CHAPTER 40

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 specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

Note For information on configuring Gigabit Ethernet interfaces, see Chapter 46, “Configuring IPv4 for Gigabit Ethernet Interfaces.”

This chapter includes the following sections:

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). See Figure 40-1.

Figure 40-1 Fibre Channel SANs Connected by FCIP

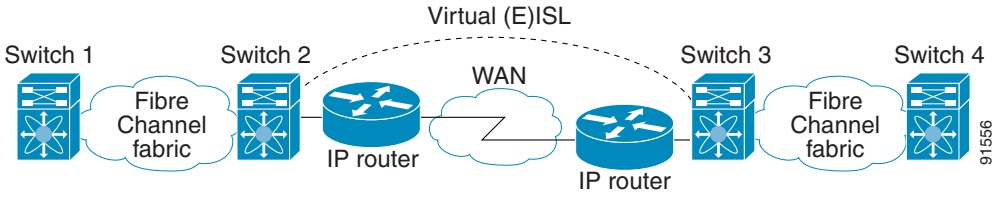

FCIP uses TCP as a network layer transport.

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

To configure IPS modules or MPS-14/2 modules for FCIP, you should have a basic understanding of the following concepts:

FCIP and VE Ports

Figure 40-2 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 40-2).

Figure 40-2 FCIP Links and Virtual ISLs

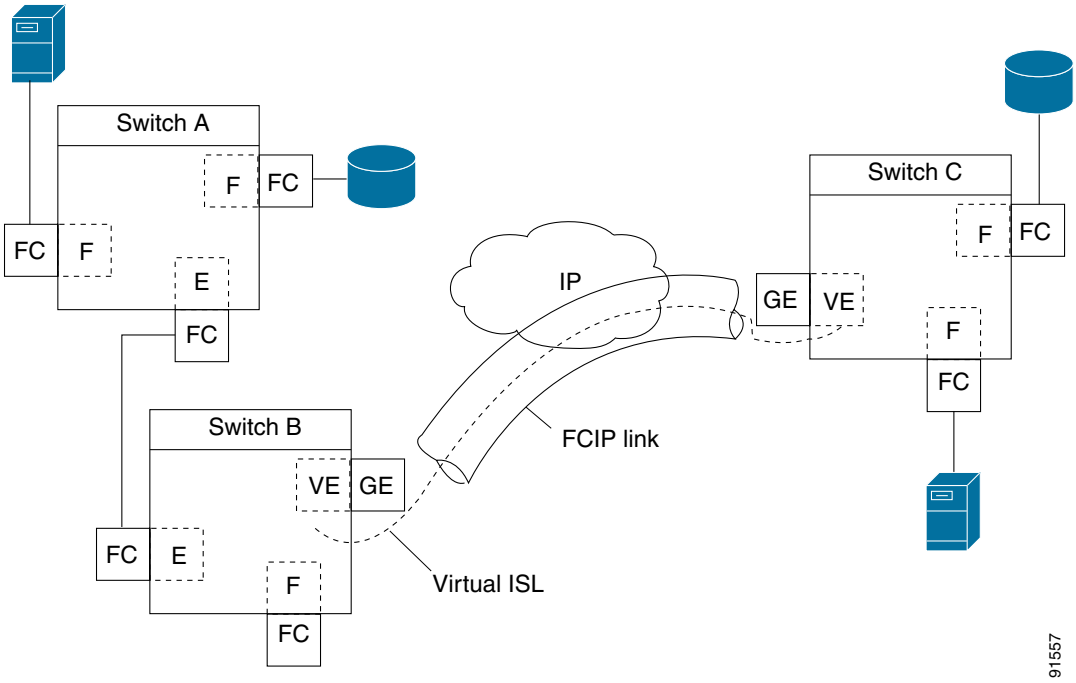

See the "Configuring E Ports" section on page 40-24.

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.

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 40-3).

Figure 40-3 FCIP Profile and FCIP Links

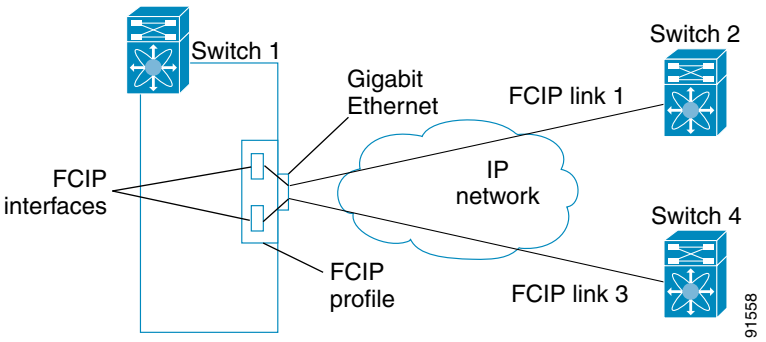

FCIP Interfaces

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 40-5
- FSPF, page 40-5
- VRRP, page 40-6
- Ethernet PortChannels, page 40-6

Fibre Channel PortChannels

Figure 40-4 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 40-4 PortChannel-Based Load Balancing

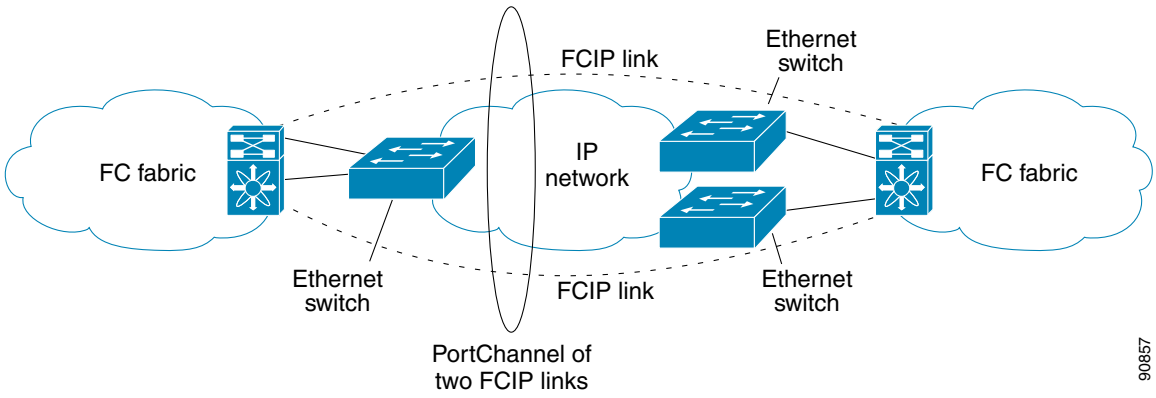

The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- 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 40-5 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 40-5 FSPF-Based Load Balancing

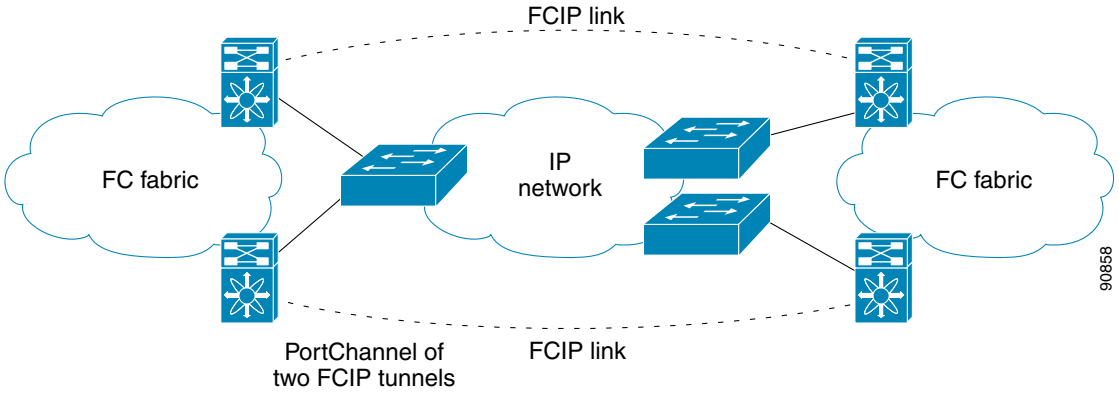

The following characteristics set FSPF solutions apart from other solutions:

- 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 link.

VRRP

Figure 40-6 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 40-6 VRRP-Based High Availability

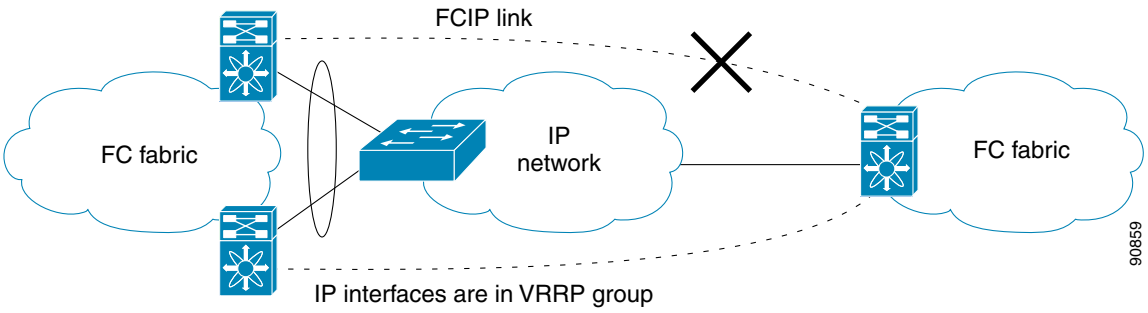

The following characteristics set VRRP solutions apart from other solutions:

- 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 40-7 displays an Ethernet PortChannel-based high- availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 40-7 Ethernet PortChannel-Based High Availability

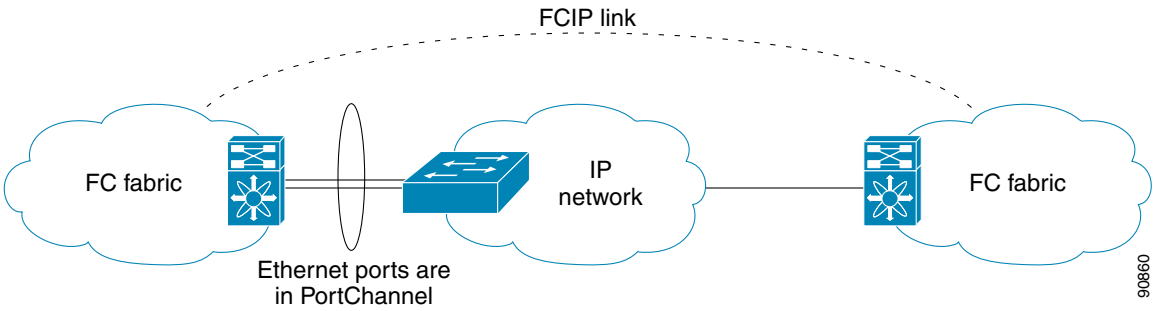

The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- 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.

- 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. On the other hand, 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 just 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 the "Configuring Gigabit Ethernet High Availability" section on page 45-5). 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 (see the "Compatibility Check" section on page 17-11). The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see Figure 40-8).

Figure 40-8 PortChannels at the Fibre Channel and Ethernet Levels

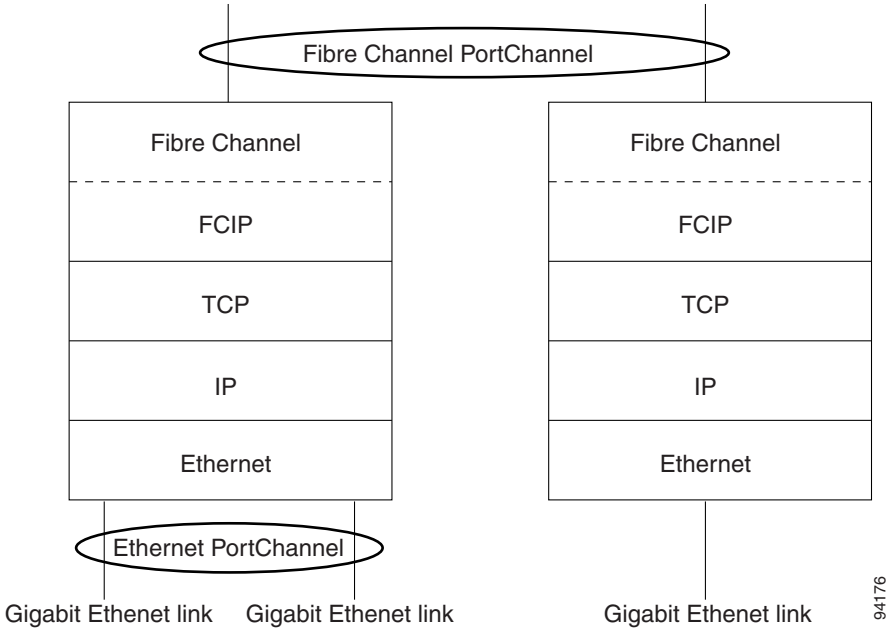

To configure Fibre Channel PortChannels, see Chapter 17, "Configuring PortChannels." To configure Ethernet PortChannels, see the "Configuring High Availability" section on page 9-1.

Configuring FCIP

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

Enabling FCIP

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 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 Chapter 3, "Obtaining and Installing Licenses").

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcip enable`	Enables FCIP on that switch.
	`switch(config)# no fcip enable`	Disables (default) FCIP on that switch.

Note If FICON is enabled/FICON VSAN is present on both the switches, the Figure 40-15 is displayed, otherwise Figure 40-14 is displayed.

Basic FCIP Configuration

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

Step 1 Configure the Gigabit Ethernet interface.

See the Chapter 46, "Configuring IPv4 for Gigabit Ethernet Interfaces."

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 40-9 shows an example configuration.

Figure 40-9 Assigning Profiles to Each Gigabit Ethernet Interface

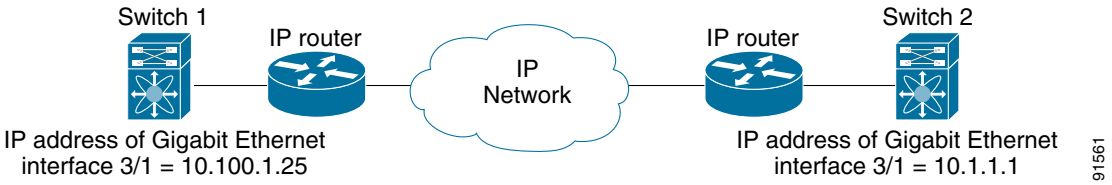

To create an FCIP profile in switch 1 in Figure 40-9, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch1(config)#`	Enters configuration mode.
Step 2	`switch1(config)# fcip profile 10` `switch1(config-profile)#`	Creates a profile for the FCIP connection. The valid range is from 1 to 255.
Step 3	`switch1(config-profile)# ip address 10.100.1.25`	Associates the profile (10) with the local IPv4 address of the Gigabit Ethernet interface (3/1).

To assign FCIP profile in switch 2 in Figure 40-9, follow these steps

	Command	Purpose
Step 1	`switch2# config terminal` `switch2(config)#`	Enters configuration mode.
Step 2	`switch2(config)# fcip profile 20` `switch2(config-profile)#`	Creates a profile for the FCIP connection.
Step 3	`switch2(config-profile)# ip address 10.1.1.1`	Associates the profile (20) with the local IPv4 address of the Gigabit Ethernet interface.

Displaying FCIP Profile Information

Example 40-1 Displays FCIP Profiles

```
switch# show fcip profile

-------------------------------------------------------------------------------
ProfileId       Ipaddr          TcpPort
-------------------------------------------------------------------------------
1               10.10.100.150   3225
2               10.10.100.150   3226
40              40.1.1.2        3225
100             100.1.1.2       3225
200             200.1.1.2       3225
```

Example 40-2 Displays the Specified FCIP Profile Information

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

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 40-10).

Figure 40-10 Assigning Profiles to Each Gigabit Ethernet Interface

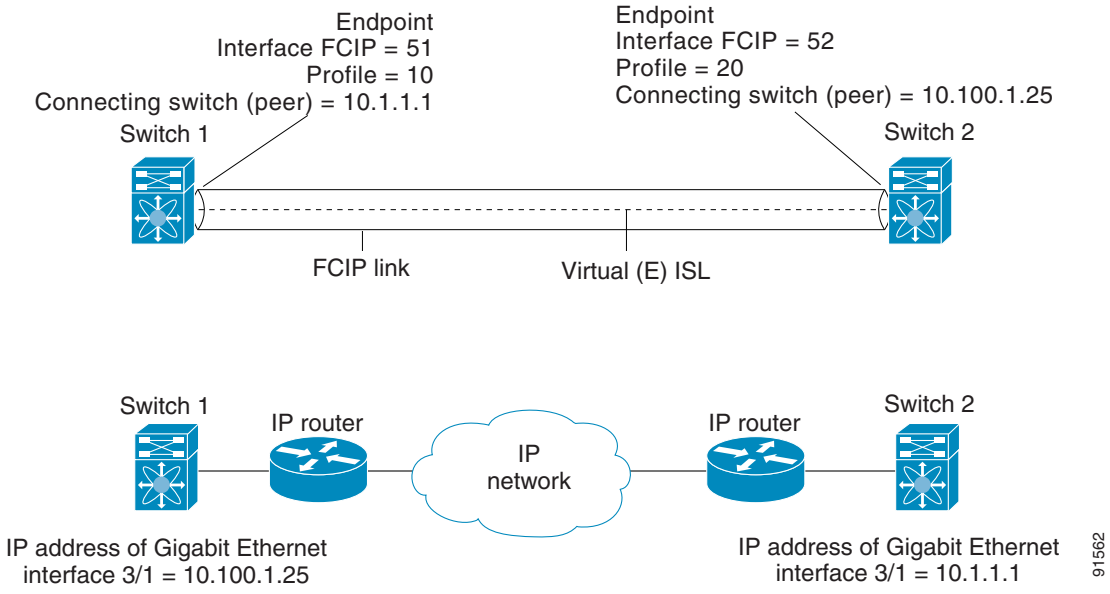

To create FCIP link endpoint in switch 1, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch1(config-if)# use-profile 10`	Assigns the profile (10) to the FCIP interface.

	Command	Purpose
Step 4	`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.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

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

	Command	Purpose
Step 1	`switch2# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch2(config)# interface fcip 52` `switch2(config-if)#`	Creates an FCIP interface (52).
Step 3	`switch2(config-if)# use-profile 20`	Binds the profile (20) to the FCIP interface.
Step 4	`switch2(config-if)# peer-info ip address 10.100.1.25`	Assigns the peer IPv4 address information (10.100.1.25 for switch 1) to the FCIP interface.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

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

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch1(config-if)# use-profile 10`	Assigns the profile (10) to the FCIP interface.
Step 4	`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.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

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

	Command	Purpose
Step 1	`switch2# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch2(config)# interface fcip 52` `switch2(config-if)#`	Creates an FCIP interface (52).
Step 3	`switch2(config-if)# use-profile 20`	Binds the profile (20) to the FCIP interface.
Step 4	`switch2(config-if)# peer-info ip address 10.100.1.25`	Assigns the peer IPv4 address information (10.100.1.25 for switch 1) to the FCIP interface.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

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.

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

Configuring TCP Listener Ports

To configure TCP listener ports, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcip profile 20` `switch(config-profile)#`	Creates the profile (if it does not already exist) and enters profile configuration submode. The valid range is from 1 to 255.

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

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

	Command	Purpose
Step 1	`switch(config-profile)# port 5000`	Associates the profile with the local port number (5000).
	`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 following TCP parameters.

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, this value is 200 milliseconds (msec).

To configure the minimum retransmit time, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp min-retransmit-time 500`	Specifies the minimum TCP retransmit time for the TCP connection to be 500 msec. The default is 200 msec and the range is from 200 to 5000 msec.
	`switch(config-profile)# no tcp min-retransmit-time 500`	Reverts the minimum TCP retransmit time to the factory default of 200 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. 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.

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

	Command	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.
	`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:

	Command	Purpose
Step 1	`switch(config-profile)# tcp max-retransmissions 6`	Specifies the maximum number of retransmissions (6). The range is from 1 to 8 retransmissions.
	`switch(config-profile)# no tcp max-retransmissions 6`	Reverts to the default of 4 retransmissions.

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:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp pmtu-enable`	Disables PMTU discovery.
	`switch(config-profile)# tcp pmtu-enable`	Enables (default) PMTU discovery with the default value of 3600 seconds.
	`switch(config-profile)# tcp pmtu-enable reset-timeout 90`	Specifies the PMTU reset timeout to 90 seconds. The range is 60 to 3600 seconds.
	`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:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp sack-enable`	Disables SACK.
	`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, keeping in mind other traffic that might be going across this link (for example, other FCIP tunnels, WAN limitations)–in other words, maximum bandwidth should be the total bandwidth minus all other traffic going across that link.

To configure window management, follow these steps:

	Command	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 msec.
	`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 msec.
	`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 msec.

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.

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.

To change the CWM defaults, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp cwm`	Disables congestion monitoring.
	`switch(config-profile)# tcp cwm`	Enables congestion monitoring and sets the burst size to its default size.
	`switch(config-profile)# tcp cwm burstsize 30`	Changes the burst size to 30 KB. The valid range is from 10 to 100 KB.
	`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.

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.

To configure the maximum jitter value, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp max-jitter`	Disables delay jitter estimation.
	`switch(config-profile)# tcp max-jitter`	Enables the delay jitter feature and sets the time to its factory default.
	`switch(config-profile)# tcp max-jitter 300`	Changes the time to 300 microseconds. The valid range is from 0 to 10000 microseconds.
	`switch(config-profile)# no tcp max-jitter 2500`	Leaves the delay jitter feature in an enabled state but changes the time to its factory default (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.

To set the buffer size, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp send-buffer-size 5000`	Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 16384 KB.
	`switch(config-profile)# no tcp send-buffer-size 5000`	Reverts the switch to its factory default. The default is 0 KB.

Displaying FCIP Profile Configuration Information

Use the **show fcip profile** command to display FCIP profile configuration information.

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

:

	Command	Purpose
Step 1	`switch(config-profile)# tcp send-buffer-size 5000`	Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 16384 KB.
	`switch(config-profile)# no tcp send-buffer-size 5000`	Reverts the switch to its factory default. The default is 0 KB.

Displaying FCIP Profile Configuration Information

Use the **show fcip profile** command to display FCIP profile configuration information.

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

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:

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

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

	Command	Purpose
Step 1	`switch# config terminal`	Enters configuration mode.
Step 2	`switch(config)# interface fcip 100`	Creates an FCIP interface (100).

Configuring Peers

To establish an FCIP link with the peer, you can use one of two options:

- Peer IP address—Configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- Special frames—Configures one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the switch WWN (sWWN) and profile ID along with the IP address.

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, follow these steps:

	Command	Purpose
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.
	`switch(config-if)# no peer-info ipaddr 10.10.1.1`	Deletes the assigned peer port information.
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.
	`switch(config-if)# no peer-info ipaddr 10.1.1.1 port 3000`	Deletes the assigned peer port information.
Step 3	`switch(config-if)# no shutdown`	Enables the interface.

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

	Command	Purpose
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.
	`switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a`	Deletes the assigned peer port information.
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.
	`switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a port 3000`	Deletes the assigned peer port information.
Step 3	`switch(config-if)# ipv6 enable`	Enables IPv6 processing on the interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Special Frames

You can alternatively establish an FCIP link with a peer using an optional protocol called *special frames*. When special frames are enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exchanged to discover and authenticate the link.

By default, the special frame feature is disabled. You must enable special frames on the interfaces on both peers to establish the FCIP link.

Note Refer to the Fibre Channel IP standards for further information on special frames.

Tip Special frame negotiation provides an additional authentication security mechanism because the link validates the WWN of the peer switch.

To enable special frames, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00`	Enables special frames and sets the peer WWN as specified. **Note** The peer WWN is the WWN of the peer switch. Use the **show wwn switch** command to obtain the peer WWN.
	`switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00`	Disables special frames (default).
Step 2	`switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155`	Enables special frames and sets the peer WWN and the profile ID (155).
	`switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155`	Disables special frames (default).
Step 3	`switch(config-if)# no shutdown`	Enables the interface.

Active Connections

You can configure the required mode for initiating a TCP connection. By default, active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection rather waits for the peer to connect to it.

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 enable the passive mode, follow these steps:

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

Number of TCP Connections

You can specify the number of TCP connections from an FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure one or two TCP connections. For example, the Cisco PA-FC-1G Fibre Channel port adapter, which has only one (1) TCP connection, interoperates with

any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, then the software handles it and proceeds with just one connection.

To specify the TCP connection attempts, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# tcp-connection 1`	Specifies the number of TCP connections. Valid values are 1 or 2.
	`switch(config-if)# no tcp-connection 1`	Reverts to the factory set default of two attempts.
Step 2	`switch(config-if)# no shutdown`	Enables the interface.

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 "NTP Configuration" section on page 5-18).

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:

	Command	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 msec.
	`switch(config-if)# no time-stamp`	Disables (default) time stamps.
Step 2	`switch(config-if)# time-stamp acceptable-diff 4000`	Configures the packet acceptance time. The valid range is from 500 to 10,000 msec.
	`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 3	`switch(config-if)# no shutdown`	Enables the interface.

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 40-11 shows a typical SAN extension over an IP network.

Figure 40-11 FCIP B Port and Fibre Channel E Port

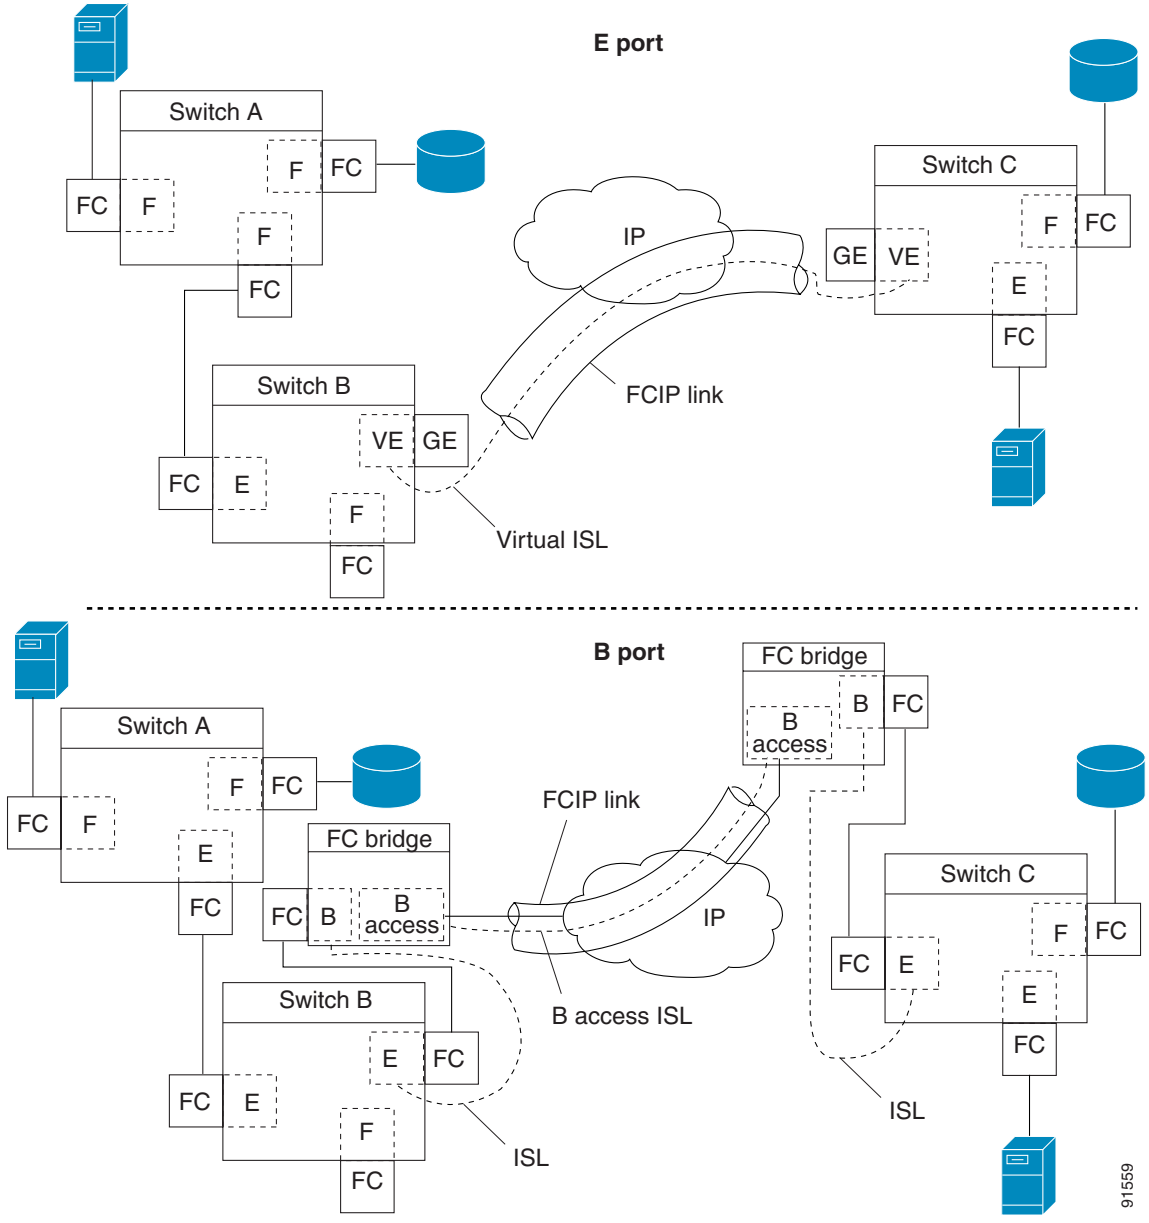

B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does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

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 40-12).

Figure 40-12 FCIP Link Terminating in a B Port Mode

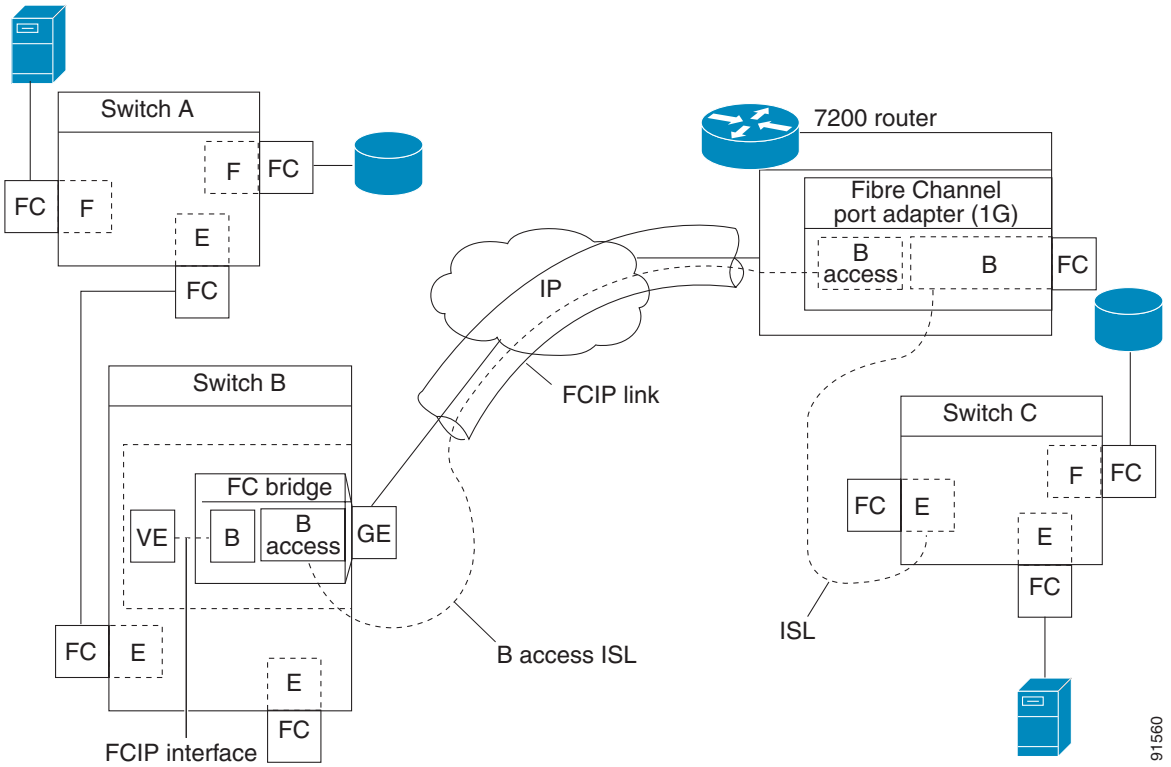

The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, eliminating the need for local bridge devices.

Configuring B Ports

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, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# bport`	Enables B port mode on the FCIP interface.
	`switch(config-if)# no bport`	Reverts to E port mode on the FCIP interface (default).
Step 2	`switch(config-if)# bport-keepalive`	Enables the reception of keepalive responses sent by a remote peer.
	`switch(config-if)# no bport-keepalive`	Disables the reception of keepalive responses sent by a remote peer (default).

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, follow these steps:

	Command	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.
	`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 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 Chapter 20, “Configuring and Managing VSANs”).
- Trunk mode and trunk allowed VSANs (see Chapter 16, “Configuring Trunking”).
- PortChannels (see Chapter 39, “Configuring Port Security”):
  - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
  - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF (see Chapter 26, “Configuring Fibre Channel Routing Services and Protocols”).
- Fibre Channel domains (fcdomains) (see Chapter 18, “Configuring Domain Parameters.”).
- Importing and exporting the zone database from the adjacent switch (see Chapter 24, “Configuring and Managing Zones”).

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. See Example 40-3 through Example 40-6.

Example 40-3 Displays the FCIP Summary

```
switch# show fcip summary
--------------------------------------------------------------------------------
Tun prof    Eth-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
--------------------------------------------------------------------------------
10  91   GE4/1      3.3.3.2         UP   N N N  N    N  1000M/1000M  2000
11  11   GE3/1.601  30.1.1.2       DOWN  N N N  N    N   1000M/500M  1000
12  12   GE3/1.602  30.1.2.2       DOWN  N N N  N    N   1000M/500M  1000
13  0              0.0.0.0        DOWN  N N N  N    N
14  0              0.0.0.0        DOWN  N N N  N    N
15  0              0.0.0.0        DOWN  N N N  N    N
16  0              0.0.0.0        DOWN  N N N  N    N
17  0              0.0.0.0        DOWN  N N N  N    N
18  0              0.0.0.0        DOWN  N N N  N    N
19  0              0.0.0.0        DOWN  N N N  N    N
20  92   GE4/2      3.3.3.1         UP   N N N  N    N  1000M/1000M  2000
21  21   GE3/2.601  30.1.1.1       DOWN  N N N  N    N   1000M/500M  1000
22  22   GE3/2.602  30.1.2.1       DOWN  N N N  N    N   1000M/500M  1000
```

Example 40-4 Displays the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 10
fcip10 is up
    Hardware is GigabitEthernet
    Port WWN is 20:d0:00:0c:85:90:3e:80
    Peer port WWN is 20:d4:00:0c:85:90:3e:80
    Admin port mode is auto, trunk mode is on
    Port mode is E, FCID is 0x720000
    Port vsan is 91
    Speed is 1 Gbps
    Using Profile id 91  (interface GigabitEthernet4/1)
    Peer Information
      Peer Internet address is 3.3.3.2 and port is 3225
    Write acceleration mode is off
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is 256 KBytes
    IP Compression is disabled
    Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
    TCP Connection Information
      50529025 Active TCP connections
         Local 0.0.0.7:6, Remote 0.0.0.200:0
      0 host table full 0 target entries in use
      211419104 Attempts for active connections, 1500 close of connections
    TCP Parameters
      Path MTU 124160 bytes
      Current retransmission timeout is 124160 ms
      Round trip time: Smoothed 127829 ms, Variance: 14336
```

```
      Advertized window: Current: 0 KB, Maximum: 14 KB, Scale: 14336
      Peer receive window: Current: 0 KB, Maximum: 0 KB, Scale: 51200
      Congestion window: Current: 14 KB, Slow start threshold: 49344 KB
      Current Send Buffer Size: 206463 KB, Requested Send Buffer Size: 429496728
3 KB
      CWM Burst Size: 49344 KB
    5 minutes input rate 491913172779207224 bits/sec, 61489146597400903 bytes/se
c, 0 frames/sec
    5 minutes output rate 491913175298921320 bits/sec, 61489146912365165 bytes/s
ec, 14316551 frames/sec
      5702 frames input, 482288 bytes
         5697 Class F frames input, 481736 bytes
         5 Class 2/3 frames input, 552 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      5704 frames output, 482868 bytes
         5698 Class F frames output, 482216 bytes
         6 Class 2/3 frames output, 652 bytes
         0 Error frames
```

Example 40-5 Displays Detailed FCIP Interface Standard Counter Information

```
switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    5 minutes input rate 207518944 bits/sec, 25939868 bytes/sec, 12471 frames/sec
    5 minutes output rate 205340328 bits/sec, 25667541 bytes/sec, 12340 frames/sec
      2239902537 frames input, 4658960377152 bytes
         18484 Class F frames input, 1558712 bytes
         2239884053 Class 2/3 frames input, 4658958818440 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      2215051484 frames output, 4607270186816 bytes
         18484 Class F frames output, 1558616 bytes
         2215033000 Class 2/3 frames output, 4607268628200 bytes
         0 Error frames
```

Example 40-6 Displays the FCIP Interface Description

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

The txbytes 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 (see Example 40-7).

Example 40-7 Displays Brief FCIP Interface Counter Information

```
switch# show interface fcip 3 counters brief
--------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   -----------------------------  -----------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
--------------------------------------------------------------------------------
fcip3              9        0                     9        0
```

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 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.

In Figure 40-13, the WRITE command 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 40-13 FCIP Link Write Acceleration

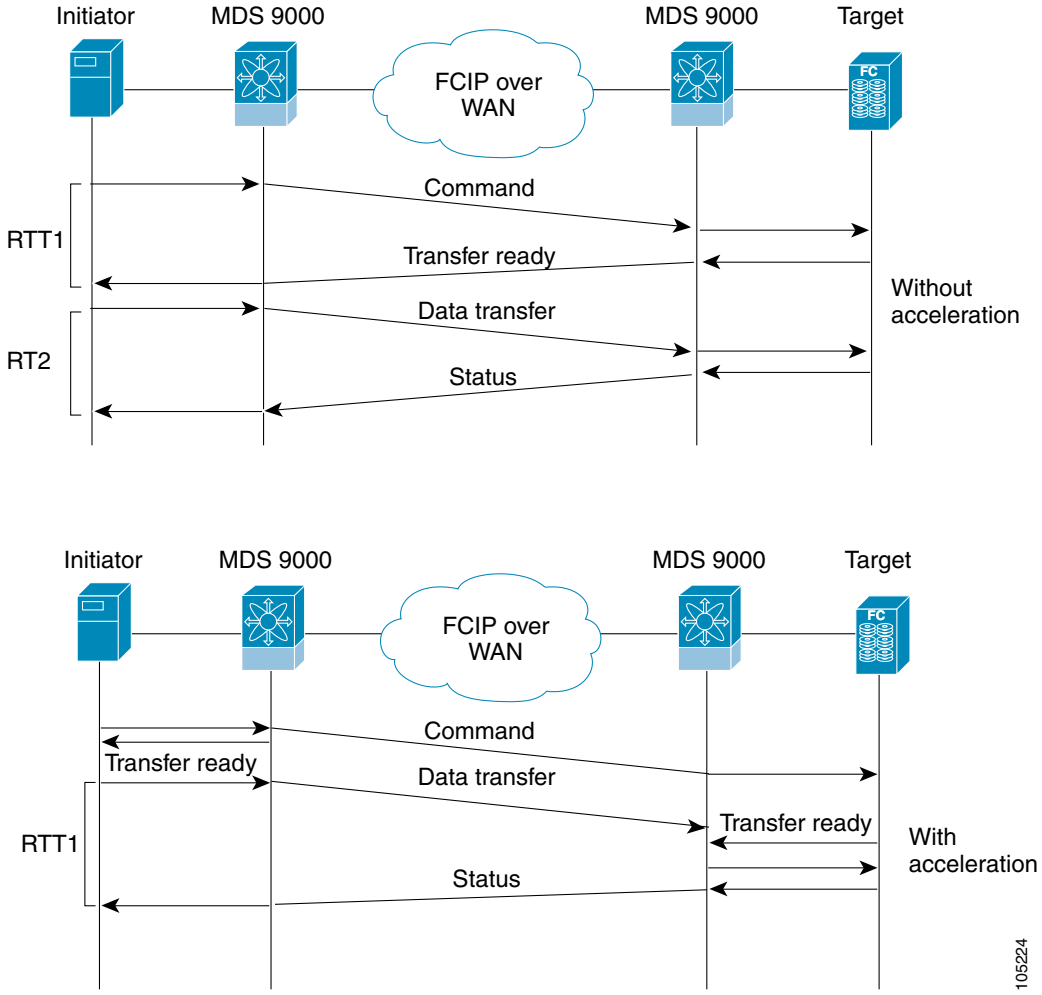

Tip FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel configured with channel mode active. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration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

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

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 FCIP write acceleration with FCIP ports as members of PortChannels in Cisco MDS SAN-OS Release 2.0(1b) and later are incompatible with the FCIP write acceleration in earlier releases.

Configuring FCIP Write Acceleration

To enable write acceleration, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch1(config-if)# write-accelerator`	Enables write acceleration.
	`switch1(config-if)# no write-accelerator`	Disables write acceleration (default).

Displaying Write Acceleration Activity Information

Example 40-8 through Example 40-10 show how to display information about write acceleration activity.

Example 40-8 Displays 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
```

Example 40-9 Displays 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
```

Example 40-10 Displays Detailed FCIP Interface Write Acceleration Counter Information, if Enabled

```
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

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS SAN-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 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 Figure 40-14 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 40-14 FCIP Link Tape Acceleration for Write Operations

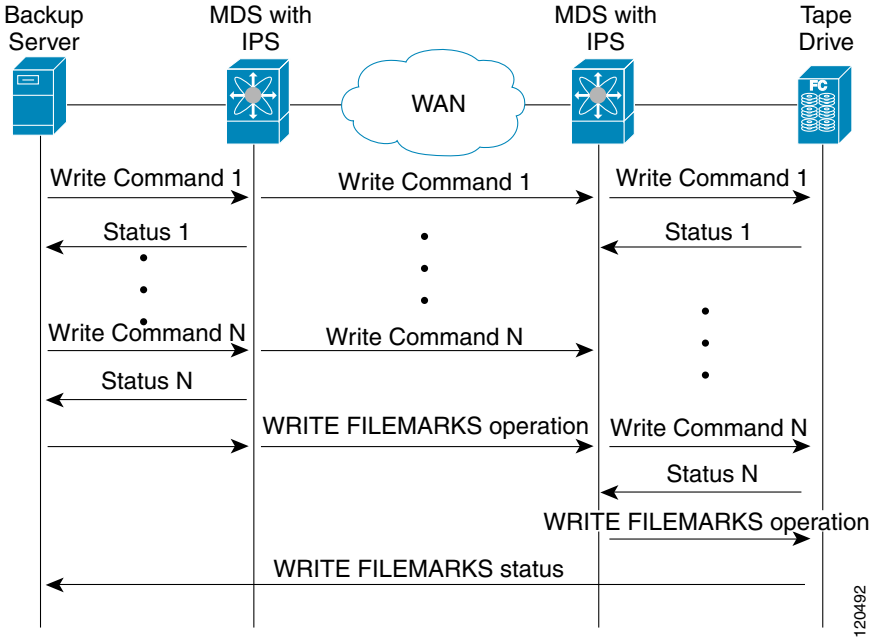

At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

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. The workaround is to keep the FCIP link disabled for a couple of minutes before enabling the link. Note that this does not apply to tape environments where the tape drives are either direct FC attached or exported as LUN 0.

The Cisco SAN-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 SAN-OS software.

In an example of tape acceleration for read operations, the restore server in Figure 40-15 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 40-15 FCIP Link Tape Acceleration for Read Operations

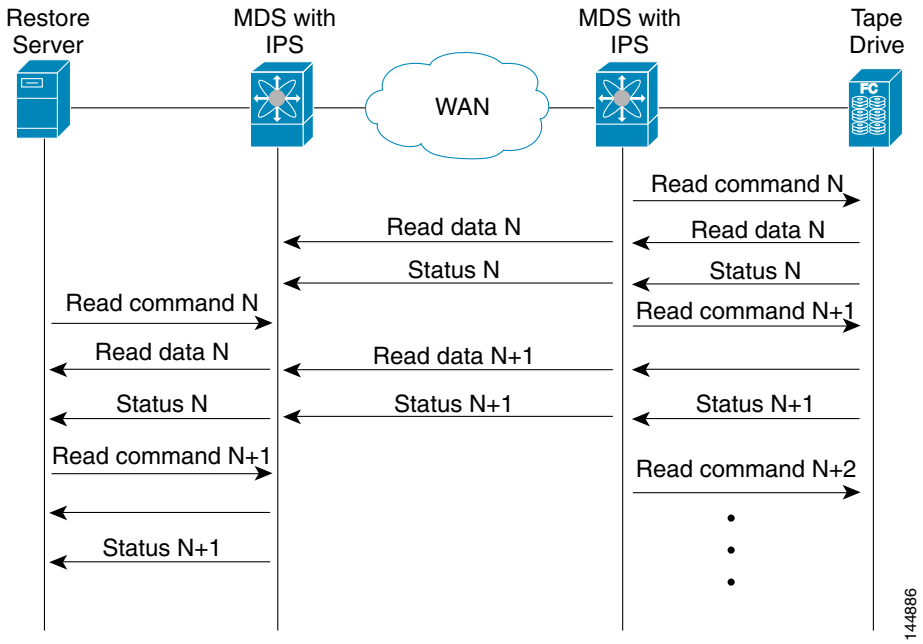

The Cisco SAN-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco SAN-OS software recovers from any other errors.

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.

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 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.

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 40-16 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 40-16 FCIP LUN Mapping Example

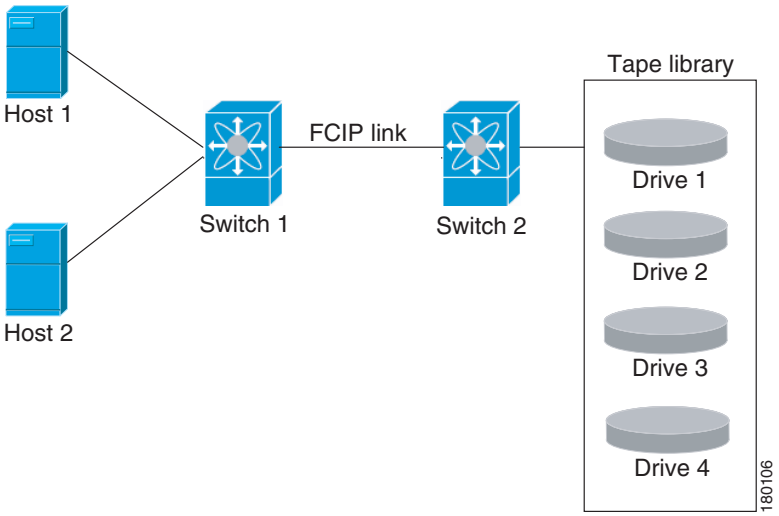

For the mappings described in Table 40-1 and Table 40-2, Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

Table 40-1 describes correct tape library LUN mapping.

Table 40-1 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 40-2 describes incorrect tape library LUN mapping.

Table 40-2 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 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 Table 40-3.

Table 40-3 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, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 5` `switch1(config-if)#`	Creates an FCIP interface (5).

	Command	Purpose
Step 3	`switch1(config-if)# write-accelerator tape-accelerator`	Enables tape acceleration (and write acceleration—if not already enabled).
	`switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size auto`	Enables tape acceleration with automatic flow control (default).
	`switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size 2048`	Sets tape acceleration flow control buffer size to 2 MB.
	`switch1(config-if)# no write-accelerator tape-accelerator`	Disables tape acceleration (default) and resets the FCIP tunnel. **Note** The write acceleration feature remains enabled.
	`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 tunnel.
	`switch1(config-if)# no write-accelerator`	Disables both the write acceleration and tape acceleration features and resets the FCIP tunnel.

Displaying Tape Acceleration Activity Information

Example 40-11 through Example 40-14 show how to display information about tape acceleration activity.

Example 40-11 Displays Information About Tapes for Which Exchanges are Tape Accelerated

```
switch# show fcip tape-session summary
--------------------------------------------------------------------------------
Tunnel    Tunnel End    tape-fcid    lun        vsan    num-hosts
--------------------------------------------------------------------------------
  1       host-end       EF0001      0x0002     0001      1
  2       targ-end       650001      0x0003     0010      2
--------------------------------------------------------------------------------
```

Example 40-12 Displays 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
```

Example 40-13 Displays 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
```

Example 40-14 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 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).

Note The **"auto"** (default) mode picks 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).

Table 40-4 lists the modes used for different cards.

Table 40-4 Algorithm Classification

Mode	IPS Card	MPS 14/2 Card	18+4/9222i
mode1	SW	HW	SW
mode2	SW	**SW**	HW
mode3	SW	SW	SW

Table 40-5 Performance Settings

Bandwidth	IPS Card	MPS 14/2 Card	18+4/9222i
>25Mbps	mode1	mode1	mode2/mode3
10-25Mbps	mode2	mode2	mode2/mode3
10Mbps	mode3	mode3	mode2/mode3

Note The Cisco **MDS 9216**i and **9222i** Switches also supports the IP compression feature. The integrated supervisor module has the same hardware components that are available in the **MPS-14/2** module.

Caution The compression modes in Cisco SAN-OS Release 2.0(1b) and later are incompatible with the compression modes in Cisco SAN-OS Release 1.3(1) and earlier.

Tip While upgrading from Cisco SAN-OS Release 1.x to Cisco SAN-OS Release 2.0(1b) or later, we recommend that you disable compression before the upgrade procedure, and then enable the required mode after the upgrade procedure.

If both ends of the FCIP link are running Cisco SAN-OS Release 2.0(1b) or later and you enable compression at one end of the FCIP tunnel, be sure to enable it at the other end of the link.

Configuring FCIP Compression

To enable FCIP compression, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fcip 51` `switch(config-if)#`	Creates an FCIP interface (51).

	Command	Purpose
Step 3	`switch(config-if)# ip-compression mode3`	Enables high compression for low bandwidth links.
	`switch(config-if)# ip-compression mode3`	Defaults to using the **auto** mode.
	`switch(config-if)# no ip-compression`	Disables (default) the FCIP compression feature.

Displaying FCIP Compression Information

Example 40-15 and Example 40-16 show how to display FCIP compression information.

Example 40-15 Displays 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
```

Example 40-16 Displays the Compression Engine Statistics for the MPS-14/2 Module

```
switch#  show ips stats hw-comp all
HW Compression Statistics for port GigabitEthernet3/1
    Compression stats
      0 input bytes,  0 output compressed bytes
      0 input pkts,    0 output compressed pkts
    Decompression stats
      0 input compressed bytes, 0 output bytes
      0 input compressed pkts,   0 output pkts
    Passthru stats
      0 input bytes, 0 output bytes
      0 input pkts,   0 output pkts
    Miscellaneous stats
      32 min input pktlen,  32 max input pktlen
      28 min output pktlen, 28 max output pktlen
      0 len mismatch,      0 incomplete processing
      0 invalid result,    0 invalid session drop
      0 comp expanded
HW Compression Statistics for port GigabitEthernet3/2
    Compression stats
      0 input bytes,  0 output compressed bytes
      0 input pkts,    0 output compressed pkts
    Decompression stats
      0 input compressed bytes, 0 output bytes
      0 input compressed pkts,   0 output pkts
    Passthru stats
      0 input bytes, 0 output bytes
      0 input pkts,   0 output pkts
    Miscellaneous stats
      32 min input pktlen,  32 max input pktlen
      28 min output pktlen, 28 max output pktlen
      0 len mismatch,      0 incomplete processing
      0 invalid result,    0 invalid session drop
```

```
0 comp expanded
```

Default Settings

Table 40-4 lists the default settings for FCIP parameters.

Table 40-6 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	1Gbps
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