



Configuring IPv4 for Gigabit Ethernet Interfaces

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. This chapter describes how to configure IPv4 addresses and other IPv4 features.

This chapter includes the following topics:

- [Overview of IPv4, on page 1](#)
- [Basic Gigabit Ethernet Configuration for IPv4, on page 2](#)
- [Configuring Autonegotiation, on page 3](#)
- [Configuring the MTU Frame Size, on page 3](#)
- [Configuring Promiscuous Mode, on page 4](#)
- [Verifying Gigabit Ethernet Connectivity, on page 5](#)
- [VLANs Support in Cisco MDS NX-OS, on page 5](#)
- [Configuring Static IPv4 Routing, on page 7](#)
- [Displaying the IPv4 Route Table, on page 7](#)
- [IPv4-Access Control Lists, on page 7](#)
- [Address Resolution Protocol Cache, on page 9](#)
- [Displaying IPv4 Statistics, on page 10](#)
- [Default Settings for IPv4 Parameters, on page 10](#)

Overview of IPv4

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each Fibre Channel module with IPS ports or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each Fibre Channel module with IPS ports 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 through the single Fibre Channel host.



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

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each Fibre Channel module with IPS ports or MSM-18/4 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 each Fibre Channel module with IPS ports or MSM-18/4 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 through the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems do not require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.



Note The Gigabit Ethernet interfaces on the MSM-18/4 module do not support EtherChannel.

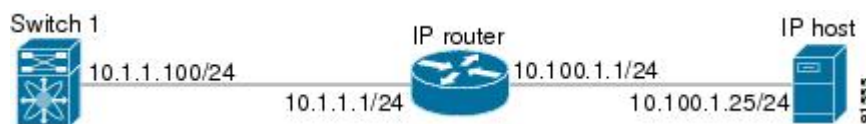


Tip Gigabit Ethernet ports on any Fibre Channel module with IPS ports or MSM-18/4 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port. They should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration for IPv4

The following figure shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 1: Gigabit Ethernet IPv4 Configuration Example



Note 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 Cisco IOS or the **set port host** command in the Catalyst OS.

To configure the Gigabit Ethernet interface for the example in the above image, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface gigabitethernet 2/2	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 IPv4 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 Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation 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 autonegotiation feature.



Note When using DS-SFP-GE-T (copper SFPs) on Gigabit Ethernet interfaces in a DS-X9316-SSNK9 module, auto-negotiation should be disabled.

To configure autonegotiation, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# switchport auto-negotiate	Enables autonegotiation for this Gigabit Ethernet interface (default).
Step 4	switch(config-if)# no switchport auto-negotiate	Disables autonegotiation 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:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2	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.

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:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# switchport promiscuous-mode on	Enables promiscuous mode for this Gigabit Ethernet interface. The default is off .
Step 4	switch(config-if)# switchport promiscuous-mode off	Disables (default) promiscuous mode for this Gigabit Ethernet interface.
Step 5	switch(config-if)# no switchport promiscuous-mode	Disables (default) the promiscuous mode for this Gigabit Ethernet interface.

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.

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 7-1](#)). The ping command sends echo request packets out to a remote device at an IP address that you specify.

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

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
```

VLANs Support in Cisco MDS NX-OS

This section describes virtual LAN (VLAN) support in Cisco MDS NX-OS and includes the following topics:

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.

If the Fibre Channel module with IPS ports or MSM-18/4 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:

**Note**

- The Ethernet switch port connected to the Fibre Channel module with IPS ports or MSM-18/4 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:

slot-number / port-number.VLAN-ID

Configuring a VLAN Subinterface

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

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.101 255.255.255.0	Enters the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# 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 the following table).

Table 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.

Interface 1	Interface 2	Same Subnet Allowed	Notes
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 the above table also apply to Ethernet PortChannels.

Configuring Static IPv4 Routing

To configure static IPv4 routing (see the above figure) through the Gigabit Ethernet interface, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1	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 IPv4 address of the router connected to the Gigabit Ethernet interface.

Displaying the IPv4 Route Table

The **ip route interface** command takes the Gigabit Ethernet interface as a parameter and returns the route table for the interface.

Displaying an 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.

IPv4-Access Control Lists

This section describes the guidelines for IPv4 access control lists and how to apply them to Gigabit Ethernet interfaces.

This section includes the following topics:

Gigabit Ethernet Guidelines

Follow these guidelines when configuring for Gigabit Ethernet interfaces:

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



Note

Other protocols such as 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 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** option is ignored when you apply containing this option to Gigabit Ethernet interfaces.
 - If a 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 IPv4-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.



Note

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. For information on configuring IPv4-ACLs, see the Cisco Fabric Manager Security Configuration Guide and Cisco MDS 9000 Family NX-OS Security Configuration Guide.

Applying on Gigabit Ethernet Interfaces

To apply on a Gigabit Ethernet interface, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 3/1	Configures a Gigabit Ethernet interface (3/1).
Step 3	switch(config-if)# ip access-group SampleName	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for both ingress and egress traffic (if the association does not exist already).

	Command or Action	Purpose
Step 4	switch(config-if)# ip access-group SampleName1 in	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for ingress traffic.
Step 5	switch(config-if)# ip access-group SampleName2 out	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for egress traffic (if the association does not exist already).

Address Resolution Protocol Cache

Cisco MDS NX-OS supports Address Resolution Protocol (ARP) cache for Gigabit Ethernet interface configured for IPv4. This section includes the following topics:

Displaying ARP Cache

You can display the ARP cache on Gigabit Ethernet interfaces.

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

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.

Displaying ARP Cache

```
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
.
.
.
```

Clearing ARP Cache

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.

Clearing an ARP Cache Entry

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

Clearing All ARP Cache Entries

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

Displaying IPv4 Statistics

Use the **show ips stats ip interface gigabitethernet** to display and verify IP v4 statistics. This command takes the main Ethernet interface as a parameter and returns the IPv4 statistics for that interface.

Use the physical interface, not the subinterface, to display IPv4 statistics.

Displaying IPv4 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
```

Default Settings for IPv4 Parameters

The below table lists the default settings for IPv4 parameters.

Table 2: Default IPv4 Parameters

Parameters	Default
IPv4 MTU frame size	1500 bytes for all Ethernet ports
Autonegotiation	Enabled
Promiscuous mode	Disabled