.

The recovery process uses a **recovery-interval** to check if the traffic going to congested device in the low-priority VL can be moved back to the normal VL. The following is the process that is used for recovery:

1. A device is identified as a congested device when the port monitor counter detects a rising threshold. After the device is identified as a congested device, the traffic destined to the congested device is moved to the low-priority VL.
2. When port monitor detects a falling threshold for the congested device, a recovery interval (15 minutes by default) is started. During this interval, if the port monitor counter stays at or below the falling threshold continuously, then the device is no longer marked as congested device and the traffic destined to the device is moved from the low-priority VL to normal VL.

   However, if port monitor detects an event threshold that is more than the falling threshold before the expiry of the recovery interval, the interval is discarded and the device remains classified as a congested device. The recovery timer is again started when the next falling threshold is detected by the port monitor. The recovery interval can be configured. For more information, see Configuring Congestion Isolation Recovery, on page 274.
3. Also, the Congestion Isolation Recovery feature allows you to determine the number of times the traffic destined to a congested device can toggle between congestion isolated and recovered, which is known as the *number of occurrences*. If the traffic destined to a congested device toggles between congestion isolated and recovered for the specified number of occurrences within a specified duration called **isolate-duration**, then on the last occurrence the device would be marked as a congestion isolated device and would not be allowed to recover until the isolate-duration expires. Isolate duration is a recurring interval and starts when a port monitor policy is activated.

For example, let us consider a device P1 that is detected as a congested device. The traffic destined to the device is moved to low-priority VL and has recovered back after sometime. The traffic destined to the device P1 keeps being detected as slow and then recovering. In such cases, you can configure the number of such transitions or occurrences known as *number of occurrences* for a specified **isolate-duration**. Suppose you have chosen this value to be 3 and the isolate duration to be 24 hours. When an event threshold that is more the falling threshold is detected for P1 for the third time in the first 2 hours of activating the isolate duration,

P1 is marked as congested device. The flows will be moved to low-priority VL for the remaining 22 hours and any subsequent falling threshold detected is ignored. The device P1 would remain as congested device until the end of 22 hours after which the device is recovered and monitored for an event threshold that is more than the falling threshold again. However, you can manually recover flows from low-priority VL to normal VL. For more information, see Configuring Excluded List of Congested Devices, on page 272.

Note The **isolate-duration** starts only after the corresponding port monitor policy is activated.

The following is the list of counters that you can use to trigger the Congestion Isolation Recovery port-monitor portguard action (cong-isolate-recover):

- credit-loss-reco
- tx-credit-not-available
- tx-slowport-oper-delay
- txwait

Fabric Notifications—FPIN and Congestion Signals

Fabric Notifications are used to notify end devices of performance impacting conditions and behaviors that affect the normal flow of IO such as link integrity degradation and congestion. The information provided by the notifications may be used by the end devices to modify their behavior to address the reported conditions. The functions include notifications in the form of Fabric Performance Impact Notifications (FPINs) ELS (Extended Link Service) requests and Congestion Signal primitives. End devices must register their FPIN and Congestion Signal capabilities using Register Diagnostic Functions (RDF) and Exchange Diagnostic Capabilities (EDC) ELS requests. RDF, EDC, FPIN and Congestion Signals are supported starting with Cisco MDS NX-OS release 8.5(1).

Registration: Register Diagnostic Functions (RDF) and Exchange Diagnostic Capabilities (EDC) ELS requests are used to exchange capabilities between an end device and a peer switch. Both RDFs and EDCs are normally initiated by end devices after they have successfully logged into the fabric. RDF requests the switch to register the diagnostic capabilities of the end device for FPINs. The switch responds with the FPIN types it supports. The MDS currently supports three types of FPINS:

- Congestion
- Peer-Congestion
- Link Integrity

EDC requests the switch to register the end device for link level congestion signal primitives when congestion events are detected by the switch on the attached port. MDS supports both Alarm and Warning level Congestion Signals.

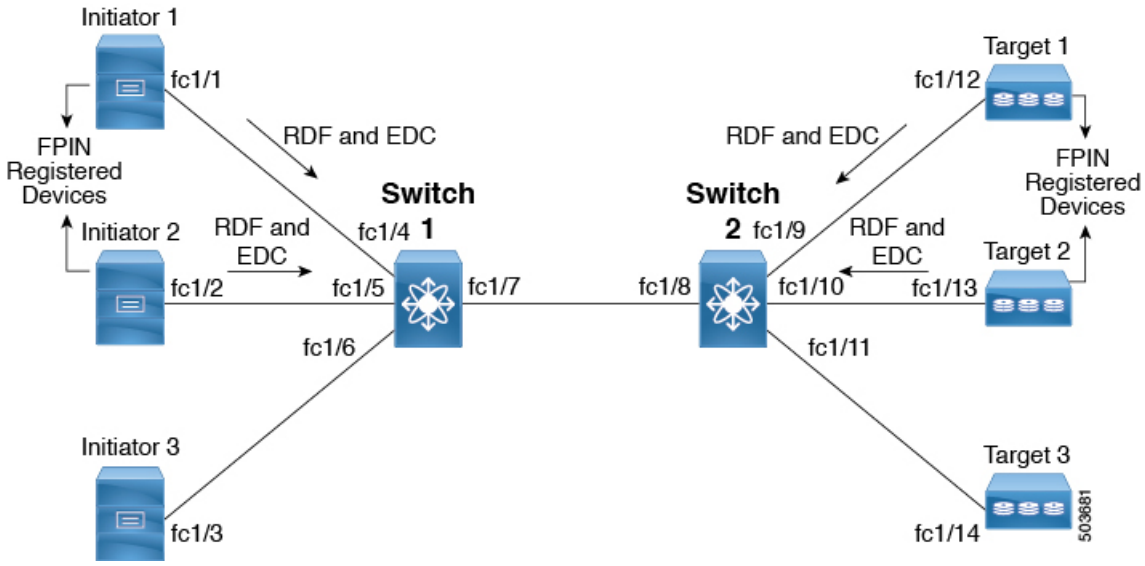
When a link integrity degradation or congestion event is detected in the fabric, the switch generates an FPIN to the device on the link affected by the event if it has registered. Then the switch generates FPIN events to all devices that are zoned with the device on the link with the event if they also have registered for FPINs.

Feature FPM must be enabled on the switch to support FPINs and Congestion Signals. If feature FPM is not enabled then the switch responds with an *LS Reject* to the RDF and EDC requests.

The following operations for fabric notifications are supported:

- Figure 3: RDF and EDC ELS Exchange, on page 231 displays a sample topology where Initiator 1, Initiator 2, Target 1, and Target 2 are registered for FPIN via RDF and EDC. Initiator 3 and Target 3 are not registered for FPIN.

Figure 3: RDF and EDC ELS Exchange



- Notifications: FPIN ELS alerts registered end devices about events that impact performance and contains the descriptions of the event occurrences.

  The following are the types of events for which FPIN is generated:

  - Congestion: A congestion condition that is detected at an F port will be notified to the connected end device.

    You can use the following list of counters to trigger FPIN congestion events:

    - TxWait
    - Tx-Overutilization

    Prior to Cisco MDS NX-OS 9.4(3), FPIN congestion events were triggered when the counter reached the rising-threshold, and the event severity was always set to *Warning*.

    In Cisco MDS NX-OS 9.4(3), 9.4(3a), and 9.4(3b), FPIN congestion events continue to be triggered upon the counter reaching the rising threshold; however, the event severity was changed to *Alarm*.

    Starting from Cisco MDS NX-OS 9.4(4), Congestion FPINs can be sent with two severities: Warning and Alarm. A Congestion FPIN with a severity of *Warning* is sent when the warning threshold is crossed in SMA or port-monitor. A Congestion FPIN with a severity of *Alarm* is sent when the alarm threshold is crossed in SMA, or the rising threshold is crossed in port-monitor.

    A Congestion and Peer Congestion FPIN with an event type of *Clear* are sent when the falling threshold is crossed.

  - Peer congestion: A congestion condition that is detected at an F port will be sent to all the devices zoned with the congested port. The Peer Congestion FPIN provides details such as the type of

congestion condition, the PWWN of the port where the congestion was detected, the PWWN of the device connected to the detecting port, and a list of affected device PWWNs. If the list of affected PWWNs is empty, it indicates that only the device connected to the detecting port is impacted by the event.

- FPIN Peer Congestion events are triggered for the same counters and thresholds that trigger FPIN congestion events. FPIN Peer Congestion events do not contain a severity.
- Link integrity: A condition that checks for port integrity. The information included in the Link Integrity FPIN includes the reason, such as, link failure, loss of signal, and so on, and a threshold value that was exceeded.

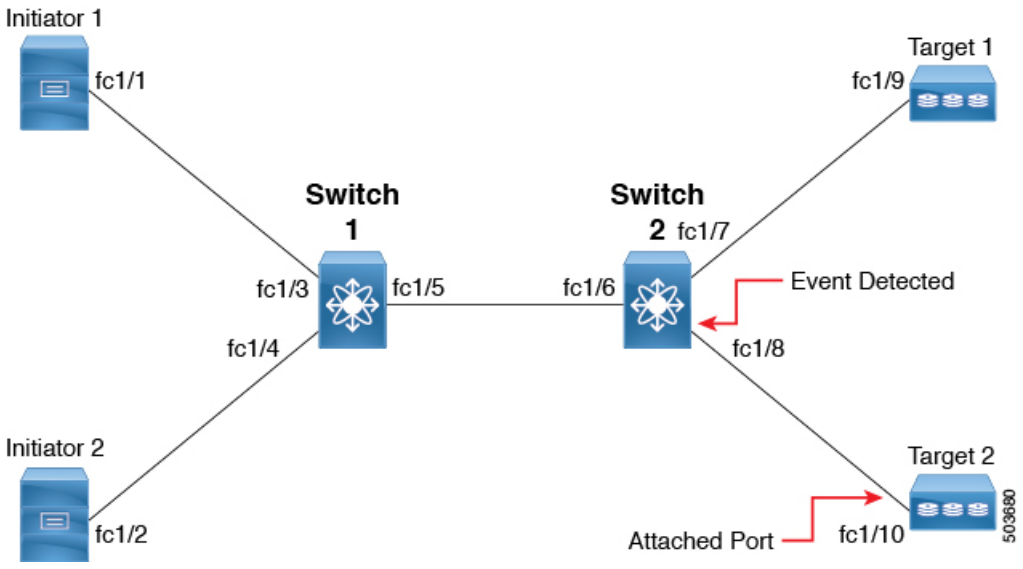
  The following is the list of counters that you can use to trigger the link integrity events:

  - link-loss
  - sync-loss
  - signal-loss
  - invalid-words
  - invalid-crc

Note The Congestion Isolation Recovery feature is not supported on these counters. For more information, see Congestion Isolation Recovery, on page 229.

Figure 4: FPIN Events, on page 232 displays a sample topology where all the devices are configured in a single zone. An event is detected at port fc1/8 and Target 2 is the attached port or peer port.

Figure 4: FPIN Events



The following provides how the information is shared between the devices when events are detected:

- Congestion: When a congestion event is detected at port fc1/8, an FPIN congestion descriptor is sent to Target 2.
- Peer congestion: When a congestion event is detected at port fc1/8 and FPIN peer congestion event is sent to Initiator 1, Initiator 2, and Target 1 containing the pWWN list of Target 2.
- Link integrity: When a link integrity event is detected at port fc1/8, FPIN link integrity is sent to Initiator 1, Initiator 2, and Target 1 with the pWWN list of Target 2 and FPIN link integrity is also sent to Target 2 with the pWWN list of Initiator 1, Initiator 2, and Target 1.



Note Cisco MDS port does not handle FPINs received from end devices. Instead, they are discarded.

- Signals: Signals: Congestion Signal primitives may be sent to and end device by an attached switch port indicating the transmission resources on the switch port are consumed above a threshold. This can occur when TxWait occurs or when Tx-Overutilization occurs.

  The following types of congestion signal primitives are supported and are configurable in the port monitor policy for the TxWait and Tx-Overutilization counters:

  - Warning congestion signal: This signal is sent when the TxWait or Tx-Overutilization condition on a port exceeds the configured warning threshold.
  - Alarm congestion signal: This signal is sent when the TxWait or Tx-Overutilization conditions on a port exceeds the configured alarm threshold.

FPM receives notifications about link integrity degradation and congestion from port monitor when its counters detect a configured rising threshold.

The following port monitor counters support FPIN as portguard actions to check the link degradation:

- LinkFailures
- SyncLoss
- SigLoss
- Invalid TxWords
- InvalidCRCs

The following port monitor counters support FPIN portguard actions for congestion and peer congestion:

- TxWait
- Tx-Overutilization

TxWait and Tx-Overutilization also support configuring Congestion Signals.

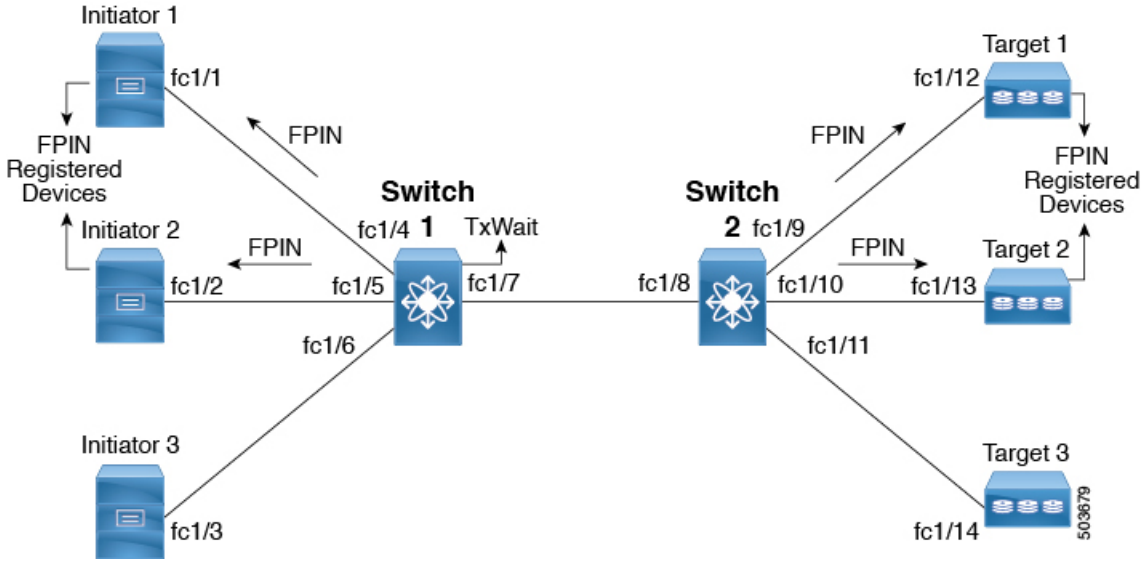
Recovery of congestion events is also notified through FPIN to the end device. Recovery of congestion events is notified from port monitor when a counter value remains below the falling threshold for the **recovery-interval**. For information about configuring recovery interval for FPIN, see Configuring the Port-Monitor Portguard Action for FPIN, on page 277.

For configuring FPIN and congestion signal fabric notifications, see Configuring EDC Congestion Signal, on page 279.

FPM can manually classify a device as congested and also exclude a device from detection of link integrity degradation and congestion. For more information, see Configuring Fabric Notifications, on page 276.

Figure 5: Fabric Notifications, on page 234 displays a sample topology where end devices Initiator 1, Initiator 2, Target 1, and Target 2 are registered for FPIN via RDF and EDC. Initiator 3 and Target 3 are not registered for FPIN. When Initiator 1 becomes slow and TxWait is seen on fc1/4, FPIN is sent to all zoned end devices of Initiator 1 that are registered for FPIN and not to devices that are not registered for FPIN.

Figure 5: Fabric Notifications



Changed Switch RDF Capabilities

End devices might need to update their RDF registration details on their connected peer switch if the switch RDF capabilities change. Switch RDF capabilities may change after a nondisruptive switch upgrade or when the FPM feature is enabled. Forcing end device to update the capabilities they register for can be done by either:

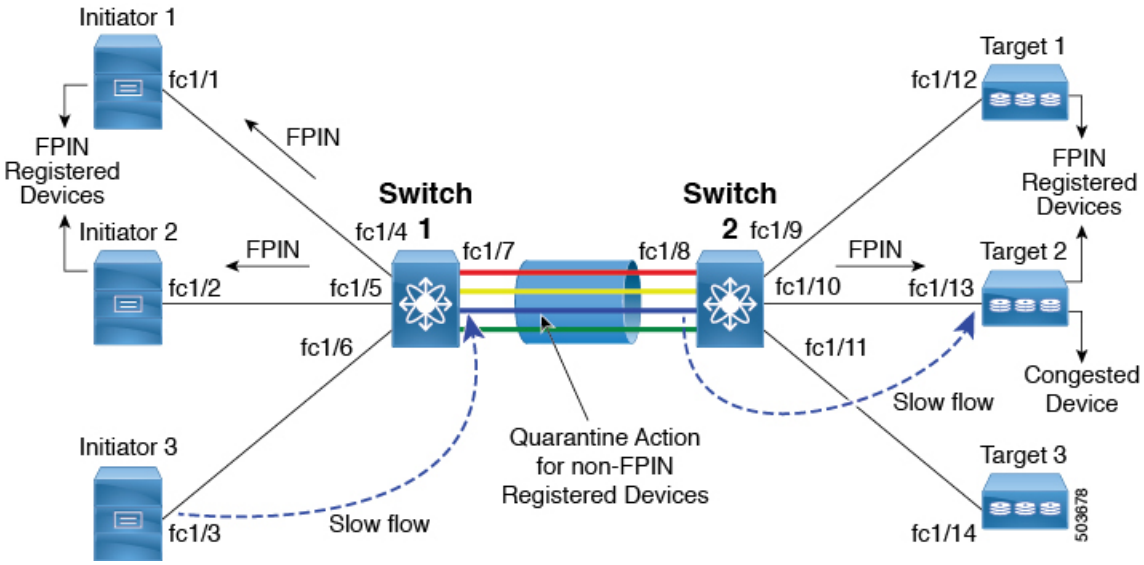
- a disruptive link flap of the link between the end device and the switch, or
- nondisruptively, using the **send fc rdf** command. This method causes the switch to send an RDF request to the attached end devices indicating to the end device that the FPIN capabilities of the MDS have changed. The end devices can then register or reregister for the set of RDF capabilities they currently support.

FPIN and ER_RDY

FPIN can also work in conjunction with the ER_RDY feature to isolate flows to the low-priority VL if end devices have not registered themselves with RDF for fabric notifications. The recovery of flows from low-priority VL to normal VL happens when port monitor notifies FPM about the recovery. For FPIN to work with the ER_RDY feature, you need to enable the ER_RDY feature. For more information, see Enabling Extended Receiver Ready, on page 269.

Figure 6: FPIN and ER_RDY, on page 235 displays a sample topology where Initiator 1, Initiator 2, Target 1, and Target 2 are registered for FPIN through RDF. Also, Initiator 1 is zoned with Target 1 and Target 2, Initiator 2 is zoned with Target 2 and Target 3, and Initiator 3 is zoned to Target 2 and Target 3. Initiator 3 and Target 3 are not registered for FPIN. Congestion is detected at Target 2 and all the zoned devices of Target 2 that are registered for FPIN are notified about the congested device. Initiator 3 is not registered for FPIN and as we have ER_RDY enabled, the flow from Initiator 3 to Target 2 uses the low-priority VL.

Figure 6: FPIN and ER_RDY



Dynamic Ingress Rate Limiting

Dynamic Ingress Rate Limiting (DIRL) is used to automatically limit the rate of ingress commands and other traffic to reduce or eliminate the congestion that is occurring in the egress direction. DIRL does this by reducing the rate of IO solicitations such that the data generated by these IO solicitations matches the ability of the end device to actually process the data without causing any congestion. As the device's ability to handle the amount of solicited data changes, DIRL, will dynamically adjust seeking to supply it the maximum amount of data possible without the end device causing congestion. After the end device recovers from congestion, DIRL will automatically stop limiting the traffic that is sent to the switch port.

In case of slow drain and over utilization, the assumption is that if the rate of IO solicitation requests is reduced then this will make a corresponding reduction in the amount of data solicited and being sent to the end device. By reducing the amount of data this will resolve both the slow drain and over utilization cases.

DIRL is comprised of two functions and can perform equally well on congestion caused both slow drain and over utilization:

- Port monitor: Detects slow drain and over utilization conditions and if the portguard action is set as **DIRL**, it notifies FPM. Port monitor portguard action **DIRL** can be configured on the following counters:
  - txwait: Use for detection of slow drain.
  - tx-datarate: Used for detection of over utilization.
  - tx-datarate-burst: Use for detection of over utilization.

- FPM: DIRL actions are taken by FPM as notified by port monitor. On detecting a rising threshold from port monitor, FPM does rate reduction causing the rate of ingress traffic to be reduced. On detecting the value of a counter being below the falling threshold continuously for the DIRL recovery interval, FPM does rate recovery.

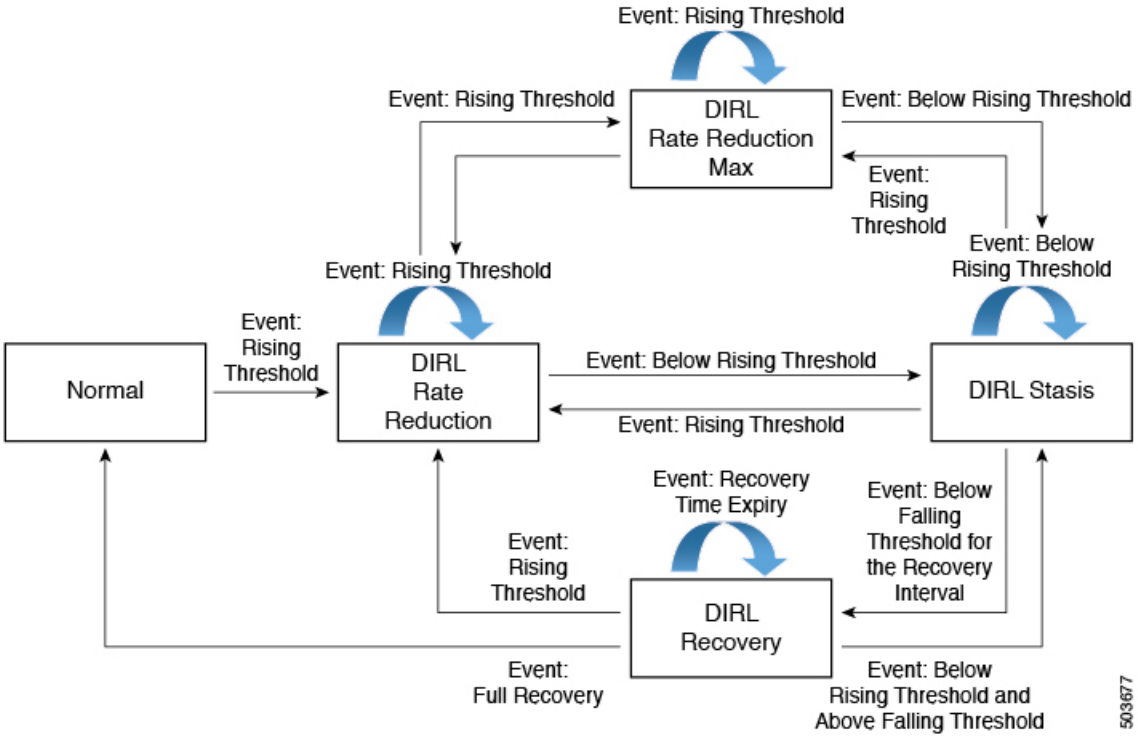
After the port monitor policy is configured with the DIRL portguard action and activated, all non-default F ports are monitored by default and FPM is notified if congestion is detected on any of these ports. However, you can manually exclude certain interface from being monitored. For more information, see Configuring Excluded List of Congested Devices, on page 272.

Note If an interface is configured using static ingress rate limit by using the **switchport ingress-rate** *limit* command, then DIRL will not function for that port. However, a port that is subject to DIRL can be overridden by static ingress rate limit.

The following are the different transition states of DIRL:

- Normal: The state in which a port is functioning normally and state before it enters DIRL Rate Reduction. After full recovery, port returns to the Normal state.
- DIRL Rate Reduction: The state in which an event rising threshold triggers the DIRL rate reduction process.
- DIRL Rate Reduction Maximum: The state in which the DIRL rate reduction has reached its maximum value and more rising thresholds events are detected.
- DIRL Stasis: The state in which an event below rising threshold and above falling threshold is detected. This state will transition to DIRL Recovery state when an event below falling threshold is detected for the configured **recovery-interval**.
- DIRL Rate Recovery: The state in which the DIRL rate recovery happens on detecting an event below falling threshold for the configured **recovery-interval**. This state will transition to the Normal state after the port recovers completely from DIRL. This state is a recurring state and there will be multiple rate recoveries before the ports are completely recovered from DIRL. This state will transition to the DIRL Stasis state when an event below rising threshold and above falling threshold is detected.

Figure 7: Different States of DIRL



Let us consider the following example where the DIRL rate recovery process has started on port fc4/12 after detecting an event rising threshold:

```
switch# show fpm ingress-rate-limit events interface fc4/12
------------------------------------------------------------------------------------------------------------------------
Interface|Counter|  Event |       Action  |Operating | Input  |Output|Current|Applied|
Time
         |       |        |               |port-speed| rate   | rate |  rate |  rate |

                                         |   Mbps   | Mhps   | Mbps |limit %|limit %|
------------------------------------------------------------------------------------------------------------------------
fc4/12    txwait  rising    rate-reduction  16000.00   8853.37 8853.10 77.010  31.563  Mon
 Jan 18 22:34:44 2021
fc4/12    txwait  recovery  rate-recovery   16000.00   8369.35 8369.35 61.608  77.010  Mon
 Jan 18 22:34:37 2021
fc4/12    txwait  recovery  rate-recovery   16000.00   6697.13 6697.16 49.287  61.608  Mon
 Jan 18 22:33:37 2021
fc4/12    txwait  recovery  rate-recovery   16000.00   5359.97 5359.95 39.429  49.287  Mon
 Jan 18 22:32:36 2021
fc4/12    txwait  recovery  rate-recovery   16000.00   4288.87 4288.86 31.543  39.429  Mon
 Jan 18 22:31:36 2021
fc4/12    txwait  rising    rate-reduction  16000.00   8847.91 8848.01 100.000 31.543  Mon
 Jan 18 22:30:24 2021
```

The following are the actions that are initiated by DIRL depending on the type of event detected on the port:

Note The events are listed in reverse chronological order with the most current event at the top.

1. An event rising threshold is detected on the port and DIRL is initiated for the port. The port ingress traffic rate is reduced to 50% of its current rate.
2. In the next polling interval, the **recovery-interval** expires without detecting a rising threshold. The port ingress traffic is increased by 25% of its current capacity.
3. In the next polling interval, the **recovery-interval** expires without detecting a rising threshold. The port ingress traffic is increased by 25% of its current capacity.
4. In the next polling interval, the **recovery-interval** expires without detecting a rising threshold. The port ingress traffic is increased by 25% of its current capacity.
5. In the next polling interval, the **recovery-interval** expires without detecting a rising threshold. The port ingress traffic is increased by 25% of its current capacity.
6. In the next polling interval, an event rising threshold is detected on the port and DIRL is initiated for the port. The port ingress traffic is reduced again to 50% of its current rate.

Static Ingress Port Rate Limiting

A static port rate limiting feature helps control the bandwidth for individual Fibre Channel ports using the **switchport ingress-rate** *limit* command. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by slowing the rate of B2B credits transmitted from the FC port to the adjacent device. Port rate limiting works on all Fibre Channel ports. Prior to Cisco MDS NX-OS Release 8.5(1), the rate limit ranges from 1 to 100%. From Cisco MDS NX-OS Release 8.5(1), the limit ranges from 0.0126 to 100%. The default rate limit is 100%.

Starting from Cisco MDS NX-OS Release 8.5(1), the FPM feature needs to be configured before configuring the dynamic or static ingress port rate limiting feature on all Cisco MDS switches except Cisco MDS 9250i and MDS 9148S switches. Prior to Cisco MDS NX-OS Release 8.5(1) or on Cisco MDS 9250i and MDS 9148S switches, static ingress port rate limiting can be configured on all Cisco MDS switches and modules only if the QoS feature is enabled.

What is Smart Monitoring and Alerting (SMA)

Smart Monitoring and Alerting (SMA) performs early detection of error or traffic pattern and congestion. SSMA provides notifications through alerts and helps perform the required actions. SMA helps monitor and detect important switch events or conditions in real-time.

SMA feature is a registration-based feature. You can register this feature when you

- receive SMA feature status notification,
- restart the switch, or
- enable or disable the feature.

What are the benefits of SMA

Benefit	Description
Unified Monitoring	Single feature that enables monitoring of all aspects of the switch.

Benefit	Description
Timely Detection	• Real time continuous monitoring of data streams and system attributes for errors and anomalies in counters. • Only port counters and protocol counters are monitored for unexpected patterns.
Proactive Notification	• Notifications or alerts are sent by SMA whenever an event of interest is detected. • Alerts such as snmp-traps, system logs or OBFL logs.
Ease of Configuration	• Easy to configure the SMA feature by using the building blocks such as the default policy or entity groups.
Maintain System Health	• SMA helps in maintainin the system health by constantly monitoring the switch that enable in timely actions
History of Events	• SMA maintains a history of events and aims to provide an overall system health report.
Customization of priorities and thresholds	• SMA provides flexibility to set custom thresholds, priorities or conditions based on specific requirements. • Define thresholds based on understanding of the system, • Enables tailored monitoring and alerting such as configuring the monitoring for per port.

How does the SMA component architecture work

The SMA architecture provides a centralized system for monitoring, managing, and alerting across different operational layers.

Summary

The key components involved in the process are:

- **SMA SERVER**: The central component responsible for processing data, managing actions, and generating notifications.
- **Monitor Client(s)**: Components that collect and send monitoring data to the SMA SERVER.
- **Action Client(s)**: Components that initiate or receive commands for specific actions from the SMA SERVER.
- **libsma**: A library that facilitates communication between clients and the SMA SERVER.
- **SUP (Supervisory Layer)**: The higher-level operational layer where the SMA SERVER and its direct clients reside.

- **LC (Local Component Layer)**: The lower-level operational layer where local clients and SMA agents communicate with the central SMA SERVER.
- **LC SMA**: A local SMA agent or component within the LC layer that communicates with the central SMA SERVER.
- **Alert(s), Syslog, Trap**: Various mechanisms used by the SMA SERVER to output notifications and logs.

Result

The SMA architecture enables comprehensive oversight and control of distributed system components, ensuring efficient operation and timely notification of events.

How the SMA workflow operates

The SMA workflow describes a centralized system for managing, monitoring, and acting upon events within an infrastructure, while also providing robust alerting capabilities. The key components involved in the process are:

Summary

- **SMA (System Management Agent):** A central management entity that provides a unified monitoring infrastructure, easy configuration options, and maintains a history of events.
- **Config Input:** The source for providing policy configurations to the SMA, typically through a Command Line Interface (CLI).
- **Monitor Clients (LC-SMA, Zone Server, FCNS, F-port Server):** Various systems or services that register with the SMA to receive configuration updates and send event triggers.
- **Action Clients (FPM, Port Manager):** Systems or modules that register with the SMA to receive and execute specific actions.
- **Alerting (Syslog, SNMP TRAP, OBFL):** External systems or protocols used by the SMA to send notifications for logging and alert generation.

Result

The SMA workflow provides a comprehensive and centralized approach to infrastructure management, enabling automated monitoring, dynamic configuration, proactive event handling, and efficient notification through various alerting mechanisms.

How SMA policy architecture works

Policy architecture provides a structured framework for defining and applying rules to manage system and network behavior. It organizes elements into logical groups and associates specific monitoring conditions with automated actions.

Summary

The key components involved in the process are:

- **Policy**: The overarching framework that defines rules and responses for managing system or network behavior.
- **Entity Group**: A logical collection of network elements or systems (for example: edge ports, core ports, system) to which specific monitoring and action rules are applied.
- **Monitor Group**: A set of predefined conditions or metrics that are continuously observed within an Entity Group to detect specific states or events (for example: slow-drain, link integrity, configuration scale, environment and so on).
- **Action Group**: A collection of automated responses or operations triggered when conditions defined by a Monitor Group are met (e.g., FPIN, Syslog, Port-Guard, Trap).

Result

The policy architecture enables a modular, scalable, and automated approach to managing and responding to events across various network and system entities, ensuring consistent application of rules and efficient operational control.

Supported hardware

SMA is supported on these MDS switches:

- Cisco MDS 9700 48-Port 64-Gbps Fibre Channel Switches
- Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switches
- Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switches
- Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switches
- Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switches
- Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switches
- Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel Switches
- Cisco MDS 9396V 64-Gbps 96-Port Fibre Channel Switches
- Cisco MDS 9220i
- Cisco MDS 9000 24/10-Port SAN Extension Module for Cisco MDS 9700 Series Multilayer Directors

Supported software

SMA is supported on:

- Cisco MDS 9000 NX-OS Release 9.4(4)
- Cisco MDS 9000 NX-OS Release 9.4(3)
- Cisco MDS 9000 NX-OS Release 9.4(2a)
- Cisco MDS 9000 NX-OS Release 9.4(1)

Migration of PMON feature to SMA feature

You can now migrate the PMON to SMA.

Migration command

The following command is used to migrate PMON to SMA.

```
PMONmigratetoSMA [--src PMON-policy-name] [--dst SMA-policy-name] [--protocol-counters]
```

Guidelines for migration

The guidelines for migration are as described.

- A new user-defined SMA policy is created when migrating PMON policy to SMA policy.
- If there is no PMON policy available on the switch, then SMA policy cannot be created.
- The migrated SMA policy is not activated by default, even though the PMON policy is active while migration.
- Ensure that PMON policy name is with in 64 characters as SMA policy name cannot be more than 64 characters.

PMON to SMA migration

When you migrate PMON policy to SMA policy the counters from PMON to SMA are migrated. The following table provide the mapping for PMON counter to SMA counter.

PMON counters mapped to SMA counters

Table 31: PMON to SMA counters mapping

PMON counter name	SMA counter name
Link Loss	link-loss
Sync Loss	sync-loss
Signal Loss	signal-loss
Invalid Words	rx-invalid-words
Invalid CRC's	rx-invalid-crc
State Change	state-change
TX Discards	tx-discards
LR RX	rx-lr
LR TX	tx-lr
Timeout Discards	tx-timeout-discards
Credit Loss Reco	credit-loss-recovery

PMON counter name	SMA counter name
TX Credit Not Available	tx-credit-not-available
RX Datarate	rx-datarate
TX Datarate	tx-datarate
TX-Slowport-Oper-Delay	tx-slowport-oper-delay
TXWait	tx-wait
RX Datarate Burst	rx-datarate-burst
TX Datarate Burst	tx-datarate-burst
Input Errors	rx-input-errors
SFP TX Power Low Warning	tx-xcvr-power-low-warning
SFP RX Power Low Warning	rx-xcvr-power-low-warning
ASIC Error Pkt from Port	No equivant counter available.
ASIC Error Pkt to xbar	No equivant counter available.
ASIC Error Pkt from xbar	No equivant counter available.
tx-slowport-count	No equivant counter available.
PMON- Alerts	SMA - Alerts
syslog	syslog
rmon	snmp-trap
Event	severity-level
PMON- PortGuard	SMA - Actions
FPIN	fpin-congestion/fpin-link-integrity
DIRL	dirl
cong-isolate	congestion-isolate
cong-isolate-recover	congestion-isolate
errordisable	error-disable
flap	flap

Guidelines and limitations for counters mapping

- XBAR counters are not migrated from PMON policy to SMA policy
- tx-slowport-count counter is not migrated as it is not supported in SMA for any hardware.

- No period is available for tx-slowport-oper-delay counter in SMA.
- If any counter has Absolute as threshold type in PMON policy, except tx-slowport-oper-delay counter, In the migrated SMA policy, same counter will be configured as periodic counter with the same config
- If any default action group with same actions/alerts exists in SMA, same will be used without creating a user defined action group.
- Default entity groups will be used while migrating PMON policy to SMA policy.
- In PMON, if polling interval is more than 600 then in the migrated SMA policy it is mapped to 600.
- In the migrated SMA policy, alarm percentage is mapped to PMON optional parameter **percentage** of SFP counters and falling percentage will be 50%. Change range in SMA for falling threshold is 49 -149.
- Rising and falling threshold count of PMON configuration is ignored.
- If warning threshold is not supported in PMON then in the migrated SMA policy, warning threshold is not configured.

SMA test CLI

The SMA test CLI is a command-line interface tool that:

- generates dummy syslog and SNMP trap alerts,
- allows for the validation of monitoring automation without an active SMA policy, and
- simulates threshold breaches for specific counters and entities.

The tool helps verify that external monitoring software correctly receives and processes alerts before a user deploys a live monitoring policy.

Threshold value comparisons for test alerts

Threshold value comparisons

When using the SMA test CLI, you must ensure that the reported value correctly corresponds to the threshold type to successfully generate an alert:

- For alarm or warning alerts, ensure the reported value is greater than or equal to the threshold value.
- For falling alerts, ensure the reported value is less than or equal to the threshold value.

Generate dummy alerts for monitoring validation

Generate dummy syslog or SNMP trap alerts to test monitoring automation.

Perform this task to ensure that external monitoring systems are correctly configured to receive SMA notifications. You do not need an active SMA policy to use this command.

Follow these steps to generate dummy alerts:

Procedure

Step 1 Enter the SMA test command: **sma test counter**

Example:

```
sma test counter test-sma threshold { alarm } 80 value 60 interface fc/05 alert syslog snmp-trap
```

Step 2 Specify the counter for which the alert is needed.

Step 3 Choose a threshold level from **alarm**, **warning**, or **falling**.

Step 4 Provide a threshold value and the value that represents the breach.

Step 5 Specify the interface or entity for which the counter has breached the threshold.

Step 6 Select the alert type as **syslog**, **snmp-trap**, or both.

The system generates a dummy alert based on the provided parameters. The entity-group is hardcoded as "eg-test-cli" for these simulated alerts.

SMA test command parameters

The following parameters are used with the SMA test CLI:

Table 32: SMA test command parameters

Parameter	Description
counter	The specific counter name (for example, `credit-loss-recovery`, `sync-loss`, `tx-wait`).
threshold	The threshold type to simulate: `alarm`, `warning`, or `falling`.
threshold-value	The threshold for the threshold type.
value	The actual value to report in the alert, which must represent a breach of the threshold.
interface	The entity associated with the alert (for example, `fc1/3`, `port-channel212`, or `switch`).
alert	The notification method: `syslog`, `snmp-trap`, or both.

Congestion isolation actions

A congestion isolation action is a mitigation response that

- restricts the impact of slow-drain or oversubscribed devices on the fabric,
- re-prioritizes traffic destined for affected devices using Virtual Lanes, and
- notifies end-devices to take corrective measures.

Isolation action methods

FPM supports the following methods to isolate and mitigate fabric congestion:

Table 33: Action methods

Action method	Description
FPIN	Sends Fabric Performance Impact Notifications to registered devices. It includes descriptors for congestion, peer congestion, and link integrity. FPIN works only with Port-Monitor Portguard Action FPIN as actions can be configured in both PMON and SMA.
FIB UP Programming	Programs the Forwarding Information Base to overwrite the User Priority (UP) of frames destined for a slow device. • This moves the traffic to a lower-priority Virtual Lane (VL 2) • This moves the traffic from a normal lane (VL 3) to a lower-priority Virtual Lane (VL 2)
Dynamic Ingress Rate Limiting (DIRL)	Reduces the ingress rate of a port by a specific percentage to match the processing speed of the connected device, preventing buffer buildup.
Congestion Signaling	Sends primitive signals (ARB-F1 for warnings or ARB-F7 for alarms) directly to the attached port to indicate congestion severity.

TxWait and Tx-Overutilization counters also support congestion isolation actions to mitigate fabric congestion.

Guidelines and Limitations for Congestion Management

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 output files 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 100 ms 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 100 ms there is no entry.
- Rxwait—Includes interfaces that record 100 ms 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 100 ms 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 fc1/8 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 20-second interval ending at 01/12/18 11:37:55.

```
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. This value can be safely reduced to 200 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 no-credit-drop value:

- 200 ms—Safe value for most fabrics
- 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** commands 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 for 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

Host Bus Adapter Extended Receiver Ready

Beginning with Cisco MDS NX-OS Release 9.3(1):

- Host Bus Adapter Extended Receiver Ready (HBA ER_RDY) is supported on F and NP ports.
- ER_RDY is currently effective between E-ports for isolating traffic unique to the slow devices to a separate virtual link (VL2). In Cisco MDS NX-OS Release 9.3(1), VLs are extended to F and NP ports.
- In HBA ER_RDY mode, initiators use the priority field in FC header to map traffic to a particular VL.
- HBA ER_RDY mode for F and NP ports are negotiated using the Fabric Login (FLOGI) ELS.
- Presently the switch supports four VLs.
    - Three VLs are supported on F port in ER_RDY mode.
    - Host Bus Adaptor (HBA) supports three VLs (VL1, VL2 and VL3).
    - VL0 is not exposed to host as it used only for inter-switch control traffic.
    - VL1 is not used for any traffic profile.
    - VL2 is used for traffic destined to a slow device.
    - VL3 is used for normal traffic.
- HBA maps negotiated priority levels to a VL and the corresponding priority range for each VL as specified in FLOGI ACC.
- HBA Rx credits per VL is programmed as negotiated in FLOGI ACC.
- HBA adds priority value in the Priority field in FC2 header when traffic is originated in ER_RDY mode. For normal traffic, HBA uses priority 0.
- NP-ports and F-ports (server interfaces) can come up in ER_RDY mode. However currently FPIN and Priority Update Notification (PUN) are not supported in NPV mode.
- When switch detects a slow device in the fabric, FPIN is sent to the devices zoned to the slow devices with Priority Update Notification (PUN) descriptor and other supported descriptors. Host use the priority value mentioned in the PUN to send the traffic to slow device. This scenario is applicable only to HBAs.
- The switch does priority to VL mapping at the F-port ingress and select a VL for traffic. Priority 0 is mapped to normal VL (VL3) and priority 2 is mapped to slow VL (VL2) in the switch.
- HBA ER_RDY feature is disabled by default. R_RDY is the default flow control mode for all ports.
- HBA ER_RDY flow control mode should be enabled on all the switches in the fabric. The E, F, and NP ports should be operationally up in ER_RDY mode end-to-end to get the complete benefit of this feature.
- ER_RDY and VMID do not work together.
- ER_RDY and zone QoS are mutually exclusive.
- A flap is needed for port to come up in ER_RDY after enabling the feature.
- **switchpo rt vl-credit** command is not supported for F/NP ports.

- ER_RDY is only supported on specific HBAs. Targets will always come up in R_RDY.
- HBA ER_RDY is supported only on Fibre Channel ports for the following
  - Cisco MDS 9000 Series 24/10 SAN Extension Module (DS-X9334-K9) (Fibre Channel ports only)
  - Cisco MDS 9700 Series with Cisco MDS 9700 64-Gbps Fibre Channel Switching Module (DS-X9748-3072K9)
  - Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9)
  - Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switch
  - Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switch
  - Cisco MDS 9220i Multiservice Fabric Switch
  - Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switch
  - Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch
  - Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel Switch
  - Cisco MDS 9396V 64-Gbps 96-Port Fibre Channel Switch
- 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 effective only between supported switches and R_RDY flow-control mode is used between unsupported switches.
- The **system fc flow-control er_rdy logical- type{core| edge| all}** command should be used to enable ER_RDY for E/F and NP/All ports.
- If you have enabled ER_RDY in releases prior to Cisco MDS NX-OS Release 9.3(1) using **system fc flow-control er_rdy** command and when you upgrade to Cisco MDS NX-OS Release 9.3(1), the running configuration would display this command as **system fc flow- control er_rdy logical-type core**.
- To enable ER_RDY for the first time for E ports, use **system fc flow-control er-rdy logical-type core**. The **system fc flow-control er-rdy logical-type core** command is applicable only on E ports for releases prior to Cisco MDS NX-OS Release 9.3(1).
- If you need to configure ER-RDY for F ports, use the **system fc flow-control er-rdy logical-type edge** command. For the F ports to come up in ER-RDY, the links must be flapped.
- ISSD fails when:
  - ER-RDY is configured on F or NP ports of a switch.
  - ER-RDY is operational in any of the F or N ports.
- For ISSD, the modified command **system fc flow-control er-rdy logical-type core** will revert to its initial form **systemfc flow-control er-rdy** after ISSD without user intervention. Follow the steps for a successful ISSD:
  - Disable F/NP port ER-RDY using **system fc flow-control r_rdy** command.
  - Flap all the F/NP ports which have come up in ER-RDY mode. To find the ports that are in ER-RDY mode, use the **show flow-control er_rdy** command.

Extended Receiver Ready

- ER_RDY is supported only on Fibre Channel ports on
  - Cisco MDS 9000 Series 24/10 SAN Extension Module (DS-X9334-K9) (Fibre Channel ports only)
  - Cisco MDS 9700 Series with Cisco MDS 9700 16-Gbps Fibre Channel Switching Module (DS-X9448-768K9)
  - Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module (DS-X9648-1536K9)
  - Cisco MDS 9700 Series with Cisco MDS 9700 64-Gbps Fibre Channel Switching Module (DS-X9748-3072K9)
  - Cisco MDS 9396S 16G Multilayer Fabric Switch
  - Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switch
  - Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switch
  - Cisco MDS 9220i Multiservice Fabric Switch
  - Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switch
  - Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch
  - Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel Switch
- 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 flow-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.
- If you have enabled ER_RDY in releases prior to Cisco MDS NX-OS Release 9.3(1) using **system fc flow-control er_rdy** command and when you upgrade to Cisco MDS NX-OS Release 9.3(1), the running configuration would display this command as **system fc flow- control er_rdy logical-type core**.
- For migration purposes, port channels can have their member links in both R_RDY and ER_RDY flow-control modes. This is to facilitate non disruptive 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.
- 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.
- From Cisco MDS NX-OS Release 8.5(1), use IOD only if your environment cannot support out-of-order frame delivery. To achieve In-Order Delivery (IOD), enable IOD using the **in-order-guarantee vsan** *id*. When a flow moves from normal VL to slow VL or vice versa, to achieve IOD functionality traffic disruption may be seen. Lossless IOD is not guaranteed.

  Prior to Cisco MDS NX-OS Release 8.5(1), 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 8191 on a Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module, MDS 9132T, MDS 9148T, MDS 9220i, and MDS 9396T switches.

- 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.
- ER_RDY feature is disabled by default. R_RDY is the default flow control mode for all ports.
- F_CTL(17) bit cannot be set for ER_RDY packets as the hardware does not provide support.

Congestion Isolation

- Congestion Isolation is disabled by default.
- The port monitor portguard action for Congestion Isolation is not supported on E (core) ports. Consequently, it should only be configured on a *logical-type edge* port-monitor policy.

  If you are upgrading to Cisco MDS NX-OS Release 8.5(1) or later release and if you have the cong-isolate portguard action configured on a *logical-type core* policy, then you must remove this policy before upgrading.
- 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.
- Prior to Cisco MDS NX-OS Release 8.5(1), 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

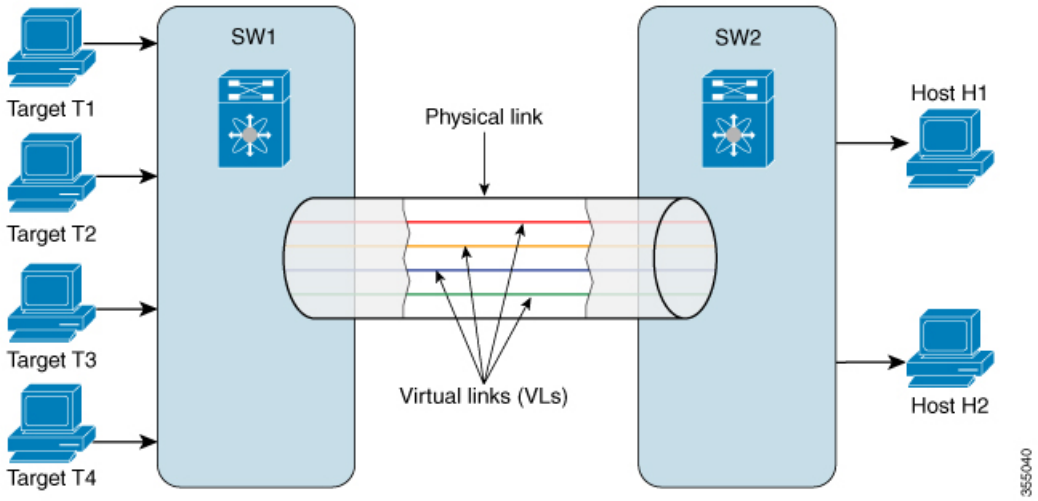
is capable of supporting this feature but does not have it enabled, then the FCNS notification is rejected and the following messages are displayed at the originating switch:

- %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 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 8: 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 and Congestion Isolation Recovery features are enabled. To downgrade to an unsupported release:
  1. If **cong-isolate** or **cong-isolate-recover** port monitor portguard action is configured in a port monitor policy, remove the action from the policy.
  2. Remove any devices that are manually included or excluded as slow-drain devices.
  3. Disable the Congestion Isolation feature.
  4. Reset the flow-control mode to R_RDY.
  5. Flap all the ISLs.
  6. Display the ISLs currently functioning in R_RDY mode.
  7. 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.

Guidelines and Limitations for Fabric Notifications

- Fabric Notifications is supported only on Fibre Channel ports. Fabric Notifications are only supported when the FPM feature is enabled.
- Fabric Notifications is supported only on Cisco MDS 9132T, MDS 9148T, MDS 9220i, MDS 9396S, MDS 9396T, MDS 9706, MDS 9710, and MDS 9718 switches.
- Fabric Notifications is not supported on Cisco MDS 9250i and MDS 9148S switches.
- Fabric Notifications is supported on MDS 9706, MDS 9710, and MDS 9718 switches using 48-port 32-Gbps Fibre Channel Switch module and 48-port 64-Gbps Fibre Channel Switch module.
- In Cisco MDS NX-OS Release 8.5(1), Fabric Notifications is not supported on switches that are operating in the Cisco NPV mode.
- Devices that support FPIN and/or Congestion Signals must register with RDF and EDC for using the Fabric Notifications capabilities.
- Starting from Cisco MDS NX-OS 9.4(5), when FPM is disabled the switch notifies all registered hosts. The switch uses an RDF to indicate changed capabilities. When a host responds with RDF registration, the switch rejects it to indicate FPM is disabled.
- Starting from Cisco MDS NX-OS 9.4(3), oversubscription notification is supported in PMON.
- Starting from Cisco MDS NX-OS 9.4(4), oversubscription notification is supported in SMA.
- In Cisco MDS NX-OS 9.4(5) or earlier releases, FPIN notifications were only sent to the target port after a successful PLOGI (Port Login) from the Fabric Controller.

- Starting with Cisco MDS NX-OS 9.4(5), FPIN notifications will be sent to the target port by the Fabric Controller even if PLOGI (Port Login) has not occurred.
- When FPM feature is disabled, all FIB UP programming is cleaned for all local zoned devices for which it was programmed.
- Fabric Notifications does not monitor devices that are behind vfc interfaces.
- Fabric Notifications supports only Tx of congestion signals and not Rx.
- Fabric Notifications supports following FPIN capabilities:
  - FPIN Link Integrity:
    - Link Failure
    - Loss-of-Synchronization
    - Loss-of-Signal
    - Invalid Transmission Word
    - Invalid CRC
  - FPIN Congestion:
    - Credit Stall
    - Oversubscription is supported from the MDS NX-OS 9.4(4) release.
  - FPIN Peer Congestion:
    - Credit Stall
    - Oversubscription is supported from the MDS NX-OS 9.4(4) release.
- Fabric Notifications do not support following FPIN capabilities:
  - FPIN Link Integrity:
    - Primitive Sequence Protocol Error
  - FPIN Congestion:
    - Oversubscription is not supported prior to MDS NX-OS 9.4(4) release.
    - Lost Credit
  - FPIN Peer Congestion:
    - Oversubscription is not supported prior to MDS NX-OS 9.4(4) release.
    - Lost Credit
  - FPIN Delivery Failure:
    - Timeout
    - Unable to Route

- If the logical type of the port is changed using the **switchport logical-type** command after the device is marked as congested, the device will not be marked as normal automatically. The device needs to be recovered using the **fpm congested-device recover pwwn** *pwwn* **vsan** *id* command.
- For devices that are not registered for FPIN, all the flows destined to slow devices are moved to the low-priority VL. After the slow devices recover from congestion, the flows are moved back to the normal VL.
- Ensure that the portguard action that is configured for slow drain counters is consistent across switches in a fabric.
- Portguard action is initiated from the switch where congestion is detected.
- Port monitor does not take action on devices that part of the excluded list. For more information, see Configuring Excluded List of Congested Devices, on page 272.
- FPIN is not supported on devices that are part of Inter VSAN Routing (IVR) zoneset.
- If you are upgrading to Cisco MDS NX-OS Release 8.5(1) or later release and if the Congestion Isolation feature is enabled, ensure that you disable the Congestion Isolation feature and then enable FPM after upgrading. After upgrading, the port monitor configurations are cleared and it starts detecting events afresh. For enabling the Congestion Isolation feature, see Configuring Congestion Isolation, on page 269.

Guidelines and Limitations for DIRL

- DIRL is supported on the following:

Table 34: List of DIRL Supported Devices in Switch and NPV modes

Device	DIRL in Switch Mode	DIRL in NPV Mode	Line Cards Supported
Cisco MDS 9706	Yes	NA	32 and 64 Gbps
Cisco MDS 9710	Yes	NA	32 and 64 Gbps
Cisco MDS 9718	Yes	NA	32 and 64 Gbps
Cisco MDS 9396T	Yes	Yes*	NA
Cisco MDS 9396S	No	Yes*	NA
Cisco MDS 9250i	No	NA	NA
Cisco MDS 9220i	Yes	NA	NA
Cisco MDS 9148S	No	Yes*	NA
Cisco MDS 9148T	Yes	Yes*	NA
Cisco MDS 9148V	Yes	Yes*	NA
Cisco MDS 9132T	Yes	Yes*	NA
Cisco MDS 9124V	Yes	Yes*	NA

***** DIRL in NPV mode is supported with MDS Release 9.3(2) and later.

- If you are upgrading to Cisco MDS NX-OS Release 8.5(1) or later release and if you have configured the port ingress rate limiting on one or more interfaces, any static ingress rate limiting must be removed using the **no switchport ingress-rate** prior to upgrading to Cisco MDS NX-OS Release 8.5(1) or later release.

  After upgrading to Cisco MDS NX-OS Release 8.5(1) or later, static ingress rate limiting can once again be configured on any interface, if necessary. However, if static ingress rate limiting is configured for an interface, then this interface will not be subject to DIRL.

- Beginning with Cisco MDS NX-OS Release 9.3(1) DIRL is supported in NPV mode with the following behavior:
  - Target ports are excluded by default in switch mode.
  - In NPV mode target ports are not excluded by default as NPV switches do not have access to the FCNS database locally to determine FC4 features/types. Additionally, NPV switches only contain initiator ports and are not recommended to be connected to targets. Due to this, if target ports do exist on a NPV switch, rate limit actions are applied on these ports as well as initiator ports. To exclude specific target ports, use the **fpm dirl exclude list** command.
- In NPV mode, DIRL applies rate limiting by default to both initiator and target ports.
- DIRL is supported only on F ports.
- The following table shows the maximum (lowest) ingress rate limits set by DIRL for each link speed.

Table 35: Maximum Ingress Rates by Hardware Type and Operational Speed

Operational Link Speed	Maximum (lowest) Ingress Rate Limit
64 Gbps	0.01250% (0.4 Gbps)
32 Gbps	0.01250% (0.4 Gbps)
16 Gbps	0.02435% (0.4 Gbps)
8 Gbps	0.04870% (0.4 Gbps)
4 Gbps	0.09741% (0.4 Gbps)

DIRL limitations are:

- DIRL is not supported on Cisco MDS 9250i switches.
- DIRL is not supported on Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module and Cisco MDS 9700 24/10-Port SAN Extension Module.
- DIRL is not supported on switches operating in Cisco NPV mode until Cisco MDS NX-OS Release 9.3(1).

Guidelines

Follw these guidelines for SMA feature.

- SMA is enabled by default. You cannot disable SMA.
- SMA supports a maximum of 5 policies (1 default and 4 user-defined).
- SMA processes is supported on both Core and NPV switches.

Limitations and restrictions

Following are the limitations and restrictions for the SMA feature.

- 1MB of memory for storing the history which accumulate up to 13K entries
- Only port counters and protocol counters are monitored for unexpected patterns.
- Tx-OverUtilization counters is not supported for MDS 24/10 port SAN Extension Module in MDS 9700 series.
- FPIN and DIRL port guard action is not supported in NPV
- Following PMON counters are not supported.
  - ASIC Error Pkt from Port
  - SIC Error Pkt to xbar
  - ASIC Error Pkt from xbar
- The PMON counter TX-Slowport-Oper-Delay is monitorred with threshold type Absoulte but in SMA the counter threshold type is Delta

Configuring Congestion Management

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
- 64-Gbps modules or switches:
  - Cisco MDS 9124V 24-Port 64-Gbps Fibre Channel Switching Module
  - Cisco MDS 9148V 48-Port 64-Gbps Fibre Channel Switching Module
- 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 36: Slow Port Monitor Support on Fibre Channel and FCoE Switching Modules, on page 260 displays the congestion detection features supported on different Fibre Channel and FCoE switching modules for the Cisco MDS NX-OS Release 8.x.

Table 36: 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

Function	Module and Switch Support	
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, from Cisco MDS NX-OS Release 8.2(1) onwards.
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:

Procedure

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, ISLs (E ports) and trunking F and NP ports (TF and TNP ports) will use the core timeout value and non-trunking F ports (F and NP ports) 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 Edge Ports
--------------------------------------------------------------------------------------------------------------
|         Counter         | Threshold  | Interval |      Warning       |     Thresholds
   |             Rising/Falling actions              | Congestion-signal |
|                         |    Type    |  (Secs)
|-------------------|-------------------|-----------------------------------------|-----------------|
|                         |            |          | Threshold | Alerts |  Rising  | Falling
  | Event |     Alerts       |     PortGuard        | Warning | Alarm   |
--------------------------------------------------------------------------------------------------------------
| Credit Loss Reco        | Delta      | 1        | none      | n/a    | 1        | 0
   | 4     | syslog,rmon      | none                | n/a     | n/a     |
| TX Credit Not Available | Delta      | 1        | none      | n/a    | 10%      | 0%
   | 4     | syslog,rmon      | none                | n/a     | n/a     |
| TX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%
   | 4     | syslog,rmon      | none                | n/a     | n/a     |
--------------------------------------------------------------------------------------------------------------
```

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; hence, 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
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 : Not Active
Oper status  : Not Active
Port type    : All Edge Ports
--------------------------------------------------------------------------------------------------------------
|         Counter         | Threshold  | Interval |      Warning       |     Thresholds
   |             Rising/Falling actions              | Congestion-signal |
|                         |    Type    |  (Secs)
```

```
|------------------|------------------|-----------------------------------------|-----------------|
|                          |           |          | Threshold | Alerts |  Rising  | Falling
  | Event |     Alerts       |     PortGuard        | Warning | Alarm   |
-------------------------------------------------------------------------------------------------
| Credit Loss Reco         | Delta      | 1        | none      | n/a    | 1        | 0
   | 4     | syslog,rmon      | none                  | n/a     | n/a     |
| TX Credit Not Available | Delta      | 1        | none      | n/a    | 20%      | 0%
   | 4     | syslog,rmon      | none                  | n/a     | n/a     |
-------------------------------------------------------------------------------------------------
```

Configuring Other Congestion Related Port Monitor Counters

The following port-monitor counters related to SAN congestion can be configured:

Table 37: 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.

Counter Name	Description
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

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:

Procedure

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:

Procedure

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:

Procedure

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-Drop 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:

Procedure

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**

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 269.
2. Configure the Congestion Isolation feature. For more information, see Configuring Congestion Isolation, on page 271.
3. Configure a port-monitor policy with one or more counters containing the portguard action *cong-isolate*. For more information, see Configuring Congestion Isolation, on page 271.

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

Flap the ISLs connecting the local and adjacent switches to enable ER_RDY flow-control mode on the ISLs.

Procedure

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 Enable ER_RDY flow-control mode:

Option	Description
ISL ER_RDY	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.
HBA ER_RDY	switch(config)# **system fc flow-control er_rdy logical-type{core| edge | all}** **Note** • The core option enables ER_RDY flow-control for E/NP ports. • The edge option enables ER_RDY flow-control for F ports. • The all option enables ER_RDY flow-control for all ports.

Step 4 Select a Fibre Channel interface and enter interface configuration submode:

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

Step 5 Gracefully shut down the interface and administratively disable traffic flow:

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

Step 6 Enable traffic flow on the interface:

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

Step 7 Return to privileged executive mode:

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

Step 8 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 Congestion Isolation, on page 271.
2. Disable the Congestion Isolation feature. For more information, see Configuring Congestion Isolation, on page 271.

Procedure

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

To configure Congestion Isolation, perform these steps:

Before you begin

Configure Extended Receiver Ready. For more information, see Enabling Extended Receiver Ready, on page 269.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable Congestion Isolation:

Prior to Cisco MDS NX-OS Release 8.5(1)

switch(config)# **feature congestion-isolation**

Cisco MDS NX-OS Release 8.5(1) and later releases

switch(config)# **feature fpm**

Step 3 Specify the counter parameters for the portguard to take Congestion Isolation action on a port:

Prior to Cisco MDS NX-OS Release 8.5(1)

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**

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

Cisco MDS NX-OS Release 8.5(1) and later releases

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* **portguard cong-isolate**

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

Note
Absolute counters do not support portguard actions. However, the tx-slowport-oper-delay absolute counter supports Congestion Isolation portguard action.

Step 4 Activate the specified port-monitor policy:

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

Configuring Excluded List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly exclude a device from congestion actions, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device exclude mode:

switch(config)# **fpm congested-device exclude list**

Step 3 Exclude a device from congestion actions:

switch(config-congested-dev-exc)# **member pwwn** *pwwn* **vsan** *id*

Configuring Static List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly configure a device as congested, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device static mode:

switch(config)# **fpm congested-device static list**

Step 3 Configure a device as congested:

switch(config-congested-dev-static)# **member pwwn** *pwwn* **vsan** *id* **credit-stall**

Recovering a Congested Device

Use this procedure to recover congested device that was detected by port monitor.

From Cisco MDS NX-OS Release 8.5(1), to recover a device from congestion, perform this step:

Procedure

Recover a device from congestion:

switch# **fpm congested-device recover pwwn** *pwwn* **vsan** *id*

Including or Excluding a Congested Device

Prior to Cisco MDS NX-OS Release 8.5(1), to explicitly include a device as congested such that it is identified as a congested device by the port monitor, or exclude a device that is identified as a congested device by the port monitor, perform these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Explicitly include a device as congested or exclude a device from being detected as congested:

switch# **congestion-isolation {exclude | include} pwwn** *pwwn* **vsan** *vsan-id*

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.

Prior to Cisco MDS NX-OS Release 8.5(1), to manually remove an interface from being detected as slow, perform these steps:

Procedure

Remove an interface from being detected as slow by the port monitor:

switch#: **congestion-isolation remove interface** *slot*/*port*

Configuring Congestion Isolation Recovery

To configure the Congestion Isolation Recovery feature, perform these steps:

Before you begin

Enable Extended Receiver Ready. For more information, see Enabling Extended Receiver Ready, on page 269.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable FPM:

switch(config)# **feature fpm**

Step 3 Specify the policy name and enter port monitoring policy configuration mode:

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

Step 4 Specify the counter parameters for the portguard to take Congestion Isolation Recovery 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-recover**

Note

Absolute counters do not support portguard actions. However, the tx-slowport-oper-delay absolute counter supports Congestion Isolation Recovery portguard actions.

Step 5 Return to configuration mode:

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

Step 6 (Optional) Change recovery-interval:

switch(config)# **port-monitor cong-isolation-recover recovery-interval** *seconds*

Step 7 (Optional) Specify isolate-duration:

switch(config)# **port-monitor cong-isolation-recover isolate-duration** *hours* **num-occurrence** *number*

Step 8 Activate the specified port-monitor policy:

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

Step 9 (Optional) You can manually exclude a device to be detected as a slow device.

See Configuring Excluded List of Congested Devices, on page 272.

Configuring Static List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly configure a device as congested, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device static mode:

switch(config)# **fpm congested-device static list**

Step 3 Configure a device as congested:

switch(config-congested-dev-static)# **member pwwn** *pwwn* **vsan** *id* **credit-stall**

Configuring Excluded List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly exclude a device from congestion actions, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device exclude mode:

switch(config)# **fpm congested-device exclude list**

Step 3 Exclude a device from congestion actions:

switch(config-congested-dev-exc)# **member pwwn** *pwwn* **vsan** *id*

Recovering a Congested Device

Use this procedure to recover congested device that was detected by port monitor.

From Cisco MDS NX-OS Release 8.5(1), to recover a device from congestion, perform this step:

Procedure

Recover a device from congestion:

switch# **fpm congested-device recover pwwn** *pwwn* **vsan** *id*

Configuring Fabric Notifications

Enabling FPM

Feature FPM must be enabled on the switch to support FPINs and Congestion Signals. If the switch receives an RDF or EDC and feature FPM is not enabled then the switch will reply with a *LS_Reject*.

To enable FPM, perform these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enable FPM:

switch# **feature fpm**

Disabling FPM

To disable FPM, perform these steps:

If the switch receives an Register Diagnostic Functions (RDF) or Exchange Diagnostic Capabilities (EDC) and feature FPM is disabled, then the switch replies with a *LS_Reject*.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Disable FPM:

switch# **no feature fpm**

Configuring the Port-Monitor Portguard Action for FPIN

To configure the port-monitor portguard action for FPIN, perform these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enable FPM:

switch(config)# **feature fpm**

Step 3 Specify the policy name and enter port monitoring policy configuration mode:

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

Step 4 Specify the counter parameters for the portguard for FPIN:

switch(config-port-monitor)# **counter** {**invalid-crc** | **invalid-words** | **link-loss** | **signal-loss** | **sync-loss** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *event-id* **warning-threshold** *count2* **falling-threshold** *count3* **portguard FPIN**

Step 5 Return to configuration mode:

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

Step 6 Activate the specified port-monitor policy:

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

Step 7 (Optional) Specify the recovery interval. By default, the recovery interval is set to 900 seconds (15 minutes).

switch(config)# **port-monitor fpin recovery-interval** *seconds*

Step 8 (Optional) Specify the isolate duration:

switch(config)# **port-monitor fpin isolate-duration** *hours* **num-occurrence** *number*

Configuring Static List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly configure a device as congested, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device static mode:

switch(config)# **fpm congested-device static list**

Step 3 Configure a device as congested:

switch(config-congested-dev-static)# **member pwwn** *pwwn* **vsan** *id* **credit-stall**

Configuring Excluded List of Congested Devices

From Cisco MDS NX-OS Release 8.5(1), to explicitly exclude a device from congestion actions, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter congested device exclude mode:

switch(config)# **fpm congested-device exclude list**

Step 3 Exclude a device from congestion actions:

switch(config-congested-dev-exc)# **member pwwn** *pwwn* **vsan** *id*

Recovering a Congested Device

Use this procedure to recover congested device that was detected by port monitor.

From Cisco MDS NX-OS Release 8.5(1), to recover a device from congestion, perform this step:

Procedure

Recover a device from congestion:

switch# **fpm congested-device recover pwwn** *pwwn* **vsan** *id*

Configuring FPIN Notification Interval

To change the default FPIN notification interval, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Change the FPIN notification interval:

switch(config)# **fpm fpin period** *seconds*

By default, the FPIN notification interval is three minutes.

Configuring EDC Congestion Signal

To configure the EDC interval for sending congestion signal, perform these steps:

Before you begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

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 congestion signals:

switch(config-port-monitor)# **counter txwait warning-signal-threshold** *count1* **alarm-signal-threshold** *count2* **portguard congestion-signals**

Step 4 (Optional) Exit the port monitor configuration mode:

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

Step 5 (Optional) Specify the EDC switch-side period for sending congestion signal. By default, the switch-side congestion signal period is set to 1 second.

switch(config)# **fpm congestion-signal period** *seconds*

Triggering HBA RDF Capability Registration from Switch

Before you begin

Ensure the end device HBAs support RDF.

Ensure that FPM is enabled. For more information, see Enabling FPM, on page 276

You can use the **show fpm registration summary** command to show the summary of the device registrations. For more information, see Verifying RDF, EDC and FPIN Information, on page 309.

Note This CLI command is supported starting from NX-OS 9.4(4) release.

Procedure

Initiate RDF registration process for the devices in the specified VSAN using any one of the following options:

- Initiate RDF registration process for specific devices in the specified VSAN.

```
switch# send fc rdf fcid <fcid>  vsan <vsan_id>
```

- Initiate RDF registration process for all devices logged into the switch in the specified VSAN.

```
switch# send fc rdf fcid all vsan <vsan_id>
```

Configuring DIRL

Before You Begin

Enable FPM. For more information, see Enabling FPM, on page 276.

Configuring the Port-Monitor Portguard Action for DIRL

To configure the port-monitor portguard action for DIRL, perform these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enable FPM:

switch(config)# **feature fpm**

Step 3 Specify the policy name and enter port monitoring policy configuration mode:

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

Step 4 Specify the counter parameters for the portguard for DIRL:

switch(config-port-monitor)# **counter** {**tx-datarate** | **tx-datarate-burst** | **txwait**} **poll-interval** *seconds* {**absolute** | **delta**} **rising-threshold** *count1* **event** *event-id* **warning-threshold** *count2* **falling-threshold** *count3* **portguard DIRL**

Step 5 Return to configuration mode:

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

Step 6 Activate the specified port-monitor policy:

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

Step 7 (Optional) Specify the recovery interval. By default, the recovery interval is set to 60 seconds .

switch(config)# **port-monitor dirl recovery-interval** *seconds*

Configuring DIRL Rate Reduction and Recovery Percentages

To configure the DIRL rate reduction percentages, perform these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 (Optional) Specify the ingress rate reduction and recovery percentages:

switch(config)# **fpm dirl reduction** *percentage* **recovery** *percentage*

What to do next

To configure ingress port rate limit, see Configuring Static Ingress Port Rate Limiting, on page 283.

Excluding Interfaces from DIRL Rate Reduction

To exclude an interface from DIRL rate reduction, perform these steps:

Note Interfaces having devices with FC4-feature as *init* are monitored by default. If other interfaces need to be monitored, use the **no member fc4-feature target** command.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Enter DIRL exclude list mode:

switch(config)# **fpm dirl exclude list**

Step 3 Specify an interface:

switch(config-dirl-excl)# **member interface fc** *slot/port*

Step 4 Specify the interface to be excluded from DIRL rate reduction:

switch(config-dirl-excl)# **member** {**fc4-feature target** | **interface fc** *slot/port*}

Recovering Interfaces from DIRL Rate Reduction

To recover interface from DIRL rate reduction, perform these steps:

Procedure

Recover interface from DIRL rate reduction:

switch# **fpm dirl recover interface fc** *slot/port*

Configuring Static Ingress Port Rate Limiting

To configure the static port rate limiting value, follow these steps

Before you begin

From Cisco MDS NX-OS Release 8.5(1), you need to enable FPM before configuring the port rate limiting value. For more information, see Enabling FPM, on page 276.

Procedure

Step 1 Enter configuration mode:

switch# **configure**

Step 2 Select the interface to specify the static ingress port rate limit:

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

Step 3 Configure the static port rate limit for the selected interface:

switch(config-if)# **switchport ingress-rate** *limit*

Create a SMA custom entity group

Follow these steps to create a entity group.

Procedure

Step 1 Configure a name to entity-group

Step 2 Add entity-type for the entity group

Step 3 Add sub entity-type for the entity group

What to do next

Add the entity group to a SMA policy.

Create or delete custom entity group

You need a entity-group to act as a filter to monitor similar kind of entities for the SMA poilcy. Entities such as all edge ports, all core ports, ports are filtered with attributes like FDMI, FCNS, device-alias, FCID, FPIN and so on.

The following conditions are applicable for adding or deleting entity-groups.

- Multiple entity-groups can be added to a SMA policy.
- Entity-type or sub-entity cannot be added or deleted when the entity-group is present in an active policy.
- You cannot delete a entity-group that is part of an inactive policy.

Procedure

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Create an entity-group for using in the SMA policy.

```
switch(config)# [no] sma entity-group name entity-group-name
```

Step 3 Add the entity-type to the entity-group.

```
switch(config-sma-entity-group)# [no] entity-type fc-port {all | core | edge | interface}
```

Create a SMA custom action group

Follow these steps to create action groups.

Procedure

Step 1 Configure a name to action-group

Step 2 Configure actions for each level (alarm, warning, falling)

What to do next

Add the action group to a SMA policy.

Create or delete SMA custom action-group

Action groups provides configuration of actions for a SMA policy. Actions such as flap, error-disable, syslog, call-home, FPIN, SNMP-trap and so on.

Procedure

Step 1 Enter configuration mode.

Example:

```
switch# configure terminal
```

Step 2 Create an action-group based on the counter group for usage in the SMA policy.

Example:

```
switch(config)# [no] sma action-group name action-group-name counter-group {cg-all | cg-congestion
| cg-datarate | cg-fcport-all | cg-link-integrity | cg-protocol-all | cg-slowdrain }
```

Step 3 Create an action-group based on the counter for usage in the SMA policy.

Example:

```
switch(config)# [no] sma action-group name action-group-name  counter {cg-all | cg-congestion |
cg-datarate | cg-fcport-all | cg-link-integrity | cg-protocol-all | cg-slowdrain }
```

Configure SMA Policy

You can create SMA policy by using one of the following options.

- Clone the default policy to create a custom policy or
- Activate the default policy or
- Create a new user defined SMA policy.

Note

- Custom entity-groups and action-groups are created outside the policy before using them under policy.
- Pre-defined entity-groups such as eg-fcport-all, eg-fcport-edge,eg-fcport-core entity-groups are used to configure port counters.
- Pre-defined entity-group such as eg-switch entity-group are used to configure protocol counters.

Procedure

Step 1 Configure a name to the policy

Step 2 Configure entity-group (pre-defined / user defined)

Step 3 Configure counters or counter-groups for entity-group

Step 4 Configure monitoring group with these thresholds.

- default or
- user defined period values for alarm or warning or falling.

Step 5 Configure action group

Associate either pre-defined or user defined action group that is created for existing counter or counter-group.

Create a user-defined SMA Policy

You can create a user-defined SMA policy.

Before you begin

Ensure the following are user-defined

- entity-groups
- action-groups

Procedure

Step 1 Enter configuration mode.

```
switch# configure terminal
```

Step 2 Configure the policy.

```
switch(config)# [no] sma policy name policy_name
```

Step 3 Configure the entity-group

```
switch(config)# [no] entity-group name entity-group-name [{before|after} name ref-entity-group-name]
```

Step 4 Configure the counter-group containing the monitor-levels, intervals and action groups.

```
switch(config)# [no] {counter counter-name|counter-group countergrp-name} monitor-level {default|
period period alarm alarm-threshold [warning warn-threshold] falling fall-threshold} action-group
{ag-default|ag-none|action-group-name}
```

Activate or deactivate the default SMA Policy

Activate the default SMA policy available in MDS 9000.

Before you begin

You can activate a default policy on when these conditions are met.

- No active policy must exist in order to activate a policy.

- When the no form of the command is used, all the configuration details are to be provided.
- Policy must be active for the policy to be deactivated.

Procedure

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Activate or deactivate SMA policy.

```
switch(config)# [no] sma policy activate name {default | policy_name}
```

Creating or deleting a name to user defined SMA Policy

Provide a name for the user-defined policy.

Procedure

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Create an user defined policy.

```
switch(config)# [no] sma policy name policy_name
```

Note
When the no form of the command is used, all the configuration details are to be provided.

Clone or copy an default SMA Policy

Clone or copy an existing SMA policy.

Before you begin

You cannot copy or clone a policy that is active in MDS 9000.

Procedure

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Copy the SMA policy.

```
switch(config)#  sma policy name copy default policy_name
```

Configuration Examples for Congestion Management

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 HBA ER_RDY flow-control mode:

```
switch# configure terminal
switch(config)# system fc flow-control er_rdy logical-type{core| edge | all}
Use the CLI show flow-control r_rdy to list the ports that are still in R_RDY mode. The
core option enables ER_RDY flow-control for E/NP ports. The edge option enables ER_RDY
flow-control for F ports. The all option enables ER_RDY flow-control for all ports.
```

This example shows how to disable HBA 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 ER_RDY flow-control mode:

```
switch# configure terminal
switch(config)# system fc flow-control er_rdy logical-type core
Use the CLI show flow-control r_rdy to list the ports that are still in R_RDY mode. The
core option enables ER_RDY flow-control for E and NP ports.
```

This example shows how to disable HBA 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 ISL ER_RDY flow-control mode for releases prior to Cisco MDS NX-OS Release 9.3(1) :

```
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 ISL 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 for releases prior to Cisco MDS NX-OS Release 8.5(1):

```
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 enable Congestion Isolation in Cisco MDS NX-OS Release 8.5(1) and later release:

```
switch# configure terminal
switch(config)# feature fpm
```

This example shows how to disable Congestion Isolation for releases prior to Cisco MDS NX-OS Release 8.5(1):

```
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 disable Congestion Isolation in Cisco MDS NX-OS Release 8.5(1) and later releases:

```
switch# configure terminal
switch(config)# no feature fpm
```

This example shows how to manually configure a device as a congested device for releases prior to Cisco MDS NX-OS Release 8.5(1). The configured device will be permanently treated as a congested device until it is removed from congestion isolation. 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 manually configure a device as a congested device in Cisco MDS NX-OS Release 8.5(1) and later releases. The configured device will be permanently treated as a congested device until it is removed from congestion isolation. 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)# fpm congested-device static list
switch(config-congested-dev-static)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 4 credit-stall
```

This example shows how to configure a device that is to be excluded from automatic congestion isolation by the port monitor for releases prior to Cisco MDS NX-OS Release 8.5(1). 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 congested 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
```

This example shows how to configure a device that is to be excluded from automatic congestion isolation by the port monitor in Cisco MDS NX-OS Release 8.5(1) and later releases. 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 congested 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)# fpm congested-device exclude list
switch(config-congested-dev-exc)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 4
```

Congested devices can be identified either via the port monitor or manually included or excluded. 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 congested device.

This example shows how to manually remove an interface from being detected as slow in the port monitor:

Prior to Cisco MDS NX-OS Release 8.5(1)

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)
================================================
```

From Cisco MDS NX-OS Release 8.5(1)

1. Identifying the interface that you want to remove from being detected as slow.

```
switch# show fpm congested-device database local
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
-------------------------------------------------------------------------------------------------
PWWN                     | FCID     | Event type   | Detect type | Detect Time
-------------------------------------------------------------------------------------------------
21:00:f4:e9:d4:54:ac:f8 | 0x7d0000 | credit-stall | local-pmon  | Thu Jan 28 05:08:31
2021
```

2. Remove the interface from being marked as slow.

```
switch# configure
switch(config)# fpm congested-device exclude list
switch(config)# member pwwn 21:00:f4:e9:d4:54:ac:f8 vsan 50
```

3. Verifying if the interface is removed from being detected as slow.

```
switch# show fpm congested-device database local
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
-------------------------------------------------------------------------------------------------
No congested devices found
```

Configuring Examples for Congestion Isolation Recovery

This example shows how to configure the *isolate-duration* to 24-hours and the number of rising threshold occurrences to be detected in this interval to 3:

```
switch# configure
switch(config)# port-monitor cong-isolation-recover isolate-duration 24 num-occurrence 3
```

This example shows how to configure the *recovery-interval* to 15 minutes:

```
switch# configure
switch(config)# port-monitor cong-isolation-recover recovery-interval 15
```

This example shows how to manually include the device with pWWN 10:00:00:00:c9:f9:16:8d in VSAN 2 as a slow device:

```
switch# configure
switch(config)# fpm congested-device static list
switch(config-congested-dev-static)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 2 credit-stall
```

This example shows how to manually exclude the device with pWWN 10:00:00:00:c9:f9:16:8d in VSAN 2 as a slow device:

```
switch# configure
switch(config)# fpm congested-device exclude list
switch(config-congested-dev-exc)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 2
```

Configuring Examples for Fabric Notifications

This example shows how to enable FPM on a switch:

```
switch# configure
switch(config)# feature fpm
```

This example shows how to disable FPM on a switch:

```
switch# configure
switch(config)# no feature fpm
```

This example shows how to explicitly configure a device with pWWN 10:00:00:00:c9:f9:16:8d in VSAN 2 as congested:

```
switch# configure
switch(config)# fpm congested-device static list
switch(config-congested-dev-static)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 2 credit-stall
```

This example shows how to explicitly exclude a device with pWWN 10:00:00:00:c9:f9:16:8d in VSAN 2 from congestion actions:

```
switch# configure
switch(config)# fpm congested-device exclude list
switch(config-congested-dev-exc)# member pwwn 10:00:00:00:c9:f9:16:8d vsan 2
```

This example shows how to recover the device with pWWN 10:00:00:00:c9:f9:16:8d in VSAN 2 from congestion actions:

```
switch# fpm congested-device recover pwwn 10:00:00:00:c9:f9:16:8d vsan 2
```

This example shows how to configure an FPIN notification interval of 30 seconds:

```
switch# configure
switch(config)# fpm fpin period 30
```

This example shows how to configure the EDC interval for sending congestion signal every 30 seconds:

```
switch# configure
switch(config)# fpm congestion-signal period 30
```

Configuation Examples for DIRL

This example shows how to configure DIRL to specify the ingress reduction rate to 50 percent and ingress recovery rate to 30 percent:

```
switch# configure
switch(config)# fpm dirl reduction 50 recovery 30
```

This example shows how to exclude DIRL based on interface:

```
switch# configure
switch(config)# fpm dirl exclude list
switch(config-dirl-excl)# member interface fc 1/1
switch(config-dirl-excl)# member interface fc 1/1
```

This example shows how to include FC4-type target connected device interfaces in DIRL:



Note Interfaces connected to targets (devices registered with FC4-type target) are excluded by default. To include them, use the **no member fc4-feature target** command.

```
switch# configure
switch(config)# fpm dirl exclude list
switch(config-dirl-excl)# no member fc4-feature target
```

This example shows how to recover interface fc1/1 which is under DIRL to normal:

```
switch# fpm dirl recover interface fc 1/1
```

This example shows how to configure port-monitor policy containing TxWait and Tx-Datarate counters for DIRL. The rate-reduction value is set to 10% and the rate-recovery value is also set to 10%:

```
switch(config)#feature fpm
switch(config)#fpm dirl reduction 10 recovery 10
switch(config)# port-monitor name edge-dirl
switch(config-port-monitor)# logical-type edge
switch(config-port-monitor)# no monitor counter all
switch(config)# monitor counter txwait
switch(config)# counter txwait poll-interval 1 delta rising-threshold 30 event 4
falling-threshold 10 portguard dirl
switch(config)# monitor counter tx-datarate
switch(config)# counter tx-datarate poll-interval 10 delta rising-threshold 80 event 4
```

```
falling-threshold 70 portguard dirl
switch(config)# port-monitor activate edge-dirl
```

Verifying Congestion Management

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         |
------------------------------------------------------------------
```

TxWait on FCoE or Virtual Fibre Channels (VFC)



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 2.5 microseconds and are converted to seconds in some command outputs. To convert to seconds, multiply 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)
================================================================================

Po1               Auto  NA      (8)     3    0        0        0(0)      0(0)
Po350             Auto  NA      (8)     3    0        0        0(0)      0(0)
Po351             Auto  NA      (8)     3    0        0        0(0)      0(0)
Po552             Auto  NA      (8)     3    111506   0        0(0)      5014944(12)
Po700             Auto  NA      (8)     3    0        0        0(0)      0(0)
Eth2/17           Auto  Off
Eth2/18           Auto  Off
Eth2/19           Auto  Off
Eth2/20           Auto  Off
Eth2/25           Auto  On      (8)     3    0        0        0(0)      0(0)
Eth2/26           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 detailed
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 TxWait history information for Ethernet 2/47:

```
switch# show interface e2/47 txwait-history
```

```
TxWait history for port Eth2/47:
================================
     3333333333433333333333333333333333333333333333333333333333333333
     5555655555055546375555555554556555555555465554635665555555555
     5570286961955529745559779559555555580995925559585259577986955
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)

   22222222222
   1111111111600000000000000000000000000000000000000000000000000000
   ............................................................
   3344454444320000000000000000000000000000000000000000000000000000
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)


     2
     3
     1000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
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 RxWait history information for Ethernet 1/47:

```
switch# show interface e1/47 rxwait-history

RxWait history for port Eth1/47:
  ==============================

      7887777787777778777777788778778777787778777777777777877777778
      9009999909999990999999900990990999909990999999999991999999900
  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)

              1
     4444444445570000000000000000000000000000000000000000000000000
     ..............................................................
     7777777775870000000000000000000000000000000000000000000000000
  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)


      2                  1                           1
      7                2 5                           9
      0000000000000000060600020000000000000000000000090000000000000000000001
  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
```



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

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)
```

Verifying Congestion Isolation Recovery

This example shows how to check the configured *isolate-duration*, *recovery-interval*, and number of rising threshold occurrences:

```
switch# show port-monitor

Port Monitor : enabled
DIRL :
      Recovery Interval   : 60 seconds
FPIN :
      Recovery Interval   : 900 seconds
Cong-isolate-recover :
      Recovery Interval   : 900 seconds
      Isolation Duration   : 24 hours
      Number of Isolation occurrences   : 3
--------------------------------------------------------------------------------

Policy Name : default
Admin status : Not Active
Oper status  : Not Active
Logical type : All Ports
--------------------------------------------------------------------------------------------------------------------------------------------
|        Counter        | Threshold  | Interval |     Warning      |     Thresholds     |         Rising/Falling actions     | Congestion-signal
 |
|                       |    Type    |  (Secs)
|-----------------------|------------|----------------------------------------------|-------|
|                       |            |          | Threshold | Alerts |  Rising  | Falling  | Event |     Alerts       | PortGuard| Warning |
Alarm   |
--------------------------------------------------------------------------------------------------------------------------------------------
| Link Loss               | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Sync Loss               | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Signal Loss             | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Invalid Words           | Delta      | 60       | none      | n/a    | 1        | 0        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Invalid CRC's           | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| State Change            | Delta      | 60       | none      | n/a    | 5        | 0        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX Discards             | Delta      | 60       | none      | n/a    | 200      | 10       | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| LR RX                   | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| LR TX                   | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Timeout Discards        | Delta      | 60       | none      | n/a    | 200      | 10       | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| Credit Loss Reco        | Delta      | 60       | none      | n/a    | 1        | 0        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX Credit Not Available | Delta      | 60       | none      | n/a    | 10%      | 0%       | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| RX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%      | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%      | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX-Slowport-Oper-Delay  | Absolute   | 60       | none      | n/a    | 50ms     | 0ms      | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TXWait                  | Delta      | 60       | none      | n/a    | 30%      | 10%      | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| RX Datarate Burst       | Delta      | 10       | none      | n/a    | 5@90%    | 1@90%    | 4     | syslog,rmon,obfl | none     | n/a     |
n/a     |
| TX Datarate Burst       | Delta      | 10       | none      | n/a    | 5@90%    | 1@90%    | 4     | syslog,rmon,obfl | none     | n/a     |
n/a     |
| Input Errors            | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
--------------------------------------------------------------------------------------------------------------------------------------------

Policy Name : slowdrain
Admin status : Not Active
Oper status  : Not Active
Logical type : All Edge Ports
--------------------------------------------------------------------------------------------------------------------------------------------
|        Counter        | Threshold  | Interval |     Warning      |     Thresholds     |    Rising/Falling actions          | Congestion-signal
 |
|                     |    Type    |  (Secs)  |------------------|--------------------|-------------------------------------|-----|
|                       |            |          | Threshold | Alerts |  Rising  | Falling  | Event |     Alerts       |PortGuard| Warning |
Alarm   |
--------------------------------------------------------------------------------------------------------------------------------------------
| Credit Loss Reco        | Delta      | 1        | none      | n/a    | 1        | 0        | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX Credit Not Available | Delta      | 1        | none      | n/a    | 10%      | 0%       | 4     | syslog,rmon      | none     | n/a     |
n/a     |
| TX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%      | 4     | syslog,obfl      | none     | n/a     |
n/a     |
--------------------------------------------------------------------------------------------------------------------------------------------

Policy Name : fabricmon_edge_policy
Admin status : Not Active
Oper status  : Not Active
Logical type : All Edge Ports
--------------------------------------------------------------------------------------------------------------------------------------------
|        Counter        | Threshold  | Interval |     Warning      |     Thresholds     |     Rising/Falling actions         | Congestion-signal
 |
|                     |    Type    |  (Secs)  |------------------|--------------------|-------------------------------------|-----|
|                       |            |          | Threshold | Alerts |  Rising  | Falling  | Event |     Alerts       |PortGuard| Warning |
```

```
Alarm   |
--------------------------------------------------------------------------------------------------------------------------------
| Link Loss               | Delta      | 30       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | FPIN    | n/a     | n/a
    |
| Sync Loss               | Delta      | 30       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | FPIN    | n/a     | n/a
    |
| Signal Loss             | Delta      | 30       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | FPIN    | n/a     | n/a
    |
| Invalid Words           | Delta      | 30       | none      | n/a    | 1        | 0        | 4     | syslog,rmon      | FPIN    | n/a     | n/a
    |
| Invalid CRC's           | Delta      | 30       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | FPIN    | n/a     | n/a
    |
| State Change            | Delta      | 60       | none      | n/a    | 5        | 0        | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| TX Discards             | Delta      | 60       | none      | n/a    | 200      | 10       | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| LR RX                   | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| LR TX                   | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| Timeout Discards        | Delta      | 60       | none      | n/a    | 200      | 10       | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| Credit Loss Reco        | Delta      | 1        | none      | n/a    | 1        | 0        | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| TX Credit Not Available | Delta      | 1        | none      | n/a    | 10%      | 0%       | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| RX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%      | 4     | syslog,rmon,obfl | none    | n/a     | n/a
    |
| TX Datarate             | Delta      | 10       | none      | n/a    | 80%      | 70%      | 4     | syslog,rmon,obfl | none    | n/a     | n/a
    |
| TX-Slowport-Oper-Delay  | Absolute   | 1        | none      | n/a    | 50ms     | 0ms      | 4     | syslog,rmon      | none    | n/a     | n/a
    |
| TXWait                  | Delta      | 1        | none      | n/a    | 30%      | 10%      | 4     | syslog,rmon      | FPIN    | 40%     | 60%
    |
| RX Datarate Burst       | Delta      | 10       | none      | n/a    | 5@90%    | 1@90%    | 4     | syslog,rmon,obfl | none    | n/a     | n/a
    |
| TX Datarate Burst       | Delta      | 10       | none      | n/a    | 5@90%    | 1@90%    | 4     | syslog,rmon,obfl | none    | n/a     | n/a
    |
| Input Errors            | Delta      | 60       | none      | n/a    | 5        | 1        | 4     | syslog,rmon      | none    | n/a     | n/a
    |
--------------------------------------------------------------------------------------------------------------------------------

 On falling threshold portguard actions FPIN, DIRL, Cong-Isolate-Recover will initiate auto recovery of ports.
```

Verifying RDF, EDC and FPIN Information

This example shows the number of devices registered for FPIN in each VSAN:

```
switch# show fpm fpin
C: Congestion Notification Descriptor
P: Peer Congestion Notification Descriptor
L: Link Integrity Notification Descriptor
D: Delivery Notification Descriptor
U: Priority Update Notification Descriptor
A: Alarm Signal
W: Warning Signal

VSAN: 1
--------------------------------------------------------------------------------------------
FCID                     |            RDF             | FPIN sent  | Last FPIN sent timestamp
PWWN                     | Registered | Negotiated    |   count    |
                         |          Timestamp         |            |
--------------------------------------------------------------------------------------------
0xdc06e0                 | L          | L             | L:       0 | L: --
10:00:00:10:9b:95:41:22  | Tue Feb  2 03:38:13 2021   |            |


VSAN: 50
--------------------------------------------------------------------------------------------
FCID                     |            RDF             | FPIN sent  | Last FPIN sent timestamp
PWWN                     | Registered | Negotiated    |   count    |
                         |          Timestamp         |            |
--------------------------------------------------------------------------------------------
0x7d0000                 | CPLD       | CPL           | L:       0 | L: --
21:00:f4:e9:d4:54:ac:f8  | Mon Feb  1 15:32:26 2021   | C:       0 | C: --
                         |            |               | P:       0 | P: --
```

```
0x7d0020                | CPLD       | CPL          | L:       0 | L: --
21:00:f4:e9:d4:54:ac:f9 | Mon Feb  1 15:32:27 2021 | C:       0 | C: --
                        |            |              | P:       0 | P: --
```

This example shows a summary of RDF and EDC registrations:

```
switch# show fpm registration summary
C: Congestion Notification Descriptor
P: Peer Congestion Notification Descriptor
L: Link Integrity Notification Descriptor
D: Delivery Notification Descriptor
U: Priority Update Notification Descriptor
A: Alarm Signal
W: Warning Signal

VSAN: 1
--------------------------------------------------------------------------
FCID     | PWWN                    | FPIN          | Congestion Signal
         |                         | Registrations | Registrations
--------------------------------------------------------------------------
0xdc06e0 | 10:00:00:10:9b:95:41:22 | L             | --

VSAN: 50
--------------------------------------------------------------------------
FCID     | PWWN                    | FPIN          | Congestion Signal
         |                         | Registrations | Registrations
--------------------------------------------------------------------------
0x7d0000 | 21:00:f4:e9:d4:54:ac:f8 | CPLD          | AW
0x7d0020 | 21:00:f4:e9:d4:54:ac:f9 | CPLD          | AW
```

This example shows EDC registration in detail:

```
switch# show fpm registration congestion-signal
A: Alarm
W: Warning
ms: milliseconds

VSAN: 1
-----------------------------------------------------------------------------------------------
No registered devices found

VSAN: 50
-----------------------------------------------------------------------------------------------
FCID     | PWWN                    |     Device Tx    |     Device Rx    |  Negotiated Tx
         |                         | Capa- | Interval | Capa- | Interval | Capa- | Interval

         |                         | bility|  (ms)    | bility|  (ms)    | bility|  (ms)
-----------------------------------------------------------------------------------------------
0x7d0020 | 21:00:f4:e9:d4:54:ac:f9 | AW    |       10 | AW    |       10 | AW    |     1000
0x7d0000 | 21:00:f4:e9:d4:54:ac:f8 | AW    |       10 | AW    |       10 | AW    |     1000
```

This example shows the list of devices that were detected as congested devices by port monitor:

```
switch# show fpm congested-device database local
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
```

```
-------------------------------------------------------------------------------------------------------
PWWN                    | FCID     | Event type   | Detect type | Detect Time
-------------------------------------------------------------------------------------------------------
21:00:f4:e9:d4:54:ac:f8 | 0x7d0000 | credit-stall | local-pmon  | Thu Jan 28 05:08:31 2021
```

This example shows a list of remote devices that are congested:

```
switch# show fpm congested-device database remote
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
----------------------------------
No congested devices found

VSAN: 70
----------------------------------
No congested devices found

VSAN: 80
----------------------------------
No congested devices found

VSAN: 1001
-------------------------------------------------------------------------------------------------------
PWWN                    | FCID     | Event type   | Detect type | Detect Time
-------------------------------------------------------------------------------------------------------
21:00:34:80:0d:6c:a7:63 | 0xec0000 | credit-stall | remote      | Thu Jan 28 05:12:00 2021
```

This example shows the list of devices that were manually included as congested devices:

```
switch# show fpm congested-device database static
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
-------------------------------------------------
PWWN                    | FCID     | Event type
-------------------------------------------------
21:00:f4:e9:d4:54:ac:f8 | 0x7d0000 | credit-stall
```

This example shows the list of congested devices that are excluded:

```
switch# show fpm congested-device database exclude
VSAN: 1
----------------------------------
No congested devices found

VSAN: 50
-----------------------------------
PWWN                    | FCID
-----------------------------------
21:00:f4:e9:d4:54:ac:f8 | 0x7d0000
```

Verifying DIRL

This example shows the configured DIRL reduction and recovery percentages:

```
switch# show fpm ingress-rate-limit status
dirl reduction rate:50%
dirl recovery rate:25%
--------------------------------------------------------------------------------
Interface  Current rate   Rate-limit-type  Previous action  Last update time
           limit(%)
--------------------------------------------------------------------------------
fc4/12     10.6435          dynamic          recovered        Wed Jan 27 20:23:34 2021

fc7/5      12.9567          dynamic          recovered        Wed Jan 27 20:23:34 2021
```

This example shows the configured DIRL reduction and recovery percentages for the port fc4/12:

```
switch# show fpm ingress-rate-limit status interface fc4/12
dirl reduction rate:50%
dirl recovery rate:25%
--------------------------------------------------------------------------------
Interface  Current rate  Rate-limit-type  Previous action  Last update time
           limit(%)
--------------------------------------------------------------------------------
fc4/12     10.6435          dynamic          recovered        Wed Jan 27 20:23:34 2021
```

This example shows the list of interfaces that are excluded from DIRL rate reduction:

```
switch# show fpm dirl exclude
All target device connected interface are excluded from DIRL
--------------
Interface
--------------
fc4/19
fc4/21
fc7/13
```

Action group example

Action groups examples are described.

View a action-group with a given name

Displays the action-group with the given name.

```
switch(config)# show sma action-group name
<action-group-name>
```

View default and user-defined action-group

Displays the action-group of class as system or user. Class system refers to system-defined and user refers to user-defined.

```
switch(config)# show sma action-group class {system | user
}
```

SMA policy examples

SMA policy examples are described.

SMA default policy

```
switch# show sma policy active
sma policy name defaultentity-group name eg-switch
counter-group cg-protocol-all monitor-level default action-group
ag-defaultentity-group name eg-fcport-edgecounter
credit-loss-recovery monitor-level default action-group
ag-error-disableentity-group name eg-fcport-allcounter-group
cg-fcport-all monitor-level default action-group ag-default
```

SMA custom policy

```
switch# show sma policy active
sma entity-group name eg-interface-rangeentity-type fc-port
interface fc1/1-10,fc1/13-20sma policy name customentity-group name
eg-switchcounter-group cg-protocol-all monitor-level default
action-group ag-defaultentity-group name eg-interface-rangecounter
tx-wait monitor-level default action-group
ag-fpin-congestioncounter-group cg-link-integrity monitor-level
default action-group ag-fpin-link-integrityentity-group name
eg-fcport-edgecounter credit-loss-recovery monitor-level default
action-group ag-error-disableentity-group name
eg-fcport-allcounter-group cg-fcport-all monitor-level default
action-group ag-default
```

Viewing the Active Policy

Displays the active policy.

```
switch(config)# show sma policy active
```

View a policy with a given name

Displays the policy with the given name.

```
switch(config)# show sma policy name
<policy-name>
```

Miscellaneous examples

Other examples are described.

View the SMA History

Displays the SMA history. 1MB of memory is reserved for history and contains an average of 11,000 entries of threshold events.

```
switch# show sma history
Action Legend:
  g: FPIN congestion             n: FPIN congestion clear
  p: FPIN link integrity         ss: Congestion signals stop
  c: Congestion isolate          r: Congestion isolate recovery
  d: DIRL                        l: DIRL recovery
  e: error disable               f: flap
  as: Alarm congestion signal    ws: Warning congestion signal

Alert Legend:
  s[0-7]: syslog, severity [0-7]
  t[0-7]: SNMP trap, level [0-7]
  o: OBFL
  c: callhome

Entity type: fc-port
-------------------------------------------------------------------------------------------------------
Counter                     | Entity     | Value               | Level    | Action | Alerts
         | Timestamp
-------------------------------------------------------------------------------------------------------
tx-credit-not-available     | fc1/1      |                   0 | falling  | none   | s5,t5
          | 2025 Mar 17 09:05:00.779 UTC
tx-credit-not-available     | fc1/1      |                 100 | alarm    | none   | s4,t4
          | 2025 Mar 17 09:04:59.775 UTC
```



Configuring Trunking

This chapter provides information about trunking and how to configure the trunking.

- Finding Feature Information, on page 318
- Information About Trunking, on page 319
- Guidelines and Limitations, on page 325
- Default Settings, on page 329
- Configuring Trunking, on page 330
- Verifying Trunking Configuration, on page 333
- Configuration Example for F Port Trunking, on page 335

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://bst.cloudapps.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.

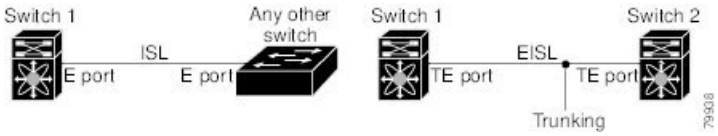
Information About Trunking

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 9: Trunking E Ports, on page 319 and Figure 10: Trunking F Ports, on page 319).

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 9: Trunking E Ports



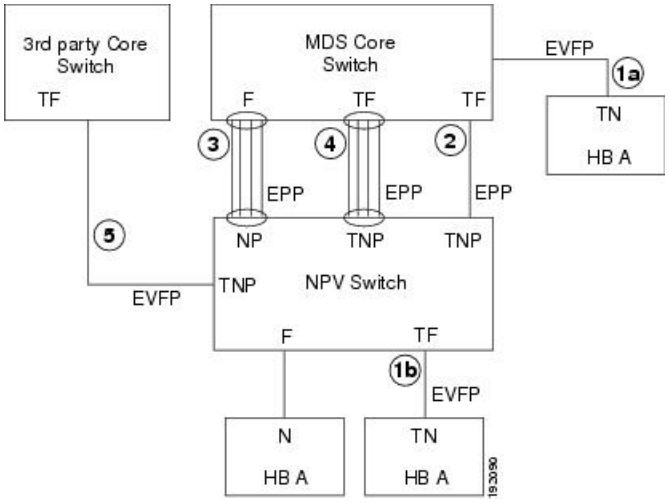
Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c_Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

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 10: Trunking F Ports, on page 319 represents the possible trunking scenarios in a SAN with MDS core switches, NPV switches, third-party core switches, and HBAs.

Figure 10: Trunking F Ports



Link Number	Link Description
1a and 1b	F port trunk with N port.[9]
2	F port trunk with NP port.
3	F port channel with NP port.
4	Trunked F port channel with NP port.
5	Trunking NP port with third-party core switch F port. 1

[9] 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 10: Trunking F Ports, on page 319) 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 10: Trunking F Ports, on page 319) 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 10: Trunking F Ports, on page 319) and that port becomes operational as a trunking NP port, it is referred to as a TNP port.
- TF port channel—If trunk mode is enabled in an F port channel (see the link 4 in Figure 10: Trunking F Ports, on page 319) and that port channel becomes operational as a trunking F port channel, it is referred to as TF port channel. 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 10: Trunking F Ports, on page 319) 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 10: Trunking F Ports, on page 319). PTP is used because PTP also supports trunking port channels.

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 FCIDs 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 38: Supported Trunking Protocols, on page 321 specifies the protocols used for trunking and channeling.

Table 38: Supported Trunking Protocols

Trunk Link	Default
TE-TE port link	Cisco EPP (PTP)
TF-TN port link[10]	FC-LS Rev 1.62 EVFP
TF-TNP port link	Cisco EPP (PTP)
E or F port channel	Cisco EPP (PCP)
TF port channel	Cisco EPP (PTP and PCP)
Third-party TF-TNP port link [11]	FC-LS Rev 1.62 EVFP

[10] These features are not currently supported.
[11] 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 39: Trunk Mode Status Between Switches , on page 322).

Table 39: 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

Note 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 11: Default Allowed-Active VSAN Configuration, on page 323) 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 11: Default Allowed-Active VSAN Configuration, on page 323).

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.

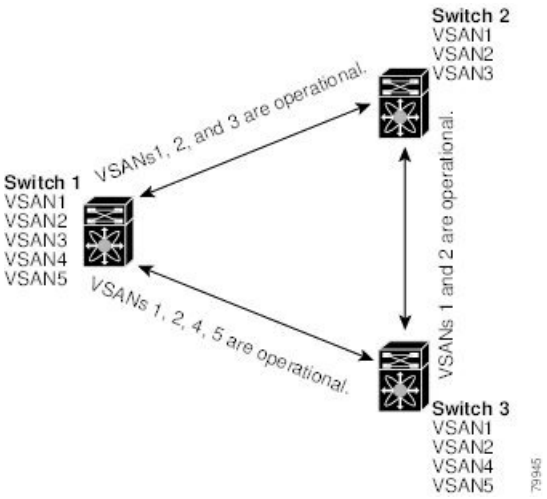
If F port trunking and channeling is enabled, or if **switchport trunk mode on** is configured in NPV mode for any interface, or if NP port channel is configured, the VSAN and VF-ID ranges available for the configuration are as described in Table 40: VSAN and VF-ID Reservations, on page 323.

Table 40: 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 11: 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.

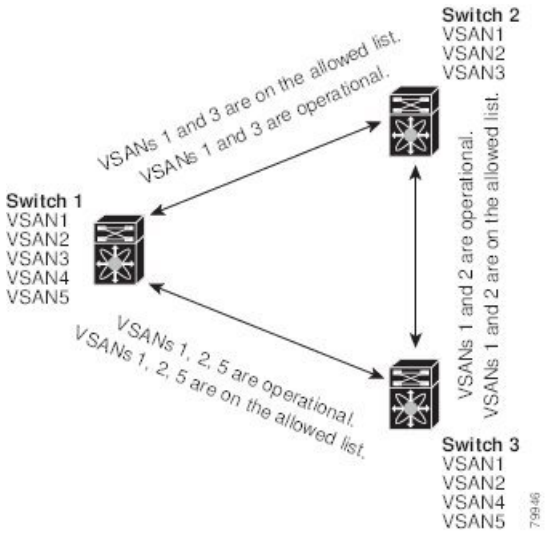
Using Figure 11: Default Allowed-Active VSAN Configuration, on page 323 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 12: Operational and Allowed VSAN Configuration, on page 324). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of 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 12: 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.
- You cannot enable the fport-channel-trunk feature in the VSAN range from 3840 to 4079. The following error is displayed `fport-channel-trunk cannot be enabled because vsans in the range 3840 to 4079 are in use (err_id 0x402A0023).`
- 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 port channels 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)
```

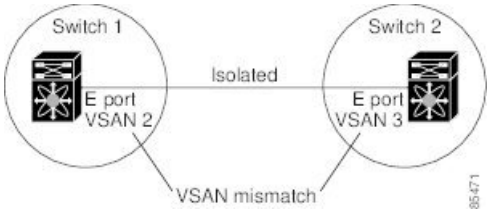
This example shows the TF port information when a server logs in on noninternal FLOGI VSAN. VSAN 2183 is moved to the up state when the server logs in to VSAN 2183.

```
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 13: VSAN Mismatch, on page 328).

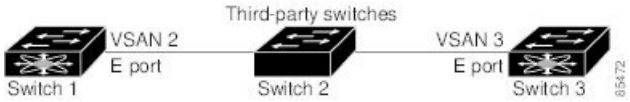
Figure 13: VSAN Mismatch



The trunking protocol detects potential VSAN merging and isolates the ports involved (see Figure 13: VSAN Mismatch, on page 328).

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 14: Third-Party Switch VSAN Mismatch, on page 328).

Figure 14: 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 41: Default Trunk Configuration Parameters , on page 329 lists the default settings for trunking parameters.

Table 41: 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.

Procedure

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.

Procedure

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:

Procedure

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:

Procedure

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 9000 Series 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 333 to Displays Per VSAN Information on Trunk Ports, on page 334.

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:

Procedure

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 Port Channels

This chapter provides information about port channels and how to configure the port channels.

- Feature History for Port Channels, on page 338
- Information About Port Channels, on page 339
- Prerequisites for Configuring Port Channels, on page 347
- Default Settings, on page 348
- Guidelines and Limitations, on page 349
- Configuring Port Channels, on page 360
- Configuring Port Channel Mode, on page 361
- Deleting Port Channels, on page 362
- Adding an Interface to a Port Channel, on page 363
- Adding a Range of Ports to a Port Channel, on page 364
- Adding an Interface using force command, on page 365
- Removing an Interface from a Port Channel, on page 366
- Verifying Port Channel Configuration, on page 367
- Configuration Examples for F and TF Port Channels, on page 371

Feature History for Port Channels

Feature Name	Release	Feature Information
Port channels	8.4(1)	The default port channel mode is changed from On to Active mode.

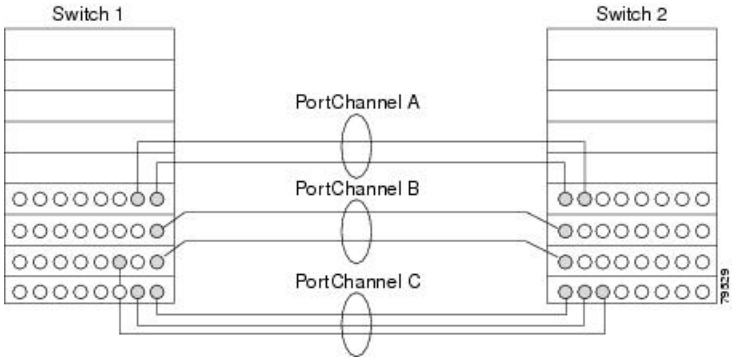
Information About Port Channels

The following sections provide information about Port Channels.

Port Channels Overview

Port channels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy (as shown in the following image). Port channels connect to interfaces across switching modules, so failure of a switching module cannot bring down the port channel link.

Figure 15: Port Channel Flexibility



Port channels on Cisco MDS 9000 Series Multilayer Switches allow flexibility in configuration. This illustrates three possible port channel configurations:

- Port channel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- Port channel B aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- Port channel 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 Port Channels

An E port channel is the aggregation of multiple E ports into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. Port channels connect interfaces across switching modules, so a failure of a switching module cannot bring down the port channel link.

Features and restrictions of a port channel:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). You can combine multiple links into a port channel.
- 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. There are two types of load balancing
  - Flow-based load balancing which is based on the source ID and destination ID.
  - Exchange-based load balancing which is based on source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic that is previously carried on this link changes over to the remaining links. If a link goes down in a port channel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. Link failure of a port channel member does not affect the routing tables and fspf cost of the port channel. Port channels may contain up to 16 physical links and may span multiple modules for added high availability.

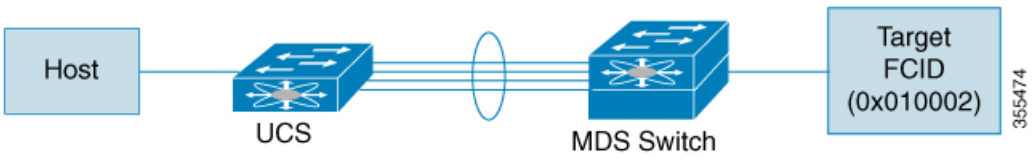
Note See *Cisco MDS 9000 Series NX-OS Fabric Configuration Guide* for information about Failover scenarios for port channels and FSPF links.

F, TF, NP, and TNP Port Channels

Note It is not recommended that you use interface, fWWN, or domain-ID based zoning for devices that are connected to the edge Cisco N-Port Virtualization (NPV) switches.

F port channels provide fault tolerance and performance benefits on connections to N-Port Virtualization (NPV) switches, including Cisco UCS Fabric Interconnects (FIs). F port channels present unique challenges to ACL TCAM programming. When F ports are aggregated into a port channel, ACL TCAM programming is repeated on each member interface. Consequently, these types of port channels multiply the amount of TCAM entries needed. Because of this, it is imperative that the member interfaces are allocated as optimally as possible, and that zoning best practices are also followed. Given that F port channels can contain 100+ host logins, TCAM can easily be exceeded, especially for fabric switches if best practices are not followed.

The following is a sample topology:



This example assumes that the port channel (PC) contains 8 interfaces, fc1/1-fc1/8.

In addition, the following two zones are active:

```
zone1
member host (host 0x010001)
member target1 (target1 0x010002)
zone2
member host (host 0x010001)
```

```
member target2 (target2 0x010003)
```

In such a scenario, the following ACL programming will be present on each member of the PC:

```
fc1/1(through fc1/8) (port-channel)
Entry#     Source ID      Mask          Destination ID          Mask              Action
1          010001         ffffff        010002(target1)         ffffff            Permit
2          010001         ffffff        010003(target2)         ffffff            Permit
3          000000         000000        000000                  000000            Drop
```

The above example shows the ACL TCAM programming that will be duplicated on each member of the F port-channel.

The following are the best practices for efficient use of TCAM with respect to F ports and F port-channels to optimize TCAM usage on a forwarding engine:

- Distribute port-channel member interfaces into different forwarding engines, especially on fabric switches.
- If TCAM usage is still too high in the case of port-channel with a large number of interfaces, then split the port-channel into two separate port-channels each with half the interfaces. This provides redundancy but reduces the number of FLOGIs per individual port-channel and thus reduces TCAM usage.
- Distribute member interfaces into separate linecards on director-class switches.
- Distribute member interfaces into forwarding engines with lower TCAM zoning region usage.
- Use single-initiator zones, single-target zones, or Smart Zoning.

Port Channels 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 port channels as follows:

- Port channels enable 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 non-trunking interface (see Figure 16: Trunking Only, on page 341 and Figure 17: Port Channeling and Trunking, on page 341). See Configuring Trunking, on page 317 for information on trunked interfaces.

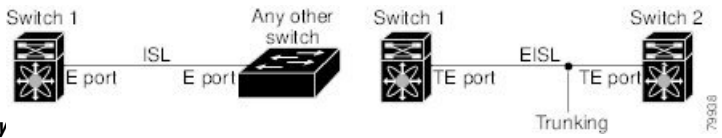
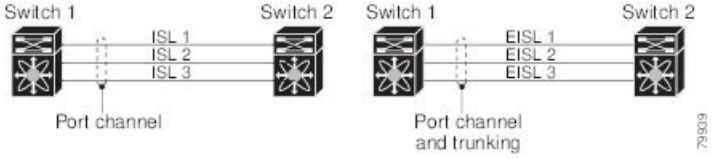


Figure 16: Trunking Only

Figure 17: Port Channeling and Trunking



Port channels and trunking are used separately across an ISL.

- Port channels—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 *Cisco MDS 9000 Series NX-OS Fabric Configuration Guide*.
- Both port channeling and trunking can be used between TE ports over EISLs.

Port Channel Modes

Note After changing the port-channel mode, each member interface must be brought down and brought back up via the **shutdown** and **no shutdown** commands for the port-channel mode to be changed. This can be done on an individual member-by-member basis such that the port-channel remains up and fully functional.

You can configure each port channel with a channel group mode parameter. Such configuration determines the port channel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- On—When the port channel created in **ON** mode, user has to run **no shutdown** command for each port channel member. . Port channels that are configured in the On mode require you to explicitly enable and disable the port channel member ports at either end if you add or remove ports from the port channel configuration. Physically verify that the local and remote ports are connected to each other.

  Beginning with Cisco MDS Release 8.4(1), the default mode is changed from On to Active.
- Active—When the port channel created in ACTIVE mode, members will be operational automatically provided the ISL(s) is operationally up. The Active port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

Note From Cisco MDS NX-OS Release 8.4(1), the CLI and the Device Manager create the port channel in Active mode on the NPIV core switches.

The following table provides a comparison between On and Active modes.

Table 42: Channel Group Configuration Differences

On Mode	Active Mode
No protocol is exchanged.	A port channel protocol negotiation is performed with the peer ports.

On Mode	Active Mode
Moves interfaces to the suspended state if its operational values are incompatible with the port channel.	Moves interfaces to the isolated state if its operational values are incompatible with the port channel.
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 port channel protocol.
When you use **channel-group***port-channel-id***force** command, the misconfigured ports are transitioned to suspended state.	When you use **channel-group***port-channel-id***force** command, the misconfigured ports are transitioned to isolated state.

Deleting Port Channels

When you delete port channels, their corresponding channel membership is also deleted. All interfaces in the deleted port channel convert to individual physical links. The ports going down gracefully indicates that no frames were lost when the interface went down Graceful Shutdown, on page 35).

The individual ports within the deleted port channel retain compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, port security and so on). You can modify these settings as required. When you delete port channel in On mode, run **no shutdown** command on each member to make it operational.

Interfaces in a Port Channel

You can add or remove a physical interface (or range of interfaces) to an existing port channel. For information about port channel support on Generation 2 switching modules, see Port Channel Limitations, on page 124.

Adding Interfaces to a Port Channel

You can add a physical interface (or range of interfaces) to an existing port channel. The compatible parameters on the configuration are mapped to the port channel. Adding an interface to a port channel increases the channel buffer size and bandwidth of the port channel.

You can add a port as a member of the port channel if the following configurations are same in the port and the port channel:

- Speed
- Mode
- Rate mode
- Port VSAN
- Trunking mode
- Allowed VSAN list or VF-ID list and so on.

After the members are added, regardless of the port channel mode (Active or On), the ports at either end are gracefully down. The ports shut down gracefully indicating that no frames were lost when the interface shut down.

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 port channel. The compatibility check is performed before a port is added to the port channel.

The check ensures that the following parameters and settings match at both ends of a port channel:

- 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, port security and so on).
- Operational parameters (remote switch WWN and trunking mode).

Adding a port fails if the capability and administrative parameters on the remote switch are incompatible with the capability and administrative parameters on 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 transitioned to a suspended or isolated state based on the configured mode:

- An interface is moved to the suspended state if the interface is configured in the **On** mode.
- An interface is moved to the isolated state if the interface is configured in the **Active** mode.

Force Add an Interface

You can force the port configuration to be overwritten by the port channel configuration. In this case, the interface is added to a port channel.

- When you use **channel-group***port-channel-id***force** command, the misconfigured ports in **On** mode port channel are transitioned to suspended state.
- When you use **channel-group***port-channel-id***force** command, the misconfigured ports in **Active** mode port channel are transitioned to isolated state.

After the members are forcefully added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down. The ports going down gracefully indicates that no frames were lost when the interface went down (see the Graceful Shutdown, on page 35) section.

Deleting an Interface from a Port Channel

When a physical interface is deleted from the port channel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the port channel status is changed to a down state. Deleting an interface from a port channel decreases the channel buffer size and bandwidth of the port channel.

Port Channel Protocols

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 configuration parameters. Any change in configuration applied to the associated port channel interface is propagated to all members of the channel group.

A protocol to exchange port channel configurations is available in all Cisco MDS switches. This simplifies port channel management with incompatible ISLs.

The port channel protocol is enabled by default.

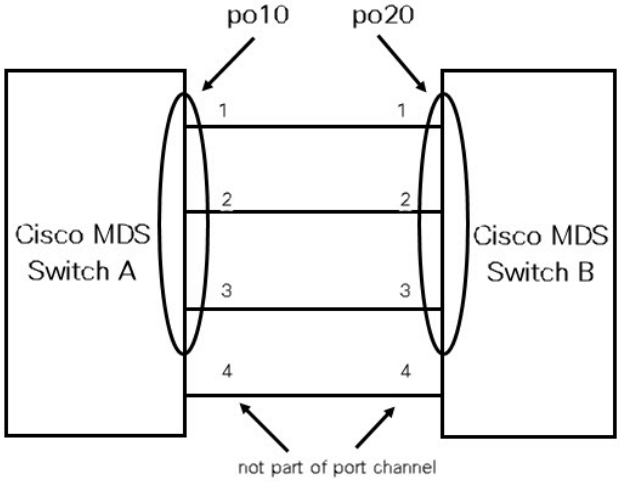
It uses exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each port uses the information that is received from the peer ports along with its local configuration and operational values to decide if it has to be part of a port channel. The protocol ensures that a set of ports is eligible to be part of the same port channel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.

The port channel protocol uses Bring-up protocol which helps in Automatically detecting the misconfigurations. This protocol synchronizes the port channel at both ends so that all frames for a given flow are carried over the physical link in both directions. This allows applications such as write acceleration, work for port channels over FCIP links.

Creating Port Channels

Assuming the link A1-B1 is operational in the port channel first, that link is considered as First Operational Port (FoP) of the port channel. When the next link is operational, for example, A2-B2, the port channel protocol identifies if this link is compatible with link A1-B1. and adds it to the port channel 10 and 20 in the respective switches. If link A3-B3 can join the the port channels, 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 18: Autocreating Channel Groups



Port channels are created manually. You can form the port channel 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. Any administrative configuration made to the port channel is applied to all ports in the port channel.

Prerequisites for Configuring Port Channels

Before configuring a port channel, consider the following guidelines:

- Configure the port channel across switching modules to implement redundancy on switching module reboots or upgrades.
- Ensure that you do not connect a port channel to different sets of switches. Port channels require point-to-point connections between the same set of switches.

If you misconfigure port channels, you get a misconfiguration message. If you receive this message, the port channel's physical links are disabled because error was detected.

Port channel error is detected if the following requirements are not met:

- Connect each switch on either side of a port channel to the same number of interfaces.
- Connect each interface to a corresponding interface on the other side (see Figure 20: Misconfigured Configurations, on page 350 for an example of an invalid configuration).
- You cannot change the physical links in a port channel after you configure the port channel. If you change the links after you configure the port channel, ensure to reconnect the links to interfaces within the port channel and reenable the links.

If all three conditions are not met, the faulty link is disabled.

Default Settings

The following table describes the default settings for port channels.

Table 43: Default Port Channel Parameters

Parameters	Default
Port channels	FSPF is enabled by default.
Create port channel	Administratively up.
Default port channel mode	Cisco MDS NX-OS Release 8.3(1) and earlier: On mode on non-NPV and NPIV core switches. Cisco MDS NX-OS Release 8.4(1) and later: Active mode on non-NPV and NPIV core switches. Active mode on NPV switches.

Guidelines and Limitations

The following sections provide information about guidelines and limitations for Port Channels.

General Guidelines and Limitations

Cisco MDS 9000 Series Multilayer switches support the following number of port channels per switch:

- A port channel number refers to the unique identifier for each channel group. This number ranges from of 1–256.



Note You cannot change the port channel number, but the member ports operate according to the properties of the port channel.

The following table describes the results of adding a member to a port channel for various configurations.

Generation 1 Port Channel Limitations

This section includes the restrictions on creation and addition of port channel members to a port channel on Generation 1 hardware:

- The 32-port 2 or 1 Gbps switching module.
- The MDS 9140 and 9120 switches.

When configuring the host-optimized ports on Generation 1 hardware, the following port channel 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 ports in the Cisco MDS 9100 Series can be included in a port channel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same port channel rules as 32-port switching modules. Only the first port of each group of 4 ports is included in a port channel.
  - 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 you configure the first port in the group as a port channel, 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 you configure any of the other three ports in a no shutdown state, you cannot configure the first port to be a port channel. The other three ports continue to remain in a no shutdown state.

Valid and Invalid Port Channel Examples

Port channels are created with default values. You can change the default configuration just like any other physical interface.

Figure 19: Valid Port Channel Configurations, on page 350 provides examples of valid port channel configurations.

Figure 19: Valid Port Channel Configurations

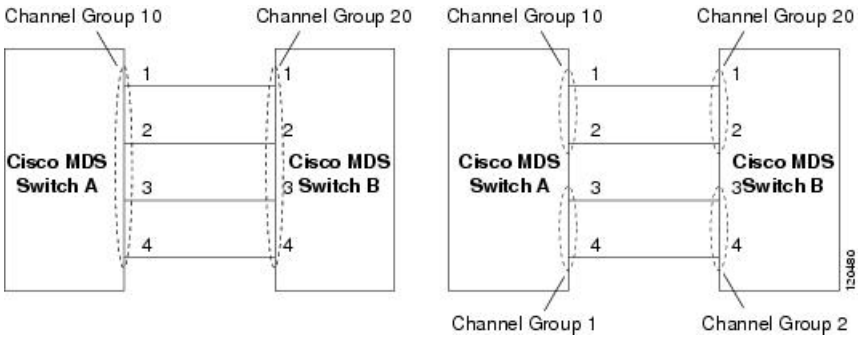
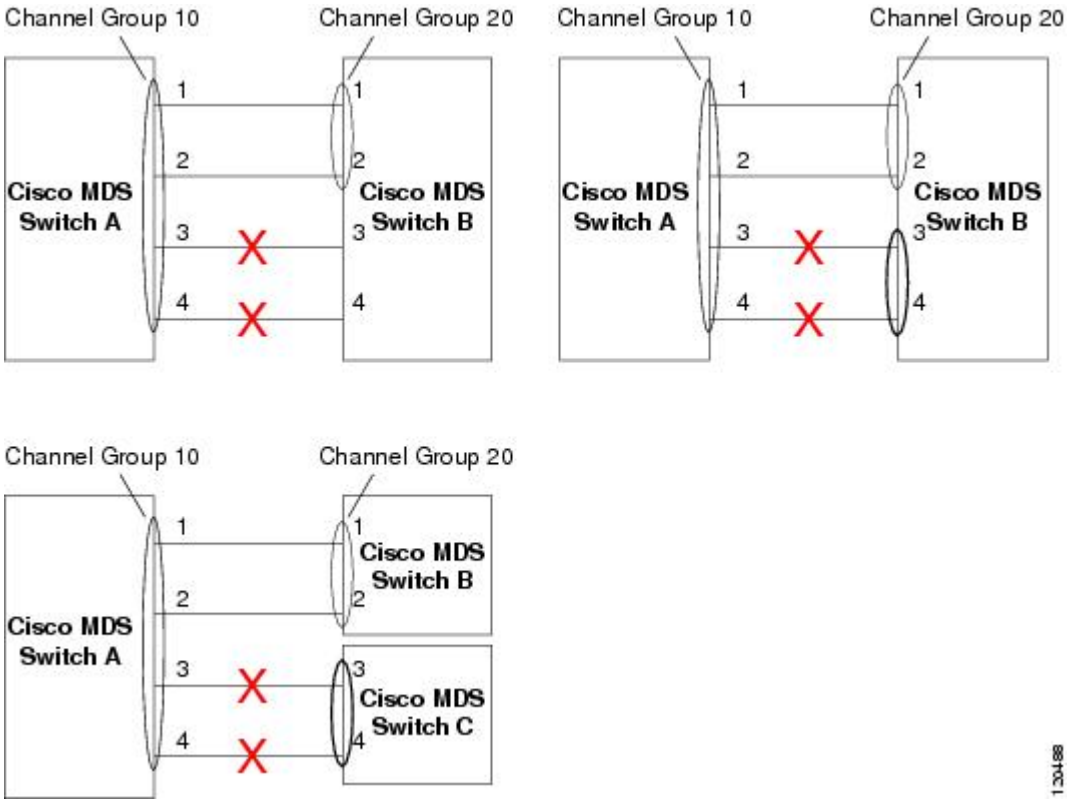


Figure 20: Misconfigured Configurations, on page 350 provides examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down because the fabric is misconfigured.

Figure 20: Misconfigured Configurations



E Port Channel Limitations

The port channel interface must be in Active mode when you configure multiple FCIP interfaces with Write-Acceleration.

F, TF, and NP Port Channel Limitations

The following guidelines and restrictions are applicable for F, TF, and NP port channels:

- On the switch with **feature npiv** configured the ports must be in F mode.
- On the switch with **feature npv** configured the ports must be in NP mode.
- On mode is not supported. Only Active-Active mode is supported. By default, the mode is Active on the NPV switches.
- Devices that are logged in through an F port channel 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 port channel member.
- Since the FLOGI payload carries only the VF bits to trigger the use of a protocol after the FLOGI exchange, those bits are overridden. If the Cisco NPV switches, the core has a Cisco WWN and tries to initiate the PCP protocol.
- The name server registration of the N ports logging in through an F port channel uses the fWWN of the port channel interface.
- DPVM configuration is not supported.
- You cannot configure the port channel port VSAN using DPVM.
- The Dynamic Port VSAN Management (DPVM) database is 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 is queried only for the physical FLOGI.

Best Practises for Forwarding Engines

Cisco MDS switches use Ternary Content Addressable Memory (TCAM) on its Fibre Channel modules. TCAM provides an Access Control List (ACL) type of function for Cisco MDS. The process that controls this functionality is called ACLTCAM. The E or TE ports (ISLs) and F (Fabric) ports have their own programming that is unique to their respective port types.

TCAM is allocated to individual forwarding engines and forwarding engines are assigned a group of ports. Director-class Fibre Channel modules have more TCAM space than fabric switches. The number of forwarding engines, the ports assigned to each forwarding engine, and the amount of TCAM allocated to each forwarding engine is hardware dependent.

The following example shows an output from Cisco MDS 9148S:

```
switch# show system internal acl tcam-soc tcam-usage
```

```
TCAM Entries:
=============
                    Region1   Region2    Region3     Region4   Region5   Region6
Mod Fwd   Dir      TOP SYS  SECURITY    ZONING      BOTTOM    FCC DIS   FCC ENA
    Eng           Use/Total Use/Total  Use/Total   Use/Total Use/Total Use/Total
--- ---  ------ ---------- --------- ------------ --------- --------- ---------
1   1    INPUT     19/407     1/407      1/2852 *    4/407     0/0       0/0
1   1    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
1   2    INPUT     19/407     1/407      0/2852 *    4/407     0/0       0/0
1   2    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
1   3    INPUT     19/407     1/407      0/2852 *    4/407     0/0       0/0
1   3    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
-------------------------------------------------
* 1024 entries are reserved for LUN Zoning purpose.
```

The above example indicates the following:

- There are three forwarding engines, 1 through 3.
- Since there are 48 ports on Cisco MDS 9148 switches, each forwarding engine handles 16 ports.
- Each forwarding engine has 2852 entries in region 3 (the zoning region) for input. This is the main region used, and consequently, has the largest amount of available entries.
- Forwarding engine 3 has only one entry that is currently in use out of the total 2852 in the zoning region.

The following example shows the output from Cisco MDS 9710 switch with a 2/4/8/10/16 Gbps Advanced Fibre Channel Module (DS–X9448–768K9):

```
F241-15-09-9710-2# show system internal acl tcam-usage
TCAM Entries:
=============
                    Region1   Region2    Region3     Region4   Region5   Region6
Mod Fwd   Dir      TOP SYS  SECURITY    ZONING      BOTTOM    FCC DIS   FCC ENA
    Eng           Use/Total Use/Total  Use/Total   Use/Total Use/Total Use/Total
--- ---  ------ ---------- --------- ------------ --------- --------- ---------
1   0    INPUT     55/19664    0/9840     0/49136*   17/19664    0/0       0/0
1   0    OUTPUT    13/4075    0/1643     0/11467     0/4075    6/1649   21/1664
1   1    INPUT     52/19664    0/9840     2/49136*   14/19664    0/0       0/0
1   1    OUTPUT     7/4078    0/1646     0/11470     0/4078    6/1652    5/1651
1   2    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   2    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   3    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   3    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   4    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   4    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   5    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   5    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
...
```

The above example indicates the following:

- There are six forwarding engines, 0 through 5.
- Since there are 48 ports on a Cisco MDS DS–X9448–768K9 module, each forwarding engine handles eight ports.
- Each forwarding engine has 49136 entries in region 3 (the zoning region) for input. This is the main region that is used, and consequently, has the largest amount of available entries.

- Forwarding engine 2 has only two entries that are currently in use out of the total 49136 in the zoning region.

The following example shows the output from Cisco MDS 9396V switch with a 2/4/8/10/16/32/64 Gbps Advanced Fibre Channel Module (DS–X9448–768K9):

```
switch9396v# show system internal acl tcam-usage
Input TCAM Entries:
=====================
                    Region1   Region2       Region3         Region4
Mod Fwd   Dir      TOP SYS  SECURITY       ZONING          BOTTOM
    Eng          Use/Total Use/Total   Use/Total(Anl)  Use/Total(Anl)
--- ---  ------ ---------- ---------- ---------------- --------------
1   0    INPUT    126/26208    0/13120     0/65536(0)    28/26208(0)
1   1    INPUT    122/26208    0/13120     2/65536(0)    27/26208(0)
1   2    INPUT    150/26208    0/13120     0/65536(0)    32/26208(0)
1   3    INPUT    126/26208    0/13120     0/65536(0)    28/26208(0)

Output TCAM Entries:
======================
    Fwd           Region1   Region2       Region3         Region4      Region5   Region6
Mod Eng/  Dir     TOP SYS  SECURITY       ZONING          BOTTOM       FCC DIS   FCC ENA
    Port         Use/Total Use/Total   Use/Total(Anl)  Use/Total(Anl) Use/Total Use/Total
    Num
--- ---  ------ ---------- ---------- ---------------- -------------- --------- ---------
1   0    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      3/51
1   1    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
1   2    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
1   3    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
1   4    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
.
..
...
1   94    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
1   95    OUTPUT     4/51       0/51       0/281  (0)     0/51  (0)       4/25      1/51
Note: Analytics Entry Count(Anl) included in Use count.
```

The above example indicates the following:

- There are four forwarding engines, 0 through 3.
- Since there are 96 ports on a Cisco MDS DS-C9396V-K9-SUP module, each forwarding engine handles twenty-four ports.
- Each forwarding engine has 65536 entries in region 3 (the zoning region) for input. This is the main region that is used, and consequently, has the largest amount of available entries.
- Forwarding engine 2 has only two entries that are currently in use out of the total 65536 in the zoning region.



Note The commands that are used to view TCAM usage on fabric switches are different from the ones used for director–class switches. For MDS 9148, MDS 9148S, and MDS 9250i fabric switches, use the **show system internal acltcam-soc tcam-usage** command. For director class switches, MDS 9396V, MDS 9396S, and 32 Gbps fabric switches, use the **show system internal acl tcam-usage** command.

The following table provides information about ports to forwarding engines mapping:

Table 44: Ports to Forwarding Engines Mapping

Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
MDS 9132T	2	1–16	0	49136	19664
		17–32	1	49136	19664
MDS 9148	3	fc1/25–36 and fc1/45–48	1	2852	407
		fc1/5–12 and fc1/37–44	2	2852	407
		fc1–4 and fc1/13–24	3	2852	407
MDS 9148S	3	fc1/1–16	1	2852	407
		fc1/17–32	2	2852	407
		fc1/33–48	3	2852	407
MDS 9148T	3	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
MDS 9250i	4	fc1/5–12 and eth1/1–8	1	2852	407
		fc1/1–4, fc1/13–20, and fc1/37–40	2	2852	407
		fc1/21–36	3	2852	407
		ips1/1–2	4	2852	407

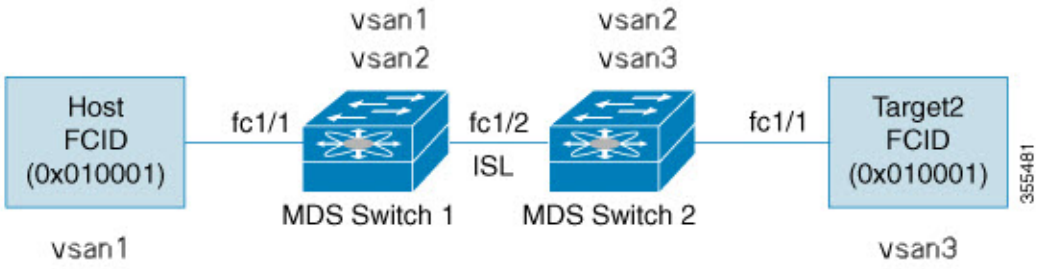
Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
MDS 9396S	12	fc1/1–8	0	49136	19664
		fc1/9–16	1	49136	19664
		fc1/17–24	2	49136	19664
		fc1/25–32	3	49136	19664
		fc1/33–40	4	49136	19664
		fc1/41–48	5	49136	19664
		fc1/49–56	6	49136	19664
		fc1/57–64	7	49136	19664
		fc1/65–72	8	49136	19664
		fc1/73–80	9	49136	19664
		fc1/81–88	10	49136	19664
		fc1/89–96	11	49136	19664
MDS 9396T	6	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
		49–64	3	49136	19664
		65-80	4	49136	19664
		81-96	5	49136	19664
DS–X9248–48K9	1	1–48	0	27168	2680
DS–X9248–96K9	2	1–24	0	27168	2680
		25–48	1	27168	2680
DS–X9224–96K9	2	1–12	0	27168	2680
		13–24	1	27168	2680
DS–X9232–256K9	4	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
		25–32	3	49136	19664

Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
DS–X9248–256K9	4	1–12	0	49136	19664
		13–24	1	49136	19664
		25–36	2	49136	19664
		37–48	3	49136	19664
DS–X9448–768K9	6	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
		25–32	3	49136	19664
		33–40	4	49136	19664
		41–48	5	49136	19664
DS–X9334–K9	3	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
DS–X9648–1536K9	3	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
DS-C9124V-K9	1	1-24	0	65536	26208
DS-C9148V-24EK9	2	1-24	0	65536	26208
		25-48	1	65536	26208
DS-C9220I-K9	1	1-12	0	49136	19664
DS-X9748-3072-K9	2	1-24	0	65536	26208
		25-48	1	65536	26208
DS-C9396V-K9	4	1-24	0	65536	26208
		25-48	1	65536	26208
		49-72	2	65536	26208
		73-96	3	65536	26208

Best Practises for E and TE Port Channels and IVR

Port channels provide Inter Switch Links (ISLs) between switches. Typically, there is minimal TCAM programming on these types of interfaces. When the Inter VSAN Routing(IVR) feature is being deployed, extensive TCAM programming can exist on ISLs because the IVR topology transitions from one VSAN to another. Most of the considerations that apply on F/TF port channels will be applicable here too.

The following is an example of a topology:



In this topology:

- Both Cisco MDS 9148S-1 and MDS 9148S-2 are in the IVR VSAN topology:

```
MDS9148S-1 vsan 1 and vsan 2
MDS9148S-2 vsan 2 and vsan 3
```

- IVR NAT is configured.
- VSAN 2 is the transit VSAN.

```
FCIDs per VSAN:
            VSAN 1  VSAN 2  VSAN 3
Host         010001  210001  550002
Target1     440002  360002  030001
```

Note Domains 0x44 in VSAN 1, 0x21 and 0x36 in VSAN 2, and 0x55 in VSAN 3 are virtual domains created by IVR NAT.

- The following is the IVR zoning topology:

```
ivr zone zone1
member host vsan 1
member target1 vsan3
```

- The following is the ACL TCAM programming for the IVR zoning topology:

```
MDS9148S-1  fc1/1(Host) - VSAN 1
Entry#     Source ID        Mask      Destination ID          Mask    Action
1          010001(host)     ffffff    440002(target1)         ffffff  Permit
           - Forward to fc1/2
```

```
       - Rewrite the following information:
         VSAN to 2
         Source ID to 210001
         Destination ID to 360002
2         000000         000000    000000                  000000  Drop
MDS9148S-1  fc1/2(ISL) - VSAN 2
Entry#    Source ID           Mask      Destination ID          Mask    Action
1         360002(Target1)     ffffff    210001(host)            ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
         VSAN to 1
         Source ID to 440002
         Destination ID to 010001
MDS9148S-2 fc1/2(ISL) - VSAN 2
Entry#    Source ID        Mask      Destination ID          Mask    Action
1         210001(host)     ffffff    360002(target1)         ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
         VSAN to 3
         Source ID to 550002
         Destination ID to 030001
MDS9148S-2  fc1/1(Target1) - VSAN 3
Entry#    Source ID           Mask      Destination ID          Mask    Action
1         030001(Target1)     ffffff    550002(host)            ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
         VSAN to 2
         Source ID to 360002
         Destination ID to 210001
2         000000          000000    000000                  000000  Drop
```

Note Besides the entries in this example, there are other entries that IVR adds to capture important frames such as PLOGIs, PRILIs, and ABTS.

The programming on the host and target1 ports is similar to the way it is without IVR, except that the FCIDs and VSANs are explicitly forwarded to an egress port and are rewritten to values that are appropriate for the transit VSAN (VSAN 2). These forwarding and rewrite entries are separate and are not included in the TCAM-usage values.

However,now, on the ISLs in both the switches, programming that did not exist earlier is present. When frames from Host to Target1 are received by Cisco MDS 9148S-2 fc1/2, they are rewritten to the values in VSAN3 where the target resides. In the reverse direction, when frames from Target1 to the Host are received by Cisco MDS 9148S-1 fc1/2, they are rewritten to the values in VSAN 1 where the Host resides. Therefore, for each VSAN transition on an ISL (that typically occurs across a transit VSAN) there is TCAM programming for each device in the IVR zone set.

Consequently,most of the best practices followed for the F and TF port channels should be followed to ensure that TCAM is utilized as efficiently as possible for the following purposes:

Note Unlike F and TF port-channels, the ACLTCAM programming on ISLs will be the same quantity regardless if the ISLs are part of a port-channel or not. If there are "n" ISLs between two MDS switches, then it doesn't matter if they are in one port-channel, two port-channels or just individual links. The ACLTCAM programming will be the same.

- Distribute port-channel member interfaces into different forwarding engines, especially on fabric switches.
- Distribute member interfaces into different linecards on director-class switches.
- Distribute member interfaces into forwarding engines with lower TCAM zoning region usage.
- Use single-initiator zones, single-target zones, or Smart Zoning.

Configuring Port Channels

To create a port channel, perform these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

Configures the specified port channel (1) using the default ON mode.

Configuring Port Channel Mode

From Cisco MDS NX-OS Release 8.4(1), the CLI and the Device Manager create the port channel in Active mode on the NPIV core switches.

Note An F port channel is supported only on Active mode.

To configure Active mode, perform these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface port-channel 1**

Configures the specified port channel (1) using the default Active mode.

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

Configures the specified port channel in default **On** mode.

Deleting Port Channels

To delete a port channel, perform these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no interface port-channel 1**

Deletes the specified port channel (1), its associated interface mappings, and the hardware associations for this port channel.

Adding an Interface to a Port Channel

To add an interface to a port channel, perform these steps:

Procedure

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 the channel-group 15 does not exist, it is created. The port is shut down if the port channel in the **On** mode.

Adding a Range of Ports to a Port Channel

To add range of ports to a port channel, perform these steps:

Procedure

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 the 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 an interface to a port channel, while NDFC SAN Controller adds the interface by force, unless specified explicitly.

Adding an Interface using force command

To force the addition of a port to a port channel, perform these steps:

Procedure

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 in port channel **ON** mode For port channel in Active mode, the interfaces recovers automatically (for consistency with earlier instances)

Step 4 switch(config-if)#**no shutdown**

Execute this command on the port channel in **On** mode to make the ports operational.

Removing an Interface from a Port Channel

To remove a physical interface (or range of physical interfaces) from a port channel, perform these steps:

Procedure

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**

Removes the physical Fibre Channel interfaces from channel group 2.

Verifying Port Channel Configuration

To display the port channel configuration information, perform one of the following tasks:

Command	Purpose
show port-channel summary	Displays a summary of port channels within the switch. A one-line summary of each port channel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP). The FOP is the primary operational interface that is selected in the port channel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a port channel and can change if the port goes down. The FOP is also identified by an asterisk (*).
show port-channel database	Displays the port channel that is configured in the **On** mode (default) 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 port channel usage.
show port-channel compatibility-parameters	Displays the port channel compatibility.
show interface fc *slot/port*	Displays interfaces information.
show port-channel database interface port-channel *number*	Displays the specified port channel interface.
how running-config interface port-channel*number*	Displays the current running configuration for port-channel interface,

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

You can view specific information about existing port channels at any time from EXEC mode. The following **show** commands provide further details on existing port channels. You can force all screen output to go to a printer or save it to a file. See Examples Displays the Port Channel Summary, on page 367 to Displays the Port Channel Summary, on page 367.

Displays the Port Channel 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 Port Channel Configured in the On Mode

```
switch# show port-channel database
port-channel1
     Administrative channel mode is on
     Last membership update succeeded
     First operational port is fc1/19
     4 ports in total, 2 ports up
     Ports:   fc1/19   [up] *
              fc1/20   [up]
              fc1/21   [down]
              fc1/22   [down

port-channel2
     Administrative channel mode is on
     Last membership update succeeded
     First operational port is fcip3
     2 ports in total, 2 ports up
     Ports: fcip1 [up]
                 fcip3 [up] *
```

Displays the Port Channel Configured in the Active Mode

```
switch# show port-channel database

port-channel1
     Administrative channel mode is active
     Last membership update succeeded
     First operational port is fc1/19
     4 ports in total, 2 ports up
     Ports:   fc1/19   [up] *
              fc1/20   [up]
              fc1/21   [down]
              fc1/22   [down

port-channel2
     Administrative channel mode is active
     Last membership update succeeded
     First operational port is fcip3
     2 ports in total, 2 ports up
     Ports: fcip1 [up]
                 fcip3 [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:
============================================
```

```
 total 1 port-channels
 port-channel 100:
     4 ports, first operational port is fc1/19
     fc1/22   [down]
     fc1/21   [down]
     fc1/19   [up]
     fc1/20   [up]
================================================
 database 1: from module 1
================================================
 total 1 port-channels
 port-channel 100:
     4 ports, first operational port is fc1/19
     fc1/19   [up]
     fc1/20   [up]
     fc1/21   [down]
     fc1/22   [down]
================================================
switch#
```

The **show port-channel usage** command displays details of the used and unused port channel numbers.

Displays the Port Channel Usage

```
switch# show port-channel usage
Totally 4 port-channel numbers used
===================================
 Used : -77 -79, 100
 Unused: 1 - 76, 80 - 99, 101 - 256
```

Displays the Port Channel 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 the Specified Port Channel Interface

```
switch# show port-channel database
 interface port-channel 100
 port-channel 100
     Administrative channel mode is active
     Last membership update succeeded
     First operational port is fc1/19
     4 ports in total, 2 ports up
     Ports:   fc1/19   [up] *
              fc1/20   [up]
              fc1/21   [down]
              fc1/22   [down]
```

Displays the Port Channel 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 Port Channels

This example shows how to configure an F port channel in shared mode and bring up the link (not supported on the MDS 91x4 switches) between F ports on the Cisco NPIV core switches and NP ports on the Cisco NPV switches:

Procedure

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 port channel 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)# exit
```

Step 4 Configure the port channel 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 8000
switch(config-if)# channel-group 1
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the port channel on the NPV switch:

Example:

```
switch(config)# interface port-channel 1
switch(config-if)# switchport mode NP
switch(config-if)# exit
```

Step 6 Configure the port channel 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 8000
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 port-channel 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
```



Configuring N Port Virtualization

This chapter provides information about N port virtualization and how to configure N port virtualization.

- Feature History for N Port Virtualization, on page 374
- Information About N Port Virtualization, on page 375
- NPV Traffic Management, on page 381
- Guidelines and Limitations, on page 383
- Configuring N Port Virtualization, on page 386
- Verifying NPV and NPIV Configuration, on page 390

Feature History for N Port Virtualization

This table lists the New and Changed features.

Table 45: New and Changed Features

Feature Name	Release	Feature Information
N Port Virtualization (NPV) Load Balancing	8.5(1)	NPV load balancing scheme is enhanced to propose a mapping of server interfaces to external interfaces based on the throughput value so that the traffic can be evenly distributed on the external interfaces. The following commands were introduced: • **show npv traffic-map proposed** • **npv traffic-map analysis clear**
N Port Identifier Virtualization	8.4(2)	The NPIV feature is enabled by default.
NP Ports	8.4(1)	Buffer-to-Buffer State Change Notification (BBSCN) allowed on NP Ports

Information About N Port Virtualization

N Port Virtualization Overview

Cisco N Port Virtualization (NPV) reduces the number of Fibre Channel domain IDs required in a fabric. Switches operating in the Cisco NPV mode do not join a fabric which eliminates the need for domain IDs for these switches. Such switches function as edge switches and pass traffic between an NPIV core switch and end devices. Cisco NPV switches cannot be standalone switches since they rely on an upstream NPIV enabled switch to provide many fabric services for them.

NPV is supported by the following Cisco MDS 9000 switches only:

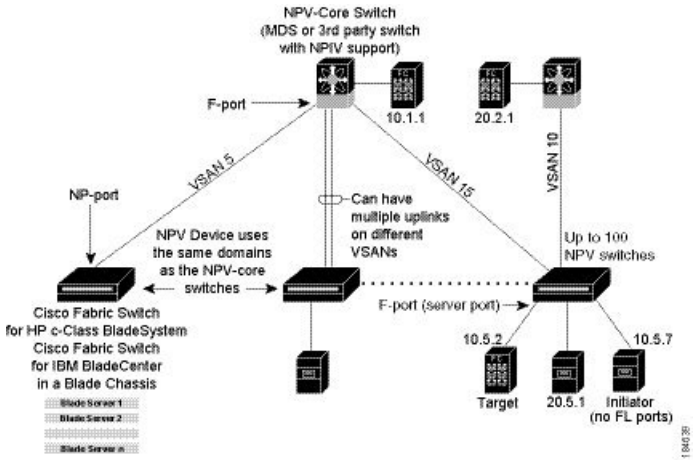
- Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switch
- Cisco MDS 9148T 32-Gbps 48-Port Fibre Channel Switch
- Cisco MDS 9396T 32-Gbps 96-Port Fibre Channel Switch
- Cisco MDS 9148S 16-Gbps Multilayer Fabric Switch
- Cisco MDS 9396T 16-Gbps Multilayer Fabric Switch
- Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch
- Cisco MDS 9148V 64-Gbps 48-Port Fibre Channel Switch
- Cisco MDS 9396V 64-Gbps 96-Port Fibre Channel Switch

Cisco NPV technology is also available on Nexus and UCS Fabric Interconnects. 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 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 many 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 switch 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 core switch to which NPV devices are connected to. NPV also allows multiple devices to attach to same port on the core switch to which NPV devices are connected to, 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.

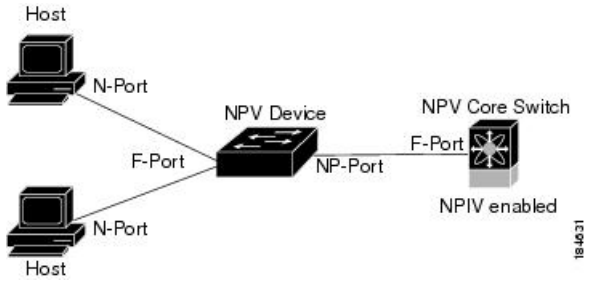
Figure 21: Cisco NPV Fabric Configuration



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 FCIDs. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of the NPIV feature on the core switch to get multiple FCIDs allocated on the NP port.

#unique_326 unique_326_Connect_42_fig_5330B6151ADF421F917437276B6DBE59 shows a more granular view of an NPV configuration at the interface level.

Figure 22: Cisco NPV Configuration-Interface View

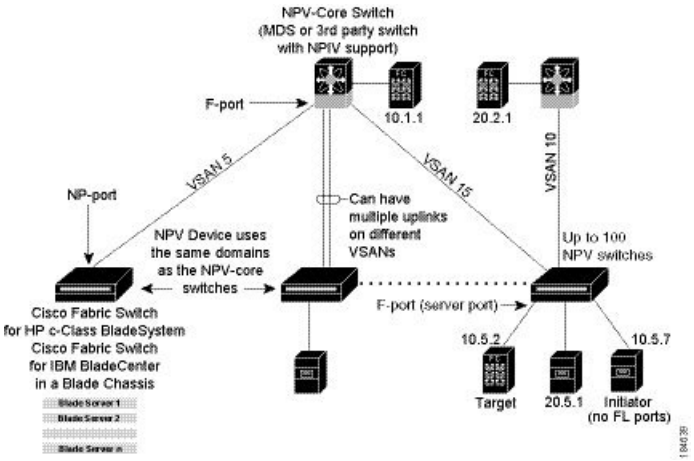


N Port Identifier Virtualization

The N port identifier virtualization (NPIV) feature provides a means to assign multiple FCIDs to a single N port. This feature allows multiple applications on the N port to use different FCIDs and allows access control, zoning, and port security to be implemented at the application level Figure 23: NPIV Example, on page 377 shows an example application using NPIV.

From Cisco MDS NX-OS Release 8.4(2), the NPIV feature is enabled by default.

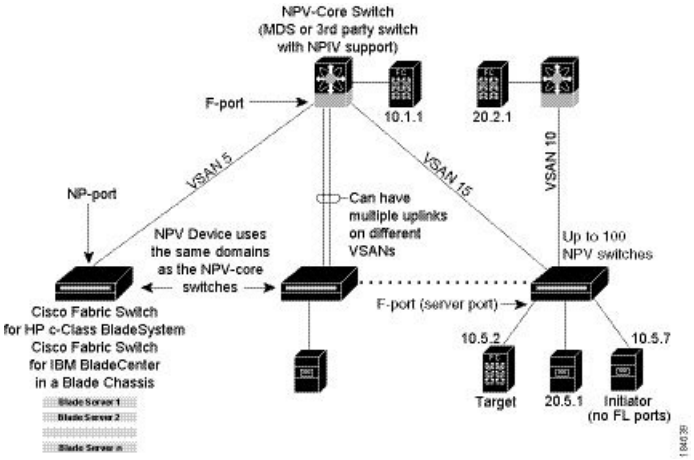
Figure 23: NPIV Example



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.

Figure 24: Cisco NPV Fabric Configuration



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 core switch to which NPV devices are connected to among multiple NPV switches. NPV also allows multiple devices to attach to same port on the core switch to which NPV devices are connected to, 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.

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. Use **feature npv** command to enable NPV. 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 NPIV 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 core switch to which NPV devices are connected to.



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, NPIV switches will enforce in-order delivery if needed and/or configured.

NP Ports

An NP port (proxy N port) is a port on a device that is in NPV mode and connected to the core switch to which NPV devices are connected to 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 core switch to which NPV devices are connected to 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 core switch to which NPV devices are connected to, and then (if the FLOGI is successful) registers itself with the core switch to which NPV devices are connected to 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 378 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 core switch to which NPV devices are connected to 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 25: Internal FLOGI Flows, on page 379 shows the internal FLOGI flows between a core switch to which NPV devices are connected to and an NPV device.

Figure 25: Internal FLOGI Flows

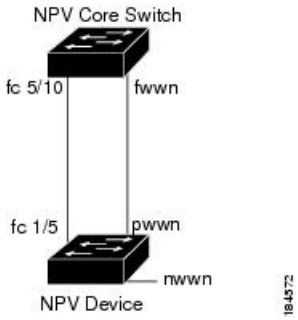


Table 46: Internal FLOGI Parameters , on page 379 identifies the internal FLOGI parameters that appear in .

Table 46: 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 core switch to which NPV devices are connected to.
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 core switch to which NPV devices are connected to).
- 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 varies depending on the switch model. By default, the first port in every four ports is selected as NP port. For example, 1st, 5th, 9th, and so on. For details about port numbers for NPV-eligible switches, see Cisco MDS 9000 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 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.

Cisco NPV Load Balancing

The Cisco NPV load balancing scheme automatically assigns traffic for each server to a logical external interface (uplink) when the server logs in to the fabric. These logical interfaces are usually F/NP port-channels but may also be individual Fibre Channel ports.

Cisco NPV switches can have multiple logical external interfaces, for example, when there are dual core switches in a single fabric. In this case, when a new server interface comes up, the external interface with the

least number of server interfaces assigned to it is selected for the new server interface. Because individual server interfaces may have different loads, selecting external interfaces solely based on the number of logged in server interfaces may lead to uneven utilization on external interfaces in the transmit, receive, or both directions.

Also, if an additional external interface is activated, the existing logged in server interfaces are not automatically rebalanced to include the new external interface. Only the server interfaces that come up after the new external interface are activated will get assigned to it.

After a server interface is logged in and assigned to a specific external interface, it cannot be moved to another external interface nondisruptively. It must first log out from the fabric which stops traffic through the server interface, and then log in on the other external interface.

The following are the challenges of this load balancing scheme when used with multiple external interfaces:

- Unable to optimally utilize the external interface bandwidth which may result in saturating bandwidth only on certain links and switches.
- Impact on the performance of the servers that are connected to an external interface that is overloaded.
- Sustained high load on any external interface may result in propagating slow drain condition to the other links in the fabric.

To improve the performance of the load balancing scheme, extra bandwidth can be added to each of the logical external interfaces. For example, in a dual core topology if there is an F/NP port-channel to each core switch, each should have sufficient bandwidth to handle the load of all server interfaces on the NPV switch. This is important in the event of a core switch failure and will also ensure that no single external interface gets over utilized.

Instead of using the traditional load balancing scheme and based on the least login count, users can now choose a new load balancing schema based on average link utilization. The **show npv traffic-map proposed** command may be used to find a mapping of server interfaces to external interfaces based on their measured loads so that server traffic can be evenly distributed on the external interfaces. This information is calculated and updated every 5 minutes. You can use this information to manually map the server interfaces to external interfaces using the **npv traffic-map server-interface** command. You can use the **npv traffic-map analysis clear** command to reset the link loads, but it does not reset the timer for calculating the loads.

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

The following sections provide information about guidelines and limitations for N Port Virtualization.

NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- For information on the number of NPV switches per NPIV switch, see the "Switch-Level Fibre Channel Configuration Limits for Cisco MDS 9000 Series Switches" section in the Cisco MDS NX-OS Configuration Limits.
- Logins that are sent from Cisco NPV switch toggle on F port-channel when the FCNS limit reaches 20,000.
- You can configure zoning for end devices that are connected to NPV switches using all available member types on an NPIV switch. However, the preferred way of zoning servers connected to any switch in NPV mode is via pWWN, device-alias, and fcalias. 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*.
- NPV switches can be connected to upstream NPIV switches using links that are not part of port channel. In this configuration, NPV uses a load balancing algorithm to automatically and efficiently assign end devices to one of the NPIV switch links when they login to the fabric. Only links in the same VSAN as the end device are considered by the algorithm. All traffic to and from that end device then uses the assigned link; VSAN load balancing is not applied to traffic on the NPV-NPIV links. If there are multiple links between an NPV devices and the upstream NPIV switch, it is possible to override the default and assign end devices to a specific link using a traffic map. There is no dynamic login rebalancing in the case of a link brought up between the NPV and NPIV switches — it is not used until an end device logs in and is assigned to it.

  There is dynamic login rebalancing in the case of link failure between the NPV and NPIV switches. If an NPV-NPIV link fails, the end device assigned to it are logged out by the NPV switch and must relogin to the fabric. The logins are then distributed over the remaining NPV-NPIV links.
- NPV switches can be connected to the NPIV switch via F port channels. In this configuration, end device logins are associated with the F port channel interface and not with any individual F port channel member. Failure of a member interface does not force end devices using the link to be logged out. Depending on the nature of the link failure, the end devices may experience some frame loss; however, if they can recover from this then they can continue normal operation using the remaining F port channel members. Likewise, if new members are added to an F port channel, all end devices utilizing it can immediately take advantage of the increased bandwidth. F port channels can also be configured for trunking (able to carry one or more VSANs). For these reasons, we recommend the use of F port channels when connecting NPV switches to the NPIV switch.
- Both servers and targets can be connected to an NPV switch. Local switching is not supported; all traffic is switched using the NPIV switch.
- NPV switches can be connected to multiple NPIV switches. In other words, different NP ports can be connected to different NPIV switches.

- Some devices will login to the fabric multiple times requesting multiple FCIDs on a single interface. To support this multiple logins, the **feature npiv** command must be enabled. This is also supported on NPV switches. Consequently, both the **feature npv** and **feature npiv** commands can be enabled on the same switch.
- You cannot configure BB_SCN on NPV switches that are using xNP ports because of interoperability issues with third-party NPIV switches.
- Nondisruptive upgrades are supported on NPV switches.
- Port security is supported on the NPIV switch for devices logged in via NPV.
- Only F, NP, and SD ports are supported on NPV switches.
- CDP is not supported on NPV devices.

NPV Traffic Management Guidelines:

- Use NPV traffic management only when the default login balancing by the NPV switch is not sufficient.
- Do not configure traffic maps for all servers. For non-configured servers, NPV will use the default login balancing.
- Ensure that the persistent FCID feature is not disabled on the upstream NPIV switch. Traffic engineering directs the associated server interface to external interfaces that lead to the same NPIV switch.
- A traffic map constrains the server interface to use the set of external interfaces specified. The server interface cannot use any other external interfaces that may be available even if all the specified external 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 NPIV switch through F port leading to traffic disruption.
- If an NPV switch is connected to multiple upstream NPIV switches, server interface traffic may be forced to only use subset of the upstream NPIV switches by specifying the set of external interfaces between the NPV switch and the desired NPIV switches in a traffic map.

NPIV Guidelines and Limitations

- If the NPIV feature was enabled using the **feature npiv** command and you are upgrading to Cisco MDS NX-OS Release 8.4(2) or later release, the NPIV feature remains enabled.
- If the NPIV feature was not enabled using the **feature npiv** command and you are upgrading to Cisco MDS NX-OS Release 8.4(2) or later release, the NPIV feature remains disabled.
- From Cisco MDS NX-OS Release 8.4(2), the NPIV feature is enabled by default. Therefore, the **feature npiv** command will not be displayed in the running configuration if this feature is enabled and the **no feature npiv** command will be displayed in the running configuration if this feature is disabled.
- If migrating an MDS from Cisco MDS NX-OS Release 8.4(2) or a later release to a release earlier than Cisco MDS NX-OS Release 8.4(2), then the behaviour of the NPIV feature depends on how it is configured and how the migration is performed. If the NPIV feature is enabled before the migration (the default configuration) and the migration is done via an ISSD downgrade, then NPIV remains enabled when the migration has completed (a nondefault configuration in these releases). If the NPIV feature is enabled

before the migration (the default configuration) and the migration is done via a reboot, then NPIV will be disabled after the migration has completed (the default configuration in these releases).

- If you are upgrading switches that have the NPIV feature disabled to Cisco MDS NX-OS Release 8.4(2) or later releases and if you are adding new switches that are running Cisco MDS NX-OS Release 8.4(2) or later releases that have the NPIV feature enabled by default to a fabric, ensure that you either disable the NPIV feature on the new switches or enable the NPIV feature on your exiting switches.

DPVM Configuration Guidelines for NPIV

The following requirements must be met before you configure DPVM on the core switch to which NPV devices are connected to:

- You must explicitly configure the WWN of the internal FLOGI in DPVM. If DPVM is configured on the core switch to which NPV devices are connected to 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 core switches to which NPV devices are connected to F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see *Cisco MDS 9000 Series NX-OS Fabric Configuration Guide*.

NPV and Port Security Configuration Guidelines

Port security is enabled on the NPIV switch on a per interface basis. To enable port security on the core switch to which NPV devices are connected to 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 core switch to which NPV devices are connected to 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 Security Configuration Guide.

Configuring N Port Virtualization

The following sections provide information about guidelines and limitations for 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 FCIDs.

Note All of the N port FCIDs are allocated in the same VSAN.

To enable or disable NPIV on the switch, perform these steps:

Procedure

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, perform the following steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode on the NPIV core switch.

Step 2 switch(config)# **feature npiv**

Enables NPIV mode on the NPIV core switch.

switch(config)# **no feature npiv**

(Optional) Disables NPIV mode on the NPIV core switch.

Step 3 switch(config)# **interface fc 2/1**

Configures the NPIV core switch port as an F port.

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

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

Changes Admin status to bring up the interfaces.

Step 4 switch(config)# **vsan database**

switch(config-vsan-db)# **vsan 8 interface fc 2/1**

Configures the port VSANs for the F port on the NPIV core switch.

Step 5 switch(config)# **feature npv**

Enables NPV mode on a NPV device. 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.

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

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

Changes Admin status to bring up the interfaces.

Step 7 switch(config-if)# **exit**

Exits interface mode for the port.

Step 8 switch(config)# **vsan database**

switch(config-vsan-db)# **vsan 9 interface fc 1/1**

Configures the port VSANs for the NP port on the NPV device.

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.

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

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

Changes Admin status to bring up the interfaces.

Step 10 switch(config)# **vsan database**

switch(config-vsan-db)# **vsan 12 interface fc 1/1 - 6**

Configures the port VSANs for the F ports on the NPV device.

Step 11 switch(config-npv)# **no feature npv**

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:

Procedure

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 port channel 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 port channel interfaces and port channel 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 Cisco NPV traffic management feature on Cisco 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:

Procedure

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 core switch to which NPV devices are connected to.

Step 3 switch (config)# **no npv auto-load-balance disruptive**

Disables disruptive load balancing on the core switch to which NPV devices are connected to.

Verifying NPV and NPIV Configuration

To display NPV configuration information, perform one of the following tasks:

Command	Purpose
show fcns database	Displays all the NPIV core 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 NPIV core 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 *Cisco MDS 9000 Series NX-OS Command Reference*.

Verifying NPV

To view all the NPIV 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 NPIV 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                                                                  EXTERNAL
INTERFACE VSAN FCID            PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
fc1/19    1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc1/1
fc1/19    1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc1/1
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, on page 394
- Information About FlexAttach Virtual pWWN, on page 395
- Guidelines and Limitations, on page 397
- Configuring FlexAttach Virtual pWWN, on page 398
- Verifying FlexAttach Virtual pWWN Configuration, on page 401
- Monitoring FlexAttach Virtual pWWN, on page 403

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://bst.cloudapps.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.

Note This feature is supported on switches in NPV mode only.

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 47: 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.

SAN Device Virtualization (SDV)	FlexAttach Virtualization
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).
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.
Non-NPV mode only	NPV mode only

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.



Note NPV switches do not have ISL (E or TE) ports, and can therefore use IPv4 or IPv6 only for CFS distribution.

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 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.
- This feature is supported on switches in NPV mode only.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Assign FlexAttach virtual pWWN automatically for the interfaces:

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:

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.
- This feature is supported on switches in NPV mode only.

Procedure

Step 1 Enter configuration mode:

switch#**configure terminal**

Step 2 Configure the FlexAttach virtual pWWN for the interface:

switch(config)# **flex-attach virtual-pwwn vpwwn interface** *interface-list*

Step 3 (Optional) Configure the FlexAttach virtual pWWN for the interface in the VSAN:

switch(config)# **flex-attach virtual-pwwn vpwwn interface** *interface* [**vsan** *vsan-range*]

Step 4 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 FDSICs 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 NPIV 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.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Map the pWWN to the virtual pWWN:

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 Command Reference.

To view and confirm that the type and value of virtual pWWNs are correct, enter the **show flex-attach virtual-pwwn** command.

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 NPIV core.

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 48: 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 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, on page 406
- Information About Port Tracking, on page 407
- Guidelines and Limitations, on page 408
- Default Settings, on page 409
- Configuring Port Tracking, on page 410
- Verifying Port Tracking Configuration, on page 415

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://bst.cloudapps.cisco.com/bugsearch and the release notes for your software release.

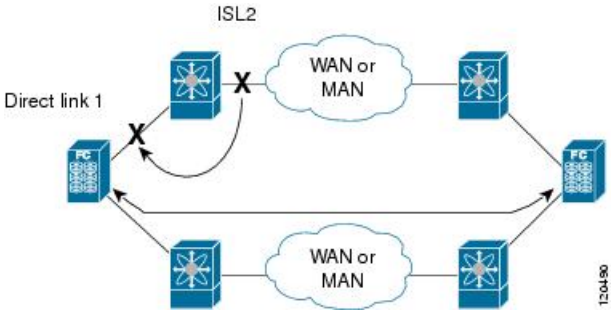
To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

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 26: Traffic Recovery Using Port Tracking, on page 407, 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 26: 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, port channel, 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 49: Default Port Tracking Parameters , on page 409 lists the default settings for port tracking parameters.

Table 49: 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:

Procedure

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:

Procedure

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**

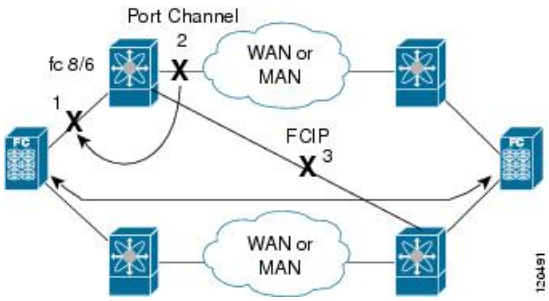
(Optional) Removes the port tracking configuration that is currently applied to interface fc8/6.

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 27: Traffic Recovery Using Port Tracking, on page 411, 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 27: Traffic Recovery Using Port Tracking



Tracking Multiple Ports

To track multiple ports, perform these steps:

Procedure

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 27: Traffic Recovery Using Port Tracking, on page 411.

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 27: Traffic Recovery Using Port Tracking, on page 411.

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 27: Traffic Recovery Using Port Tracking, on page 411.

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:

Procedure

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 port channel 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 port channel 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:

Procedure

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 415 to Displays a Forced Shutdown Configuration , on page 416).

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 Port Channel 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
```