CHAPTER 36

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 includes the following sections:

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 (see Figure 36-1).

Figure 36-1 Management Access to Switches

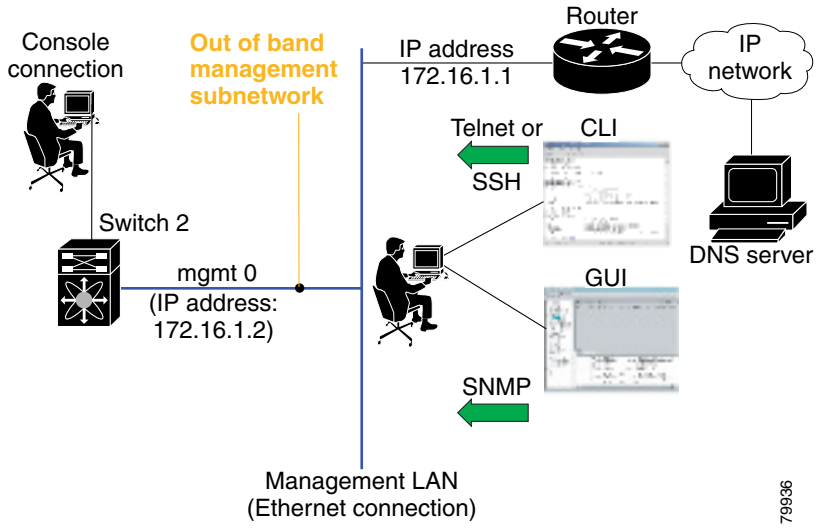

Management Interface Configuration

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.

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 some IP parameters (IP address, subnet mask) so that the switch is reachable. You can manually configure the management interface from the CLI.

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 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 is or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

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, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	`switch(config-if)# ip address 1.1.1.1 255.255.255.0`	Enters the IP address (1.1.1.1) and IP subnet mask (255.255.255.0) for the management interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Default Gateway Configuration

The default gateway IP address 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 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 4-2 for more information on configuring the IP addresses for all entries in the switch.

Use the **IP default-gateway** command to configure the IP address for a switch's default gateway and the **show ip route** command to verify that the IP address for the default gateway is configured.

To configure default gateways, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip default-gateway 1.12.11.1`	Configures the IP address for the default gateway.

Default Network Configuration

If you 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.

See the "Initial Setup Routine" section on page 4-2 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 IP address of the gateway switch (see Figure 36-2).

Figure 36-2 Overlay VSAN Functionality

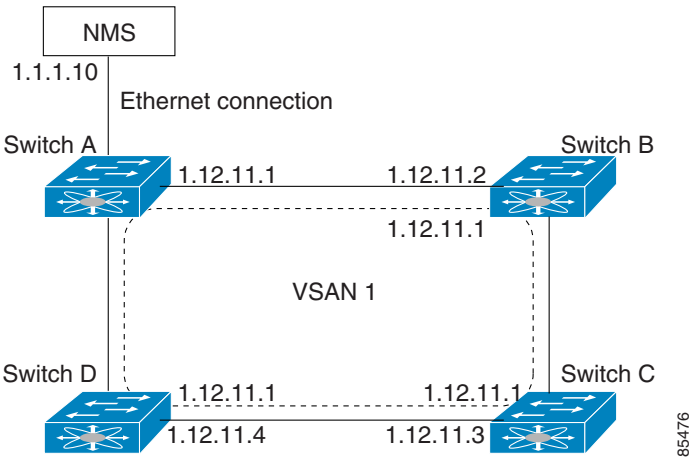

In Figure 36-2, switch A has the IP address 1.12.11.1, switch B has the IP address 1.12.11.2, switch C has the IP address 1.12.11.3, and switch D has the IP address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. 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 (see the "Configuring VSAN Interfaces" section on page 11-22).

Instead of the **ip default-gateway** command, use the **ip default-network** command when IP routing is enabled on the switch. Use the **show ip route** command to verify if the IP address for the default gateway is configured.

To configure default networks, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip default-network 190.10.1.0`	Configures the IP address for the default network (190.10.1.0).
	`switch(config)# ip route 10.0.0.0 255.0.0.0 131.108.3.4` `switch(config)# ip default-network 10.0.0.0`	Defines a static route to network 10.0.0.0 as the static default route.

IPFC Configuration

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

Configuring an IP Address in a VSAN

To configure a VSAN interface and an IP address for that interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures the interface for the specified VSAN (1).
Step 3	`switch(config-if)# ip address 10.0.0.12 255.255.255.0` `switch(config-if)#`	Configures the IP address and netmask for the selected interface.

Enabling IP Routing

By default, the IP routing feature is disabled in all switches.

To enable the IP routing feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip routing` `switch(config)#`	Enables IP routing (disabled by default).
Step 3	`switch(config)# no ip routing` `switch(config)#`	Disables IP routing and reverts to the factory settings.

IPFC Configuration Example

This section describe an example configuration for IPFC. Figure 36-3 shows an example network.

The example network has the following links:

- Switch_1 is connected to the main network by the mgmt 0 interface and to the fabric by an ISL.
- Switch_2 and Switch_3 are connected to the fabric by an ISL but are not connected to the main network.

Figure 36-3 IPFC Example Network

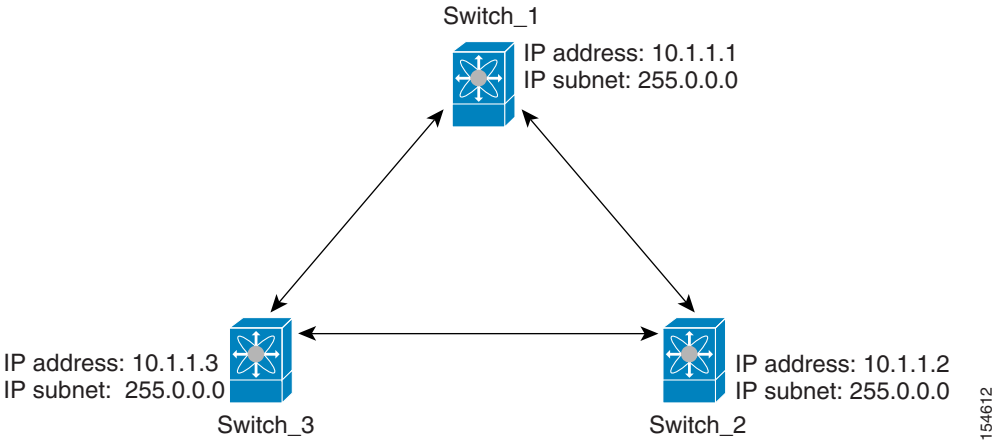

The following steps show how to configure Switch_1 in the example network in Figure 36-3:

Step 1 Create the VSAN interface and enter interface configuration submode.

```
switch_1# config t
switch_1(config)# interface vsan 1
switch_1(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_1(config-if)# ip address 10.1.1.1 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_1(config-if)# no shutdown
switch_1(config-if)# exit
switch_1(config)#
```

Step 4 Enable IPv4 routing.

```
switch_1(config)# ip routing
switch_1(config)# exit
switch_1#
```

Step 5 Display the routes.

```
switch_1# show ip route

Codes: C - connected, S - static

C 172.16.1.0/23 is directly connect, mgmt0
C 10.0.0.0./8 is directly connected, vsan1
```

The following steps show how to configure Switch_2 in the example network in Figure 36-3.

Step 1 Disable the mgmt 0 interface.

Note Configure this switch using the console connection.

```
switch_2# config t
switch_2(config)# interface mgmt 0
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 2 Create the VSAN interface and enter interface configuration submode.

```
switch_2# config t
switch_2(config)# interface vsan 1
switch_2(config-if)#
```

Step 3 Configure the IP address and subnet mask.

```
switch_2(config-if)# ip address 10.1.1.2 255.0.0.0
```

Step 4 Enable the VSAN interface and exit interface configuration submode.

```
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 5 Enable IPv4 routing.

```
switch_2(config)# ip routing
switch_2(config)# exit
switch_2#
```

Step 6 Display the routes.

```
switch_2# show ip route

Codes: C - connected, S - static

C 10.0.0.0./8 is directly connected, vsan1
```

Step 7 Verify the connectivity to Switch_1.

```
switch_2# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.618 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.528 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.567 ms

--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 4998 ms
rtt min/avg/max/mdev = 0.528/0.570/0.618/0.057 ms
```

The following steps show how to configure Switch_3 in the example network in Figure 36-3.

Step 1 Disable the mgmt 0 interface.

Note Configure this switch using the console connection.

```
switch_3# config t
switch_3(config)# interface mgmt 0
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#

switch_3# config t
switch_3(config)# interface vsan 1
switch_3(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_3(config-if)# ip address 10.1.1.3 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#
```

Step 4 Enable IPv4 routing.

```
switch_3(config)# ip routing
switch_3(config)# exit
switch_3#
```

Step 5 Display the routes.

```
switch_3# show ip route

Codes: C - connected, S - static

C 10.0.0.0./8 is directly connected, vsan1
```

Step 6 Verify the connectivity to Switch_1.

```
switch_3# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=1.19 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.510 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.653 ms

--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2008 ms
rtt min/avg/max/mdev = 0.510/0.787/1.199/0.297 ms
```

Configuring IP Static Routes

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

If your configuration does not need an external router, you can use static routing.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IP 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.

To configure a static route, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

Step 2	`switch(config)# IP route <network IP address> <netmask> <next hop IP address> distance <number> interface <vsan number>`
	
	`For example:`
	`switch(config)# IP route 10.0.0.0 255.0.0.0 20.20.20.10 distance 10 interface vsan 1`
	`switch(config)#`
		Configures the static route for the specified IP address, subnet mask, next hop, and distance, and VSAN or management interface.

Displaying and Clearing ARPs

Address Resolution Protocol (ARP) entries in Cisco MDS 9000 Family switches can be displayed, deleted, or cleared. The ARP feature is enabled on all switches.

- Use the **show arp** command to display the ARP table.

```
switch# show arp
Protocol Address           Age (min)  Hardware Addr  Type  Interface
Internet  171.1.1.1               0  0006.5bec.699c  ARPA  mgmt0
Internet  172.2.0.1               4  0000.0c07.ac01  ARPA  mgmt0
```

- Use the **no arp** command in configuration mode to remove an ARP entry from the ARP table.

```
switch(config)# no arp 172.2.0.1
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
```

Displaying IP Interface Information

Use the following **show** commands to view configured IP interface information (see Examples 36-1 to 36-4).

Example 36-1 Displays the VSAN Interface

```
switch# show interface vsan1
vsan1 is up, line protocol is up
    WWPN is 10:00:00:05:30:00:59:1f, FCID is 0x9c0100
    Internet address is 10.1.1.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
```

Note You can see the output for this command only if you have previously configured a virtual network interface (see the "Configuring an IP Address in a VSAN" section on page 36-5).

Example 36-2 Displays the Connected and Static Route Details

```
switch# show ip route

Codes: C - connected, S - static

Default gateway is 172.22.95.1

C 172.22.95.0/24 is directly connected, mgmt0
C 10.1.1.0/24 is directly connected, vsan1
```

Example 36-3 Displays Configured Routes

```
switch# show ip route configured
Destination          Gateway             Mask Metric        Interface

         default      172.22.95.1          0.0.0.0      0            mgmt0
        10.1.1.0          0.0.0.0    255.255.255.0      0            vsan1
     172.22.95.0          0.0.0.0    255.255.255.0      0            mgmt0
```

Example 36-4 Displays the IP Routing Status

```
switch# show ip routing
ip routing is disabled
```

Overlay VSAN Configuration

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:

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 (see Figure 36-4).

Figure 36-4 Overlay VSAN Configuration Example

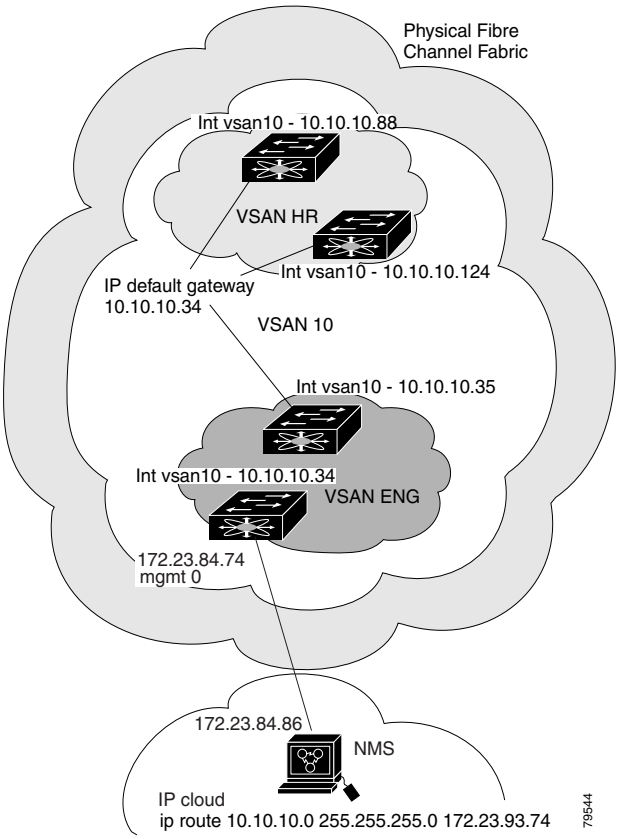

Note To configure the management interface displayed in Figure 36-4, set the default gateway to an IP address on the Ethernet network.

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in Figure 36-4), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch-config-vsan-db#`	Configures the VSAN database.
Step 3	`switch--config-vsan-db# vsan 10 name MGMT_VSAN`	Defines the VSAN in the VSAN database on all of the switches in the Fibre Channel fabric.
Step 4	`switch--config-vsan-db# exit` `switch(config)#`	Exits the VSAN database mode.
Step 5	`switch(config)# interface vsan10` `switch(config-if)#`	Creates a VSAN interface (VSAN 10).
Step 6	`switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0`	Assigns an IP address and netmask for this switch.
Step 7	`switch(config-if)# no shut`	Enables the configured interface.
Step 8	`switch--config-if# end` `switch#`	Exits to EXEC mode.
Step 9	`switch# exit`	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.

To configure the NMS station displayed in Figure 36-4, follow this step:

	Command	Purpose
Step 1	`nms# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.

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 route on the Fibre Channel switches and the IP cloud (see Figure 36-5).

Figure 36-5 Multiple VSANs Configuration Example

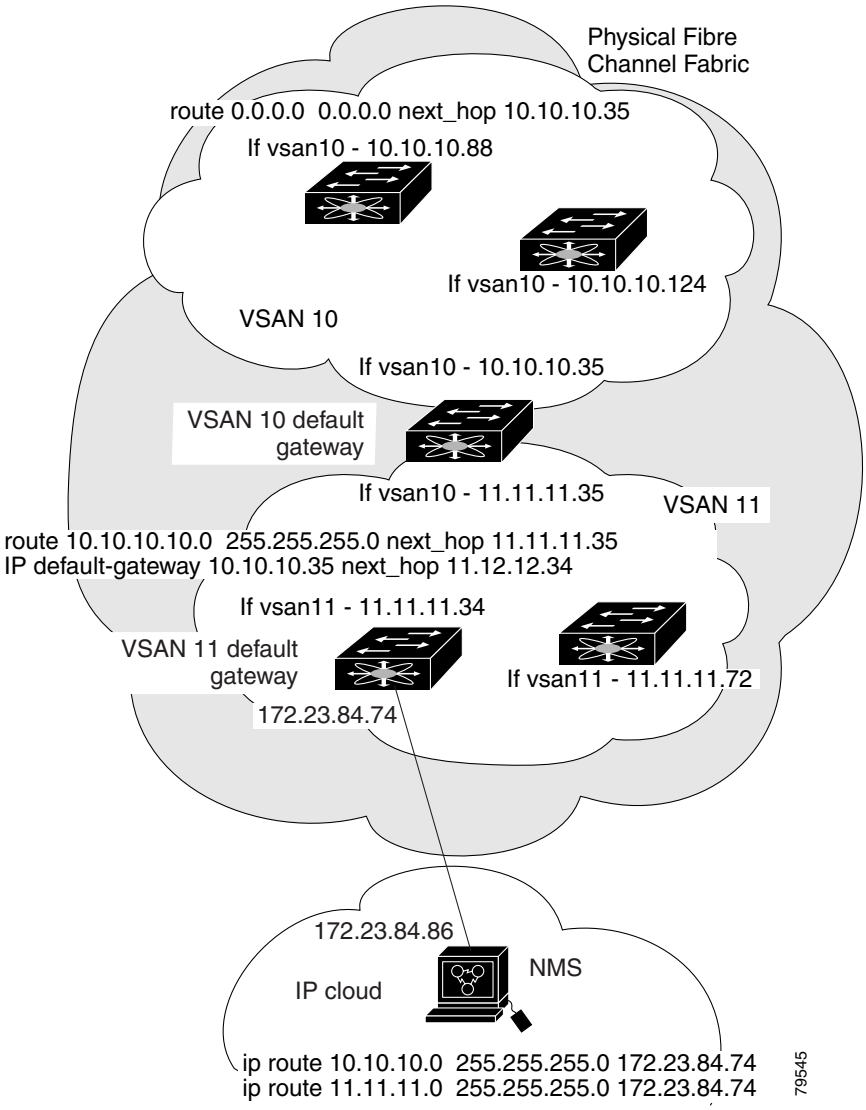

To configure an overlay VSAN (using the example in Figure 36-5), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch-config-vsan-db#`	Configures the VSAN database.
Step 3	`switch-config-vsan-db# vsan 10 name MGMT_VSAN_10` `switch-config-vsan-db#`	Defines the VSAN in the VSAN database on all of the switches in VSAN 10.
Step 4	`switch-config-vsan-db# exit` `switch(config)#`	Exits the database 10 mode.
Step 5	`switch-config-vsan-db# vsan 11 name MGMT_VSAN_11` `switch-config-vsan-db#`	Defines the VSAN in the VSAN database on all of the switches in VSAN 11.
Step 6	`switch-config-vsan-db# exit` `switch(config)#`	Exits the VSAN database 11 mode.
Step 7	`switch(config)# interface vsan10` `switch(config-if)#`	Enters the VSAN 10 interface configuration mode for VSAN 10.
Step 8	`switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0` `switch(config-if)#`	Assigns an IP address and netmask for this switch.
Step 9	`switch(config-if)# no shut`	Enables the configured interface for VSAN 10.
Step 10	`switch--config-if# exit` `switch(config)#`	Exits the VSAN 10 interface mode.
Step 11	`switch(config)# interface vsan11` `switch(config-if)#`	Enters the VSAN 11 interface configuration mode.
Step 12	`switch(config-if)# ip address 11.11.11.0 netmask 255.255.255.0` `switch(config-if)#`	Assigns an IP address and netmask for this switch
Step 13	`switch(config-if)# no shut`	Enables the configured interface for VSAN 11.
Step 14	`switch-config-if# end` `switch#`	Exits to EXEC mode.
Step 15	`switch# exit`	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.
Step 16	`NMS# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the IP cloud.
Step 17	`NMS# route ADD 11.11.11.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route for VSAN 11 on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.
Step 18	`switch# route 10.10.10.0 255.255.255.0 next_hop 11.11.11.35`	Defines the route to reach subnet 10 from subnet 11.

The 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.

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 36-6, 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 36-6 VRRP Functionality

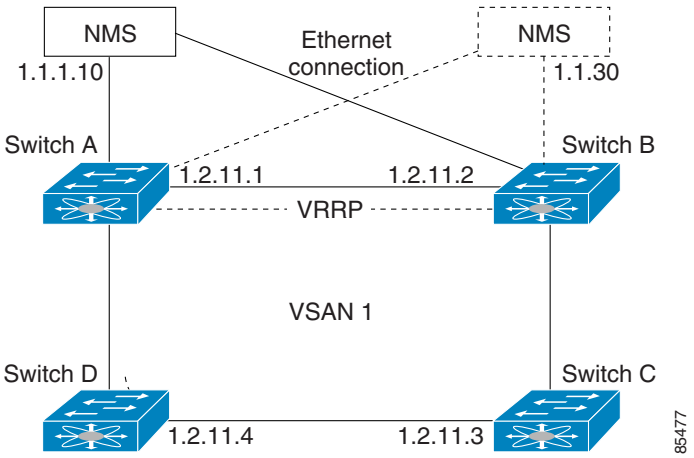

In Figure 36-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 36-7 Redundant Gateway

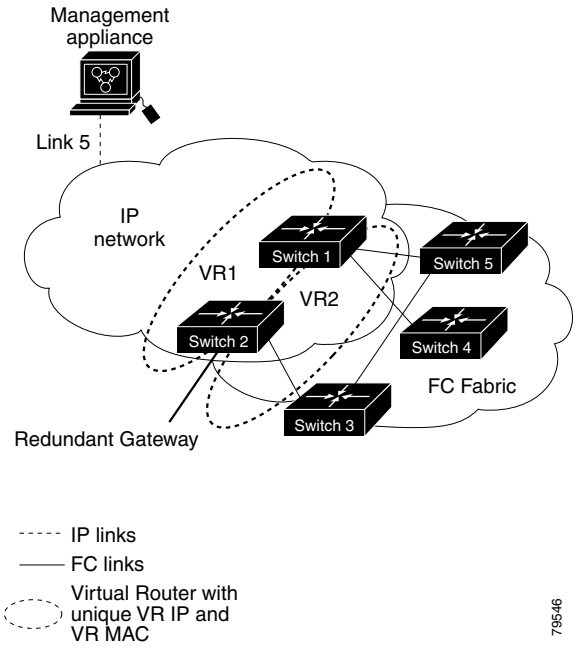

Virtual Router Addition and Deletion

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

To create or remove a VR, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)`	Creates a VR ID 250.
	`switch(config-if-vrrp)# no vrrp 250` `switch(config-if)`	Removes a VR ID 250.

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 before attempting to enable a VR.

To enable or disable a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch(config-if-vrrp)# no shutdown`	Enables VRRP configuration.
	`switch(config-if-vrrp)# shutdown`	Disables VRRP configuration.

Virtual Router IP Address Addition

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.

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 IP address, the VRRP router will accept these packets when it is the master.

To configure an IP address for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# interface ip address 10.0.0.12 255.255.255.0xi`	Configures an IP address. The IP address must be configured before the VRRP is added.
Step 4	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates VR ID 250.
Step 5	`switch(config-if-vrrp)# address 10.0.0.10`	Configures the IP address (10.0.0.10) for the selected VR. **Note** This IP address should be in the same subnet as the IP address of the interface.
	`switch(config-if-vrrp)# no address 10.0.0.10`	Removes the IP address (10.0.0.10) for the selected VR.

	Command	Purpose
Step 6	`switch(config-if-vrrp)# address 10.0.0.10 secondary`	Configures the IP address (10.0.0.10) as secondary for the selected VR. **Note** The **secondary** option should be used only with applications that require VRRP routers to accept the packets sent to the virtual router's IP address and deliver to them. For example, iSNS requires this option (see the "Enabling the iSNS Server" section on page 35-64).
	`switch(config-if-vrrp)# no address 10.0.0.10 secondary`	Removes the IP address (10.0.0.10) as secondary for the selected VR.

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.

To set the priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# priority 2` `switch(config-if-vrrp)#`	Configures the priority for the selected VRRP. **Note** Priority 255 cannot be preempted.

Time Interval for Advertisement Packets

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.

To set the priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# advertisement-interval 15`	Sets the interval time in seconds between sending advertisement frames.

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.

To enable or disable preempting, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# preempt`	Enables the higher priority backup virtual router to preempt the lower priority master virtual router. **Note** This preemption does not apply to the primary IP address.
	`switch(config-if-vrrp)# no preempt`	Disables (default) the preempt option and allows the master to keep its priority level.

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.

To set an authentication option for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# authentication text password`	Assigns the simple text authentication option and specifies the password for this option.
	`switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003`	Assigns MD5 authentication option and specifies the key and the unique SPI value for this option. The SPI and the valid range is 0x100 to 0xFFFFFFFF.
	`switch(config-if-vrrp)# no authentication`	Assigns the no authentication option, which is the default.

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.

Note For interface tracking to function, you must enable preemption on the interface. See the "Priority Preemption" section on page 36-20.

To track the interface priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# preempt`	Enables priority preemption.
Step 5	`switch(config-if-vrrp)# track interface mgmt 0 priority 2`	Specifies the priority of the virtual router to be modified based on the state of the management interface.
	`switch(config-if-vrrp)# no track`	Disables the tracking feature.

Displaying VRRP Information

Use the **show vrrp vr** command to display configured VRRP information (see Examples 36-5 to 36-8).

Example 36-5 Displays VRRP Configured Information

```
switch# show vrrp vr 7 interface vsan 2 configuration
vr id 7 configuration
admin state down
priority 100
no authentication
advertisement-Interval 1
preempt yes
tracking interface vsan1 priority 2
protocol IP
```

Example 36-6 Displays VRRP Status Information

```
switch# show vrrp vr 7 interface vsan 2 status
vr id 7 status
MAC address 00:00:5e:00:01:07
Operational state: init
```

Example 36-7 Displays VRRP Statistics

```
switch# show vrrp vr 7 interface vsan 2 statistics
vr id 7 statistics
Become master 0
Advertisement 0
Advertisement Interval Error 0
Authentication Failure 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Authentication Type 0
Mismatch Authentication 0
Invalid Packet Length 0
```

Example 36-8 Displays VRRP Cumulative Statistics

```
switch# show vrrp statistics
Invalid checksum 0
Invalid version 0
Invalid VR ID 0
```

Clearing VRRP Statistics

Use the **clear vrrp** command to clear all the software counters for the specified virtual router (see Example 36-9).

Example 36-9 Clears VRRP Information

```
switch# clear vrrp 7 interface vsan2
switch#
```

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 due to one of the following reasons:

- The IP address or the switch name is wrongly configured
- The DNS server is not reachable due to 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.

To configure a DNS server, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ip domain-lookup`	Enables the IP Domain Naming System (DNS)-based host name-to-address translation.
	`switch(config)# no ip domain-lookup`	Disables (default) the IP DNS-based host name-address translation and reverts to the factory default.
Step 3	`switch(config)# no ip domain-name cisco.com`	Disables the domain name and reverts to the factory default.
	`switch(config)# ip domain-name cisco.com`	Enables (default) the default domain name feature used to complete unqualified host names. Any IP host name that does not contain a domain name (that is, any name without a dot), will have the dot and cisco.com appended to it before being added to the host table.

Step 4	`switch(config)# ip domain-list harvard.edu`
	`switch(config)# ip domain-list stanford.edu`
	`switch(config)# ip domain-list yale.edu`
		Defines a filter of default domain names to complete unqualified host names, use the **ip domain-list** global configuration command. You can define up to 10 domain names in this filter. To delete a name from a filter, use the **no** form of this command.

	`switch(config)# no ip domain-list`	Deletes the defined filter and reverts to factory default. No domains are configured by default.

Note If you have not configured a domain list, the domain name that you specified with the **ip domain-name** global configuration command is used. If you did configure a domain list, the default domain name is not used. The **ip domain-list** command is similar to the **ip domain-name** command, except that with the **ip domain-list** command you can define a list of domains, each to be tried in turn.

Step 5	`switch(config)# ip name-server 15.1.0.1 15.2.0.0`	Specifies the first address (15.1.0.1) as the primary server and the second address (15.2.0.0) as the secondary sever. You can configure a maximum of six servers.
	`switch(config)# no ip name-server`	Deletes the configured server(s) and reverts to factory default. No server is configured by default.

Note Alternatively, you can configure the DNS entry using the switch names (instead of IP addresses). The configured switch name automatically looks up the corresponding IP address.

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see Example 36-10).

Example 36-10 Displays Configured Host Details

```
switch# show hosts
Default domain is cisco.com
Domain list: ucsc.edu harvard.edu yale.edu stanford.edu
Name/address lookup uses domain service
Name servers are 15.1.0.1 15.2.0.0
```

Default Settings

Table 36-1 lists the default settings for FSPF features.

Table 36-1 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.