



FCIP

FCIP (Fibre Channel over IP) is an IETF standards based protocol for connecting Fibre Channel SANs over IP based networks. FCIP encapsulates the FCP frames in a TCP/IP packet which is then sent across an IP network. FCIP can interconnect geographically dispersed SANs using the IP backbone network. In short, FCIP creates an FC tunnel over an existing IP network. The MDS 9200 and MDS 9500 series switches support FCIP, using the IPS-8, IPS-4 and the 14+2 blades.

This chapter explains how to enable and configure FCIP and includes the following sections:

- [Enabling FCIP, page 9-1](#)
- [Configuring FCIP, page 9-1](#)
- [Tuning FCIP, page 9-5](#)

Enabling FCIP

You must enable FCIP before attempting to configure it on the switch.



Warning

If you do not run the FCIP enable command, further FCIP configuration is not possible. This command enables additional FCIP configuration options in the CLI.

```
MDS1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
MDS1(config)# FCIP enable
MDS1(config)#^Z
MDS1#
```

Configuring FCIP

The recipe below shows the configuration of FCIP on the Cisco MDS 9000 switch. [Figure 9-1](#) shows the topology.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 9-1 FCIP Topology



The topology consists of two MDS 9509 switches, each with one IPS-8 blade. The Gigabit Ethernet port 8/8 on the switch sjc7-9509-5 is connected to a Catalyst 6509. The interface has an IP address of 172.22.36.81, with a subnet mask of 255.255.254.0, and the gateway address of 172.22.36.1. Another switch, sjc7-9509-6, has the Gigabit Ethernet port 8/3 also connected to a Catalyst 6509. This interface has an IP address of 172.22.34.58, with a subnet of 255.255.254.0, and the gateway address of 172.22.34.1. In the recipe, the FCIP tunnel is established between the switch sjc7-9509-5 (GE 8/8) and sjc7-9509-6 (GE 8/3).



Warning

The IP address for the ports on the IPS blade must be on a different subnet and Ethernet segment than the management interface. The different segment can be achieved either through physical separation or using VLANs. Having different segments is critical to get FCIP working on the MDS 9000 Family switch.

To make FCIP operational between the switches, follow these steps:

- Step 1** Configure the Gigabit Ethernet port on the MDS 9000 switch sjc7-9509-5 by assigning an IP address and a subnet mask. The address and subnet mask allow the Gigabit Ethernet interface to communicate with the network.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-6(config)# interface gigabitethernet 8/3
sjc7-9509-6(config-if)# ip address 172.22.34.58 255.255.254.0
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-5#
```

- Step 2** Configure IP routes so that the two Gigabit Ethernet interfaces on switches sjc7-9509-5 and sjc7-9509-6 can communicate with each other.

In the preceding step, the Gigabit Ethernet ports are in two different subnets. As a result, explicit routes are required to allow the two switches to communicate with each other.



Note

The recommended best practice for configuring an IP route is to have a host route to each of the two Gigabit Ethernet interfaces. In other words, you should allow the two Gigabit Ethernet interfaces to communicate only with each other. You can accomplish this by designating a subnet mask of 255.255.255.255 when you add the route.

- Step 3** Use the following syntax to configure a route to allow the Gigabit Ethernet port 8/8 on the switch sjc7-9509-5 to communicate with the Gigabit Ethernet port 8/3 on switch sjc7-9509-6:

```
sjc7-9509-5# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-5(config)# ip route 172.22.34.58 255.255.255.255 172.22.36.1 interface
gigabitethernet 8/8
sjc7-9509-5(config)# ^Z
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
sjc7-9509-5#
```

In this configuration, you use the gateway 172.22.36.1 and interface GE 8/8 on switch sjc7-9509-5 to reach 172.22.34.58.

- Step 4** Similarly, use the following syntax to configure a route for GE 8/3 on switch sjc7-9509-6 to communicate with the port GE 8/8 on switch sjc7-9509-5:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-6(config)# ip route 172.22.36.81 255.255.255.255 172.22.34.1 interface
gigabitethernet 8/3
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

In this configuration, you use the gateway 172.22.34.1 and interface GE 8/3 on switch sjc7-9509-6 to reach 172.22.36.81.

- Step 5** Ping the Gigabit Ethernet interfaces to ensure that the Gigabit Ethernet ports can communicate with each other.

- From the switch prompt on sjc7-9509-5, ping the IP address of the Gigabit Ethernet interface 8/3 on switch sjc7-9509-6.
- From the switch prompt on sjc7-9509-6, ping the IP address of the Gigabit Ethernet interface 8/8 on switch sjc7-9509-5.

```
sjc7-9509-5# ping 172.22.34.58
PING 172.22.34.58 (172.22.34.58) 56(84) bytes of data.
64 bytes from 172.22.34.58: icmp_seq=1 ttl=254 time=1.04 ms
64 bytes from 172.22.34.58: icmp_seq=2 ttl=254 time=0.624 ms
64 bytes from 172.22.34.58: icmp_seq=3 ttl=254 time=0.678 ms
64 bytes from 172.22.34.58: icmp_seq=4 ttl=254 time=0.580 ms

--- 172.22.34.58 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3008ms
rtt min/avg/max/mdev = 0.580/0.732/1.046/0.184 ms
sjc7-9509-5#
```

```
sjc7-9509-6# ping 172.22.36.81
PING 172.22.36.81 (172.22.36.81): 56 data bytes
64 bytes from 172.22.36.81: icmp_seq=0 ttl=254 time=0.7 ms
64 bytes from 172.22.36.81: icmp_seq=1 ttl=254 time=0.6 ms
64 bytes from 172.22.36.81: icmp_seq=2 ttl=254 time=0.5 ms
64 bytes from 172.22.36.81: icmp_seq=3 ttl=254 time=0.6 ms
```

```
--- 172.22.36.81 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 0.5/0.6/0.7 ms
sjc7-9509-6#
```



Note

It is critical to check the connectivity between the two Gigabit Ethernet ports on each MDS 9000 Family switch IPS blade before proceeding further. A successful ping test should be sufficient to verify connectivity.

- Step 6** Measure the round-trip time between the two Gigabit Ethernet interfaces because you will need the value later in the configuration.

```
sjc7-9509-5# ips measure-rtt 172.22.34.58 interface gigabitethernet 8/8
Round trip time is 691 micro seconds (0.69 milli seconds)
sjc7-9509-5#
```

```
sjc7-9509-6# ips measure-rtt 172.22.36.81 interface gigabitethernet 8/3
Round trip time is 743 micro seconds (0.74 milli seconds)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 7** sjc7-9509-6#
Configure the FCIP profile on both the switches by creating the FCIP profile.

The profile defines the characteristics for the FCIP tunnel that you set up. The measured round-trip time is needed during the profile configuration. Round the value to the nearest whole number. In this case, the 690 microseconds is rounded to 1 millisecond. The IP address in the configuration that follows is the IP address assigned to the Gigabit Ethernet interface on the switch.



Tip

We recommend that the FCIP profile and FCIP interface numbers at both ends of a FCIP tunnel be the same for easy management.

```

sjc7-9509-5# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-5(config)# fcip profile 10
sjc7-9509-5(config-profile)# ip address 172.22.36.81
sjc7-9509-5(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth-mbps 200
round-trip-time-ms 1
sjc7-9509-5(config-if)# ^Z
sjc7-9509-5#

```

```

sjc7-9509-6# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-6(config)# fcip profile 10
sjc7-9509-6(config-profile)# ip address 172.22.34.58
sjc7-9509-6(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth-mbps 200
round-trip-time-ms 1
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#

```

- Step 8** Configure the FCIP interface on both switches.

In the FCIP interface configuration that follows, the profile used and the peer information (remote Gigabit Ethernet IP address) are specified. Compression and write acceleration can also be configured.

```

sjc7-9509-5# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-5(config)# interface fcip 1
sjc7-9509-5(config-if)# use-profile 10
sjc7-9509-5(config-if)# peer-info ipaddr 172.22.34.58
sjc7-9509-5(config-if)# no shut
sjc7-9509-5(config-if)# ^Z
sjc7-9509-5#

```

```

sjc7-9509-6# conf t
Enter configuration commands, one per line. End with CNTL/Z.
sjc7-9509-6(config)# interface fcip 1
sjc7-9509-6(config-if)# use-profile 10
sjc7-9509-6(config-if)# peer-info ipaddr 172.22.36.81
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#

```

The FCIP tunnel should be up and running now.

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 9 Use the **show interface fcip 1 brief** command to show the status of the FCIP link between the two switches.

```
sjc7-9509-5# show interface fcip 1 brief
```

```
-----
Interface Vsan Admin Admin Status Oper Profile Eth Int Port-channel
          Mode Trunk          Mode
          Mode
-----
fcip1     1     auto on   trunking TE    10 GigabitEthernet8/8 --
sjc7-9509-5#
```

```
sjc7-9509-6# show interface fcip 1 brief
```

```
-----
Interface Vsan Admin Admin Status Oper Profile Eth Int Port-channel
          Mode Trunk          Mode
          Mode
-----
fcip1     1     auto on   trunking TE    10 GigabitEthernet8/3 --
sjc7-9509-6#
```

Tuning FCIP

Implementing an FCIP tunnel is more than just creating the FCIP link and modifying the VSAN allowed list. To achieve the greatest efficiency from the link, some parameters that are specific to the underlying connection may need to be tuned. Therefore, if the FCIP link is running over a slow, 1 MB connection, the FCIP link should be tuned differently than one running over a low-latency, 1 GB connection. This section provides insight into what parameters can be used and how to use them for achieving increased efficiency and utilization from the FCIP connection.



Note

Your results may vary due to network conditions (existing utilization) and storage array and host type.

Figure 9-2 shows the topology that is used throughout this chapter:

Figure 9-2 FCIP Topology



TCP Tuning: Latency and Available Bandwidth

The latency of the link is the amount of time that it takes a packet to go from one end of the FCIP link to the other. Latency can be due to many factors, including distance and the number of devices that the packet must traverse. Even the fastest routers and switches incur some amount of latency.

Send documentation comments to mdsfeedback-doc@cisco.com.

Latency cannot be eliminated; however protocols can be tuned and MDS features such as FCIP Write Acceleration (see [FCIP Write Acceleration, page 9-7](#)) can be enabled to eliminate the effects of it. These features can be modified in the FCIP profile.

Available bandwidth is the amount of bandwidth that the FCIP link can use on the network. You need to define a maximum and a minimum value for the FCIP link to use in the FCIP profile.

- The maximum available bandwidth value is the maximum amount of bandwidth that the FCIP link should use on the network.
- The minimum available bandwidth value is used as a guideline for the minimum value. Of course, if there are serious problems on the network (such as dropped packets or congestion), the link will go slower than the minimum value. We recommend that the minimum value be set to what the minimum acceptable bandwidth that the application (EMC SRDF, IBM PPRC, and so on) requires.



Tip

For dedicated links, set the minimum and maximum available bandwidth values to be the same.

[Table 9-1](#) contains some common WAN links and their speeds. These circuits are most often used as the underlying network for an FCIP link. For example, the underlying network may be a OC3, but you may only be able to use 100 MB of that link.



Tip

When deploying FCIP, you should always involve the LAN and WAN teams to find out about the connectivity that they are providing you. If there are performance issues, they can often help you troubleshoot them from the network's standpoint. Involve them earlier rather than later.

Table 9-1 Common WAN Circuit Speeds

Circuit Name	Speed
T1	1.544 Mbps
T3	43.232 Mbps
OC-3	155 Mbps
OC-12	622 Mbps
OC-48	2.5 Gbps
OC-192	9.6 Gbps

To specify the minimum and maximum available bandwidth and the network latency, perform the following procedure. For this example, we assume that the LAN/WAN links are a dedicated Gigabit Ethernet link.

Step 1 Measure the round-trip time between the two FCIP end points.

```
sjc7-9509-5# ips measure-rtt 172.22.34.58 interface gigabitethernet 8/8
Round trip time is 5691 micro seconds (5.69 milli seconds)
```

```
sjc7-9509-6# ips measure-rtt 172.22.36.81 interface gigabitethernet 8/3
Round trip time is 5743 micro seconds (5.74 milli seconds)
```

Step 2 Modify the profile by specifying the minimum and maximum bandwidth.

```
sjc7-9509-5# conf t
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# fcip profile 2
sjc7-9509-5(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth
-mbps 1000 round-trip-time-us 5691

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# fcip profile 2
sjc7-9509-6(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth
-mbps 1000 round-trip-time-us 5743

```

FCIP Write Acceleration

To alleviate the effects of latency, the IPS-8 and IPS-4 blades support software compression. The configuration is shown below.



Note

Only in SAN-OS Release 2.0(1b) and higher should write acceleration be used with PortChannels.

Step 1 Enable write acceleration on switch sjc7-9509-5.

```

sjc7-9509-5# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# interface fcip 1
sjc7-9509-5(config)# write-accelerator
sjc7-9509-5(config)# ip-compression mode high-throughput
sjc7-9509-5(config)# ^Z
sjc7-9509-5#

```

Step 2 Enable write acceleration on switch sjc7-9509-6.

```

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface fcip 1
sjc7-9509-6(config)# write-accelerator
sjc7-9509-6(config)# ip-compression mode high-throughput
sjc7-9509-6(config)# ^Z
sjc7-9509-6#

```

FCIP Compression

FCIP based compression should be enabled when the bandwidth requirements of the application (such as SRDF, TrueCopy, or PPRC) exceed the bandwidth that the LAN or WAN link can provide.

For SAN-OS versions prior to 2.0, software compression was available on the IP Services Module using two compression modes, a high compression ratio, and high throughput. For environments where the bandwidth speed of the WAN is relatively low (such as T3/DS 3, 45 Mbps and lower speeds), the high compression ratio option should be enabled. For WAN links up to OC-3, 155 Mbps, the high throughput option should be enabled.

In SAN-OS version 2.0, three new modes of compression were introduced, along with hardware based compression using the 14+2 module. These modules are summarized in [Figure 9-3](#).

Send documentation comments to mdsfeedback-doc@cisco.com.



Note The numbers displayed in Figure 9-3 are derived using the Canterbury Corpus test suite. Your actual performance numbers may vary depending on device type and data pattern. However, the numbers shown here should be used as a guide to determine which compression mode is most appropriate.

Figure 9-3 Approximate Application Throughput with MDS 2.0 Compression

