

Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services 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 it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

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

This chapter includes the following sections:

- [Services Modules, page 37-1](#)
- [Supported Hardware, page 37-4](#)
- [Configuring Gigabit Ethernet Interfaces, page 37-4](#)
- [IPS Module Core Dumps, page 37-19](#)
- [Default Settings, page 37-20](#)

Services Modules

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

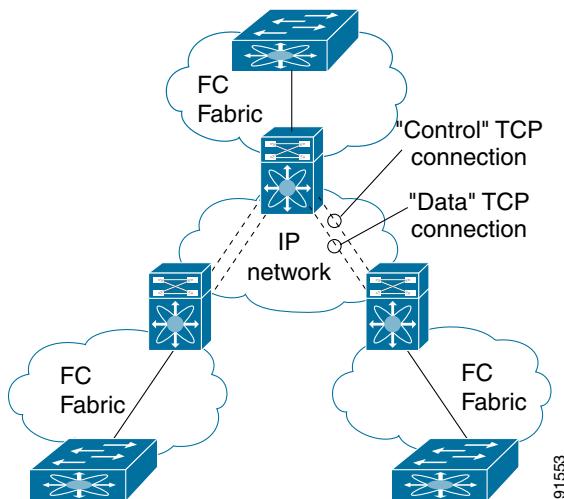
- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2)

Gigabit Ethernet ports in these modules can be configured to support FCIP protocol, iSCSI protocol, or both protocols simultaneously.

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

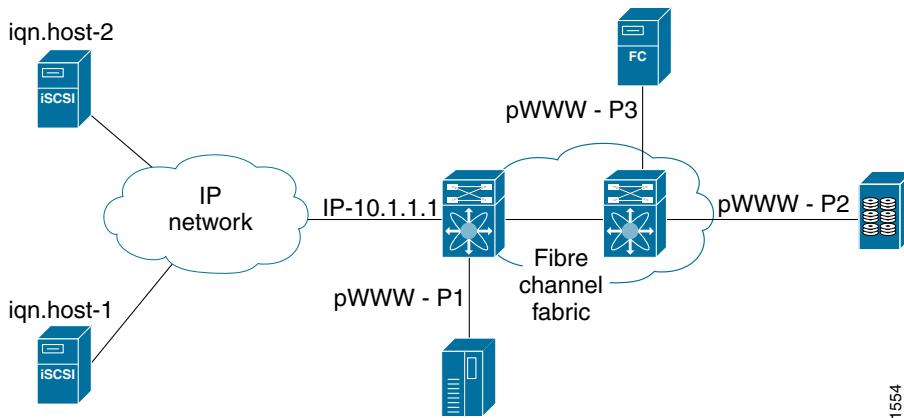
- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 37-1](#) shows how the IPS module is used in different FCIP scenarios.

Figure 37-1 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 Cisco MDS 9000 Family switch 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 37-2](#) depicts the iSCSI scenarios in which the IPS module is used.

Figure 37-2 iSCSI Scenarios



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

Module Status Verification

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

```
switch# show module
Mod Ports Module-Type Model Status
--- -----
1 0 Caching Services Module DS-X9560-SMAP ok
2 8 IP Storage Services Module DS-X9308-SMIP ok <-----IPS-8 module
4 16 2x1GE IPS, 14x1/2Gbps FC Module DS-X9216i-K9-SUP ok <-----MPS-14/2 module
5 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 active *
6 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 ha-standby
9 4 IP Storage Services Module DS-X9304-SMIP ok <-----IPS-4 module

Mod Sw Hw World-Wide-Name(s) (WWN)
--- -----
1 2.0(1) 0.201 20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
2 2.0(1) 0.201 20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
4 2.0(1) 0.201 20:c1:00:05:30:00:07:1e to 20:d0:00:05:30:00:07:1e
5 2.0(1) 0.0 --
6 2.0(1) 0.0 --
9 2.0(1) 0.1 22:01:00:05:30:00:07:1e to 22:04:00:05:30:00:07:1e

Mod Application Image Description Application Image Version
--- -----
1 svc-node1 1.3 (5M)
1 svc-node2 1.3 (5M)

Mod MAC-Address(es) Serial-Num
--- -----
1 00-05-30-01-49-c2 to 00-05-30-01-4a-46 JAB073907EP
2 00-05-30-00-9d-d2 to 00-05-30-00-9d-de JAB064605a2
4 00-05-30-01-7f-32 to 00-05-30-01-7f-38 JAB081405AM
5 00-05-30-00-2c-4e to 00-05-30-00-2c-52 JAB06350B1M
6 00-05-30-00-19-66 to 00-05-30-00-19-6a JAB073705GL
9 00-0d-bc-2f-d6-00 to 00-0d-bc-2f-d6-08 JAB080804TN
```

* this terminal session

IPS Module Upgrade



Caution

A software upgrade is only disruptive for the IPS module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

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

MPS-14/2 Module Upgrade



Caution A software upgrade is only partially disruptive for the MPS-14/2 module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and 2 Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one of more of the following hardware:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).



Note In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered as 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

Configuring Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 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.



Note For information about configuring FCIP, see [Chapter 33, “Configuring FCIP”](#). For information about configuring iSCSI, see [Chapter 35, “Configuring iSCSI”](#).

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 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.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed via the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.

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



Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.

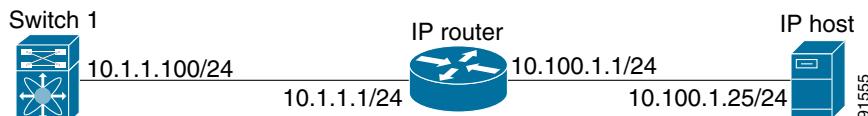


Tip Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration

Figure 37-3 shows an example of a basic Gigabit Ethernet configuration.

Figure 37-3 Gigabit Ethernet Configuration Example



The port on the Ethernet switch to which the MDS Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the ethernet switch) should be disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. FoR Cisco Ethernet switches, use either the **switchport host** command in IOS or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

To configure the Gigabit Ethernet interface for the example in Figure 37-3, 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 (slot 2, 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. |

Configuring Interface Descriptions

See the “Interface Descriptions” section on page 11-12 for details on configuring the switchport description for any interface.

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

Configuring Beacon Mode

See the “Beacon Mode” section on page 11-17 for details on configuring the beacon mode for any interface.

Configuring Auto-Negotiation

By default, auto-negotiation is enabled all Gigabit Ethernet interface. You can enable or disable auto-negotiation for a specified Gigabit Ethernet interface. When auto-negotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the auto-negotiation feature.

To configure auto-negotiation, 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 (slot 2, port 2). |
| Step 3 switch(config-if)# switchport auto-negotiate | Enables auto-negotiation for this Gigabit Ethernet interface (default). |
| switch(config-if)# no switchport auto-negotiate | Disables auto-negotiation for this Gigabit Ethernet interface. |

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note The minimum MTU size is 576 bytes.



Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

You do not need to explicitly issue the **shutdown** and **no shutdown** commands.

To configure the MTU frame size, 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 (slot 2, port 2). |
| Step 3 switch(config-if)# switchport mtu 3000 | Changes the MTU size to 3000 bytes. The default is 1500 bytes. |

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

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

To configure the promiscuous mode, 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 (slot 2, port 2). |
| Step 3 switch(config-if)# switchport promiscuous-mode on switch(config-if)# switchport promiscuous-mode off switch(config-if)# no switchport promiscuous-mode | Enables promiscuous mode for this Gigabit Ethernet interface. The default is off . Disables (default) promiscuous mode for this Gigabit Ethernet interface. Disables (default) the promiscuous mode for this Gigabit Ethernet interface. |

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.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.



Note

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name (the <slot-number>/<port-number>. <VLAN-ID>).

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

Configuring the VLAN Subinterface

To configure a VLAN subinterface (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 (slot 2, 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 10.1.1.101 255.255.255.0 | 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 interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 37-1](#)).

Table 37-1 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 ID 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 IDs cannot be configured in the same subnet. |
| Gigabit Ethernet 1/1 | Gigabit Ethernet 1/1.100 | No | A subinterface 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 37-1](#) also apply to Ethernet PortChannels.

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

Configuring Static IP Routing

To configure static IP routing (see [Figure 37-3](#)) 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** command takes the Gigabit Ethernet interface as a parameter and returns the route table for the interface. See [Example 37-1](#).

Example 37-1 Displays the IP 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

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.



Note

If the connection fails, verify the following, and ping the IP host again:

- 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.
- The Gigabit Ethernet interface is in the `up` state.

Use the **ping** command to verify the Gigabit Ethernet connectivity (see [Example 37-2](#)). 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 2-14).

Use the **show interface gigabitethernet** command to verify if the Gigabit Ethernet interface is up.

Example 37-2 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
```

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

```
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
```

Gigabit Ethernet IP-ACL Guidelines



Tip If IP-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See the “IP Access Control Lists” section on page 29-1 for information on configuring IP-ACLs.

Follow these guidelines when configuring IP-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note Other protocols like, User Datagram Protocol (UDP) and HTTP, are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IP-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established**, **precedence**, and **fragments** options are ignored when you apply IP-ACLs (containing these options) to Gigabit Ethernet interfaces.
 - If an IP-ACL rule applies to a pre-existing TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B and an IP-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Applying IP-ACLs on Gigabit Ethernet Interfaces

To apply an IP-ACL on an Gigabit Ethernet interface, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 3/1 switch(config-if)# | Configures a Gigabit Ethernet interface (3/1). |
| Step 3 | switch(config-if)# ip access-group SampleName | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for both ingress and egress traffic (if the association does not exist already). |

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

| Command | Purpose |
|---|---|
| <pre>switch(config-if)# ip access-group SampleName1 in</pre> <pre>switch(config-if)# ip access-group SampleName2 out</pre> | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for ingress traffic. |
| | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for egress traffic (if the association does not exist already). |

Displaying ARP Caches

You can display the ARP cache on Gigabit Ethernet interfaces.



Note Use the physical interface, not the subinterface, for all ARP cache commands.

Use the **show ips arp interface gigabitetherent** 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 37-3](#).

Example 37-3 Displays ARP Caches

```
switch# show ips arp interface gigabitetherent 7/1
Protocol      Address    Age (min)   Hardware Addr  Type     Interface
Internet      20.1.1.5      3          0005.3000.9db6  ARPA    GigabitEtherent7/1
Internet      20.1.1.10     7          0004.76eb.2ff5  ARPA    GigabitEtherent7/1
Internet      20.1.1.11     16         0003.47ad.21c4  ARPA    GigabitEtherent7/1
Internet      20.1.1.12      6          0003.4723.c4a6  ARPA    GigabitEtherent7/1
Internet      20.1.1.13     13         0004.76f0.ef81  ARPA    GigabitEtherent7/1
Internet      20.1.1.14      0          0004.76e0.2f68  ARPA    GigabitEtherent7/1
Internet      20.1.1.15      6          0003.47b2.494b  ARPA    GigabitEtherent7/1
Internet      20.1.1.17      2          0003.479a.b7a3  ARPA    GigabitEtherent7/1
...
...
```

Clearing ARP Caches

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache.

Use the **clear ips arp** command to clear the ARP cache. See [Example 37-4](#) and [Example 37-5](#).

Example 37-4 Clearing One ARP Cache Entry

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

Example 37-5 Clearing All ARP Cache Entries

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

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

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 37-6](#).

Example 37-6 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
    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 37-7 Displays the Gigabit Ethernet 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 gigabitetherne**t command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 37-8](#).



Note Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

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

Example 37-8 Displays Ethernet MAC Statistics

```
switch# show ips stats mac interface gigabitetherinet 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 gigabitetherinet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns DMA bridge statistics for that interface. See [Example 37-9](#).



Note

Use the physical interface, not the subinterface, to display DMA-bridge statistics.

Example 37-9 Displays DMA-Bridge Statistics

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

Use the **show ips stats ip interface gigabitetherinet** to display and verify IP statistics. This command takes the main Ethernet interface as a parameter and returns the IP statistics for that interface. See [Example 37-10](#).

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



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

Example 37-10 Displays IP Statistics

```
switch# show ips stats ip interface gigabitetherinet 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 gigabitetherinet** 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 [Example 37-11](#) and [Example 37-12](#).

Example 37-11 Displays TCP Statistics

```
switch# show ips stats tcp interface gigabitetherinet 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
```

Example 37-12 Displays Detailed TCP Statistics

```
switch# show ips stats tcp interface gigabitetherinet 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
```

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

```

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:                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 37-13](#).

Example 37-13 Displays ICMP Statistics

```

switch# show ips stats icmp interface gigabitethernet 2/1
ICMP Statistics for port GigabitEthernet2/1
  0 ICMP messages received
  0 ICMP messages dropped due to errors
  ICMP input histogram
    0 destination unreachable
    0 time exceeded
    0 parameter problem
    0 source quench
    0 redirect
    0 echo request
    0 echo reply
    0 timestamp request
    0 timestamp reply
    0 address mask request
    0 address mask reply
  ICMP output histogram
    0 destination unreachable
    0 time exceeded
    0 parameter problem
    0 source quench
    0 redirect
    0 echo request
    0 echo reply
    0 timestamp request
    0 timestamp reply
    0 address mask request
    0 address mask reply

```

Configuring Gigabit Ethernet High Availability

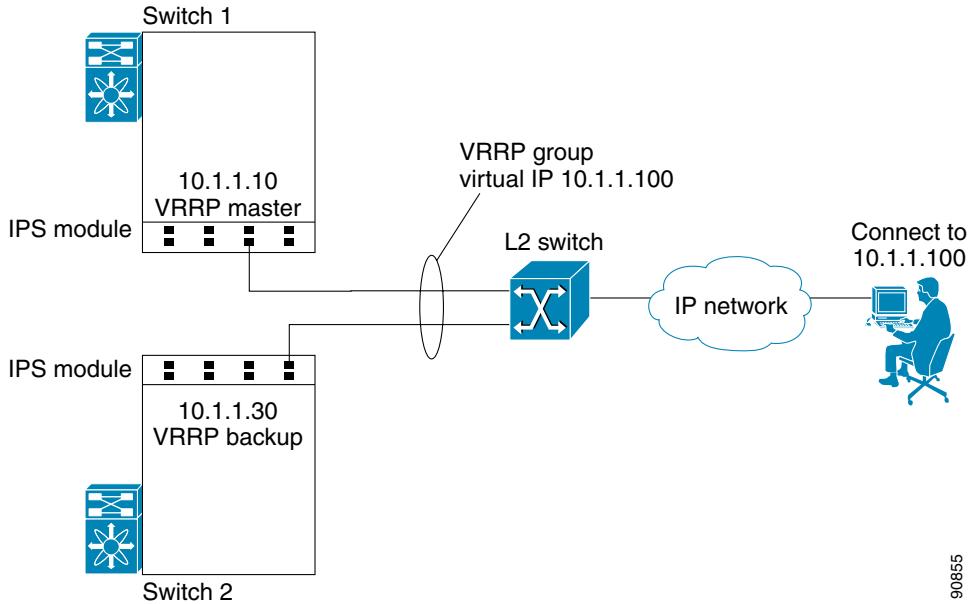
Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

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

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 37-4).

Figure 37-4 VRRP Scenario



90855

In Figure 37-4, 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 or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch
- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces

See the “The Virtual Router Redundancy Protocol” section on page 36-16.

Configuring VRRP for Gigabit Ethernet Interfaces

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 gigabitehernet 2/2 switch(config-if) # | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, 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. |

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

| Command | Purpose |
|---|---|
| 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 IPS Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

About Ethernet PortChannel Aggregation

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.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that FCIP link.



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

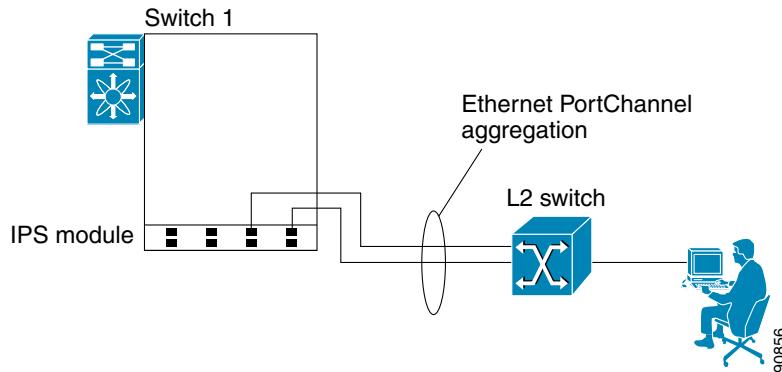
Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 37-5](#)).



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

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

Figure 37-5 Ethernet PortChannel Scenario



In Figure 37-5, Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.



Note 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.

Configuring Ethernet PortChannels

The PortChannel configuration specified in Chapter 13, “Configuring PortChannels” also applies to Ethernet PortChannel configurations.

To configure Ethernet PortChannels, follow these steps:

| Command | Purpose |
|---|---|
| Step 1 switch1# config 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. |
| Step 7 switch(config-if)# no shutdown | Enables the selected interface. |

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

| Command | Purpose |
|---|--|
| Step 8 switch(config)# interface gigabitethernet 9/4 switch(config-if)# | Configures the specified Gigabit Ethernet interface (slot 9, port 4). |
| Step 9 switch(config-if)# channel-group 10 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 switch(config-if)# no shutdown | Enables the selected interface. |

**Note**

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

- The interface already has an IP address assigned.
- The subinterfaces are configured on that interface.
- The interface already has an associated IP-ACL rule and the PortChannel does not.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS module or MPS-14/2 module.

See the “[Configuring CDP](#)” section on page [4-32](#).

IPS Module 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 useful to obtain a 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. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

Use the **show cores** command to list these files.

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module due to its large space requirement. They are copied directly to an external TFTP server.

Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

Default Settings

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

To configure IPS core dumps on the IPS module, follow these steps:

| Command | Purpose |
|--|---|
| Step 1 switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 switch(config)# ips core dump full ips core dump full' successfully set for module 9 | Configures a dump of the full core generation for all IPS modules in the switch. |
| switch(config)# no ips core dump full ips core dump partial' successfully set for module 9 | Configures a dump of the partial core generation for the IPS module in slot 9. |

Default Settings

Table 37-2 lists the default settings for Gigabit Ethernet parameters.

Table 37-2 Default Gigabit Ethernet Parameters

| Parameters | Default |
|-------------------|-----------------------------------|
| IP MTU frame size | 1500 bytes for all Ethernet ports |
| Auto-negotiation | Enabled. |
| Promiscuous mode | Disabled |