



Configuring IP Storage

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

This chapter includes the following sections:

- [IP Storage Services Module, page 17-2](#)
- [Configuring Gigabit Ethernet Interfaces, page 17-4](#)
- [Configuring FCIP, page 17-17](#)
- [Configuring iSCSI, page 17-38](#)
- [Default IP Storage Settings, page 17-76](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)


IP Storage Services Module

The IPS services module (IPS module) allows you to use FCIP and iSCSI features. It integrates seamlessly into the Cisco MDS 9000 Family, and supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports. Each port can run FCIP and iSCSI protocols simultaneously.


- **FCIP**—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 17-1](#) depicts the FCIP scenarios in which the IPS module is used.

Figure 17-1 FCIP Scenarios



- **iSCSI**—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a MDS 9000 IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. [Figure 17-2](#) depicts the iSCSI scenarios in which the IPS module is used.

Figure 17-2 iSCSI Scenarios



Send documentation comments to mdsfeedback-doc@cisco.com

Verifying the Module Status

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

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    16     1/2 Gbps FC Module        DS-X9016             ok
4    8      IP Storage Module        DS-X9308-SMIP       ok <-----IPS module
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      ha-standby

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.1(1)      0.3         20:41:00:05:30:00:86:5e to 20:50:00:05:30:00:86:5e
4    1.1(1)      0.2         20:c1:00:05:30:00:86:5e to 20:c8:00:05:30:00:86:5e
5    1.1(1)      0.602      --
6    1.1(1)      0.602      --

Mod  MAC-Address(es)                Serial-Num
---  -
2    00-05-30-00-9f-62 to 00-05-30-00-9f-66  JAB064505YV
4    00-05-30-00-a1-ae to 00-05-30-00-a1-ba  JAB0649059h
5    00-05-30-00-9f-f6 to 00-05-30-00-9f-fa  JAB06350B1M
6    00-05-30-00-9f-f2 to 00-05-30-00-9f-f6  JAB06350B1F

* this terminal session
```

[Send documentation comments to mdsfeedback-doc@cisco.com](#)

Configuring Gigabit Ethernet Interfaces

This section includes the following topics:

- [About Gigabit Ethernet Interfaces, page 17-4](#)
- [Basic Gigabit Ethernet Configuration, page 17-4](#)
- [About VLANs for Gigabit Ethernet, page 17-5](#)
- [VLAN Configuration, page 17-6](#)
- [Interface Subnet Requirements, page 17-6](#)
- [Managing IP Routing, page 17-7](#)
- [Verifying Gigabit Ethernet Connectivity, page 17-7](#)
- [Managing ARP Caches, page 17-8](#)
- [Displaying Statistics, page 17-8](#)
- [Gigabit Ethernet High Availability, page 17-13](#)
- [Configuring CDP, page 17-16](#)
- [IPS Core Dumps, page 17-16](#)

About Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On the IPS module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

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




Tip

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

Basic Gigabit Ethernet Configuration

Figure 17-3 depicts a basic Gigabit Ethernet configuration.

Figure 17-3 Gigabit Ethernet Configuration



Send documentation comments to mdsfeedback-doc@cisco.com

To configure the Gigabit Ethernet interface for the scenario in [Figure 17-3](#), follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.100 255.255.255.0	Enters the IP address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

You can configure the switch to receive and transfer large (or jumbo) frames on a port. The default IP MTU frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased to 9000 bytes. In this example, the size was set to 3000 bytes. Independent of the MTU size, the IPS module does not pack multiple IP frames (converted to FCIP or to iSCSI).



Note

The minimum MTU size for a port running iSCSI is 620 bytes.

To configure MTU frame size, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config-if)# switchport mtu 3000 switch(config-if)#	Changes the IP maximum transmission unit (MTU) to 3000. The default is 1500.

About VLANs for Gigabit Ethernet

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

The Gigabit Ethernet port can be configured as a trunking port and uses the IEEE 802.1Q standard for VLAN encapsulation.



Note

If the IPS module is connected to a Cisco Ethernet switch, verify the following requirements:

- the Ethernet switch port is configured as a trunking port, and
- the encapsulation is set to 802.1Q and not ISL, which is the default.

Send documentation comments to mdsfeedback-doc@cisco.com

VLAN Configuration

To configure a VLAN subinterface (the VLAN ID), follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2.100 switch(config-if)#	Specifies the subinterface on which 802.1Q is used (slot2, port 2, VLAN ID 100). Note The subinterface number, 100 in this example, is the VLAN ID. The VLAN ID ranges from 1 to 4093.
Step 3	switch(config-if)# ip address 10.1.1.100 255.255.255.0	Enters the IP address (10.1.1.100) and IP mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

Interface Subnet Requirements

Gigabit Ethernet interface (major), subinterfaces (VLAN tags) and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 17-1](#)).

Table 17-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN tag can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN tags cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A VLAN tag cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	



Note

The configuration requirements in [Table 17-1](#) also applies to Ethernet PortChannels.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Managing IP Routing

To configure static IP routing through the Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1 switch(config-if)#	Enters the IP subnet (10.100.1.0 255.255.255.0) of the IP host and configures the next hop 10.1.1.1, which is the IP address of the router connected to the Gigabit Ethernet interface.

Displaying the IP Route Table

The **show ips ip route interface ethernet** command takes the ethernet interface as a parameter and returns the route table for the interface. See [Example 17-1](#).

Example 17-1 Displays the Route Table

```
switch# show ips ip route interface gig 8/1
Codes: C - connected, S - static
No default gateway
C 10.1.3.0/24 is directly connected, GigabitEthernet8/1
```

Connected (C) identifies the subnet in which the interface is configured (directly connected to the interface). Static (S) identifies the static routes that go through the router.

Verifying Gigabit Ethernet Connectivity

The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the “Using the ping Command” section on page 2-18).

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch using the **ping** command. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly. See [Example 17-2](#).

Example 17-2 Verifying Gigabit Ethernet Connectivity

```
switch# ping 10.100.1.25
PING 10.100.1.25 (10.100.1.25): 56 data bytes
64 bytes from 10.100.1.25: icmp_seq=0 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=1 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=2 ttl=255 time=0.1 ms
--- 10.100.1.25 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.1 ms
```



Note

If the connection fails, verify the following, and repeat the **ping** command:

- the IP address for the destination (IP host) is correctly configured,
- the host is active (powered on),
- the IP route is configured correctly,
- the IP host has a route to get to the Gigabit Ethernet interface subnet, and
- the Gigabit Ethernet interface is in the **up** state (use the **show interface gigabitethernet** command).

Send documentation comments to mdsfeedback-doc@cisco.com

Managing ARP Caches

Use the **show ips arp interface gigabitethernet** command to display the ARP cache on the Gigabit Ethernet interfaces. This command takes the Ethernet interface as a parameter and returns the ARP cache for that interface. See [Example 17-3](#).

Example 17-3 Displays ARP Caches

```
switch# show ips arp interface gigabitethernet 7/1
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet      20.1.1.5     3            0005.3000.9db6 ARPA    GigabitEthernet7/1
Internet      20.1.1.10    7            0004.76eb.2ff5 ARPA    GigabitEthernet7/1
Internet      20.1.1.11    16           0003.47ad.21c4 ARPA    GigabitEthernet7/1
Internet      20.1.1.12    6            0003.4723.c4a6 ARPA    GigabitEthernet7/1
Internet      20.1.1.13    13           0004.76f0.ef81 ARPA    GigabitEthernet7/1
Internet      20.1.1.14    0            0004.76e0.2f68 ARPA    GigabitEthernet7/1
Internet      20.1.1.15    6            0003.47b2.494b ARPA    GigabitEthernet7/1
Internet      20.1.1.17    2            0003.479a.b7a3 ARPA    GigabitEthernet7/1
...
```

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache. See [Examples 17-4](#) and [17-5](#).

Example 17-4 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 17-5 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

Displaying Statistics

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

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface Gigabit Ethernet** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 17-6](#).

Example 17-6 Displays the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
GigabitEthernet8/1 is up          <-----The interface is in the up state.
  Hardware is GigabitEthernet, address is 0005.3000.a98e
  Internet address is 10.1.3.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
```


Send documentation comments to mdsfeedback-doc@cisco.com

```

5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
3343 packets input, 406582 bytes
  0 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
8 packets output, 336 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors

```

Example 17-7 Displays the Gigabit Ethernet's Subinterface

```

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

```

Displaying Ethernet MAC Statistics

The **show ips stats mac interface gigabitethernet** command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 17-8](#).

Example 17-8 Displays Ethernet MAC Statistics

```

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

```

Displaying DMA-Bridge Statistics

You can display direct memory access (DMA) device statistics using the **show ips stats dma-bridge interface gigabitethernet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 17-9](#).

Example 17-9 Displays DMA-Bridge Statistics

```

switch# show ips stats dma-bridge interface gigabitethernet 7/1
Dma-bridge ASIC Statistics for port GigabitEthernet7/1
  Hardware Egress Counters

```

Send documentation comments to mdsfeedback-doc@cisco.com

```

231117 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
Hardware Ingress Counters
218255 Good, 0 protocol error, 0 header checksum error
0 FC CRC error, 0 iSCSI CRC error, 0 parity error
Software Egress Counters
231117 good frames, 0 bad header cksum, 0 bad FIFO SOP
0 parity error, 0 FC CRC error, 0 timestamp expired error
0 unregistered port index, 0 unknown internal type
0 RDL ok, 0 RDL drop (too big), 0 RDL ttl_1
3656368645 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
Software Ingress Counters
218255 Good frames, 0 header cksum error, 0 FC CRC error
0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
0 out of memory drop, 0 queue full drop
0 RDL ok, 0 RDL drop (too big)
Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]

```

This output shows all Fibre Channel frames that ingress or egress from the Gigabit Ethernet port.

Displaying TCP/IP Statistics



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

Use the **show ips stats ip interface gigabitethernet** to display and verify IP statistics. This command takes the main Gigabit Ethernet interface as a parameter and returns IP statistics for that interface. See [Example 17-10](#).

Example 17-10 Displays IP Statistics

```

switch# show ips stats ip interface gigabitethernet 4/1
Internet Protocol Statistics for port GigabitEthernet4/1
168 total received, 168 good, 0 error
0 reassembly required, 0 reassembled ok, 0 dropped after timeout
371 packets sent, 0 outgoing dropped, 0 dropped no route
0 fragments created, 0 cannot fragment

```

Use the **show ips stats tcp interface gigabitethernet** to display and verify TCP statistics. This command takes the main ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The **detail** option shows all information maintained by the interface. See [Examples 17-11](#) and [17-12](#).

Example 17-11 Displays TCP Statistics

```

switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
Connection Stats
0 active openings, 3 accepts
0 failed attempts, 12 reset received, 3 established
Segment stats
163 received, 355 sent, 0 retransmitted
0 bad segments received, 0 reset sent

TCP Active Connections

```

Local Address	Remote Address	State	Send-Q	Recv-Q
0.0.0.0:3260	0.0.0.0:0	LISTEN	0	0

Send documentation comments to mdsfeedback-doc@cisco.com

Example 17-12 Displays Detailed TCP Statistics

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

  TCP Active Connections
    Local Address      Remote Address      State      Send-Q      Recv-Q
    0.0.0.0:3260       0.0.0.0:0          LISTEN     0           0
```

Use the **show ips stats icmp interface gigabitethernet** to display and verify IP statistics. This command takes the main ethernet interface as a parameter and returns the ICMP statistics for that interface. See [Example 17-13](#).

Example 17-13 Displays ICMP Statistics

```
switch# show ips stats icmp interface gigabitethernet 4/1
ICMP Statistics for port GigabitEthernet4/1
  5 ICMP messages received
  0 ICMP messages dropped due to errors
  ICMP input histogram
    5 echo request
  ICMP output histogram
```

Send documentation comments to mdsfeedback-doc@cisco.com

5 echo reply

Send documentation comments to mdsfeedback-doc@cisco.com

Gigabit Ethernet High Availability

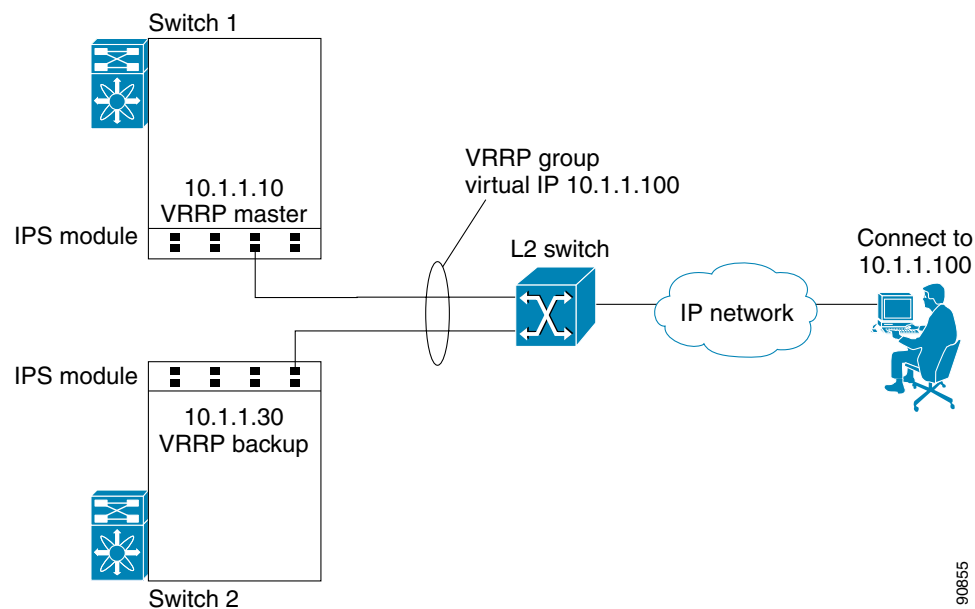
Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

Configuring VRRP

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services (see the “Configuring VRRP” section on page 16-12).

VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 17-4).

Figure 17-4 VRRP Scenario



In Figure 17-4, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module
- Interfaces across IPS modules in one switch
- Interfaces across IPS modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels
- Subinterfaces

Send documentation comments to mdsfeedback-doc@cisco.com

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

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.10 255.255.255.0	Enters the IP address (10.1.1.10) and IP mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp 100 switch(config-if-vrrp)	Creates a VR ID 100.
Step 6	switch(config-if-vrrp)# address 10.1.1.100	Configures the virtual IP address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IP address must be in the same subnet as the IP address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IP address.
Step 7	switch(config-if-vrrp)# priority 10	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	switch(config-if-vrrp)# no shutdown	Enables the VRRP protocol on the selected interface.



Note

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

Configuring Ethernet PortChannels

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

The data traffic from one TCP connection always travels on the same physical links. An Ethernet switch connecting to the MDS Gigabit Ethernet port can implement load balancing based on its IP address, its source-destination MAC address, or its IP and port. If Ethernet-based load balancing cannot be implemented for iSCSI scenarios based on the IP and port, multiple iSCSI initiators are required to take advantage of the Ethernet PortChannel feature.



Note

The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3aa protocol.


Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 17-5](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

PortChannel members must be one of these combinations: ports 1-2, ports 3-4, ports 5-6, or ports 7-8.

Figure 17-5 Ethernet PortChannel Scenario



In [Figure 17-5](#), Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel.

**Note**

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that FCIP link.

PortChannel configuration specified in [Chapter 11, “Configuring PortChannels”](#) also apply to Ethernet PortChannel configurations.

PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

To configure Ethernet PortChannels, follow these steps:

	Command	Purpose
Step 1	switch1# conf terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface port-channel 10 switch(config-if)#	Configures the specified PortChannel (10).
Step 3	switch(config-if)# ip address 10.1.1.1 255.255.255.0	Enters the IP address (10.1.1.1) and IP mask (255.255.255.0) for the PortChannel. Note A PortChannel does not have any members when first created.
Step 4	switch(config-if)# no shutdown	Enables the interface.
Step 5	switch(config)# interface gigabitethernet 9/3 switch(config-if)#	Configures the specified Gigabit Ethernet interface (slot 9, port 3).
Step 6	switch(config-if)# channel-group 10 gigabitethernet 9/3 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do “no shutdown” at both ends to bring them up switch(config-if)#	Adds Gigabit Ethernet interfaces 9/3 to channel group 10. If channel group 10 does not exist, it is created. The port is shut down.

Send documentation comments to mdsfeedback-doc@cisco.com

	Command	Purpose
Step 7	<code>switch(config-if)# no shutdown</code>	Enables the selected interface.
Step 8	<code>switch(config)# interface gigabitethernet 9/4</code> <code>switch(config-if)#</code>	Configures the specified Gigabit Ethernet interface (slot 9, port 4).
Step 9	<code>switch(config-if)# channel-group 10</code> gigabitethernet 9/4 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up	Adds Gigabit Ethernet interfaces 9/4 to channel group 10. The port is shut down.
Step 10	<code>switch(config-if)# no shutdown</code>	Enables the selected interface.



Note

Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:

- if the interface already has an IP address assigned, or
- if subinterfaces are configured on that interface.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. See the [“Configuring CDP” section on page 3-33](#).

IPS Core Dumps

IPS core dumps are different from the system’s kernel core dumps for other modules. When the IPS module’s operating system (OS) unexpectedly resets, it is sometimes useful to obtain a full copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Core dumps take up significant space and hence the level of what gets stored can be configured using one of the two options:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files).
- Full core dumps—Each full core dump consists of 75 parts (75 files). This dump cannot be saved on the supervisor module due to its large space requirement. If you choose this option, then you must configure an external TFTP server using the **system cores tftp:** command (see [“Configuring Core and Log Files” section on page 26-6](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring FCIP

This section includes the following topics:

- [About FCIP, page 17-17](#)
- [Basic FCIP Configuration, page 17-20](#)
- [Advanced FCIP Profile Configuration, page 17-22](#)
- [Advanced FCIP Interface Configuration, page 17-26](#)
- [E Port Configurations, page 17-32](#)
- [Displaying FCIP Information, page 17-33](#)
- [FCIP High Availability, page 17-35](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 17-37](#)

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 17-6](#).

Figure 17-6 *Fibre Channel SANs Connected by FCIP*

FCIP uses Transmission Control Protocol (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).

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

- [FCIP and VE Ports, page 17-18](#)
- [FCIP Link, page 17-18](#)
- [FCIP Profiles, page 17-19](#)
- [FCIP Interface, page 17-19](#)

Send documentation comments to mdsfeedback-doc@cisco.com

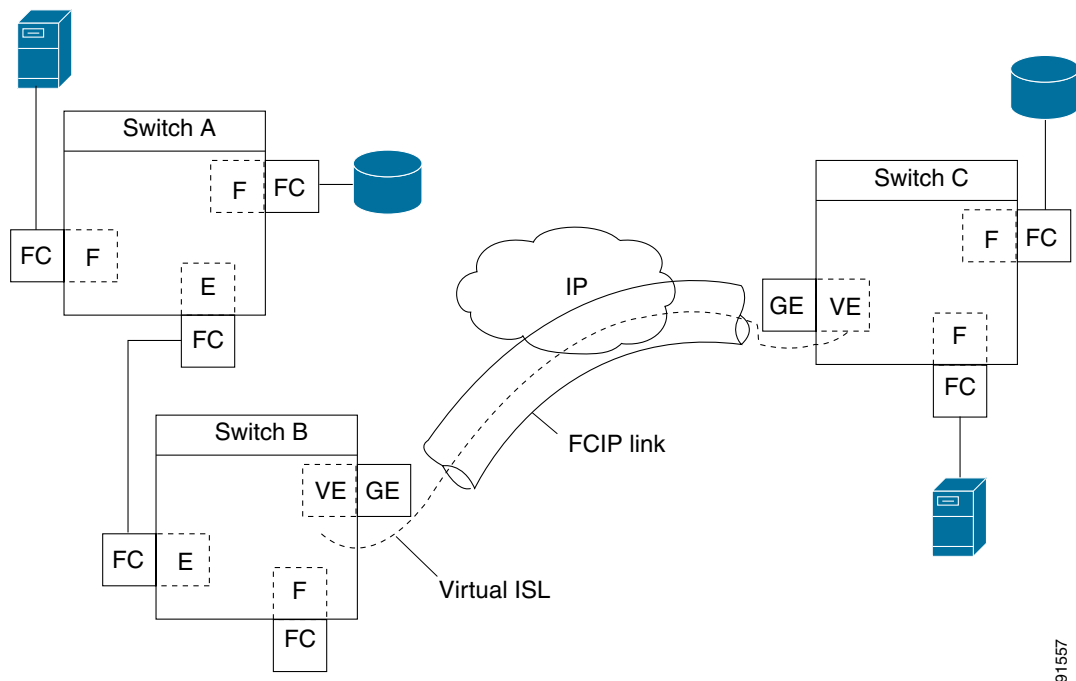
FCIP and VE Ports

Figure 17-7 provides the internal model of FCIP with respect to Fibre Channel inter switch links (ISLs) and Cisco's enhanced ISLs (EISLs). See the “E Port” section on page 9-3.

FCIP defines virtual E (VE) ports, which 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 a 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 17-7).

Figure 17-7 FCIP Links and Virtual ISLs



91557

FCIP Link

FCIP links consist of one or more TCP connections between two FCIP link end points. 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 second connection is used only for Fibre Channel control frames, i.e. switch-to-switch protocol frames (all Class F) frames. This arrangement is used to provide low latency for all control frames.

To enable FCIP on the IPS module, a FCIP profile and FCIP interface (interface FCIP) must be configured.

Send documentation comments to mdsfeedback-doc@cisco.com

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 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 terminates (see [Figure 17-8](#)).

Figure 17-8 FCIP Profile and FCIP Links

FCIP Interface

The FCIP interface is the local end point 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 terminates 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 mdsfeedback-doc@cisco.com

Basic FCIP Configuration


To configure a FCIP link, perform this procedure on both switches.

-
- Step 1** Configure the Gigabit Ethernet interface.
 - Step 2** Create a FCIP profile, assign the Gigabit Ethernet interface's IP address to the profile. See the [“Creating FCIP Profiles” section on page 17-20](#).
 - Step 3** Create a FCIP interface, assign the profile to the interface. See the [“Creating FCIP Links” section on page 17-21](#).
 - Step 4** Configure the peer IP address for the FCIP interface. See the [“Creating FCIP Links” section on page 17-21](#).
 - Step 5** Enable the interface. See the [“Creating FCIP Links” section on page 17-21](#).
-

Creating FCIP Profiles

To create a FCIP profile, you must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile (see [Figure 17-9](#)).

Figure 17-9 Assigning Profiles to Each Gigabit Ethernet Interface



To create a FCIP profile in switch 1, 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 IP address of the Gigabit Ethernet interface (3/1).

To assign FCIP profile in switch 2, 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 IP address of the Gigabit Ethernet interface.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Creating FCIP Links

When two FCIP link end points are created, a FCIP link is established between the two IPS modules. To create a FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) a FCIP link to that peer switch (see Figure 17-10).

Figure 17-10 Assigning Profiles to Each Gigabit Ethernet Interface



To create a FCIP link end point 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 a 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 IP 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 a FCIP link end point in switch 2, follow these steps:

	Command	Purpose
Step 1	switch2# config terminal switch2(config)#	Enters configuration mode.
Step 2	switch2(config)# interface fcip 52 switch2(config-if)#	Creates a 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 IP address information (10.100.1.25 for switch 1) to the FCIP interface
Step 5	switch1(config-if)# no shutdown	Enables the interface.

Send documentation comments to mdsfeedback-doc@cisco.com

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.

- [Configuring TCP Listener Ports, page 17-22](#)
- [Configuring TCP Parameters, page 17-22](#)

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

To enter the `switch(config-profile)#` prompt, follow these steps:

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

Configuring TCP Listener Ports

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	<code>switch(config-profile)# port 5000</code>	Associates the profile with the local port number (5000).
	<code>switch(config-profile)# no port</code>	Reverts to the default 3225 port.

Configuring TCP Parameters

This section provides details on the TCP parameters that can be configured to control TCP behavior in a switch. The following TCP parameters can be configured.

- [Minimum Retransmit Timeout, page 17-23](#)
- [Keepalive Timeout, page 17-23](#)
- [Maximum Retransmissions, page 17-23](#)
- [Path MTU, page 17-24](#)
- [SACK, page 17-24](#)
- [Window Management, page 17-24](#)
- [Buffer Size, page 17-25](#)
- [Quality of Service, page 17-25](#)
- [Monitoring Window Congestion, page 17-26](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Minimum Retransmit Timeout

The **tcp minimum-retransmit-time** option controls the minimum amount of time TCP waits before retransmitting. By default, this value is 300 milliseconds.

To configure the minimum retransmit time, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp min-retransmit-time 500</code>	Specifies the minimum TCP retransmit time for the TCP connection in milliseconds (500). The default is 300 milliseconds and the range is from 250 to 5000 milliseconds.
	<code>switch(config-profile)# no tcp min-retransmit-time 500</code>	Reverts the minimum TCP retransmit time to the factory default of 300 milliseconds.

Keepalive Timeout

The **tcp keepalive-timeout** option enables you to configure the interval between which the TCP connection verifies if the FCIP link is functioning. This ensures that a FCIP link failure is detected quickly even when there is no traffic.

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

The first interval during which the connection is idle is 60 seconds (default). When the connection is idle for 60 seconds, 8 keepalive probes are sent at 1-second intervals. If no response is received for these 8 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 from the default of 60 seconds. This interval is identified using the **keepalive-timeout** option. The valid range is from 1 to 7200 seconds.

To configure the keep alive timeout, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp keepalive-timeout 120</code>	Specifies the keepalive timeout interval for the TCP connection in seconds (120). The default is 60 seconds. The range is from 1 to 7200 seconds.
	<code>switch(config-profile)# no tcp keepalive-timeout 120</code>	Reverts the keepalive-timeout to 60 seconds.

Maximum Retransmissions

The **tcp max-retransmissions** option specifies 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	<code>switch(config-profile)# tcp max-retransmissions 6</code>	Specifies the maximum number of retransmissions (6). The default is 4 and the range is from 1 to 8 retransmissions.
	<code>switch(config-profile)# no tcp max-retransmissions 6</code>	Reverts to the default of 4 retransmissions.

Send documentation comments to mdsfeedback-doc@cisco.com

Path MTU

Path MTU (PMTU) is the minimum MTU on the IP network between the two end points 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 default timeout of 3600 seconds. If TCP reduces the size of the max 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	<code>switch(config-profile)# no tcp pmtu-enable</code>	Disables PMTU discovery.
	<code>switch(config-profile)# tcp pmtu-enable</code>	Enables (default) PMTU discovery with the default value of 3600 seconds.
	<code>switch(config-profile)# tcp pmtu-enable reset-timeout 90</code>	Specifies the PMTU reset timeout to 90 seconds. The default is 3600 seconds and the range is from 60 to 3600 seconds.
	<code>switch(config-profile)# no tcp pmtu-enable reset-timeout 600</code>	Leaves the PMTU in an enabled state but changes the timeout to the default of 3600 seconds.

SACK

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 time. 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	<code>switch(config-profile)# no tcp sack-enable</code>	Disables SACK.
	<code>switch(config-profile)# tcp sack-enable</code>	Enables SACK (default).

Window Management

The optimal TCP window size is computed using three options.

- The **maximum-bandwidth** option configures the maximum available end-to-end bandwidth in the path (900 Mbps in the configuration example).
- The **minimum-available-bandwidth** option configures the minimum slow start threshold.
- The **round-trip-time** option is the estimated round trip time across the IP network to reach the FCIP peer end point (10 milliseconds in the configuration example). If the round-trip-time value is under-estimated, the TCP window size will be too small to reach the maximum available bandwidth. If the round-trip-time is overestimated, the TCP window size will be too big. If the maximum available bandwidth is correct, this will cause increase in latency and potential packet drop in the network but will not affect the speed.

The **maximum-bandwidth** option and the **round-trip-time** option together determine the window size.

Send documentation comments to mdsfeedback-doc@cisco.com

The **minimum-available-bandwidth** option and the **round-trip-time** option together determine the threshold below which TCP aggressively increases its size. After it reaches the threshold the software uses standard TCP rules to reach the maximum available bandwidth. The defaults are max-bandwidth = 1G, min-available-bandwidth = 2 Mbps, and round-trip-time is 10ms

To configure window management, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code>	Configures the maximum available bandwidth at 900 Mbps, the minimum slow start threshold as 300 Mbps, and the round trip time as 10 milliseconds.
	<code>switch(config-profile)# no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code>	Reverts to the factory defaults. The defaults are max-bandwidth = 1G, min-available-bandwidth = 2 Mbps and round-trip-time is 10ms.
	<code>switch(config-profile)# tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200</code>	Configures the maximum available bandwidth at 2000 Kbps, the minimum slow start threshold as 2000 Kbps, and the round trip time as 200 microseconds.

Buffer Size

The **send-buffer-size** option defines 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 buffer size is 0 KB.

To set the buffer size, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp send-buffer-size 5000</code>	Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 8192 KB.
	<code>switch(config-profile)# no tcp send-buffer-size 5000</code>	Reverts the switch to its factory default (0 KB).

Quality of Service

The **qos control** option 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 control values, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp qos control 3 data 5</code>	Configures the control TCP connection and data connection to mark all packets on that DSCP value.
	<code>switch(config-profile)# no tcp qos control 3 data 5</code>	Reverts the switch to its factory default (no packets).

Send documentation comments to mdsfeedback-doc@cisco.com

Monitoring Window Congestion

By configuring the congestion window monitoring (CWM) option, you can influence the rate at which TCP ramps up the transmitted bandwidth after an idle period as listed below:

- If the traffic is transmitted in burst sizes that are smaller than the configured CWM value, you can send the whole traffic burst immediately, provided no drops occurred.
- If the traffic burst is larger than the configured CWM value, some traffic will not be sent immediately.
- If the end-to-end path between the two Cisco MDS 9000 Family switches is 1G, you can set the maximum burst size.
- If the router connecting to the IPS port does not have sufficient buffering, you can use the smallest available value to decrease the burst size.

By default the **tcp cwm** option is enabled and the default burst size is 10KB.



Tip

We recommend that this feature remain enabled to realize optimal performance.

To change the CWM defaults, follow these steps:

Step 1	Command	Purpose
	switch(config-profile)# no tcp cwm	Disables congestion monitoring.
	switch(config-profile)# tcp cwm	Enables congestion monitoring and sets the defaults burst size at 10 KB.
	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 cp cwm burstsize 25	Leaves the CWM feature in an enabled state but changes the burst size to the default of 10 KB.

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface. To do so, you must first create the interface and enter the `config-if` submode.

- [Configuring Peers, page 17-27](#)
- [Configuring Active Connection, page 17-28](#)
- [Configuring the Number of TCP Connections, page 17-29](#)
- [Enabling Time Stamps, page 17-29](#)
- [B Port Interoperability Mode, page 17-30](#)

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 a FCIP interface (100).

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Peers

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

- [Peer IP Address, page 17-27](#)—used to configure both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- [Special Frames, page 17-27](#)—used to configure one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the port 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.

To assign the peer information based on the IP address, port number, or a profile ID, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# peer-info ipaddr 10.1.1.1</code>	Assigns an IP address to configure the peer information. Since no port is specified, the default port number, 3225, is used.
	<code>switch(config-if)# no peer-info ipaddr 10.10.1.1</code>	Deletes the assigned peer port information.
Step 2	<code>switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000</code>	Assigns the IP address and sets the peer TCP port to 3000. The valid port number range is from 0 to 65535.
	<code>switch(config-if)# no peer-info ipaddr 10.1.1.1 port 2000#</code>	Deletes the assigned peer port information.
Step 3	<code>switch(config-if)# peer-info profile_id 20</code>	Assigns the peer profile ID to connect to 20. The valid range is from 1 to 255.
	<code>switch(config-if)# no peer-info profile_id 500</code>	Deletes the assigned peer profile information.
Step 4	<code>switch(config-if)# no shutdown</code>	Enables the interface.

Special Frames

You can alternatively establish a FCIP link with a peer using an optional protocol called special frames. You can enable or disable the **special-frame** option. On the peer side, the **special-frame** option must be enabled in order to establish the FCIP link. When the **special-frame** option is 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.



Note

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

Send documentation comments to mdsfeedback-doc@cisco.com

To enable special frames, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# special-frame peer-wnn 12:12:34:45:ab:bc:cd:00</code>	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 .
	<code>switch(config-if)# no special-frame peer-wnn 12:12:34:45:ab:bc:cd:00</code>	Disables special frames (default).
Step 2	<code>switch(config-if)# special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code>	Enables special frames and sets the peer WWN as specified by the profile ID (155).
	<code>switch(config-if)# no special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code>	Disables special frames (default).
Step 3	<code>switch(config-if)# no shutdown</code>	Enables the interface.

Configuring Active Connection

Use the **passive-mode** option to configure the required mode for initiating an IP 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 and merely 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 will not be initiated.

To enable the passive mode, follow these steps:

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

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring the Number of TCP Connections

Use the **tcp-connection** option to specify the number of TCP connections from a FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure 1 or 2 TCP connections.

For example, the Cisco PA-FC-1G Fibre Channel port adapter which has only 1 (one) TCP connection interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit and you can change the configuration on the switch using the **tcp-connection 1** command. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, the software handles it gracefully and moves on 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. Two (2) is the default and the maximum number of TCP connection attempts.
	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.

Enabling Time Stamps

Use the **time-stamp** option to enable or disable FCIP time stamps on a packet. The **time stamp** option instructs the switch to discard packets that are outside the specified time. By default, the **time-stamp** option is disabled.

The **acceptable-diff** option 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 if a packet arrives within a 1000 millisecond interval (+ or -1000 milliseconds), that packet is accepted.



Note

If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the “[NTP Configuration](#)” section on page 3-18).

To enable or disable the **time-stamp** option, follow these steps:

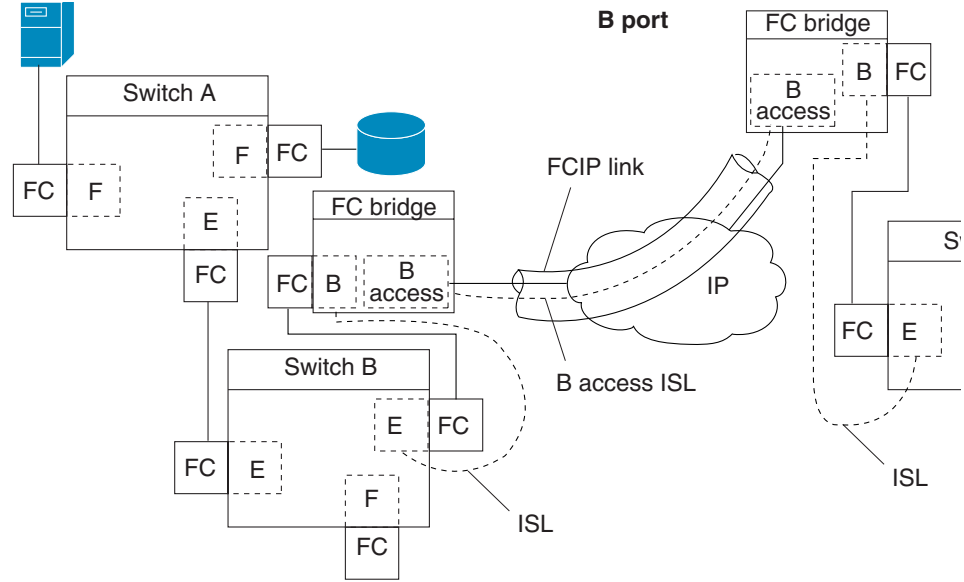
	Command	Purpose
Step 1	switch(config-if)# time-stamp Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link	Enables time stamp checking for received packets with a default acceptable time difference of 1000 milliseconds.
	switch(config-if)# no time-stamp	Disables (default) time stamps.
Step 2	switch(config-if)# time-stamp acceptable-diff 4000	Configures the acceptable time within which a packet is accepted. The default difference is a 1000 millisecond interval from the network time. The valid range is from 1 to 60,000 milliseconds.
	switch(config-if)# no time-stamp acceptable-diff 500	Deletes the configured time difference and reverts the difference to factory defaults.
Step 3	switch(config-if)# no shutdown	Enables the interface.

Send documentation comments to mdsfeedback-doc@cisco.com

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 17-11](#) depicts a typical SAN extension over an IP network.

Figure 17-11 FCIP B Port and Fibre Channel E Port



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

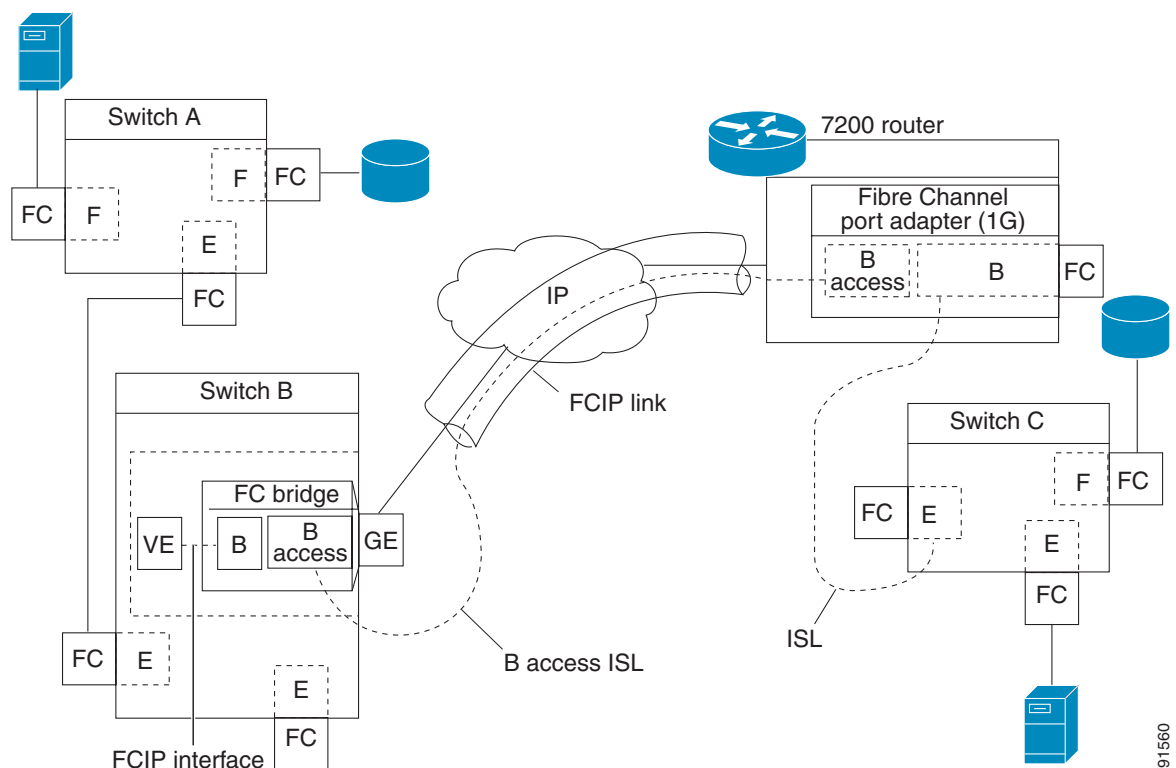
Send documentation comments to mdsfeedback-doc@cisco.com

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 which 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 a FCIP link, B ports use a B access ISL*.

The IPS module supports 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 which completes the end-to-end E port connectivity requirement (see Figure 17-12).

Figure 17-12 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module allows remote B port SAN extenders to communicate directly with an Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring B Ports

When a 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	<code>switch(config-if)# bport</code>	Enables B port mode on the FCIP interface.
	<code>switch(config-if)# no bport</code>	Reverts to E port mode on the FCIP interface (default).
Step 2	<code>switch(config-if)# bport-keepalive</code>	Enables the reception of keepalive responses sent by a remote peer.
Step 3	<code>switch(config-if)# no bport-keepalive</code>	Disables the reception of keepalive responses sent by a remote peer (default).

E Port Configurations

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available FCIP interfaces:

- VSANs (see [Chapter 8, “Configuring and Managing VSANs”](#))
 - FCIP interfaces can be a member of any VSAN.
- Trunk mode (see [Chapter 10, “Configuring Trunking”](#))
 - Trunk mode can be configured.
 - Trunk allowed VSANs can be configured
- PortChannels (see [Chapter 11, “Configuring PortChannels”](#))
 - 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 15, “Configuring Fibre Channel Routing Services and Protocols”](#))
- Fibre Channel domains (fcdomains—see [Chapter 19, “Configuring Domain Parameters”](#))
- Zone merge (see [Chapter 12, “Configuring and Managing Zones”](#))
 - Importing the zone database from the adjacent switch.
 - Exporting the zone database from the adjacent switch.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Displaying FCIP 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 Examples 17-14 to 17-19.

Example 17-14 Displays the FCIP Interface

```
switch# show interface fcip 3
fcip3 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:ca:00:05:30:00:07:1e
  Peer port WWN is 20:ca:00:00:53:00:18:1e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1,10)
  Trunk vsans (operational) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (10)
  Trunk vsans (initializing) ()
  Using Profile id 3 (interface GigabitEthernet4/3)
  Peer Information
    Peer Internet address is 43.1.1.1 and port is 3225
    Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  B-port mode disabled
  TCP Connection Information
    2 Active TCP connections
      Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
      Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
    30 Attempts for active connections, 0 close of connections
  TCP Parameters
    Path MTU 1500 bytes
    Current retransmission timeout is 300 ms
    Round trip time: Smoothed 10 ms, Variance: 5
    Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
    Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
    Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
  5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
  5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
    808 frames input, 75268 bytes
      808 Class F frames input, 75268 bytes
      0 Class 2/3 frames input, 0 bytes
      0 Error frames timestamp error 0
    806 frames output, 74712 bytes
      806 Class F frames output, 74712 bytes
      0 Class 2/3 frames output, 0 bytes
      0 Error frames 0 reass frames
```

Example 17-15 Displays Detailed FCIP Interface Counter Information

```
switch# show interface fcip 3 counters
fcip3
  TCP Connection Information
    2 Active TCP connections
      Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
      Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
    30 Attempts for active connections, 0 close of connections
  TCP Parameters
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

Path MTU 1500 bytes
Current retransmission timeout is 300 ms
Round trip time: Smoothed 10 ms, Variance: 5
Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
814 frames input, 75820 bytes
  814 Class F frames input, 75820 bytes
  0 Class 2/3 frames input, 0 bytes
  0 Error frames timestamp error 0
812 frames output, 75264 bytes
  812 Class F frames output, 75264 bytes
  0 Class 2/3 frames output, 0 bytes
  0 Error frames 0 reass frames

```

Example 17-16 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 Mbits/s	Total Frames	Rate Mbits/s	Total Frames
fcip3	9	0	9	0

Example 17-17 Displays the FCIP Interface Description

```
switch# show interface fcip 51 description
```

```

FCIP51
  Sample FCIP interface

```

Example 17-18 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 17-19 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

```

Send documentation comments to mdsfeedback-doc@cisco.com

Minimum available bandwidth is 15000 kbps
Estimated round trip time is 1000 usec

FCIP High Availability

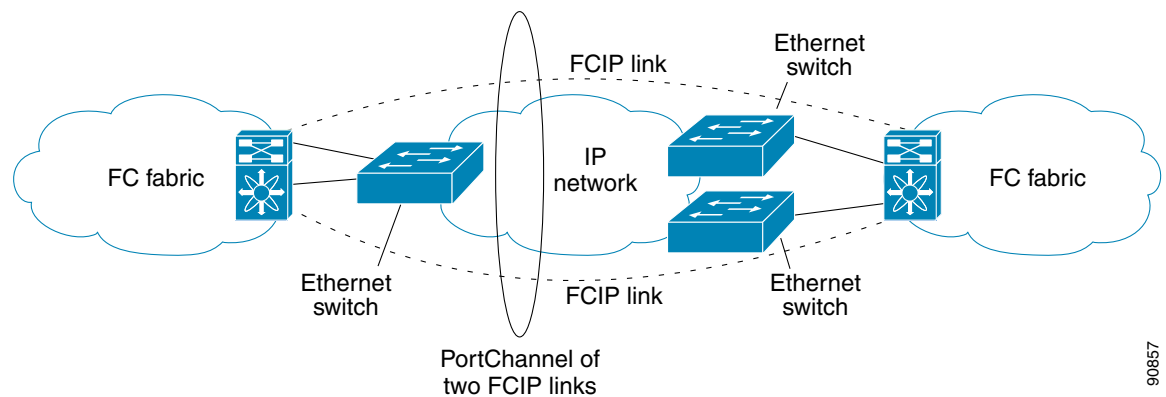
The following high availability solutions are available for FCIP configurations:

- [Fibre Channel PortChannels, page 17-35](#)
- [FSPF, page 17-36](#)
- [VRRP, page 17-36](#)
- [Ethernet PortChannels, page 17-37](#)

Fibre Channel PortChannels

[Figure 17-13](#) 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 17-13 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.


90857

Send documentation comments to mdsfeedback-doc@cisco.com

FSPF

Figure 17-14 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 17-14 FSPF-Based Load Balancing



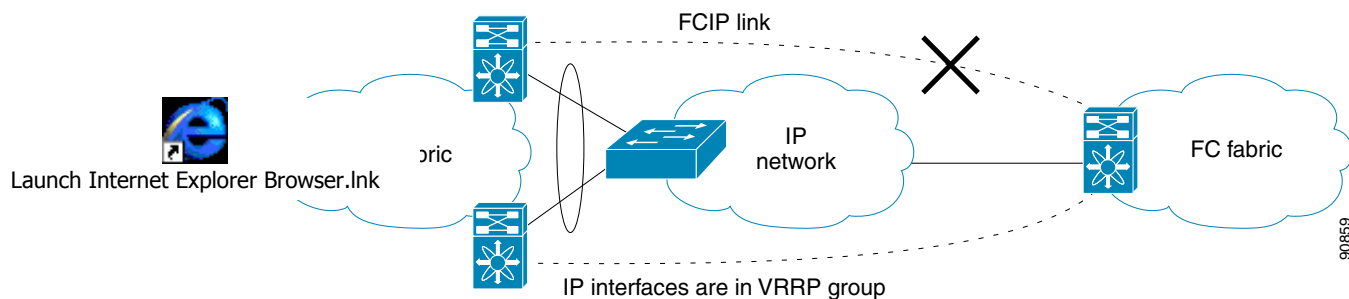
The following characteristics set FSPF solutions apart from other solutions:

- 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 17-15 displays a 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 17-15 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Ethernet PortChannels

Figure 17-16 displays a Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 17-16 Ethernet PortChannel-Based High Availability



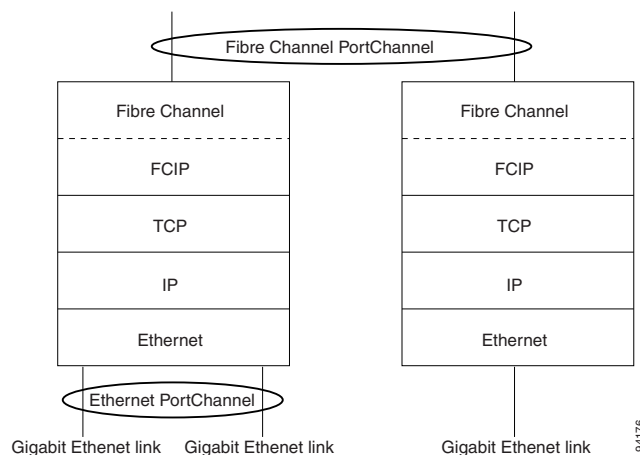
The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- 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 appears like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer Ethernet-level redundancy, Fibre Channel PortChannels offer (E)ISL-level redundancy. FCIP is unaware of any Ethernet PortChannels or Fibre Channel PortChannels. Fibre Channel PortChannels are unaware of any Ethernet PortChannels, and there is no mapping between the two (see [PortChannels at the Fibre Channel and Ethernet Levels](#), page 17-37).

Figure 17-17 PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see [Chapter 11, “Configuring PortChannels.”](#) To configure Ethernet PortChannels, refer to the [“Configuring Ethernet PortChannels”](#) section on page 17-14.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring iSCSI


This section includes the following topics:

- [About iSCSI, page 17-38](#)
- [Presenting Fibre Channel Targets as iSCSI Targets, page 17-41](#)
- [iSCSI Virtual Target Configuration Examples, page 17-44](#)
- [Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 17-46](#)
- [Access Control in iSCSI, page 17-49](#)
- [User Authentication Using iSCSI, page 17-51](#)
- [Assigning VSAN Membership to iSCSI Hosts, page 17-48](#)
- [Displaying iSCSI Information, page 17-53](#)
- [iSCSI High Availability, page 17-62](#)
- [iSCSI Authentication Setup Guidelines, page 17-64](#)

About iSCSI

The IPS module provides transparent SCSI routing. IP hosts using iSCSI protocol can transparently access iSCSI targets on the Fibre Channel network. [Figure 17-18](#) provides an example of a typical configuration of iSCSI hosts with access to a Fibre Channel SAN.


Figure 17-18 Typical IP to Fibre Channel SAN Configuration



Send documentation comments to mdsfeedback-doc@cisco.com


The IPS module enables you to create virtual iSCSI targets and maps them to physical Fibre Channel targets available in the Fibre Channel SAN. It presents the Fibre Channel targets to IP hosts as if the physical targets were attached to the IP network (see Figure 17-19).

Figure 17-19 iSCSI View



In conjunction with presenting Fibre Channel targets to iSCSI hosts, the iSCSI feature presents each iSCSI host as a Fibre Channel host, i.e. Host Bus Adaptor (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network (see Figure 17-20).

Figure 17-20 Fibre Channel SAN View



Note

Refer to the IETF standards for IP storage at <http://www.ietf.org>, for information on the iSCSI protocol.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Routing iSCSI Requests and Responses

The iSCSI feature consists of routing iSCSI requests and responses between hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see [Figure 17-21](#)).

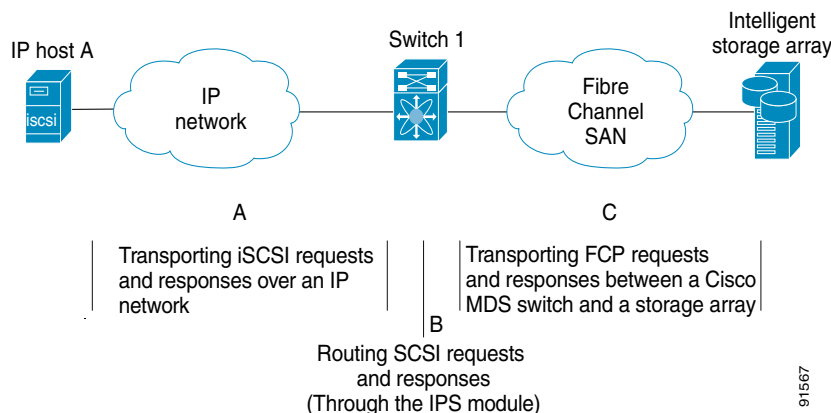
Figure 17-21 Routing iSCSI Requests and Responses for Transparent iSCSI Routing

Each iSCSI host that requires access to storage via the IPS module needs to have a compatible iSCSI driver installed. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the perspective of a host operating system, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver for a peripheral channel in the host. From the perspective of the storage device, each IP host appears as a Fibre Channel host.

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions (see [Figure 17-21](#)):

- Transporting iSCSI requests and responses over an IP network between hosts and the IPS module.
- Routing SCSI requests and responses between hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). This routing is performed by the IPS module.
- Transporting FCP requests or responses between the IPS module and Fibre Channel storage devices.

Figure 17-22 Transparent SCSI Routing Actions



Note

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN.

Send documentation comments to mdsfeedback-doc@cisco.com

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module presents physical Fibre Channel targets as iSCSI targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- **Dynamic Importing**—used if all logical units (LUs) in all Fibre Channel storage targets are made available to iSCSI hosts (subject to VSAN and zoning).
- **Static Importing**—used if iSCSI hosts are restricted to subsets of LUs in the Fibre Channel targets and additional iSCSI access control is needed (see the “[Access Control in iSCSI](#)” section on page 17-49). Also, static import allows automatic failover if the Fibre Channel targets’ LU is reached by redundant Fibre Channel ports (see the “[High Availability Static Importing](#)” section on page 17-43).



Note

The IPS module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have a configured name. Un mapped targets are advertised with the name created by the conventions explained in this section.

Dynamic Importing

To enable dynamic importing of Fibre Channel targets into iSCSI, use the **iscsi import target fc** command.

The IPS module maps each physical Fibre Channel target port as one iSCSI target. That is, all LU accessible via the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the storage target.

For example, if an iSCSI target was created for Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0 through 2, those LUNs would become available to an IP host as LUNs 0 through 2 as well.

The iSCSI target node name is created automatically using the iSCSI I qualified name (IQN) format. The IPS module creates an IQN formatted iSCSI node name using the following conventions:

- IPS ports that are not part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
```

- IPS ports that are part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
```

- Ports that are part of a PortChannel use this format:

```
iqn.1987-05.com.cisco:05.PC-<port-ch-intf#>-<port-ch-sub-intf#>.<Target-pWWN>
```



Note

In this format, each IPS port in a Cisco MDS 9000 Family switch creates a different iSCSI target node name for the same Fibre Channel target.

Send documentation comments to mdsfeedback-doc@cisco.com

To dynamically import Fibre Channel targets, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi import target fc	IPS modules dynamically map each Fibre Channel target in the Fibre Channel SAN to the IP network. The automatically-created iSCSI target node names use the IQN format. Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms.

Static Importing

You can manually (statically) create an iSCSI target and assign a node name to it. A statically-mapped iSCSI target can either contain the whole FC target port, or it can contain one or more LUs from a Fibre Channel target port.

To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.abc switch(config-(iscsi-tgt))#	Creates the iSCSI target name iqn.abc.
Step 3	switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06	Maps a virtual target node to a Fibre Channel target. One iSCSI target cannot contain more than one Fibre Channel target. Don't specify the LUN if you wish to map the whole Fibre Channel target to an iSCSI target. All Fibre Channel target LUNs are exposed to iSCSI. Use the LUN option to map different Fibre Channel LUNs to different iSCSI virtual targets. If you have already mapped the whole Fibre Channel target, you will not be able to use this option.
	switch(config-(iscsi-tgt))# pwwn 26:00:00:00:00:11:00:11 fc-lun 1 iscsi-lun 1	Maps the whole target using LUN mapping options.

Creating iSCSI Targets

To create a static iSCSI target for the entire Fibre Channel target port, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.abc switch(config-(iscsi-tgt))#	Creates the iSCSI target name iqn.abc.

Send documentation comments to mdsfeedback-doc@cisco.com

Advertising iSCSI Targets

You can limit the Gigabit Ethernet interfaces over which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-(iscsi-tgt))# advertise interface GigabitEthernet 2/5</code>	Advertises the virtual target only on the specified interface. By default, it is advertised on all interfaces in all IPS modules.
	<code>switch(config-(iscsi-tgt))# no advertise interface GigabitEthernet 2/5</code>	Removes this interfaces from the list of interfaces from which this target is advertised.

High Availability Static Importing

Physical Fibre Channel targets are configured to have LUs visible over two Fibre Channel N ports—one in active mode and another in passive mode. When the active port fails, the passive port takes over. Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the passive port becomes active and the iSCSI session switches to use the new active port. If both the primary and secondary pWWNs are available, then both pWWNs can be used—each session may use either pWWN (see [Figure 17-23](#)).

Figure 17-23 Mapping LUNs to be Visible

In [Figure 17-23](#), you can create a virtual iSCSI target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

To create a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	<code>switch# config terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# iscsi virtual-target name abc</code>	Creates the iSCSI target name iqn.abc.
Step 3	<code>switch(config-(iscsi-tgt))# pwn 26:00:01:02:03:04:05:06</code> <code>secondary-pwn 26:00:01:02:03:10:11:12</code>	Configures the secondary port for this virtual target.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

iSCSI Virtual Target Configuration Examples


This section provides three examples of virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a virtual iSCSI target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 17-24](#)).

```
iscsi virtual-target name iqn.abc-1
pWWN 28:00:01:02:03:04:05:06
```

Figure 17-24 Assigning iSCSI Node Names




91569

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see [Figure 17-25](#)).

```
iscsi virtual-target name iqn.abc-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 1
iscsi virtual-target name iqn.abc-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 1
iscsi virtual-target name iqn.abc-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 1
```

Figure 17-25 Mapping LUNs to iSCSI a Node Name



91570


Send documentation comments to mdsfeedback-doc@cisco.com

Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see [Figure 17-26](#)).

```
iscsi virtual-target name iqn.abc-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.abc-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.abc-3
  pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
  pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
```

Figure 17-26 Mapping LUNs to Multiple iSCSI Node Names



Send documentation comments to mdsfeedback-doc@cisco.com

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The iSCSI hosts are mapped to virtual Fibre Channel hosts in one of two ways (see [Figure 17-20](#)):

- **Dynamic Mapping** (default)—used if no access control is done on the Fibre Channel target. An iSCSI host may use different pWWNs each time it connects to a Fibre Channel target.
- **Static Mapping**—used if an iSCSI host should always have the same pWWN or nWWN each time it connects to a Fibre Channel target.

Dynamic Mapping

When an iSCSI host connects to the IPS module using the iSCSI protocol, a virtual N port is created for the host. The nWWNs and pWWNs are dynamically allocated from the switch's Fibre Channel WWN pool. The IPS module registers this N port in the Fibre Channel SAN. The IPS module continues using that nWWN and pWWN to represent this iSCSI host until it no longer has a connection to any iSCSI target via that IP storage port.

At that point, the virtual Fibre Channel host is taken offline from the Fibre Channel SAN and the nWWNs and pWWNs are released back to the switch's Fibre Channel WWN pool. These addresses become available for assignment to other iSCSI hosts requiring access to Fibre Channel SANs.

When a dynamically mapped iSCSI initiator has multiple sessions to multiple Fibre Channel targets, each session can use the same pWWN and nWWN as long as it uses the same node name in the iSCSI login message. If the host has multiple network interfaces (and the same IP address), and each IP address is treated as different hosts, then the **switchport initiator id ip-address** command is used to identify an iSCSI initiator. This command uses the IP address instead of the initiator name.

All dynamic iSCSI initiators are members of the default VSAN (VSAN 1).

Identifying Initiators

An iSCSI initiator is identified in one of two ways:

- The iSCSI node name (**switchport initiator id name** command)—If the node name is used, an initiator with multiple IP addresses (multiple interface cards—NICs or multiple network interfaces) has one virtual N port.
- The IP address (**switchport initiator id ip-address** command)—If the IP address is used, a virtual N port is created for each NIC or network interface.

By default, the switch uses the iSCSI node name to identify the initiator. You can change this default so the switch identifies the initiator using the IP address.

To identify the initiator using the IP address, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	switch(config-if)# switchport initiator id ip-address	Identifies the iSCSI initiator based on the IP address.
	switch(config-if)# switchport initiator id name	Identifies the iSCSI initiator based on the initiator node name.

Send documentation comments to mdsfeedback-doc@cisco.com

Static Mapping

With dynamic mapping, each time the iSCSI host connects to the IPS module a new Fibre Channel N port is created and the nWWNs and pWWNs allocated for this N port may be different. Use the static mapping method if you need to obtain the same nWWNs and pWWNs for the iSCSI host each time it connects to the IPS module.

You can implement static mapping in one of two ways: system assignment or manual assignment.

- System assignment—When a static mapping configuration is created, one nWWN and/or one or more pWWNs are allocated from the switch's Fibre Channel WWN pool and the mapping is kept permanent. This assignment uses the **system-assign** option.
- Manual assignment—You can specify your own unique WWN using the **manual-assign** option. Each time the iSCSI session is created, the same nWWN/pWWN that was initially created is used.



Tip We recommend using the **system-assign** option. If you manually assign a WWN, you must uniquely associate the WWN to a single device (see the [“Configuring World Wide Names” section on page 24-16](#)).

Static mapping can be used on the IPS module to access intelligent Fibre Channel storage arrays that have access control and LUN mapping/masking configuration based on the initiator's pWWNs and/or nWWNs.



Note If an iSCSI host connects to multiple IPS ports, each port independently creates one virtual N port for the host. If static mapping is used, enough pWWNs should be configured for as many IPS ports to which a host connects.

To configure static mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)#	Configures an iSCSI initiator using the iSCSI name of the initiator node. The maximum name length is restricted to 255 alphanumeric characters. The minimum length is 16.
	switch(config)# no iscsi initiator name iqn.1987-02.com.cisco.initiator	Deletes the configured iSCSI initiator.

To configure static mapping (using the **ip-address** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator ip address 10.50.0.0 switch(config-iscsi-init)#	Configures an iSCSI initiator. using the IP address of the initiator node.
	switch(config)# no iscsi initiator ip address 10.50.0.0	Deletes the configured iSCSI initiator.

Send documentation comments to mdsfeedback-doc@cisco.com

To assign the WWN for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-(iscsi-init))# static nwwn system-assign</code>	Uses the switch's WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent.
	<code>switch(config-(iscsi-init))# nwwn 20:00:00:05:30:00:59:11</code>	Assigns the user provided WWN as nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node.
Step 2	<code>switch(config-(iscsi-init))# static pwwn system-assign 2</code>	Uses the switch's WWN pool to allocate two pWWNs for this iSCSI initiator and keeps it persistent. The range is from 1 to 64.
	<code>switch(config-(iscsi-init))# pwwn 21:00:00:20:37:73:3b:20</code>	Assigns the user provided WWN as pWWN for the iSCSI initiator.

Assigning VSAN Membership to iSCSI Hosts

An iSCSI host can reside in multiple VSANs based on the configuration. By default, a host is only in VSAN 1 (default VSAN). The IPS module creates one Fibre Channel virtual N port in each VSAN to which the host belongs.

An iSCSI host can become a member of one or more VSANs.

To assign VSAN membership for iSCSI hosts, follow these steps:

	Command	Purpose
Step 1	<code>switch# config terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# iscsi initiator name ign.1987-02.com.cisco.initiator</code> <code>switch(config-(iscsi-init))#</code>	Configures an iSCSI initiator.
Step 3	<code>switch(config-(iscsi-init))# vsan 3</code>	Assigns the iSCSI initiator node to a specified VSAN. Note You can assign this host to one or more VSANs.
	<code>switch(config-(iscsi-init))# no vsan 5</code>	Removes the iSCSI node from the specified VSAN.



Note

By default, an iSCSI initiator is only present in the default VSAN (VSAN 1). When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Access Control in iSCSI

You can control access to each statically-mapped iSCSI target by specifying a list of IPS ports on which it will be advertised and specifying a list of iSCSI initiator node names allowed to access it. Fibre Channel zoning-based access control and iSCSI-based access control are the two mechanisms by which access control can be provided for iSCSI. Both methods can be used simultaneously.



Note

This access control is in addition to the existing Fibre Channel access control. The iSCSI initiator has to be in the same VSAN and zone as the physical Fibre Channel target.

Fibre Channel Zoning-Based Access Control

Zoning is an access control mechanism within a VSAN. The switch's zoning implementation extends the VSAN and zoning concepts from the Fibre Channel domain to also cover the iSCSI domain. This extension includes both iSCSI and Fibre Channel features and provides a uniform, flexible access control across a SAN. Static and dynamic are the two Fibre Channel zoning access control mechanisms.

- **Static**—statically map the iSCSI host to Fibre Channel virtual N port(s). This creates a permanent nWWNs and pWWNs. Next, configure the assigned pWWN into zones, similar to adding a regular Fibre Channel host's pWWN to a zone.
- **Dynamic**—add the iSCSI host's initiator node name as a member of a zone. When the IP host's Fibre Channel virtual N port is created and the Fibre Channel address (nWWNs and pWWNs) is assigned, Fibre Channel zoning is enforced.

To register an iSCSI initiator in the zone database, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# zone name iSCSIzone vsan 1 switch(config-zone)	Creates a zone name for the iSCSI devices in the IPS module to be included.
Step 3	switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1	Adds the device as specified by the node name.
	switch(config-zone)# no member iqn.1987-02.com.cisco.initiator1	Deletes the specified device.
	switch(config-zone)# member symbolic-nodename 10.50.1.1	Adds the device as specified by the IP address.
	switch(config-zone)# no member 10.50.1.1	Deletes the specified device.

iSCSI-Based Access Control

For static iSCSI targets, you can manually configure a list of iSCSI initiators that are allowed to access it. The iSCSI initiator is identified by the iSCSI node name or the IP address of the iSCSI host.

By default, static virtual iSCSI targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a virtual iSCSI target to be accessed by all hosts. The initiator access list can contain one or more initiators. Each initiator is identified by one of the following:

- iSCSI node names
- IP addresses
- IP subnets

Send documentation comments to mdsfeedback-doc@cisco.com

To configure access control in iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.abc switch(config- (iscsi-tgt))#	Creates the iSCSI target name iqn.abc.
Step 3	switch(config- (iscsi-tgt))# pWWN 26:00:01:02:03:04:05:06 switch(config- (iscsi-tgt))#	Maps a virtual target node to a Fibre Channel target.
Step 4	switch(config- (iscsi-tgt))# initiator iqn.1987-02.com.cisco.initiator1 permit	Allows the specified iSCSI initiator node to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	switch(config- (iscsi-tgt))# no initiator iqn.1987-02.com.cisco.initiator1 permit	Prevents the specified initiator node from accessing virtual targets.
	switch(config- (iscsi-tgt))# initiator ip address 10.50.1.1 permit	Allows the specified IP address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	switch(config- (iscsi-tgt))# no initiator ip address 10.50.1.1 permit	Prevents the specified IP address from accessing virtual targets.
	switch(config- (iscsi-tgt))# initiator ip address 10.50.1.1 255.255.255.0 permit	Allows all initiators in this subnetwork to access this virtual target.
	switch(config- (iscsi-tgt))# no initiator ip address 10.50.1.1 255.255.255.0 permit	Prevents all initiators in this subnetwork from accessing virtual targets.
	switch(config- (iscsi-tgt))# all-initiator-permit	Allows all initiator nodes to access this virtual target.
	switch(config- (iscsi-tgt))# no all-initiator-permit	Prevents any initiator from accessing virtual targets (default).

Enforcing Access Control

IPS modules use both iSCSI node name-based and Fibre Channel zoning-based access control lists to enforce access control during iSCSI discovery and iSCSI session creation.

- iSCSI discovery—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section.
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module verifies if the specified iSCSI target (in the session login request) is a static mapped target, and if true, verifies if the IP host's iSCSI node name is allowed to access the target. If the IP host does not have access, its login is rejected.

The IPS module, then creates a Fibre Channel virtual N port (the N port may already exist) for this IP host and does a Fibre Channel name server query for the FCID of the Fibre Channel target pWWN that is being accessed by the IP host. It uses the IP host virtual N port's pWWN as the requester of the name server query. Thus, the name server does a zone-enforced query for the pWWN and responds to the query.

If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

User Authentication Using iSCSI

The IPS module supports the iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. When iSCSI authentication is enabled, the iSCSI hosts must provide user name and password information each time an iSCSI session is established.



Note

Only the Challenge Handshake Authentication Protocol (CHAP) authentication method is supported.

The IPS module also supports iSCSI hosts to challenge the IPS module to authenticate itself.

The **aaa auth iscsi radius** command enables RADIUS authentication for the iSCSI host. If RADIUS authentication is not enabled or RADIUS servers are unavailable, the local database is used (see the “Configuring RADIUS Authentication” section on page 14-14).

To configure RADIUS authentication for an iSCSI user, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# aaa authentication iscsi radius	Uses RADIUS for iSCSI authentication method.
	switch(config)# aaa authentication iscsi local	Configures the switch to only use the local password database for iSCSI CHAP authentication.

Authentication Mechanism

During an iSCSI login, both the iSCSI initiator and target have the option to authenticate each other. By default, the IPS module allows either CHAP authentication or no authentication from iSCSI hosts. If CHAP authentication should always be used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used, issue the **iscsi authentication none** command.



Note

The authentication for a Gigabit Ethernet interface or subinterface configuration overrides the authentication for the global interface configuration.

To configure the authentication mechanism for iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi authentication chap	Configures CHAP as the default authentication mechanism globally for the Cisco MDS switch. CHAP authentication is required for all iSCSI sessions. The validation is done using RADIUS or local authentication.

Send documentation comments to mdsfeedback-doc@cisco.com

To configure the authentication policy for iSCSI sessions to a particular interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface GigabitEthernet 2/1.100 switch(config-if)#	Selects the Gigabit Ethernet interface.
Step 3	switch(config-if)# iscsi authentication none	Specifies that no authentication is required for iSCSI sessions to the selected interface.

The IPS module verifies the iSCSI host authentication in one of two ways: the local password database or RADIUS (see the “[User Authentication](#)” section on page 14-2). If local authentication is used, the **username iscsi-user password iscsi** command assigns a password and a user name for a new user. If the user name does not exist it will be created.



Note

The **iscsi** keyword is mandatory to identify iSCSI users.

To configure iSCSI users for local authentication, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# username iscsiuser password ffsffsfssfs345353554535 iscsi	Configures a user name (iscsiuser) and password (ffsffsfssfs345353554535) in the local database for iSCSI login authentication.
Note	The iscsi keyword is required at the end to identify the user.	

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying iSCSI Information

This section includes the following topics:

- [Displaying iSCSI Interfaces, page 17-53](#)
- [Displaying Global iSCSI Information, page 17-55](#)
- [Displaying iSCSI Sessions, page 17-55](#)
- [Displaying iSCSI Initiators, page 17-56](#)
- [Displaying iSCSI Virtual Targets, page 17-59](#)
- [Displaying IPS Statistics, page 17-60](#)
- [Displaying iSCSI User Information, page 17-61](#)

Displaying iSCSI Interfaces

Use the **show iscsi interface** command to view the summary, counter, description, and status of the iSCSI interface. Use the output to verify the administrative mode, the interface status, TCP parameters currently used, and brief statistics. See [Example 17-20](#).

Example 17-20 Displays the iSCSI Interface Information

```
switch# show interface iscsi 2/1
iscsi2/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:41:00:05:30:00:50:de
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 7, Number of TCP connection: 7
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is disabled
    Keepalive-timeout is 1 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 8
    Sack is disabled
    Minimum available bandwidth is 0 kbps
    Estimated round trip time is 0 usec
  5 minutes input rate 265184 bits/sec, 33148 bytes/sec, 690 frames/sec
  5 minutes output rate 375002168 bits/sec, 46875271 bytes/sec, 33833 frames/sec
  iSCSI statistics
    6202235 packets input, 299732864 bytes
      Command 6189718 pdus, Data-out 1937 pdus, 1983488 bytes, 0 fragments
    146738794 packets output, 196613551108 bytes
      Response 6184282 pdus (with sense 4), R2T 547 pdus
      Data-in 140543388 pdus, 189570075420 bytes
```

The **show iscsi stats** command can be used to view brief or detailed iSCSI statistics per iSCSI interface. See Examples [17-21](#) and [17-22](#).

Example 17-21 Displays iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 4/1
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
iSCSI statistics
  1196 packets input, 173680 bytes
    Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
  output 1802 packets, 647152 bytes
    Response 483 pdus (with sense 0), R2T 25 pdus
    Data-in 685 pdus, 554696 bytes
```

Example 17-22 Displays Detailed iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 4/1 detail
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  iSCSI statistics
    1196 packets input, 173680 bytes
      Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
    output 1802 packets, 647152 bytes
      Response 483 pdus (with sense 0), R2T 25 pdus
      Data-in 685 pdus, 554696 bytes
  iSCSI Forward:
    Command: 483 PDUs (Rcvd: 483)
    Data-Out (Write): 104 PDUs (Rcvd 104), 0 fragments, 106496 bytes
  FCP Forward:
    Xfer_rdy: 25 (Rcvd: 25)
    Data-In: 685 (Rcvd: 719), 554696 bytes
    Response: 483 (Rcvd: 534), with sense 0
    TMF Resp: 0

  iSCSI Stats:
    Login: attempt: 25, succeed: 25, fail: 0, authen fail: 0
    Rcvd: NOP-Out: 556, Sent: NOP-In: 556
      NOP-In: 0, Sent: NOP-Out: 0
      TMF-REQ: 0, Sent: TMF-RESP: 0
      Text-REQ: 6, Sent: Text-RESP: 6
      SNACK: 0
      Unrecognized Opcode: 0, Bad header digest: 0
      Command in window but not next: 0, exceed wait queue limit: 0
      Received PDU in wrong phase: 0
      SCSI Busy responses: 0
  FCP Stats:
    Total: Sent: 726
      Received: 1366 (Error: 0, Unknown: 0)
    Sent: PLOGI: 17, Rcvd: PLOGI_ACC: 17, PLOGI_RJT: 0
      PRLI: 17, Rcvd: PRLI_ACC: 17, PRLI_RJT: 0, Error resp: 0
      LOGO: 12, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
      PRLO: 12, Rcvd: PRLO_ACC: 0, PRLO_RJT: 0
      ABTS: 0, Rcvd: ABTS_ACC: 0
      TMF REQ: 0
      Self orig command: 51, Rcvd: data: 34, resp: 51
    Rcvd: PLOGI: 20, Sent: PLOGI_ACC: 5, PLOGI_RJT: 15
      LOGO: 5, Sent: LOGO_ACC: 5, LOGO_RJT: 0
      PRLI: 5, Sent: PRLI_ACC: 5, PRLI_RJT: 0
      PRLO: 0, Sent: PRLO_ACC: 0, PRLO_RJT: 0
      ABTS: 0

  iSCSI Drop:
    Command: Target down 0, Task in progress 0, LUN map fail 0
      CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
      Persistent Resv 0, No task: 0
    Data-Out: 0, Data CRC Error: 0
    TMF-Req: 0, No task: 0
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
FCP Drop:
  Xfer_rdy: 0, Data-In: 0, Response: 0

Buffer Stats:
  Buffer less than header size: 0, Partial: 53, Split: 79
  Pullup give new buf: 0, Out of contiguous buf: 0, Unaligned m_data: 0
```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See [Example 17-23](#)

Example 17-23 Displays the Current Global iSCSI Configuration and State.

```
switch# show iscsi global
iSCSI Global information
  Authentication: NONE
  Import FC Target: Enabled
  Number of target nodes: 5
  Number of portals: 8
  Number of sessions: 6
  Failed session: 0, Last failed initiator name:
```

Displaying iSCSI Sessions

Use the **show iscsi session** command to view details about the current iSCSI sessions in the switch. Without parameters, this command displays all sessions. The output can be filtered by specificity an initiator, a target, or both.

[Example 17-24](#) displays one iSCSI initiator configured based on the iqn name (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on it's IP address (10.10.100.199).

Example 17-24 Displays Brief Information of All iSCSI Sessions

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  Session #1
    Discovery session, ISID 00023d000043, Status active

  Session #2
    Target VT1
    VSAN 1, ISID 00023d000046, Status active, no reservation

  Session #3
    Target VT2
    VSAN 1, ISID 00023d000048, Status active, no reservation

Initiator 10.10.100.199
  Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT2
    VSAN 1, ISID 246700000000, Status active, no reservation

  Session #2
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
Session #3
Target ign.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
VSAN 1, ISID 246e00000000, Status active, no reservation
```

[Example 17-25](#) and [Example 17-26](#) display the iSCSI initiator configured based on its IP address (10.10.100.199).

Example 17-25 Displays Brief Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1
Initiator 10.10.100.199
Initiator name ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
Session #1
Target VT1
VSAN 1, ISID 246b00000000, Status active, no reservation
```

Example 17-26 Displays Detailed Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1 detail
Initiator 10.10.100.199 (oasis-qa)
Initiator name ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
Session #1 (index 3)
Target VT1
VSAN 1, ISID 246b00000000, TSIH 384, Status active, no reservation
Type Normal, ExpCmdSN 39, MaxCmdSN 54, Barrier 0
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
Stats:
  PDU: Command: 38, Response: 38
  Bytes: TX: 8712, RX: 0
Number of connection: 1
Connection #1
  Local IP address: 10.10.100.200, Peer IP address: 10.10.100.199
  CID 0, State: LOGGED_IN
  StatSN 62, ExpStatSN 0
  MaxRecvDSLength 1024, our_MaxRecvDSLength 1392
  CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
  AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
  Version Min: 2, Max: 2
  FC target: Up, Reorder PDU: No, Marker send: No (int 0)
  Received MaxRecvDSLen key: No
```

Displaying iSCSI Initiators

Use the **show iscsi initiator** command to display information about all initiators connected to a iSCSI interface in the switch. The information can be filtered to display only the desired iSCSI initiator by specifying the initiator name. Detailed output of the iscsi initiator can be obtained by specifying the **detail** option. The **iscsi-session** (and optionally **detail**) parameter displays only iSCSI session information. The **fcip-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See [Examples 17-27](#) and [17-28](#).

Example 17-27 Displays Information About Connected iSCSI Initiators

```
switch# show iscsi initiator
iSCSI Node name is ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
Initiator ip addr (s): 10.10.100.116
iSCSI alias name: AVANTI12-W2K
```


Send documentation comments to mdsfeedback-doc@cisco.com

```

Node WWN is 22:01:00:05:30:00:10:e1 (configured)
Member of vsans: 1, 2, 10
Number of Virtual n_ports: 1
Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
  Interface iSCSI 4/1, Portal group tag: 0x180
  VSAN ID 1, FCID 0x6c0202
  VSAN ID 2, FCID 0x6e0000
  VSAN ID 10, FCID 0x790000

iSCSI Node name is 10.10.100.199
  iSCSI Initiator name: ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  iSCSI alias name: oasis-qa
Node WWN is 22:03:00:05:30:00:10:e1 (configured)
Member of vsans: 1, 5
Number of Virtual n_ports: 1
Virtual Port WWN is 22:00:00:05:30:00:10:e1 (configured)
  Interface iSCSI 4/1, Portal group tag: 0x180
  VSAN ID 5, FCID 0x640000
  VSAN ID 1, FCID 0x6c0203

```

Example 17-28 Display Detailed Information About the iSCSI Initiator

```

switch# show iscsi initiator ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1

  Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag is 0x180
    VSAN ID 1, FCID 0x6c0202
    1 FC sessions, 1 iSCSI sessions
    iSCSI session details
      Target: VT1
      Statistics:
        PDU: Command: 0, Response: 0
        Bytes: TX: 0, RX: 0
        Number of connection: 1
      TCP parameters
        Local 10.10.100.200:3260, Remote 10.10.100.116:4190
        Path MTU: 1500 bytes
        Retransmission timeout: 310 ms
        Round trip time: Smoothed 160 ms, Variance: 38
        Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
        Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
        Congestion window: Current: 1 KB

  FCP Session details
    Target FCID: 0x6c01e8 (S_ID of this session: 0x6c0202)
    pWWN: 21:00:00:20:37:62:c0:0c, nWWN: 20:00:00:20:37:62:c0:0c
    Session state: CLEANUP
    1 iSCSI sessions share this FC session
    Target: VT1
    Negotiated parameters
      RcvDataFieldSize 1392 our_RcvDataFieldSize 1392
      MaxBurstSize 0, EMPD: FALSE
      Random Relative Offset: FALSE, Sequence-in-order: Yes
    Statistics:
      PDU: Command: 0, Response: 0

```

Send documentation comments to mdsfeedback-doc@cisco.com

Use the **show fcns database** (and optionally **detail**) to display the Fibre Channel name server entry for the Fibre Channel N port created for iSCSI initiators in the SAN. See Examples 17-29 and 17-31.

Example 17-29 Displays Fibre Channel Name Server Database

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x6c0001      NL    21:00:00:04:cf:4c:52:c1 (Seagate)          scsi-fcp:target
0x6c01e8      NL    21:00:00:20:37:62:c0:0c (Seagate)          scsi-fcp:target
0x6c0202      N     22:04:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
0x6c0203      N     22:00:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
0x6c0301      NL    21:00:00:20:37:a6:be:32 (Seagate)          scsi-fcp:target
Total number of entries = 5

VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x6e0000      N     22:04:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
Total number of entries = 1

VSAN 5:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x640000      N     22:00:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
Total number of entries = 1
```

Example 17-30 Displays Detailed Information for a Fibre Channel N Port Created for An iSCSI Initiator Identified by it's IQN Name

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
node-wwn               :22:03:00:05:30:00:10:e1
class                  :2,3
node-ip-addr           :10.10.100.199
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :10.10.100.199
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :20:c1:00:05:30:00:10:de
hard-addr              :0x000000

Total number of entries = 1
```

Example 17-31 Displays Detailed Information for a Fibre Channel N Port created for An iSCSI Initiator Identified by it's IP Address

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
node-wwn          :22:03:00:05:30:00:10:e1
class             :2,3
node-ip-addr      :10.10.100.199
ipa               :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :10.10.100.199
port-type         :N
port-ip-addr      :0.0.0.0
fabric-port-wwn   :20:c1:00:05:30:00:10:de
hard-addr         :0x000000
```

Total number of entries = 1

Use the **show iscsi initiator configured** to display information about all the configured iSCSI initiators. Specifying the name shows information about the desired initiator. See [Example 17-32](#).

Example 17-32 Display Information About Configured Initiators

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Member of vsans: 1, 2, 10
  Node WWN is 22:01:00:05:30:00:10:e1
  No. of PWWN: 5
    Port WWN is 22:04:00:05:30:00:10:e1
    Port WWN is 22:05:00:05:30:00:10:e1
    Port WWN is 22:06:00:05:30:00:10:e1
    Port WWN is 22:07:00:05:30:00:10:e1
    Port WWN is 22:08:00:05:30:00:10:e1

iSCSI Node name is 10.10.100.199
  Member of vsans: 1, 5
  Node WWN is 22:03:00:05:30:00:10:e1
  No. of PWWN: 4
    Port WWN is 22:00:00:05:30:00:10:e1
    Port WWN is 22:09:00:05:30:00:10:e1
    Port WWN is 22:0a:00:05:30:00:10:e1
    Port WWN is 22:0b:00:05:30:00:10:e1
```

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the FC targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See [Example 17-33](#).

Example 17-33 Displays Exported Targets

```
switch# show iscsi virtual-target
target: VT1
  * Port WWN 21:00:00:20:37:62:c0:0c
  Configured node
  all initiator permit is enabled

target: VT2
  Port WWN 21:00:00:04:cf:4c:52:c1
  Configured node
  all initiator permit is disabled

target: iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
  Port WWN 21:00:00:20:37:a6:be:32 , VSAN 1
  Auto-created node
```

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Displaying IPS Statistics

The **show ips stats tcp interface** command displays information about the underlying transport for iSCSI. See Examples 17-34 and 17-35.

Example 17-34 Displays iSCSI Stats (brief)

```
switch# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
  Connection Stats
    0 active openings, 6 accepts
    0 failed attempts, 0 reset received, 6 established
  Segment stats
    640780835 received, 150953931 sent, 12 retransmitted
    0 bad segments received, 0 reset sent
  TCP Active Connections

```

Local Address	Remote Address	State	Send-Q	Recv-Q
10.48.69.250:3260	10.48.69.226:1026	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.231:1026	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.231:1033	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.226:1038	ESTABLISH	0	0
0.0.0.0:3260	0.0.0.0:0	LISTEN	0	0

Example 17-35 Displays SCSI Stats (detail)

```
switch# show ips stats tcp interface gigabitethernet 2/1 detail
TCP Statistics for port GigabitEthernet2/1
  TCP send stats
    150953931 segments, 2755572300 bytes
    53986369 data, 82341597 ack only packets
    4 control (SYN/FIN/RST), 0 probes, 14625949 window updates
    12 segments retransmitted, 576 bytes
    12 retransmitted while on ethernet send queue, 0 packets split
    118741734 delayed acks sent
  TCP receive stats
    640780835 segments, 640325552 data packets in sequence, 925034009772 bytes in
sequence
    0 predicted ack, 615117910 predicted data
    0 bad checksum, 0 multi/broadcast, 0 bad offset
    0 no memory drops, 0 short segments
    0 duplicate bytes, 0 duplicate packets
    0 partial duplicate bytes, 0 partial duplicate packets
    0 out-of-order bytes, 0 out-of-order packets
    0 packet after window, 0 bytes after window
    0 packets after close
    25656078 acks, 2755572210 ack bytes, 0 ack toomuch, 5786 duplicate acks
    0 ack packets left of snd_una, 0 non-4 byte aligned packets
    12100 window updates, 0 window probe
    29 pcb hash miss, 17 no port, 0 bad SYN, 0 paws drops
  TCP Connection Stats
    0 attempts, 6 accepts, 6 established
    4 closed, 4 drops, 0 conn drops
    0 drop in retransmit timeout, 4 drop in keepalive timeout
    0 drop in persist drops, 0 connections drained
  TCP Miscellaneous Stats
    21635776 segments timed, 21642712 rtt updated
    12 retransmit timeout, 0 persist timeout
    8494 keepalive timeout, 8490 keepalive probes
  TCP SACK Stats
    0 recovery episodes, 0 data packets, 0 data bytes
    0 data packets retransmitted, 0 data bytes retransmitted
    0 connections closed, 0 retransmit timeouts
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
TCP SYN Cache Stats
  6 entries, 6 connections completed, 0 entries timed out
  0 dropped due to overflow, 0 dropped due to RST
  0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
  0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
  0 hash collisions, 0 retransmitted
TCP Active Connections
  Local Address      Remote Address      State      Send-Q    Recv-Q
  10.48.69.250:3260   10.48.69.226:1026   ESTABLISH   0          0
  10.48.69.250:3260   10.48.69.231:1026   ESTABLISH   0          0
  10.48.69.250:3260   10.48.69.231:1033   ESTABLISH   0          0
  10.48.69.250:3260   10.48.69.226:1038   ESTABLISH   0          0
  0.0.0.0:3260        0.0.0.0:0          LISTEN      0          0
```

The **show ips stats buffer** command displays information about the iSCSI buffers. See Example 17-36

Example 17-36 Displays iSCSI Buffers

```
switch# show ips stats buffer interface gigabitethernet 4/2
Mbuf Statistics for port GigabitEthernet4/2
Free Mbufs           : 83221
Mbuf high watermark   : 124830
Mbuf low watermark    : 20805
Mbuf alloc failures   : 0
Total clusters        : 2304
Free Clusters          : 80145
Clusters high watermark : 87381
Clusters low watermark : 79059
Clusters alloc failures : 0
Free shared mbufs      : 0
Shared Mbuf alloc failures : 0
Free shared clusters   : 0
Shared clusters alloc failures: 0

Ether channel Statistics for port GigabitEthernet4/2
TCP segments sent      : 0
TCP segments received  : 0
Xmit packets sent      : 0
Xmit packets received  : 0
Config packets sent    : 0
Config packets received : 0
MPQ packets send errors : 0
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See Example 17-37.

Example 17-37 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdfg

username:user2
secret:cshadhdhsadadjajdjas
```

Send documentation comments to mdsfeedback-doc@cisco.com

iSCSI High Availability

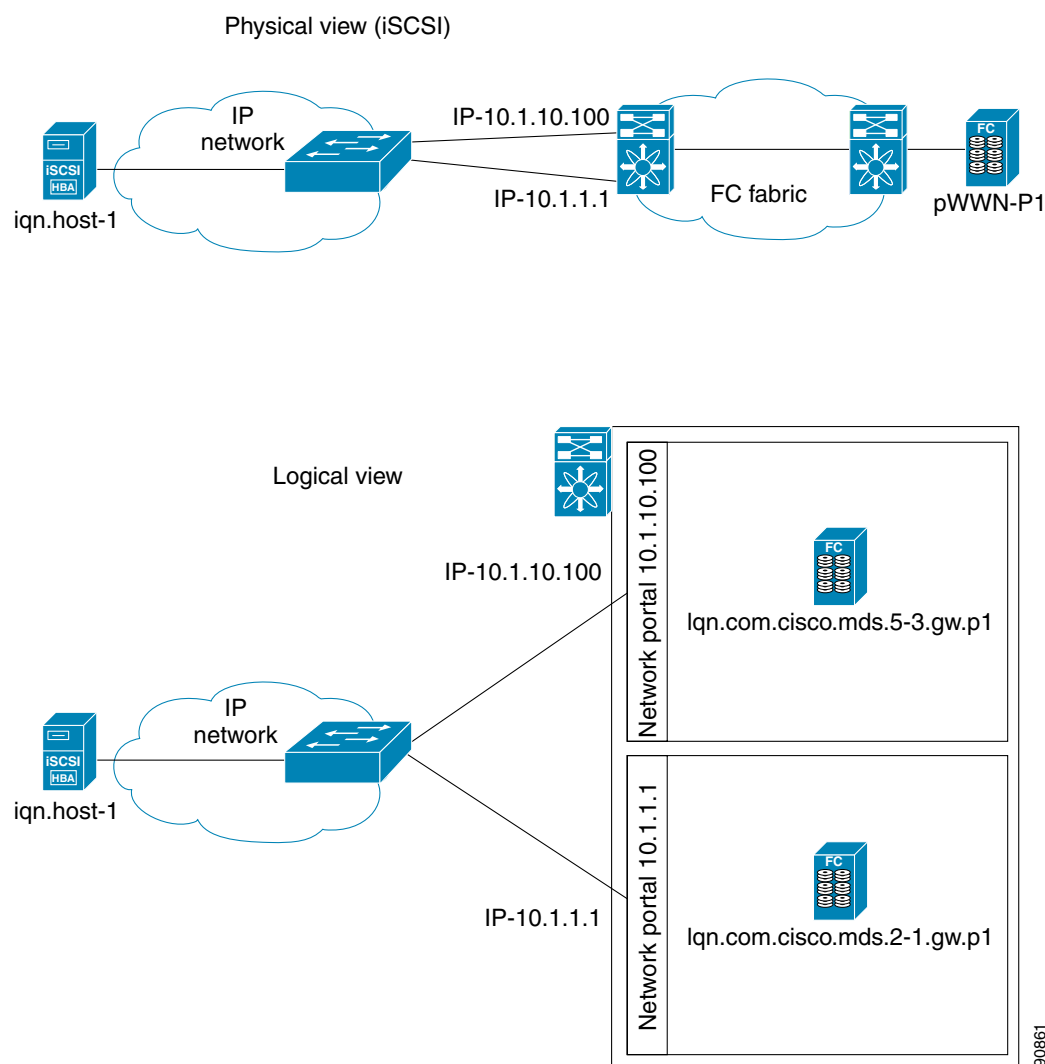
The following high availability features are available for iSCSI configurations:

- [Multiple IPS Ports Connected to the Same IP Network](#), page 17-62
- [VRRP-Based High Availability](#), page 17-63
- [Ethernet PortChannel-Based High Availability](#), page 17-64

Multiple IPS Ports Connected to the Same IP Network

Figure 17-27 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 17-27 Multiple Gigabit Ethernet Interfaces in the Same IP Network



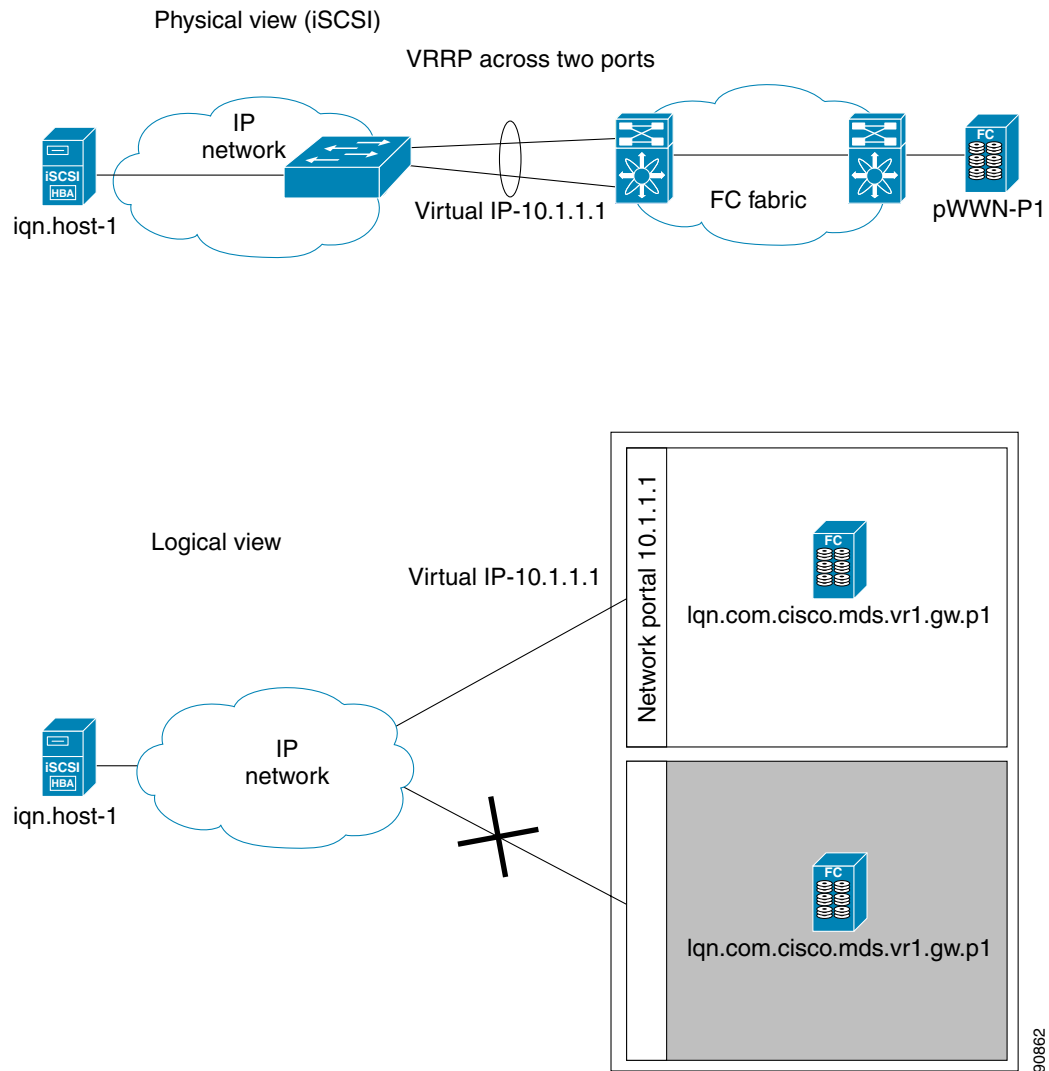
Send documentation comments to mdsfeedback-doc@cisco.com

In Figure 17-27, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 17-28 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 17-28 VRRP-Based iSCSI High Availability



In Figure 17-28, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Send documentation comments to mdsfeedback-doc@cisco.com

Ethernet PortChannel-Based High Availability




Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that iSCSI link.

Figure 17-29 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 17-29 Ethernet PortChannel-Based iSCSI High Availability



In Figure 17-29, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the virtual iSCSI target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

iSCSI Authentication Setup Guidelines

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- [No Authentication, page 17-64](#)
- [CHAP with Local Password Database, page 17-65](#)
- [CHAP with External RADIUS Server, page 17-65](#)
- [Scenario 1, page 17-66](#)
- [Scenario 2, page 17-71](#)



Note

This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before issuing any command.

No Authentication

To configure a network with no authentication set the iSCSI authentication method to none.

```
switch(config)# iscsi authentication none
```


Send documentation comments to mdsfeedback-doc@cisco.com

CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

- Step 1** Set the AAA authentication to use the local password database for iSCSI protocol.

```
switch(config)# aaa authentication iscsi local
```

- Step 2** Set the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 3** Configure the user names and passwords for iSCSI users.

```
switch(config)# username iscsi-user password abcd iscsi
```



Note If you do not specify the **iscsi** option, the user name is assumed to be a MDS switch login user instead of an iSCSI user.

- Step 4** Verify the global iSCSI authentication setup.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----verify
  Import FC Target: Disabled
  ...
```

CHAP with External RADIUS Server

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

- Step 1** Setup the authentication verification for iSCSI protocol to go to RADIUS server.

```
switch(config)# aaa authentication iscsi radius
```

- Step 2** Configure the RADIUS server IP address.

```
switch(config)# radius-server host 10.1.1.10
```

- Step 3** Password for the MDS as RADIUS client to the RADIUS server.

```
switch(config)# radius-server key mds-1
```

- Step 4** Setup the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 5** Verify that the global iSCSI authentication setup is CHAP.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----- Verify
  Import FC Target: Disabled
  ...
```

- Step 6** Verify that the RADIUS shared secret is MDS-1.

```
switch# show radius-server
Global RADIUS shared secret:mds-1 <----- Verify
  ...
```

Send documentation comments to mdsfeedback-doc@cisco.com

To configure an iSCSI RADIUS server, follow these steps:

- Step 1** Configure the RADIUS server to allow access from the MDS switch's management Ethernet IP address.
- Step 2** Configure the shared secret for the RADIUS server to authenticate the MDS switch.
- Step 3** Configure the iSCSI users and passwords on the RADIUS server.
- Step 4** Verify that the switch can communicate using the switch's management Ethernet IP address.

```
switch# show iscsi global
global iSCSI authentication method is CHAP <----- Authentication is CHAP
...
```

- Step 5** Verify that the RADIUS server can communicate using the switch's management Ethernet IP address.

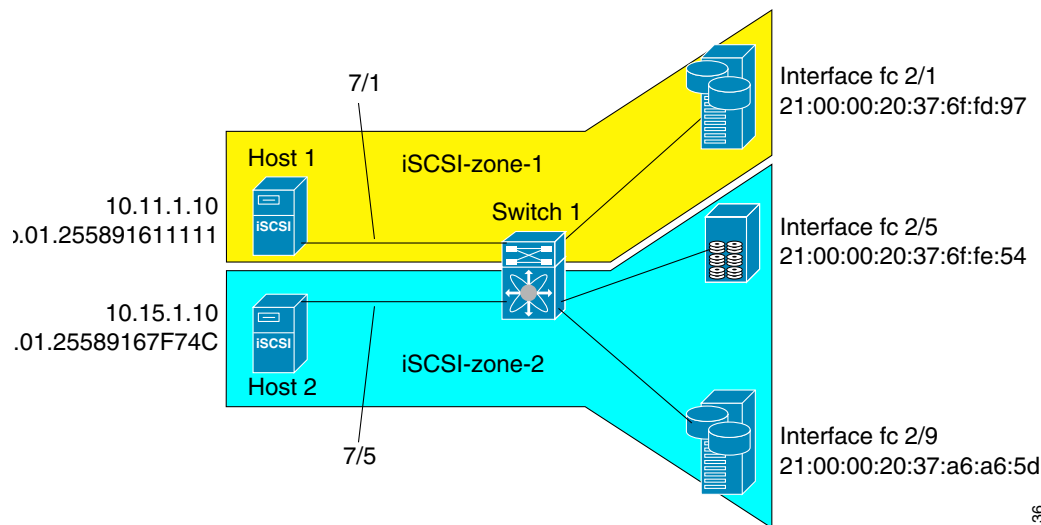
```
switch# show radius-server
Global RADIUS shared secret:mds-1 <----- Secret is mds-1
retransmission count:1
timeout value:1
following RADIUS servers are configured:
    10.1.1.100: <----- RADIUS server's IP address
...
```

Scenario 1

Sample scenario 1 assumes the following configuration (see [Figure 17-30](#)):

- Access control using Fibre Channel zoning.
- No target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator identified using IP address (Host 1 = 10.11.1.10 and Host 2 = 10.15.1.11).
- iSCSI initiator identified using node name (Host 1 = iqn.1987-05.com.cisco:01.e41695d16b1a and Host 2 = iqn.1987-05.com.cisco:01.25589167f74c).

Figure 17-30 iSCSI Scenario 1



34136

Send documentation comments to mdsfeedback-doc@cisco.com

To configure scenario 1 (see [Figure 17-30](#)), follow these steps:

- Step 1** Configure null authentication for all iSCSI hosts.

```
switch(config)# iscsi authentication none
```

- Step 2** Configure iSCSI to dynamically import all FC targets into the iSCSI SAN using auto-generated iSCSI target names.

```
switch(config)# iscsi import target fc
```

- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```

- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by the IP address, and enable the interface.

```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```

- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name, and enable the interface.

```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

- Step 7** Verify the available Fibre Channel targets (see [Figure 17-30](#)).

```
switch# show fcns database
VSAN 1:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x6d0001	NL	21:00:00:20:37:6f:fd:97	(Seagate)	scsi-fcp:target
0x6d0101	NL	21:00:00:20:37:6f:fe:54	(Seagate)	scsi-fcp:target
0x6d0201	NL	21:00:00:20:37:a6:a6:5d	(Seagate)	scsi-fcp:target

Total number of entries = 3

- Step 8** Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.



Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

```
switch(config)# zone name iscsi-zone-1 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fd:97
switch(config-zone)# member symbolic-nodename 10.11.1.10
```

Send documentation comments to mdsfeedback-doc@cisco.com

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two FC targets in it.



Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

```
switch(config)# zone name iscsi-zone-2 vsan 1
switch(config-zone)# member pwwn 21:00:00:20:37:6f:fe:54
switch(config-zone)# member pwwn 21:00:00:20:37:a6:a6:5d
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c
```

Step 10 Create a zoneset and add the two zones as members.

```
switch(config)# zoneset name zoneset-iscsi vsan 1
switch(config-zoneset)# member iscsi-zone-1
switch(config-zoneset)# member iscsi-zone-2
```

Step 11 Activate the zoneset.

```
switch(config)# zoneset activate name zoneset-iscsi vsan 1
```

Step 12 Display the active zoneset.



Note The iSCSI hosts has not connected so they do not have a FCID yet.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97] <-----Target
      symbolic-nodename 10.11.1.10 <-----iSCSI host (host 1)

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54] <-----Target
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d] <-----Target
      symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c <---iSCSI host (host 2)
```

Step 13 Check all iSCSI hosts to verify their connectivity.

Step 14 Show all the iSCSI sessions (use the **detail** option for detailed information).

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.25589167f74c
Initiator ip addr (s): 10.15.1.11
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54
```



Note The last part of the auto-created target name is the FC target's pWWN.

```
VSAN 1, ISID 00023d000001, Status active, no reservation
```

```
Session #2
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d
VSAN 1, ISID 00023d000001, Status active, no reservation
```

```
Initiator 10.11.1.10
Initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97
VSAN 1, ISID 00023d000001, Status active, no reservation
```

Send documentation comments to mdsfeedback-doc@cisco.com

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.25589167f74c <-----
  Initiator ip addr (s): 10.15.1.11
  iSCSI alias name: oasis11.cisco.com
  Node WWN is 20:02:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:03:00:0b:fd:44:68:c2 (dynamic)
  Interface iSCSI 7/5, Portal group tag: 0x304
  VSAN ID 1, FCID 0x6d0300

iSCSI Node name is 10.11.1.10 <-----
  iSCSI Initiator name: iqn.1987-05.com.cisco:01.e41695d16b1a
  iSCSI alias name: oasis10.cisco.com
  Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:05:00:0b:fd:44:68:c2 (dynamic)
  Interface iSCSI 7/1, Portal group tag: 0x300
  VSAN ID 1, FCID 0x6d0301
```

Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Initiator ID based on IP address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FCIDs are resolved.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x6d0301 [symbolic-nodename 10.11.1.10] <-----

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d]
    * fcid 0x6d0300 [symbolic-nodename
iqn.1987-05.com.cisco:01.25589167f74c]<-----
```

FCID resolved for iSCSI host

FCID for iSCSI host

Step 17 The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE
-----
0x6d0001      NL    21:00:00:20:37:6f:fd:97 (Seagate)         scsi-fcp:target
0x6d0101      NL    21:00:00:20:37:6f:fe:54 (Seagate)         scsi-fcp:target
0x6d0201      NL    21:00:00:20:37:a6:a6:5d (Seagate)         scsi-fcp:target
0x6d0300      N     20:03:00:0b:fd:44:68:c2 (Cisco)          scsi-fcp:init isc..w
0x6d0301      N     20:05:00:0b:fd:44:68:c2 (Cisco)          scsi-fcp:init isc..w
```

Send documentation comments to mdsfeedback-doc@cisco.com

Step 18 Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

```
switch# show fcns database fcid 0x6d0300 detail vsan 1
```

```
-----
VSAN:1      FCID:0x6d0300
-----
```

```
port-wwn (vendor)      :20:03:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:02:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr           :10.15.1.11  <-----
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name     :
```

**IP address of the iSCSI
host**

iSCSI gateway node

```
symbolic-node-name
```

```
:ign.1987-05.com.cisco:01.25589167f74c<-----
```

**iSCSI initiator ID is
based on the registered
node name**

```
port-type             :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:91:00:0b:fd:44:68:c0
hard-addr              :0x000000
Total number of entries = 1
```

```
switch# show fcns database fcid 0x6d0301 detail vsan 1
```

```
-----
VSAN:1      FCID:0x6d0301
-----
```

```
port-wwn (vendor)      :20:05:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:04:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr           :10.11.1.10
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name     :
```

iSCSI gateway node

```
symbolic-node-name     :10.11.1.10  <-----
```

**iSCSI initiator ID is
based on the IP address
registered in
symbolic-node-name
field**

```
port-type             :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:81:00:0b:fd:44:68:c0
hard-addr              :0x000000
```


Send documentation comments to mdsfeedback-doc@cisco.com

Scenario 2

Sample scenario 2 (see [Figure 17-31](#)) assumes the following configuration:

- Access control based on Fibre Channel zoning.
- Target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator assigned to different VSANs.

Figure 17-31 iSCSI Scenario 2



To configure scenario 1 (see [Figure 17-31](#)), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all FC targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iscsi initiators by the IP address, and enable the interface.
- ```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iscsi initiators by IP address, and enable the interface.

```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

- Step 7** Add static configuration for each iSCSI initiator.

```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
switch(config-(iscsi-init))# static pwwn system-assign 1
switch(config-(iscsi-init))# static nwwn system-assign

switch(config)# iscsi initiator ip address 10.15.1.11
switch(config-(iscsi-init))# static pwwn system-assigned 1
switch(config-(iscsi-init))# vsan 2
```

- Step 8** View the configured initiators.



**Note** The WWNs are assigned by the system. The initiators are members of different VSANs.

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
 Member of vsans: 1
 Node WWN is 20:03:00:0b:fd:44:68:c2
 No. of PWWN: 1
 Port WWN is 20:02:00:0b:fd:44:68:c2

iSCSI Node name is 10.15.1.11
 Member of vsans: 2
 No. of PWWN: 1
 Port WWN is 20:06:00:0b:fd:44:68:c2
```

- Step 9** Create a zone with Host 1.

```
switch(config)# zone name iscsi-zone-1 vsan 1
```

- Step 10** Add three members to the zone named *iscsi-zone-1*.



**Note** Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN can be used. In this case, the pWWN is persistent.

- Based on the symbolic node name.

```
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
```

- Based on the persistent pWWN assigned to the initiator. You can obtain the pWWN from the **show iscsi initiator** output.

```
switch(config-zone)# member pwwn 20:02:00:0b:fd:44:68:c2
switch(config-zone)# member pwwn 21:00:00:20:37:6f:fd:97
```

- Step 11** Create zone with Host 2 and two Fibre Channel targets.



**Note** If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

```
switch(config)# zone name iscsi-zone-2 vsan 2
```



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 12** Activate the zoneset in VSAN 2

```
switch(config)# zoneset activate name iscsi-zoneset-v2 vsan 2
Zoneset activation initiated. check zone status
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
 zone name iscsi-zone-2 vsan 2
 * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
 * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
 pwwn 20:06:00:0b:fd:44:68:c2
```

**Step 13** Start the iSCSI clients on both hosts and verify that sessions come up.

**Step 14** Display the iSCSI sessions to verify the Fibre Channel target and the configured WWNs.

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.e41695d16b1a
 Initiator ip addr (s): 10.11.1.10
 Session #1
 Discovery session, ISID 00023d000001, Status active

 Session #2
 Target
iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97<----- To FC target
 VSAN 1, ISID 00023d000001, Status active, no reservation
```

**Step 15** Display the iSCSI initiator to verify the configured nWWN and pWWN.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
 Initiator ip addr (s): 10.11.1.10
 iSCSI alias name: oasis10.cisco.com

 Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----- The configured nWWN
 Member of vsans: 1
 Number of Virtual n_ports: 1

 Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<----- The configured pWWN
 Interface iSCSI 7/1, Portal group tag: 0x300
 VSAN ID 1, FCID 0x680102
```

**Step 16** Check the Fibre Channel name server

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x680001 NL 21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102 N 20:02:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w <----- iSCSI initiator in
 name server
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 17** Verify the details of the iSCSI initiator's FCID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1

VSAN:1 FCID:0x680102

port-wwn (vendor) :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:03:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.11.1.10
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:81:00:0b:fd:44:68:c0
iSCSI alias name: oasis10.cisco.com

Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----- The configured nWWN
Member of vsans: 1
Number of Virtual n_ports: 1

Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<----- The configured pWWN
Interface iSCSI 7/1, Portal group tag: 0x300
VSAN ID 1, FCID 0x680102
```

**Step 18** Check the Fibre Channel name server

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x680001 NL 21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target
0x680102 N 20:02:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w <-- ---- iSCSI initiator in name server
```

**Step 19** Verify the details of the iSCSI initiator's FCID in the name server

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1

VSAN:1 FCID:0x680102

port-wwn (vendor) :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:03:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.11.1.10
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:81:00:0b:fd:44:68:c0
hard-addr :0x000000
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 20** Verify that zoning has resolved the FCID for the iSCSI client.

```
switch# show zoneset active vsan 1
zoneset name iscsi-zoneset-v1 vsan 1
zone name iscsi-zone-1 vsan 1
* fcid 0x680001 [pwwn 21:00:00:20:37:6f:fd:97]
* fcid 0x680102 [pwwn 20:02:00:0b:fd:44:68:c2]
```

**Step 21** Do the same to verify that the second initiator is connected to the two Fibre Channel targets in VSAN 2.

```
switch# show iscsi session initiator 10.15.1.11
Initiator 10.15.1.11
Initiator name ign.1987-05.com.cisco:01.25589167f74c
Session #1
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <-- Session to
VSAN 2, ISID 00023d000001, Status active, no reservation first target

Session #2
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d <-- Session to
VSAN 2, ISID 00023d000001, Status active, no reservation second
target

switch# show iscsi initiator
iSCSI Node name is 10.15.1.11 <--- Initiator ID is the IP address
iSCSI Initiator name: ign.1987-05.com.cisco:01.25589167f74c
iSCSI alias name: oasis11.cisco.com

Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <----- Dynamic
Member of vsans: 2 <--- vsan membership WWN as
Number of Virtual n_ports: 1 static WWN
not
assigned

Virtual Port WWN is 20:06:00:0b:fd:44:68:c2 (configured) <----- Static
Interface iSCSI 7/5, Portal group tag: 0x304 pWWN for
VSAN ID 2, FCID 0x750200 the initiator

switch# show fcns database vsan 2
VSAN 2:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x750001 NL 21:00:00:20:37:6f:fe:54 (Seagate) scsi-fcp:target
0x750101 NL 21:00:00:20:37:a6:a6:5d (Seagate) scsi-fcp:target

0x750200 N 20:06:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w <-- iSCSI
Total number of entries = 3 initiator
entry in
name server
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```
switch# show fcns database fcid 0x750200 detail vsan 2
```

```

VSAN:2 FCID:0x750200

port-wwn (vendor) :20:06:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:04:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.15.1.11
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :10.15.1.11
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:91:00:0b:fd:44:68:c0
hard-addr :0x000000
Total number of entries = 1
```

```
switch# show zoneset active vsan 2
```

```
zoneset name iscsi-zoneset-v2 vsan 2
 zone name iscsi-zone-2 vsan 2
 * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
 * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
```

```
* fcid 0x750200 [pwwn 20:06:00:0b:fd:44:68:c2] <-----
```

**FCID  
resolved for  
iSCSI  
initiator**

## Default IP Storage Settings

Table 17-2 lists the default settings for Gigabit Ethernet parameters.

**Table 17-2 Default Gigabit Ethernet Parameters**

| Parameters        | Default                           |
|-------------------|-----------------------------------|
| IP MTU frame size | 1500 bytes for all Ethernet ports |

Table 17-3 lists the default settings for FCIP parameters.

**Table 17-3 Default FCIP Parameters**

| Parameters                | Default            |
|---------------------------|--------------------|
| TCP default port for FCIP | 3225               |
| minimum-retransmit-time   | 300 milliseconds.  |
| keepalive-timeout         | 60 seconds.        |
| max-retransmissions       | 4 retransmissions. |
| PMTU discovery            | Enabled.           |
| pmtu-enable reset-timeout | 3600 seconds.      |
| SACK                      | Enabled.           |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 17-3 Default FCIP Parameters**

| Parameters                              | Default                     |
|-----------------------------------------|-----------------------------|
| max-bandwidth                           | 1G.                         |
| min-available-bandwidth                 | 2 Mbps.                     |
| round-trip-time                         | 10ms.                       |
| buffer size                             | 0 KB.                       |
| Control TCP and data connection         | No packets are transmitted. |
| TCP congestion window monitoring        | Enabled                     |
| Burst size                              | 10KB.                       |
| TCP connection mode                     | active mode is enabled.     |
| special-frame                           | Disabled.                   |
| FCIP timestamp                          | Disabled.                   |
| acceptable-diff range to accept packets | + or - 1000 milliseconds.   |
| B port keepalive responses              | Disabled                    |

Table 17-4 lists the default settings for iSCSI parameters.

**Table 17-4 Default iSCSI Parameters**

| Parameters                                         | Default                                                                                                             |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Number of TCP connections                          | One per iSCSI session.                                                                                              |
| Fibre Channel targets to iSCSI                     | Not imported.                                                                                                       |
| Advertising iSCSI target                           | Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces |
| iSCSI hosts mapping to virtual Fibre Channel hosts | Dynamic mapping.                                                                                                    |
| Dynamic iSCSI initiators                           | Members of the default VSAN (VSAN 1).                                                                               |
| Identifying initiators                             | iSCSI node names.                                                                                                   |
| Advertising static virtual targets                 | No initiators allowed to access a virtual target (unless explicitly configured).                                    |
| iSCSI login authentication                         | CHAP or none authentication mechanism.                                                                              |
| Ethernet PortChannel IP address usage              | Source and destination IP addresses.                                                                                |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***