



## CHAPTER 2

# Configuring Interfaces

---

This chapter describes the basic interface configuration to get your switch up and running.

This chapter includes the following topics:

- [Information About Interfaces section, page 2-1](#)
- [Prerequisites for Interfaces section, page 2-17](#)
- [Guidelines and Limitations section, page 2-17](#)
- [Default Settings section, page 2-19](#)
- [Configuring Interfaces section, page 2-19](#)
- [Verifying Interfaces Configuration section, page 2-27](#)

## 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.

This section includes the following topics:

- [Interface Description section, page 2-2](#)
- [Interface Modes section, page 2-2](#)
- [Interface States section, page 2-6](#)
- [Graceful Shutdown section, page 2-9](#)
- [Port Administrative Speeds section, page 2-10](#)
- [Beacon LEDs section, page 2-11](#)
- [Speed LEDs section, page 2-11](#)
- [Bit Error Thresholds section, page 2-11](#)
- [SFP Transmitter Types section, page 2-12](#)
- [TL Ports section, page 2-12](#)
- [TL Port ALPA Caches section, page 2-14](#)
- [Port Guard section, page 2-14](#)
- [Port Monitor section, page 2-15](#)

- [Port Monitor Port Guard section, page 2-15](#)
- [Port Group Monitor section, page 2-15](#)
- [Local Switching section, page 2-16](#)
- [Management Interfaces section, page 2-17](#)
- [VSAN Interfaces section, page 2-17](#)

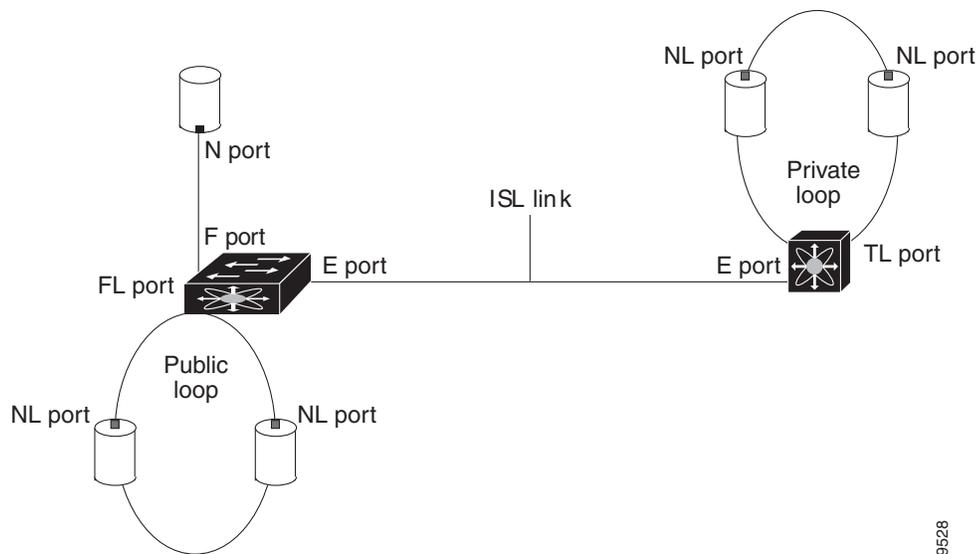
## Interface Description

For the Fibre Channel interfaces, you can configure the description parameter to provide a recognizable name for the interface. Using a unique name for each interface allows you to quickly identify the 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 that 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, ST port, and B port (see [Figure 2-1](#)). 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 2-1** Cisco MDS 9000 Family Switch Port Modes



### Note

Interfaces are created in VSAN 1 by default. See the *Cisco MDS 9000 Family 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 like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

**Note**

---

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

---

Each interface is briefly described in the sections that follow.

## E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may 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 to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see [Chapter 1, “Configuring PortChannels .”](#)).

**Note**

---

We recommend that you configure E ports on 16-port modules. If you must configure an E port on a 32-port oversubscribed module, then you can only use the first port in a group of four ports (for example, ports 1 through 4, 5 through 8, and so forth). The other three ports cannot be used.

---

## F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may 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 service.

## FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may 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 service.

**Note**

---

FL port mode is not supported on 4-port 10-Gbps switching module interfaces.

---

## 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.

**Note**


---

A Cisco Nexus 5000 Series switch in NPV mode that runs Cisco NX-OS Release 4.2(1) or later releases supports trunking F port mode on NP ports. You can enable either, or both, VSAN trunking and an F port on an NP port.

---

For more details about NP ports and NPV, see [Chapter 1, “Configuring N Port Virtualization.”](#)

## TL Port

In translative loop port (TL port) mode, an interface functions as a translative loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the [“TL Port ALPA Caches”](#) section on page 2-14).

**Tip**


---

We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

---

**Note**


---

TL port mode is not supported on Generation 2 switching module interfaces.

---

## TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They 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 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 Family switches (see [Chapter 1, “Configuring Trunking .”](#)). TE ports support class 2, class 3, and class F service.

## TF Port

In trunking F port (TF port) mode, an interface functions as a trunking expansion port. It may 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 HBA to carry tagged frames. TF ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of F ports to support VSAN trunking.

In TF port mode, all 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 Family (see [Chapter 1, “Configuring Trunking .’](#)). TF ports support class 2, class 3, and class F service.

## TNP Port

In trunking NP port (TNP port) mode, an interface functions as a trunking expansion port. It may be connected to a trunked F port (TF port) to create a link to a core NPV switch from an NPV switch 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 Family. It monitors network traffic that passes through a Fibre Channel interface. This monitoring 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 nonintrusive and does not affect switching of network traffic for any SPAN source ports (see the *Cisco MDS 9000 Family 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 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

**Note**

ST port mode is not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

## Fx Port

Interfaces configured as Fx ports can operate in either F port 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.

## B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*).

## Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, TE port, or TF port. 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 or FL port mode depending on the N port 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 Family, it may become operational in TE port mode (see [Chapter 1, “Configuring Trunking .’](#)).

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



### Note

Fibre Channel interfaces on Storage Services Modules (SSMs) cannot be configured in auto mode.

## Interface States

The 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 2-1](#).

**Table 2-1 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

The operational state indicates the current operational state of the interface as described in [Table 2-2](#).

**Table 2-2** Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic as desired. 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 or TF mode.

## Reason Codes

Reason codes are dependent on the operational state of the interface as described in [Table 2-3](#).

**Table 2-3** 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 <a href="#">Table 2-4</a> .



### Note

Only some of the reason codes are listed in [Table 2-4](#).

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 2-4](#).

**Table 2-4 Reason Codes for Nonoperational States**

<b>Reason Code (long version)</b>	<b>Description</b>	<b>Applicable Modes</b>
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. For example: <ul style="list-style-type: none"> <li>• Configuration failure.</li> <li>• Incompatible buffer-to-buffer credit configuration.</li> </ul> To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shut down or enable the interface.	
FC 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.	

**Table 2-4 Reason Codes for Nonoperational States (continued)**

<b>Reason Code (long version)</b>	<b>Description</b>	<b>Applicable Modes</b>
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. It may happen 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
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel 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 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 information about IOD, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*).
- If the Min\_LS\_interval interval is higher than 10 seconds. For information about FSPF global configuration, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.


**Note**


---

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

For internal ports on the Cisco Fabric Switch for HP c\_Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter, a port speed of 1 Gbps is not supported. Auto-negotiation is supported between 2 Gbps and 4 Gbps only. Also, if the BladeCenter is a T chassis, then port speeds are fixed at 2 Gbps and auto-negotiation is not enabled.

## Autosensing

Autosensing 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 the 4-Gbps switching modules, and 8 Gbps on the 8-Gbps switching modules. When autosensing 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 that it does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports that are configured for autosensing.


**Tip**


---

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

---

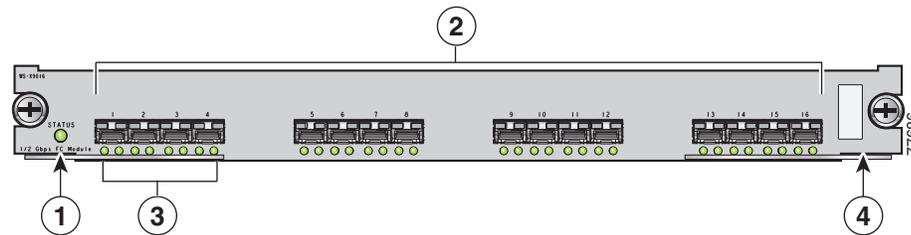
You can set the frame format to EISL for all frames transmitted by the interface in SD port mode. If you sent the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources. See the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* to configure frame encapsulation on an interface.

## Beacon LEDs

Figure 2-2 displays the status, link, and speed LEDs in a 16-port switching module.

**Figure 2-2 Cisco MDS 9000 Family Switch Interface Modes**



<b>1</b>	Status LED <sup>1</sup>	<b>3</b>	Link LEDs <sup>1</sup> and speed LEDs <sup>2</sup>
<b>2</b>	1/2-Gbps Fibre Channel port group	<b>4</b>	Asset tag <sup>3</sup>

1. See the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.

2. See the “Speed LEDs” section on page 2-11.

3. Refer to the Cisco MDS 9000 Family hardware installation guide for your platform.

## Speed LEDs

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off or solid green—Beacon mode is disabled.
- Flashing green—The beacon mode is enabled. The LED flashes at one-second intervals.



### Note

Generation 2, Generation 3, and Generation 4 modules and fabric switches do not have speed LEDs.

## Bit Error Thresholds

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

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or 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 sync loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached. You can enter a **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.

## SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed. [Table 2-5](#) defines the acronyms used for SFPs.

**Table 2-5 SFP Transmitter Acronym Definitions**

Definition	Acronym
<b>Standard transmitters defined in the GBIC specifications</b>	
short wavelaser	swl
long wavelaser	lwl
long wavelaser cost reduced	lwcr
electrical	elec
<b>Extended transmitters assigned to Cisco-supported SFPs</b>	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530
CWDM-1550	c1550
CWDM-1570	c1570
CWDM-1590	c1590
CWDM-1610	c1610

See the [“Displaying SFP Transmitter Types”](#) section on page 2-28.

## TL Ports

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop. The legacy devices are used in Fibre Channel networks, and devices outside the loop may need to communicate with them. The communication functionality is provided through TL ports. See the [“Interface Modes”](#) section on page 2-2.

TL port mode is not supported on the following hardware:

- Generation 2 switching module interfaces
- Cisco MDS 9124 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

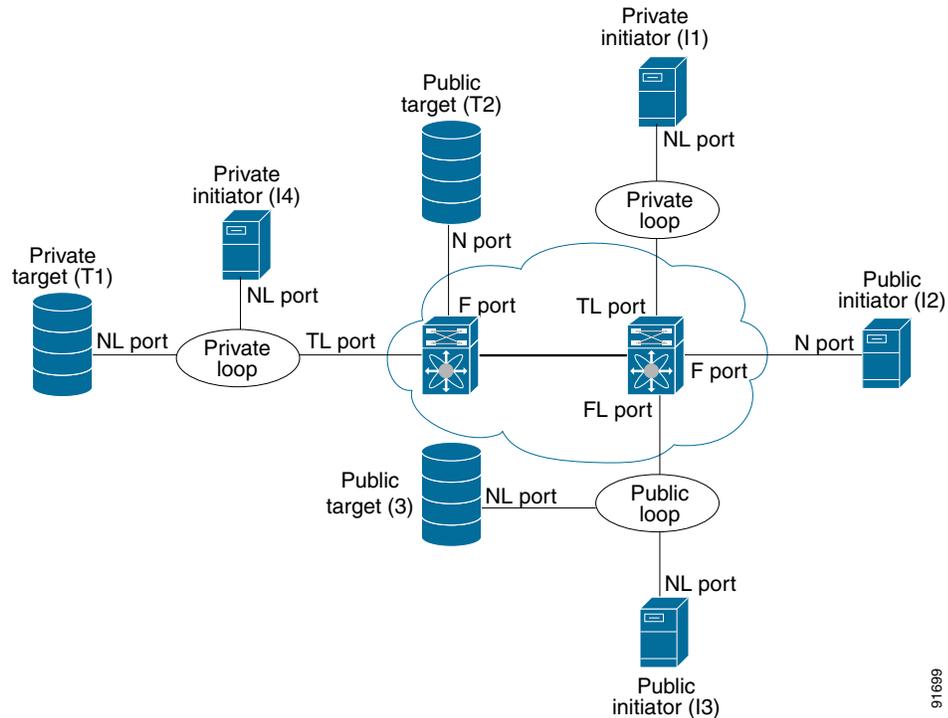
Table 2-6 lists the TL port translations supported in Cisco MDS 9000 Family switches.

**Table 2-6 Supported TL Port Translations**

Translation from	Translation to	Example
Private initiator	Private target	From I1 to T1 or vice versa
Private initiator	Public target — N port	From I1 to T2 or vice versa
Private initiator	Public target — NL port	From I4 to T3 or vice versa
Public initiator — N port	Private target	From I2 to T1 or vice versa
Public initiator — NL port	Private target	From I3 to T1 or vice versa

Figure 2-3 shows examples of TL port translation support.

**Figure 2-3 TL Port Translation Support Examples**



916599

## TL Port ALPA Caches

Although TL ports cannot be automatically configured, you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco NX-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco NX-OS software discards an inactive cache entry (if available) to make space for the new entry. See the “TL Port” section on page 2-4 for more information on TL ports.

Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* to manage the TL Port ALPA cache.

## Port Guard

The port guard feature is intended for use in environments where the system and application environment does not adapt quickly and efficiently to a port going down and back up, or to a port rapidly cycling up and down, which can happen in some failure modes. For example, if a system takes five seconds to stabilize after a port goes down, but the port is going up and down once a second, a more severe failure in the fabric might occur.

The port guard feature gives the SAN administrator the ability to prevent this issue from occurring in environments that are vulnerable to these problems. The port can be configured to stay down after the first failure or after a specified number of failures in a specified time period. This allows the SAN administrator to intervene and control the recovery, avoiding any problems caused by the cycling.

Using the port guard feature, you can restrict the number of error reports and bring a malfunctioning port to down state dynamically. A port can be configured to go into error-disabled state for specific types of failures.

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 to the number of link-down link failures. This means a port is brought to down state when it reaches the maximum number of allowed link failures or the number of specific causes.

The causes of link failure can be any of the following:

- ESP trustsec-violation
- Bit-errors
- Signal loss
- Sync loss
- Link reset
- Credit loss
- Additional causes might be the following:
  - Not operational (NOS).
  - Too many interrupts.
  - Cable is disconnected.
  - Hardware recoverable errors.

- The connected device rebooted (F ports only).
- The connected linecard rebooted (ISL only).

## Port Monitor

Port monitor helps to monitor the performance and the status of ports and generate alerts when problems occur. You can configure the thresholds for various counters and trigger an event when the values cross the threshold settings.

The default port monitor policy has the following threshold values:

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	Event	% Falling Threshold	Event
Link Loss	Delta	60	5	4	1	4
Sync Loss	Delta	60	5	4	1	4
Protocol Error	Delta	60	1	4	0	4
Signal Loss	Delta	60	5	4	1	4
Invalid Words	Delta	60	1	4	0	4
Invalid CRCs	Delta	60	5	4	1	4
RX Performance	Delta	60	2147483648	4	524288000	4
TX Performance	Delta	60	2147483648	4	524288000	4

## Port Monitor Port Guard

Port monitor port guard is a feature that disables or shuts down a port when an event occurs. Depending on the configuration, when an event occurs the port is either error-disabled or flapped.

Port monitor port guard is a different or separate feature that functions based on the configuration of the **errordisable** command.

## Port Group Monitor

Each line card or module has a predefined set of ports which share the same backplane bandwidth called port groups. While oversubscription is a feature, the port group monitor feature helps to monitor the spine bandwidth utilization. An alarm syslog is generated so that you can provision the ports across port groups evenly to manage the oversubscription better.

When the port group monitor feature is enabled and a policy consisting of polling interval in seconds, and the raising and falling thresholds in percentage are specified, port group monitor generates a syslog if a port group traffic goes above the specified percentage of the maximum supported bandwidth for that port group (for rx and for tx) and another syslog if the value falls below the specified threshold.

The default port group policy has the following threshold values:

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

## Local Switching

Local switching can be enabled in Generation 4 modules, which allows traffic to be switched directly with a local crossbar when the traffic is directed from one port to another on the same line card. By using local switching, an extra switching step is avoided, which decreases the latency.

When using local switching, note the following guidelines:

- All ports need to be in shared mode, which usually is the default state. To place a port in shared mode, enter the **switchport ratemode shared** command.
- E ports are not allowed in the module because they must be in dedicated mode.

## Slow Drain Device Detection and Congestion Avoidance

All data traffic between end devices in a SAN fabric is carried by Fibre Channel Class 3. In some cases, the traffic is carried by Class 2 services that use link-level, per-hop-based, and buffer-to-buffer flow control. These classes of service do not support end-to-end flow control. When there are slow devices attached to the fabric, the end devices do not accept the frames at the configured or negotiated rate. The slow devices lead to ISL credit shortage in the traffic destined for these devices and they congest the links. The credit shortage affects the unrelated flows in the fabric that use the same ISL link even though destination devices do not experience slow drain.

This feature provides various enhancements to detect slow drain devices that are causing congestion in the network and also provides a congestion avoidance function.

This feature is focused mainly on the edge ports that are connected to slow drain devices. The goal is to avoid or minimize the frames being stuck in the edge ports due to slow drain devices that are causing ISL blockage. To avoid or minimize the stuck condition, configure lesser frame timeout for the ports. No-credit timeout drops all packets once the slow drain is detected using the configured thresholds. The lesser frame timeout value helps to alleviate the slow drain condition that affects the fabric by dropping the packets on the edge ports sooner than the time they actually get timed out (500 ms). This function frees the buffer space in ISL, which can be used by other unrelated flows that do not experience slow drain condition.



### Note

This feature is used mainly for edge ports that are connected to slow edge devices. Even though this feature can be applied to ISLs as well, we recommend that you apply this feature only for edge F ports and retain the default configuration for ISLs as E and TE ports. This feature is not supported on Generation 1 modules.

## Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection on the mgmt0 interface, you must configure either the IP version 4 (IPv4) parameters (IP address, subnet mask, and default gateway) or the IP version 6 (IPv6) parameters so that the switch is reachable.

Before you begin to configure the management interface manually, obtain the switch's IPv4 address and subnet mask, or the IPv6 address.

The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 10/100/1000 Mbps. Autosensing supports both the speed 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 is auto and the default duplex mode is auto.

**Note**

You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

## VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexisting VSANs.

## Prerequisites for Interfaces

Before you begin configuring the interfaces, ensure that the modules in the chassis are functioning as designed. For information about verifying the module status, refer to the *Cisco NX-OS Fundamentals Configuration Guide*.

## Guidelines and Limitations

This section includes the following topics:

- [Generation 1 Interface Configuration Guidelines section, page 2-17](#)
- [Private Loop Configuration Guidelines section, page 2-18](#)
- [VSAN Interface Configuration Guidelines section, page 2-19](#)

## Generation 1 Interface Configuration Guidelines

The Generation 1 interfaces configuration guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module interfaces
- The Cisco MDS 9140 and 9120 switch interfaces




---

**Note** Due to the hardware design of the MDS 9134 switch, we do not support interface out-of-service action on either of its two 10-Gigabit ports. This is because no internal port hardware resource is released when an out-of-service action is performed on these 10-Gigabit ports.

---

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8, and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8, and so on) are not usable and remain shutdown.
- 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.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules.
- The 32-port switching module does not support FICON.




---

**Note** We recommend that you configure your E ports on a 16-port switching module. If you must configure an E port on a 32-port host-optimized switching module, the other three ports in that 4-port group cannot be used.

---




---

**Note** In the Cisco MDS 9100 Series, the groups of ports that are located on the left and outlined in white are full line rate. The other ports are host-optimized. Each group of 4 host-optimized ports have the same features as for the 32-port switching module.

---

## Private Loop Configuration Guidelines

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxy to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxy to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

## VSAN Interface Configuration Guidelines

Follow these guidelines when creating or deleting VSAN interfaces:

- 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 Family NX-OS IP Services Configuration Guide*.

## Default Settings

Table 2-7 lists the default settings for interface parameters.

**Table 2-7**      **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

This section includes the following topics:

- [Common Interface Configuration section, page 2-20](#)
- [Configuring Fibre Channel Interfaces section, page 2-20](#)
- [Setting the Interface Administrative State section, page 2-20](#)
- [Configuring Interface Modes section, page 2-21](#) [Configuring Port Administrative Speeds section, page 2-22](#)
- [Configuring Port Speed Group section, page 2-22](#)

- [Configuring the Interface Description section, page 2-23](#)
- [Specifying a Port Owner section, page 2-23](#)
- [Configuring Beacon Mode section, page 2-23](#)
- [Disabling Bit Error Threshold section, page 2-24](#)
- [Configuring Switch Port Attribute Default Values section, page 2-24](#)
- [Configuring TL Ports section, page 2-24](#)
- [Configuring Port Guard section, page 2-25](#)
- [Configuring Management Interfaces section, page 2-26](#)
- [Creating VSAN Interfaces section, page 2-26](#)

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

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

## Common Interface Configuration

Some configuration settings are similar for Fibre Channel, management, and VSAN interfaces. You can configure interfaces from DCNM-SAN by expanding **Switches > FC Interfaces** and selecting either the Physical or Logical interface type from the Physical Attributes pane.

## Configuring Fibre Channel Interfaces

For the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter, you can configure a range of interfaces among internal ports or external ports, but you cannot mix both interface types within the same range. For example, “bay 1-10, bay 12” or “ext 0, ext 15-18” are valid ranges, but “bay 1-5, ext 15-17” is not.

## Setting the Interface Administrative State

To disable or enable an interface using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**. You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab.
  - Step 3** Click **Mode admin**.  
You see the drop-down box.

- Step 4** Set the status or mode to the required status.
  - Step 5** (Optional) Set other configuration parameters using the other tabs.
  - Step 6** Click **Apply Changes**.
- 

## Configuring Interface Modes

To configure the interface mode using DCNM-SAN, follow these steps:

- Step 1** Expand **Switches > FC Interfaces > Physical**.  
You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab.
  - Step 3** Click **Mode Admin**. Set the desired interface mode from the Admin drop-down menu.
  - Step 4** (Optional) Set other configuration parameters using the other tabs.
  - Step 5** Click **Apply Changes** icon.
- 

## Configuring 10-Gbps FC Mode

The 48-port 8-Gbps Advanced Fibre Channel module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9) can switch between two speed modes—the 1-, 2-, 4-, 8-Gbps or 10-Gbps. By default, the modules are online in the 1-, 2-, 4-, and 8-Gbps modes when they are loaded for the first time. There are two ways to change the ports to the 10-Gbps speed mode:

- Using the **10G-speed mode** command, which is the recommended method.
- Using the generic speed configuration **switchport speed** command which has certain constraints.

The following conditions apply when the ports in the module can be configured to 10-Gbps speed mode:

- The ports in the module can be configured to 10-Gbps speed only when the DS-13SLT-FAB3 module bandwidth is 256-G. Any other combination of fabric modules will not let the ports come up in 10-Gbps.
- When in 10-Gbps mode, the ports in the module that are not 10-Gbps capable are disabled and will be in out-of-service state.
- The ports function only in full rate mode. They cannot be moved to shared rate mode.
- The ports cannot be configured in any other speed.
- Ports that are capable of 10-Gbps that are disabled or out-of-service cannot be put back in service using the **no out-of-service** command. To put these ports back in service, all ports in the module first have to be moved to the out-of-service state. Then they can be brought back to the in service state.
- Local switching must be disabled, otherwise, ports cannot be configured in dedicated mode.

Only certain ports on the 48-port and 32-port 8-Gbps Advanced Fibre Channel modules are 10-Gbps capable. When running in 10-Gbps mode, the non-10-Gbps ports cannot be operational. They have to be either in shut state or out-of-service state.

To configure the interface mode using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**.  
You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab.
  - Step 3** Click **Mode Admin**. Set the desired interface mode from the Admin drop-down menu.
  - Step 4** (Optional) Set other configuration parameters using the other tabs.
  - Step 5** Click **Apply Changes** icon.
- 

## Configuring Port Administrative Speeds

### Restrictions

- Changing the port administrative speed is a disruptive operation.

To configure the administrative speed of the interface using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**.  
You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab.
  - Step 3** Click **Speed Admin**. Set the desired speed from the drop-down menu.  
The number indicates the speed in megabits per second (Mbps). You can set the speed to 1-Gbps, 2-Gbps, 4-Gbps, 8-Gbps, autoMax2G, autoMax4G, or auto (default).



**Note** On a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2), you can configure the 8-Gbps administrative speed only on a M1060 switch module. You can configure the speed to 1-Gbps, 2-Gbps, or 4-Gbps on all switch modules on a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2) or earlier releases.

- Step 4** Click **Apply Changes**.
- 

## Configuring Port Speed Group

To configure the administrative speed of the interface using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**.  
You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab.
  - Step 3** Click **SpeedGroup**. Set the desired speed group from the drop-down menu.

You can select any of the speed groups from the menu list—notApplicable, tenG, oneTwoFourEightG, or twoFourEightSixteenG.



**Note** For a DS-X9248-256K9 or DS-X9232-256K9 line card, the speed group must be set to tenG.

**Step 4** Click **Apply Changes**.

---

## Configuring the Interface Description

Interface descriptions enable you to identify the traffic or the use for that interface. The interface description can be any alphanumeric string.

## 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 port guard and port owner features are available for all ports regardless of the operational mode.

To specify or remove the port owner using DCNM-SAN, follow these steps:

---

- Step 1** Expand **Switches > FC Interfaces > Physical**.  
You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab and then select the switch/port.
  - Step 3** In the Owner text box, enter a port owner and the purpose for which port is used.
- 

To specify or remove the port owner using Device Manager, follow these steps:

---

- Step 1** Double-click the interface in the modules panel.
  - Step 2** Click the **General** tab.
  - Step 3** In the Owner text box, enter a port owner and the purpose for which the port is used.
  - Step 4** Click **Apply**.
- 

## 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. Configuring the beacon mode has no effect on the operation of the interface.

To enable beacon mode for a specified interface or range of interfaces using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > Ethernet Interfaces > Physical > IPS** (the Gigabit Ethernet Interfaces).  
You see the interface configuration in the Information pane.
  - Step 2** Click the Beacon Mode and enable this option for the selected switch.
  - Step 3** Click **Apply Changes**.
- 

### Troubleshooting Tips

- The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. 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.

## Disabling Bit Error Threshold



### Note

Even if you disable the switch port ignore bit-error threshold for an interface, the switch generates a syslog message when bit-error threshold events are detected.

---

## Configuring Switch Port Attribute Default Values

You can configure attribute 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.

Refer to the Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide to configure switch port attributes.

## Configuring TL Ports

Private loops require setting the interface mode to TL.

To configure the TL interface mode using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**. You see the interface configuration in the Information pane.
  - Step 2** Click the **General** tab and click **Mode Admin**.
  - Step 3** Set the Mode Admin drop-down menu to the required status.
  - Step 4** (Optional) Set other configuration parameters using the other tabs.
  - Step 5** Click **Apply Changes**.
-

## Configuring Port Guard

To enable port guard using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical > Port Guard** from the Physical Attributes pane.  
You see the interfaces listed in the Information pane.
  - Step 2** Click the **Link Down** tab and then select a switch or port.
  - Step 3** Check the check box in the Enable column.
  - Step 4** (Optional) Enter the Duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if the link flaps even once. Otherwise, the link is brought to down state if the link flaps for the number of flaps within the duration.
  - Step 5** Click **Apply Changes** to activate the configuration.
  - Step 6** Click the **TrustSec Violation** tab, and then select a switch or port.
  - Step 7** Check the check box in the Enable column.
  - Step 8** (Optional) Enter the duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if a trustsec violation occurs even once. Otherwise, the link is brought to down state if there is trustsec violation for the number of flaps within the duration.
  - Step 9** Click the **Bit Errors, Signal Loss, Sync Loss, Link-reset, and Credit Loss** tabs and complete the port guard configuration.
  - Step 10** Click **Apply Changes** to activate the configuration.
- 

To enable port guard for single or multiple interfaces using Device Manager, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical > Port Guard** from the Physical Attributes pane.  
You see the FC Interfaces listed.
  - Step 2** Click the **Link Down** tab, and then select the switch or port.
  - Step 3** Check the check box in the Enable column.
  - Step 4** (Optional) Enter the duration in seconds and the number of flaps. If the values are 0, the port goes into a down state even if the link flaps once. Otherwise, the link goes into a down state if the link flaps for the number of flaps within the duration.
  - Step 5** Click **Apply Changes** to activate the configuration.
  - Step 6** Click the **TrustSec Violation** tab, and then select the switch or port.
  - Step 7** Check the check box in the Enable column.
  - Step 8** (Optional) Enter the Duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if a trustsec violation occurs even once. Otherwise, the link is brought to down state if a trustsec violation occurs for the number of flaps within the duration.
  - Step 9** Click **Apply Changes** to activate the configuration.
-

**Note**


---

By default, the port monitor port guard is disabled. To enable this feature, you must explicitly configure the port monitor port guard feature on a particular counter by performing Step 3 or Step 4.

---

## Configuring Management Interfaces

To configure the management interface using DCNM-SAN, follow these steps:

- 
- Step 1** Select a VSAN in the Logical Domains pane.
  - Step 2** In the Physical Attributes pane, expand **Switches > Management**.  
You see the interface configuration in the Information pane.
  - Step 3** Click the **IP Addresses** tab and set the Interface, IP Address/Mask field.
  - Step 4** Click the **General** tab and set the Status Admin to **up**.
  - Step 5** (Optional) Set other configuration parameters using the other tabs.
  - Step 6** Click **Apply Changes**.
- 

## Creating VSAN Interfaces

To create a VSAN interface using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > Management**.
  - Step 2** Click **Create Row**.  
You see the Create Interface dialog box.
  - Step 3** Select the switch and VSAN ID for which you want to configure a VSAN interface.

**Note**


---

You can only create a VSAN interface for an existing VSAN. If the VSAN does not exist, you cannot create a VSAN interface for it.

---

- Step 4** Set IPAddress/Mask to the IP address and subnet mask for the new VSAN interface.
  - Step 5** Click **Create** to create the VSAN interface or click **Close** to close the dialog box without creating the VSAN interface.
-

## Verifying Interfaces Configuration

This section includes the following topics:

- [Monitoring a Port Group section, page 2-28 section, page 2-27](#)
- [Obtaining Interface Statistics section, page 2-27](#)
- [Displaying SFP Transmitter Types section, page 2-28](#)
- [Monitoring a Port Group section, page 2-28](#)

## Displaying the Owned Ports

To display the interfaces owned using Device Manager, follow these steps:

- 
- Step 1** From the menu bar, click the **Ports All** drop-down button.
- Step 2** From the drop-down list, select **Owned**.
- 

## Obtaining Interface Statistics

You can use DCNM-SAN or Device Manager to collect interface statistics on any switch. These statistics are collected at intervals that you can set.



**Note** In DCNM-SAN, you can collect interface statistics by expanding **ISLs** and selecting **Statistics** from the Physical Attributes pane.

---

To obtain and display interface counters using Device Manager, follow these steps:

- 
- Step 1** From the menu bar, click **Interface**. Select **Monitor**.  
Select any of the Interfaces that are displayed. For example, Virtual FC Enabled.
- Step 2** Set both the number of seconds at which you want to poll the interface statistics and how you want the data represented in the Interval drop-down menus. For example, click **10s** and **LastValue/sec**.
- Step 3** Select any tab to view those related statistics.
- Step 4** (Optional) Click the **Pencil** icon to reset the cumulative counters.
- Step 5** (Optional) Click the **Save** icon to save the gathered statistics to a file or select the **Print** icon to print the statistics.
- Step 6** Click **Close** when you are finished gathering and displaying statistics.
-

## Displaying SFP Transmitter Types

To show the SFP types for an interface using DCNM-SAN, follow these steps:

- 
- Step 1** Expand **Switches > FC Interfaces > Physical**. You see the interface configuration in the Information pane.
  - Step 2** Click the **Physical** tab to see the transmitter type for the selected interface.
- 

## Monitoring a Port Group

To monitor a particular group using Device Manager, follow these steps:

- 
- Step 1** Right-click any port group module and select **Check Oversubscription**.  
The **Check Oversubscription** table is displayed.
  - Step 2** From the **Monitor** drop-down list, select one particular group to monitor.  
The Device Manager displays the monitoring table of the selected group with counters on each interval and displays the line chart automatically. From the Monitoring table, you can also choose the **Bar chart** icon to view the selected group as bar charts.
-