CHAPTER 22

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 (NMS):

- 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.

This chapter contains the following topics:

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running IP protocol over a 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.

Figure 22-1 Management Access to Switches

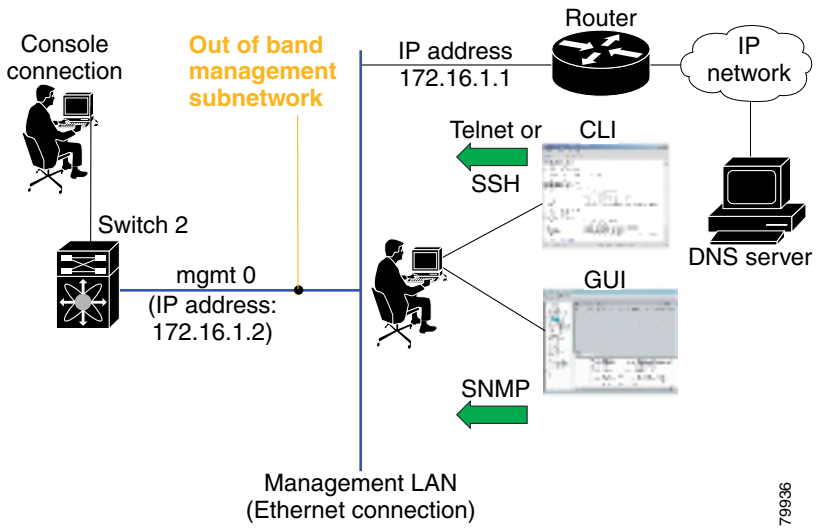

Configuring the Ethernet Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP network management sessions. You can also configure the supervisor module Ethernet interface and VSAN interfaces as management ports. This section focuses on the Ethernet management port (mgmt0). You can remotely configure the switch through the management port. To configure a connection remotely, you must configure the IP parameters (IP address and subnet mask) from the CLI so that the switch is reachable.

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

Configuring the Default Gateway

The IP address for a switch's default gateway should be configured along with the IP static routing commands (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 IP addresses will be used regardless of the default-gateway being enabled or disabled. If these IP address are configured and 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.

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 IP address of the gateway switch.

Figure 22-2 Overlay VSAN Functionality

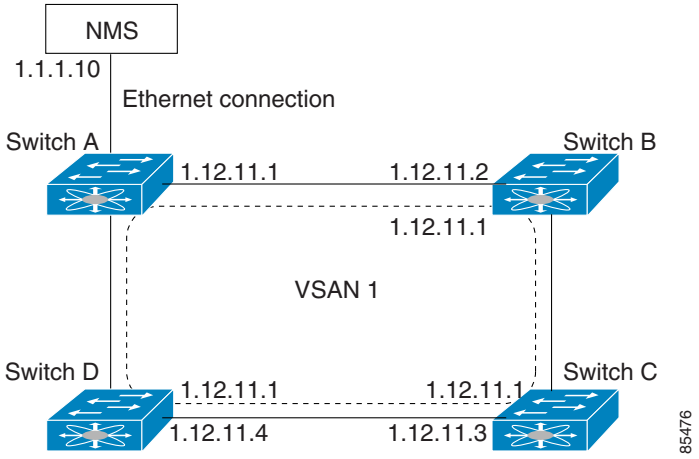

In Figure 22-2, Switch A has the IP address 1.12.11.1, Switch B has the IP address 1.12.11.2, Switch C has the IP address 1.12.11.3, and Switch D has the IP address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IP 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.

Configuring the Default Network

When IP routing is enabled on the switch, assign the IP default network address. The switch considers routes to that network as the last resort. If the IP default network address is not available, the switch uses the IP default gateway address. For every network configured with the IP 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 IP addresses will be used regardless of the default gateway being enabled or disabled. If these IP address are configured and 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.

Configuring an IP Route

To configure an IP route or identify the default gateway, perform the following steps.

Step 1 From the Device Manager, choose **IP > Routes**.

You see the IP Routes window.

Step 2 To create a new IP route or identify the default gateway on a switch, click **Create**.

You see the Create IP Routes window.

Step 3 Complete the fields in this window and click **OK** to add an IP route.

Step 4 To configure a static route, enter the destination network ID and subnet mask in the Dest and Mask fields. To configure a default gateway, enter the IP address of the seed switch in the Gateway field.

IP Access Control Lists

IP Access control lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related MDS out-of-band management traffic and in-band traffic based on IP addresses (Layer 3 and Layer 4 information).

You can use IP-ACLs to control transmissions on an interface.

IP-ACL Configuration Guidelines

Follow these guidelines when configuring IP-ACLs in any switch or director in the Cisco MDS 9000 Family:

- IP-ACLs cannot be configured for Gigabit Ethernet interfaces (IPS modules) or for Fibre Channel interfaces.
- IP-ACLs can only be configured on the management interface and VSAN interfaces.

- An IP-ACL is a sequential collection of permit and deny conditions that apply to IP flows. Each IP packet is tested against the conditions in the list. The first match determines if the software accepts or rejects the rule. Because the software stops testing conditions after the first match, the order of the conditions in the list is critical. If no conditions match, the software rejects that rule.
- An IP protocol can be configured using an integer ranging from 0 to 255 to represent a particular IP protocol. Alternatively, you can specify the name of a protocol: **icmp, ip, tcp**, or **udp.** IP includes Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), and other protocols.
- The source/source-wildcard and destination/destination-wildcard is specified in one of two ways:
  - Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Using the **any** option as an abbreviation for a source/source-wildcard or destination/destination-wildcard (0.0.0.0/255.255.255.255)

To configure an IP-ACL, you must complete the following tasks:

- Create an IP-ACL by specifying a name and access condition.

  All lists use the source and destination address for matching operations. You can configure finer granularity using optional keywords
- Apply the access list to specified interfaces.

Creating IP-ACLs

You can specify IP- ACLs using a assigned name. Each IP-ACL can have a maximum of 256 entries. Each entry is a unique filter applied to a specified interface. Each switch can have a maximum of 64 IP-ACLs.

Traffic coming into the switch is compared to IP-ACL entries based on the order that the entries occur in the switch. New statements are added to the end of the list. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the list, the traffic is denied. For this reason, you should have the frequently hit entries at the top of the list. There is an implied deny for traffic that is not permitted. A single-entry IP-ACL with only one **deny** entry has the effect of denying all traffic.

Adding Entries to an Existing IP-ACL

After you create an IP- ACL, you place subsequent additions at the end of the IP-ACL. You cannot insert entries in the middle of an IP-ACL. Each configured entry is automatically added to the end of a IP-ACL.

Comparing Ports

Use the following operators to compare the source and destination ports:

- eq = equal
- gt = greater than
- lt = less than
- range = range of ports

Port numbers range from 0 to 65535 for TCP and UDP ports. displays the port numbers for associated TCP and UDP ports.

Table 22-1 TCP and UDP Port Numbers

Protocol	Port	Number
TCP **Note** If the TCP connection is already established, use the **established** option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.	ftp	20
ftp-data	21	
ssh	22	
telnet	23	
smtp	25	
tasacs-ds	65	
www	80	
sftp	115	
http	143	
radius	1812	
wbem-http	5988	
wbem-https	5989	
UDP	dns	53
tftp	69	
ntp	123	
snmp	161	
snmp-trap	162	
syslog	514	

ICMP packets are filtered by the ICMP message type or the message code. Both values range from 0 to 255. displays the value for each associated ICMP type.

Table 22-2 ICMP Type Value

ICMP Type	Value
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

Applying IP-ACLs

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to the switch's interface.

Tip Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IP-ACL to E0 on Switch 1 instead of an outbound list to E1 on Switch 3.

Figure 22-3 Denying Traffic on the Inbound Interface

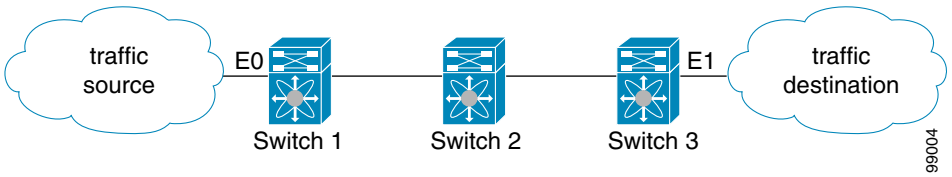

The access group controls access to an interface. Each interface can only be associated with one access list per direction. The ingress direction can have a different ACL than the egress direction. The access group becomes active on creation.

Tip We recommend creating all rules in an access list, before creating the access group that uses this access -list.

Caution If you create an access group before an access-list, all packets in that interface are dropped, because the access list is empty.

The terms in, out, source, and destination are used as referenced by the switch.

- In—Traffic that is arriving on the interface and which will go through the switch; the source would be where it's been and the destination is where it's going (on the other side of the router).

The access-group configuration for the ingress traffic applies to both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source would be where it's been (on the other side of the router) and the destination is where it's going.

The access-group configuration for the egress traffic applies only to local traffic.

Using the **log-deny** option at the end of the individual ACL entries shows the ACL number and whether the packet was permitted or denied, in addition to port-specific information. This option causes an information logging message about the packet that matches the dropped entry (or entries).

For the input ACL, the log displays the raw MAC information. The keyword "MAC=" does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not dumped to the log.

The following is an example of an input ACL log dump.

```
Jul 17 20:38:44 excal-2%KERN-7-SYSTEM_MSG:%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following is an example of an output ACL log dump.

```
Jul 17 20:38:44 excal-2%KERN-7-SYSTEM_MSG:%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2
DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00 TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277
SEQ=1280
```

Configuring IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN.

Configuring 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.

To configure an overlay VSAN, follow these steps:

To configure the management interface displayed in Figure 22-4, set the default gateway to an IP address on the Ethernet network.

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 default gateway (route) and the IP address on switches that point to the NMS.

Figure 22-4 Overlay VSAN Configuration Example

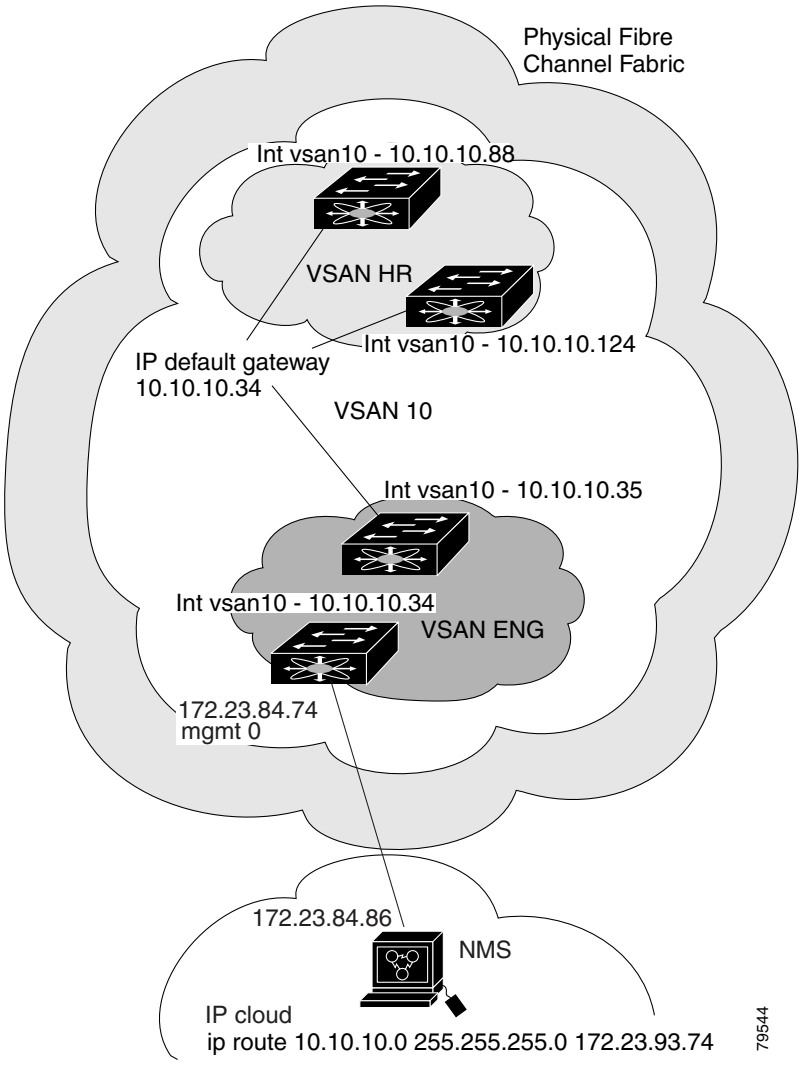

Configuring Multiple VSANs

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 an overlay VSAN, 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 route on the Fibre Channel switches and the IP cloud.

Figure 22-5 Multiple VSANs Configuration Example

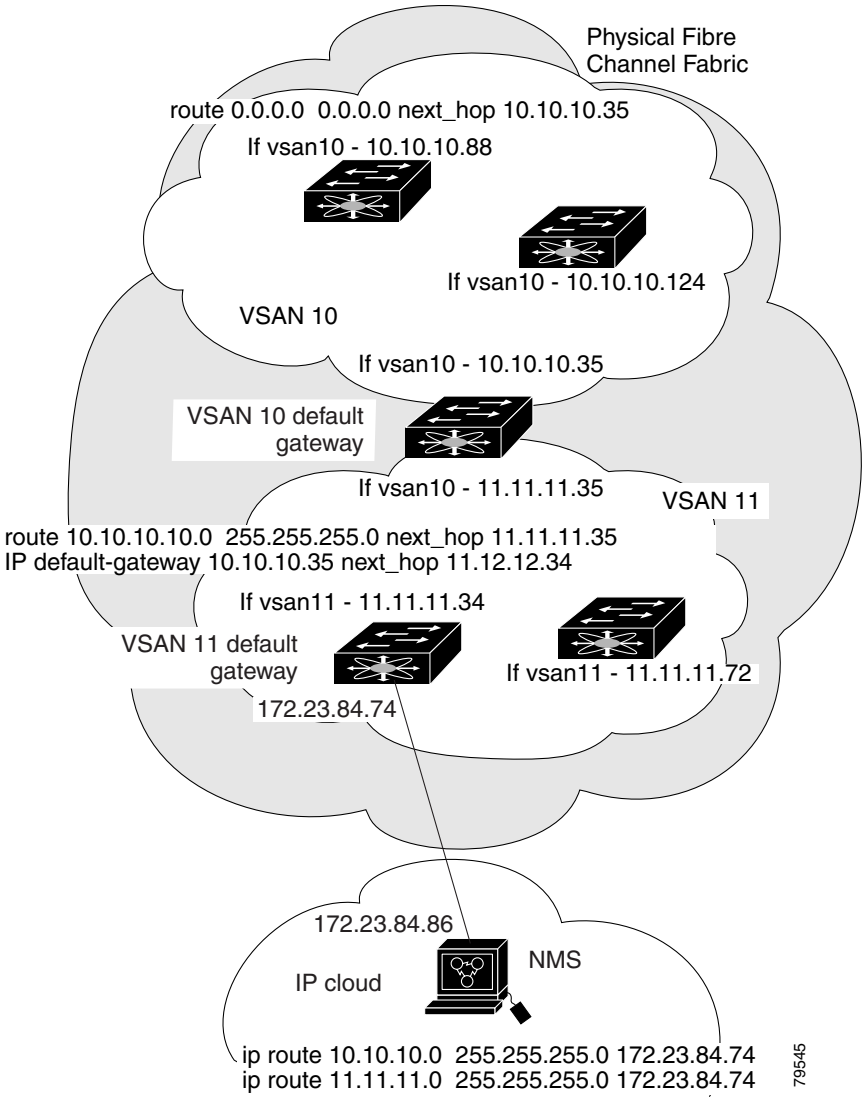

Managing IPFC Connectivity with Multiple VSANs

To configure IPFC from the Device Manager, choose VSAN from the FC menu and click the General tab.

Configuring VRRP

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.

VRRP Features

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.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with a different virtual router IP mapping.
- 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.

VRRP Functionality

In Figure 22-6, switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches don't have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 22-6 VRRP Functionality

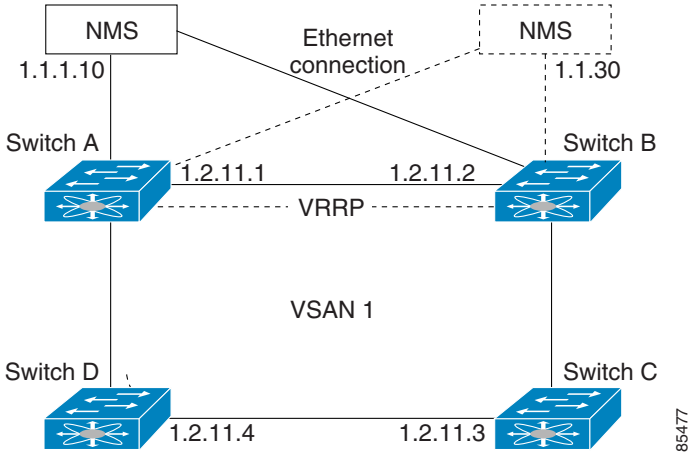

In Figure 22-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 22-7 Redundant Gateway

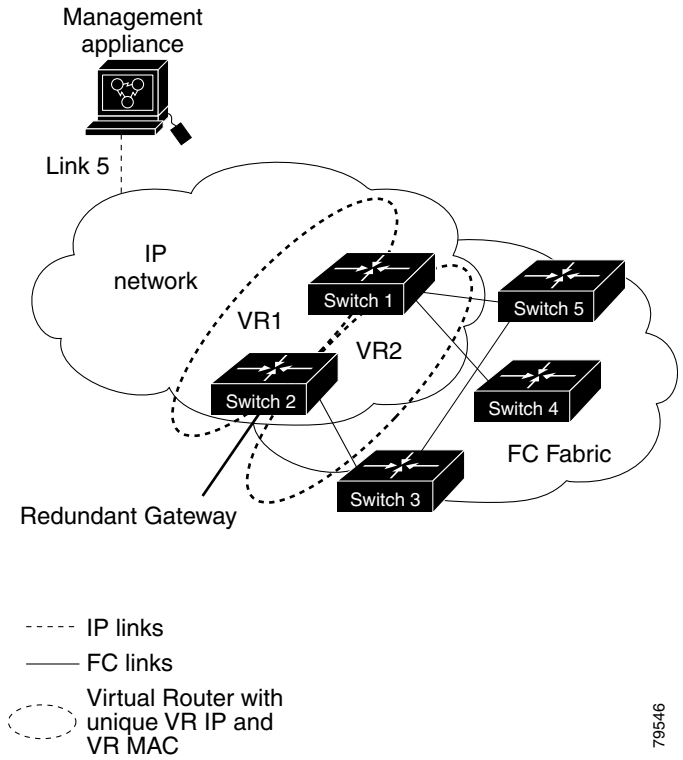

Creating or Removing a Virtual Router

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

Enabling a Virtual Router

By default, a virtual router is always disabled (**shutdown**). VRRP can be configured only if this state is disabled. Be sure to configure at least one IP address before attempting to enable a VR.

Adding an IP Address for a Virtual Router

One primary IP address and multiple secondary addresses can be configured for a switch. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address.

Viewing IP Address Information

To view IP addresses of the switches in the current fabric from the Fabric Manager, choose Switches from the menu tree.

The Information pane displays IP address information for multiple switches.

Managing IP Addresses for VRRP

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**.

Setting 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 a switch with the primary IP address.

Setting the Time Interval for the Advertisement Packet

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

Preempting the Master Virtual Router

By default, the preempt option is enabled. An owner with priority 255 cannot be preempted. If two priorities match, the owner with the highest priority preempts the master virtual router.

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 Authentication for the Virtual Router

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.

All VRRP configurations must be duplicated

Setting the Priority Based on Interface State

The tracking feature is disabled by default. When you specify the tracking option, the priority of the virtual router is changed based on the state of another interface in the switch. When the tracked interface is down, the priority of the virtual router is changed to a lower priority value. When the tracked interface is up, the priority of the virtual router is restored to its original value. You can track one of two interfaces on a switch in the Cisco MDS 9000 Family: a specified VSAN interface or a management interface.

Configuring VRRP Operations Attributes

To configure VRRP operations attributes from Device Manager, follow these steps:

Step 1 Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.

Step 2 Configure operations attributes for the virtual router.

Step 3 To create a new VRRP entry, click **Create**. You see the Create VRRP Entry window.

Step 4 Complete the fields in this window to create a new VRRP entry, and click **OK** or **Apply**.

Default Settings

Table 22-3 lists the default settings for IP features.

Table 22-3 Default IPFC Settings

Parameters	Default
VSAN IP interface configuration	No IP address is assigned by default.
IP routing	Disabled.
Domain lookup	Disabled.
Domain name	Enabled.
Domain list	No domains are configured.
Name server	No servers are configured.
Virtual router	Disabled (shutdown).
Virtual router priority for switches with secondary IP address	100.
Virtual router priority for switches with primary IP address	255.
Time interval between advertisement frames	1 second.
Preempting master VR	Enabled.
VRRP security authentication	No authentication.
Interface state tracking	Disabled.

Note ICMP redirect packets are always rejected.

Enabling or Disabling IP Forwarding

To view or change the IP forwarding configuration of the switches in the current fabric, perform the following steps.

Step 1 Choose **IP > Forwarding** from the Fabric Manager menu tree.

Step 2 To enable IP forwarding for a specific switch, check the **RoutingEnabled** check box.

Viewing Information and Statistics

You can monitor a variety of information and statistics about IP services from Fabric Manager and Device Manager.

This section includes the following topics:

Viewing VRRP Statistics

To monitor VRRP statistics, click the **Statistics** tab on the VRRP dialog box. The VRRP dialog box with the Statistics tab selected is displayed.

Viewing TCP Information and Statistics

To view TCP information from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu.

To monitor TCP statistics from the Fabric Manager, choose **IP > Mgmt Statistics** from the menu tree and click the **TCP** tab. To monitor TCP statistics from the Device Manager, choose **Statistics** from the IP menu and view the **TCP** tab.

Viewing UDP Information and Statistics

To view User Datagram Protocol (UDP) information, from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu and click the **UDP** tab.

To monitor UDP traffic from the Fabric Manager, choose **IP > Mgmt Statistics** from the menu tree and click the **UDP** tab. From Device Manager, choose **Statistics** from the IP menu and click the **UDP** tab.

The Fabric Manager Information pane displays TCP traffic information for multiple switches. The Device Manager dialog box displays information for a single switch.

Viewing IP Statistics

To monitor IP statistics from the Fabric Manager, choose **IP >Mgmt Statistics** from the menu tree and click the **IP** tab. From Device Manager, select **Statistics** from the IP menu and click the **IP** tab.

The Fabric Manager Information pane displays IP statistics for multiple switches. The Device Manager dialog box displays information for a single switch.

Viewing ICMP Statistics

To monitor statistics for ICMP packets received, select **IP > Mgmt Statistics** from the menu tree and click the **ICMP In** tab. To monitor statistics for ICMP packets transmitted from the Fabric Manager, select **IP > Mgmt Statistics** from the menu tree and click the **ICMP Out** tab.

To monitor ICMP statistics from Device Manager, select **Statistics** from the IP menu and click the **ICMP** tab.

The Fabric Manager Information pane displays information for multiple switches. The Device Manager dialog box displays information for a single switch.

In the Device Manager, a prefix (In or Out) identifies whether the packets are received or transmitted. In the Fabric Manager, separate tabs on the Information pane are provided for incoming and outbound ICMP traffic and this prefix is omitted.