CHAPTER 17

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 17-1).

Figure 17-1 Management Access to Switches

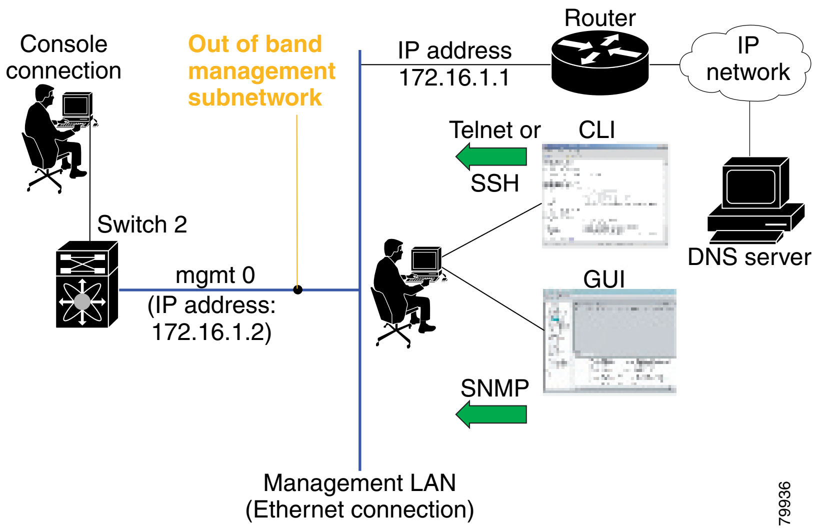

Configuring the Ethernet Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP network management sessions. You can also configure the supervisor module's 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.

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.

Configuring the Default Gateway

Use the **IP default-gateway** command to configure the IP address for a switch's default gateway. This 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 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 3-2).

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 17-2).

Figure 17-2 Overlay VSAN Functionality

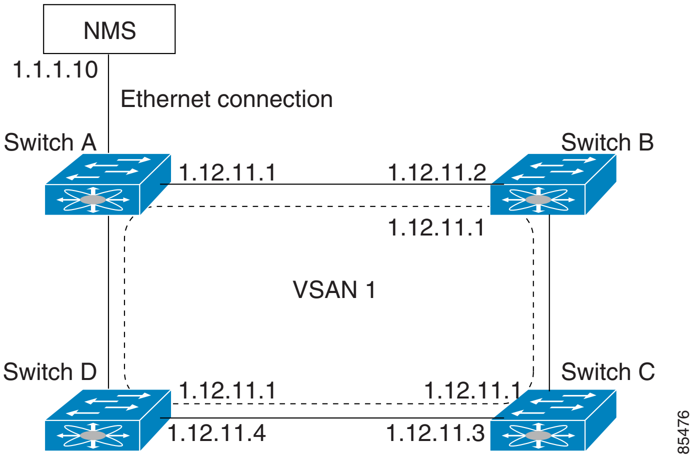

In Figure 17-2, switch A has the IP address 1.12.11.1, switch B has the IP address 1.12.11.2, switch C has the IP address 1.12.11.3, and switch D has the IP address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. 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 9-17).

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`**
	`switch(config)#`
		Configures the IP address for the default gateway (1.12.11.1).

Use the **show ip route** command to verify that the IP address for the default gateway is configured.

Configuring the Default Network

Unlike the **ip default-gateway** command, use the **ip default-network** command when IP routing is enabled on the switch. 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 3-2).

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`**
	`switch(config)#`
		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`**
	`switch(config)#`
		Defines a static route to network 10.0.0.0 as the static default route.

Use the **show ip route** command to verify if the IP address for the default gateway is configured.

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:
  1. 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
  2. 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.

To create an IP-ACL, follow these steps:

	Command	Purpose
Step 1	switch# **config t**	Enters configuration mode.
Step 2	switch(config)# **ip access-list List1 permit ip any any**	Configures an IP-ACL called List1 and permits IP traffic from any source address to any destination address.
	switch(config)# **no ip access-list List1 permit ip any any**	Removes the IP-ACL called List1.
Step 3	switch(config)# **ip access-list List1 deny tcp any any**	Updates List1 to deny TCP traffic from any source address to any destination address.

To define an IP-ACL that permits a specified network, follow these steps:

	Command	Purpose
Step 1	switch# **config t**	Enters configuration mode.
Step 2	switch(config)# **ip access-list List1 permit udp 192.168.32.0 0.0.7.255**	Defines an IP-ACL that permits this network. Subtracting 255.255.248.0 (normal mask) from 255.255.255.255 yields 0.0.7.255.

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.

To add entries to an existing IP-ACL, follow these steps:

	Command	Purpose
Step 1	switch# **config t**	Enters configuration mode.
Step 2	switch(config)# **ip access-list List1 permit tcp host 10.1.1.2 host 172.16.1.1 eq telnet** switch(config)# **ip access-list List1 permit tcp host 10.1.1.2 host 172.16.1.1** switch(config)# **ip access-list List1 permit udp host 10.1.1.2 host 172.16.1.1** switch(config)# **ip access-list List1 permit ip 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.25**5	Permits all IP traffic from and to the specified networks. **Note** In this example, the last entry is sufficient. You do not need the first three entries.

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

To use the operand and port options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any`	Denies TCP traffic from 1.2.3.0 through source port 5 to any destination.

Port numbers range from 0 to 65535 for TCP and UDP ports.Table 17-1 displays the port numbers for associated TCP and UDP ports.

Table 17-1 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. Table 17-2 displays the value for each associated ICMP type.

Table 17-2 ICMP Type Value

ICMP Type[1]	Value
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

[1] ICMP redirect packets are always rejected.

Removing Entries from an Existing IP-ACL

Use the **no permit** and **no deny** commands to remove entries from a configured IP-ACL.

To remove configured entries from an IP-ACL, follow these steps:

	Command	Purpose
Step 1	`switch#` **`config t`**	Enters configuration mode.
Step 2	`switch(config)#` **`no ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any`**	Removes this entry from the IP-ACL.
	`switch(config)#` **`no ip access-list x3 deny ip any any`**	Removes this entry from the IP-ACL.
	`switch(config)#` **`no ip access-list x3 permit ip any any`**	Removes this entry from the IP-ACL.

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 (see Figure 17-1).

Figure 17-3 Denying Traffic on the Inbound Interface

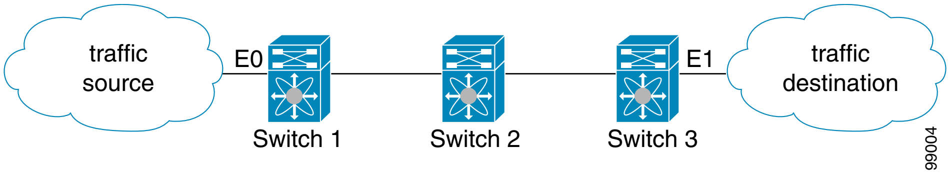

The **access-group** command controls access to an interface. Each interface can only be associated with one access list per direction. The ingress direction can have a different ACL than the egress direction. The access group becomes active on creation.

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).

  Tip 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.

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

To create an access group, follow these steps:

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

Step 2	`switch(config)# `**`interface mgmt0`**
	`switch(config-if)#`
		Configures a management interface (mgmt0).

Step 3	`switch(config-if)# `**`ip access-group SampleName`**	Creates an access group called SampleName for both the ingress and egress traffic (default)
	`switch(config-if)# `**`no ip access-group NotRequired`**	Deletes the access group called NotRequired.
Step 4	`switch(config-if)# `**`ip access-group SampleName1 in`**	Creates an access group called SampleName (if it does not already exist) for ingress traffic.
	`switch(config-if)# `**`no ip access-group SampleName1 in`**	Deletes the access group called SampleName for ingress traffic.
	`switch(config-if)# `**`ip access-group SampleName2 out`**	Creates an access group called SampleName (if it does not already exist) for local egress traffic.
	`switch(config-if)# `**`no ip access-group SampleName2 out`**	Deletes the access group called SampleName for local egress 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.

Below 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
```

Below 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
```

Displaying IP-ACLs

Use the **show ip access-list** command to view the contents of configured access lists. Each access list can have several filters.

Example 17-1 Displays Configured IP-ACLs

```
switch# show ip access-list usage
Access List Name/Number          Filters IF   Status    Creation Time
-------------------------------- ------- ---- --------- -------------
abc                              3       7    active    Tue Jun 24 17:51:40 2003
x1                               3       1    active    Tue Jun 24 18:32:25 2003
x3                             0       1   not-ready Tue Jun 24 18:32:28 2003
```

Example 17-2 Displays a Summary of the Specified IP-ACL

```
switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)
```

Clearing IP-ACL Counters

Use the **clear** command to clear the counters for a specified IP-ACL entry. Note that you cannot use this command to clear the counters for each individual filter.

```
switch# clear ip access-list counters abc permit ip 10.1.1.0 0.0.0.255
```

Configuring IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN using the **ip address** command.

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.

Configuring IP Static Routes

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

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.

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.

Viewing and Clearing ARPs

Address Resolution Protocol (ARP) entries can be viewed (**show arp**), deleted (**no arp**), or cleared (**clear arp-cache**) in Cisco MDS 9000 Family switches 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
switch(config)#
```

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

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

Displaying IP Interface Information

Use the following **show** commands to view configured IP interface information (see Examples 17-3 to 17-6).

Example 17-3 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 17-11).

Example 17-4 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 17-5 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 17-6 Displays the IP Routing Status

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

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:

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 17-4).

Figure 17-4 Overlay VSAN Configuration Example

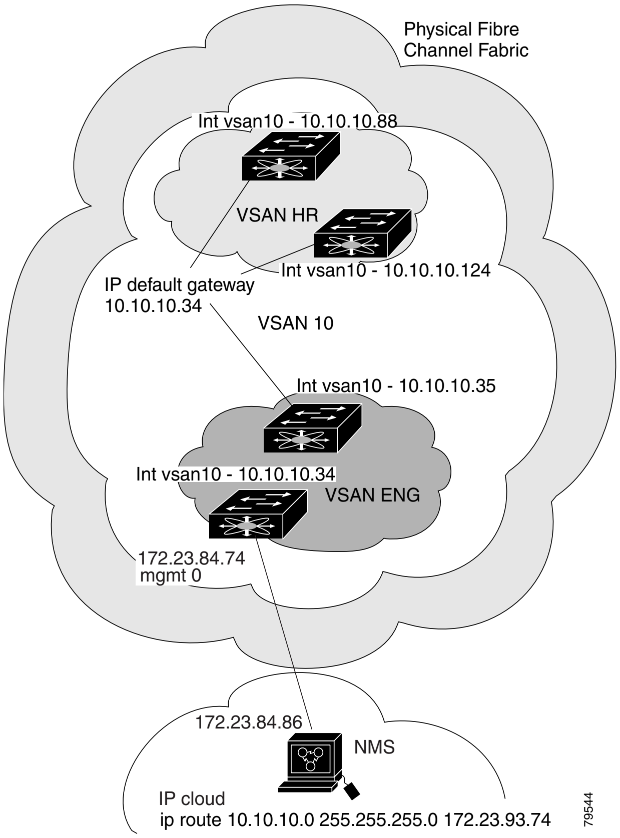

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in Figure 17-4), 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** switch--config-vsan-db#	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.x netmask 255.255.255.0** switch(config-if)#	Assigns an IP address and netmask on all switches in the fabric.
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 17-4, follow this step:

	Command	Purpose
Step 1	switch# **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.

Note To configure the management interface displayed in Figure 17-4, set the default gateway to an IP address on the Ethernet network.

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 (see Figure 17-5).

Figure 17-5 Multiple VSANs Configuration Example

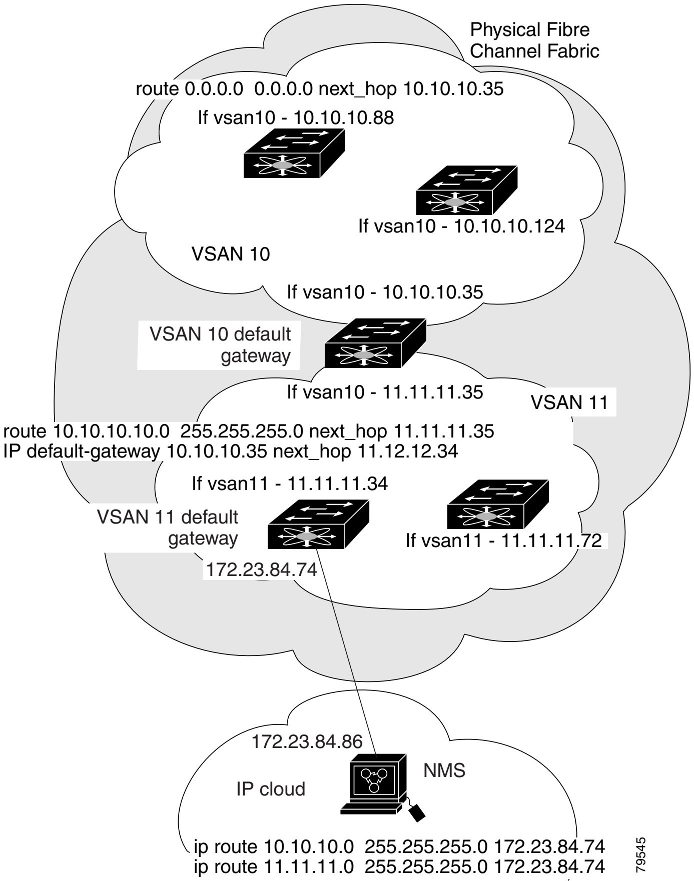

To configure an overlay VSAN (using the example in Figure 17-5), follow these steps:

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

Step 2	`switch(config)# vsan database`
	`switch-config-vsan-db#`
		Configures the VSAN database.

	Command	Purpose
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.x netmask 255.255.255.0** switch(config-if)#	Assigns an IP address and netmask on all switches in VSAN 10.
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.x netmask 255.255.255.0** switch(config-if)#	Assigns an IP address and netmask on all of the switches in VSAN 11.
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.

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 17-6, 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 17-6 VRRP Functionality

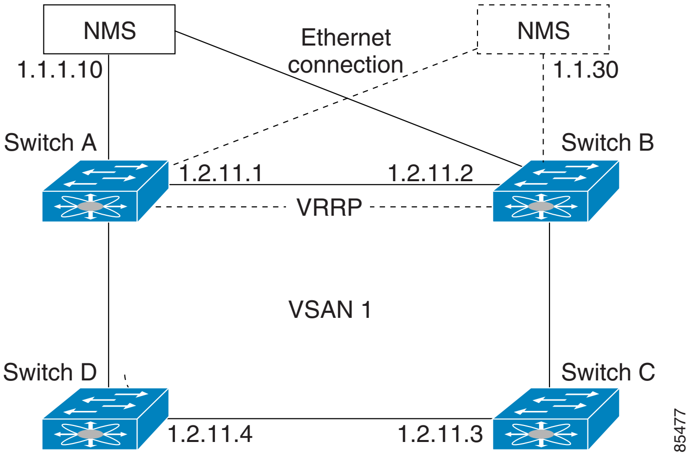

In Figure 17-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 17-7 Redundant Gateway

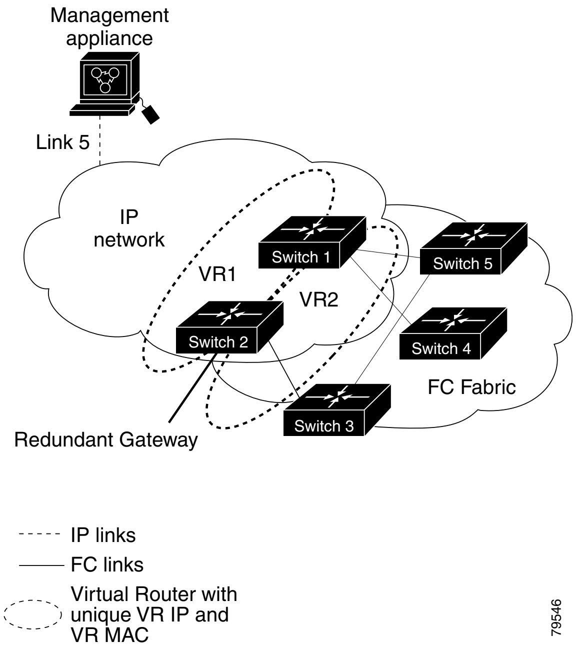

Creating or Removing a Virtual Router

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.

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.

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.

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.

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.

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.

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).

	Command	Purpose

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.

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.

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.

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.

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 the preempt option and allows the master to keep its priority level.

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.

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.

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.

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).

	Command	Purpose
Step 3	switch(config-if)# **vrrp 250** switch(config-if-vrrp)**#**	Creates a virtual router.
Step 4	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 17-7 to 17-10).

Example 17-7 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 17-8 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 17-9 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 17-10 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 17-11).

Example 17-11 Clears VRRP Information

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

Configuring DNS Server

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

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 list 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 list. To delete a name from a list, use the **no** form of this command.

	`switch(config)# no ip domain-list`	Deletes the defined list 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.

	Command	Purpose
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.

The DNS server may be dropped after two attempts due to the following reasons:

- if the IP address or the switch name is wrongly configured
- if 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.

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see Example 17-12).

Example 17-12 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 17-3 lists the default settings for IP features.

Table 17-3 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.

Table 17-3 Default IPFC Settings (continued)

Parameters	Default
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.