



CHAPTER 20

Configuring Interfaces

A switch's main function is to relay frames from one data link to another. To do that, 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 chapter describes the basic interface configuration to get your switch up and running. It includes the following sections:

- [Common Interface Configuration, page 20-1](#)
- [Fibre Channel Interfaces, page 20-2](#)
- [TL Ports for Private Loops, page 20-18](#)
- [Buffer Credits, page 20-20](#)
- [Management Interfaces, page 20-23](#)
- [VSAN Interfaces, page 20-24](#)
- [Default Settings, page 20-26](#)

See [Chapter 12, “Initial Configuration,”](#) and [Chapter 43, “Configuring IP Services,”](#) for more information on configuring mgmt0 interfaces.

See [Chapter 53, “Configuring IPv4 for Gigabit Ethernet Interfaces”](#) and [Chapter 46, “Configuring IPv6 for Gigabit Ethernet Interfaces”](#) for more information on configuring Gigabit Ethernet interfaces.



Tip

Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed. See the [“Verifying the Module Status”](#) section on page 12-2.

Common Interface Configuration

Some configuration settings are similar for Fibre Channel, management, and VSAN interfaces. You can configure interfaces from Fabric Manager by expanding **Switches > Interfaces** and selecting the interface type from the Physical Attributes pane.

[Figure 20-1](#) shows a sample of what you might see in the Information pane for Fibre Channel Interfaces.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Figure 20-1 Fibre Channel Interface Configuration

Switch	Interface	Mode Admin	Mode Oper	Port VSAN	Dynamic VSAN	Description	Speed Admin	Speed Oper	Rate Mode	Status Service	Status Admin	Status Oper	FailureCause	Was Enabled	LastChange
172.22.46.180	Fc1/1	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/2	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/3	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/4	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/5	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/6	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/7	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/8	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/9	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/10	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/11	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/12	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/13	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc1/14	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	adminDown	False	n/a
172.22.46.180	Fc3/1	FX	auto	1	n/a		auto	n/a	shared	in	down	down	adminDown	False	n/a
172.22.46.180	Fc3/2	FX	auto	1	n/a		auto	n/a	shared	in	down	down	adminDown	False	n/a

Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but not limited to) modes, frame encapsulation, states, SFPs, and speeds.

This section includes the following topics:

- [32-Port Switching Module Configuration Guidelines, page 20-2](#)
- [About Interface Modes, page 20-3](#)
- [Configuring Interface Modes, page 20-12](#)
- [Configuring Administrative Speeds, page 20-12](#)
- [About Interface Descriptions, page 20-13](#)
- [About Frame Encapsulation, page 20-13](#)
- [About Receive Data Field Size, page 20-13](#)
- [Configuring Receive Data Field Size, page 20-13](#)
- [Identifying the Beacon LEDs, page 20-14](#)
- [About Beacon Mode, page 20-15](#)
- [About Bit Error Thresholds, page 20-15](#)
- [About SFP Transmitter Types, page 20-16](#)
- [Displaying SFP Transmitter Types, page 20-16](#)

32-Port Switching Module Configuration Guidelines

The 32-port switching module guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

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

Send documentation comments to mdsfeedback-doc@cisco.com

- 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 and 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 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 left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

About 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 20-2](#)). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 20-2 Cisco MDS 9000 Family Switch Port Modes



Note

Interfaces are created in VSAN 1 by default. See [Chapter 1, “Configuring and Managing VSANs.”](#)

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.

A brief description of each interface mode follows.

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 23, “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.

Send documentation comments to mdsfeedback-doc@cisco.com

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 a 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.

For more details about NP ports and NPV, see [Chapter 21, “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 [“About TL Port ALPA Caches”](#) section on page 20-19).



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.

Send documentation comments to mdsfeedback-doc@cisco.com

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 (see [Chapter 24, “Configuring Trunking”](#)). TE ports support class 2, class 3, and class F service.

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 merely 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 [Chapter 52, “Monitoring Network Traffic Using SPAN”](#)).

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 [“Configuring SPAN” section on page 52-6](#)).



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.

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 44-1 on page 44-2 depicts a typical SAN extension over an IP network.

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 [Chapter 44](#), “Configuring IP Storage”).

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, or TE 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 24](#), “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.

About 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 20-1](#).

Table 20-1 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 20-2](#).

Table 20-2 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Table 20-2 **Operational States (continued)**

Operational State	Description
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE mode.

Reason Codes

Reason codes are dependent on the operational state of the interface as described in [Table 20-3](#).

Table 20-3 **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 Table 20-4 .



Note

Only some of the reason codes are listed in [Table 20-4](#).

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 20-4](#).

Send documentation comments to mdsfeedback-doc@cisco.com

Table 20-4 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 SAN-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"> • Configuration failure. • Incompatible buffer-to-buffer credit configuration. 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.	

Send documentation comments to mdsfeedback-doc@cisco.com

Table 20-4 Reason Codes for Nonoperational States (continued)

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. 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.	

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.

Send documentation comments to mdsfeedback-doc@cisco.com

Graceful Shutdown

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

The Cisco SAN-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 SAN-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 SAN-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 (see [“In-Order Delivery”](#) section on page 25-15).
- If the Min_LS_interval interval is higher than 10 seconds ().



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.

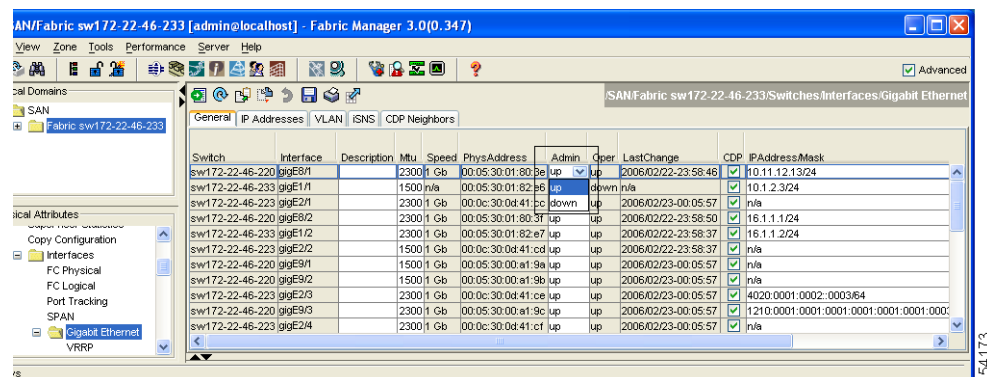
Setting the Interface Administrative State

To disable or enable an interface using Fabric Manager, follow these steps:

- Step 1** Either expand **Switches > Interfaces** and then select **Gigabit Ethernet** or expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.
- Step 2** Click the **General** tab.
- Step 3** Click **Admin**.

You see the drop-down box shown in [Figure 20-3](#).

Figure 20-3 Changing the Administrative Status of a Switch



- Step 4** Set the status to down (disable) or up (enable).

Send documentation comments to mdsfeedback-doc@cisco.com

- Step 5** Optionally, set other configuration parameters using the other tabs.
- Step 6** Click **Apply Changes**.
-

Configuring Interface Modes



Note To ensure that ports that are part of ISLs do not get changed to port mode F, configure the ports in port mode E, rather than in Auto mode.

To configure the interface Fabric Manager, follow these steps:

- Step 1** Expand **Switches > Interfaces** then select **FC 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** Optionally, set other configuration parameters using the other tabs.
- Step 5** Click **Apply Changes** icon.
-

Configuring Administrative Speeds

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



Caution Changing the administrative speed is a disruptive operation.

To configure administrative speed of the interface using Fabric Manager, follow these steps:

- Step 1** Expand **Switches > Interfaces** then select **FC 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, or **auto** (default).
- Step 4** Click **Apply Changes**.
-

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.

Send documentation comments to mdsfeedback-doc@cisco.com

Autosensing

Autosensing speed is enabled on all 4-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. 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 Fibre Channel switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 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.



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.

About Interface Descriptions

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

About Frame Encapsulation

You can set the frame format to EISL for all frames transmitted by the interface in SD port mode. If you set the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format, irrespective of the SPAN source(s). See the “[Monitoring Network Traffic Using SPAN](#)” section on [page 52-1](#).

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to configure frame encapsulation on an interface.

About Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Configuring Receive Data Field Size

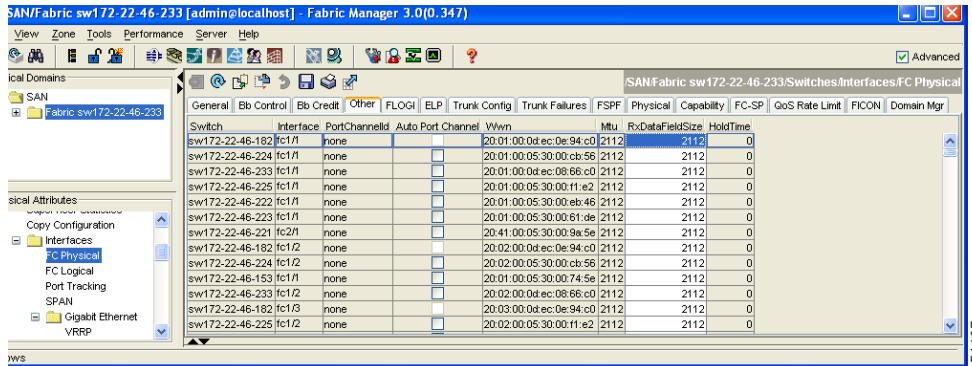
You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure the receive data field size using Fabric Manager, follow these steps:

-
- Step 1** Expand **Switches > Interfaces** and then select **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **Other** tab and set the RxDataFieldSize field (see [Figure 20-4](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 20-4 Changing Rx Data Size




- Step 3** Optionally, set other configuration parameters using the other tabs.
- Step 4** Click **Apply Changes**.

Identifying the Beacon LEDs

Figure 20-5 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 20-5 Cisco MDS 9000 Family Switch Interface Modes



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

1. See the “Identifying Module LEDs” section on page 19-9.
2. See the “About Speed LEDs” section on page 20-14.
3. Refer to the Cisco MDS 9000 Family hardware installation guide for your platform.

About 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:

Send documentation comments to mdsfeedback-doc@cisco.com

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

About 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.

Configuring Beacon Mode

To enable beacon mode for a specified interface or range of interfaces using Fabric Manager, follow these steps:

-
- Step 1** Expand **Switches > Interfaces** and then select **Gigabit Ethernet**.
You see the interface configuration in the Information pane.
- Step 2** Enable the Beacon Mode option for the selected switch.
- Step 3** Click **Apply Changes**.
-



Note

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.

About 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 to reenable the interface.

Send documentation comments to mdsfeedback-doc@cisco.com

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

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to disable the bit error threshold for an interface.



Note

Regardless of disabling the switchport ignore bit-error threshold for an interface, the switch generates a syslog message when bit error threshold events are detected.

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 CLI Configuration Guide* to configure switch port attributes.

About SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed. [Table 20-5](#) defines the acronyms used for SFPs (see the “[Displaying SFP Transmitter Types](#)” section on page 20-16).

Table 20-5 SFP Transmitter Acronym Definitions

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

Displaying SFP Transmitter Types

To show the SFP types for an interface using Fabric Manager, follow these steps:

Send documentation comments to mdsfeedback-doc@cisco.com

-
- Step 1** Expand **Switches > Interfaces** and then select **FC 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.
-

About Gathering Interface Statistics

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

Gathering Interface Statistics

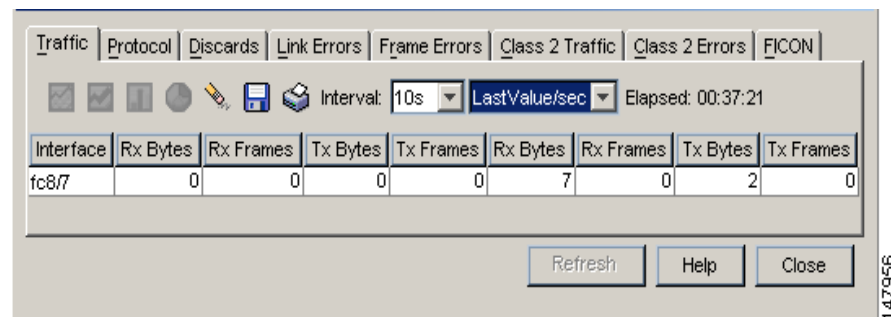


Note In Fabric Manager, you can collect interface statistics by expanding **Switches > ISLs** and selecting **Statistics** from the Physical Attributes pane.

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

-
- Step 1** Right-click an interface and select **Monitor**.
You see the Interface Monitor dialog box.
- 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 shown in [Figure 20-6](#) to view those related statistics.

Figure 20-6 Device Manager Interface Monitor Dialog Box



- Step 4** Optionally, click the **Pencil** icon to reset the cumulative counters.
- Step 5** Optionally, 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.
-

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

TL Ports for Private Loops

Private loops require setting the interface mode to TL. This section describes TL ports and includes the following sections:

- [About TL Ports, page 20-18](#)
- [Configuring TL Ports, page 20-19](#)
- [About TL Port ALPA Caches, page 20-19](#)

About TL Ports

TL port mode is not supported on the following:

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

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 [“About Interface Modes” section on page 20-3](#).

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxied 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.


[Table 20-6](#) lists the TL port translations supported in Cisco MDS 9000 Family switches. [Figure 20-7](#) shows examples of TL port translation support.

Table 20-6 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

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 20-7 TL Port Translation Support Examples



Configuring TL Ports

To configure the TL interface mode using Fabric Manager, follow these steps:

-
- Step 1** Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.
 - Step 2** Choose the **General** tab and click **Mode Admin**.
 - Step 3** Set the Mode Admin drop-down menu to **TL**.
 - Step 4** Optionally, set other configuration parameters using the other tabs.
 - Step 5** Click **Apply Changes**.
-

About 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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 SAN-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 SAN-OS software discards an inactive cache entry (if available) to make space for the new entry. See the “TL Port” section on page 20-5 for more information on TL ports.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to manage the TL Port ALPA cache.

Buffer Credits

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination. This section describes the different buffer credits available on the Cisco MDS Family switches and includes the following topics:

- [About Buffer-to-Buffer Credits, page 20-20](#)
- [Configuring Buffer-to-Buffer Credits, page 20-20](#)
- [About Performance Buffers, page 20-21](#)
- [Configuring Performance Buffers, page 20-21](#)
- [About Extended BB_credits, page 20-21](#)
- [Configuring Extended BB_credits, page 20-23](#)

About Buffer-to-Buffer Credits

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

The receive BB_credit (`fcrxbbcredit`) value may be configured for each FC interface. In most cases, you do not need to modify the default configuration.



Note

The receive BB_credit values depend on the module type and the port mode, as follows:

- For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
- For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.
- For Generation 2 switching modules, see the “Buffer Pools” section on page 22-8.



Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Configuring Buffer-to-Buffer Credits

To configure BB_credits for a Fibre Channel interface using Fabric Manager, follow these steps:

Send documentation comments to mdsfeedback-doc@cisco.com

-
- Step 1** Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.
 - Step 2** Choose the **Bb Credit** tab.
You see the buffer credits.
 - Step 3** Set any of the buffer-to-buffer credits for an interface.
 - Step 4** Click **Apply Changes**.
-

About Performance Buffers



Note

Performance buffers are 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.

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

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

The default performance buffer value is 0. If you set the performance buffer value to 0, the built-in algorithm is used. If you do not specify the performance buffer value, 0 is automatically used.

Configuring Performance Buffers

To configure performance buffers for a Fibre Channel interface using Fabric Manager, follow these steps:

-
- Step 1** Expand **Switches > Interfaces** and then select **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **BB Credit** tab.
You see performance buffer information in the columns Perf Bufs Admin and Perf Bufs Oper.
 - Step 3** Set the performance buffers for an interface.
 - Step 4** Click **Apply Changes**.
-

About Extended `BB_credits`

You can use the extended `BB_credits` flow control mechanism in addition to `BB_credits` for long haul links.

This section includes the following topics:

Send documentation comments to mdsfeedback-doc@cisco.com

- [Extended BB_credits on Generation 1 Switching Modules, page 20-22](#)
- [Extended BB_credits on Generation 2 Switching Modules, page 20-23](#)


Extended BB_credits on Generation 1 Switching Modules

The BB_credits feature allows you to configure up to 255 receive buffers on Generation 1 switching modules. To facilitate BB_credits for long haul links, you can configure up to 3,500 receive BB_credits on a Fibre Channel port on a Generation 1 switching module.

To use this feature on Generation 1 switching modules, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license (see [Chapter 10, “Obtaining and Installing Licenses”](#)).
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see [Figure 20-8](#)).

Figure 20-8 Port Group Support for the Extended BB_Credits Feature



The port groups that support extended credit configurations are as follows.

- Any one port in ports 1 to 4 (identified as Group 1 in [Figure 20-2](#)).
- Any one port in ports 5 to 8 (identified as Group 2 in [Figure 20-2](#)).
- Any one port in ports 9 to 12 (identified as Group 3 in [Figure 20-2](#)).



Note The last two Fibre Channel ports (port 13 and port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature (see [Figure 20-2](#)).

- Explicitly enable this feature in the required Cisco MDS switch.
- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
 - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.



Note The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

- If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.
- Be aware that changing the BB_credit value results in the port being disabled and then reenabled.
- Disable (explicitly) this feature if you need to nondisruptively downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

Extended BB_credits on Generation 2 Switching Modules

To use this feature on Generation 2 switching modules, you must meet the following requirements:

- You see the interface configuration in the Information pane (see [Figure 20-9](#)).
- Obtain the Enterprise package (ENTERPRISE_PKG) license (see [Chapter 10, “Obtaining and Installing Licenses”](#)).
- Configure this feature in any port on a Generation 2 switch module. See the [Extended BB_Credits, page 22-15](#) for more information on extended BB_credits on Generation 2 switching modules.

**Note**

Extended BB_credits are not supported on the Cisco MDS 9124 Fabric Switch.

Configuring Extended BB_credits

To configure extended BB_credits for an MDS-14/2 interface, for a Generation 2 switching module interface, or for an interface in a Cisco MDS 9216i switch using Fabric Manager, follow these steps:

- Step 1** Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.
- Step 2** Choose the **BB Credit** tab.
- Step 3** In the **Extended** column, set the extended BB_credits for the selected interface.
- Step 4** Click **Apply Changes**.

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.

This section describes the management interfaces and includes the following topics:

- [About Management Interfaces, page 20-23](#)
- [Configuring Management Interfaces, page 20-24](#)

About Management Interfaces

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

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Configuring Management Interfaces

To configure the mgmt0 Ethernet interface using Fabric Manager, follow these steps:

-
- Step 1** Select a VSAN in the Logical Domains pane.
 - Step 2** Expand **Switches > Interfaces** and then select **Management**.
You see the interface configuration in the Information pane.
 - Step 3** Click the **General** tab.
 - Step 4** Set the IP Address/Mask field.
 - Step 5** Set Admin to **up**.
 - Step 6** Optionally, set other configuration parameters using the other tabs.
 - Step 7** Click **Apply Changes**.
-

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 nonexistent VSANs.

This section describes VSAN interfaces and includes the following topics:

- [About VSAN Interfaces, page 20-24](#)
- [Creating VSAN Interfaces, page 20-25](#)

About VSAN Interfaces

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.

Send documentation comments to mdsfeedback-doc@cisco.com



Tip

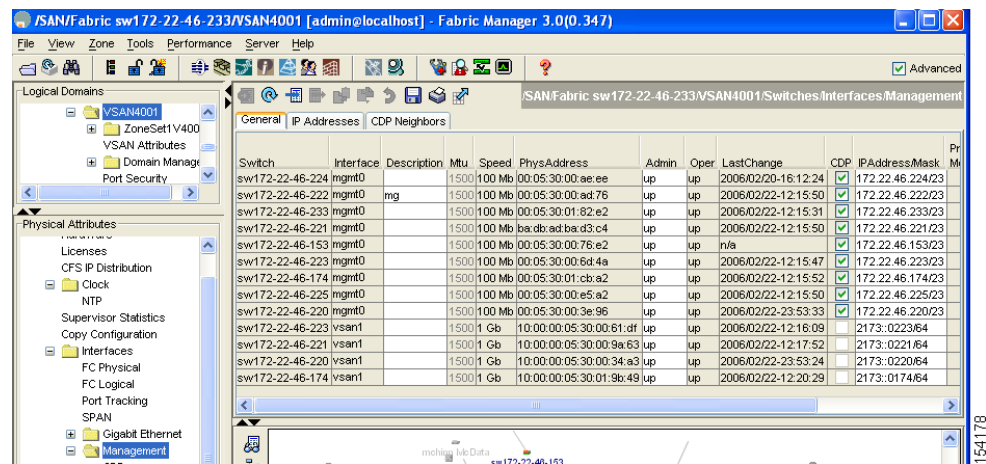
After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature (see [Chapter 43, “Configuring IP Services”](#)).

Creating VSAN Interfaces

To create a VSAN interface using Fabric Manager, follow these steps:

- Step 1** Expand **Switches > Interfaces** and then select **Management**.


Figure 20-9 General Management Tab



- Step 2** Click **Create Row**.

You see the Create Interface dialog box (see [Figure 20-10](#)).

Figure 20-10 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.

Send documentation comments to mdsfeedback-doc@cisco.com

- Step 5** Click **Create** to create the VSAN interface or click **Close** to close the dialog box without creating the VSAN interface.
-

Default Settings

Table 20-7 lists the default settings for interface parameters.

Table 20-7 *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)
Trunk-allowed VSANs	1 to 4093
Interface VSAN	Default VSAN (1)
Beacon mode	Off (disabled)
EISL encapsulation	Disabled
Data field size	2112 bytes