



CHAPTER 51

Configuring IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMSs):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding or in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.



Note

For information about configuring IPv6, see [Chapter 54, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

This chapter includes the following sections:


- [Traffic Management Services, page 51-2](#)
- [Management Interface Configuration, page 51-2](#)
- [Default Gateway, page 51-3](#)
- [IPv4 Default Network Configuration, page 51-4](#)
- [IPFC, page 51-5](#)
- [IPv4 Static Routes, page 51-5](#)
- [Overlay VSANs, page 51-6](#)
- [Multiple VSAN Configuration, page 51-7](#)
- [Virtual Router Redundancy Protocol, page 51-8](#)
- [DNS Server Configuration, page 51-12](#)
- [Default Settings, page 51-13](#)

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

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running the IP protocol over an FC interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric (see [Figure 51-1](#)).

Figure 51-1 Management Access to Switches



Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure IP version 4 (IPv4) parameters (IP address, subnet mask) or an IP version 6 (IPv6) address and prefix length so that the switch is reachable. For information on configuring IPv6 addresses, see [Chapter 54](#), “Configuring IPv6 for Gigabit Ethernet Interfaces.”

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.



Note

The port on the Ethernet switch to which the MDS management 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 MDS 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.

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

To configure the mgmt0 Ethernet interface using Device Manager for IPv6, follow these steps:

-
- Step 1** Select **Interface > Mgmt > Mgmt0**.
 - Step 2** Enter the description.
 - Step 3** Select the administrative state of the interface.
 - Step 4** Check the **CDP** check box to enable CDP.
 - Step 5** Enter the IP address mask.
 - Step 6** Click **Apply** to apply the changes.
-

Default Gateway

You can configure a default gateway IPv4 address on your Cisco MDS 9000 Family switch.

This section includes the following topics:

- [About the Default Gateway, page 51-3](#)
- [Configuring the Default Gateway, page 51-3](#)

About the Default Gateway

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes (IP default network, destination prefix, and destination mask, and next hop address).

**Tip**

If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

See the [“Initial Setup Routine” section on page 2-2](#) for more information on configuring the IP addresses for all entries in the switch.

Configuring the Default Gateway

To configure an IP route or identify the default gateway using Device Manager, follow these steps:

-
- Step 1** Choose **IP > Routes**.
You see the IP Routes window.
 - Step 2** Create a new IP route or identify the default gateway on a switch by clicking **Create**.

Send documentation comments to mdsfeedback-doc@cisco.com

You see the Create IP Routes window.

- Step 3** Complete the fields in this window. Configure a static route, by entering the destination network ID and subnet mask in the Dest and Mask fields. Configure a default gateway by entering the IP address of the seed switch in the Gateway field.
- Step 4** Click **Create** to add the IP route.

IPv4 Default Network Configuration

If you assign the IPv4 default network address, the switch considers routes to that network as the last resort. If the IPv4 default network address is not available, the switch uses the IPv4 default gateway address. For every network configured with the IPv4 default network address, the switch flags that route as a candidate default route, if the route is available.




Tip

If you configure the static route IP forwarding and the default network details, these IPv4 addresses will be used regardless of the default gateway being enabled or disabled. If these IPv4 addresses are configured and not available, the switch will fall back to using the default gateway IPv4 address, if you have configured it. Be sure to configure IPv4 addresses for all entries in the switch if you are using IPv4.

See the “[Initial Setup Routine](#)” section on page 2-2 for more information on configuring the IP addresses for all entries in the switch.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IPv4 address of the gateway switch (see [Figure 51-2](#)).

Figure 51-2 Overlay VSAN Functionality



In [Figure 51-2](#), switch A has the IPv4 address 1.12.11.1, switch B has the IPv4 address 1.12.11.2, switch C has the IPv4 address 1.12.11.3, and switch D has the IPv4 address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IPv4 address 1.1.1.10 to connect to the gateway

Send documentation comments to mdsfeedback-doc@cisco.com

switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IPv4 address (1.12.11.1) in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet world, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface (see the “[VSAN Interfaces](#)” section on page 20-24).

IPFC

IPFC provides IP forwarding or in-band switch management over a Fibre Channel interface (rather than out-of-band using the Gigabit Ethernet mgmt 0 interface). You can use IPFC to specify that IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.

Once the VSAN interface is created, you can specify the IP address for that VSAN. You can assign an IPv4 address or an IPv6 address.

**Note**

See the [Chapter 54, “Configuring IPv6 for Gigabit Ethernet Interfaces”](#) for information about configuring IPv6 on the Cisco MDS 9000 Family switches.

IPFC Configuration Guidelines

Follow these guidelines to configure IPFC:

1. Create the VSAN to use for in-band management, if necessary.
2. Configure an IPv4 address and subnet mask for the VSAN interface.
3. Enable IPv4 routing.
4. Verify connectivity.

IPv4 Static Routes

If your network configuration does not need an external router, you can configure IPv4 static routing on your MDS switch.

**Note**

For information about IPv6 static routing, see the “[Configuring IPv6 for Gigabit Ethernet Interfaces](#)” section on page 54-1.

Static routing is a mechanism to configure IPv4 routes on the switch. You can configure more than one static route.

Send documentation comments to mdsfeedback-doc@cisco.com

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IPv4 routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

Overlay VSANs

This section describes overlay VSANs and how to configure them.

This section includes the following topics:

- [About Overlay VSANs, page 51-6](#)
- [Configuring Overlay VSANs, page 51-6](#)

About Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.


Configuring Overlay VSANs

To configure an overlay VSAN, follow these steps:

-
- Step 1** Add the VSAN to the VSAN database on all switch in the fabric.
 - Step 2** Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side.
 - Step 3** Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.
 - Step 4** Configure the default gateway (route) and the IPv4 address on switches that point to the NMS (see [Figure 51-3](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 51-3 Overlay VSAN Configuration Example



Multiple VSAN Configuration

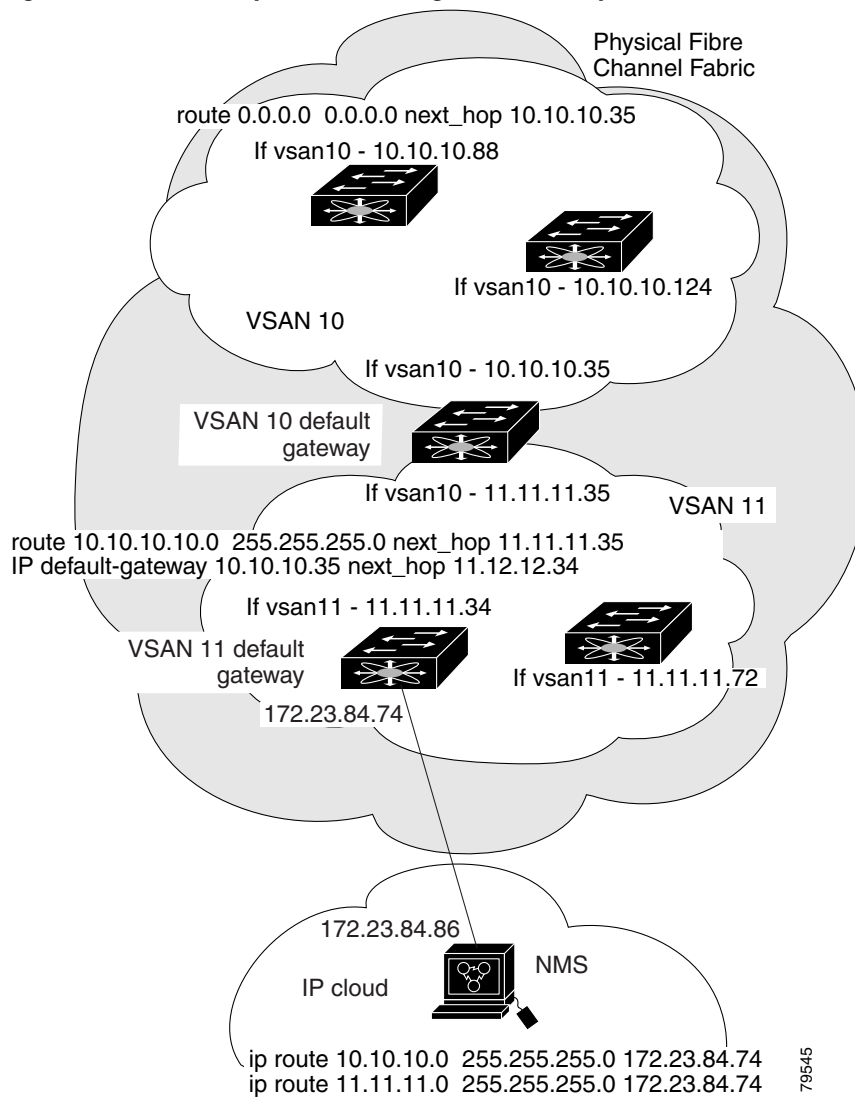
More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure multiple VSANs, follow these steps:

- Step 1** Add the VSAN to the VSAN database on any switch in the fabric.
- Step 2** Create a VSAN interface for the appropriate VSAN on any switch in the fabric.
- Step 3** Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.
- Step 4** Define the multiple static routes on the Fibre Channel switches and the IP cloud (see [Figure 51-4](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 51-4 Multiple VSAN Configuration Example



Virtual Router Redundancy Protocol

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. This section provides details on the VRRP feature.

This section includes the following topics:

- [About VRRP, page 51-9](#)
- [Configuring VRRP, page 51-10](#)

Send documentation comments to mdsfeedback-doc@cisco.com

About VRRP

VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:


- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and the draft-ietf-vrrp-ipv6 specification.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- Interface Mgmt 0 supports only one VRRP group. All other interface supports up to 7 virtual router groups, including both IPv4 and IPv6 combined.
- VR IDs can be reused in multiple VSANs with different virtual router IP mapping.
- Both IPv4 and IPv6 is supported.
- The management interface (mgmt 0) supports only one virtual router group. All other interfaces each support up to seven virtual router groups, including both IPv4 and IPv6 combined. Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.



Note If you are using IPv6, you must either configure an IPv6 address on the interface or enable IPv6 on the interface. For more information about IPv6, see [Chapter 54, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

In [Figure 51-5](#), switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have an IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches do not have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.


Figure 51-5 VRRP Functionality



Send documentation comments to mdsfeedback-doc@cisco.com

In [Figure 51-6](#), the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 51-6 Redundant Gateway



Configuring VRRP

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

- [Adding and Deleting Virtual Router, page 51-10](#)
- [Virtual Router Initiation, page 51-11](#)
- [Adding Virtual Router IP Addresses, page 51-11](#)
- [Priority for the Virtual Router, page 51-11](#)
- [Time Interval for Advertisement Packets, page 51-11](#)
- [Priority Preemption, page 51-11](#)
- [Virtual Router Authentication, page 51-12](#)
- [Priority Based on Interface State Tracking, page 51-12](#)

Adding and Deleting Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.



Note

The total number of VRRP groups that you can configure on a Gigabit Ethernet port, including main interfaces and subinterfaces, cannot exceed seven. This limitation applies to both IPv4 and IPv6 groups.

Send documentation comments to mdsfeedback-doc@cisco.com

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address, either IPv4 or IPv6, before attempting to enable a VR.

Adding Virtual Router IP Addresses

One virtual router IP address can be configured for a virtual router. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address. You can configure either an IPv4 address or an IPv6 address.

According to the VRRP specification, the master VRRP router drops the packets addressed to the virtual router's IP address because the virtual router is only intended as a next-hop router to forward packets. In MDS switches however, some applications require that packets addressed to virtual router's IP address be accepted and delivered to them. By using the **secondary** option to the virtual router IPv4 address, the VRRP router will accept these packets when it is the master.

To manage IP addresses for virtual routers from Device Manager, follow these steps:

-
- Step 1** Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.
 - Step 2** Click the **IP Addresses** tab on the VRRP dialog box.
 - Step 3** To create a new VRRP entry, click **Create**. You see the Create VRRP IP Addresses window.
 - Step 4** Complete the fields in this window to create a new VRRP IP Address, and click **OK** or **Apply**.
-

Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for switches with the primary IP address.

Time Interval for Advertisement Packets

The valid time range for an advertisement packet on an interface using IPv4 is between 1 and 255 seconds. The default value is 1 (one) second. If the switch has the primary IP address, this time must be specified.

Priority Preemption

You can enable a higher priority backup virtual router to preempt the lower priority master virtual router.



Note

If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.



Note

The VRRP preemption is not supported on IP storage Gigabit Ethernet interfaces.

Send documentation comments to mdsfeedback-doc@cisco.com

Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.



Note All VRRP configurations must be duplicated.



Note VRRP router authentication does not apply to IPv6.

Priority Based on Interface State Tracking

Interface state tracking changes the priority of the virtual router based on the state of another interface in the switch. When the tracked interface is down, the priority reverts to the priority value for the virtual router (see the [“Priority for the Virtual Router”](#) section on page 51-11). When the tracked interface is up, the priority of the virtual router is restored to the interface state tracking value. You can track the state of either a specified VSAN interface or the management interface (mgmt 0). The interface state tracking feature is disabled by default.



Note For interface state tracking to function, you must enable preemption on the interface. See the [“Priority Preemption”](#) section on page 51-11.

DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts because of one of the following reasons:

- The IP address or the switch name is wrongly configured.
- The DNS server is not reachable because external reasons (reasons beyond our control).



Note When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

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

Default Settings

Table 51-1 lists the default settings for DNS features.

Table 51-1 **Default DNS Settings**

Parameters	Default
Domain lookup	Disabled.
Domain name	Disabled.
Domains	None.
Domain server	None.
Maximum domain servers	6.

Table 51-2 lists the default settings for VRRP features.

Table 51-2 **Default VRRP Settings**

Parameters	Default
Virtual router state	Disabled.
Maximum groups per VSAN	255.
Maximum groups per Gigabit Ethernet port	7.
Priority preemption	Disabled.
Virtual router priority	100 for switch with secondary IP addresses. 255 for switches with the primary IP address.
Priority interface state tracking	Disabled.
Advertisement interval	1 second for IPv4. 100 centiseconds for IPv6.

Send documentation comments to mdsfeedback-doc@cisco.com