



CHAPTER 2

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 the MDS 9222i switch, MSM-18/4 module, MDS 9216i switch, MPS-14/2 module, 16-Port Storage Services Node (SSN-16), and IPS modules on MDS 9200 Series directors.

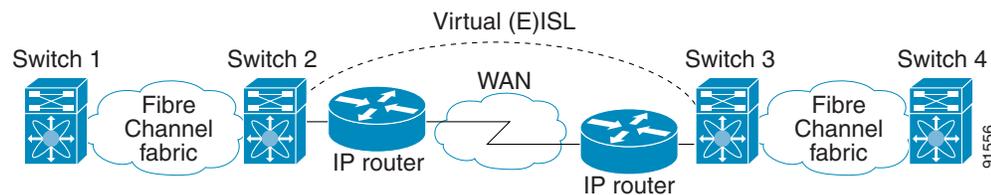
This chapter includes the following sections:

- [About FCIP, page 2-1](#)
- [Configuring FCIP, page 2-7](#)
- [Using the FCIP Wizard, page 2-8](#)
- [Default Settings, page 2-34](#)

About 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 [Figure 2-1](#)).

Figure 2-1 Fibre Channel SANs Connected by FCIP



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

Send documentation comments to fm-docfeedback@cisco.com

**Note**

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

This section includes the following topics:

- [FCIP Concepts, page 2-2](#)
- [FCIP High-Availability Solutions, page 2-4](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 2-7](#)

FCIP Concepts

To enable FCIP on the IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured:

- [FCIP and VE Ports, page 2-2](#)
- [FCIP Links, page 2-3](#)
- [FCIP Profiles, page 2-4](#)
- [FCIP Interfaces, page 2-4](#)

FCIP and VE Ports

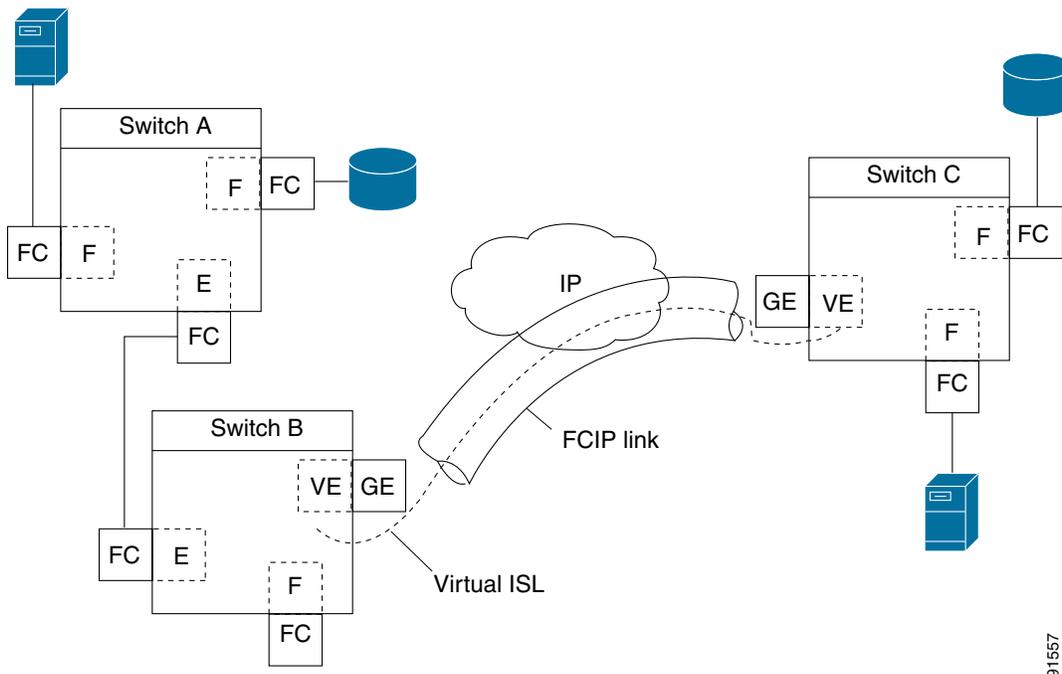
[Figure 2-2](#) describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

FCIP virtual E (VE) ports behave 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 [Figure 2-2](#)).

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-2 FCIP Links and Virtual ISLs



91557

See the “Configuring E Ports” section on page 2-25 for more information.

FCIP Links

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 IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior 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 behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

Send documentation comments to fm-docfeedback@cisco.com

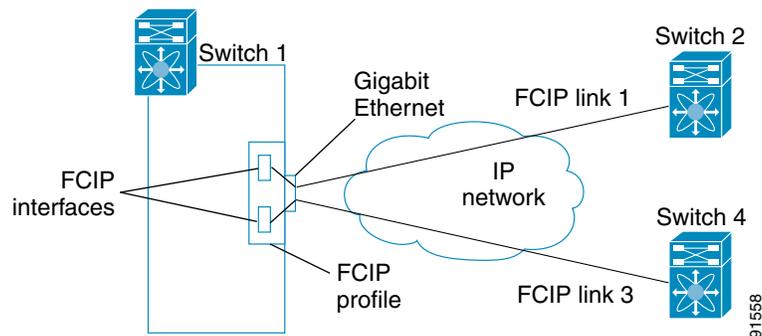
FCIP Profiles

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 behavior of the underlying TCP connections for all FCIP links that use this profile

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see [Figure 2-3](#)).

Figure 2-3 FCIP Profile and FCIP Links



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 Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- 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:

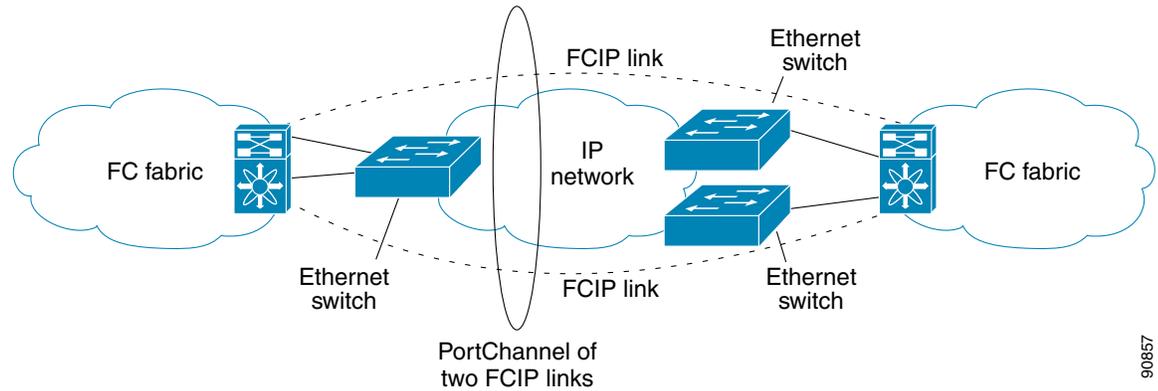
- [Fibre Channel PortChannels](#), page 2-5
- [FSPF](#), page 2-5
- [VRRP](#), page 2-6
- [Ethernet PortChannels](#), page 2-6

Send documentation comments to fm-docfeedback@cisco.com

Fibre Channel PortChannels

Figure 2-4 provides an example of a PortChannel-based load-balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 2-4 PortChannel-Based Load Balancing



90857

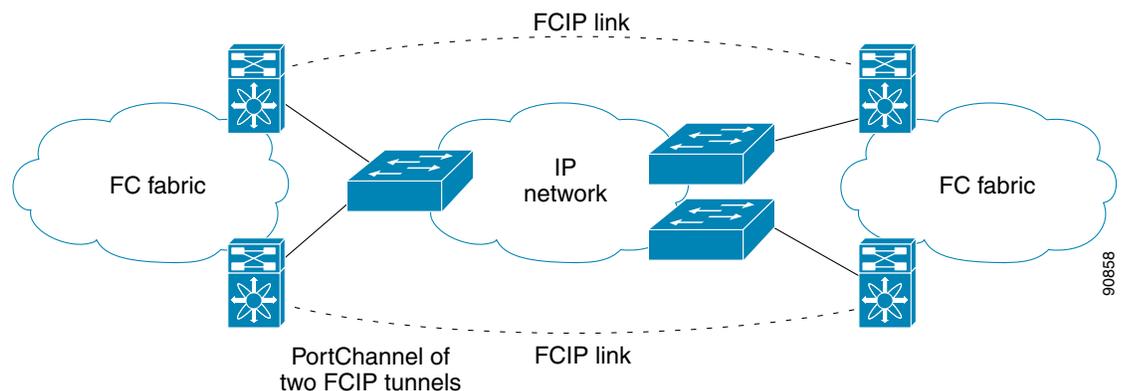
The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

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

FSPF

Figure 2-5 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 2-5 FSPF-Based Load Balancing



90858

The following characteristics set FSPF solutions apart from other solutions:

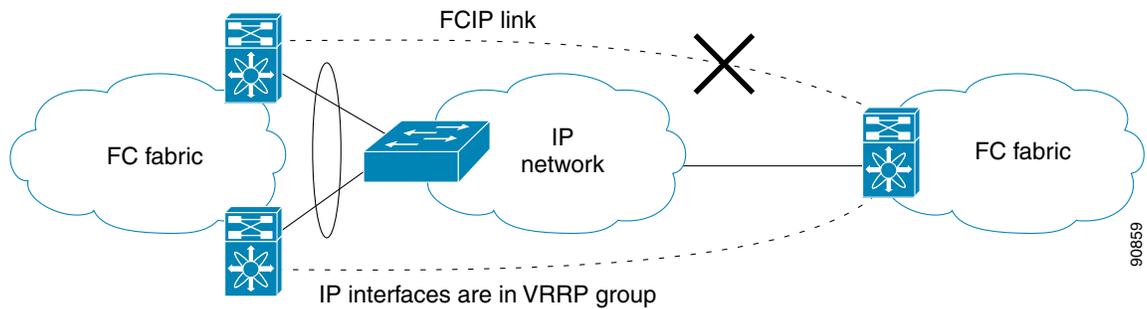
- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP link.

Send documentation comments to fm-docfeedback@cisco.com

VRRP

Figure 2-6 displays a Virtual Router Redundancy Protocol (VRRP)-based high availability FCIP configuration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 2-6 VRRP-Based High Availability



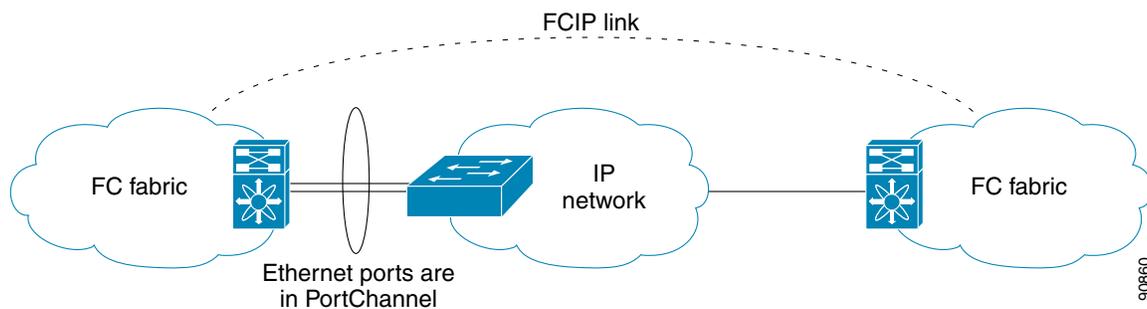
The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Ethernet PortChannels

Figure 2-7 displays an Ethernet PortChannel-based high-availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 2-7 Ethernet PortChannel-Based High Availability



The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link-level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet link.

Send documentation comments to fm-docfeedback@cisco.com

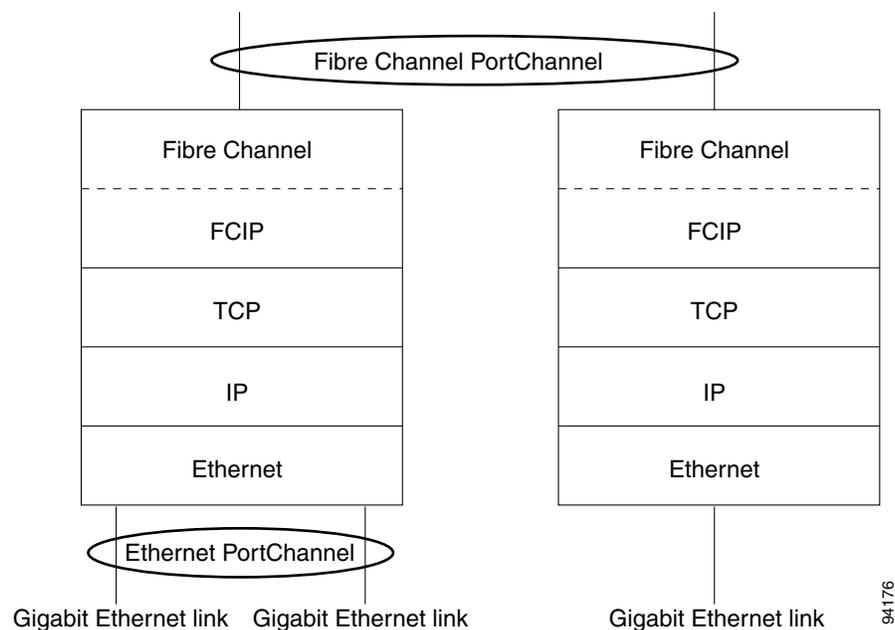
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer link redundancy between the Cisco MDS 9000 Family switch's Gigabit Ethernet ports and the connecting Ethernet switch. Fibre Channel PortChannels also offer (E)ISL link redundancy between Fibre Channel switches. FCIP is an (E)ISL link and is only applicable for a Fibre Channel PortChannel. Beneath the FCIP level, an FCIP link can run on top of an Ethernet PortChannel or on one Gigabit Ethernet port. This link is totally transparent to the Fibre Channel layer.

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet PortChannel (see [Chapter 7, “Configuring IPv4 for Gigabit Ethernet Interfaces”](#) for more information). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortChannel (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 PortChannel as long as it passes the compatibility check. The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see [Figure 2-8](#)).

Figure 2-8 PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see the *Cisco Fabric Manager Interfaces Configuration Guide*.

To configure Ethernet PortChannels, see the *Cisco Fabric Manager High Availability and Redundancy Configuration Guide*.

Configuring FCIP

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

Send documentation comments to fm-docfeedback@cisco.com

- [Enabling FCIP, page 2-8](#)
- [Basic FCIP Configuration, page 2-15](#)
- [Verifying Interfaces and Extended Link Protocol, page 2-16](#)
- [Checking Trunk Status, page 2-17](#)
- [Advanced FCIP Profile Configuration, page 2-17](#)
- [Advanced FCIP Interface Configuration, page 2-20](#)
- [Configuring E Ports, page 2-25](#)
- [Configuring E Ports, page 2-25](#)
- [Advanced FCIP Features, page 2-26](#)

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 for the FCIP feature are available only when FCIP is enabled on a switch.

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.

To use the FCIP feature, you need to obtain the SAN extension over IP package license (SAN_EXTN_OVER_IP or SAN_EXTN_OVER_IP_IPS4) (see the *Cisco Family NX-OS Licensing Guide*). By default, the MDS 9222i and 9216i switches are shipped with the SAN extension over IP package license.

Using the FCIP Wizard



Note

In Cisco MDS SAN-OS Release 2.0 and later and NX-OS Release 4.x, there is an additional login prompt to log into a switch that is not a part of your existing fabric.

To create and manage FCIP links with Fabric Manager, use the FCIP Wizard. Make sure that the IP services module is inserted in the required Cisco MDS 9000 Family switch, and that the Gigabit Ethernet interfaces on these switches are connected, and then verify the connectivity. The procedures for creating FCIP links using the FCIP Wizard are as follows:

- Select the endpoints.
- Choose the interfaces' IP addresses.
- Specify link attributes.
- (Optional) Enable FCIP write acceleration or FCIP compression.

To create FCIP links using the FCIP Wizard, follow these steps:

Send documentation comments to fm-docfeedback@cisco.com

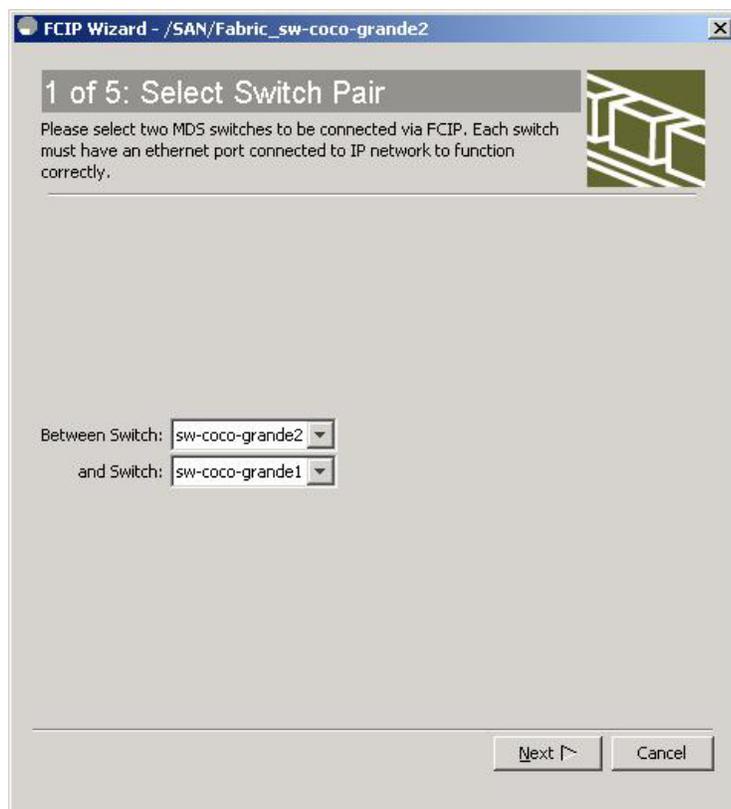
- Step 1** Click the **FCIP Wizard** icon in the Fabric Manager toolbar (See [Figure 2-9](#)).

Figure 2-9 FCIP Wizard



You see the switch selections as shown in [Figure 2-10](#).

Figure 2-10 Switch Selections



- Step 2** Choose the switches that act as endpoints for the FCIP link and click **Next**.
- Step 3** Choose the Gigabit Ethernet ports on each switch that will form the FCIP link.
- Step 4** If both Gigabit Ethernet ports are part of MPS-14/2 modules, check the **Enforce IPSEC Security** check box and set the **IKE Auth Key** (see [Figure 2-11](#)). See the *Cisco Fabric Manager Security Configuration Guide* for information on IPsec and IKE.

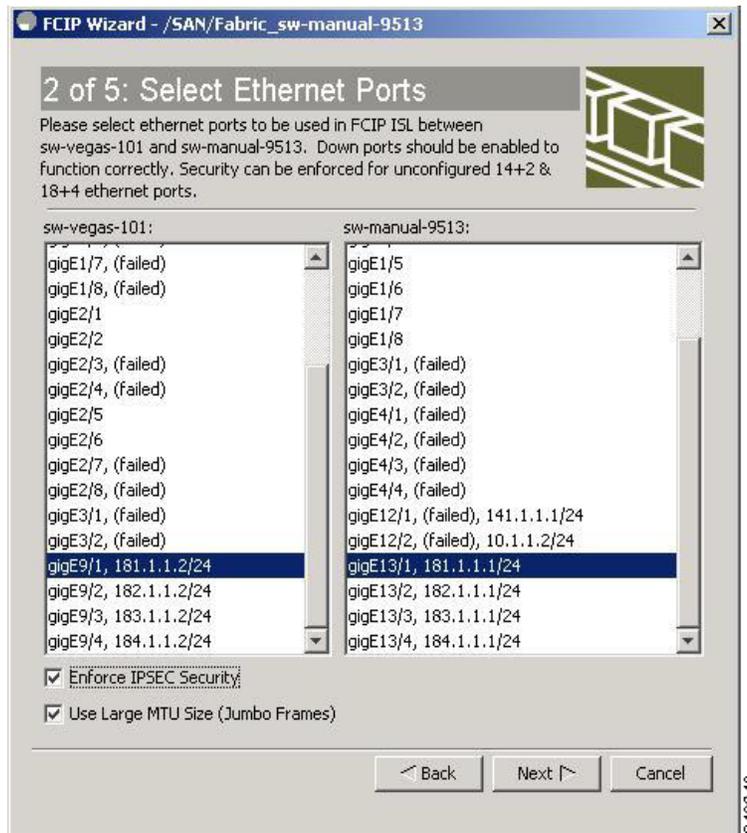
Check the **Use Large MTU Size (Jumbo Frames)** option to use jumbo size frames of 2300. Since Fibre Channel frames are 2112, we recommended that you use this option. If you uncheck the box, the FCIP Wizard does not set the MTU size, and the default value of 1500 is set.

Send documentation comments to fm-docfeedback@cisco.com



Note In Cisco MDS 9000 SAN-OS, Release 3.0(3), by default the **Use Large MTU Size (Jumbo Frames)** option is not selected.

Figure 2-11 Enabling IPsec on an FCIP link



Step 5 Click **Next**. You see the **IP Address/Route** input screen.

Step 6 Select **Add IP Route** if you want to add an IP route, otherwise retain the defaults (see [Figure 2-12](#)).

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-12 Specify IP Address/Route

FCIP Wizard - /SAN/Fabric_sw-manual-9513

3 of 5: Specify IP Address/Route

Please supply Ethernet Port IP Address. Specify Route if the Port addresses are in different subnet.
Note: The changes to IP Address and IP Route Addition will be applied on pressing the Next button.

Switch sw-vegas-101 (gigE9/1)

IP Address/Mask: 181.1.1.2/24 e.g. 10.1.1.1/24
Dest/Mask: e.g. 10.1.0.0/16

Add IP Route: Gateway: e.g. 11.1.1.1
Metric: 0 e.g. 0,32768

Switch sw-manual-9513 (gigE13/1)

IP Address/Mask: 181.1.1.1/24 e.g. 10.1.1.1/24
Dest/Mask: e.g. 10.1.0.0/16

Add IP Route: Gateway: e.g. 11.1.1.1
Metric: 0 e.g. 0,32768

Back Next Cancel

240376

Step 7 Click **Next**. You see the TCP connection characteristics.

Step 8 Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link (see [Figure 2-13](#)).

You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-13 Specifying Tunnel Properties

The screenshot shows the 'FCIP Wizard - /SAN/Fabric' window at step 4 of 5, 'Specify Tunnel Properties'. The window title is 'FCIP Wizard - /SAN/Fabric'. The main heading is '4 of 5: Specify Tunnel Properties'. Below the heading is a sub-heading 'Please supply the following parameters to tune the TCP connections. If Write Acceleration is enabled, ensure that flows will not load balanced across multiple ISLs.' and a small graphic of a server rack. The configuration area includes:

- Max Bandwidth:** A text box containing '1000' and a dropdown menu showing '1..1000 Mb'.
- Min Bandwidth:** A slider between 'Shared' and 'Dedicated' with a green marker, and a text box containing '500' followed by 'Mb'.
- Estimated RTT (RoundTrip Time):** A text box containing '1000' and a dropdown menu showing '0..300000 us', with a 'Measure' button to its right.
- Checkboxes:** Three unchecked checkboxes: 'Write Acceleration', 'Enable Optimum Compression', and 'Enable XRC Emulator'.

 At the bottom of the window are three buttons: 'Back', 'Next', and 'Cancel'. A vertical number '274886' is visible on the right side of the window frame.

- Step 9** Check the **Write Acceleration** check box to enable FCIP write acceleration on this FCIP link. See the “[FCIP Write Acceleration](#)” section on page 2-26.
- Step 10** Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the “[FCIP Compression](#)” section on page 2-33.
- Step 11** Check the Enable XRC Emulator check box to enable XRC emulator on this FCIP link. For more information on XRC Emulator, see the *Cisco Fabric Manager Fabric Configuration Guide*.
- Step 12** Click **Next**.
- Step 13** Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link (see [Figure 2-14](#)).



Note If FICON is enabled/FICON VSAN is present on both the switches, the [Figure 2-16](#) is displayed, otherwise [Figure 2-17](#) is displayed.

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-14 Create FCIP ISL

FCIP Wizard - /SAN/Fabric_sw-manual-9513

5 of 5: Create FCIP ISL

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.
NOTE: the FCIP link may take time to appear in map.

Between Switch sw-vegas-101 (fcip13 over gigE9/1[Secure])

And Switch sw-manual-9513 (fcip13 over gigE13/1[Secure])

Attributes

Port VSAN: 1 1..4093

Trunk Mode: nonTrunk trunk auto

Back Finish Cancel

240945

Figure 2-15 Enter FICON Port Address

FCIP Wizard - /SAN/Fabric_c-186

5 of 5: Create FCIP ISL

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.
NOTE: the FCIP link may take time to appear in map.

Between Switch v-190 (fcip6 over gigE2/1)

FICON Port Address 0xe0..0xf9

And Switch V-172.22.31.184 (fcip6 over gigE3/1)

FICON Port Address 0xe0..0xf9

Attributes

Port VSAN: 1 1..4093

Trunk Mode: nonTrunk trunk auto

FICON Tape Acceleration: VSAN List:

Back Finish Cancel

240947

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-16 Create FCIP ISL

FCIP Wizard - /SAN/Fabric_sw-manual-9513

5 of 5: Create FCIP ISL

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.
NOTE: the FCIP link may take time to appear in map.

Between Switch sw-vegas-101 (fcip13 over gigE9/1[Secure])

And Switch sw-manual-9513 (fcip13 over gigE13/1[Secure])

Attributes

Port VSAN: 1 1..4093

Trunk Mode: nonTrunk trunk auto

Back Finish Cancel

240946

Figure 2-17 Enter FICON Port Address

FCIP Wizard - /SAN/Fabric_c-186

5 of 5: Create FCIP ISL

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.
NOTE: the FCIP link may take time to appear in map.

Between Switch v-190 (fcip6 over gigE2/1)

FICON Port Address 0xe0..0xf9

And Switch V-172.22.31.184 (fcip6 over gigE3/1)

FICON Port Address 0xe0..0xf9

Attributes

Port VSAN: 1 1..4093

Trunk Mode: nonTrunk trunk auto

FICON Tape Acceleration: VSAN List:

Back Finish Cancel

240947

Send documentation comments to fm-docfeedback@cisco.com

Step 14 Click **Finish** to create this FCIP link.

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. Each Gigabit Ethernet interface can have three active FCIP links at one time.

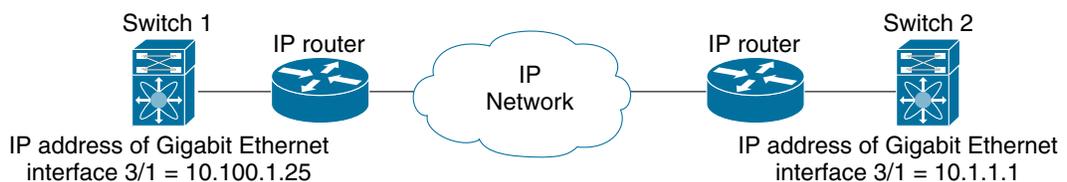
To configure an FCIP link, follow these steps on both switches:

-
- Step 1** Configure the Gigabit Ethernet interface.
 - Step 2** Create an FCIP profile, and then assign the Gigabit Ethernet 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 a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile. You can assign IPv4 or IPv6 addresses to the interfaces. [Figure 2-18](#) shows an example configuration.

Figure 2-18 Assigning Profiles to Each Gigabit Ethernet Interface



To create an FCIP profile in switch 1, follow these steps:

-
- Step 1** Verify that you are connected to a switch that contains an IPS module.
 - Step 2** From Fabric Manager, choose **Switches > ISLs > FCIP** in the Physical Attributes pane. From Device Manager, choose **FCIP** from the IP menu.
 - Step 3** Click the **Create Row** button in Fabric Manager or the **Create** button on Device Manager to add a new profile.
 - Step 4** Enter the profile ID in the ProfileId field.
 - Step 5** Enter the IP address of the interface to which you want to bind the profile.
 - Step 6** Modify the optional TCP parameters, if desired. Refer to Fabric Manager Online Help for explanations of these fields.

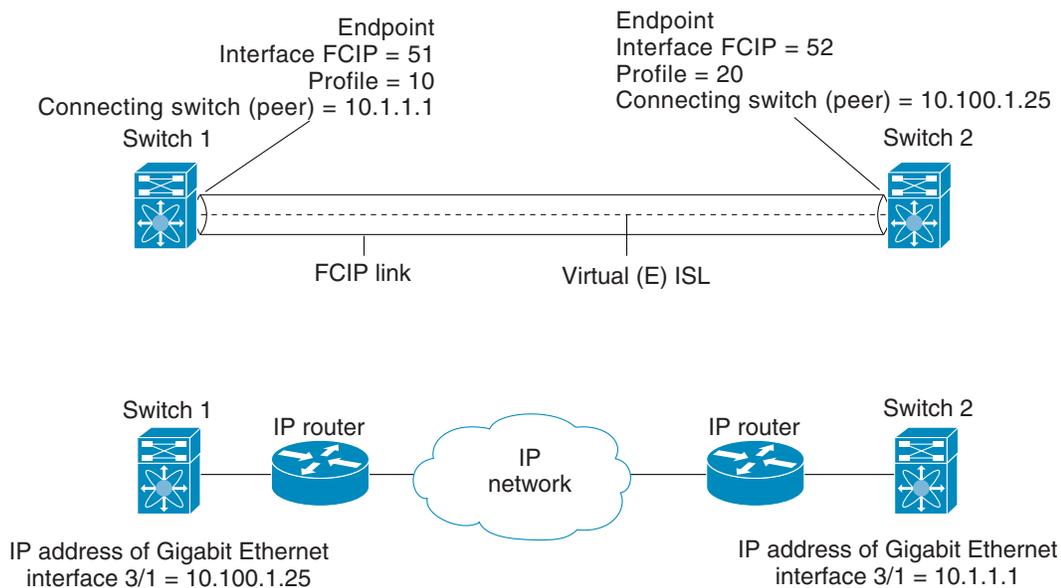
Send documentation comments to fm-docfeedback@cisco.com

- Step 7** (Optional) Click the **Tunnels** tab and modify the remote IP address in the Remote IPAddress field for the endpoint to which you want to link.
- Step 8** Enter the optional parameters, if required. See the “[FCIP Profiles](#)” section on page 2-4 for information on displaying FCIP profile information.
- Step 9** Click **Apply Changes** icon to save these changes.

Creating FCIP Links

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules or MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (see [Figure 2-19](#)).

Figure 2-19 Assigning Profiles to Each Gigabit Ethernet Interface



91562

Verifying Interfaces and Extended Link Protocol

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All 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 Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—Trunking mode and trunk-allowed VSAN list.

Send documentation comments to fm-docfeedback@cisco.com

To verify the FCIP interfaces and Extended Link Protocol (ELP) on Device Manager, follow these steps:

-
- Step 1** Make sure you are connected to a switch that contains an IPS module.
 - Step 2** Select **FCIP** from the Interface menu.
 - Step 3** Click the **Interfaces** tab if it is not already selected. You see the FCIP Interfaces dialog box.
 - Step 4** Click the **ELP** tab if it is not already selected. You see the FCIP ELP dialog box.
-

Checking Trunk Status

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configuration at the two ends of an ISL, between two switches, determines the trunking state of the link and the port modes at both ends.

To check the trunk status for the FCIP interface on Device Manager, follow these steps:

-
- Step 1** Make sure you are connected to a switch that contains an IPS module.
 - Step 2** Select **FCIP** from the IP menu.
 - Step 3** Click the **Trunk Config** tab if it is not already selected. You see the FCIP Trunk Config dialog box. This shows the status of the interface.
 - Step 4** Click the **Trunk Failures** tab if it is not already selected. You see the FCIP Trunk Failures dialog box.
-

Launching Cisco Transport Controller

Cisco Transport Controller (CTC) is a task-oriented tool used to install, provision, and maintain network elements. It is also used to troubleshoot and repair NE faults.

To launch CTC using Fabric Manager, follow these steps:

-
- Step 1** Right-click an ISL carrying optical traffic in the fabric.
 - Step 2** Click **Element Manager**.
 - Step 3** Enter the URL for the Cisco Transport Controller.
 - Step 4** Click **OK**.
-

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 section includes the following topics:

Send documentation comments to fm-docfeedback@cisco.com

- [Configuring TCP Parameters, page 2-18](#)

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, page 2-18](#)
- [Keepalive Timeout, page 2-18](#)
- [Maximum Retransmissions, page 2-18](#)
- [Path MTUs, page 2-19](#)
- [Selective Acknowledgments, page 2-19](#)
- [Window Management, page 2-19](#)
- [Monitoring Congestion, page 2-19](#)
- [Estimating Maximum Jitter, page 2-20](#)
- [Buffer Size, page 2-20](#)

Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds (msec).

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 This command 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 intervals. If no response is received for these eight probes and the connection remains idle throughout, that FCIP link is automatically closed.

**Note**

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

Maximum Retransmissions

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

Send documentation comments to fm-docfeedback@cisco.com

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.

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.

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 tunnels, WAN limitations). Maximum bandwidth should be the total bandwidth minus all other traffic going across that link.

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.

Send documentation comments to fm-docfeedback@cisco.com

- 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.



Note

The default burst size is 50 KB.



Tip

We recommend that this feature remain 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.

Estimating Maximum Jitter

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

You can configure the maximum estimated jitter in microseconds by the packet sender. The estimated variation should not include network queuing delay. By default, this parameter is enabled in Cisco MDS switches when IPS modules or MPS-14/2 modules are present.

The default value is 1000 microseconds for FCIP interfaces.

Buffer Size

You can define the required additional buffering—beyond the normal send window size—that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default FCIP buffer size is 0 KB.



Note

Use the default if the FCIP traffic is passing through a high throughput WAN link. If you have a mismatch in speed between the Fibre Channel link and the WAN link, then time stamp errors occur in the DMA bridge. In such a situation, you can avoid time stamp errors by increasing the buffer size.

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, page 2-21](#)
- [Assigning a Peer IP Address, page 2-21](#)
- [Configuring Active Connections, page 2-22](#)

Send documentation comments to fm-docfeedback@cisco.com

- [Enabling Time Stamp Control, page 2-22](#)
- [FCIP B Port Interoperability Mode, page 2-23](#)
- [Quality of Service, page 2-25](#)

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

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. The peer IP switch information initiates (creates) 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 establish an FCIP link with the peer, you can use 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.

Assigning a Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can also specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection. You can specify an IPv4 address or an IPv6 address.

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

-
- Step 1** Expand **ISLs** and select **FCIP** in the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab. You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the **NumTCPCon** field to the number of TCP connections from this FCIP link.
 - Step 8** (Optional) Check the **Enable** check box in the Time Stamp section and set the Tolerance field.
 - Step 9** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

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

-
- Step 1** From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane.

Send documentation comments to fm-docfeedback@cisco.com

You see the FCIP profiles and links in the Information pane.

From Device Manager, choose **IP > FCIP**. You see the FCIP dialog box.

- Step 2** Click the **Tunnels** tab. You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the **NumTCPCon** field to the number of TCP connections from this FCIP link.
 - Step 8** (Optional) Check the **Enable** check box in the Time Stamp section and set the Tolerance field.
 - Step 9** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

Configuring Active Connections

You can configure the required 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.

Enabling Time Stamp Control

You can instruct the switch to discard packets that are outside the specified time. 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. If a packet arrives within a 2000 millisecond interval (+ or -2000 msec) from the network time, that packet is accepted.



Note

The default value for packet acceptance is 2000 microseconds. If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the *Cisco NX-OS Fundamentals Configuration Guide* for more information).



Tip

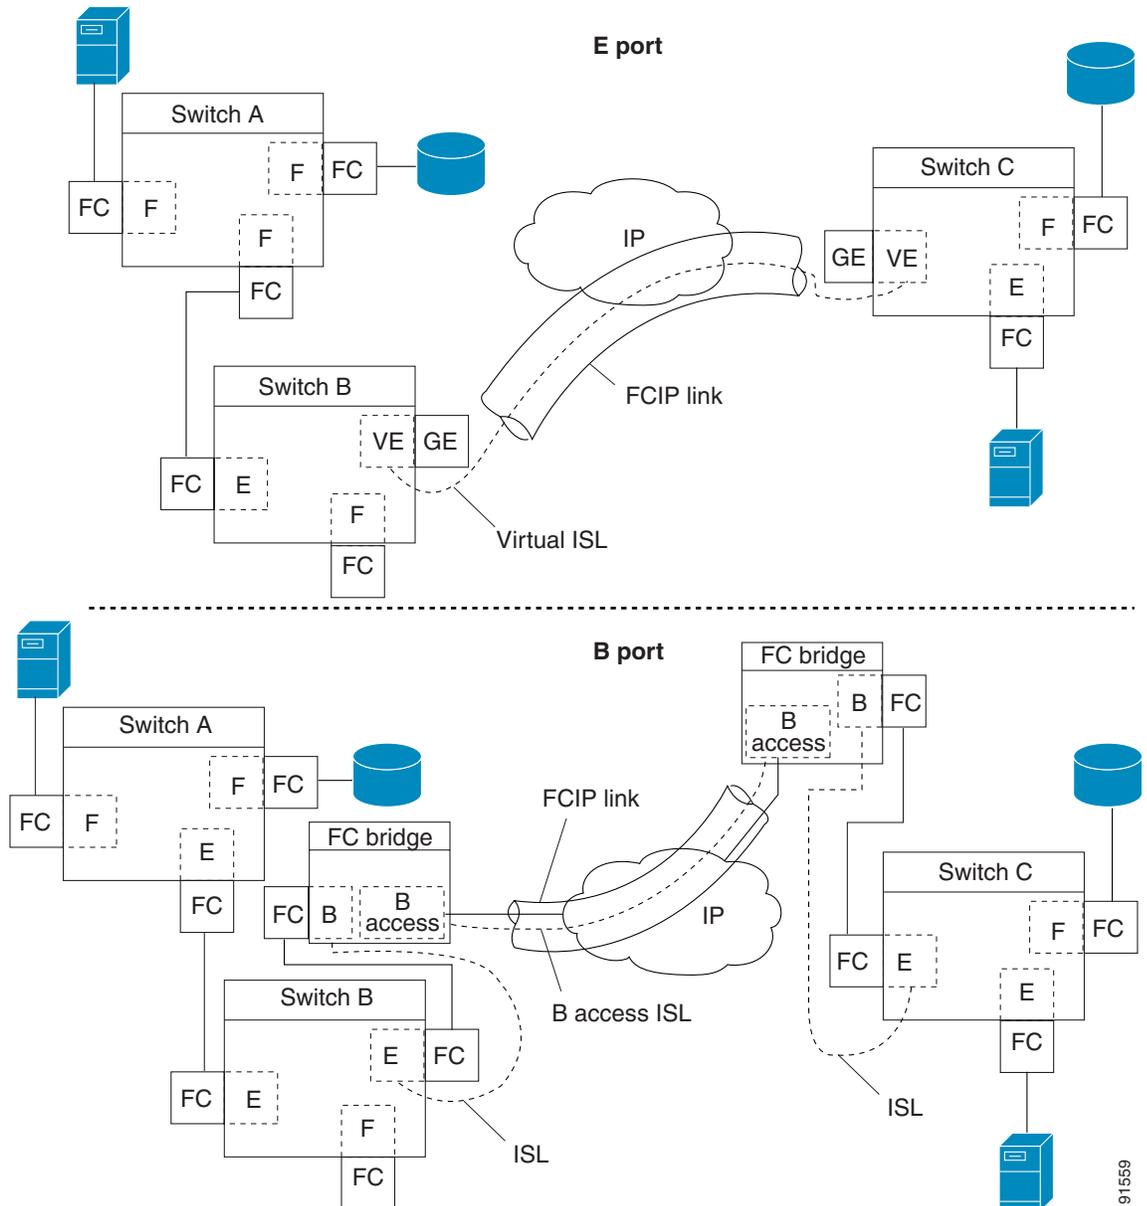
Do not enable time stamp control on an FCIP interface that has tape acceleration or write acceleration configured.

[Send documentation comments to fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)

FCIP B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. [Figure 2-20](#) shows a typical SAN extension over an IP network.

Figure 2-20 FCIP B Port and Fibre Channel E Port



B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not

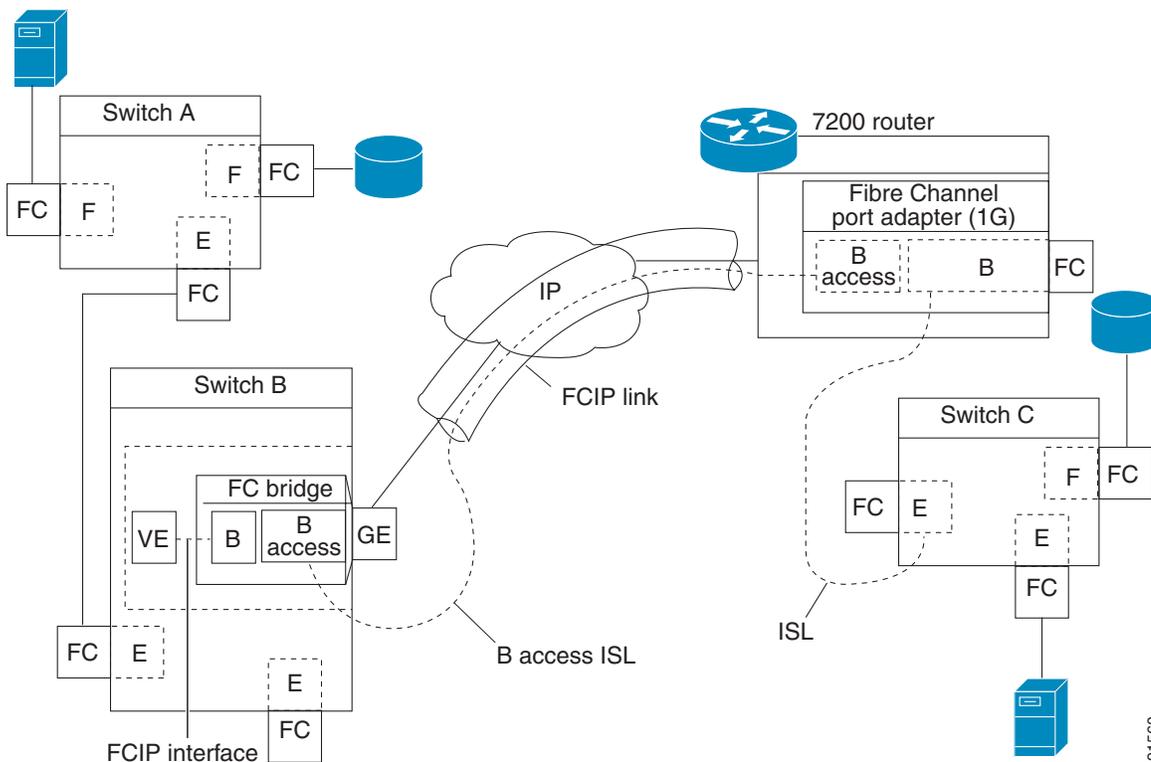
Send documentation comments to fm-docfeedback@cisco.com

interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over an FCIP link, B ports use a B access ISL*.

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see [Figure 2-21](#)).

Figure 2-21 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, eliminating the need for local bridge devices.

Configuring B Ports

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and

Send documentation comments to fm-docfeedback@cisco.com

Fibre Channel fabric shortest path first (FSPF) routing. The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface.

When an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To enable B port mode using Fabric Manager, follow these steps:

-
- Step 1** Choose **ISLs > FCIP** from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**. You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab.
You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the NumTCPCon field to the number of TCP connections from this FCIP link.
 - Step 8** Check the **Enable** check box in the B Port section of the dialog box and optionally check the **KeepAlive** check box if you want a response sent to an ELS Echo frame received from the FCIP peer.
 - Step 9** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

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.

Configuring E Ports

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN
- Trunk mode and trunk allowed VSANs
- PortChannels
- FSPF

Send documentation comments to fm-docfeedback@cisco.com

- Fibre Channel domains (fcdomains)
- Importing and exporting the zone database from the adjacent switch

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN
See the *Cisco Fabric Manager Fabric Configuration Guide*.
- Trunk mode and trunk allowed VSANs
See the *Cisco Fabric Manager Interfaces Configuration Guide*.
- PortChannels
 - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
FCIP links and Fibre Channel links cannot be combined in one PortChannel.
See the *Cisco Fabric Manager Security Configuration Guide*.
- FSPF
See the *Cisco Fabric Manager Fabric Configuration Guide*
- Fibre Channel domains (fcdomains)
See the *Cisco Fabric Manager System Management Configuration Guide*.
- Importing and exporting the zone database from the adjacent switch
See the *Cisco Fabric Manager System Management Configuration Guide*.

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, page 2-26](#)
- [Configuring FCIP Write Acceleration, page 2-28](#)
- [FCIP Tape Acceleration, page 2-28](#)
- [Configuring FCIP Tape Acceleration, page 2-33](#)
- [FCIP Compression, page 2-33](#)

FCIP Write Acceleration

The FCIP write acceleration feature enables you to significantly improve 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 minimizing the impact of WAN latency for write operations.



Note

The write 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 tunnel the write acceleration feature will be turned operationally off.



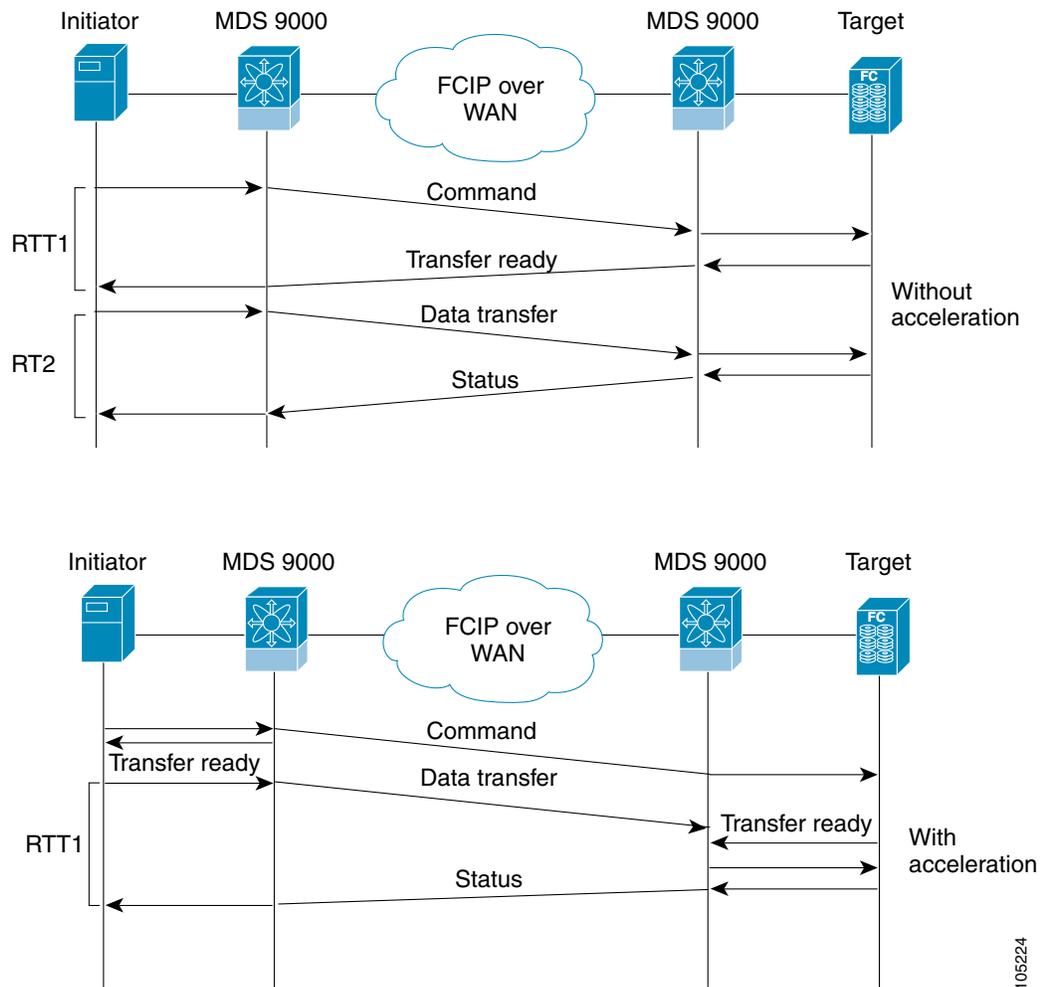
Note

IBM Peer to Peer Remote Copy (PPRC) is not supported with FCIP Write Acceleration.

Send documentation comments to fm-docfeedback@cisco.com

The WRITE command (see Figure 2-22), without write acceleration requires two round-trip transfers (RTT), while the 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 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 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 2-22 FCIP Link Write Acceleration



FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel configured with channel mode active. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or failed WRITE or READ operations.



Do not enable time stamp control on an FCIP interface with write acceleration configured.

Send documentation comments to fm-docfeedback@cisco.com

**Note**

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).

**Caution**

In Cisco MDS SAN-OS Release 2.0(1b) and later and NX-OS Release 4.x, FCIP write acceleration with FCIP ports as members of PortChannels are not compatible with the FCIP write acceleration in earlier releases.

Configuring FCIP Write Acceleration

You can enable FCIP write acceleration when you create the FCIP link using the FCIP Wizard.

To enable write acceleration on an existing FCIP link, follow these steps:

- Step 1** Choose **ISLs > FCIP** from the Physical Attributes pane on Fabric Manager.
You see the FCIP profiles and links in the Information pane.
On Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
- Step 2** Click the **Tunnels (Advanced)** tab.
You see the FCIP link information (see [Figure 2-23](#)).

Figure 2-23 *FCIP Tunnels (Advanced) Tab*

Switch	ProfileId	Interface	Timestamp Enable	Timestamp Tolerance	NumConn	Passive	QoS Control	QoS Data	IP Compression	Write Accelerator	Write Accelerator Oper
sw172-22-46-174	3	fcip3	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	none	<input checked="" type="checkbox"/>	False
sw172-22-46-174	4	fcip4	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	none	<input type="checkbox"/>	False
sw172-22-46-174	7	fcip7	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	high-comp-ratio(1.3)	<input type="checkbox"/>	False
sw172-22-46-174	8	fcip9	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	high-throughput(1.3)	<input type="checkbox"/>	False

- Step 3** Check or uncheck the **Write Accelerator** check box.
- Step 4** Choose the appropriate compression ratio from the **IP Compression** drop-down list.
- Step 5** Click the **Apply Changes** icon to save these changes.

FCIP Tape Acceleration

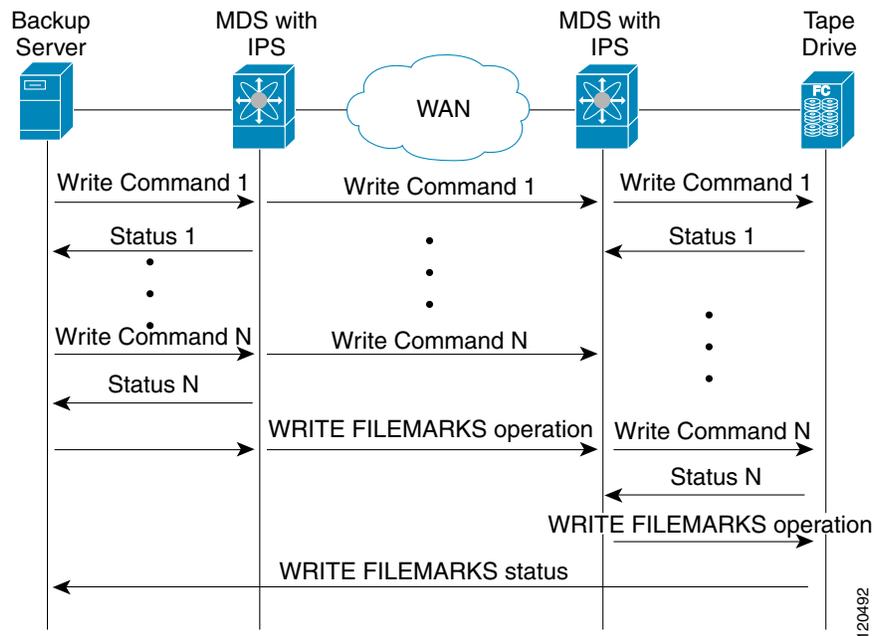
The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations. The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS NX-OS provides both tape write and read acceleration.

Send documentation comments to fm-docfeedback@cisco.com

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 tunnel over a long-distance WAN link. It impacts backup, restore, and restore 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 (see [Figure 2-24](#)) 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 tunnel 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 2-24 FCIP Link Tape Acceleration for Write Operations



At the tape end of the FCIP tunnel, 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

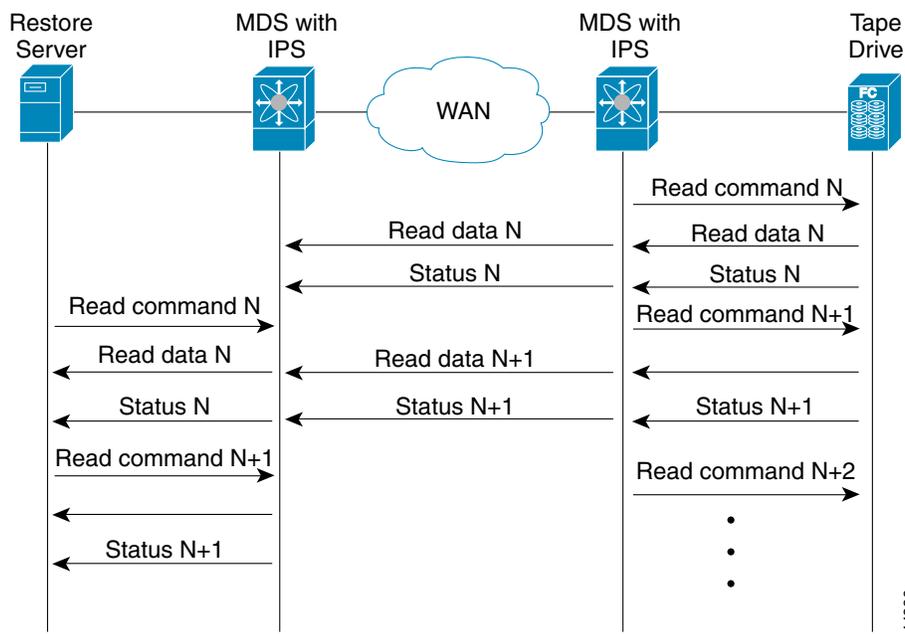
In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape library environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these sessions. You need to 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.

Send documentation comments to fm-docfeedback@cisco.com

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 (see [Figure 2-25](#)) 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 tunnel 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 2-25 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 tunnel, the tape acceleration feature is turned operationally off.



Tip

FCIP tape acceleration does not work if the FCIP port is part of a PortChannel 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.

Send documentation comments to fm-docfeedback@cisco.com**Caution**

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 tunnel, the tunnel 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).

**Tip**

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

**Tip**

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

**Note**

If one end of the FCIP tunnel is running Cisco MDS SAN-OS Release 3.0(1) or later and NX-OS Release 4.x, and the other end is running Cisco MDS SAN-OS Release 2.x, and tape acceleration is enabled, then the FCIP tunnel will run only tape write acceleration, not tape-read acceleration.

**Note**

In Cisco MDS NX-OS Release 4.2(1), the FCIP Tape Acceleration feature is not supported on FCIP back-to-back connectivity between MDS switches.

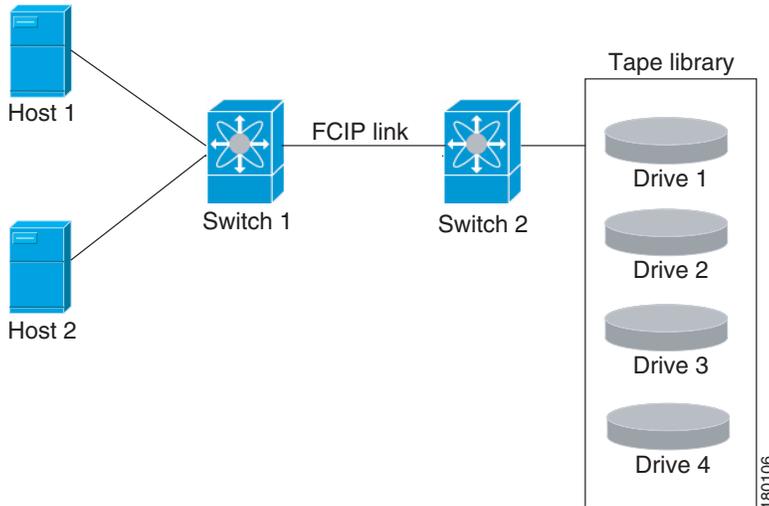
Tape Library LUN Mapping for FCIP Tape Acceleration

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.

[Figure 2-26](#) 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.

Send documentation comments to fm-docfeedback@cisco.com

Figure 2-26 FCIP LUN Mapping Example



For the mappings described in [Table 2-1](#) and [Table 2-2](#), Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

[Table 2-1](#) describes correct tape library LUN mapping.

Table 2-1 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

[Table 2-2](#) describes incorrect tape library LUN mapping.

Table 2-2 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 Drive 1 and Drive 2, 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 [Table 2-3](#).

Send documentation comments to fm-docfeedback@cisco.com

Table 2-3 Correct LUN Mapping Example with Multiple Host Access

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

Configuring FCIP Tape Acceleration

To enable FCIP tape acceleration using Fabric Manager, follow these steps:

-
- Step 1** From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab. You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the profile ID in the ProfileID field and the tunnel ID in the TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **TapeAccelerator** check box.
 - Step 7** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

FCIP Compression

The FCIP compression feature allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled. When enabled, the software defaults to using the **auto** mode (if a mode is not specified).

Mode1 and Mode3 compression modes are deprecated in Cisco MDS NX-OS Release 5.0(1a) and later. The MSM-18/4 and SSN-16 modules support Auto and Mode2 compression modes. Both 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.



Note

The **auto** mode (default) selects the appropriate compression scheme based on the card type and bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters).

Send documentation comments to fm-docfeedback@cisco.com

If both ends of the FCIP link are running NX-OS Release 4.x and 5.0(1a) and later, and you enable compression at one end of the FCIP tunnel, be sure to enable it at the other end of the link.

Default Settings

Table 2-4 lists the default settings for FCIP parameters.

Table 2-4 **Default FCIP Parameters**

Parameters	Default
TCP default port for FCIP	3225
minimum-retransmit-time	200 msec
Keepalive timeout	60 sec
Maximum retransmissions	4 retransmissions
PMTU discovery	Enabled
pmtu-enable reset-timeout	3600 sec
SACK	Enabled
max-bandwidth	1 Gbps
min-available-bandwidth	500 Mbps
round-trip-time	1 msec
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 is enabled
special-frame	Disabled
FCIP timestamp	Disabled
acceptable-diff range to accept packets	+/- 2000 msec
B port keepalive responses	Disabled
Write acceleration	Disabled
Tape acceleration	Disabled