

Configuring FCIP

Cisco MDS 9000 Family IP Storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch can connect separated SAN islands using Fibre Channel over IP (FCIP).

Note FCIP is supported on Fibre Channel module with IPS ports on Cisco 24/10 port SAN Extension Module on Cisco MDS 9700 Series switches, MDS 9250i Multiservice Fabric Switch, and MDS 9220i Fabric Switch.

This chapter includes the following sections:

- Feature Information, on page 1
- Overview of FCIP, on page 2
- Configuring FCIP, on page 6
- Advanced FCIP Profile Configuration, on page 12
- Advanced FCIP Interface Configuration, on page 19
- Advanced FCIP Features, on page 30
- Configuring FCIP Links for Maximum Performance, on page 43
- Verifying FCIP Configuration, on page 53
- Default Settings for FCIP Parameters, on page 53

Feature Information

This section briefly describes the new and updated features for releases.

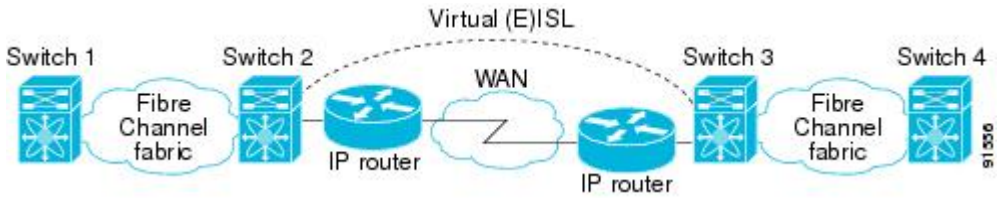
Table 1: Feature Information Table

Feature	Release	Description
FCIP MAC statistics	9.4(3)	Added the Ethernet MAC statistics for an IPStorage interface used for FCIP
tcp min-retransmit-time *milliseconds*	9.4(2)	The default TCP minimum retransmit time is changed to 50 from 200 milliseconds. The TCP minimum retransmit time range is changed to 50 to 5000 milliseconds from 200 to 5000 milliseconds.

Overview of FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). The switch can connect separated SAN islands using Fibre Channel over IP (FCIP). (See the following figure).

Figure 1: Fibre Channel SANs Connected by FCIP



FCIP uses TCP as a network layer transport. The DF bit is set in the TCP header.

For more information about FCIP protocols, refer to the IETF standards for IPStorage at http://www.ietf.org. Also refer to the Fibre Channel standards for switch backbone connection at http://www.t11.org (see FC-BB-2).

This section includes the following topics:

FCIP Concepts

To configure Fibre Channel module with IPS ports for FCIP, you should have a basic understanding of the following concepts:

- FCIP and VE Ports, on page 2
- FCIP Links, on page 3
- FCIP Profiles, on page 4
- FCIP Interfaces, on page 4

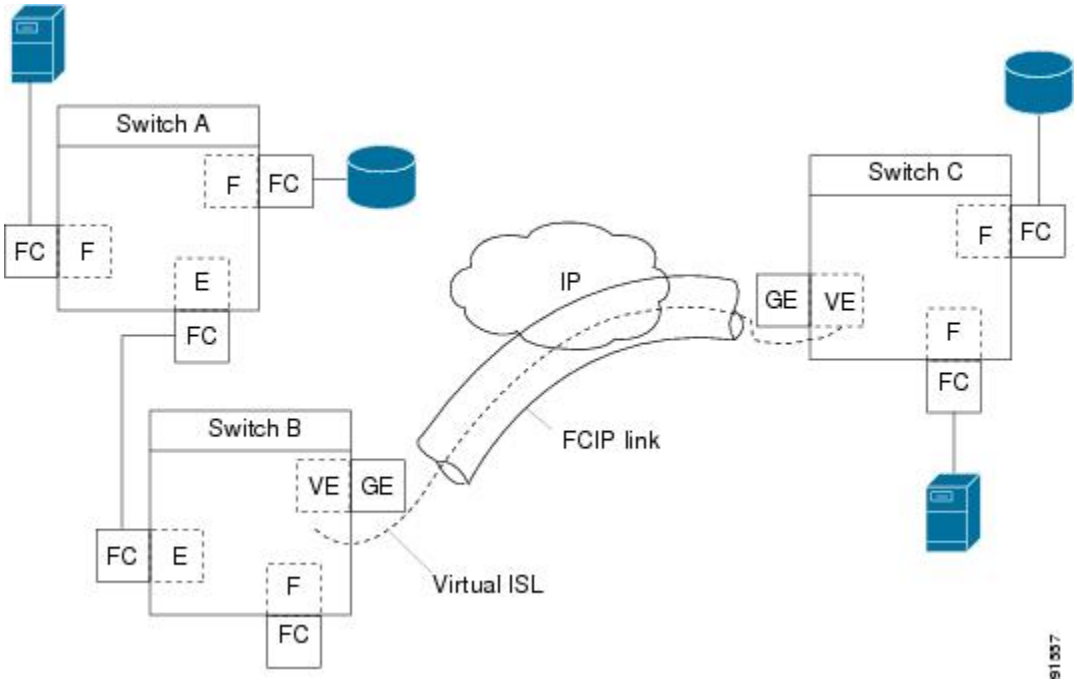
FCIP and VE Ports

The below figure shows the internal model of FCIP in relation to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

FCIP virtual E (VE) ports operate exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see the below figure).

Figure 2: FCIP Links and Virtual ISLs



See the Configuring E Ports, on page 23 for more information.

FCIP Links

Note In-order delivery (IOD) is not supported on FCIP ISLs where WA and TA are enabled.

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link:

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch protocol frames (all Class F). This arrangement provides low latency for all control frames.

To enable FCIP on the Fibre Channel module with IPS ports, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization operation is identical to a normal E port. This operation is independent of the link being FCIP or pure Fibre Channel, and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port operation is identical to E port operation for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

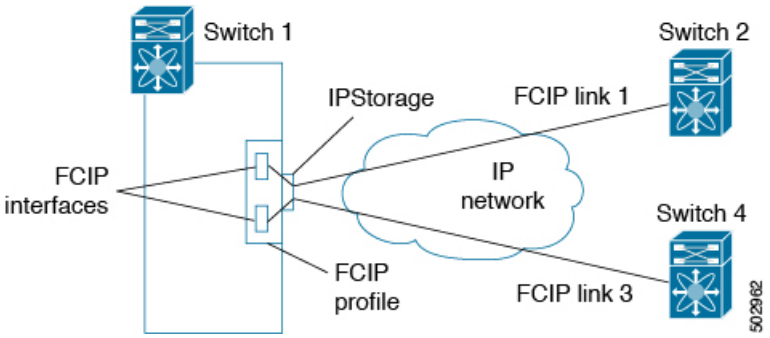
FCIP Profiles

The FCIP profile contains information about the local IP address and TCP parameters. The profile defines the following information:

- The local connection points (IP address and TCP port number)
- The operation of the underlying TCP connections for all FCIP links that use this profile

The FCIP profile's local IP address determines the IPStorage port where the FCIP links terminate (see the below figure).

Figure 3: FCIP Profile and FCIP Link



FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which IPStorage port initiates the FCIP links and defines the TCP connection operation.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk-allowed VSAN list.

FCIP High-Availability Solutions

The following high-availability solutions are available for FCIP configurations:

Fibre Channel Port Channels

Port channels comprised of FCIP interfaces behave the same way as Fibre Channel port channels. They offer the same benefits of link redundancy between Fibre Channel switches as native FC port channels. Beneath

the FCIP level, an FCIP link can run on top of a IPStorage port. This link is totally transparent to the Fibre Channel layer.

The Fibre Channel Port Channel (to which FCIP link can be a part of) does not have a restriction on which (E)ISL links can be combined in a Fibre Channel Port Channel as long as it passes the compatibility check.

Note When port channel is configured with members on Cisco MDS 48-Port 64-Gbps Fibre Channel Switching Module (DS-X9748-3072K9), packet drops may be expected when IOD is configured and a member of the port channel is added, removed, or flaps.

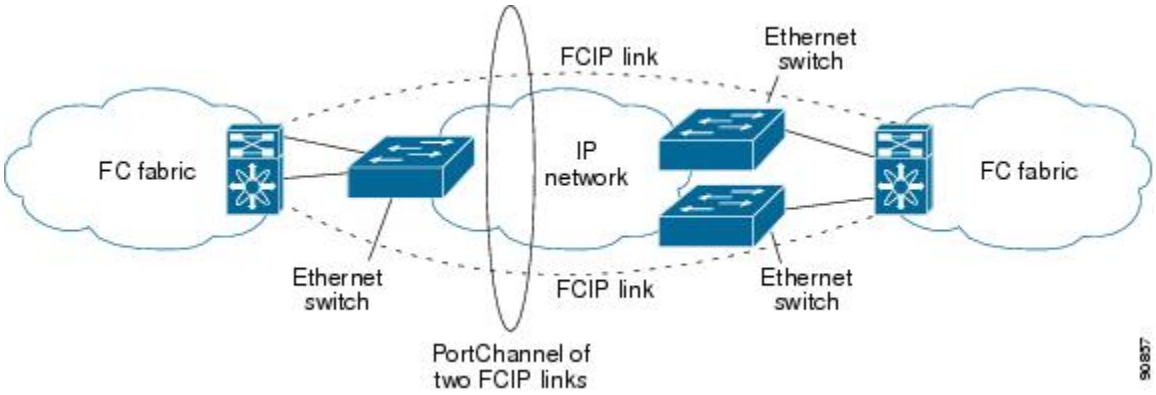
For information, see the Cisco Fabric Manager Interfaces Configuration Guide and Cisco MDS 9000 Series NX-OS Interfaces Configuration Guide.

The maximum number of Fibre Channel ports that can be put into a Fibre Channel Port Channel is 16.

To configure Fibre Channel Port Channels, see the Cisco MDS 9000 Series NX-OS Interfaces Configuration Guide and Cisco Fabric Manager Interfaces Configuration Guide.

The below figure provides an example of a port channel-based load-balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 4: Port Channel-Based Load Balancing



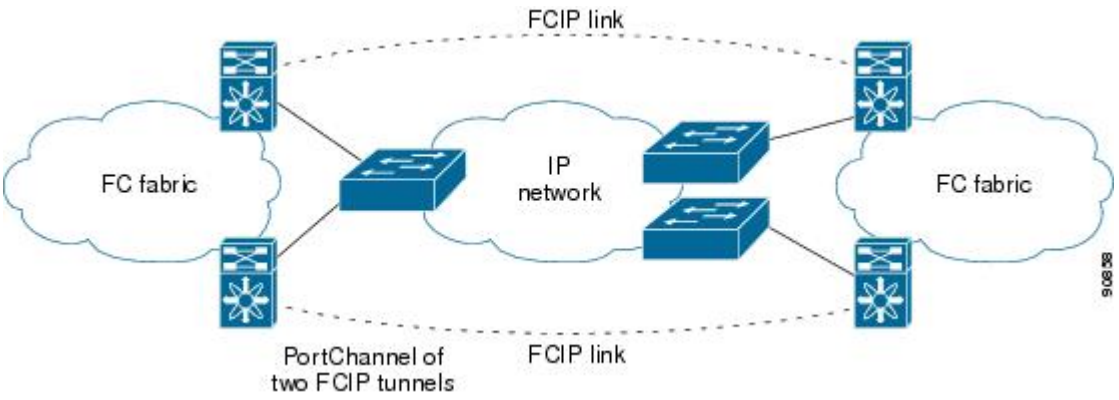
The following characteristics set Fibre Channel port channel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the port channel should terminate on the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the port channel.

FSPF

The below figure displays a FPSF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 5: FSPF-Based Load Balancing



The following characteristics set FSPF solutions apart from other solutions:

- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP links.

Configuring FCIP

This section describes how to configure FCIP and includes the following topics:

Enabling FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs).

To begin configuring the FCIP feature, you must explicitly enable FCIP on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification operations commands for the FCIP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

By default, the Cisco MDS 9700 series switches, MDS 9250i switches, and MDS 9220i switches are shipped with the SAN extension over IP package license.

To enable FCIP on any participating switch, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **feature fcip**	Enables FCIP on that switch.

	Command or Action	Purpose
Step 3	(Optional) switch(config)# **no feature fcip**	Disables (default) FCIP on that switch.

Basic FCIP Configuration

Once you have created FCIP links using the FCIP wizard, you may need to modify parameters for these links. This includes modifying the FCIP profiles as well as the FCIP link parameters.

- For Cisco MDS 9220i Switch, each 1 Gbps or 10 Gbps IPStorage interface can have three FCIP links configured at a time and the 40 Gbps IPStorage interface can have four FCIP links configured at a time.
- For Cisco MDS 9250i Switch, each IPStorage interface can have six FCIP links configured at a time.
- For Cisco MDS 24/10-Port SAN Extension Module, each 1 Gbps or 10 Gbps IPStorage interface can have three FCIP links configured at a time and the 40 Gbps IPStorage interface can have four FCIP links configured at a time.



Note When using IPsec and IKE, each IPStorage interface on the Fibre Channel module with IPS ports must be configured in its own IP subnet. If there are multiple IPStorage interfaces configured with IP address or network-mask in the same IP subnet, IKE packets might not be sent out to the correct IPS port and the IPsec link will not come up.



Note You can configure FCIP with IPsec and NAT on Cisco MDS 9220i switches.

To configure an FCIP link, follow these steps on both switches:

Procedure

Step 1 Configure the IPStorage interface.

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

Step 2 Create an FCIP profile and then assign the IPStorage interface's IP address to the profile.

Step 3 Create an FCIP interface and then assign the profile to the interface.

Step 4 Configure the peer IP address for the FCIP interface.

Step 5 Enable the interface.

Creating FCIP Profiles

You must assign a local IP address of an IPStorage interface or subinterface to the FCIP profile to create an FCIP profile. You can assign IPv4 or IPv6 addresses to the interfaces. The below figure shows an example configuration.

Figure 6: Assigning Profiles to Each IPStorage Interface



Procedure

Step 1 To create an FCIP profile in switch 1 in the above figure, follow these steps:

a) switch1# **configure terminal**

Enters configuration mode.

b) switch1(config)# **fcip profile 10**

Creates a profile for the FCIP connection. The valid range is from 1 to 255.

c) switch1(config-profile)# **ip address 10.100.1.25**

Associates the profile (10) with the local IPv4 address of the IPStorage interface (3/1).

d) switch1# **show fcip profile**

Displays FCIP profile information.

Step 2 To assign an FCIP profile in switch 2 in the above figure, follow these steps:

a) switch2# **configure terminal**

Enters configuration mode.

b) switch2(config)# **fcip profile 20**

Creates a profile for the FCIP connection.

c) switch2(config-profile)# **ip address 10.1.1.1**

Associates the profile (20) with the local IPv4 address of the IPStorage interface.

d) switch2# **show fcip profile**

Displays FCIP profile information.

Displaying FCIP Profile Information

Displaying Summary of FCIP Profiles

```
switch# show fcip profile

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

```
ProfileId Ipaddr   TcpPort
------------------------------------------------------------------------------
1         20.1.1.1 3225
2         20.1.1.1 2000
3         20.1.1.1 3000
4         20.1.1.1 4000
5         20.1.1.1 5000
6         20.1.1.1 6000
7         30.1.1.1 3225
8         31.1.1.1 3225
9         32.1.1.1 3225
10        33.1.1.1 3225
11        34.1.1.1 3225
12        35.1.1.1 3225
```

Displaying Detailed FCIP Profile Information

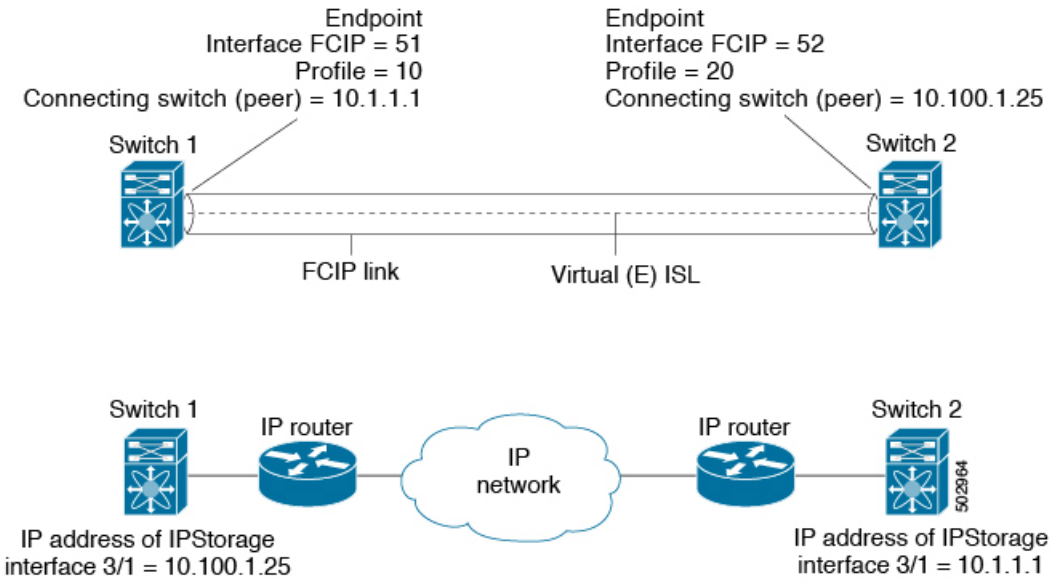
```
switch# show fcip profile 1
FCIP Profile 1
Internet Address is 20.1.1.1 (interface IPStorage1/1)
Tunnels Using this Profile: fcip1
Listen Port is 3225
TCP parameters
SACK is enabled
PMTU discovery is enabled, reset timeout is 3600 sec
Keep alive is 60 sec
Minimum retransmission timeout is 200 ms
Maximum number of re-transmissions is 4
Send buffer size is 16384 KB
Maximum allowed bandwidth is 5000000 kbps
Minimum available bandwidth is 4000000 kbps
Configured round trip time is 1000 usec
Congestion window monitoring is enabled, burst size is 50 KB
Auto jitter detection is enabled
```

Creating FCIP Links

When two FCIP link endpoints are created, an FCIP link is established between the two IPS ports. The peer IP address specifies the address of the remote FCIP endpoint. This allows the creation of an FCIP link to that peer switch once the FCIP interface is enabled.

The below figure shows an example configuration of an FCIP link.

Figure 7: Assigning Profiles to Each IPStorage Interface



Procedure

Step 1 To create an FCIP link endpoint in switch 1, follow these steps:

a) switch1# **configure terminal**

Enters configuration mode.

b) switch1(config)# **interface fcip 51**

Creates an FCIP interface (51).

c) switch1(config-if)# **use-profile 10**

Assigns the profile (10) to the FCIP interface.

d) switch1(config-if)# **peer-info ipaddr 10.1.1.1**

Assigns the peer IPv4 address information (10.1.1.1 for switch 2) to the FCIP interface.

e) switch1(config-if)# **no shutdown**

Enables the interface.

Step 2 To create an FCIP link endpoint in switch 2, follow these steps:

a) switch2# **configure terminal**

Enters configuration mode.

b) switch2(config)# **interface fcip 52**

Creates an FCIP interface (52).

c) switch2(config-if)# **use-profile 20**

Assigns the profile (20) to the FCIP interface.

d) switch2(config-if)# **peer-info ipaddr 10.100.1.25**

Assigns the peer IPv4 address information (10.100.1.25 for switch 1) to the FCIP interface.

e) switch2(config-if)# **no shutdown**

Enables the interface.

Displaying FCIP MAC Statistics

The following example shows how to display Ethernet MAC statistics for an IPStorage interface used for FCIP from Cisco MDS 9000 NX-OS Release 9.4(3):

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

Rx packets input: 0
Rx bytes: 0
Rx multicast frames: 0
Rx broadcast frames: 0
Rx errors: 0
Rx queue drops: 0
Rx if-down drops: 0
Rx RED drops: 0
Rx bad ether type drops: 0
Rx bad protocol drops: 0
Tx packets output: 0
Tx bytes: 0
Tx underruns: 0
Tx multicast frames: 0
Tx broadcast frames: 0
Tx errors: 0
Tx collisions: 0
Tx ARP drops: 0
Tx if-down drops: 0
MAC Error Stats
Rx packets with FCS or align errors: 0
Rx packets with length < min: 0
Rx packets with length < min & FCS error: 0
Rx packets with length > max: 0
Rx packets with length > max & FCS error: 0
Rx packets with PCI error: 0
Rx packets with GMX/SPX/PCI error: 0
Rx packets dropped by RED buffer exhaustion: 0
Rx free buffers went below RED threshold count: 0
Rx L2 multicast packets: 0
Rx L2 broadcast packets: 0
Rx packets dropped by dest MAC filter: 0
Rx packets dropped due to a full receive FIFO: 0
```

```
Rx total MAC errors: 0
Tx packets dropped in PKO queue: 0
Tx packets sent with single collision: 0
Tx packets sent with multiple collisions: 0
Tx packets dropped due to excessive collisions: 0
Tx packets delayed count: 0
Tx total MAC errors: 0
```

Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

This sections includes the following topics:

Note FCIP configuration options can be accessed from the `switch (config-` profile`)#` submode prompt.

Configuring TCP Listener Ports

Procedure

Step 1 To configure TCP listener ports, follow these steps:

a) switch# **configure terminal**

Enters configuration mode.

b) switch(config)# **fcip profile 20**

Creates the profile (if it does not already exist) and enters profile configuration submode. The valid range is from 1 to 511.

The default TCP port for FCIP is 3225. You can change this port by using the **port** command.

Step 2 To change the default FCIP port number (3225), follow these steps:

a) switch(config-profile)# **port 5000**

Associates the profile with the local port number (5000).

b) switch(config-profile)# **no port**

Reverts to the default 3225 port.

Configuring TCP Parameters

You can control TCP behavior in a switch by configuring the TCP parameters that are described in this section.

Note When FCIP is sent over a WAN link, the default TCP settings may not be appropriate. In such cases, we recommend that you tune the FCIP WAN link by modifying the TCP parameters (specifically bandwidth, round-trip times, and CWM burst size).

This section includes the following topics:

Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, from Cisco MDS NX-OS Release 9.4(2) and later releases the value is 50 milliseconds and prior to Cisco MDS NX-OS Release 9.4(2) the value is 200 milliseconds.

To configure the minimum retransmit time, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **tcp min-retransmit-time 500**	Specifies the minimum TCP retransmit time for the TCP connection to be 500 milliseconds. From Cisco MDS NX-OS Release 9.4(2) and later releases, the default is 50 milliseconds and the range is from 50 to 5000 milliseconds. Prior to Cisco MDS NX-OS Release 9.4(2), the default value is 200 milliseconds and the range is from 200 to 5000 milliseconds.
Step 2	(Optional) switch(config-profile)# **no tcp min-retransmit-time 500**	Reverts the minimum TCP retransmit time to the factory default. From Cisco MDS NX-OS Release 9.4(2) and later releases, the factory default is 50 milliseconds. Prior to Cisco MDS NX-OS Release 9.4(2), the factory default value is 200 milliseconds.

Keepalive Timeout

You can configure the interval that the TCP connection uses to verify that the FCIP link is functioning. This ensures that an FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. The keepalive timeout feature can be used to tune the time taken to detect FCIP link failures.

You can configure the first interval during which the connection is idle (the default is 60 seconds). When the connection is idle for the configured interval, eight keepalive probes are sent at 1-second interval. If no response is received for these eight probes and the connection remains idle throughout, then the FCIP link is automatically closed.

Note Only the first interval (during which the connection is idle) can be changed.

To configure the first keepalive timeout interval, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **tcp keepalive-timeout 120**	Specifies the keepalive timeout interval for the TCP connection in seconds (120). The range is from 1 to 7200 seconds.
Step 2	(Optional) switch(config-profile)# **no tcp keepalive-timeout 120**	Reverts the keepalive timeout interval to the default 60 seconds.

Maximum Retransmissions

You can specify the maximum number of times a packet is retransmitted before TCP decides to close the connection.

To configure maximum retransmissions, follow these steps:

Procedure

Step 1 switch(config-profile)# **tcp max-retransmissions 6**

Specifies the maximum number of retransmissions (6). The range is from 1 to 8 retransmissions.

Step 2 (Optional) switch(config-profile)# **no tcp max-retransmissions 6**

Reverts to the default of 4 retransmissions.

Maximum Retransmission Threshold for OBFL Logging

FCIP links utilize peer to peer TCP sessions. The intermediate network is often Ethernet which is not lossless and may drop frames. These are automatically retransmitted by the TCP sender. End applications should be able to tolerate some degree of delay caused by retransmissions. However, some applications may have a lower tolerance for retransmissions than others. Excessive TCP retransmission events are logged to OBFL to assist in troubleshooting. The threshold retransmission rate that will trigger logging can be configured. This retransmission rate is measured per FCIP profile.

To configure the OBFL retransmission rate per FCIP profile, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters the global configuration mode.
Step 2	switch(config)# **fcip profile** *profile_id*	Creates a profile (if it does not already exist) and enters profile configuration submode. The range is from 1 to 255.
Step 3	• `switch(config-profile)# tcp obfl max-retransmissions_pkt_thresh` *value* • `switch(config-profile)# tcp logging onboard tcp-retransmission-threshold` *value*	Specifies the threshold for logging the TCP retransmission rate to OBFL, that is, the number of packets retransmitted per 10000 packets measured every minute. The range is from 5 to 10. The default value is 5. Starting from Cisco MDS Release 8.4(1), the **tcp obfl max-retransmissions_pkt_thresh** *value* command is deprecated. Use the **tcp logging onboard tcp-retransmission-threshold** *value* command to configure the TCP retransmission rate to OBFL. The range is from 0.01 to 0.10. The default value is 0.05%.
Step 4	switch# **show logging onboard error-stats**	Displays the excessive FCIP retransmission logs.

Path MTUs

Path MTU (PMTU) is the minimum MTU on the IP network between the two endpoints of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a timeout of 3600 seconds. If TCP reduces the size of the maximum segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

To configure PMTU, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **no tcp pmtu-enable**	Disables PMTU discovery.
Step 2	switch(config-profile)# **tcp pmtu-enable**	Enables (default) PMTU discovery with the default value of 3600 seconds.
Step 3	switch(config-profile)# **tcp pmtu-enable reset-timeout 90**	Specifies the PMTU reset timeout to 90 seconds. The range is 60 to 3600 seconds.
Step 4	switch(config-profile)# **no tcp pmtu-enable reset-timeout 600**	Leaves PMTU discovery enabled but reverts the timeout to the default of 3600 seconds.

Selective Acknowledgments

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

To configure SACK, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **no tcp sack-enable**	Disables SACK.
Step 2	switch(config-profile)# **tcp sack-enable**	Enables SACK (default).

Window Management

The optimal TCP window size is automatically calculated using the maximum bandwidth parameter, the minimum available bandwidth parameter, and the dynamically measured round-trip time (RTT).

Note The configured round-trip-time parameter determines the window scaling factor of the TCP connection. This parameter is only an approximation. The measured RTT value overrides the round trip time parameter for window management. If the configured round-trip-time is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

The **min-available-bandwidth** parameter and the measured RTT together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at minimum available bandwidth.

The **max-bandwidth-mbps** parameter and the measured RTT together determine the maximum window size.

Note Set the maximum bandwidth to match the worst-case bandwidth available on the physical link, considering other traffic that might be going across this link (for example, other FCIP links, WAN limitations). Maximum bandwidth should be the total bandwidth minus all other traffic going across that link.

Note In Cisco MDS 9250i Multiservice Fabric Switch, you can configure the TCP maximum bandwidth up to 5 Gbps. We recommend that the minimum available bandwidth is 80% of the maximum bandwidth.

To configure window management, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10**	Configures the maximum available bandwidth at 900 Mbps, the minimum slow start threshold at 300 Mbps, and the RTT at 10 milliseconds.
Step 2	(Optional) switch(config-profile)# **no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10**	Reverts to the factory defaults. The FCIP defaults are maximum bandwidth at 1 Gbps, minimum available bandwidth at 500 Mbps, and RTT at 1 milliseconds.
Step 3	switch(config-profile)# **tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200**	Configures the maximum available bandwidth at 2000 Kbps, the minimum available bandwidth at 2000 Kbps, and the RTT at 200 milliseconds.

Monitoring Congestion

By enabling the congestion window monitoring (CWM) parameter, you allow TCP to monitor congestion after each idle period. The CWM parameter also determines the maximum burst size allowed after an idle period. By default, this parameter is enabled and the default burst size is 50 KB.

The interaction of bandwidth parameters and CWM and the resulting TCP behavior is outlined as follows:

- If the average rate of the Fibre Channel traffic over the preceding RTT is less than the min-available-bandwidth multiplied by the RTT, the entire burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.
- If the average rate of the Fibre Channel traffic is greater than min-available-bandwidth multiplied by the RTT, but less than max-bandwidth multiplied by the RTT, then if the Fibre Channel traffic is transmitted in burst sizes smaller than the configured CWM value the entire burst is sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the Fibre Channel traffic is larger than the min-available-bandwidth multiplied by the RTT and the burst size is greater than the CWM value, then only a part of the burst is sent immediately. The remainder is sent with the next RTT.

The software uses standard TCP rules to increase the window beyond the one required to maintain the min-available-bandwidth to reach the max-bandwidth.



Tip We recommend that this feature remains enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

To change the CWM defaults, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-profile)# **no tcp cwm**	Disables congestion monitoring.
Step 2	switch(config-profile)# **tcp cwm**	Enables congestion monitoring and sets the burst size to its default size. The default burst size is 50 KB.
Step 3	switch(config-profile)# **tcp cwm burstsize 30**	Changes the burst size to 30 KB. The valid range is from 10 to 100 KB.
Step 4	(Optional) switch(config-profile)# **no tcp cwm burstsize 25**	Leaves the CWM feature in an enabled state but changes the burst size to its factory default.

Displaying FCIP Profile Information

Use the **show fcip profile** command to display FCIP profile information for the Cisco MDS 9250i Multiservice Fabric Switch:

```
switch# show fcip profile 1
Internet address is 209.165.200.226 (interface IPStorage5/4.101)
Listen Port is 3225
TCP parameters
SACK is enabled
PMTU discovery is enabled, reset timeout is 3600 sec
Keep alive is 60 sec
Minimum retransmission timeout is 200 ms
Maximum number of re-transmissions is 4
Maximum number of obfl re-transmission Thresh is 5
Maximum allowed bandwidth is 5000000 kbps
Minimum available bandwidth is 4000000 kbps
Configured round trip time is 1000 usec
Congestion window monitoring is enabled, burst size is 50 KB
Auto jitter detection is enabled
```

Use the **show fcip profile** command to display FCIP profile information for the 24/10 port SAN Extension module:

```
switch# show fcip profile 41
FCIP Profile 41
Internet Address is 209.165.200.225 (interface IPStorage5/4.101)
Listen Port is 3225
TCP parameters
SACK is enabled
PMTU discovery is enabled, reset timeout is 3600 sec
Keep alive is 60 sec
Minimum retransmission timeout is 200 ms
Maximum number of re-transmissions is 4
Maximum number of obfl re-transmission Thresh is 6
Maximum allowed bandwidth is 10000000 kbps
Minimum available bandwidth is 8000000 kbps
Configured round trip time is 1000 usec
Congestion window monitoring is enabled, burst size is 50 KB
Auto jitter detection is enabled
```

Advanced FCIP Interface Configuration

This section describes the options you can configure on an FCIP interface to establish connection to a peer and includes the following topics:

Configuring Peers

All the FCIP and E port parameters are configured in context to the FCIP interface. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information on the two switches at the ends of the FCIP link. The peer IP switch information causes the switch to initiate an FCIP link to that peer switch. The basic FCIP configuration uses the peer's IP address to configure the peer information. You can establish an FCIP link with the peer using the Peer IP address option. This option configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.

To enter the `config-if` submode, follow these steps:

Before you begin

To establish a peer connection, you must first create the FCIP interface and enter the `config-if` submode.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fcip 100**

Creates an FCIP interface (100).

Note
Each IPStorage interface can have three FCIP links configured at a time. For Cisco MDS 9250i, each IPStorage port can have six FCIP links configured at a time. For Cisco MDS 24/10-Port SAN Extension Module, each IPStorage port can have three FCIP links configured at a time.

Assigning Peer IP Address (IPv4)

To assign the peer information based on the IPv4 address and port number, follow these steps:

Procedure

Step 1 switch(config-if)# **peer-info ipaddr 10.1.1.1**

Assigns an IPv4 address to configure the peer information. Because no port is specified, the default port number (3225) is used.

Step 2 switch(config-if)# **peer-info ipaddr 10.1.1.1 port 3000**

Assigns the IPv4 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.

Step 3 switch(config-if)# **no shutdown**

Enables the interface.

Assigning Peer IP Address (IPv6)

To assign the peer information based on the IPv6 address and port number, follow these steps:

Procedure

Step 1 switch(config-if)# **peer-info ipaddr**

Assigns an IPv6 address to configure the peer information. Because no port is specified, the default port number (3225) is used.

Step 2 switch(config-if)# **peer-info ipaddr 2001:0db8:800:200c::417a port 3000**

Assigns the IPv6 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.

Step 3 switch(config-if)# **no shutdown**

Enables the interface.

Step 4 switch(config-if)# **interface IPStorage slot-number/plot-number**

Enters IPStorage interface configuration mode.

Step 5 switch(config-if)# **shutdown**

Disable the interface.

Step 6 switch(config-if)# **ipv6 enable**

Enables IPv6 processing on the interface.

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

Enables the interface.

Configuring Number of TCP Connections

You can specify the number of TCP connections used for an FCIP link to be either two or five connections. By default, FCIP uses two connections for each link. Connection 0 is the FCIP control connection. The remaining one or four TCP connections are used for data.

Note Make sure that the peer switch FCIP link is also configured with the same number of TCP connections, otherwise FCIP link will not come up.

Note On the Cisco MDS platform, 10 Gb IPStorage ports have different performance characteristics than 1 Gb Ethernet ports. To achieve maximum throughput on FCIP links utilizing MDS 10 Gb IPStorage ports, set the number of TCP connections to 5 on these links.

To specify the TCP connection attempts, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters global configuration mode.

Step 2 switch(config)# **interface fcip 4**

Enters FCIP interface configuration mode.

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

Disables the interface.

Step 4 switch(config-if)# **tcp-connection 5**

Specifies the number of TCP connections. Valid values are 2 or 5.

Note
To change the number of TCP connections ensure that the FCIP interface is shut down first.

Step 5 (Optional) switch(config-if)# **no tcp-connection**

Reverts to the factory set default of two TCP sessions per FCIP interface.

Step 6 switch(config-if)# **no shutdown**

Enables the interface.

Configuring Active Connections

You can configure the mode for initiating a TCP connection. By default, the active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection but waits for the peer to connect to it. By default, the switch tries two TCP connections for each FCIP link.

Note Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

To configure the passive mode, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# **passive-mode**	Enables passive mode while attempting a TCP connection.
Step 2	switch(config-if)# **no passive-mode**	Reverts to the factory set default of using the active mode while attempting the TCP connection.
Step 3	switch(config-if)# **no shutdown**	Enables the interface.

Enabling Time Stamp Control

You can configure the switch to discard packets that are outside a specified time range. When enabled, this feature specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped.

By default, time stamp control is disabled in all switches in the Cisco MDS 9000 Family. When enabled, if a packet arrives within a 2000 millisecond interval (+ or –2000 milliseconds) from the network time that packet is accepted.

Note If the **time-stamp** option is enabled, ensure to configure NTP on both switches (see the Cisco NX-OS Fundamentals Configuration Guide for more information).

Tip Do not enable time stamp control on an FCIP interface that has tape acceleration or Write Acceleration configured.

To enable or disable the time stamp control, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# **time-stamp**`Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link`	Enables time stamp checking for received packets with a default acceptable time difference of 2000 milliseconds.
Step 2	(Optional) switch(config-if)# **no time-stamp**	Disables (default) time stamps.

	Command or Action	Purpose
Step 3	switch(config-if)# **time-stamp acceptable-diff 4000**	Configures the packet acceptance time. The valid range is from 500 to 10,000 millisecond.
Step 4	(Optional) switch(config-if)# **no time-stamp acceptable-diff 500**	Deletes the configured time difference and reverts the difference to factory defaults. The default difference is a 2000-millisecond interval from the network time.
Step 5	switch(config-if)# **no shutdown**	Enables the interface.

Quality of Service

The quality of service (QoS) parameter specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that connection.

To set the QoS values on FCIP interfaces, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# **qos control 24 data 26**	Configures the control TCP connection and data connection to mark all packets on that DSCP value. The control and data value ranges from 0 to 63.
Step 2	switch(config-if)# **no qos control 24 data 26**	Reverts the switch to its factory default (marks all control and data packets with DCSP value 0).

Configuring E Ports

You configure FCIP interfaces in the same way as you configure FC (T)E interfaces. Specifically, the following features are available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN.

  For more information, see the Cisco MDS 9000 Series Fabric Configuration Guide.

- Trunk mode and trunk allowed VSANs

  For more information, see the Cisco MDS 9000 Series Interfaces Configuration Guide.

- Port Channels
  - Multiple FCIP links can be bundled into a Fibre Channel Port Channel.
  - FCIP links and Fibre Channel links cannot be combined in one Port Channel.

For more information, see the Cisco MDS 9000 Series Security Configuration Guide.

- FSPF

  For more information, see the Cisco MDS 9000 Series Fabric Configuration Guide.

- Importing and exporting the zone database from the adjacent switch

  For more information, see the Cisco MDS 9000 Series System Management Configuration Guide.

Displaying FCIP Interface Information

Use the **show interface** commands to view the summary, counter, description, and status of the FCIP link. Use the output of these commands to verify the administration mode, the interface status, the operational mode, the related VSAN ID, and the profile used.

Displaying the FCIP Summary

```
switch# show fcip summary
--------------------------------------------------------------------------------
Tun prof IPS-if peer-ip Status T W T Enc Comp Bandwidth rtt
E A A max/min (us)
--------------------------------------------------------------------------------
1 1 IPS1/1 20.1.1.2 TRNK Y N N N A 5000M/4000M 1000
2 2 IPS1/1 20.1.1.2 TRNK Y N N N A 1000M/800M 1000
3 3 IPS1/1 20.1.1.2 DOWN N N N N N 1000M/800M 1000
4 4 IPS1/1 20.1.1.2 DOWN N N N N N 1000M/800M 1000
5 5 IPS1/1 20.1.1.2 DOWN N N N N N 1000M/800M 1000
6 6 IPS1/1 20.1.1.2 DOWN N N N N N 1000M/800M 1000
7 7 IPS1/2.1 30.1.1.2 TRNK Y N N N M2 1000M/800M 1000
8 8 IPS1/2.2 31.1.1.2 TRNK Y N N N M2 1000M/800M 1000
9 9 IPS1/2.3 32.1.1.2 DOWN N N N N N 1000M/800M 1000
10 10 IPS1/2.4 33.1.1.2 DOWN N N N N N 1000M/800M 1000
11 11 IPS1/2.5 34.1.1.2 DOWN N N N N N 1000M/800M 1000
12 12 IPS1/2.6 35.1.1.2 DOWN N N N N N 1000M/800M 1000
```

Displaying the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 1
fcip1 is trunking
Hardware is IPStorage
Port WWN is 20:2b:00:2a:6a:1b:60:70
Peer port WWN is 20:62:8c:60:4f:73:d7:80
Admin port mode is auto, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 1
Operating Speed is 5 Gbps
Trunk vsans (admin allowed and active) (1)
Trunk vsans (up) (1)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
Interface last changed at Wed Nov 27 15:12:25 2019

Using Profile id 1 (interface IPStorage1/1)
Peer Information
Peer Internet address is 10.1.1.6 and port is 3225
Write acceleration mode is configured off
```

```
Tape acceleration mode is configured off
Tape Accelerator flow control buffer size is automatic
FICON XRC Accelerator is configured off
Ficon Load Balancer configured off for all vsans
Ficon Tape acceleration configured off for all vsans
IP Compression is enabled and set for mode2
Maximum number of TCP connections is 5
QOS control code point is 0
QOS data code point is 0
TCP Connection Information
5 Active TCP connections
27 Attempts for active connections, 2 close of connections
Path MTU 1500 bytes
Current retransmission timeout is 6400 ms
Current Send Buffer Size: 87080 KB, Requested Send Buffer Size: 62500 KB
CWM Burst Size: 50 KB
CONN<0>
Data connection: Local 10.1.1.12:65489, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 7372 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 0 ms, Variance: 100 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
Data connection: Local 10.1.1.12:65487, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 487 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 870 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7573 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
Data connection: Local 10.1.1.12:65485, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 477 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 853 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7419 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
Data connection: Local 10.1.1.12:65483, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 488 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 872 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7590 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
Control connection: Local 10.1.1.12:65481, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 335 KB, Maximum: 24580 KB, Scale: 6
```

```
Peer receive window: Current: 8125 KB, Maximum: 8125 KB, Scale: 7
Congestion window: Current: 50 KB, Slow start threshold: 8069 KB
Measured RTT : 18 us Min RTT: 18 us Max RTT: 23 us
Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
5 minutes input rate 208 bits/sec, 26 bytes/sec, 0 frames/sec
5 minutes output rate 272 bits/sec, 34 bytes/sec, 0 frames/sec
1658 frames input, 174756 bytes
1658 Class F frames input, 174756 bytes
0 Class 2/3 frames input, 0 bytes
0 Reass frames
0 Error frames timestamp error 0
1663 frames output, 202836 bytes
1663 Class F frames output, 202836 bytes
0 Class 2/3 frames output, 0 bytes
0 Error frames
```

Displaying the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 1
fcip1 is trunking
Hardware is IPStorage
Port WWN is 20:62:8c:60:4f:73:d7:80
Peer port WWN is 20:2b:00:2a:6a:1b:60:70
Admin port mode is auto, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 1
Operating Speed is 10000 Mbps
Trunk vsans (admin allowed and active) (1)
Trunk vsans (up) (1)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
Interface last changed at Tue Nov 26 15:16:55 2019

Using Profile id 1 (interface IPStorage2/3)
Peer Information
Peer Internet address is 10.1.1.12 and port is 3225
Write acceleration mode is configured off
Tape acceleration mode is configured off
Tape Accelerator flow control buffer size is automatic
FICON XRC Accelerator is configured off
Ficon Load Balancer configured off for all vsans
Ficon Tape acceleration configured off for all vsans
IP Compression is enabled and set for mode2
Maximum number of TCP connections is 5
QOS control code point is 0
QOS data code point is 0
TCP Connection Information
5 Active TCP connections
23 Attempts for active connections, 5 close of connections
Path MTU 1500 bytes
Current retransmission timeout is 200 ms
Current Send Buffer Size: 149580 KB, Requested Send Buffer Size: 125000 KB
CWM Burst Size: 50 KB
CONN<0>
Data connection: Local 10.1.1.6:3225, Remote 10.1.1.12:65489
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
Peer receive window: Current: 4095 KB, Maximum: 4095 KB, Scale: 6
```

```
Congestion window: Current: 3686 KB, Slow start threshold: 3998 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 24 ms, Variance: 12 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
Data connection: Local 10.1.1.6:3225, Remote 10.1.1.12:65487
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
Peer receive window: Current: 487 KB, Maximum: 487 KB, Scale: 6
Congestion window: Current: 438 KB, Slow start threshold: 462 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 24 ms, Variance: 12 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
Data connection: Local 10.1.1.6:3225, Remote 10.1.1.12:65485
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
Peer receive window: Current: 477 KB, Maximum: 477 KB, Scale: 6
Congestion window: Current: 429 KB, Slow start threshold: 453 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 24 ms, Variance: 12 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
Data connection: Local 10.1.1.6:3225, Remote 10.1.1.12:65483
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 7
Peer receive window: Current: 488 KB, Maximum: 488 KB, Scale: 6
Congestion window: Current: 439 KB, Slow start threshold: 463 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 24 ms, Variance: 12 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
Control connection: Local 10.1.1.6:3225, Remote 10.1.1.12:65481
TCP Parameters
Advertized window: Current: 8123 KB, Maximum: 24580 KB, Scale: 7
Peer receive window: Current: 334 KB, Maximum: 334 KB, Scale: 6
Congestion window: Current: 50 KB, Slow start threshold: 373 KB
Measured RTT : 19 us Min RTT: 20 us Max RTT: 25 us
Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
5 minutes input rate 288 bits/sec, 36 bytes/sec, 0 frames/sec
5 minutes output rate 224 bits/sec, 28 bytes/sec, 0 frames/sec
1130 frames input, 137324 bytes
1130 Class F frames input, 137324 bytes
0 Class 2/3 frames input, 0 bytes
0 Reass frames
0 Error frames timestamp error 0
1132 frames output, 119060 bytes
1132 Class F frames output, 119060 bytes
0 Class 2/3 frames output, 0 bytes
0 Error frames
```

Displaying Detailed FCIP Interface Standard Counter Information

```
switch# show interface fcip 1 counters
fcip1
TCP Connection Information
5 Active TCP connections
27 Attempts for active connections, 2 close of connections
Path MTU 1500 bytes
Current retransmission timeout is 6400 ms
Current Send Buffer Size: 87080 KB, Requested Send Buffer Size: 62500 KB
CWM Burst Size: 50 KB
CONN<0>
Data connection: Local 10.1.1.12:65489, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 24580 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 7372 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 500000 us Max RTT: 0 us
Round trip time: Smoothed 0 ms, Variance: 100 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<1>
Data connection: Local 10.1.1.12:65487, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 487 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 870 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7573 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<2>
Data connection: Local 10.1.1.12:65485, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 477 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 853 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7419 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<3>
Data connection: Local 10.1.1.12:65483, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 488 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8191 KB, Maximum: 8191 KB, Scale: 7
Congestion window: Current: 872 KB, Slow start threshold: 8094 KB
Measured RTT : 500000 us Min RTT: 7590 us Max RTT: 0 us
Round trip time: Smoothed 8 ms, Variance: 4 Jitter: 150 us
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
CONN<4>
Control connection: Local 10.1.1.12:65481, Remote 10.1.1.6:3225
TCP Parameters
Advertized window: Current: 333 KB, Maximum: 24580 KB, Scale: 6
Peer receive window: Current: 8122 KB, Maximum: 8122 KB, Scale: 7
Congestion window: Current: 50 KB, Slow start threshold: 8069 KB
Measured RTT : 500000 us Min RTT: 18 us Max RTT: 0 us
Round trip time: Smoothed 1 ms, Variance: 1 Jitter: 150 us
```

```
TCP Connection Rate
Input Bytes: 0.00 MB/sec, Output Bytes: 0.00 MB/sec
Input Frames: 0/sec, Output Frames: 0/sec
5 minutes input rate 224 bits/sec, 28 bytes/sec, 0 frames/sec
5 minutes output rate 288 bits/sec, 36 bytes/sec, 0 frames/sec
1686 frames input, 177260 bytes
1686 Class F frames input, 177260 bytes
0 Class 2/3 frames input, 0 bytes
0 Reass frames
0 Error frames timestamp error 0
1691 frames output, 206072 bytes
1691 Class F frames output, 206072 bytes
0 Class 2/3 frames output, 0 bytes
0 Error frames
IP compression statistics
83360 rxbytes
55942 rxbytes compressed, 0 rxbytes non-compressed
1.49 rx compression ratio
99132 txbytes
64015 txbytes compressed, 1152 txbytes non-compressed
1.52 tx compression ratio
IP compression flow control statistics
0 bytes queued for hw compression
0 queued for hardware compression
0 queued for hardware decompression
0 slowed tcp flow control
0 accelerated tcp flow control
0 side band flow control ON
2346 side band flow control OFF
IP compression hung statistics
0 times compression engine hung detected
0 jobs replayed for hardware compression
0 jobs replayed for hardware decompression
0 compression jobs not processed during compression engine reset
0 compression response job not processed during compression engine reset
0 decompression jobs not processed during decompression engine reset
0 decompression response job not processed during decompression engine reset
```

Displaying the FCIP Interface Description

```
switch# show interface fcip 51 description
FCIP51
Sample FCIP interface
```

The transmitted bytes shown in the total *txbytes* counter is the amount of data before compression. After compression, the compressed *txbytes* bytes are transmitted with compression and the uncompressed *txbytes* bytes are transmitted without compression. A packet may be transmitted without compression, if it becomes bigger after compression.

Displaying Brief FCIP Interface Counter Information (Cisco MDS 9250i Multiservice Fabric Switch)

```
switch# show interface fcip 1-12 counters brief
--------------------------------------------------------------------------------
Interface Input (rate is 5 min avg) Output (rate is 5 min avg)
------------------------------ ------------------------------
Rate Total Rate Total
MB/s Frames MB/s Frames
--------------------------------------------------------------------------------
fcip1 191 1155974124 225 1363537690
fcip2 173 1046686124 227 1372311228
```

```
fcip3 0 0 0 0
fcip4 0 0 0 0
fcip5 0 0 0 0
fcip6 0 0 0 0
fcip7 189 1143612956 221 1339130294
fcip8 194 1167499884 218 1317700800
fcip9 0 0 0 0
fcip10 0 0 0 0
fcip11 0 0 0 0
fcip12 0 0 0 0
```

Displaying Brief FCIP Interface Counter Information (24/10 port SAN Extension Module)

```
switch# show interface fcip 41 counters brief
--------------------------------------------------------------------------------
Interface Input (rate is 5 min avg) Output (rate is 5 min avg)
------------------------------ ------------------------------
Rate Total Rate Total
MB/s Frames MB/s Frames
--------------------------------------------------------------------------------
fcip41 191 1155974124 225 1363537690
```

Advanced FCIP Features

You can significantly improve application performance by configuring one or more of the following options for the FCIP interface:

FCIP Write Acceleration

The FCIP Write Acceleration feature enables you to significantly improve the application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP Write Acceleration is enabled, WAN throughput is maximized by acknowledging the write frames from the sender at the closer FCIP switch, thereby eliminating the impact of WAN latency for write operations. The Write Acceleration feature is disabled by default and must be enabled on both sides of the FCIP link.

Note FCIP links using Write Acceleration (WA) must be ensured that all accelerated flows go through a single FCIP link (or port channel). This applies to both commands and responses in both directions. If that does not occur, then FCIP WA will fail. Consequently, FCIP WA cannot be used across FSPF equal cost paths because commands and responses could take different paths.

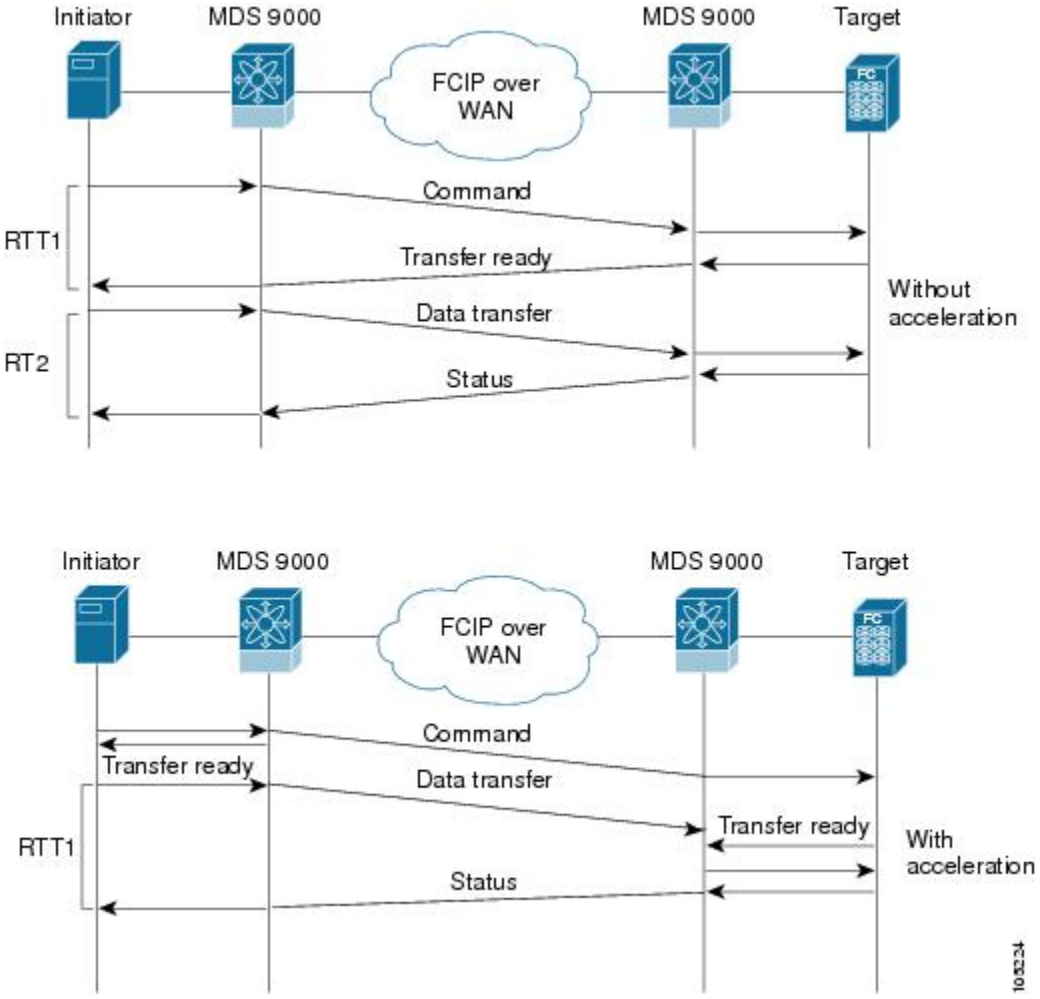
Note Ensure that all FCIP links in a port channel have the same attributes such as 2 or 5 connections, WA, TA, and so on. Otherwise, you will encounter undesirable results during an upgrade.

Note The FCIP Write Acceleration feature accelerates FC standard complaint SCSI WRITE commands only.

Note IBM Peer-to-Peer Remote Copy (PPRC) is not supported by FCIP Write Acceleration.

In the below figure, the SCSI WRITE command without Write Acceleration requires two round-trip transfers (RTT), while the SCSI WRITE command with Write Acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the SCSI WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the SCSI WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Figure 8: FCIP Link Write Acceleration



Tip FCIP Write Acceleration (WA) can be enabled for multiple FCIP links if the links are part of a port channel configured with "channel mode active". These are port channels constructed with Port Channel Protocol (PCP). Write Acceleration cannot be used across FSPF equal cost paths in FCIP deployments. Native Fibre Channel Write Acceleration can be used with port channels. Also, FCIP Write Acceleration can be used in port channels configured with channel mode active or constructed with Port Channel Protocol (PCP). FCIP WA does not work if multiple non-port channel FCIP links exist with equal cost between the initiator and the target ports. This configuration might cause either SCSI discovery failure or failed WRITE or READ operations. When FCIP WA is used the FSPF routing should ensure that a single FCIP Port-Channel or ISL is always in the path between the initiator and the target ports.

Only one FCIP port channel is supported per VSAN on FCIPs configured on Cisco MDS 9700 Series switches with Write Acceleration configured.

Tip Do not enable time stamp control on an FCIP interface with Write Acceleration configured.

Note From Cisco MDS NX-OS Release 7.3(1)DY(1) and later, FCIP Write Acceleration can be enabled when FCIP port channels are configured between a Cisco MDS 9250i switch and a Cisco MDS 24/10 port SAN Extension Module in a Cisco MDS 9700 Director.

Ensure that the following prerequisites are met before enabling Write Acceleration:

- Use the **fcip-enhanced** command on the Cisco MDS 9220i and MDS 9250i switches while creating new port channels for FCIP ports. For more information on creating port channels, see the Configuring Port Channels chapter in the Cisco MDS 9000 Series Interfaces Configuration Guide. For more information on the **fcip-enhanced** command, see the Cisco MDS 9000 Series Command Reference.
- When you configure FCIP port channels between Cisco MDS 9250i and MDS 9700 switches and if write acceleration is enabled on the FCIP ports, then ensure that you enable write acceleration on the FCIP port-channels using the **fcip-enhanced** command on the Cisco MDS 9250i switch. For more information on creating port channels, see the Configuring Port Channels chapter in the Cisco MDS 9000 Series Interfaces Configuration Guide. For more information on the **fcip-enhanced** command, see the Cisco MDS 9000 Series Command Reference.
- Use the **show port-channel database** command to ensure that only FCIP interfaces are members of port channels.
- Enable **passive-mode** on FCIP interfaces created on a Cisco MDS 24/10 port SAN Extension Module in a Cisco MDS 9700 Director. For more information on enabling passive mode, see the Configuring Active Connections, on page 21 section.

Tip In Cisco MDS NX-OS Release 8.1(1) and later, Write Acceleration on FCIP port channels should not be configured between two MDS 9700 Series switches with Cisco MDS 24/10 port SAN Extension Module under the following conditions:

- If traffic flows between an FCoE port and an FC port through an FCIP port channel.
- If traffic flows through two successive FCIP port channels, one after another.

This section includes the following topics:

Configuring FCIP Write Acceleration

To enable Write Acceleration, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch1# **configure terminal**	Enters configuration mode.
Step 2	switch1(config)# **interface fcip 51**	Creates an FCIP interface (51).
Step 3	switch1(config-if)# **write-accelerator**	Enables Write Acceleration.

Displaying Write Acceleration Activity Information

The following examples show how to display information about Write Acceleration activity.

Displaying the Exchanges Processed by Write Acceleration at the Specified Host-End FCIP Link

```
switch# show fcip host-map 100

MAP TABLE (5 entries TOTAL entries 5)

OXID | RXID | HOST FCID| TARG FCID| VSAN | Index
------+------+----------+----------+------+---------
0xd490|0xffff|0x00690400|0x00620426|0x0005|0x0000321f
0xd4a8|0xffff|0x00690400|0x00620426|0x0005|0x00003220
0xd4c0|0xffff|0x00690400|0x00620426|0x0005|0x00003221
0xd4d8|0xffff|0x00690400|0x00620426|0x0005|0x00003222
0xd4f0|0xffff|0x00690400|0x00620426|0x0005|0x00003223
```

Displaying Exchanges Processed by Write Acceleration at the Specified Target End FCIP Link

```
switch# show fcip target-map 100

MAP TABLE (3 entries TOTAL entries 3)

OXID | RXID | HOST FCID| TARG FCID| VSAN | Index
------+------+----------+----------+------+---------
0xc308|0xffff|0x00690400|0x00620426|0x0005|0x00003364
```

```
0xc320|0xffff|0x00690400|0x00620426|0x0005|0x00003365
0xc338|0xffff|0x00690400|0x00620426|0x0005|0x00003366
```

Displaying Detailed FCIP Interface Write Acceleration Counter Information

```
switch# show interface fcip 4 counters
fcip4
TCP Connection Information
...
Write Accelerator statistics
6091 packets in 5994 packets out
0 frames dropped 0 CRC errors
0 rejected due to table full
0 ABTS sent 0 ABTS received
0 tunnel synchronization errors
37 writes recd 37 XFER_RDY sent (host)
0 XFER_RDY rcvd (target)
37 XFER_RDY rcvd (host)
0 XFER_RDY not proxied due to flow control (host)
0 bytes queued for sending
0 estimated bytes queued on the other side for sending
0 times TCP flow ctrl(target)
0 bytes current TCP flow ctrl(target)
```

FCIP Tape Acceleration

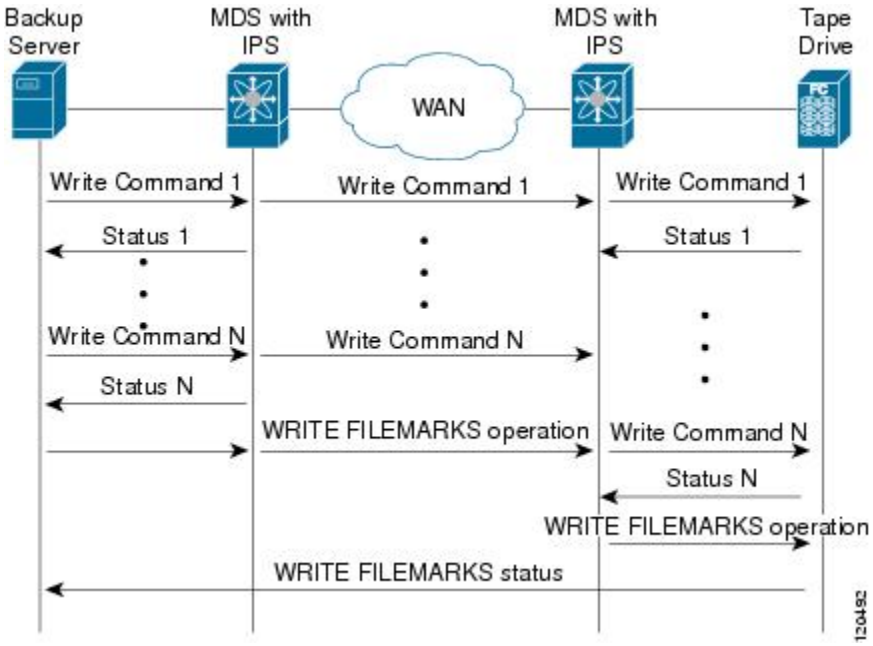
The FCIP Tape Acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP Tape Acceleration is enabled, WAN throughput is maximized by acknowledging the frames from the sender at the closer FCIP switch, thereby eliminating the impact of WAN latency for write operations. The FCIP TAPE Acceleration feature is disabled by default and must be enabled on both sides of the FCIP link.

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS NX-OS provides both tape write and read acceleration.

Applications that access tape drives normally have only one SCSI WRITE or READ operation outstanding to it. This single command process limits the benefit of the tape acceleration feature when using an FCIP link over a long-distance WAN link. It impacts backup and performance because each SCSI WRITE or READ operation does not complete until the host receives a good status response from the tape drive. The FCIP tape acceleration feature helps solve this problem. It improves tape backup, archive, and restore operations by allowing faster data streaming between the host and tape drive over the WAN link.

In an example of tape acceleration for write operations, the backup server in the below figure issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP link in the same time period compared to the time taken to send data without proxying. The proxy method improves the performance on WAN links.

Figure 9: FCIP Link Tape Acceleration for Write Operations



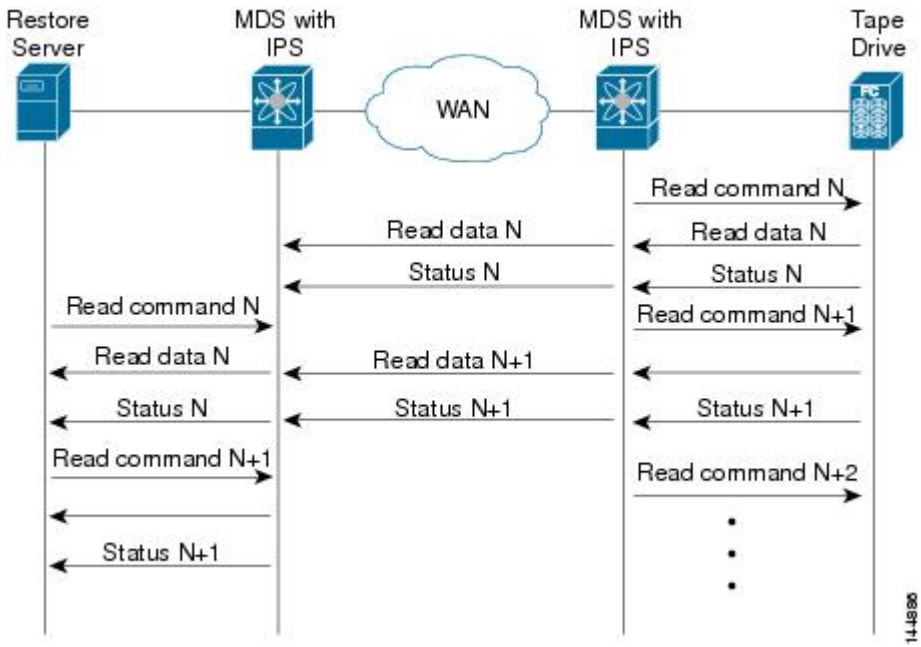
At the tape end of the FCIP link, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

Note In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape library environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these sessions. You need to keep the FCIP link disabled for a couple of minutes before enabling the link. This does not apply to tape environments where the tape drives are either direct FC attached or exported as LUN 0.

The Cisco NX-OS provides reliable data delivery to the remote tape drives using TCP/IP over the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco NX-OS software.

In an example of tape acceleration for read operations, the restore server in the below figure issues read operations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI read operations from the host, sends out SCSI read operations on its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI read operations from the host, sends out the cached data. This method results in more data being sent over the FCIP link in the same time period compared to the time taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Figure 10: FCIP Link Tape Acceleration for Read Operations



The Cisco NX-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco NX-OS software recovers from any other errors.

Note The FCIP tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP link, the tape acceleration feature is operationally off.

Tip FCIP tape acceleration does not work if the FCIP port is part of a port channel or if there are multiple paths between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.

When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if an FCIP interface is up in a FICON VSAN, tape acceleration cannot be enabled on that interface.

Note When you enable the tape acceleration feature for an FCIP link, the link is reinitialized and the write and read acceleration feature is also automatically enabled.

In tape acceleration for writes, after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch by not proxying the Transfer Ready. On completion of a write operation when some data buffers are freed, the local Cisco MDS switch resumes the proxying. Likewise, in tape acceleration for reads, after a certain amount of data has been buffered at the local Cisco MDS switch, the read operations to the tape drive are flow controlled by

the remote Cisco MDS switch by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow control buffer size (the maximum buffer size is 12 MB).

We recommend that you use the default option for flow-control buffering.

Note FCIP tape acceleration will work with five TCP connections.

Tip Do not enable time-stamp control on an FCIP interface with tape acceleration configured.

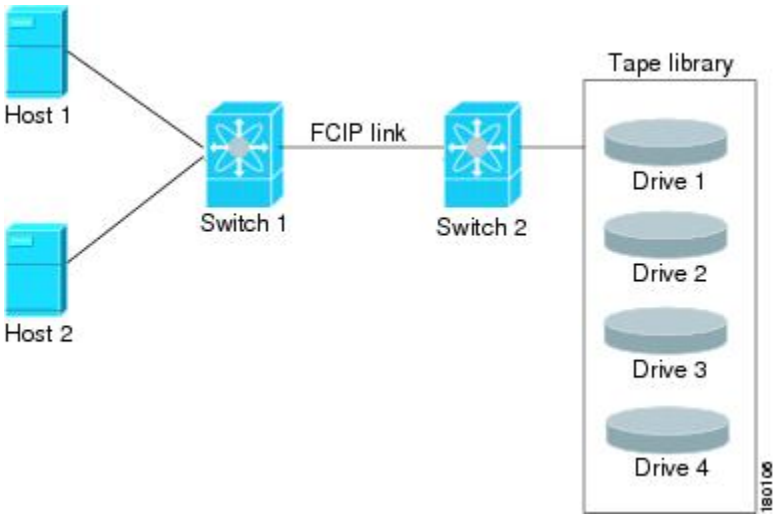
This section includes the following topics:

Tape Library LUN Mapping for FCIP Tape Acceleration

If a tape library provides logical unit (LU) mapping and FCIP tape acceleration is enabled, you must assign a unique LU number (LUN) to each physical tape drive accessible through a target port.

The following figure shows tape drives connected to Switch 2 through a single target port. If the tape library provides LUN mapping, then all the four tape drives should be assign unique LUNs.

Figure 11: FCIP LUN Mapping



For the mappings described in Table 2: Correct LUN Mapping Example with Single Host Access, on page 38 and Table 3: Incorrect LUN Mapping Example with Single Hosts Access, on page 38, Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

The following table describes correct tape library LUN mapping.

Table 2: Correct LUN Mapping Example with Single Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 3	Drive 3
	LUN 4	Drive 4

The following table describes incorrect tape library LUN mapping.

Table 3: Incorrect LUN Mapping Example with Single Hosts Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 1	Drive 3
	LUN 2	Drive 4

Another example setup is when a tape drive is shared by multiple hosts through a single tape port. For instance, Host 1 has access to Drive1 and Drive2, and Host 2 has access to Drive 2, Drive 3, and Drive 4. A correct LUN mapping configuration for such a setup is shown in the following table.

Table 4: Correct LUN Mapping Example with Multiple Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2

Host	LUN Mapping	Drive
Host 2	LUN 2	Drive 2
	LUN 3	Drive 3
	LUN 4	Drive 4

Configuring FCIP Tape Acceleration



Note In an FCIP tape acceleration link, if the trunk mode is **on** for TA enabled links, then the trunk mode allowed VSAN should be configured such that each VSAN's traffic passes through only one link. If the traffic passes through multiple links, it may cause traffic failures.

To enable FCIP tape acceleration, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch1# **configure terminal**	Enters configuration mode.
Step 2	switch1(config)# **interface fcip 5**	Creates an FCIP interface (5).
Step 3	switch1(config-if)# **write-accelerator tape-accelerator**	Enables tape acceleration (and Write Acceleration—if not already enabled).
Step 4	switch1(config-if)# **write-accelerator tape-accelerator flow-control-buffer-size auto**	Enables tape acceleration with automatic flow control (default).
Step 5	switch1(config-if)# **write-accelerator tape-accelerator flow-control-buffer-size 2048**	Sets tape acceleration flow control buffer size to 2 MB.
Step 6	(Optional) switch1(config-if)# **no write-accelerator tape-accelerator**	Disables tape acceleration (default) and resets the FCIP link. **Note** The Write Acceleration feature remains enabled.
Step 7	(Optional) switch1(config-if)# **no write-accelerator tape-accelerator flow-control-buffer-size 2048**	Changes the flow control buffer size to the default value of automatic. The tape acceleration and Write Acceleration features remain enabled. This command does not reset the FCIP link.
Step 8	(Optional) switch1(config-if)# **no write-accelerator**	Disables both the Write Acceleration and tape acceleration features and resets the FCIP link.

Displaying Tape Acceleration Activity Information

The following examples show how to display information about tape acceleration activity.

Displaying Information About Tapes for Which Exchanges are Tape Accelerated

```
switch# show fcip tape-session summary

--------------------------------------------------------------------------------
Tunnel Tunnel End host-fcid tape-fcid lun vsan TCP Connection
--------------------------------------------------------------------------------

16 host-end 0x7c0006 0x390006 0x0000 3000 0
16 host-end 0x7c0004 0x390004 0x0000 3000 0
16 host-end 0x7c0003 0x390003 0x0000 3000 0
16 host-end 0x7c0007 0x390007 0x0000 3000 0
16 host-end 0x7c0005 0x390005 0x0000 3000 0
16 host-end 0x7c0000 0x390000 0x0000 3000 0
16 host-end 0x7c0002 0x390002 0x0000 3000 0
```

Displaying Information About Tapes for Which Exchanges are Tape Accelerated at the Host-End FCIP Link

```
switch# show fcip tape-session tunnel 1 host-end

HOST TAPE SESSIONS (1 entries TOTAL entries 1)

Host Tape Session #1
FCID 0xEF0001, VSAN 1, LUN 0x0002
Outstanding Exchanges 0, Outstanding Writes 0
Target End Write Buffering 0 Bytes, Auto Max Writes 3
Flags 0x0, FSM state Non TA Mode
Cached Reads 0
First index 0xfffffff7, Last index 0xfffffff7, RA index 0x0000f99a
Current index=0xfffffffe, Els Oxid 0xfff7
Hosts 1
FCID 0x770100
```

Displaying Information About Tapes for Which Exchanges are Tape Accelerated at the Target-End FCIP Link

```
switch# show fcip tape-session tunnel 1 targ-end

TARGET TAPE SESSIONS (1 entries TOTAL entries 1)

Target Tape Session #1
FCID 0xEF0001, VSAN 1, LUN 0x0002
Outstanding Exchanges 0, Outstanding Writes 0
Host End Read Buffering 0 Bytes, Auto Max Read Blocks 3
Flags 0x800, Timer Flags 0x0
FSM State Default, Prev FSM State Bypass
Relative Block offset 0
First index 0xfffffff7, Last index 0xfffffff7, RA index 0x0000f99a
Current index=0xfffffffe, Els Oxid 0xfff7
Hosts 1
FCID 0x770100
```

Displays Detailed FCIP Interface Tape Acceleration Counter Information, if Enabled

```
switch# show interface fcip 1 counters
fcip1
TCP Connection Information
....
Tape Accelerator statistics
1 Host Tape Sessions
0 Target Tape Sessions
Host End statistics
Received 31521 writes, 31521 good status, 0 bad status
Sent 31517 proxy status, 4 not proxied
Estimated Write buffer 0 writes 0 bytes
Received 31526 reads, 10 status
Sent 31516 cached reads
Read buffer 0 reads, 0 bytes
Host End error recovery statistics
Sent REC 0, received 0 ACCs, 0 Rejects
Sent ABTS 0, received 0 ACCs
Received 31 RECs, sent 2 ACCs, 0 Rejects
Received 0 SRRs, sent 0 ACCs, 0 Rejects
Received 0 TMF commands
Target End statistics
Received 0 writes, 0 good status, 0 bad status
Write Buffer 0 writes, 0 bytes
Received 0 reads, 0 good status, 0 bad status
Sent 0 reads, received 0 good status, 0 bad status
Sent 0 rewinds, received 0 good status, 0 bad status
Estimated Read buffer 0 reads, 0 bytes
Target End error recovery statistics
Sent REC 0, received 0 ACCs, 0 Rejects
Sent SRR 0, received 0 ACCs
Sent ABTS 0, received 0 ACCs
Received 0 TMF commands
```

FCIP Compression

The FCIP compression option allows IP packets to be compressed on an FCIP link. This option is enabled per FCIP link. The FCIP maximum bandwidth configuration is calculated on the size of the FC frames. Since compression is done at the IP layer and FC is above the IP layer, enabling compression will not cause FCIP interfaces to transmit more than without it (that is, not increase FC throughput). Rather, it helps in reducing the amount of IP traffic sent over the IP network (that is, reduces the IP load). By default, FCIP compression is disabled. When enabled without specifying a compression mode, a compression algorithm appropriate for the FCIP service engine and link speed (as configured in the TCP parameters of the FCIP profile) is automatically used.



Note The main purpose of the FCIP compression feature is to increase throughput on low bandwidth IP connections which are hitting maximum bandwidth usage. Depending on the IP Services switch or module used, the performance of the built-in IP compression engine can be significantly lower than the maximum line rate of IPStorage ports. Therefore, configuring FCIP compression on high bandwidth links can be detrimental to the overall application level performance.

The Cisco MDS 9250i switch and Cisco 24/10 port SAN Extension Module supports Auto, Mode1 and Mode2 compression modes. All of these modes internally use the hardware compression engine in the module. Auto

mode is enabled by default. Mode2 uses a larger batch size for compression than Auto-mode, which results in higher compression throughput. However, Mode2 incurs a small latency due to the compression throughput. For those deployments where aggressive throughput is most important, Mode2 can be used. Mode1 gives the best compression ratio when compared to all other modes. For those deployments where compression ratio is most important, Mode1 can be used.

Note If both ends of an FCIP link are running Cisco MDS NX-OS Release 8.1(1) or later, and you enable compression at one end of the FCIP link, be sure to enable it at the other end of the link also.

Note When using FCIP compression, the rates specified in **tcp max-bandwidth-xxxx** and **min-available-bandwidth-xxxx** in the FCIP profile are in compressed bites.

This section includes the following topics:

Configuring FCIP Compression

To enable FCIP compression, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fcip 51**	Creates an FCIP interface (51).
Step 3	switch(config-if)# **ip-compression mode2**	Enables high compression.
Step 4	(Optional) switch(config-if)# **ip-compression auto**	Defaults to using the **auto** mode.
Step 5	(Optional) switch(config-if)# **no ip-compression**	Disables (default) the FCIP compression feature.

Displaying FCIP Compression Information

The following examples show how to display FCIP compression information.

Displaying Detailed FCIP Interface Compression Information, if Enabled

```
switch# show interface fcip 4 counters
fcip4
TCP Connection Information
...
IP compression statistics
208752 rxbytes, 208752 rxbytes compressed
5143584 txbytes
0 txbytes compressed, 5143584 txbytes non-compressed
1.00 tx compression ratio
```

Displaying the Compression Engine Statistics for the 9250i

```
switch# (config-vsan-db)# show ips stats hw-comp interface IPStorage 1/1
HW Compression statistics for port IPStorage1/1
Compression stats
10444189094728 input bytes, 2822607905236 output compressed bytes
85952406 input pkts, 85952065 output compressed pkts
Decompression stats
8596899248 input compressed bytes, 27669956608 output bytes
45879853 input compressed pkts, 45879669 output pkts
Passthru stats
0 input bytes, 0 output bytes
0 input pkts, 0 output pkts
Miscellaneous stats
0 min input pktlen, 638570 max input pktlen
0 min output pktlen, 185641 max output pktlen
0 len mismatch, 0 incomplete processing
0 invalid result, 0 invalid session drop
0 comp expanded
Errors stats
0 decomp tx error, 0 post comp error
0 post decomp error, 0 comp packets expanded
```

Displaying the Compression Engine Statistics for 24/10 port SAN Extension Module

```
switch# show ips stats hw-comp interface iPStorage 5/1

HW Compression statistics for port IPStorage5/1
Compression stats
53280732 input bytes, 44561835 output compressed bytes
544700 input pkts, 544700 output compressed pkts
Decompression stats
41760802 input compressed bytes, 49574684 output bytes
511886 input compressed pkts, 511886 output pkts
Passthru stats
0 input bytes, 0 output bytes
0 input pkts, 0 output pkts
Miscellaneous stats
0 min input pktlen, 3816 max input pktlen
0 min output pktlen, 1485 max output pktlen
0 len mismatch, 0 incomplete processing
0 invalid result, 0 invalid session drop
0 comp expanded
Errors stats
0 decomp tx error, 0 post comp error
0 post decomp error, 0 comp packets expanded
```

Configuring FCIP Links for Maximum Performance

This section describes how to configure FCIP links for optimum performance between two Cisco MDS 9250i switches, or 24/10 port SAN Extension Modules. We recommend that the maximum and minimum bandwidth parameters in an FCIP profile be the same on both the sides.

Note

- FCIP links with a **tcp max-bandwidth-mbps** of 33 Mbps or less will result in FSPF cost of 30000. This makes the interface unusable. Starting from Cisco MDS NX-OS Releases 8.2(1), the FSPF cost for low bandwidth FCIP links is set to 28999. Because this value is less than the FSPF maximum cost of 30000, and it will allow the traffic to be routed across the interface. It also allows additional FC or FCoE hops (including the FCIP hop) in the end-to-end path. The total FSPF cost of these additional hops should not exceed 1000, because the path will not be usable. If the FSPF cost of 28999 is not applicable for a specific topology, it should be manually configured using the **fspf cost** interface configuration command. To check the FSPF cost of an interface, use the **show fspf interface** command.

  For more information on FSPF Cost, see the Cisco MDS Fabric Configuration Guide.

- FSPF equal cost multipath (ECMP) cannot be used to load balance the traffic across different types of ISLs. In other words, two or more equal cost ISLs of a different type between two switches (e.g. FCIP + FC, or FCIP + FCoE) is not supported.

- On Cisco MDS 24/10 Port SAN Extension Module, configuring multiple ECMP port channels with FCIP members in the same VSAN is not a valid configuration. If this is configured, then traffic will flow through only one of the port channels.

- Maximum FCIP performance is achieved with no packet loss between the FCIP TCP endpoints. FCIP will perform best when the TCP retransmission rate is below 0.05%. To determine if the TCP retransmission rate is above 0.05% use the **show logging onboard error-stats** command to display 60 second intervals of time, where the TCP retransmission rate is above 0.05% on a per FCIP interface and per TCP connection basis.

This section includes the following topics:

Configuring FCIP Links for Maximum Performance on a Cisco MDS 9250i Switch

This section includes the following topics:

Configuring FCIP performance in 1 Gbps mode

To achieve maximum FCIP performance in 1 Gbps mode, the following configuration is recommended:

1. Create an FCIP link on the IPStorage port.

Note If more than one FCIP link is bound to an IPStorage interface at 1 Gbps, the combined maximum bandwidth of all links bound to that interface must not exceed 1 Gbps.

2. Set the TCP maximum and minimum bandwidth as 1000 Mbps and 800 Mbps respectively.

Note If the TCP maximum bandwidth is set to any value more than 1000 Mbps, we recommend that you set the number of TCP connections to five.

3. Configure two TCP connections on each FCIP link.

4. Set the MTU size to 2500 for the IPStorage port.
5. Enable compression on each FCIP link.

To achieve maximum FCIP performance in 1 Gbps mode, follow these configuration steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcip profile***profile-id*	Configures an FCIP profile and enters FCIP profile configuration mode.
Step 3	switch(config-profile)# **ip address**ip-address	Assigns an IP address to the FCIP profile. The assigned IP address can be an IPv4 or an IPv6 address.
Step 4	switch(config-profile)# **tcp max-bandwidth-mbps** *1000* **min-available-bandwidth-mbps** *800* **round-trip-time-ms** *milliseconds*	Sets the maximum and minimum available bandwidth of the FCIP link to 1000 Mbps and 800 Mbps respectively, and configures the round-trip time in milliseconds.
Step 5	switch(config-profile)# **exit**	Exits FCIP profile configuration mode and returns to global configuration mode.
Step 6	switch(config)# **interface fcip** *interface-number*	Enters FCIP interface configuration mode.
Step 7	switch(config-if)# **use-profile** *profile-id*	Binds the specified profile to the FCIP link.
Step 8	switch(config-if)# **peer-info ipaddr** *ip-address*	Configures the peer IP address (IPv4 or IPv6).
Step 9	switch(config-if)# **tcp-connections** *2*	Sets the number of TCP connections to 2. This value must be the same at the peer end.
Step 10	switch(config-if)# **ip-compression mode2**	Sets the compression algorithm to **mode2** for the interface. The other modes that can be set are **auto** and **mode1**.
Step 11	switch(config-if)# **no shutdown**	Enables the FCIP interface.
Step 12	switch (config-if)# **exit**	Exits FCIP interface configuration mode and returns to global configuration mode.
Step 13	switch(config)# **interface IPStorage** *slot-number/port-number*	Enters IPStorage interface configuration mode.
Step 14	switch(config-if)# ip address *ip-address subnet-mask*	Assigns an IP address to the interface.
Step 15	switch(config-if)# **switchport mtu 2500**	Sets the MTU size to 2500 for the interface. The valid range for MTU is from 576 to 9216.
Step 16	switch(config-if)# **no shutdown**	Enables the interface.
Step 17	switch(config-if)# **end**	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring FCIP performance in 10 Gbps mode

To achieve maximum FCIP performance in 10 Gbps mode, the following configuration is recommended:

1. Create an FCIP link on the IPStorage port.

   If more than two FCIP links are bound to an IPStorage interface at 10 Gbps, the combined maximum bandwidth of all links bound to that interface must not exceed 10 Gbps.
2. Set the TCP maximum and minimum bandwidth to 5000 Mbps and 4000 Mbps respectively (default value).
3. Configure five TCP connections on each FCIP link.
4. Set the MTU size to 2500 on the IPStorage port.

To achieve maximum FCIP performance in 10 Gbps mode, follow these configuration steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcip profile** *profile-id*	Configures an FCIP profile and enters FCIP profile configuration mode.
Step 3	switch(config-profile)# **ip address** *ip-address*	Assigns an IP address to the FCIP profile. The assigned IP address can be an IPv4 or an IPv6 address.
Step 4	switch(config-profile)# **tcp max-bandwidth-mbps***5000***min-available-bandwidth-mbps***4000 round-trip-time-ms***milliseconds*	Sets the maximum and minimum available bandwidth of the FCIP link to 5000 Mbps and 4000 Mbps respectively, and configures the round-trip time in milliseconds.
Step 5	switch(config-profile)# **exit**	Exits FCIP profile configuration mode and returns to global configuration mode.
Step 6	switch(config)# **interface fcip interface-number**	Enters FCIP interface configuration mode.
Step 7	switch(config-if)# **use-profile***profile-id*	Binds the specified profile to the FCIP link.
Step 8	switch(config-if)# **peer-info ipaddr***ip-address*	Configures the peer IP address (IPv4 or IPv6).
Step 9	switch(config-if)# **tcp-connections***5*	Sets the number of TCP connections to 5. This value must be the same at the peer end.
Step 10	switch(config-if)# **no shutdown**	Enables the FCIP interface.
Step 11	switch(config-if)# **exit**	Exits FCIP interface configuration mode and returns to global configuration mode.
Step 12	switch(config)# **interface IPStorage***slot-number/port-number*	Enters IPStorage interface configuration mode.
Step 13	switch(config-if)# **ip address***ip-address subnet-mask*	Assigns an IP address to the interface.

	Command or Action	Purpose
Step 14	switch(config-if)# **switchport mtu 2500**	Sets the MTU size to 2500 for the interface. The valid range for MTU is from 576 to 9216.
Step 15	switch(config-if)# **no shutdown**	Enables the interface.
Step 16	switch(config-if)# **end**	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring FCIP Links for Maximum Performance on Cisco MDS 24/10 port SAN Extension Module and M9220i switch

This section includes the following topics:

Configuring FCIP performance in 1 Gbps mode

To achieve maximum FCIP performance in 1 Gbps mode, the following configuration is recommended:

1. Create an FCIP link on the IPStorage port.



Note If more than one FCIP link is bound to an IPStorage interface at 1 Gbps, the combined maximum bandwidth of all links bound to that interface must not exceed 1 Gbps.

2. Set the TCP maximum and minimum bandwidth as 1000 Mbps and 800 Mbps respectively.



Note If the TCP maximum bandwidth is set to a value that is more than 1000 Mbps, we recommend that you set the number of TCP connections to five.

3. Configure two TCP connections on each FCIP link.
4. Set the MTU size to 2500 for the IPStorage port.

To achieve maximum FCIP performance in 1 Gbps mode, follow these configuration steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcip profile** *profile-id*	Configures an FCIP profile and enters FCIP profile configuration mode.
Step 3	switch(config-profile)# **ip address** *ip-address*	Assigns an IP address to the FCIP profile. The assigned IP address can be an IPv4 or an IPv6 address.

	Command or Action	Purpose
Step 4	switch(config-profile)# **tcp max-bandwidth-mbps** *1000* **min-available-bandwidth-mbps** *800* **round-trip-time-ms** *milliseconds*	Sets the maximum and minimum available bandwidth of the FCIP link to 1000 Mbps and 800 Mbps respectively, and configures the round-trip time in milliseconds.
Step 5	switch(config-profile)# **exit**	Exits FCIP profile configuration mode and returns to global configuration mode.
Step 6	switch(config)# **interface fcip** *interface-number*	Enters FCIP interface configuration mode.
Step 7	switch(config-if)# **use-profile** *profile-id*	Binds the specified profile to the FCIP link.
Step 8	switch(config-if)# **peer-info ipaddr** *ip-address*	Configures the peer IP address (IPv4 or IPv6).
Step 9	switch(config-if)# **tcp-connections** *2*	Sets the number of TCP connections to 2. This value must be the same at the peer end.
Step 10	switch(config-if)# **ip-compression** *mode2*	Sets the compression algorithm to **mode2** for the interface. The other modes that can be set are **auto** and **mode1** .
Step 11	switch (config-if)# **no shutdown**	Enables the FCIP interface.
Step 12	switch (config-if)# **exit**	Exits FCIP interface configuration mode and returns to global configuration mode.
Step 13	switch(config)# **interface IPStorage** *slot-number/port-number*	Enters IPStorage interface configuration mode.
Step 14	switch(config-if)# **ip address** *ip-address subnet-mask*	Assigns an IP address to the interface.
Step 15	switch(config-if)# **switchport mtu** *2500*	Sets the MTU size to 2500 for the interface. The valid range for MTU is from 576 to 9216.
Step 16	switch(config-if)# **no shutdown**	Enables the interface.
Step 17	switch(config-if)# **end**	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring FCIP performance in 10 Gbps mode

To achieve maximum FCIP performance in 10 Gbps mode, the following configuration is recommended:

In the releases prior to Cisco MDS NX-OS Release 9.3(1), the maximum throughput on a single FCIP tunnel is 7 Gbps. For maximum performance, create two fcip tunnels on 10 Gbps IPS ports.

1. Create an FCIP link on the IPStorage port.

   If more than two FCIP links are bound to an IPStorage interface at 10 Gbps, the combined maximum bandwidth of all links bound to that interface must not exceed 10 Gbps.

2. Set the TCP maximum and minimum bandwidth to 10000 Mbps and 8000 Mbps respectively (default value).

3. Configure five TCP connections on each FCIP link.

4. Set the MTU size to 2500 on the IPStorage port.

To achieve maximum FCIP performance in 10 Gbps mode, follow these configuration steps:

Procedure

Step 1 switch# **configure terminal**

Enters global configuration mode.

Step 2 switch(config)# **fcip profile** *profile-id*

Configures an FCIP profile and enters FCIP profile configuration mode.

Step 3 switch(config-profile)# **ip address** *ip-address*

Assigns an IP address to the FCIP profile. The assigned IP address can be an IPv4 or an IPv6 address.

Step 4 switch(config-profile)# **tcp max-bandwidth-mbps** *10000* **min-available-bandwidth-mbps** *8000* **round-trip-time-ms** *milliseconds*

Sets the maximum and minimum available bandwidth of the FCIP link to 10000 Mbps and 8000 Mbps respectively, and configures the round-trip time in milliseconds.

Step 5 switch(config-profile)# **exit**

Exits FCIP profile configuration mode and returns to global configuration mode.

Step 6 switch(config)# **interface fcip** *interface-number*

Enters FCIP interface configuration mode.

Step 7 switch(config-if)# **use-profile** *profile-id*

Binds the specified profile to the FCIP link.

Step 8 switch(config-if)# **peer-info ipaddr** *ip-address*

Configures the peer IP address (IPv4 or IPv6).

Step 9 switch(config-if)# **tcp-connections***5*

Sets the number of TCP connections to 5. This value must be the same at the peer end.

Step 10 switch(config-if)# **no shutdown**

Enables the FCIP interface.

Step 11 switch(config-if)# **exit**

Exits FCIP interface configuration mode and returns to global configuration mode.

Step 12 switch(config)# **interface IPStorage** *slot-number/port-number*

Enters IPStorage interface configuration mode.

Step 13 switch(config-if)# **ip address** *ip-address subnet-mask*

Assigns an IP address to the interface.

Step 14 switch(config-if)# **switchport mtu 2500**

Sets the MTU size to 2500 for the interface. The valid range for MTU is from 576 to 9216.

Step 15 switch(config-if)# **no shutdown**

Enables the interface.

Step 16 switch(config-if)# **end**

Exits interface configuration mode and returns to privileged EXEC mode.

Configuring FCIP performance in 25 Gbps mode

To achieve maximum FCIP performance in 25 Gbps mode, the following configuration is recommended:

The maximum bandwidth of FCIP tunnel is 10 Gbps. Therefore, configure two FCIP tunnels on each 25 Gbps IPS port.



Note For maximum performance, create two VLAN sub interfaces on each 25 Gbps IPStorage port. On each VLAN subinterface, create FCIP tunnel.

1. Create VLAN sub interfaces on the 25 Gbps IPStorage port.
2. Create an FCIP link on each VLAN sub interface.
3. Set the TCP maximum and minimum bandwidth to 10000 Mbps (10 Gbps) and 8000 Mbps (8Gbps) respectively (default value).
4. Configure five TCP connections on each FCIP link.
5. Set the MTU size to 2500 on the IPStorage port.

To achieve maximum FCIP performance in 25 Gbps mode, follow these configuration steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface IPStorage** *slot-number/port-number VLAN-ID*	Configure VLAN subinterface (VLAN ID) on 25Gbps IPStorage port.
Step 3	switch(config-if)# **ip address** *ip-address subnet-mask*	Assigns an IP address to the VLAN subinterface.
Step 4	switch(config-if)# **switchport mtu** *2500*	Sets the MTU size to 2500 for the VLAN subinterface. The valid range for MTU is from 576 to 9216.
Step 5	switch(config-if)# **no shutdown**	Enables the VLAN subinterface.
Step 6	switch(config-if)# **end**	Exits interface configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
Step 7	switch(config)# **fcip profile** *profile-id*	Configures an FCIP profile and enters FCIP profile configuration mode.
Step 8	switch(config-profile)# **ip address** *ip-address*	Assign IP address of VLAN subinterface configured in Step 3 to the FCIP profile.
Step 9	switch(config-profile)# **tcp max-bandwidth-mbps** *10000* **min-available-bandwidth-mbps** *8000* **round-trip-time-ms** *milliseconds*	Sets the maximum and minimum available bandwidth of the FCIP link to 10000 Mbps and 8000 Mbps respectively, and configures the round-trip time in milliseconds.
Step 10	switch(config-profile)# **exit**	Exits FCIP profile configuration mode and returns to global configuration mode.
Step 11	switch(config)# **interface fcip** *interface-number*	Enters FCIP interface configuration mode.
Step 12	switch(config-if)# **use-profile** *profile-id*	Binds the specified profile to the FCIP link.
Step 13	switch(config-if)# **peer-info ipaddr** *ip-address*	Configures the peer IP address (IPv4 or IPv6).
Step 14	switch(config-if)# **tcp-connections** *5*	Sets the number of TCP connections to 5. This value must be the same at the peer end.
Step 15	switch(config-if)# **no shutdown**	Enables the FCIP interface.
Step 16	switch(config-if)# **exit**	Exits FCIP interface configuration mode and returns to global configuration mode.

Configuring FCIP performance in 40 Gbps mode

To achieve maximum FCIP performance in 40 Gbps mode, the following configuration is recommended:

- The maximum bandwidth of FCIP tunnel is 10 Gbps. Therefore, configure four FCIP tunnels on the 40 Gbps IPS port.
- For maximum performance, create four VLAN sub interfaces on the 40 Gbps IPStorage port. On each VLAN subinterface, create FCIP tunnel.
  - Releases prior to Cisco MDS NX-OS Rel 9.3(1), the maximum throughput on a single FCIP tunnel on Cisco MDS 9220i is 7 Gbps. On the 40 Gbps IPS port, the maximum throughput is 28.5 Gbps when 4 FCIP tunnels are created on separate VLAN subinterfaces.
  - From the Release Cisco MDS NX-OS Rel 9.3(1), the maximum throughput on a single FCIP tunnel on Cisco MDS 9220i is 10Gbps. On the 40 Gbps IPS port, the maximum throughput is 34.6 Gbps when 4 FCIP tunnels are created on separate VLAN subinterfaces.

To achieve maximum FCIP performance in 40 Gbps mode, follow these configuration steps to configure four FCIP tunnels:

1. Create VLAN sub interfaces on the 40 Gbps IPStorage port.
2. Create an FCIP link on each VLAN sub interface.

3. Set the TCP maximum and minimum bandwidth to 10000 Mbps (10 Gbps) and 8000 Mbps (8Gbps) respectively (default value).
4. Configure five TCP connections on each FCIP link.
5. Set the MTU size to 2500 on the IPStorage port.

To achieve maximum FCIP performance in 40 Gbps mode, follow these configuration steps:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface IPStorage** *slot-number/port-number VLAN-ID*	Configure VLAN subinterface (VLAN ID) on 40 Gbps IPStorage port.
Step 3	switch(config-if)# **ip address** *ip-address subnet-mask*	Assigns an IP address to the VLAN subinterface.
Step 4	switch(config-if)# **switchport mtu 2500**	Sets the MTU size to 2500 for the VLAN subinterface. The valid range for MTU is from 576 to 9216.
Step 5	switch(config-if)# **no shutdown**	Enables the VLAN subinterface.
Step 6	switch(config-if)# **end**	Exits interface configuration mode and returns to privileged EXEC mode.
Step 7	switch(config)# **fcip profile** *profile-id*	Configures an FCIP profile and enters FCIP profile configuration mode.
Step 8	switch(config-profile)# **ip address** *ip-address*	Assign IP address of VLAN subinterface configured in Step 3 to the FCIP profile.
Step 9	switch(config-profile)# **tcp max-bandwidth-mbps** *10000* **min-available-bandwidth-mbps** *8000* **round-trip-time-ms** *milliseconds*	Sets the maximum and minimum available bandwidth of the FCIP link to 10000 Mbps and 8000 Mbps respectively, and configures the round-trip time in milliseconds.
Step 10	switch(config-profile)# **exit**	Exits FCIP profile configuration mode and returns to global configuration mode.
Step 11	switch(config)# **interface fcip***interface-number*	Enters FCIP interface configuration mode.
Step 12	switch(config-if)# **use-profile** *profile-id*	Binds the specified profile to the FCIP link.
Step 13	switch(config-if)# **peer-info ipaddr***ip-address*	Configures the peer IP address (IPv4 or IPv6).
Step 14	switch(config-if)# **tcp-connections***5*	Sets the number of TCP connections to 5. This value must be the same at the peer end.
Step 15	switch(config-if)# **no shutdown**	Enables the FCIP interface.
Step 16	switch(config-if)# **exit**	Exits FCIP interface configuration mode and returns to global configuration mode.

Verifying FCIP Configuration

To verify FCIP configurations, use the following commands:

Command	Purpose
show fcip host-map *fcip-id*	Displays the information for a specified map.
show fcip profile *profile-id*	Displays the information for a profile.
show fcip summary	Displays summary information.
show interface fcip	Displays the status of an interface.
show interface fcip *fcip-id* **counters**	Displays the interface counter information.
show interface fcip *fcip-id* **counters brief**	Displays brief information of interface counters.
show interface fcip *fcip-id* **description**	Displays the interface description.
show fcip target-map	Displays information for a specified target map.
show fcip tape-session summary	Displays tape session information.
show fcip tape-session tunnel 1 host-end	Displays information for a specified FCIP link ID.
show fcip tape-session tunnel 1 targ-end	Displays information for the target end.
show ips stats hw-comp all	Displays hardware compression statistics.
show port-channel database	Displays port channel database information.

Default Settings for FCIP Parameters

The following table lists the default settings for FCIP parameters.

Table 5: Default FCIP Parameters

Parameters	Default
TCP default port for FCIP	3225
Base port for TCP connections to peer	2
Minimum retransmit time	200 milliseconds
Keepalive timeout	60 sec
Maximum retransmissions	4 retransmissions
PMTU discovery	Enabled
PMTU enable reset timeout	3600 seconds
SACK	Enabled
Maximum bandwidth (Cisco MDS 24/10 port SAN Extension Module)	10 Gbps
Maximum bandwidth (Cisco MDS 9220i Fabric Switch)	10 Gbps

Maximum bandwidth (Cisco MDS 9250i Multiservice Fabric Switch)	5 Gbps
Minimum available bandwidth (Cisco MDS 24/10 port SAN Extension Module)	8 Gbps
Minimum available bandwidth (Cisco MDS 9220i Fabric Switch)	8 Gbps
Minimum available bandwidth (Cisco MDS 9250i Multiservice Fabric Switch)	4 Gbps
Round trip time	1 millisecond
Buffer size	0 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
TCP connection mode	Active mode
FCIP timestamp	Disabled
Acceptable difference range to accept packets	+/– 2000 milliseconds
B port keepalive responses	Disabled
Write Acceleration	Disabled
Tape Acceleration	Disabled