
Chapter 6

Software Overview

This chapter provides an overview of the software features supported on the HP 9308M, HP 9304M, and HP 6308M-SX routing switches and the HP 6208M-SX switch.

- For configuration details for these features, see chapter 8 in this guide and the *Advanced Configuration and Management Guide*, included in PDF format on the Product Documentation CD-ROM included with your switch or routing switch product.
- For detailed information about CLI commands shown in this chapter, see the *Command Line Interface Reference*.
- For information about the protocols, RFCs, and standards supported by the software, see the “Software Specifications” appendix in the *Advanced Configuration and Management Guide*.
- For an overview of the hardware, see “Hardware Overview” on page 7-1.

Software Feature Summary

This section lists the flash image files (system software) that HP devices can run and the features that are supported in each type of flash image. HP devices run one of three types of flash images:

- Routing switch code
- Switch code

This section describes the features provided in each type of software and how to determine the type of software an HP device is running.

Flash Images

The flash image (system software) that is running on a device determines the software features that are supported by that device. Table 6.6.1 lists the flash images that can be used on each HP device.

Table 6.1: HP Flash Software Images

Product	Flash image	Description
HP 9308M	HPRxxxxx.BIN	Routing switch code
HP 9304M	H2Rxxxxx.BIN (Redundant Management Module)	
HP 6308M-SX	HPRxxxxx.BIN	Routing switch code

Table 6.1: HP Flash Software Images (Continued)

Product	Flash image	Description
HP 6208M-SX	HPSxxxxx.BIN	Switch code

NOTE: Some features are supported only on specific products or require specific hardware configurations. See the chapters describing those features or contact HP or your HP reseller for information.

Determining the Flash Version a Device Is Running

To determine the flash image running on an HP device, do one of the following.

USING THE CLI

Enter the following command: **show version**

NOTE: You can enter this command from any CLI access level.

USING THE WEB MANAGEMENT INTERFACE

1. Enter the device's IP address in the Location or Address field of your browser.
2. When the device's chassis is displayed, click on a blank area of the chassis' management module to display the Device Information screen. The software is listed in the Running Image Version field.

NOTE: You can access the version information whether you have read-write ("set") or read-only ("get") access.

Feature List

Table 6.6.2 lists the major software features available in the types of flash software listed in Table 6.6.1. Some features are supported only in certain flash software. For each feature, the table indicates the types of flash code in which the feature is supported. Table 6.6.2 uses the following labels to indicate the flash code types:

- **Router** – Routing Switch. A device capable of performing Layer 2, Layer 3, and Layer 4 switching and Layer 3 routing. The following HP devices are configured as routing switches:
 - HP 9308M
 - HP 9304M
 - HP 6308M-SX
- **Switch** – Switch. A device capable of performing Layer 2 switching. The following HP device is configured as a switch:
 - HP 6208M-SX

NOTE: Some features are supported only on specific products. Footnotes at the end of the table list exceptions such as this one.

Some features require specific hardware configurations. See the chapters describing those features or contact HP or your HP reseller for information.

Table 6.2: HP Software Features

Feature	Supported on...		See page...
	Router	Switch	
Access and Management Features			
Command-line and web-based management interfaces	X	X	6-6
Simple Network Management Protocol (SNMP)-based management application	X	X	6-6
Multiple levels of access control	X	X	6-8
RADIUS authentication	X	X	6-8
Dynamic configuration	X	X	6-8
Soft reboot (reboot flash image without resetting the system)	X	X	6-9
Scheduled system reload	X	X	6-9
Telnet	X	X	6-9
Trivial File Transfer Protocol (TFTP)	X	X	6-9
Simple Network Time Protocol (SNTP)	X	X	6-9
Domain Name Server (DNS) resolver	X	X	6-10
Remote Monitoring (RMON)	X	X	6-10
SNMPv2c	X	X	6-10
SNMP alarms and trap log	X	X	6-10
SyslogD client	X	X	6-10
Ping and trace-route facilities	X	X	6-11
Port mirroring	X	X	6-11
Layer 2 Switching Features			
MAC switching	X	X	6-11
Static MAC entries	X	X	6-12
Spanning Tree Protocol	X	X	6-12
Trunk groups	X	X	6-12
Port-based Virtual LANs (VLANs)	X	X	6-13
802.1p VLAN tagging	X	X	6-13
MAC filters	X	X	6-13
Address-lock filters		X	6-13
Dynamic Host Configuration Protocol (DHCP) Assist		X	6-14
IP Multicast Containment	X	X	6-14

Table 6.2: HP Software Features (Continued)

Feature	Supported on...		See page...
	Router	Switch	
Layer 3 Switching Features			
Protocol-based Virtual LANs (VLANs)	X	X	6-14
Layer 3 Routing Features			
Multi-netting	X		6-15
Multi-port subnets (integrated switch-routing)	X	X	6-15
Static IP routes, Address Resolution Protocol (ARP) entries, and Reverse ARP (RARP) entries	X		6-16
IP/RIP routing	X		6-16
IP/OSPF routing	X		6-17
IP route and protocol-port filters	X		6-17
IP/RIP filters	X		6-17
IPX routing	X		6-17
IPX route and socket filters	X		6-18
IPX/RIP and IPX/SAP filters	X		6-18
AppleTalk routing	X		6-18
AppleTalk zone and network filters	X		6-18
IP Multicast Routing (PIM and DVMRP)	X		6-18
IP/RIP and IP/OSPF redistribution filters	X		6-19
User Datagram Protocol (UDP) Helper	X		6-19
Layer 4 Switching Features			
TCP/UDP access policies	X	X	6-19
Load Balancing and Redundancy Features			
Selectable Quality of Service (QoS) (Layers 2, 3, and 4)	X	X	6-20
Router-based health checking	X		6-20
Virtual Router Redundancy Protocol (VRRP)	X		6-20
Server Redundancy Protocol (SRP)	X		6-21

Showing System Defaults

You can display the defaults for system parameters using either of the following methods.

USING THE CLI

To display the default information, enter the following command from any level of the CLI:

```
show default [values]
```

If you specify “default” but not the optional “values”, the default states for parameters that can either be enabled or disabled are displayed. If you also specify “values”, the default values for parameters that take a numeric value are displayed.

Here is an example of the information displayed by the **show default** command on an HP 9308M routing switch.

```
HP9300# show default
spanning tree disabled
auto sense port speed      port untagged      port flow control on
no username assigned      no password assigned  boot sys flash primary
system traps enabled      sntp disabled      radius disabled
rip disabled              ospf disabled      bgp disabled

when ip routing enabled :
ip irdp enabled           ip load-sharing enabled  ip proxy arp enabled
ip rarp enabled           ip bcast forward enabled
dvmrp disabled           pim/dm disabled
vrrp disabled            srp disabled

when rip enabled :
rip type:v2 only         rip poison rev enabled

ipx disabled             appletalk disabled
```

See the *Command Line Interface Reference* for additional examples.

USING THE WEB MANAGEMENT INTERFACE

You cannot display the system defaults using the Web management interface.

NOTE: You can display and configure the sizes of various tables such as the MAC, ARP, and IP tables by selecting the Parameters link from the System configuration sheet. See “Displaying and Modifying System Parameter Default Settings” on page 8-58.

Access and Management Features

The following sections describe the access and management features listed in Table 6.6.2.

Management Interfaces

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches and the HP 6208M-SX switch can be managed using any of the following interfaces:

- Command Line Interface (CLI) – a text-based interface accessible through a direct serial connection or a Telnet session.
- Web management interface – A GUI-based management interface accessible through an HTTP (web browser) connection.

Command Line Interface (CLI)

The CLI is a text-based operator interface that allows you to configure a system with a PC or terminal without special software.

Up to five read-only Telnet sessions can operate concurrently. Only one read-write Telnet session is allowed at a time.

Web Management Interface

A Web management interface is supported on web browsers Netscape Navigator™ versions 2.0 or later, and Microsoft Internet Explorer™ versions 3.0 or later. No application software is required.

To use the Web management interface, open a web browser and enter the IP address of the device in the Location or Address field. The web browser contacts the device and displays a login dialog, as shown in Figure 6.1.




Figure 6.1 Web Management interface login dialog

- For read-write access, enter “set” in the User Name field and a read-write community string that you have configured on the device in the Password field, as shown in Figure 6.1. (For security, the software displays asterisks when you type your password.)
- For read-only access, enter “get” in the User Name field and “public” (the default read-only community string) or a read-only community string you have configured in the Password field.

NOTE: The software does not contain a default read-write SNMP community string. You must configure a read-write string before you can make configuration changes using the Web management interface. See “Configuring the SNMP Community Strings” on page 3-5.

Figure 6.2 shows an example of the Web management display of an HP 9308M routing switch.

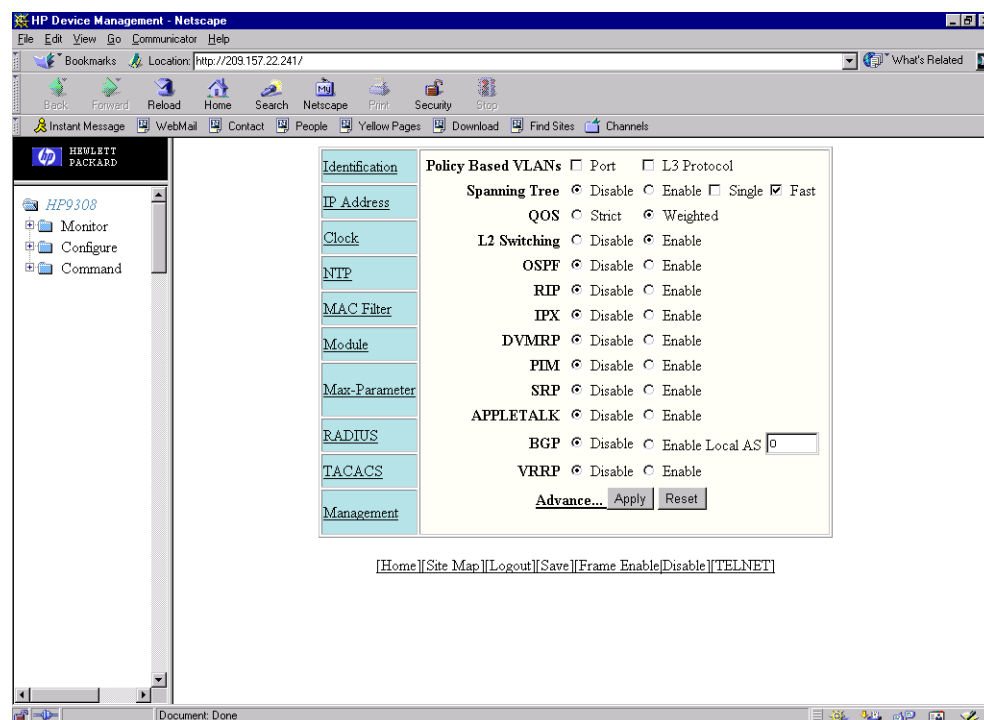


Figure 6.2 Example of Web management interface

The configuration and management procedures in this guide include instructions for the Web management interface.

- To display general system information, click on a blank area of the device's management module. If the chassis display is disabled as shown in this example, click on the object shown in the chassis window. The object contains the product name.
- To display information about a specific port, click on the port on the front panel display. (This option is available only when you enable display of the front panel. See the note below.)
- Click on the links in the left-hand frame or on the bottom of the display to view statistics or to view and change configuration parameters.

NOTE: The Web management interface automatically refreshes the system information at regular intervals, including the link LEDs for the ports. To streamline performance, display of the device's front panel is disabled by default. To enable front panel display, select the [Preference](#) link, select the Enable radio button for Front panel display, then click Apply. Select Reload or Refresh on your browser's tool bar to immediately see the effect of the change.

Multiple Levels of Access Control

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches and the HP 6208M-SX switch provide multiple levels of access to allow system administrators complete configuration control while protecting the system from unauthorized changes.

CLI Access

Three levels of password protection offer a range of access points for various users within the network. The three levels are:

- **Super user** – This setting allows a user unlimited access to all levels of the CLI. This level is generally reserved for system administrators within the network. The super user is also the only one who can assign a password access level to another user.
- **Configure port** – This level allows a user to configure interface parameters only and to view any show command displays.
- **Read only** – A user at this password level will only be able to view show command displays within the CLI. No configuration is allowed at this password level.

Web Management Interface Access

By default, access through the Web management interface is controlled by passwords associated with the “get” (read-only) and “set” (read-write) SNMP community strings. The default password for “get” is “public”. There is not default password for “set”. You can configure SNMP community strings using CLI commands. See “Configuring the SNMP Community Strings” on page 3-5. You also can use locally configured user names and passwords to control access through the Web management interface. See “Local Access Control” on page 6-8.

Local Access Control

You can configure up to 16 user names and passwords to control access to a device. The passwords and user names can be used for accessing devices using the CLI and the Web management interface. For each management platform, you configure an authentication-method list that specifies sources the device can consult to authenticate an access attempt and the order in which to consult the sources. For example, you can configure an authentication-method list to authenticate CLI management access based on a local access list first (user names and passwords you have configured), then a RADIUS server, then the enable passwords.

See “Configuring Local User Accounts” on page 3-13 and “Configuring Authentication-Method Lists” on page 3-24.

TACACS and TACACS+ Authentication

You can secure CLI access to the switch or routing switch by configuring the device to consult a Terminal Access Controller Access Control System (TACACS) or TACACS+ server to authenticate user names and passwords. See “Configuring for TACACS/TACACS+ Authentication” on page 3-15.

NOTE: TACACS/TACACS+ authentication is not supported for Web management.

RADIUS Authentication

You can further secure CLI access to the switch or routing switch by configuring the device to consult a Remote Access Dial In User Service (RADIUS) server to authenticate user names and passwords. See “Configuring for RADIUS Authentication” on page 3-20.

NOTE: RADIUS authentication is not supported for Web management.

Dynamic Configuration

Dynamic configuration enables you to make configuration changes without rebooting the system. Many of the configuration changes you can make to the devices do not require a reboot and take effect immediately. You can make the changes without causing network outages. The individual configuration chapters describing each feature area (chapters 7-18) list the parameters that can be dynamically changed.

Soft Reboot

When you upgrade the software image on a device, you do not need to power down the system to use the new software. You can boot the new software immediately from the primary flash, secondary flash, a TFTP server, or a BootP server.

You also can use this feature to test new versions of flash code before replacing the previous flash image.

For more details on the boot commands and on copying software to and from devices, refer to “Updating Software Images and Configuration Files” on page 5-1.

Scheduled System Reload

Although the dynamic configuration feature (see “Dynamic Configuration” on page 6-8) allows many parameter changes to take effect immediately without a system reset, other parameters do require a system reset.

To place these parameters into effect, you must save the configuration changes to the configuration file, then reload the system. The management interfaces provide an option to immediately reset the system. Alternatively, you can use the scheduled system reload feature to configure the system to reload its flash code at a specific time (based on the system clock or SNTP time) or after a specific amount of time has passed.

See “Updating Software Images and Configuration Files” on page 5-1.

Telnet

As described in “Management Interfaces” on page 6-6, HP devices allow you to access the CLI through a Telnet connection. To establish the Telnet connection, you need the following:

- An IP address on the device. See “Assign a Permanent IP Address” on page 2-13 for information.
- A third-party terminal emulation application installed on a PC or workstation that has network access to the HP device.

Trivial File Transfer Protocol (TFTP)

The HP devices allow you to use TFTP to copy files to and from the flash memory modules on the management module. You can use TFTP to perform the following operations:

- Upgrade boot or flash code.
- Archive boot or flash code or a configuration file on a TFTP server.
- Load the system using flash code and a configuration file stored on a TFTP server. (This occurs as part of the BootP or DHCP process.)

NOTE: Certain boot upgrades may require you to install new firmware. Contact your HP reseller or HP for information.

See “Updating Software Images and Configuration Files” on page 5-1 for more information about using TFTP on the devices.

Simple Network Time Protocol (SNTP)

The HP devices can use either of two time and date sources:

- An on-board system clock.
- An external SNTP server. The server can be on the same sub-net or a different sub-net.

If you have access to an SNTP server, HP recommends that you use the SNTP server as the time and date source. Using an SNTP server ensures that all devices that use the SNTP server have a consistent time and date. In addition, the settings on the system time counter are not retained across power cycles. The counter has to be reset following each power-up. If the device is configured to reference an SNTP server, the device automatically sets its time counter according to the SNTP server after a system reset.

Regardless of the time and date source you use, you can configure the time zone of the time and date. You also can enable daylight savings time, which is disabled by default.

See “Setting the System Clock” on page 8-14 for more information about setting the time and date.

Domain Name Server (DNS) Resolver

The DNS Resolver feature allows you to use just a host name rather than a fully-qualified domain name when you use Telnet, ping, and trace-route commands. To configure the feature, you specify the domain name, then specify the IP addresses of up to four DNS servers that have authority for the domain.

For example, if you define the domain “newyork.com” on a device, you can initiate a ping to a host on that domain by specifying only the host name in the command. You do not need to specify the host’s entire domain name.

As an example, here are two CLI commands. The first command uses only the host name. The second command uses the fully-qualified domain name for the host.

```
HP9300# ping nyc01
HP9300# ping nyc01.newyork.com
```

See “Enabling Domain Name Server (DNS) Resolver” on page 8-7 for information about configuring this feature.

SNMPv2c Support

The HP devices support SNMPv2c, including support for GetBulk requests. The SNMPv2c support is enabled by default and cannot be disabled. Thus, you do not need to perform any configuration on the device to use the feature.

To use this enhancement, you need an SNMP management application that is capable of sending GetBulk requests. See the documentation for your application for more information.

NOTE: The SNMPv2c support does not include support of SNMPv2c traps.

Remote Monitoring (RMON) Statistics

The HP devices include an RMON agent that supports the following groups. The group numbers come from the RMON specification (RFC 1757).

- Statistics (RMON Group 1) – Current packet and error statistics for each port.
- History (RMON Group 2) – Samplings of packet and error statistics captured at regular intervals. You can configure the sampling rate and the number of “buckets” in DRAM for storing the samplings.
- Alarms (RMON Group 3) – A list of alarm events, which indicate that a threshold level for a specific part of the device has been exceeded. You can select the system elements you want RMON to monitor and the thresholds for triggering the alarms.
- Events (RMON Group 9) – A log of system events (such as port-state change to up or down, and so on) and alarms. RMON Group 9 also specifies the action to be taken if an alarm threshold is exceeded.

See the “Network Monitoring” appendix in the *Advanced Configuration and Management Guide* for information about setting and displaying the RMON statistics.

Syslog Logging

In addition to the event and alarm logs provided by RMON, HP devices contain a Syslog agent that can write log messages to a local buffer and optionally to a third-party SyslogD server. The Syslog feature can write messages at the following severity levels:

- Emergencies
- Alerts
- Critical

- Errors
- Warnings
- Notifications
- Informational
- Debugging

Syslog is disabled by default.

When you enable the Syslog feature, the device automatically writes the Syslog messages to a local buffer. If you specify the IP address or name of a SyslogD server, the device also writes the messages to the SyslogD server. The default facility for messages written to the server is “user”. You can change the facility if needed. You also can change the number of entries that can be stored in the local buffer. The default is 50. HP devices do not have a limit to the number of messages that can be logged on a remote SyslogD server.

NOTE: You can specify only one facility.

See “Configuring the Syslog Service” on page 8-16 for configuration information.

Ping and Traceroute Facilities

After you configure an IP address for the device, you can test the device’s network connections using the following facilities:

- Ping – You can send a test packet to a host’s IP address or host name. If the packet reaches the host, the host generally sends a reply packet to let you know the host received your ping. If the host does not reply within a specified interval, the device re-attempts the ping up to a specified number of times.
- Traceroute – You can trace the IP path to a host. The traceroute feature displays a list of all the intervening router hops the trace-route request traversed to reach the host.

See “Verifying Proper Connections” on page 2-20.

Port Mirroring

The mirror port feature lets you connect a protocol analyzer to a port on a device to observe the traffic flowing into and out of another port on the same device. To use this feature, you specify the port you want to monitor and the port into which you are plugging the protocol analyzer.

NOTE: Only one mirror port can be active on a switch or routing switch at a time. By default, no mirror port is assigned.

For more information, see “Assigning a Mirror Port and a Monitor Port” on page 8-61.

Layer 2 Switching Features

The following sections describe the Layer 2 switching features listed in Table 6.6.2.

MAC Switching

The HP devices support MAC switching. **MAC switching** enables intelligent wire-speed bridging of Layer 2 packets. The first time a device receives a packet from a given MAC destination, the device makes an entry in its Layer 2 cache. The entry consist of the packet’s source MAC address and the port on which the device received the packet.

When the device receives a bridge packet destined for the cached address, the device does not need to send the packet as a broadcast through all the ports within the broadcast domain. Instead, the device can intelligently send the packet only through the port to which the destination device is connected. Thus, even though Layer 2 domains are typically broadcast domains, MAC switching enhances performance in the domain by reducing the amount of broadcast traffic in the domain.

In addition, HP routing switches that are enabled for MAC switching can switch traffic for route protocols that are not supported in the routing software. If IPX routing is disabled on a routing switch, the routing switch can switch the IPX packets instead.

To avoid accumulating stale cache entries, the HP devices use an aging mechanism. The aging mechanism removes a learned entry from the cache after the entry has remained unused for a specified interval (by default, 300 seconds). You can change or disable the aging interval.

See “Configuring Basic Layer 2 Parameters” on page 8-30 for more information about configuring MAC switching parameters.

NOTE: By default, all ports in a device belong to a common Layer 2 broadcast domain, VLAN 1. You can configure port-based VLANs (Virtual LANs) to create smaller broadcast domains that use subsets of the device’s ports. See “Port-Based Virtual LANs (VLANs)” on page 6-13.

Static MAC Entries

MAC entries that the device learns and caches are subject to an aging time. After a cached entry remains unused for the duration of the aging time, the software removes the entry from the Layer 2 cache. If you want certain MAC addresses to always be present in the device’s Layer 2 address table, you can add them as static entries.

A **static MAC entry**, like a cached (dynamic) MAC entry, maps a MAC address to the port attached to the device that has the MAC address.

Unlike cached MAC entries, static MAC entries provide the following benefits:

- You can assign a QoS priority to a static MAC entry.
- You can specify VLAN membership for a static MAC entry.
- A static entry prevents broadcast storms that can be caused when a server’s MAC entry is removed. For example, if the server goes down long enough for the server’s entry to age out, the HP device sends packets addressed to the server as broadcasts until the device relearns the cache entry for the server.

You can specify port priority (QoS) and VLAN membership (VLAN ID) for the MAC address. On switches, you also can specify the device type (router or host) for the entry.

NOTE: On the HP 9308M, HP 9304M, and HP 6308M-SX routing switches, you also can create static IP routes, ARP entries, and RARP entries. The HP 6208M-SX switch supports only static MAC addresses.

For more details on configuring static MAC addresses, see “Configuring Static MAC Entries” on page 8-34.

Spanning Tree Protocol (STP)

The **Spanning Tree Protocol (STP)** is a protocol for detecting and eliminating logical loops in a Layer 2 broadcast domain. STP is described in the IEEE 802.1d bridge protocols standard.

STP also ensures that the device uses the most efficient path when multiple paths exist between ports. Moreover, if a selected path fails, STP searches for and then establishes an alternate path to prevent or limit retransmission of data.

STP is disabled by default on routing switches but is enabled by default on the HP 6208M-SX switch.

For information about configuring STP, see “Enabling or Disabling the Spanning Tree Protocol (STP)” on page 8-30.

Trunk Groups

A **trunk group** is a set of ports that provide a high speed link between two HP devices or between an HP device and a server. A trunk group can provide a transfer rate of up to 4 Gbps of bi-directional traffic.

In addition to enabling load sharing of traffic, trunk groups provide redundant, alternate paths for traffic. Thus, if a link in a trunk group fails, the device still uses the other links in the trunk group.

A trunk group can consist of two to four ports, and up to eight ports on a Chassis device. You can configure up to 64 trunk groups on a device. (If your device has fewer than 64 ports, then the trunk group limit depends only on the number of ports.)

For configuration information, see “Configuring Trunk Groups” on page 8-38.

Port-Based Virtual LANs (VLANs)

By default, all ports in a device belong to a common Layer 2 broadcast domain. When the device sends a broadcast packet, the packet goes out all active ports. A **port-based VLAN** (Virtual LAN) is a subset of ports on a device that constitutes a Layer 2 broadcast domain.

Port-based VLANs can reduce the likelihood and severity of broadcast storms by reducing the number of ports affected by a storm. In addition, for devices such as servers that can cause broadcast storms, you can add static MAC entries for the devices and assign the static entries to a VLAN.

Each port-based VLAN maintains a separate spanning tree. (See “Spanning Tree Protocol (STP)” on page 6-12.)

For configuration information, see the “Configuring VLANs” chapter in the *Advanced Configuration and Management Guide*.

VLAN Tagging

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches and the HP 6208M-SX switch support 802.1p VLAN tagging. **VLAN tagging** is a method of identifying a packet as a member of a VLAN. VLAN tagging enables you to configure ports on multiple switches into a single VLAN. Using tagged VLANs can ease network management and ensures interoperability with other devices.

When a switch sends a packet that is a member of a tagged VLAN, the switch “tags” the packet to indicate its VLAN membership. Other switches that support VLAN tagging recognize the tag and process the packet according to its VLAN membership.

For more information, see the “Configuring VLANs” chapter in the *Advanced Configuration and Management Guide*.

MAC Filters

A **MAC filter** enables you to explicitly permit or deny switching of a Layer 2 packet received by the device. When the device receives a Layer 2 packet for switching, the device checks the packet’s contents against the defined MAC filters. If the packet matches a filter, the system takes the action specified in the filter.

- If the action is permit, the system allows the packet to be switched.
- If the action is deny, the system immediately drops the packet.

To ensure security, if a packet does not match any of the MAC filters defined on the system, the system drops the packet by default. To configure the system to permit packets by default, you must define the last MAC filter in the filter list to allow all packets.

MAC filters can evaluate packets based on criteria such as source address and mask, destination address and mask, and protocol type (IP, ARP, and so on).

See “Defining MAC Address Filters” on page 8-52 for information on configuring MAC filters.

Address-Lock Filters

An **address-lock filter** restricts the number of MAC addresses that a switch can learn from a specific port. After the switch learns the specified number of MAC addresses from the port, the switch stops learning addresses received on that port. In addition, the switch does not accept or forward traffic on the port unless the traffic contains one of the source or destination MAC addresses locked for the port.

Address-lock filters apply only to Layer 2 traffic and do not affect Layer 3 or Layer 4 traffic on the locked ports.

Unlike addresses learned from other ports, addresses learned from a locked port are not subject to aging.

See “Locking a Port To Restrict Addresses” on page 8-57 for information on configuring address-lock filters.

Dynamic Host Configuration Protocol (DHCP) Assist

DHCP Assist allows the HP 6208M-SX switch to assist a router that is performing multi-netting on its interfaces as part of its DHCP relay function. DHCP eliminates the need to manually assign IP addresses to clients. Instead of each client having a statically configured IP address, clients petition a server for IP addresses when the clients are booted.

DHCP Assist ensures that a DHCP server that manages multiple IP sub-nets can readily recognize the requester's IP sub-net, even when that server is not on the client's local LAN segment. The HP 6208M-SX switch does this by stamping the correct gateway IP address into a DHCP discovery packet on behalf of the router.

NOTE: DHCP assist applies only to the HP 6208M-SX switch. To configure a routing switch to assist DHCP packets, use the UDP Helper feature. See "User Datagram Protocol (UDP) Helper" on page 6-19.

See "Configuring DHCP Assist (switch only)" on page 8-48 for information on configuring DHCP assist.

IP Multicast Containment

IP multicast containment allows the HP 6208M-SX switch to limit switching of IP multicast packets to only those ports on the switch that are identified as IP multicast members. The switch can provide IP multicast containment in either of the following modes:

- **Passive** – The switch listens for Internet Group Membership Protocol (IGMP) packets and forwards them to the appropriate ports.
- **Active** – The switch actively sends out host queries to identify IP multicast groups on the network and inserts this information into the IGMP packets.

Routers in the network generally handle host queries. Unless your configuration does not contain a router to provide this service, use IP multicast containment in the passive mode.

Layer 3 Switching Features

The following section describes the Layer 3 switching features.

Protocol-Based Virtual LANs (VLANs)

Protocol and sub-net based VLANs increase network performance and provide managers with a high degree of network flexibility.

With sub-net VLANs, devices with a common sub-net can be resident across multiple ports of a device. This increases performance by providing a greater pool of bandwidth for all devices.

Protocol VLANs enable managers to easily and transparently group like protocols into a defined VLAN. This reduces the number of non-essential broadcasts on other ports and allows a port to belong to multiple VLANs without VLAN tagging.

You can define Layer 3 VLANs for the following protocols:

- IP protocol
- IPX protocol
- IP sub-net
- IPX network number
- AppleTalk cable range
- AppleTalk
- Decnet
- NetBIOS
- Others

For more details on the value and configuration of VLANs, see the “Configuring VLANs” chapter in the *Advanced Configuration and Management Guide*.

Routing Between VLANs

In addition to supporting the assignment of VLANs, the HP 9308M, HP 9304M, and HP 6308M-SX routing switches support routing between VLANs using virtual interfaces.

VLAN Tagging

VLAN tagging (802.1p) extends the boundaries of the VLAN by allowing creation of VLANs that cross switch boundaries. This eases network management and ensures interoperability with other devices. See “VLAN Tagging” on page 6-13.

Layer 3 Routing Features

The following sections describe the Layer 3 routing features listed in Table 6.6.2. The HP 9308M, HP 9304M, and HP 6308M-SX routing switches provide traditional Layer 3 routing at wire speeds with support for the following routing protocols:

- Internet Protocol (IP)
 - IP Routing Information Protocol (RIP)
 - Open Shortest Path First (OSPF)
- Internet Packet Exchange (IPX)
 - IPX RIP
 - IPX Service Advertisement Protocol (SAP)
- AppleTalk

The following sections describe the routing support for the protocols listed above and the additional routing features.

Multi-Netting

Multi-netting allows you to assign multiple IP or IPX protocol interfaces to the same physical port on the switch or routing switch. The routing switches support multi-netting for IP and IPX.

Multiple IP Sub-nets per Interface

Up to 64 IP sub-nets can be defined per port. IP/RIP and OSPF can be assigned to these multi-homed interfaces.

Multiple IPX Frame Type Support per Interface

Up to four different IPX network numbers and frame encapsulation types can be defined for each IPX interface. You can define and receive traffic from four separate IPX networks on a single interface. Each of the networks must have a distinct network number and one of the following encapsulation types: Ethernet SNAP, Ethernet 802.2, Ethernet 802.3, and Ethernet II.

Multi-Port Subnets (Integrated Switch-Routing)

Integrated switch routing allows a routing switch to assign and support VLANs on its interfaces as would a switch. In addition, this feature provides routing between its VLANs. This combined logical switch and router operation within a single device is what defines a routing switch as an Integrated Switch-Router, as shown in Figure 6.3.

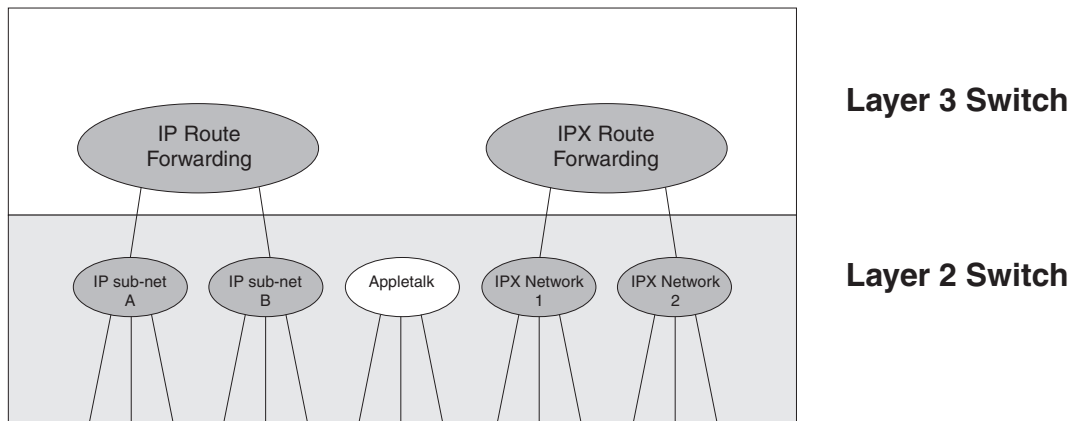


Figure 6.3 Logical representation of ISR within routing switches

Routing between the VLANs is performed without dedicating physical ports by using *virtual interfaces*. These virtual interfaces serve as a link between the configured VLANs and the routing core of the routing switches.

The ISR architecture provides the platform for support of policy-based VLANs within the routing switches.

Static IP Routes, Address Resolution Protocol (ARP) Entries, and Reverse ARP (RARP) Entries

The routing switches can learn and cache IP routes, ARP entries, and RARP entries. In addition, you can add static entries for these items.

For more details on configuring static routes, ARP, or RARP entries, see the “Configuring IP and IP/RIP” chapter in the *Advanced Configuration and Management Guide*.

IP/RIP Routing

IP is the most widely used networking protocol on the Internet and in enterprise LANs. The implementation of IP on the HP 9308M, HP 9304M, and HP 6308M-SX routing switches adheres to the IP-related RFCs listed in the “Software Specifications” appendix in the *Advanced Configuration and Management Guide*. In addition, features such as multi-netting, integrated switch-routing, and VLANs (described in the sections above) enhance the IP support.

The routing switches support the following intra-domain routing protocols:

- RIP
- OSPF

NOTE: An intra-domain protocol is a protocol that is used by routers under common administrative control. The term “domain”, used in this context, is synonymous with “autonomous system”. In contrast, Border Gateway Protocol (BGP) is an example of an inter-domain protocol. BGP is used by routers in one domain to exchange information with routers in other domains.

RIP

RIP is a distance-vector protocol. It uses a cost value associated with each route to express the preferability of that route. Generally, the cost is equivalent to the number of hops in the route, but the routing switches allow you to bias the preferability of a route by changing its cost. You also can configure the routing switches to prefer one route over another equal cost route.

By default, the routing switches using RIP propagate route information to other RIP routers by sending route updates every 30 seconds. You can change this update interval if needed.

You can enable routing switches to use RIP version 1, RIP version 1 with version 2 compatibility, or RIP version 2 to manage IP routes. The default is version 2.

As described in “Static IP Routes, Address Resolution Protocol (ARP) Entries, and Reverse ARP (RARP) Entries” on page 6-16, you also can make static route entries if needed.

See the “Configuring IP and IP/RIP” chapter in the *Advanced Configuration and Management Guide* for information.

OSPF Routing

OSPF is a link-state routing protocol. Each router that runs OSPF uses information from its own interfaces and from other OSPF routers to build a topological map of the network. OSPF routers exchange link-state databases and then periodically send link-state advertisements to notify other routers of route changes.

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches are configured to be compliant with RFC 1583 OSPF V2 (RFC 1583) by default. You also can configure routing switches to run the latest OPSF standard, RFC 2178.

See the “Configuring OSPF” chapter in the *Advanced Configuration and Management Guide* for information.

IP Access and QoS Filters

You can control the IP traffic that a routing switch receives and forwards by defining **IP access policies**. An IP access policy can filter on source IP address, destination IP address, UDP port number, or TCP port number. For example, if you want to permit Telnet access only to specific IP addresses, you can create permit policies for those IP addresses.

You also can use IP access policies to specify the Quality of Service (QoS) packets that certain Layer 4 session should receive. A Layer 4 session is a combination of the source and destination addresses and the TCP or UDP port number. For more information about QoS, see “Selectable Quality of Service (QoS)” on page 6-20.

You assign policies to individual ports by defining access policy groups. An access policy group identifies a list of policies and a set of ports to which the policies are applied. Access policies are applied in the order you list them in the access policy group.

IP Route Filters

You can use IP route filters to control the following:

- Routes learned (cached) by the routing switch. IP route filters applied to inbound traffic affect the routes that the routing switch learns.
- Routes advertised by the routing switch. Filters assigned to outbound traffic affect the routes that the routing switch advertises.

You specify whether the filter is applied to incoming or outgoing traffic by adding individual filters to filter groups and assigning the groups to specific ports.

For details on IP and IP/RIP filters and how to configure them, see the “Configuring IP and IP/RIP” chapter in Book 2.

You can control the RIP neighbors from which the routing switch learns RIP updates by defining RIP neighbor filters. Neighbor filters either permit or deny RIP updates from the specified neighbor.

IPX Routing

The routing switches support the Internet Packet Exchange (IPX) protocol created by Novell™. IPX is based on a client-server networking architecture.

The Routing Information Protocol (RIP) and the Service Advertisement Protocol (SAP) are two key components of Novell NetWare and its IPX protocol suite. By default, Novell NetWare versions 3.x and 4.x broadcast RIP and SAP updates at 60 second intervals.

Up to four different IPX network numbers and frame encapsulation types can be defined for each IPX interface on the routing switches. Therefore, you can define and receive traffic from four separate IPX networks on a single

interface. Each of the networks must have a distinct network number and encapsulation type (Ethernet SNAP, Ethernet 802.2, Ethernet 802.3 and Ethernet II).

IPX Forward Filters

You can define IPX filters to control client access to servers. For example, if you want to restrict access to a print server to specific users, you can define a filter group containing filters that check for the source IPX addresses and nodes of those users. The filter explicitly permits users that match a filter to access the print server specified by the destination address, destination node, and socket number of the print server.

For details on IPX filtering and how to configure the filters, see the “Configuring IPX” chapter in the *Advanced Configuration and Management Guide*.

IPX/RIP and IPX/SAP Filters

In addition to controlling client access to servers, you can control the following:

- Client access to other IPX networks. You control client access to other IPX networks by filtering IPX/RIP routes received or advertised by the routing switches.
- Client access to services. You control service by filtering IPX/SAP service advertisements sent by the routing switch.

For information on configuring IPX/RIP and IPX/SAP filters, see the “Configuring IPX” chapter in the *Advanced Configuration and Management Guide*.

AppleTalk Routing

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches support Phase II AppleTalk routing. The implementation supports all the following AppleTalk protocols:

- EtherTalk Link Access Protocol (ELAP) – AppleTalk physical layer protocol
- Datagram Delivery Protocol (DDP) – AppleTalk equivalent of IP/UDP
- AppleTalk Echo Protocol (AEP) – AppleTalk equivalent of IP/ICMP
- AppleTalk Transaction Protocol (ATP) – AppleTalk equivalent of IP/TCP

NOTE: A sub-set of ATP is implemented to support ZIP on HP routing switches.

- Name Binding Protocol (NBP) – AppleTalk equivalent of IP/DNS

AppleTalk Zone and Network Filters

Zone filters and network filters enable you to control access to AppleTalk networks and individual nodes:

- Zone filters – Explicitly permit or deny access to specific zones on specific ports
- Network filters – Explicitly permit or deny access to specific networks on specific ports

IP Multicast Routing (PIM and DVMRP)

Multicast protocols allow a group or channel to be accessed over different networks by multiple stations (clients) for the receipt and transmit of multicast routing. Distribution of stock quotes, video transmissions, such as news services or remote classrooms and video conferencing, are all examples of multicast routing.

The routing switches support the following IP multicast protocols:

- Distance Vector Multicast Routing Protocol (DVMRP) – a broadcast and pruning multicast protocol that delivers IP multicast datagrams to its intended receivers.
- Internet Group Membership Protocol (IGMP) – a protocol used by DVMRP routers to advertise multicast groups to the routers that are distributing the multicasts.

- Protocol Independent Multicast Dense (PIM-DM) protocol – an alternative to DVMRP that uses the router's IP route table rather than maintaining a separate multicast route table as DVMRP does.

DVMRP and PIM can concurrently operate on different ports of the routing switches.

For both versions of IP multicast, the routing switches support IP tunneling. IP tunneling allows routing switches that are performing IP multicast to send multicast traffic through routers that do not support either PIM or DVMRP multicasting.

For more details on configuring the routing switches for IP multicast, see the “Configuring IP Multicast” chapter in the *Advanced Configuration and Management Guide*.

IP/RIP and IP/OSPF Redistribution Filters

The routing switches allow you to configure parameters for redistributing routes among the following routing protocols:

- IP/RIP
- IP/OSPF

For example, a routing switch running OSPF and RIP can pass a route learned through RIP to OSPF. The routing switch associates a metric and other parameters with a route when the routing switch redistributes the route to other protocols. You can modify these parameters and permit or deny routes from being distributed using route distribution filters.

You define the filters for each of the protocols that redistributes the routes. For example, if you want to control how the routing switch redistributes routes learned through RIP to OSPF, you use IP/RIP commands or Web management screens to define the filters.

User Datagram Protocol (UDP) Helper

The routing switches can relay UDP packets to their destination for a specific application even when the destination server is not on the local LAN segment. For example, a routing switch can relay UDP packets for the following applications to their destination nodes: bootps, domain, and tftp. This feature is especially useful for configuring routing switches to help DHCP packets reach their intended server and client.

For details on UDP helper and its configuration, see the “Configuring IP and IP/RIP” chapter in the *Advanced Configuration and Management Guide*.

NOTE: UDP Helper is supported only on routing switches. To configure the HP 6208M-SX switch to help BootP\DHCP packets, use the DHCP Assist feature. See “Configuring DHCP Assist (switch only)” on page 8-48.

TCP/UDP Access Policies

TCP/UDP access policies (sometimes called session filters) allow you to filter packets for specific Layer 4 sessions. For example, you can use session filters to prohibit specific users from using UDP port 80 (HTTP for web traffic). The switch and the routing switches support TCP/UDP access policies. For syntax information, see the “Policies and Filters” appendix in the *Advanced Configuration and Management Guide*.

Load Balancing and Redundancy Features

The following sections describe the load balancing and redundancy features listed in Table 6.6.2.

Selectable Quality of Service (QoS)

The HP 9308M, HP 9304M, and HP 6308M-SX routing switches and the HP 6208M-SX switch provide the following types of selectable Quality of Service (QoS):

- Layer 2 flow control – The devices provide support for the Full Duplex Flow Control specification, 802.3x.
- Layer 2 and Layer 3 802.1p – 802.1q QoS support provides benefits beyond the local switch or routing switch by supporting and recognizing standard-based virtual LAN (VLAN) tagging, in addition to providing support for flow control, port priority and IP multicast traffic reduction.
- Layer 2 and Layer 3 packet-based priority for individual ports, VLANs, and static MAC entries
- Layer 4 session packet-based priority
- AppleTalk sockets

See the “Quality of Service (QoS)” chapter in the *Advanced Configuration and Management Guide* for more information.

Router-Based Health Checking

The routing switches contain a Layer 4 HTTP health check, which you can use to check the health of an HTTP application on a web server. The server can be a web host that has Layer 2 connectivity to the routing switch or it can be a third-party Server Load Balancer that is configured to load balance traffic for the web site. Based on the results of the health check, the routing switch inserts a host route for the web server in its route table and advertises the route. If the HTTP application on the web server fails the health check, the routing switch removes the host route from its route table.

You can use this feature to support web load balancing implementations, in which the same HTTP site is configured on multiple servers. The health check helps ensure that clients are directed to the closest instance of the web site.

See the “Route Health Injection” chapter in the *Advanced Configuration and Management Guide* for more information about this feature.

Virtual Router Redundancy Protocol (VRRP)

The **Virtual Router Redundancy Protocol (VRRP)**, described in RFC 2338, allows routing switches to be configured together as a virtual router. Generally, a host configured to use a default router will lose its connection to the rest of the network if the default router becomes unavailable. However, if you configure several routers as a VRRP virtual router, and then use the virtual router as the default router for the hosts, the hosts receive uninterrupted service even if one of the routers within the virtual router becomes unavailable.

One of the routers in the virtual router is the “active” or “master” router and handles the traffic sent to the virtual router’s MAC address or IP address. The other routers remain in standby mode while the active router is functioning.

If the active router becomes unavailable, one of the standby routers becomes the new active router. The new active router uses the same virtual MAC address and virtual IP address as the previous master, so hosts are unaware that a router has become unavailable. As far as the hosts are concerned, the MAC address and IP address of the virtual router is still alive. You can fix the link or router problem off-line while network service continues uninterrupted.

In addition to the standard redundancy support described in RFC 2338, the implementation of VRRP on the HP 9308M, HP 9304M, and HP 6308M-SX routing switches enables you to track the status of both the in and out ports for host traffic. The track port feature ensures that if an out port goes down, even if the in port is still up, VRRP lowers the router priority and thus causes a renegotiation for the Master.

For more details on VRRP and its configuration, see the “Configuring VRRP” chapter in the *Advanced Configuration and Management Guide*.

Server Redundancy Protocol (SRP)

In addition to VRRP, the routing switches continue to provide support for redundant routing switch configurations through **Standby Router Protocol (SRP)**. SRP provides many of the same features as VRRP, but SRP can be used only with the HP 9308M, HP 9304M, and HP 6308M-SX routing switches or with third-party routers that support Cisco Systems' Hot Standby Router Protocol (HSRP).

The routing switch software continues to provide SRP support for backward compatibility on routing switches that are already configured to use the protocol. If you have routing switches that are running SRP, you do not need to reconfigure them for VRRP. However, if you are planning to configure your routing switches to use a redundancy protocol, HP recommends that you use VRRP. Using VRRP allows you to include third-party routers in the virtual router.

NOTE: HP recommends that you do not use VRRP and SRP on the same device.

NOTE: The virtual interface feature might not be supported on third-party routers. See the documentation for those routers for information.

For more details on SRP and its configuration, see the "Configuring SRP" chapter in the *Advanced Configuration and Management Guide*.

