



Configuring IP Storage Services

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note

FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches, Cisco MDS 9500 Directors, and Cisco MDS 9700 Directors.

The Cisco MDS 24/10 port SAN Extension Module for MDS 9700, and the 18/4 Multiprotocol Services (MSM-18/4) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MSM-18/4 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series, and the Cisco MDS 24/10 port SAN Extension Module can be used in any of the Cisco MDS 9700 series switches.

This chapter includes the following sections:

- [Feature Information, page 6-ccxlvi](#)
- [IP Storage Modules, page 6-ccxlvi](#)
- [Supported Hardware, page 6-ccli](#)
- [Configuring Gigabit Ethernet Interfaces for IPv4, page 6-ccli](#)
- [IPS Module Core Dumps, page 6-ccli](#)
- [Configuring Gigabit Ethernet High Availability, page 6-cclvi](#)
- [Configuring CDP, page 6-cclviii](#)
- [Changing Link Speed on IP Storage Interfaces, page 6-cclix](#)
- [Default Settings for IP Storage Services Parameters, page 6-cclxv](#)

Feature Information

This section briefly describes the new and updated features for releases, starting from Cisco MDS NX-OS Release 6.2(13).

Table 6-1 Feature Information Table

Feature	Release	Description
Changing Link Speed on Cisco MDS 24/10 port SAN Extension Module, page 6-cclxi	7.3(0)DY(1)	This feature enables users to change the link speed on IP Storage interfaces between 1 Gbps and 10 Gbps on the Cisco MDS 24/10 port SAN Extension Module.
Changing Link Speed on Cisco MDS 9250i Multiservice Fabric Switch, page 6-cclix	6.2(13)	This feature enables users to change the link speed on IP Storage interfaces between 1 Gbps and 10 Gbps on the Cisco MDS 9250i Multiservice Fabric Switch.


IP Storage Modules

The IP Storage services module (IPS module) and the MSM-18/4 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management.

Gigabit Ethernet ports in these modules can be configured to support the FCIP protocol, the iSCSI protocol, or both protocols simultaneously:


- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 6-1](#) shows how the IPS module is used in different FCIP scenarios.

Figure 6-1 FCIP Scenarios



- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. Figure 6-2 depicts the iSCSI scenarios in which the IPS module is used.

Figure 6-2 iSCSI Scenarios



Module Status Verification

To verify the status of the module using Fabric Manager, follow these steps:

-
- Step 1** Select a switch in the Fabric pane.
 - Step 2** Open the **Switches** folder and select **Hardware** in the Physical Attributes pane. You see the status for all modules in the switch in the Information pane.
-

After inserting the module, verify the status of the module using the **show module** command:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
-----
1    48     2/4/8/10/16 Gbps Advanced FC Module DS-X9448-768K9      ok
4    34     1/10/40G IPS,2/4/8/10/16G FC Module DS-X9334-K9        ok
5    0      Supervisor Module-3        DS-X97-SF1-K9      active *
6    0      Supervisor Module-3        DS-X97-SF1-K9      ha-standby
9    34     1/10/40G IPS,2/4/8/10/16G FC Module DS-X9334-K9        ok
10   48     2/4/8/10/16 Gbps Advanced FC Module DS-X9448-768K9      ok

Mod  Sw          Hw
-----
1    7.3(0)DY(1)  0.301
4    7.3(0)DY(1)  0.304
5    7.3(0)DY(1)  1.1
6    7.3(0)DY(1)  1.1
9    7.3(0)DY(1)  0.402
10   7.3(0)DY(1)  1.1

Mod  MAC-Address(es)                Serial-Num
-----

```

```

1 54-7f-ee-d7-bc-70 to 54-7f-ee-d7-bc-73 JAE164302O2
4 00-8e-73-39-39-e0 to 00-8e-73-39-39-ef JAE200806T0
5 3c-0e-23-c4-71-86 to 3c-0e-23-c4-71-98 JAE17510BAE
6 3c-0e-23-c4-74-f0 to 3c-0e-23-c4-75-02 JAE17510BC1
9 9c-57-ad-2a-7d-e0 to 9c-57-ad-2a-7d-ef JAE201100XU
10 1c-df-0f-79-42-e8 to 1c-df-0f-79-42-eb JAE172009XM

```

```
Mod Online Diag Status
```

```

--- -----
1 Pass
4 Pass
5 Pass
6 Pass
9 Pass
10 Pass

```

Xbar	Ports	Module-Type	Model	Status
1	0	Fabric Module 1	DS-X9710-FAB1	ok
2	0	Fabric Module 1	DS-X9710-FAB1	ok
3	0	Fabric Module 1	DS-X9710-FAB1	ok

Xbar	Sw	Hw
1	NA	1.2
2	NA	1.2
3	NA	1.2

Xbar	MAC-Address(es)	Serial-Num
1	NA	JAE18070AR0
2	NA	JAE180602PF
3	NA	JAE18070ANJ

* this terminal session

IPS Module Upgrade

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.



Caution

A software upgrade is only disruptive for the IPS module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

Cisco MDS 9250i Switch

The Cisco MDS 9250i switches have 40 Fibre Channel ports (nondisruptive upgrade) and two IP Storage ports (disruptive upgrade). Cisco MDS 9250i switches use a rolling upgrade install mechanism for the two IP Storage ports where each module in a given switch can only be upgraded in sequence.



Caution

A software upgrade is only partially disruptive for the Cisco MDS 9250i switch. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

24/10 port SAN Extension Module

The 24/10 port SAN Extension Modules have 24 Fibre Channel ports (nondisruptive upgrade) and ten IP Storage ports (disruptive upgrade). 24/10 port SAN Extension Modules use a rolling upgrade install mechanism for the ten IP Storage ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each 24/10 port SAN Extension Module in a switch requires a 5-minute delay before the next module is upgraded.



Caution

A software upgrade is only partially disruptive for the 24/10 port SAN Extension Module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

MSM-18/4 Module Upgrade

The MSM-18/4 modules have 18 Fibre Channel ports (nondisruptive upgrade) and four Gigabit Ethernet ports (disruptive upgrade). MSM-18/4 modules use a rolling upgrade install mechanism for the four Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MSM-18/4 module in a switch requires a 5-minute delay before the next module is upgraded.



Caution

A software upgrade is only partially disruptive for the MSM-18/4 module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

Supported Hardware

You can configure the FCIP and iSCSI features using one or more of the following hardware:

- Cisco MDS 24/10 port SAN Extension Module on Cisco MDS 9700 Series Director switches.
- MSM-18/4 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).
- Cisco MDS 9250i Multiservice Fabric Switch.

Configuring Gigabit Ethernet Interfaces for IPv4

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module, connectivity is provided in the form of Gigabit Ethernet interfaces on Cisco MDS 9500 series switches, and in the form of IP storage ports on Cisco MDS 9250i switches and Cisco MDS 9700 series switches with 24/10 port SAN Extension modules that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.

**Note**

To configure IPv6 on a Gigabit Ethernet interface, see the *Cisco Fabric Manager Security Configuration Guide*. For information about configuring FCIP, see [Chapter 2, “Configuring Fibre Channel over IP.”](#) For information about configuring iSCSI, see [Chapter 4, “Configuring Internet Small Computer Systems Interface.”](#)


**Tip**

Gigabit Ethernet ports on any IPS module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration

[Figure 6-3](#) shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 6-3 Gigabit Ethernet IPv4 Configuration Example

**Note**

The port on the Ethernet switch to which the Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the ethernet switch) should be disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in Cisco IOS or the **set port host** command in Catalyst OS.

IPS Module Core Dumps

IPS core dumps are different from the system’s kernel core dumps for other modules. When the IPS module’s operating system (OS) unexpectedly resets, it is useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

In Cisco MDS 9700 Series Switches with 24/10 port SAN Extension Modules, each partial core dump consists of five parts (five files). All five files are saved in the active supervisor module.

Use the **show cores** command to list these files.

- Full core dumps—Each full core dump of Cisco MDS 9250i Switches and SSN-16 modules consists of 64 parts (64 files), and each full core dump of Cisco MDS 9700 Series Switches with 24/10 port SAN Extension Modules consists of 67 parts (67 files). The IPS core dump for MSM-18/4 modules consists of 32 parts. This dump cannot be saved on the supervisor module because of its large space requirement. They are copied directly to an external TFTP server.

Use the **system cores tftp** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

To configure IPS core dumps on the IPS module, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# ips core dump full ips core dump full' successfully set for module 9	Configures a dump of the full core generation for all IPS modules in the switch.
	switch(config)# no ips core dump full ips core dump partial' successfully set for module 9	Configures a dump of the partial core (default) generation for the IPS module in slot 9.

To configure the Gigabit Ethernet interface for the scenario in [Figure 6-3](#), follow these steps:

-
- Step 1** From Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.
- From Device Manager, right-click the Gigabit Ethernet port that you want to configure and choose **Configure...** You see the Gigabit Ethernet configuration dialog box.
- Step 2** Click the **General** tab in Fabric Manager, or click the **GigE** tab in Device Manager to display the general configuration options for the interface.
- Step 3** Set the description and MTU value for the interface. The valid value for the MTU field can be a number in the range from 576 to 9000.
- Step 4** Set **Admin** up or down and check the **CDP** check box if you want this interface to participate in CDP.
- Step 5** Set **IpAddress/Mask** with the IP address and subnet mask for this interface.
- Step 6** From Fabric Manager, click the **Apply Changes** icon to save these changes, or click the **Undo Changes** icon to discard changes.
- From Device Manager, click **Apply** to save these changes, or click **Close** to discard changes and close the Gigabit Ethernet configuration dialog box.
-

Configuring Interface Descriptions

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the switch port description for any interface.

Configuring Beacon Mode

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the beacon mode for any interface.

Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note

The minimum MTU size is 576 bytes.



Tip

MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name: *slot-number / port-numberVLAN-ID*.

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 6-2](#)).

Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	

**Note**

The configuration requirements in [Table 6-2](#) also apply to Ethernet PortChannels.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.

**Note**

If the connection fails, verify the following, and ping the IP host again:

- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the `up` state.

Gigabit Ethernet IPv4-ACL Guidelines

**Tip**

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established**, **precedence**, and **fragments** options are ignored when you apply IPv4-ACLs (containing these options) to Gigabit Ethernet interfaces.
 - If an IPv4-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B, and an IPv4-ACL specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.


Configuring Gigabit Ethernet High Availability

Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see [Figure 6-4](#)).

Figure 6-4 VRRP Scenario



In Figure 6-4, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MSM-18/4 module
- Interfaces across IPS modules or MSM-18/4 modules in one switch
- Interfaces across IPS modules or MSM-18/4 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces

**Note**

You can configure no more than seven VRRP groups, both IPv4 and IPv6, on a Gigabit Ethernet interface, including the main interface and all subinterfaces.

Configuring VRRP for Gigabit Ethernet Interfaces

To configure VRRP for Gigabit Ethernet interfaces using IPv4, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.10 255.255.255.0	Assigns the IPv4 address (10.1.1.10) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp 100 switch(config-if-vrrp)	Creates VR ID 100.
Step 6	switch(config-if-vrrp) # address 10.1.1.100	Configures the virtual IPv4 address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IPv4 address must be in the same subnet as the IPv4 address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IPv4 address.
Step 7	switch(config-if-vrrp) # priority 10	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	switch(config-if-vrrp) # no shutdown	Enables the VRRP protocol on the selected interface.

To configure VRRP for Gigabit Ethernet interfaces using IPv6, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ipv6 address 2001:0db8:800:200c::417a/64	Assigns the IPv6 address for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp ipv6 100 switch(config-if-vrrp-ipv6)	Creates VR ID 100.
Step 6	switch(config-if-vrrp-ipv6) # address 2001:0db8:800:200c::417a	Assigns single primary link-local IPv6 address or one of the multiple secondary IPv6 addresses. Note If this IPv6 address is the same as the physical IPv6 address, this switch is automatically the owner of this IPv6 address.
Step 7	switch(config-if-vrrp-ipv6) # priority 10	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	switch(config-if-vrrp-ipv6) # no shutdown	Enables the VRRP protocol on the selected interface.

**Note**

If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

**Note**

The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual IPv4 IP address is also the IPv4 IP address for the interface, then preemption is implicitly applied.

**Note**

If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS module or MSM-18/4 module.

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

Changing Link Speed on IP Storage Interfaces

Changing Link Speed on Cisco MDS 9250i Multiservice Fabric Switch

The Cisco MDS 9250i Multiservice Fabric Switch has two IP storage interfaces that support 1 Gbps and 10 Gbps link speeds. By default, IP storage interfaces are configured at 10 Gbps link speed.



Note

Switching between different link speeds is supported on Cisco 10 Gbps IP storage platforms starting from Cisco MDS NX-OS Release 6.2(13). An ISSD to a release earlier than Cisco MDS NX-OS Release 6.2(13) when any of the IP storage ports are configured at 1 Gbps, is disallowed. Reconfigure such ports back to the default link speed of 10 Gbps before attempting such a downgrade.

To configure 1 Gbps link speed on an IP storage interface, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface IPStorage <i>slot-number/port-number-range</i> switch(config-if)#	Enters IPStorage interface configuration mode.
Step 3	switch(config-if)# shutdown	Administratively disables the interface and stops traffic through the interface.
Step 4	switch(config-if)# switchport speed 1000 ¹ This speed change will disrupt FCIP/iSCSI traffic for 5 mins on all IPStorage ports. If FCIP tunnels are configured please make sure max-bw <= 1000 Mbps and tcp-connections set to 2. Do you want to continue(y/n) ? [n]	Sets the link speed of the interface and all subinterfaces to 1000 Mbps (1 Gbps). Note This command causes all IP storage ports on the selected FCIP engine to be reset. This may cause traffic disruption for up to 5 minutes. By default, <i>n</i> is selected. Press Enter to abort the command. Enter <i>y</i> and press Enter to continue. ²
Step 5	switch(config-if)# no shutdown	Administratively enables the interface.
Step 6	switch(config-if)# end switch#	Exits IPStorage interface configuration mode and returns to privileged EXEC mode.
Step 7	switch# show ips status	Displays the operational speed of the IP storage port.

1. Configuring the link speed of an interface generates the following port software failure syslog message:

```
%IF_DOWN_SOFTWARE_FAILURE: %$VSAN 1%$ Interface fcip is down (Port software failure)
```

2. If the conditions specified in the warning message are not met, the configured link speed is still applied. However, issues such as packet drops, retransmissions, and FCIP tunnel flaps may occur.

To configure 10 Gbps link speed on an IP storage interface, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface IPStorage <i>slot-number/port-number-range</i> switch(config-if)#	Enters IPStorage interface configuration mode.
Step 3	switch(config-if)# shutdown	Administratively disables the interface and stops traffic through the interface.
Step 4	switch(config-if)# switchport speed 10000 ¹ "This speed change will disrupt FCIP/iSCSI traffic for 5 mins on all IPStorage ports. Do you want to continue(y/n) ? [n]"	Sets the link speed of the interface and all subinterfaces to 10000 Mbps (10 Gbps). Note This command causes all IP storage ports on the selected FCIP engine to be reset. This may cause traffic disruption for up to 5 minutes. By default, <i>n</i> is selected. Press Enter to abort the command. Enter <i>y</i> and press Enter to continue. ²
Step 5	switch(config-if)# no shutdown	Administratively enables the interface.
Step 6	switch(config-if)# end switch#	Exits IPStorage interface configuration mode and returns to privileged EXEC mode.
Step 7	switch# show ips status	Displays the operational speed of the IP storage port.

1. Configuring the link speed of an interface generates the following port software failure syslog message:
%IF_DOWN_SOFTWARE_FAILURE: %\$VSAN 1%\$ Interface fcip is down (Port software failure)
2. If the conditions specified in the warning message are not met, the configured link speed is still applied. However, issues such as packet drops, retransmissions, and FCIP tunnel flaps may occur.

If there is a mismatch between the configured link speed and the small form-factor pluggable (SFP) speed capabilities, the port goes into an Error Disabled state and a corresponding syslog message is logged. In such a scenario, either the configured link speed or the SFP should be changed. If the link speed is changed, even if the port is already enabled, the **shutdown** and **no shutdown** commands must be explicitly issued for the change to be applied.

For more information about supported 1 Gbps SFPs for a Cisco MDS 9250i Multiservice Fabric Switch, see the [Cisco MDS 9000 Family Pluggable Transceivers Data Sheet](#).

For information about configuring FCIP tunnels with IP storage interfaces at 1 Gbps speed, see the [Configuring FCIP](#) chapter.

Changing Link Speed on Cisco MDS 24/10 port SAN Extension Module

To configure 1 Gbps link speed on an IP storage interface, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface IPStorage <i>slot-number/port-number-range</i> switch(config-if)#	Enters IPStorage interface configuration mode. Note The values for <i>port-number-range</i> can be 1-4 and 5-8 only.
Step 3	switch(config-if)# 1G-speed-mode This speed change will disrupt FCIP/iSCSI traffic for 60 seconds on selected IPStorage ports.If FCIP tunnels are configured please make sure max-bw <= 1000 Mbps and tcp-connections set to 2. Do you wish to continue(y/n)? [n]	Sets the link speed of the interface and all subinterfaces to 1000 Mbps (1 Gbps) and administratively enables the interface.
Step 4	switch(config-if)# end switch#	Exits IPStorage interface configuration mode and returns to privileged EXEC mode.
Step 5	switch# show ips status	Displays the operational speed of the IP storage port.

To configure 10 Gbps link speed on an IP storage interface, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface IPStorage <i>slot-number/port-number-range</i> switch(config-if)#	Enters IPStorage interface configuration mode. Note The values for <i>port-number-range</i> can be 1-4 and 5-8 only.
Step 3	switch(config-if)# 10G-speed-mode This speed change will disrupt FCIP/iSCSI traffic for 60 seconds on select IPStorage ports. Do you wish to continue(y/n)? [n]	Sets the link speed of the interface and all subinterfaces to 10000 Mbps (10 Gbps) and administratively enables the interface.
Step 4	switch(config-if)# end switch#	Exits IPStorage interface configuration mode and returns to privileged EXEC mode.
Step 5	switch# show ips status	Displays the operational speed of the IP storage port.

Displaying Statistics

This section provides examples to verify Gigabit Ethernet and TCP/IP statistics on the IP storage ports.

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface gigabitethernet** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 6-1](#) and [Example 6-2](#).

Example 6-1 Displaying the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
GigabitEthernet8/1 is up
  Hardware is GigabitEthernet, address is 0005.3000.a98e
  Internet address is 10.1.3.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  3343 packets input, 406582 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  8 packets output, 336 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Example 6-2 Displaying the Gigabit Ethernet Subinterface

```
switch# show interface gigabitethernet 4/2.100
GigabitEthernet4/2.100 is up
  Hardware is GigabitEthernet, address is 0005.3000.abcb
  Internet address is 10.1.2.100/24
  MTU 1500 bytes
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 packets input, 0 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  1 packets output, 46 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Use the **show interface IPStorage** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 6-3](#).

Example 6-3 Displaying the IP Storage Interface

```
switch# show interface ipStorage 4/1
IPStorage4/1 is up
  Hardware is IPStorage, address is 008e.7339.39e7
  Internet address is 10.197.141.81/24
  MTU 2500 bytes
  Port mode is IPS
  Speed is 10 Gbps
  Beacon is turned off
  Auto-Negotiation is turned on
  5 minutes input rate 77012744 bits/sec, 9626593 bytes/sec, 112755 frames/sec
  5 minutes output rate 2762915176 bits/sec, 345364397 bytes/sec, 175258 frames/sec
  71187036 packets input, 6078261484 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  110617842 packets output, 217860230652 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

```
switch# show interface ipStorage 5/1-10 brief
```

Interface	Status	IP Address	Speed	MTU
IPStorage5/1	up	1.1.1.1/24	10 Gbps	2500
IPStorage5/2	up	2.2.2.2/24	10 Gbps	2500
IPStorage5/3	up	3.3.3.3/24	10 Gbps	2500
IPStorage5/4	up	4.4.4.4/24	10 Gbps	2500
IPStorage5/5	up	6811::3456/64	1 Gbps	2300
IPStorage5/6	up	9.9.9.1/24	1 Gbps	2500
IPStorage5/7	up	7.7.7.1/24	1 Gbps	2500
IPStorage5/8	up	8.8.8.1/24	1 Gbps	2500
IPStorage5/9	outOfServc	--	auto	1500
IPStorage5/10	outOfServc	--	auto	1500



Note In Cisco MDS NX-OS Release 7.3(0)DY(1), 40GE IP Storage interfaces are not supported.

Displaying Ethernet MAC Statistics

The `show ips stats mac interface gigabitethernet` command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 6-4](#).



Note

Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

Example 6-4 Displaying Ethernet MAC Statistics

```
switch# show ips stats mac interface gigabitethernet 8/1
Ethernet MAC statistics for port GigabitEthernet8/1
Hardware Transmit Counters
  237 frame 43564 bytes
  0 collisions, 0 late collisions, 0 excess collisions
  0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
Hardware Receive Counters
  427916 bytes, 3464 frames, 0 multicasts, 3275 broadcasts
  0 bad, 0 runt, 0 CRC error, 0 length error
  0 code error, 0 align error, 0 oversize error
Software Counters
  3429 received frames, 237 transmit frames
  0 frames soft queued, 0 current queue, 0 max queue
  0 dropped, 0 low memory
```

Displaying TCP Statistics

Use the `show ips stats tcp interface gigabitethernet` to display and verify TCP statistics. This command takes the main Ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The `detail` option shows all information maintained by the interface. See [Example 6-5](#) and [Example 6-6](#).

Example 6-5 Displaying TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
Connection Stats
```

```

    0 active openings, 3 accepts
    0 failed attempts, 12 reset received, 3 established
Segment stats
    163 received, 355 sent, 0 retransmitted
    0 bad segments received, 0 reset sent
TCP Active Connections
  Local Address          Remote Address          State      Send-Q  Recv-Q
  0.0.0.0:3260          0.0.0.0:0              LISTEN    0       0

```

Example 6-6 Displaying Detailed TCP Statistics

```

switch# show ips stats tcp interface gigabitethernet 4/1 detail
TCP Statistics for port GigabitEthernet4/1
TCP send stats
  355 segments, 37760 bytes
  222 data, 130 ack only packets
  3 control (SYN/FIN/RST), 0 probes, 0 window updates
  0 segments retransmitted, 0 bytes
  0 retransmitted while on ethernet send queue, 0 packets split
  0 delayed acks sent
TCP receive stats
  163 segments, 114 data packets in sequence, 6512 bytes in sequence
  0 predicted ack, 10 predicted data
  0 bad checksum, 0 multi/broadcast, 0 bad offset
  0 no memory drops, 0 short segments
  0 duplicate bytes, 0 duplicate packets
  0 partial duplicate bytes, 0 partial duplicate packets
  0 out-of-order bytes, 1 out-of-order packets
  0 packet after window, 0 bytes after window
  0 packets after close
  121 acks, 37764 ack bytes, 0 ack toomuch, 4 duplicate acks
  0 ack packets left of snd_una, 0 non-4 byte aligned packets
  8 window updates, 0 window probe
  30 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
TCP Connection Stats
  0 attempts, 3 accepts, 3 established
  3 closed, 2 drops, 0 conn drops
  0 drop in retransmit timeout, 1 drop in keepalive timeout
  0 drop in persist drops, 0 connections drained
TCP Miscellaneous Stats
  115 segments timed, 121 rtt updated
  0 retransmit timeout, 0 persist timeout
  12 keepalive timeout, 11 keepalive probes
TCP SACK Stats
  0 recovery episodes, 0 data packets, 0 data bytes
  0 data packets retransmitted, 0 data bytes retransmitted
  0 connections closed, 0 retransmit timeouts
TCP SYN Cache Stats
  15 entries, 3 connections completed, 0 entries timed out
  0 dropped due to overflow, 12 dropped due to RST
  0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
  0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
  0 hash collisions, 0 retransmitted
TCP Active Connections
  Local Address          Remote Address          State      Send-Q  Recv-Q
  0.0.0.0:3260          0.0.0.0:0              LISTEN    0       0

```

Use the **show ips stats icmp interface gigabitethernet** to display and verify IP statistics. This command takes the main Ethernet interface as a parameter and returns the ICMP statistics for that interface. See [Example 6-7](#).

Example 6-7 Displaying ICMP Statistics

```

switch# show ips stats icmp interface gigabitethernet 2/1
ICMP Statistics for port GigabitEthernet2/1
  0 ICMP messages received
  0 ICMP messages dropped due to errors
ICMP input histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply
ICMP output histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply

```

Displaying IP Storage Ports Speed

Use the **show ips status** command to verify the programmed speed of an IP storage port.

Example 6-8 Displays IP Storage Port Speed

```

switch# show ips status
Port 1/1 READY 10G
Port 1/2 READY 1G

```

Default Settings for IP Storage Services Parameters

Table 6-3 lists the default settings for IP storage services parameters.

Table 6-3 Default Gigabit Ethernet Parameters

Parameters	Default
IPS core size	Partial

