Chapter 8
Configuring Basic Features

This chapter describes how to configure basic, non-protocol features on the HP 9304M, 9308M, and 6308M-SX routing switches, and on the 6208M-SX switch.

A summary of all CLI commands (including syntax) described in this chapter can be found in “Command Line Interface Commands” on page B-1.

This chapter contains procedures for configuring the following parameters:

- Basic System Parameters – see “Configuring Basic System Parameters” on page 8-3
- Basic Port Parameters – see “Configuring Basic Port Parameters” on page 8-23
- Basic Layer 2 Parameters – see “Configuring Basic Layer 2 Parameters” on page 8-30
- Basic Layer 3 Parameters – see “Configuring Basic Layer 3 Parameters” on page 8-60
- Layer 4 Quality-of-Service (QoS) parameters – see “Configuring Layer 4 Quality of Service Parameters” on page 8-61
- System defaults – see “Modifying System Parameter Default Settings” on page 8-69
- Mirror ports (for traffic diagnosis and troubleshooting) – see “Assigning a Mirror Port and a Monitor Port” on page 8-72

The HP 9304M, 9308M, and 6308M-SX routing switches, and on the 6208M-SX switch are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the advanced features such as VLANs or routing protocols for the routing switch must first be enabled at the system (global) level before they can be configured.

- If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.
- If you use the Web management interface, you configure the system level parameters on the System configuration sheet by selecting the System link on the main menu. Figure 8.1 shows an example of the System configuration sheet on a 9304M or 9308M routing switch.

NOTE: Before assigning or modifying any routing switch parameters, you must assign the IP sub-net (interface) addresses for each port.

NOTE: This chapter does not describe how to configure Virtual LANs (VLANs). For VLAN configuration information, see “Configuring VLANs” on page 17-1.

Using the Web Management Interface for Basic Configuration Changes

The Web management interface enables you to easily make numerous configuration changes by entering or changing information on configuration panels such as the one shown in Figure 8.1. This example is for a routing switch. The 6208M-SX switch does not have routing options but does have some additional options not available on routing switches.

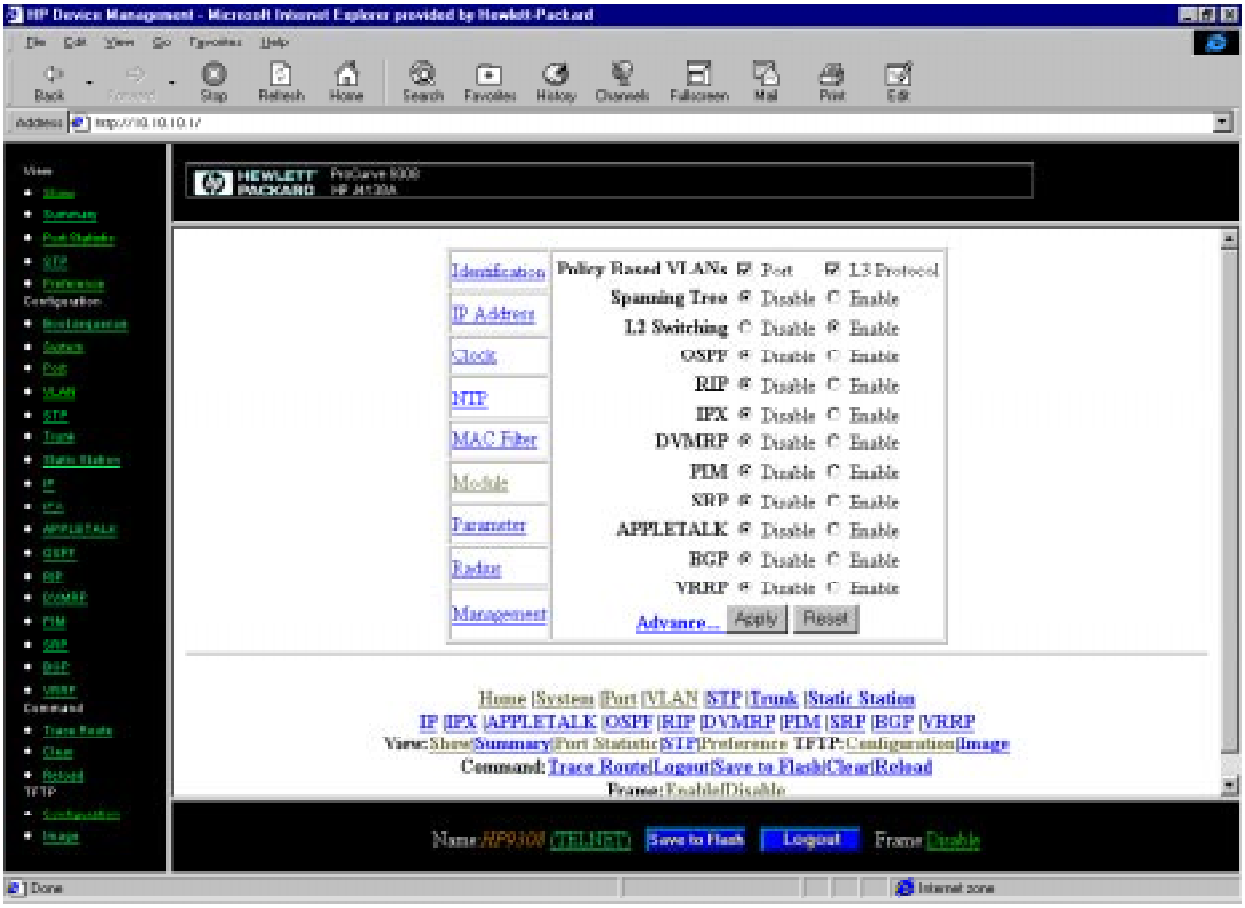

Figure 8.1 System configuration sheet for routing switch

You can perform the following configuration tasks from the System configuration sheet:

- Enter system administration information.
- Review or modify the IP, mask, and gateway addresses (6208M-SX switch only).
- Assign IP sub-net (interface) addresses and masks (routing switches only).
- Assign DHCP gateway lists for DHCP Assist operation (6208M-SX switch only).
- Configure Domain Name Server (DNS) Resolver.
- Assign a SNMP trap receiver station to collect traps.
- Modify SNMP traps generated.
- Modify the community string.
- Modify SNMP traps saved to local event log.
- Define a MAC address filter.

- Set the system clock.
- Establish a reference network time protocol (NTP) server.
- Enable port-based and/or layer 3 protocol VLANs.
- Enable or disable IP Multicast Traffic Reduction (6208M-SX switch only).
- Enable or disable IGMP (6208M-SX switch only).
- Enable or disable protocol—OPSF, IP/RIP, IPX, DVMRP, PIM, SRP, VRRP, BGP4, AppleTalk (routing switches only).
- Assign Layer 4 QoS Priority (6208M-SX switch only).

NOTE: Layer 4 priority for routing switches is set using the IP policy command found at the global CONFIG level of the CLI and the IP configuration sheet for the Web management interface.

- Enable or disable Spanning Tree Protocol.
- Enable or disable SNMP operation.
- Enable or disable IEEE 802.1q VLAN tagging.
- Enable or disable layer 2 switching (routing switches only).
- Enable or disable Telnet.
- Change the aging period (switch age time) for entries in the address table.
- Assign a mirror port.
- Modify system parameters.
- Add or delete module (chassis only).
- Modify tag type.
- Modify telnet timeout period.
- Modify broadcast limit.
- Enable or disable management using the Web management interface.
- Apply base (system) default values (6208M-SX switch only).

The procedures in this chapter describe how to configure these parameters.

Configuring Basic System Parameters

The procedures in this section describe how to configure the following basic system parameters:

- System name, contact, and location – see “Entering System Administration Information” on page 8-4
- IP addresses – see “Configuring the IP Address Information” on page 8-4
- Domain Name System (DNS) resolver – see “Enabling Domain Name Server (DNS) Resolver” on page 8-6
- SNMP trap server and individual traps – see “Configuring Simple Network Management (SNMP) Parameters” on page 8-8
- System time using a Simple Network Time Protocol (SNTP) server or local system counter – see “Specifying a Simple Network Time Protocol (SNTP) Server” on page 8-11 and “Setting the System Clock” on page 8-13
- Syslog server and local syslog buffer parameters – see “Configuring the Syslog Service” on page 8-15
- Default Gigabit negotiation mode (for chassis devices) – “Changing the Default Gigabit Negotiation Mode” on page 8-20
- Broadcast, multicast, or unknown-unicast limits, if required to support slower third-party devices – see “Limiting Broadcast, Multicast, or Unknown-Unicast Rates” on page 8-21

Entering System Administration Information

You can configure a system name, contact, and location for a device and save the information locally in the configuration file for future reference. This information is not required for system operation but is suggested. When you configure a system name, the name replaces the default system name in the CLI command prompt.

The name, contact, and location each can be up to 32 alphanumeric characters.

USING THE CLI

Here is an example of how to configure a switch or routing switch name, system contact, and location:

```
HP9300(config)# hostname oakland
HP9300(config)# snmp-server contact jack london
HP9300(config)# snmp-server location oakcabldg519
HP9300(config)# end
HP9300# write memory
```

syntax: hostname <string>

syntax: snmp-server contact <string>

syntax: snmp-server location <string>

USING THE WEB MANAGEMENT INTERFACE

Here is an example of how to configure a switch or routing switch name, system contact, and location:

1. Select the System link from the main menu.
2. Select Identification to display the Identification panel.
3. Enter system name, contact, and location information.
4. Click the Apply button to save the changes.

Configuring the IP Address Information

To support management of the switch or routing switch by Telnet or an SNMP station, you must define an IP address and network mask. In addition, for the 6208M-SX switch, you must define a default gateway.

6208M-SX Switch

To support management of the switch using Telnet or an SNMP station, you must define an IP address, mask, and gateway for the switch. Network management applications use SNMP, so you cannot access the switch using a network management application until you have configured the IP information.

NOTE: The Web management interface also uses SNMP, but the SNMP packets are encapsulated in HTTP packets.

USING THE CLI

To assign an IP address, mask, and gateway to a switch to support Telnet and SNMP management:

```
HP6208(config)# ip address 192.22.3.44 255.255.255.0
HP6208(config)# ip default-gateway 192.22.33.100
```

Syntax: enable [<password>]

Syntax: configure terminal

Syntax: ip address <ip-addr> <mask>

or

ip address <ip-addr>/<mask-bits>

Syntax: ip default-gateway <ip address>

NOTE: You can use the syntax, **ip address <ip address /sub-net mask length>** if you know the sub-net mask length. In the above example, you could enter **ip address 192.22.3.44/24**.

USING THE WEB MANAGEMENT INTERFACE

You need a direct serial connection to the Console port to configure the switch's IP address, sub-net mask, and default gateway. After you configure this information, you can view or modify the information on the Web management interface using the System configuration sheet.

To modify the IP address, mask, and gateway for a switch:

1. Select the System link from the main menu.
2. Select the IP Address link.
3. Enter the IP address and network mask.
4. Enter the default gateway if applicable.
5. Click the Apply button to assign the changes.
6. Select the Save To Flash link below the panel to save the new IP address in the system-config file. Select Yes when prompted.

Routing Switches

Before attaching equipment or a management station to ports on the routing switch, you must assign individual sub-net IP addresses and masks to each of the ports. By default no IP addresses are assigned. You can assign up to 24 IP addresses to each routing switch port, loopback interface, and virtual interface.

USING THE CLI

To assign an IP address and mask to a routing switch interface:

```
HP9300(config)# int e 1/5
HP9300(config-if-1/5)# ip address 192.22.3.44 255.255.255.0
```

Syntax: enable [<password>]

Syntax: configure terminal

Syntax: ip address <ip-addr> <mask> [secondary]

or

ip address <ip-addr>/<mask-bits> [secondary]

Use the **secondary** parameter if you have already configured an IP address in the same sub-net on the interface.

NOTE: You can use the syntax **ip address <ip-addr>/<network-mask-length>** if you know the sub-net mask length. In the above example, you could enter **ip address 192.22.3.44/24**.

USING THE WEB MANAGEMENT INTERFACE

To assign an IP address and mask to a routing switch interface:

1. Select the IP Address link from the System configuration sheet. (To display this sheet, select System from the menu underneath or on the left side of the chassis window.)
2. If the routing switch already has IP addresses, the addresses are listed. Select Add IP Address to display the Router IP Address panel.

   NOTE: The Loopback link lets you add a loopback interface. Loopback interfaces can be quite useful for exchanging routing information. See "Adding a Loopback Interface" on page 12-12.

3. Select the interface from the Slot and Port field's pulldown menus.
4. Enter the IP address.

5. Enter the network mask.
6. Select the Secondary option if this is not the first IP address assigned to this interface.
7. Click the Add button to assign the address to the interface.
8. Select the Save To Flash link below the panel to save the new IP address in the system-config file. Select Yes when prompted.

Enabling Domain Name Server (DNS) Resolver

The Domain Name Server (DNS) resolver feature lets you use a host name to perform Telnet, ping, and traceroute commands. You can also define a DNS domain on a switch or routing switch and thereby recognize all hosts within that domain. After you define a domain name, the switch or routing switch automatically appends the appropriate domain to the host and forwards it to the domain name server.

For example, if the domain "newyork.com" is defined on a switch or routing switch and you want to initiate a ping to host "NYC01" on that domain, you need to reference only the host name in the command instead of the host name and its domain name. For example, you could enter either of the following commands to initiate the ping:

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

Defining a DNS Entry

You can define up to four DNS servers for each DNS entry. The first entry serves as the primary default address. If a query to the primary address fails to be resolved after three attempts, the next gateway address is queried (also up to three times). This process continues for each defined gateway address until the query is resolved. The order in which the default gateway addresses are polled is the same as the order in which you enter them.

USING THE CLI

Suppose you want to define the domain name of newyork.com on a switch and then define four possible default DNS gateway addresses. To do so, enter the following commands:

```
HP9300(config)# ip dns domain-name newyork.com
```

```
HP9300(config)# ip dns server-address 209.157.22.199 205.96.7.15 208.95.7.25
201.98.7.15
```

Syntax: ip dns server-address <ip-addr> [<ip-addr>] [<ip-addr>] [<ip-addr>]

In this example, the first IP address in the **ip dns server-address...** command becomes the primary gateway address and all others are secondary addresses. Because IP address 201.98.7.15 is the last address listed, it is also the last address consulted to resolve a query.

USING THE WEB MANAGEMENT INTERFACE

To map a domain name server to multiple IP addresses:

1. On a switch – Select the System link from the main menu to display the System configuration sheet, then select the DNS link to display the panel shown in Figure 8.2.
2. On a routing switch – Select the IP link from the main menu to display the IP configuration sheet, then select the DNS link to display the panel shown in Figure 8.2.
3. Enter the name of the domain name server in the Domain Name field.
4. Enter an IP address for each system that will serve as a gateway to the domain name server.

NOTE: The first address entered will be the primary DNS gateway address. The other addresses will be used in chronological order, left to right, should the primary address not be available.

---

5. Click the Apply button after you have entered all the gateway addresses to assign the changes.

6. Select the Save To Flash link below the panel to save the new IP address in the system-config file. Select Yes when prompted.

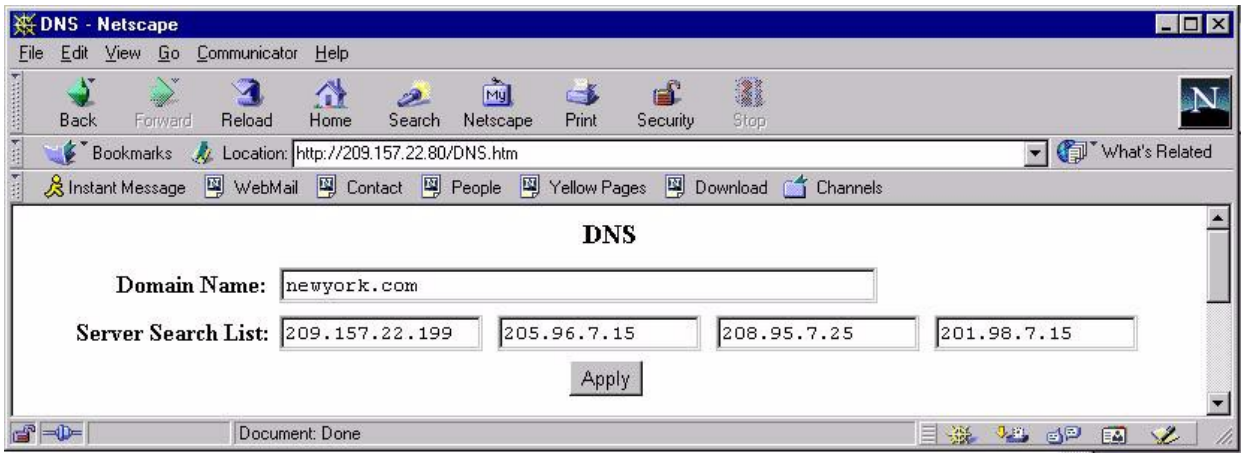

Figure 8.2 DNS resolver configuration panel

Using a DNS Name To Initiate a Trace Route

EXAMPLE: Suppose you want to trace the route from a 9304M to a remote server identified as NYC02 on domain newyork.com. Because the newyork.com domain is already defined on the switch, you need to enter only the host name, NYC02, as noted below.

USING THE CLI

```
HP9300# traceroute nyc02
```

Syntax: trace-route <IP-addr> [minttl <value> maxttl <value> timeout <value>]

The only required parameter is the IP address of the host at the other end of the route. See “traceroute” on page B-68 for information about the parameters.

After you enter the command, a message indicating that the DNS query is in process and the current gateway address (IP address of the domain name server) being queried appear on the screen:

```
Type Control-c to abort
Sending DNS Query to 209.157.22.199
Tracing Route to IP node 209.157.22.80
To ABORT Trace Route, Please use stop-traceroute command.
 Traced route to target IP node 209.157.22.80:
   IP Address          Round Trip Time1    Round Trip Time2
  207.95.6.30         93 msec             121 msec
```

NOTE: In the above example, 209.157.22.199 is the IP address of the domain name server (default DNS gateway address), and 209.157.22.80 represents the IP address of the NYC02 host.

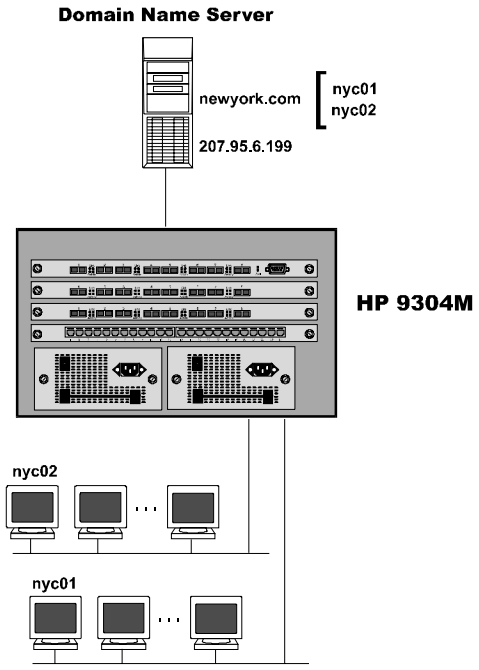

Figure 8.3 Querying a host on the newyork.com domain

USING THE WEB MANAGEMENT INTERFACE

1. Select the Trace Route link from the menu beneath the currently displayed panel. The Trace Route panel is displayed.
2. Enter the host name or IP address in the Target Address field.

   NOTE: You can use the host name only if you have already configured the DNS resolver for the domain that contains the host.

3. Optionally change the minimum and maximum TTLs and the Timeout.
4. Select Start to begin the trace. The trace results are displayed below the Start and Abort buttons.

Configuring Simple Network Management (SNMP) Parameters

Use the procedures in this section to perform the following configuration tasks:

- Specify an SNMP trap receiver.
- Disable individual SNMP traps. (All traps are enabled by default.)

NOTE: To add and modify "get" (read-only) and "set" (read-write) community strings, see “Configuring the SNMP Community Strings” on page 2-26.

Specifying an SNMP Station to Collect Traps

To activate a station as a trap receiver, enter the IP address of the target SNMP station. You can assign up to ten stations to operate simultaneously as trap receivers. After assigning the SNMP station IP address, enter the community string. By default, no community string values are assigned.

The device also stores up to 100 traps in a local buffer. The traps are cleared if you reload the system. By default, the buffer can hold up to 50 traps. To change the capacity of the trap buffer, see “Changing the Number of Entries the Local Buffer Can Hold” on page 8-16.

USING THE CLI

To specify a trap receiver, enter a command such as the following:

```
HP9300(config)# snmp-server trap-receiver 192.22.3.33 public
```

syntax: [no] snmp-server trap-receiver <ip address> <communitystring>

NOTE: In the above example, “public” refers to the community string.

USING THE WEB MANAGEMENT INTERFACE

To assign a station to act as trap receiver:

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Management link.
3. Select the Trap Receiver link. The panel shown in Figure 8.4 will appear.
4. Enter the IP address of the station that is to serve as the trap receiver.
5. Enter the community string.

   NOTE: The community string is initially set with the CLI. The value you enter here must match the value you configure using the CLI.

6. Click Add to assign the changes.

Figure 8.4 Configuration panel for trap receiver

Disabling SNMP Traps

The HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch come with SNMP trap generation enabled by default for all traps. You can selectively disable one or more of the following traps.

NOTE: By default, all SNMP traps are enabled at system startup.

Switch Traps

The following traps are generated on the switch:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- Module insert (applies only to chassis devices)
- Module remove (applies only to chassis devices)

Routing Switch Traps

The following traps are generated on the routing switches:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- Module insert
- Module remove
- OSPF
- SRP
- VRRP

USING THE CLI

To stop link down occurrences from being reported, enter the following:

```
HP9300(config)# no snmp-server trap link-down
```

Syntax: [no] snmp-server trap <trap-type>

NOTE: For a list of the trap type values, see “snmp-server trap” on page B-123.

USING THE WEB MANAGEMENT INTERFACE

To modify the SNMP traps generated for a system:

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Management link.
3. Select the Trap link. The resulting panel shows the current state of all traps.

NOTE: The panel lists different traps for the switch and for routing switches.

4. Select the Disable or Enable button next to the trap you want to disable or enable.
5. Click Apply to save the changes.

Specifying a Simple Network Time Protocol (SNTP) Server

The HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch can consult SNTP servers for the current system time and date.

NOTE: The devices do not retain time and date information across power cycles. Unless you want to reconfigure the system time counter each time the system is reset, HP recommends that you use the SNTP feature.

USING THE CLI

To identify an SNTP server with IP address 208.99.8.95 to act as the clock reference for a switch or routing switch, enter the following:

```
HP9300(config)# sntp server
```

Syntax: sntp server <ip-addr> | <hostname>

By default, the devices do not change the system time for daylight savings time. To enable daylight savings time, enter the following command:

```
HP9300# clock summer-time
```

Syntax: clock summer-time

USING THE WEB MANAGEMENT INTERFACE

To identify a reference SNTP server for the system:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the NTP link from the System configuration sheet. The panel shown in Figure 8.5 appears.
3. Optionally change the polling time by editing the value in the Polling Time field, then click Apply to implement the change.
4. Select the NTP Server link to display the panel shown in Figure 8.6.

NOTE: If you have already configured an SNTP server, the server information is listed. Select the Add NTP Server link at the bottom of the panel to display the panel shown in Figure 8.6.

5. Enter the IP address of the SNTP server.
6. Select the SNTP version the server is running from the version field's pulldown menu.
7. Click Add to assign the changes.

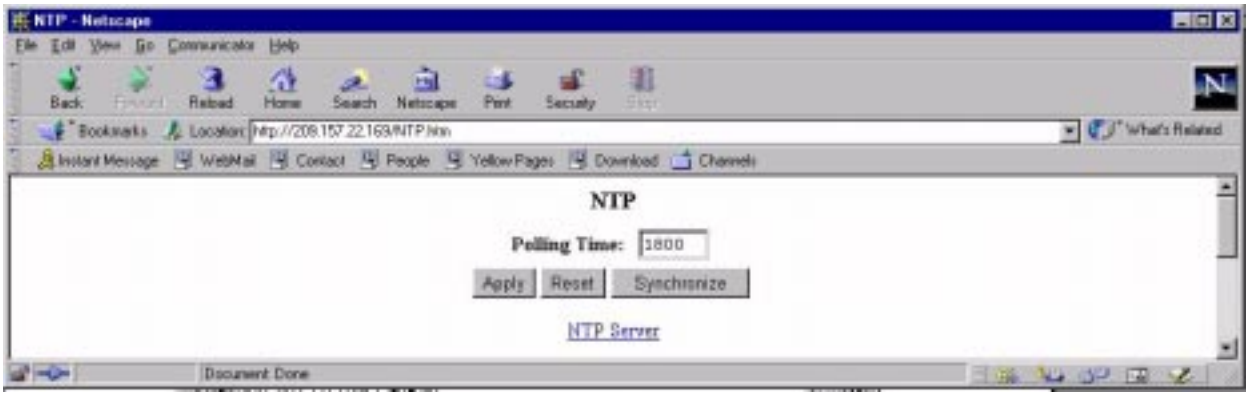

Figure 8.5 NTP panel

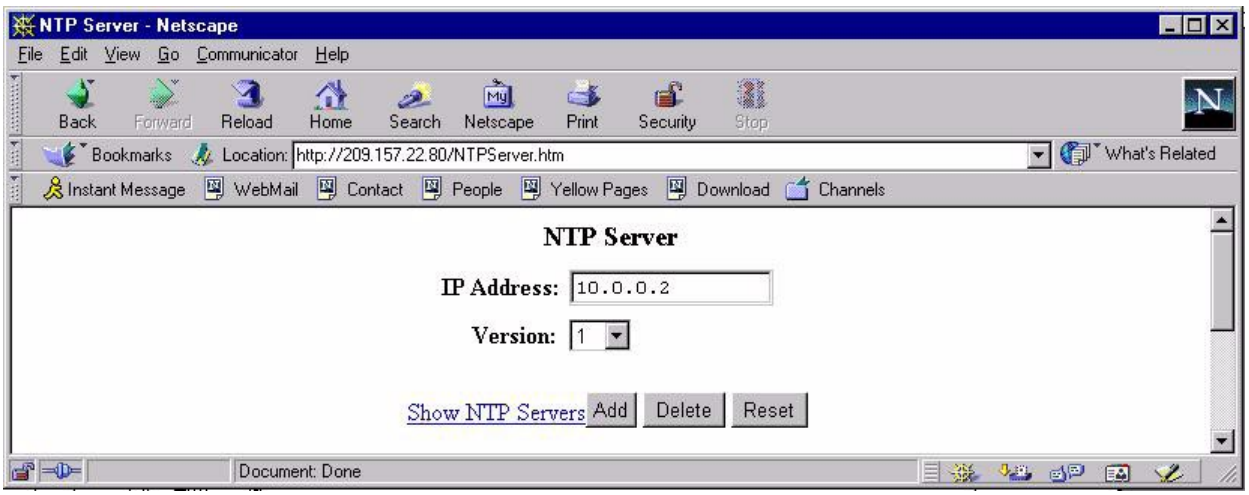

Figure 8.6 NTP Server panel

Changing the Time Zone

Although SNTP servers typically deliver the time and date in Greenwich Mean Time (GMT), you can configure the switch or routing switch to adjust the time for any one-hour offset from GMT or for one of the following U.S. time zones:

- US Pacific
- Alaska
- Aleutian
- Arizona
- Central
- East-Indiana
- Eastern
- Hawaii
- Michigan
- Mountain

- Pacific
- Samoa

The default is US Pacific.

USING THE CLI

To change the time zone to local Australian time (which is normally 10 hours ahead of Pacific standard time), enter the following command:

```
HP9300(config)# clock time-zone gmt+10
```

Syntax: clock time-zone <gmt>|<us> <time-zone>

You can enter one of the following values for <gmt> or <us>:

- US time zones: alaska, aleutian, arizona, central, east-indiana, eastern, hawaii, michigan, mountain, pacific, samoa.
- GMT time zones: gmt+12, gmt+11, gmt+10...fmt+01, gmt+00, gmt-01...gmt-10, gmt-11, gmt-12.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Clock link from the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the offset from GMT from the Time Zone field's pulldown menu.
3. Click Apply to assign the changes.

Changing the SNTP Poll Interval

By default, the switch or routing switch polls its SNTP server every 30 minutes (1800 seconds).

USING THE CLI

To configure the switch or routing switch to poll for clock updates from a SNTP server every 15 minutes, enter the following:

```
HP9300(config)# sntp poll-interval 900
```

Syntax: [no] sntp poll-interval <1-65535>

USING THE WEB MANAGEMENT INTERFACE

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the NTP link from the System configuration sheet. The panel shown in Figure 8.5 appears.
3. Change the value in the Polling Time field.
4. Click Apply button to assign the change.

Setting the System Clock

In addition to SNTP support, the HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch also allow you to set the system time counter. The time counter setting is not retained across power cycles and is not automatically synchronized with an SNTP server. The counter merely starts the system time and date clock with the time and date you specify.

NOTE: You can synchronize the time counter with your SNTP server time by entering the **sntp sync** command from the Privileged EXEC level of the CLI. You cannot perform this procedure using the Web management interface.

NOTE: Unless you identify an SNTP server for the system time and date, you will need to re-enter the time and date following each reboot.

For more details about SNTP, see “Specifying a Simple Network Time Protocol (SNTP) Server” on page 8-11.

USING THE CLI

To set the system time and date to 10:15:05 on October 15, 1998, enter the following command:

```
HP9300# clock set 10:15:05 10-15-98
```

Syntax: [no] clock set <hh:mm:ss> <mm-dd-yy> | <mm-dd-yyyy>

By default, the devices do not change the system time for daylight savings time. To enable daylight savings time, enter the following command:

```
HP9300# clock summer-time
```

Syntax: clock summer-time

To change the time zone from the default (Greenwich Mean Time), see “Changing the Time Zone” on page 8-12.

USING THE WEB MANAGEMENT INTERFACE

To set the system time counter for a device:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Clock link from the System configuration sheet. The panel shown in Figure 8.7 appears.
3. Select a time zone from the pulldown menu.
4. Select Enable or Disable to enable or disable daylight savings time.
5. Enter the current date.
6. Enter the current time.
7. Click Apply to apply the change.

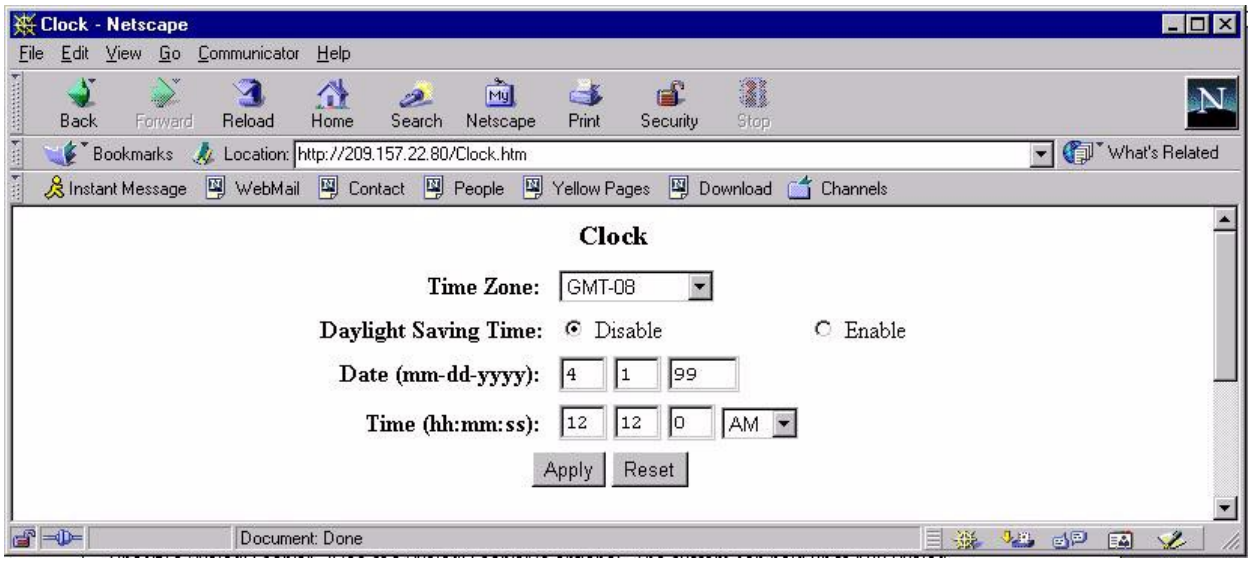

Figure 8.7 Setting the time counter

Configuring the Syslog Service

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a SyslogD server. You can configure a device to use one or two SyslogD servers. (Use of a SyslogD server is optional. The system can hold up to 100 Syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local Syslog buffer can hold.
- Display the Syslog configuration.
- Clear the local Syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No SyslogD server is specified.

Syslog Overview

The device's software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer that can hold up to 100 messages. You also can specify the IP address or host name of one or two SyslogD servers. When you specify a SyslogD server, the device writes the messages both to the system log and to the SyslogD server.

Using a SyslogD server ensures that the messages remain available even after a system reload. The device's local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the SyslogD server remain on the server.

The SyslogD service on a Syslog server receives logging messages from applications on the local host or from devices such as a routing switch or switch. SyslogD adds a time stamp to each received message and directs messages to a log file. Most Unix workstations come with SyslogD configured. Some third party vendor products also provide SyslogD running on NT.

SyslogD uses UDP port 514 and each SyslogD message thus is sent with destination port 514. Each SyslogD message is one line with syslogd message format. The message is embedded in the text portion of the SyslogD format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfield order is insensitive except that the text subfield should be the last field in the message. All the subfields are optional.

To enable Syslog logging and set logging parameters, use one of the following methods.

USING THE CLI

To enable Syslog parameters using the CLI, enter the following commands at the global CONFIG level:

```
HP9300(config)# logging on
```

Syntax*:* logging on | off

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No SyslogD server is specified.

Specifying a SyslogD Server

To specify a SyslogD server, enter a command such as the following:

```
HP9300(config)# logging 10.0.0.99
```

Syntax: logging <ip-addr> | <server-name>

NOTE: You can specify a server name only if you have already configured the DNS Resolver feature. See "Enabling Domain Name Server (DNS) Resolver" on page 8-6.

Specifying a Second SyslogD Server

To specify a second SyslogD server, enter the logging <IP-addr> command again, as in the following example:

```
HP9300(config)# logging 10.0.0.69
```

Syntax: logging <ip-addr> | <server-name>

NOTE: If you accidentally try to configure a third SyslogD server, the software displays an error message.

Disabling Logging of a Message Level

To change the message level, disable logging of specific message levels. You must disable the message levels on an individual basis. For example, to disable logging of debugging and informational messages, enter the following commands:

```
HP9300(config)# no logging buffered debugging
HP9300(config)# no logging buffered informational
```

Syntax: [no] logging buffered <level>| <num-entries>

The <level> parameter can have one of the following values:

- **alerts**
- **critical**
- **debugging**
- **emergencies**
- **errors**
- **informational**
- **notifications**
- **warnings**

The commands in the example above change the log level to notification messages or higher. The software will not log informational or debugging messages. The changed message level also applies to the SyslogD servers.

Changing the Number of Entries the Local Buffer Can Hold

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store. For example:

```
HP9300(config)# logging buffered 100
```

The default number of messages is 50. The value can be 50 – 100. The change takes effect after you restart the system.

Changing the Log Facility

The SyslogD daemon on the SyslogD server uses a facility to determine where to log the messages from the device. The default facility for messages the device sends to the SyslogD server is "user". You can change the facility using the following command.

NOTE: You can specify only one facility. If you configure the device to use two SyslogD servers, the device uses the same facility on both servers.

```
HP9300(config)# logging facility local0
```

***Syntax**:* logging facility <facility-name>

The <facility-name> can be one of the following:

- kern – kernel messages
- user – random user-level messages
- mail – mail system
- daemon – system daemons
- auth – security/authorization messages
- syslog – messages generated internally by syslogd
- lpr – line printer subsystem
- news – netnews subsystem
- uucp – uucp subsystem
- sys9 – cron/at subsystem
- sys10 – reserved for system use
- sys11 – reserved for system use
- sys12 – reserved for system use
- sys13 – reserved for system use
- sys14 – reserved for system use
- cron – cron/at subsystem
- local0 – reserved for local use
- local1 – reserved for local use
- local2 – reserved for local use
- local3 – reserved for local use
- local4 – reserved for local use
- local5 – reserved for local use
- local6 – reserved for local use
- local7 – reserved for local use

Displaying the Syslog Configuration

To display the Syslog parameters currently in effect on a device, enter the following command from any level of the CLI:

```
HP9300> show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 2 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Log Buffer (50 entries):
at 0 days 0 hours 1 minutes 0 seconds, level informational
Interface Ethernet1/6, changed state to up
at 0 days 0 hours 0 minutes 0 seconds, level informational
Cold start
```

***Syntax**:* show logging

The value "ACDMEINW" indicates message levels that are enabled. Each letter represents a message type and is identified by the key below the value. This example shows two log entries, an informational message stating that Ethernet interface 1/6 came up and another informational message stating that the system was cold started. The system time for each event is listed. (In this example, the system does not use an SNTP server and the local system clock has not been set.)

USING THE WEB MANAGEMENT INTERFACE

To configure Syslog parameters using the Web management interface, use the following procedure:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select Management from the System configuration sheet to display the Management panel.
3. Select System Log to display the panel shown in Figure 8.8.
4. Enable or disable logging. This setting affects both logging to the local buffer and logging to a SyslogD server if you configure the system to use one.
5. Optionally change the number of entries the local Syslog buffer can hold. The buffer size can be from 50 – 100. The default is 50.

   NOTE: A change in the buffer size takes effect only after you restart the system. The buffer size does not affect how many entries the device can log on a SyslogD server. The number of entries the device can log on the server depends on the server's configuration.

6. Enter the IP address of your SyslogD server, if you want the device to log messages on the SyslogD server as well as in the local buffer.
7. Select the messages facility. The default is User. For list of values, display the pulldown menu or see "Changing the Log Facility" on page 8-17.
8. Select the message levels you want the device to log.
9. Click Apply to assign the changes.

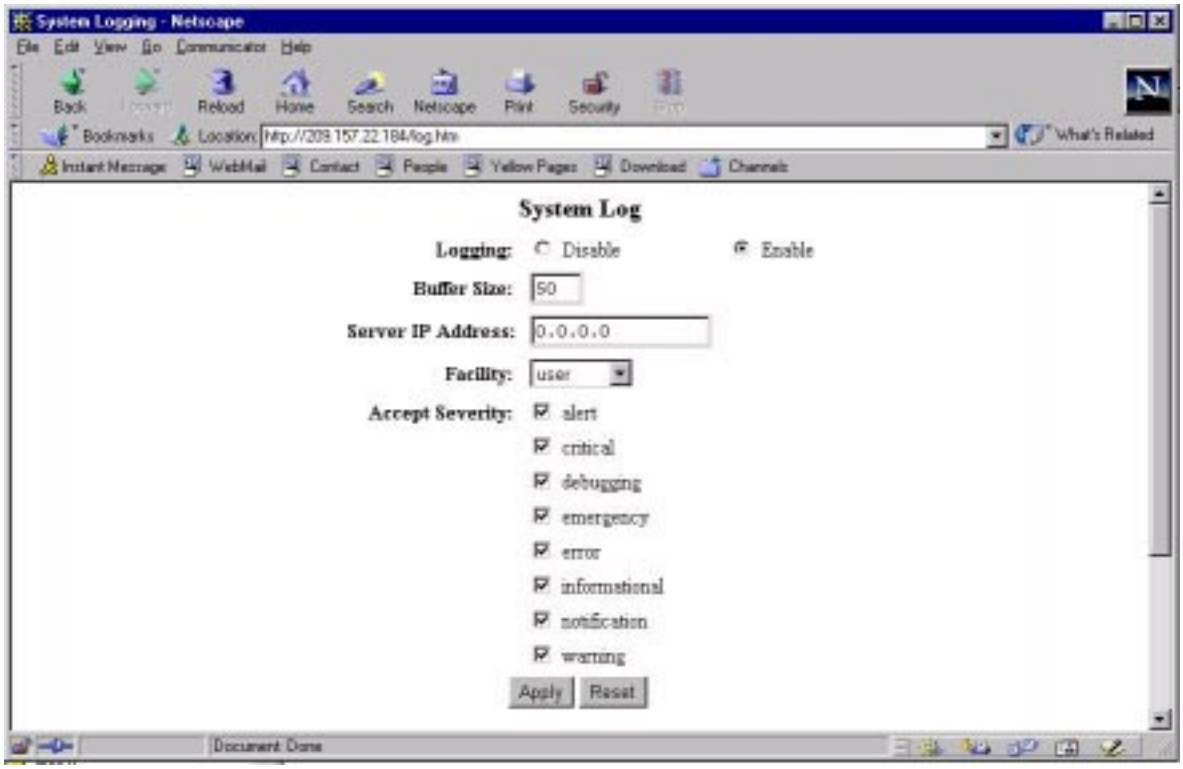

Figure 8.8 System Log configuration panel

Clearing the Syslog Messages from the Local Buffer

To clear the Syslog messages stored in the device's local buffer, enter the following command from the Privileged EXEC level the CLI:

```
HP9300# clear logging
```

Syntax: clear logging

USING THE WEB MANAGEMENT INTERFACE

To clear Syslog messages using the Web management interface, use the following procedure:

1. Select Clear to display the Clear panel. The Clear link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the checkbox next to System Logging to place a checkmark in the box.
3. Click Apply to clear the log.

Changing the Default Gigabit Negotiation Mode

You can configure the default Gigabit negotiation mode to be one of the following:

- Negotiate-full-auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default.
- Auto-Gigabit – The port tries to perform a handshake with the other port to exchange capability information.
- Negotiation-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

Although the standard for 100Base-Tx ports provides an option for a negotiating port to link with a non-negotiating port, the 802.3x standard for Gigabit ports does not provide this option. As a result, unless the ports at both ends of a Gigabit Ethernet link use the same mode (either auto-Gigabit or negotiation-off), the ports cannot establish a link. An administrator must intervene to manually configure one or both sides of the link to enable the ports to establish the link.

The HP 9304M, 9308M, 6308M-SX, and 6208M-SX software provides a solution by changing the default negotiation behavior for Gigabit Ethernet ports from what the behavior was in earlier software releases. The new default behavior allows a port to establish a link with another port whether the other port is configured for auto-Gigabit or negotiation-off. By default, Gigabit Ethernet ports first attempt auto-Gigabit. If auto-Gigabit does not succeed (typically because the port at the other end is not configured for auto-Gigabit), the port switches to negotiation-off.

Backward Compatibility

When you upgrade a chassis device that is running software older than 05.2.x, the new software makes modifications to the running-config and startup-config files to ensure that the negotiation settings remain unchanged for the installed device. For new devices running 05.2.x or later, the default for all Gigabit Ethernet ports is negotiate-full-auto.

To provide the backward compatibility, the software places a line in the running-config file to identify the software version that generated the file. For software release 05.2.x, the version line is as follows: “version 05.2.*x*”, where *x* is the specific version. When you save configuration changes to the startup-config file, the software assumes, based on the presence of the version line in the running-config file, that the device is running software release 05.2.x or later, which contains the change to the Gigabit Ethernet negotiation default.

If the device already has a startup-config file when you update to software release 05.2.x, the software adds the following command to the startup-config file: **gig-default neg-off**. This command sets the global negotiation mode to negotiation-off, the default behavior in software releases earlier than 05.2.x. By setting the default mode to negotiation-off, the new software ensures that the device’s Gigabit Ethernet links continue to operate as before. (Although you cannot set a global default for Gigabit Ethernet negotiation in software releases earlier than 05.2.x, the implicit default behavior is negotiation-off.)

If the startup-config file contains the **auto-gig** command to configure individual ports for auto-Gigabit, the command is changed to the new format, **gig-default auto-gig**. Thus, the ports continue to use the auto-Gigabit setting.

Changing the Negotiation Mode

You can change the negotiation mode globally and for individual ports. Use either of the following methods.

USING THE CLI

To change the mode globally, enter a command such as the following:

```
HP9300(config)# gig-default neg-off
```

This command changes the global setting to negotiation-off. The global setting applies to all Gigabit Ethernet ports except those for which you set a different negotiation mode on the port level.

To change the mode for individual ports, enter commands such as the following:

```
HP9300(config)# int ethernet 4/1 to 4/4
HP9300(config-mif-4/1-4/4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gigabit for ports 4/1 – 4/4.

Syntax for global setting: gig-default neg-full-auto | auto-gig | neg-off

Syntax for individual ports: gig-default neg-full-auto | auto-gig | neg-off

USING THE WEB MANAGEMENT INTERFACE

To change the global default:

1. Select the System link to display the system configuration sheet.
2. Select the Advance link to display the advanced System parameters panel.
3. Select one of the following values from the Gig Port Default field's pulldown menu:
    - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
    - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
    - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).

- Click Apply to apply the change.
- Select the Save To Flash link, then select Yes when prompted to save the changes to the system-config file.

To override the global negotiation mode for an individual port:

1. Select the Port link to display the Port configuration sheet.
2. Click on the Modify button next to the row of information for the port.
3. Select one of the following values from the Gig Port Default field's pulldown menu:
    - Default – The port uses the negotiation mode that was set at the global level.
    - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
    - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
    - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).
4. Click Apply to apply the change.
5. Select the Save To Flash link, then select Yes when prompted to save the changes to the system-config file.

Limiting Broadcast, Multicast, or Unknown-Unicast Rates

The HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch forward all traffic at wire speed. However, some third-party networking devices cannot handle high forwarding rates for these types of packets. You can limit the number of broadcast, multicast, or unknown-unicast packets a device forwards each second using the following methods.

The limits are individually configurable for broadcasts, multicasts, and unknown-unicasts.

NOTE: By default, IP Multicast (including IGMP) is disabled. You can enable it using the **ip multicast passive|active** command. As long as IP Multicast is enabled (regardless of whether it is passive or active), no IP Multicast packets (not even IGMP packets) are limited. See “Enabling or Disabling IP Multicast Traffic Reduction (switch only)” on page 8-53.

Limiting Broadcasts

To limit the number of broadcast packets a device can forward each second, use one of the following methods.

USING THE CLI

To globally limit the number of broadcast packets a 9304M forwards to 100,000 per second, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# broadcast limit 100000
HP9300(config)# write mem
```

To limit the number of broadcast packets sent on port 1/3 to 80,000, enter the following commands:

```
HP9300(config)# int ethernet 1/3
HP9300(config-if-1/3)# broadcast limit 80000
HP9300(config-if-1/3)# write mem
```

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Limiting Multicasts

To limit the number of multicast packets a device can forward each second, use one of the following methods.

USING THE CLI

To globally limit the number of multicast packets a 9304M forwards to 120,000 per second, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# multicast limit 120000
HP9300(config)# write mem
```

To limit the number of multicast packets sent on port 3/6 to 55,000, enter the following commands:

```
HP9300(config)# int ethernet 3/6
HP9300(config-if-3/6)# multicast limit 55000
HP9300(config-if-3/6)# write mem
```

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Limiting Unknown Unicasts

To limit the number unknown unicast packets a device can forward each second, use one of the following methods.

USING THE CLI

To globally limit the number of unknown unicast packets a 9304M routing switch forwards to 110,000 per second, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# unknown-unicast limit 110000
HP9300(config)# write mem
```

To limit the number of unknown unicast packets sent on port 4/2 to 40,000, enter the following commands:

```
HP9300(config)# int ethernet 4/2
HP9300(config-if-4/2)# unknown-unicast limit 40000
HP9300(config-if-4/2)# write mem
```

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Configuring Basic Port Parameters

The procedures in this section describe how to configure the following port parameters:

- Name – see “Assigning a Port Name” on page 8-25
- Speed – see “Modifying Port Speed” on page 8-26
- Mode (half-duplex or full-duplex) – see “Modifying Port Mode” on page 8-27
- Status – see “Disabling or Re-Enabling Port Status” on page 8-27
- Flow control – see “Disabling or Re-Enabling Flow Control” on page 8-28
- Gigabit negotiate mode – see “Changing the 802.3x Gigabit Negotiation Mode” on page 8-29
- QoS priority – see “Modifying Port Priority (QoS)” on page 8-29

NOTE: To modify Layer 2, Layer 3, or Layer 4 features on a port, see the appropriate section in this chapter or other chapters. For example, to modify Spanning Tree Protocol (STP) parameters for a port, see “Modifying STP Bridge and Port Parameters” on page 8-31.

The HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch have default port values that allow the devices to be fully operational at initial startup without any additional configuration. However, in some cases, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

The current port configuration for all ports is displayed when you select the Port link from the main menu. Figure 8.9 shows an example. You can easily determine a port’s state by observing the color in the Port field.

- Red – indicates there is no link.
- Green – indicates the link is good.

This example shows the port states for a 9308M that has not yet been connected to the rest of the network.

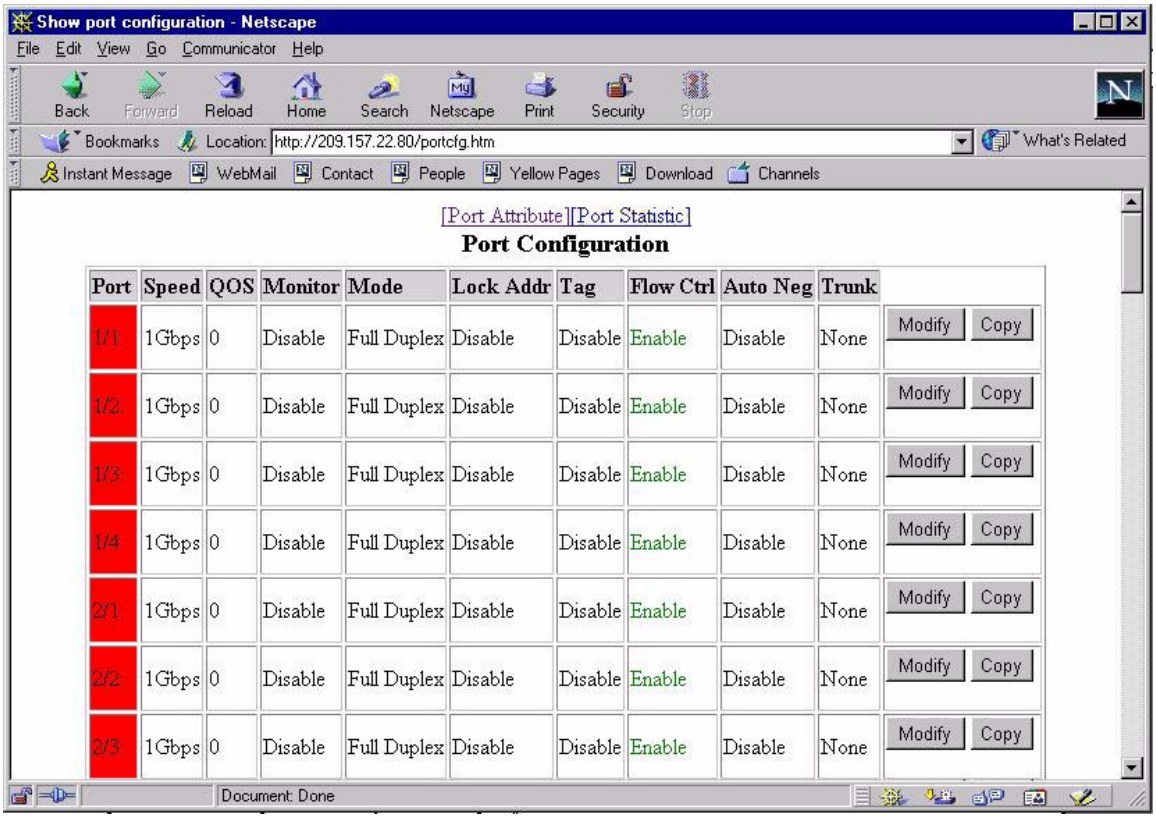

Figure 8.9 Port configuration summary sheet for a 9304M or 9308M

Click on the Copy or Modify button next to a row of port information to display a configuration panel for that port, as shown in Figure 8.10.

- Select Modify to change parameters for a port.
- Select Copy to apply a port's parameter settings to another port.

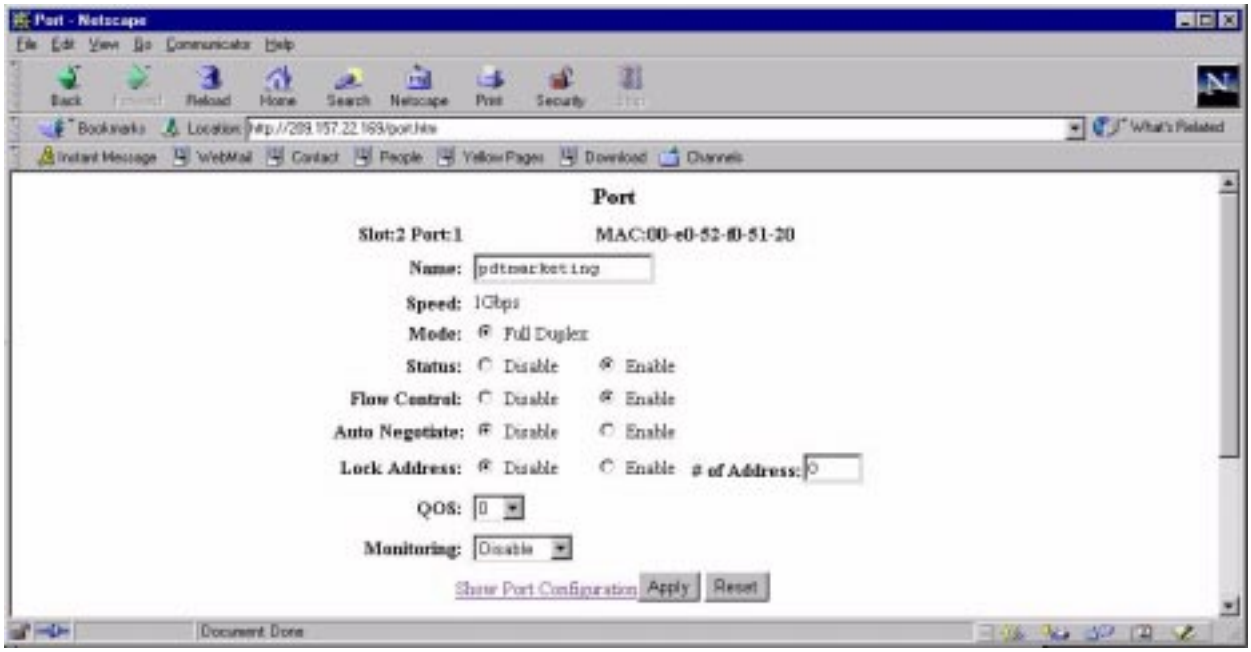

Figure 8.10 Port configuration panel for a specific port

NOTE: A slot option appears on the chassis port configuration sheet. Slot corresponds to a module slot number. See “Slot and Port Numbers” on page 6-8.

NOTE: The IEEE Tagging option appears only on the Port configuration sheet when tagging is enabled at the system level and a VLAN is defined on the system.

NOTE: The port speed option 1 Gbps is displayed only when a 1000BaseSX or 1000BaseLX Gigabit port or module is resident on the device. Additionally, only the full-duplex mode is visible. When an Ethernet port or module is being configured, the options are 10/100 Auto, 10 Mbps, and 100Mbps.

Assigning a Port Name

A port name can be assigned to help identify a segment in the network.

USING THE CLI

To assign a name to a port:

```
HP9300(config)# interface e 2/8
HP9300(config-if-2/8)# port-name pdtmarketing
```

syntax: port-name <text>

USING THE WEB MANAGEMENT INTERFACE

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click Modify next to the row of port information to display the Port panel, as shown in Figure 8.10.
3. Enter a name in the Name field.
4. Click the Apply button to assign the change.

Modifying Port Speed

Each of the 10BaseT/100BaseTX ports is designed to auto-sense and auto-negotiate the speed and mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10 Mbps or 100 Mbps. The default value for 10BaseT/100BaseTX ports is 10/100 Auto-sense.

The 100BaseFX ports operate in the full-duplex mode at 100 Mbps only and cannot be modified.

The 1000BaseSX and 1000BaseLX ports operate in the full-duplex mode at one Gigabit only and cannot be modified.

USING THE CLI

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, enter the following:

```
HP9300(config)# interface e 1/8
```

```
HP9300(config-if-1/8)# speed-duplex 10-full
```

Syntax: speed-duplex <value>

The <value> can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto

The default is auto.

USING THE WEB MANAGEMENT INTERFACE

To modify port speed:

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port you want to modify. The configuration panel shown in Figure 8.10 is displayed.
3. Click next to Full Duplex if you want to change the mode to full-duplex only.
4. Click Disable or Enable next to Auto Negotiate to enable or disable auto-negotiation.
5. Click the Apply button to assign the new configuration.

NOTE: Remember, before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to flash (using the File menu) for the changes to be preserved over a power cycle.

Modifying Port Mode

You can configure a port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic. This option is available only for 10/100 Mbps ports. The 100BaseFx, 1000BaseSx, and 1000BaseLx ports operate only at full-duplex.

USING THE CLI

Pot duplex mode and port speed are modified by the same command.

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, enter the following:

```
HP9300(config)# interface e 1/8
HP9300(config-if-1/8)# speed-duplex 10-full
```

Syntax: speed-duplex <value>

The <value> can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto

The default is auto.

USING THE WEB MANAGEMENT INTERFACE

To modify port mode:

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port you want to modify. The configuration panel shown in Figure 8.10 is displayed.
3. Click next to Full Duplex to select or de-select full duplex mode. Full-duplex mode is selected when the radio button (small circle) next to Full Duplex contains a black dot.
4. Select the Apply button to assign the new configuration.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to flash (using the File menu) for the changes to be preserved over a power cycle.

Disabling or Re-Enabling Port Status

The port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is enabled.

USING THE CLI

To disable port 1/8 on a 9308M, enter the following:

```
HP9300(config)# interface e 1/8
HP9300(config-if-1/8)# disable
```

Syntax: disable

Syntax: enable

USING THE WEB MANAGEMENT INTERFACE

To disable or enable a port:

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port you want to modify. The configuration panel shown in Figure 8.10 is displayed.
3. Select either Enable or Disable option next to the Status option.
4. Click the Apply button to assign the new configuration.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to flash (using the File menu) for the changes to be preserved over a power cycle.

Disabling or Re-Enabling Flow Control

You can configure full-duplex ports on a system to operate with or without flow control (802.3x). Flow control is enabled by default.

USING THE CLI

To disable flow control on full-duplex ports on a system, enter the following:

```
HP9300(config)# no flow-control
```

To turn the feature back on:

```
HP9300(config)# flow-control
```

Syntax: [no] flow-control

USING THE WEB MANAGEMENT INTERFACE

To disable or enable flow control on full-duplex ports on a system:

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port you want to modify. The configuration panel shown in Figure 8.10 is displayed.
3. Select either Enable or Disable next to Flow Control.
4. Click the Apply button to assign the new configuration.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to flash (using the File menu) for the changes to be preserved over a power cycle.

Changing the 802.3x Gigabit Negotiation Mode

The globally configured Gigabit negotiation mode for 802.3x flow control is the default mode for all Gigabit ports. You can override the globally configured default and set individual ports to the following:

- Negotiate-full-auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default.
- Auto-Gigabit – The port tries to perform a handshake with the other port to exchange capability information.
- Negotiation-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

USING THE CLI

To change the mode for individual ports on a chassis device, enter commands such as the following:

```
HP9300(config)# int ethernet 4/1 to 4/4
HP9300(config-mif-4/1-4/4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gigabit for ports 4/1 – 4/4.

Syntax: gig-default neg-full-auto | auto-gig | neg-off

USING THE WEB MANAGEMENT INTERFACE

To override the global 802.3x negotiation mode for an Gigabit individual port on a chassis device:

1. Select the Port link to display the Port configuration sheet.
2. Click on the Modify button next to the row of information for the port.
3. Select one of the following values from the Gig Port Default field's pulldown menu:
   - Default – The port uses the negotiation mode that was set at the global level.
   - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
   - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
   - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).
4. Click Apply to apply the change.
5. Select the Save To Flash link, then select Yes when prompted to save the changes to the system-config file.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to flash (using the File menu) for the changes to be preserved over a power cycle.

Modifying Port Priority (QoS)

You can give preference to the outbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. By default, the outbound traffic on all ports has normal priority. You can change a port's priority to one of the following:

- 0 or 1– Assigns an internal priority queue of 0. This is the default and is normal priority.
- 2 or 3 – Assigns an internal priority queue of 1.
- 4 or 5 – Assigns an internal priority queue of 2.
- 6 or 7 – Assigns an internal priority queue of 3. Priority 7 is the highest priority.

See “Quality of Service Algorithm” on page C-1 for information about how the queues work.

To change the priority of a port’s outbound traffic to a different queue, use one of the following methods.

NOTE: You also can assign VLANs, static MAC addresses, Layer 4 sessions, and AppleTalk sockets to different priority queues. See the following for information:

- VLANs – see “Configuring VLANs” on page 17-1.
- Static MAC entries – see “Configuring Static MAC Entries” on page 8-35.
- Layer 4 sessions – see “Configuring Layer 4 Quality of Service Parameters” on page 8-61.
- AppleTalk sockets – see “AppleTalk QoS Socket” on page 16-25.

USING THE CLI

To assign outbound traffic on port 5 of module 1 in a 9304M or 9308M to the highest priority, enter the following commands:

```
HP9300(config)# interface e 1/5
HP9300(config-if-1/5)# priority 7
```

Syntax: priority <0-7>

USING THE WEB MANAGEMENT INTERFACE

To change the QoS level for a port:

1. Select the Port link from the main menu to display the port summary panel.
2. Click the Modify button next to the port that is to be modified to display the configuration panel shown in Figure 8.9.
3. Select the appropriate value from the QoS field’s pulldown menu.
4. Click the Apply button to assign the change.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, you must save the configuration changes to the device’s flash memory (using the Save To Flash link) for the changes to be preserved over a power cycle.

Configuring Basic Layer 2 Parameters

The procedures in this section describe how to configure the following Layer 2 parameters. Note that some of these parameters apply only to the 9208M-SX switch, not to the routing switches.

- Spanning Tree Protocol (STP) – see “Enabling or Disabling the Spanning Tree Protocol (STP)” on page 8-31
- Layer 2 switching of unsupported router protocols (routing switches only) – see “Enabling or Disabling Layer 2 Switching (routing switches only)” on page 8-34
- Aging time for learned MAC address entries – see “Changing the Switch Age Time” on page 8-35
- Static, non-aging MAC address entries – see “Configuring Static MAC Entries” on page 8-35
- Trunk groups – see “Configuring Trunk Groups” on page 8-38
- Port-based VLANs – see “Enabling Port-Based VLANs” on page 8-37
- DHCP assist (6208M-SX switch only) – see “Configuring DHCP Assist (switch only)” on page 8-48
- Dynamic Host Configuration Protocol (DHCP) gateway list (switch only) – see “Defining a DHCP Gateway List (switch only)” on page 8-52

- IP Multicast traffic reduction (6208M-SX switch only) – see “Enabling or Disabling IP Multicast Traffic Reduction (switch only)” on page 8-53
- MAC address filters – see “Defining MAC Address Filters” on page 8-53
- Broadcast and Multicast Filters – see “Defining Broadcast and Multicast Filters” on page 8-57
- Port locks – see “Locking a Port To Restrict Addresses” on page 8-59

Enabling or Disabling the Spanning Tree Protocol (STP)

The STP (IEEE 802.1d bridge protocol) is supported on the HP 9304M, 9308M, and 6308M-SX routing switches, and the 6208M-SX switch. STP detects and eliminates logical loops in the network. STP also ensures that the least cost path is taken when multiple paths exist between ports or VLANs. If the selected path fails, STP searches for and then establishes an alternate path to prevent or limit retransmission of data.

STP must be enabled at the system level to allow assignment of this capability on the VLAN level. On the 6208M-SX switch, STP is enabled by default. On the 9304M, 9308M, and 6308M-SX routing switches, STP is disabled by default.

USING THE CLI

To enable STP for all ports on a device:

```
HP9300(config)# spanning tree
```

syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select Enable next to Spanning Tree.
3. Click the Apply button to assign the change.

Modifying STP Bridge and Port Parameters

You can modify the following STP Parameters:

- Bridge parameters—forward delay, maximum age, hello time, and priority
- Port parameters—priority and path cost

STP Bridge Parameters

You can configure the following STP parameters:

- Forward Delay: The period of time a bridge will wait (the listen and learn period) before forwarding data packets. Possible values: 4 – 30 seconds. Default is 15.
- Maximum Age: The interval a bridge will wait for receipt of a hello packet before initiating a topology change. Possible values: 6 – 40 seconds. Default is 20.
- Hello Time: The interval of time between each configuration BPDU sent by the root bridge. Possible values: 1 – 10 seconds. Default is 2.
- Priority: A parameter used to identify the root bridge in a network. The bridge with the lowest value has the highest priority and is the root. Possible values: 0 – 65,535. Default is 32,678.

STP Port Parameters

Spanning Tree Protocol port parameters priority and path cost are preconfigured with default values. If the default parameters meet your network requirements, no other action is required.

You can configure the following STP port parameters:

- Port Priority: This parameter can be used to assign a higher (or lower) priority to a port. In the event that traffic is re-routed, this parameter gives the port forwarding preference over lower priority ports within a VLAN or on the switch or routing switch (when no VLANs are configured for the system). Ports are re-routed based on their priority. The highest value is routed first. Possible values: 0 – 255. Default is 128**.**
- Path Cost: This parameter can be used to assign a higher or lower path cost to a port. This value can be used to bias traffic toward or away from a certain path during periods of rerouting. For example, if you wish to bias traffic away from a certain port, assign it a higher value than other ports within the VLAN or all other ports (when VLANs are not active on the switch or routing switch). Possible values are 0 – 65535 and the default values are 1000/port speed for half-duplex ports and (1000/port speed)/2 for full-duplex ports.

USING THE CLI

EXAMPLE: Suppose you want to enable STP on a 9308M-SX on which no port-based VLANs are active and change the hello-time from the default value of 2 to 8 seconds. Additionally, suppose you want to change the path and priority costs for port 5 only. To do so, enter the following commands.

```
HP9308(config)# span hello-time 8

HP9308(config)# span ethernet 5 path-cost 15 priority 64
```

syntax for global parameters: span [forward-delay <value>] | [hello-time <value>] | [maximum-age <time>] | [priority <value>]

syntax for port parameters: span ethernet <portnum> path-cost <value> | priority <value>

USING THE WEB MANAGEMENT INTERFACE

To modify the bridge parameters:

1. Select the STP link. The STP link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Click the Modify button next to the STP bridge summary panel. The entry panel shown in Figure 8.11 or Figure 8.12 will appear.
3. Enter the desired changes and click the Apply button to assign the changes.

To modify the STP port parameters:

1. Click the Modify button next to the STP port summary panel. The entry panel shown in Figure 8.11 or Figure 8.12 is displayed.
2. Enter the desired changes to the priority and path cost fields, seen at the bottom of the screen, and click the Apply Port STP button to save the changes.

NOTE: If you want to save the priority and path costs of one port to all other ports on the device within a VLAN, you can click the Apply To All Ports button.

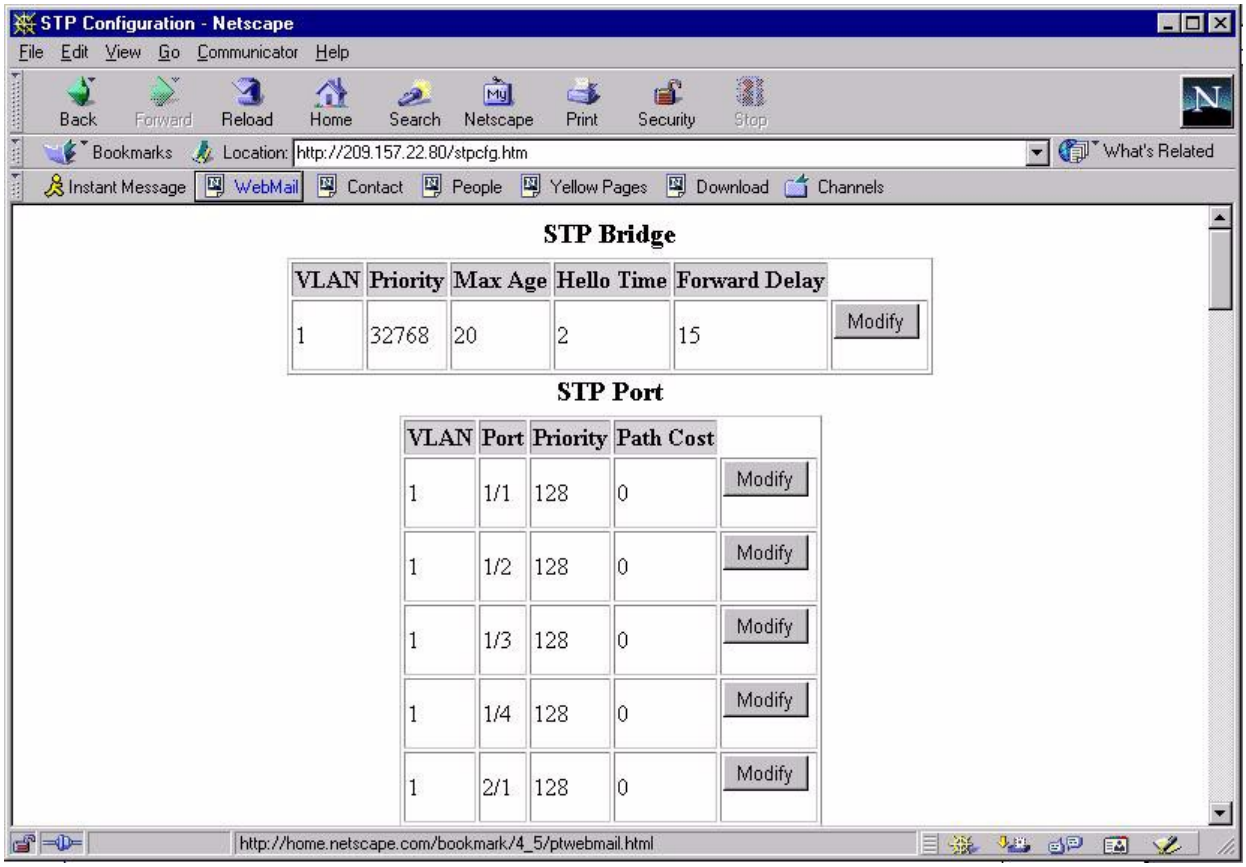

Figure 8.11 Spanning tree protocol summary showing operation with VLANs enabled

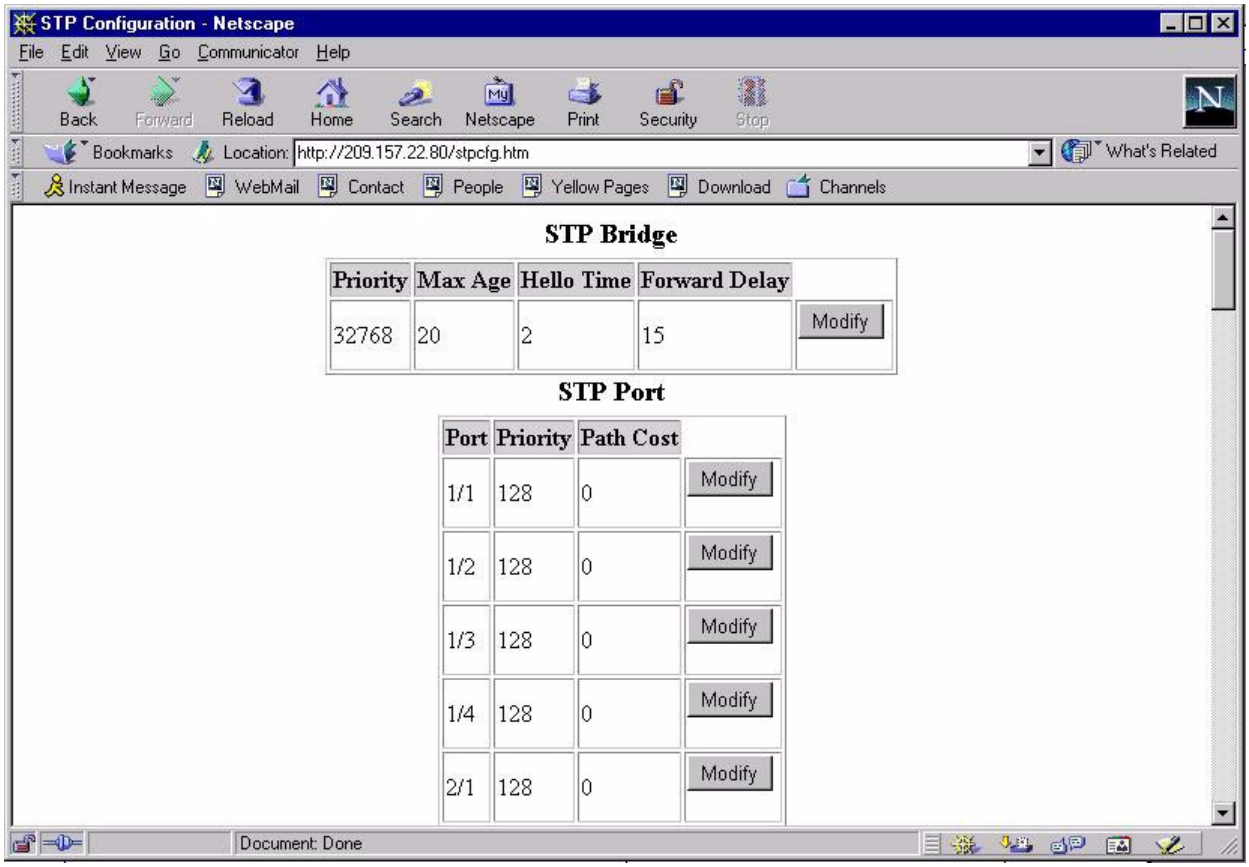

Figure 8.12 Spanning tree protocol summary showing operation without VLANs enabled

Enabling or Disabling Layer 2 Switching (routing switches only)

By default, the HP 9304M, 9308M, and 6308M-SX routing switches support Layer 2 switching. These devices switch the routing protocols that are not supported on the devices. If IPX routing is not enabled, then IPX traffic also is switched. By default IPX routing is disabled.

If you want to disable Layer 2 switching, you can configure a routing switch to route only.

NOTE: Make sure you really want to disable all Layer 2 switching operations before you use this option. Consult your reseller or HP for information.

USING THE CLI

To disable Layer 2 switching on a routing switch, enter the following:

```
HP9300(config)# route-only
HP9300(config)# exit
HP9300# write mem
HP9300# reload
```

To re-enable layer 2 switching on a routing switch, enter the following:

```
HP9300(config)# no route-only
HP9300(config)# exit
```

```
HP9300# write mem
HP9300# reload
```

Syntax: [no] route-only

USING THE WEB MANAGEMENT INTERFACE

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select Enable or Disable next to L2 Switching.
3. Click the Apply button to assign the change.

Changing the Switch Age Time

This parameter sets the aging period for ports on a device, defining how long a port address remains active in the address table. This parameter value range is from 0 – 65,535 seconds. The zero value results in no address aging. The default value for this field is 300 (seconds).

USING THE CLI

To change the aging period for MAC addresses from the default value of 300 seconds to 600 seconds:

```
HP9300(config)# mac-age-time 600
```

Syntax: [no] mac-age-time <age-time>

The <age-time> can be from 0– 65535.

USING THE WEB MANAGEMENT INTERFACE

To change the aging period for MAC addresses to 600 seconds:

1. Select the System link. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Advance link.
3. Enter the new age in the Switch Age Time field. In this example, enter 600.
4. Click the Apply button to assign the change.

Configuring Static MAC Entries

You can configure static MAC addresses on the devices.

NOTE: The routing switches also support the assignment of static IP Routes, static ARP, and static RARP entries. For details on configuring these types of static entries, see “Defining Static IP Routes” on page 9-11 and “Assigning Static ARP and RARP Entries (optional)” on page 9-12.

You can manually input the MAC address of a device to prevent it from being aged out of the system address table.

This option can be used to prevent traffic for a specific device, such as a server, from flooding the network with traffic when it is down. Additionally, the static MAC address entry is used to assign higher priorities to specific MAC addresses.

You can specify port priority (QoS) and VLAN membership (VLAN ID) for the MAC Address as well as specify device type of either router or host.

NOTE: The device type parameter "router" or "host" is not supported on routing switches when assigning static MAC addresses. This parameter is available only on the switch.

The default and maximum configurable MAC table sizes can differ depending on the device. To determine the default and maximum MAC table sizes for your device, display the system parameter values. See “Modifying System Parameter Default Settings” on page 8-69.

EXAMPLE: To add a static entry for a server with a MAC address of 1145.5563.67FF and a priority of 7 to port 2 of module 1 of a 9305M or 9308M:

USING THE CLI

```
HP9300(config)# static-mac-address 1145.5563.67FF e 1/2 priority 7
```

Syntax: static-mac-address <MAC-addr> ethernet <port-num> [priority <0-7>] [host-type|router-type]

The priority can be 0 – 7 (0 is lowest priority and 7 is highest priority).

The default priority is 0 (normal). The default type is host-type.

USING THE WEB MANAGEMENT INTERFACE

1. Select Static Station. The Static Station link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. The panel shown in Figure 8.13 will appear.

   NOTE: If static MAC address entries already exist, a summary panel will appear when you select Static Station. You can then access the static station entry panel by selecting Add Static Station.

2. Enter the 12-digit MAC address of the device requiring a static entry into the MAC Address Field. Separate each pair of digits with a dash (for example: aa-bb-cc-dd-ee-ff).
3. Enter the device VLAN ID, if it is a member of a port-based VLAN.
4. Select the interface that is attached to the device with the MAC address you are entering.
   - For the 9304M and 9308M, select the slot from the Slot field’s pulldown menu and select the port from the Port field’s pulldown menu.
   - For the 6308M-SX and 6208M-SX, select the port from the Port field’s pulldown menu.
5. Select the device type—router or host (6208M-SX switch only).
6. Select a priority level (0 – 7) from the QoS pulldown menu.
7. Select the Add button.

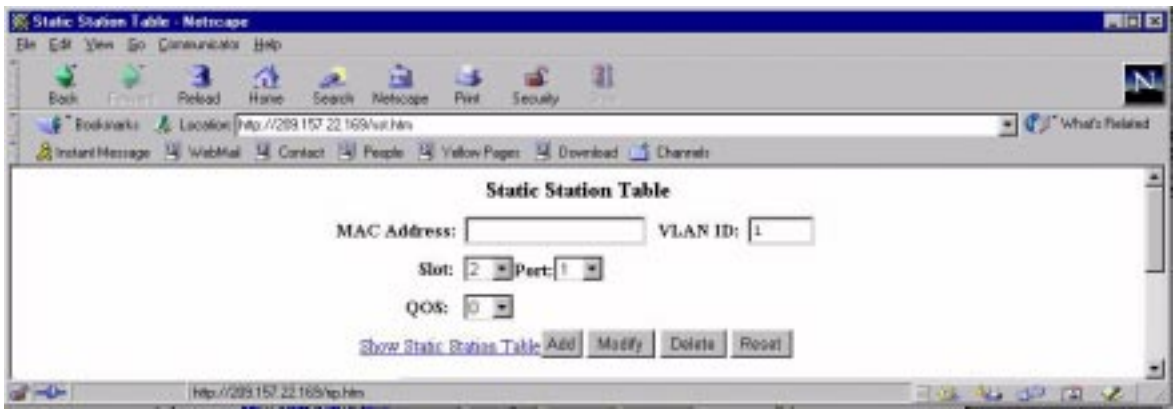

Figure 8.13 Static Station Table for assigning a static MAC address

Enabling Port-Based VLANs

Port and protocol VLANs must first be enabled at the system (global) level before they can be configured at the VLAN level. For details on configuring VLANs refer to “Configuring VLANs” on page 17-1.

USING THE CLI

When using the CLI, port and protocol-based VLANs are created by entering one of the following commands at the global CONFIG level of the CLI.

To create a port-based VLAN, enter commands such as the following:

```
HP9300(config)# vlan 222 by port
```

```
HP9300(config)# vlan 222 name Mktg
```

Syntax: vlan <num> by port

Syntax: vlan <num> name <string>

The <num> parameter specifies the VLAN ID. The valid range for VLAN IDs starts at 1 on all systems but the upper limit of the range differs depending on the device. In addition, you can change the upper limit on some devices using the **vlan max-vlans...** command. See “vlan max-vlans” on page B-131.

The <string> parameter is the VLAN name and can be a string up to 16 characters. You can use blank spaces in the name if you enclose the name in double quotes (for example, "Product Marketing".)

NOTE: The second command is optional and also creates the VLAN if the VLAN does not already exist. You can enter the first command after you enter the second command if you first exit to the global CONFIG level of the CLI.

USING THE WEB MANAGEMENT INTERFACE

To enable port-based VLANs on the switch or routing switch:

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the box next to the port option next to the policy-based VLANs heading.
3. Click the Apply button to assign the change.

Assigning IEEE (802.1q) Tagging to a Port

When a port is tagged, it allows communication among the different VLANs to which it is assigned. A common use for this might be to place an email server that multiple groups may need access to on a tagged port, which in turn, is resident in all VLANs that need access to the server.

NOTE: Tagging is disabled by default on individual ports.

NOTE: Tagging does not apply to the default VLAN.

For details on configuring port-based VLANs refer to “Configuring VLANs” on page 17-1.

USING THE CLI

When using the CLI, ports are defined as either tagged or untagged at the VLAN level.

EXAMPLE: Suppose you want to make port 5 on module 1 a member of port-based VLAN 4, a tagged port. To do so, enter the following:

```
HP9300(config)# vlan 4
```

```
HP9300(config-vlan-4)# tagged e 1/5
```

Syntax: tagged ethernet <port-num> [to <port-num> [ethernet <port-num>]]

USING THE WEB MANAGEMENT INTERFACE

To configure a port as tagged:

1. Select the Port link. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Select the Modify button next to the port that is to be modified to display the configuration panel shown in Figure 8.10.
3. Select Enable next to IEEE Tagging.

NOTE: This option appears only if you are modifying a port that is a member of a port-based VLAN other than the default VLAN. Tagging does not apply to ports that are not in a port-based VLAN and does not apply to the default VLAN.

4. Click the Apply button to assign the change.

Configuring Trunk Groups

The Trunk Group feature allows you to establish multiple high-speed load-sharing links between two switches or routing switches or between a switch or routing switch and a server. You can configure from 2 – 4 ports as a trunk group, supporting transfer rates of up to 4 Gbps of bi-directional traffic.

In addition, on the HP 9304M and 9308M, you can configure up to eight ports on two Gigabit Ethernet modules as a multi-module trunk group. Figure 8.14 shows an example of a configuration that uses trunk groups.

In addition to enabling load sharing of traffic, trunk groups provide redundant, alternate paths for traffic if any of the segments fail. The default value for this feature is disabled.

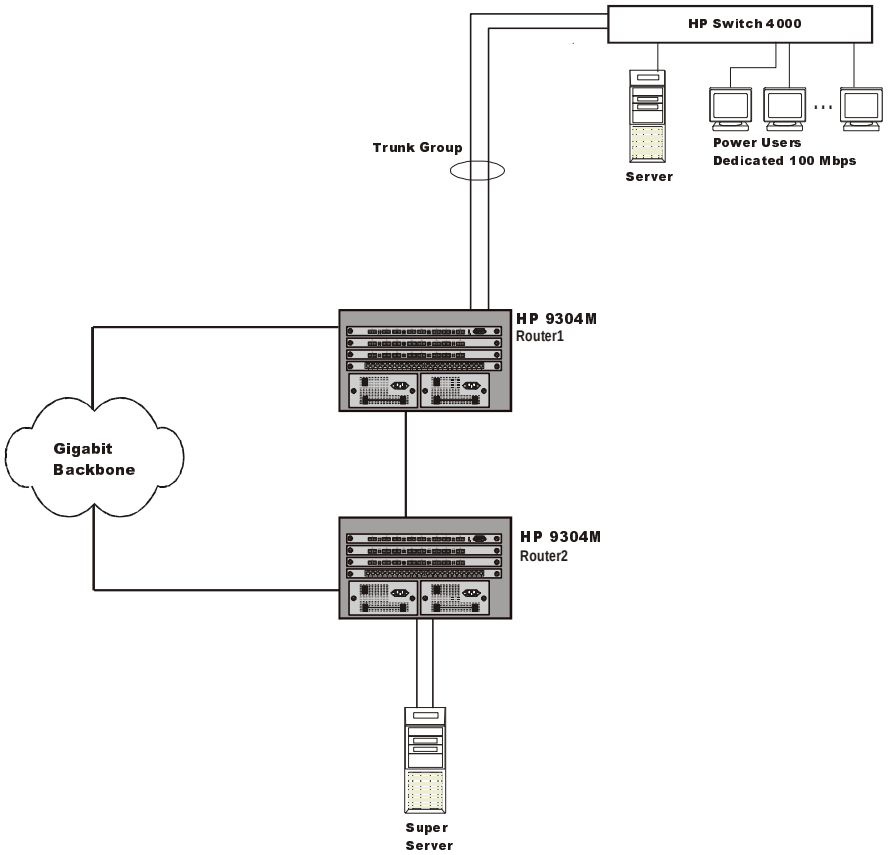

Figure 8.14 Trunk Group application within an HP routing switch network

NOTE: The ports in a trunk group make a single logical link. Therefore, all the ports in a trunk group must be connected to the same device at the other end.

Trunk Group Connectivity to a Server

To support termination of a trunk group, the server must have either multiple network interface cards (NICs) or a quad interface card installed. The trunk server is designated as a server with multiple adapters or a single adapter with multiple ports that share the same MAC and IP address. Figure 8.15 shows an example of a trunk group between a server and a switch.

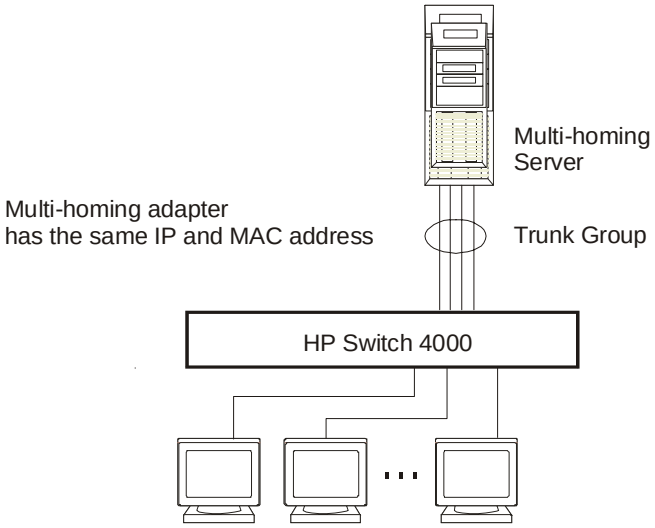

Figure 8.15 Trunk group between a server and a switch or routing switch

Trunk Group Rules

- You can configure up to four trunk groups.
- Each trunk group must start with a primary port. The primary port is always the lowest number in the following port ranges:
  - 9304M and 9308M: 1 – 4, 5 – 8, 9 – 12, 13 – 16 and 17 – 18 and 21 –24
  - 9308M-SX and 9208M-SX: 1 – 4 and 5 – 8 or 1 – 2, 3 – 4, 5 – 6, 7 – 8
- Port assignment must be contiguous. The port range cannot contain gaps. For example, you can configure ports 1, 2, 3, and 4 together as a trunk group but not ports 1, 3, 4, and 5.
- Port assignment cannot be across multiple trunk group boundaries. For example, on the 6308M-SX or 6208M-SX, ports 4 and 5 cannot be in the same trunk group.
- All the ports must be connected to the same device at the other end.
- All trunk group member properties must match the lead port of the trunk group with respect to the following parameters:
  - Port tag type (untagged or tagged port)
  - Port speed and duplex
  - QoS priority

  To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.

Additional Trunk Group Rules for Gigabit Ethernet Modules on a 9304M or 9308M

- You can configure a multi-slot trunk group on two Gigabit Ethernet modules.
- You can configure a maximum of eight ports in the trunk group.
- You can configure up to two groups of ports to make the trunk group and the groups must be alike. For example, you can group two sets of two ports together or two sets of four ports together but you cannot group a set of two ports with a set of four ports. Each group of ports can contain two or four ports.
- Each group of ports must begin with a primary port. On Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.
- When you specify the ports in the trunk group, you must specify them in ascending numerical order, beginning with the primary port. For example, to specify a group containing ports 1/1 – 1/4 and 3/1 – 3/4, you must specify them in the order shown. You cannot specify 3/1 – 3/4 first.
- Port configuration for each trunk group is based on the configuration of the primary port. To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.

Trunk Group Load Sharing

When you configure a trunk group, you specify whether the trunk group is a "switch" trunk group or a "server" trunk group:

- Switch trunk group – Use this type of trunk group to connect a switch or routing switch to another switch or routing switch.
- Server trunk group – Use this type of trunk group to connect a switch or routing switch to a file server or single host device.

The HP9304M, 9308M, 6308M-SX, and 6208M-SX devices load share across the ports in the trunk group. The type of load sharing depends on whether the device is a switch or a routing switch, the type of trunk group (switch or server), and the type of traffic. Table 8.1 lists the types of trunk group load sharing. Load sharing applies only to packets sent from the switch or routing switch across the trunk group.

Table 8.1: Trunk Group Load Sharing

Device Type	Trunk Group Type	Traffic Type	Load-Sharing Basis
Switch	Switch	All traffic	Destination MAC address
	Server	IP	Source and destination IP address
		All other	Source MAC address
Routing Switch	Switch	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other	Load sharing rules for switch trunk groups on switches apply
	Server	IP	Source IP address
		IPX	Source IPX address
		AppleTalk	Source AppleTalk address
		All other	Load sharing rules for server trunk groups on switches apply

Configuring a Trunk Group

The following steps must also be followed in configuring trunk groups:

1. Disconnect the cables from those ports on both devices that will be connected by the trunk group. Do not configure the trunk groups with the cables connected.

   NOTE: If you connect the cables before configuring the trunk groups and then rebooting, the traffic on the ports can create a spanning tree loop.

2. Configure the trunk group on one of the two switches or routing switches involved in the configuration. Save this configuration to flash and reboot the system.

   NOTE: HP recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

3. If the device at the other end of the trunk group is another switch or routing switch, repeat Step 2 for the other device.
4. When both devices are reset (re-booted) and operational, reconnect the cables to those ports that are now configured as trunk groups, starting with the first port (lead port) of each trunk group.
5. To verify the connection is operational, use the **show trunk** command.

Example 1: Configuring the Trunk Groups Shown in Figure 8.14

To configure the trunk groups shown in Figure 8.14, enter the following commands. Notice that the commands are entered on multiple devices.

USING THE CLI

To configure the trunk group link between the upper 9304M and the HP Switch 4000:

NOTE: The text shown in italics in the CLI example below shows messages echoed to the screen in answer to the CLI commands entered.

```
HP9300(config)# trunk switch e5 to 7
Trunk 2 is created for next power cycle.
Please save configuration to flash and reboot.
HP9300(config)# write mem
Write startup-config in progress.
.Write startup-config done.
HP9300(config)# exit
HP9300# reload
```

To configure the trunk group link between the lower 9304M and the server:

```
HP9300(config)# trunk server e2 to 4
Trunk 0 is created for next power cycle.
Please save configuration to flash and reboot.
HP9300(config)# write mem
Write startup-config in progress.
.Write startup-config done.
HP9300(config)# exit
HP9300# reload
```

Syntax: trunk <server|switch> ethernet <port number> to <port number>

You then configure the trunk group on the HP Switch 4000. See the documentation for this HP Switch 4000 for information.

USING THE NETWORK MANAGEMENT INTERFACE

To configure ports 5 – 8 as a trunk group between two switches, two routing switches, or a switch or routing switch and a server:

1. Select Trunk from the main menu. The Trunk Group configuration sheet is displayed on the screen as seen in Figure 8.16.

   NOTE: If trunk groups already exist, a table listing the configured trunk groups will appear instead. In this case, select Add Trunk.

   ---

2. Select a port range (for example, 5 – 8). On chassis devices, the port numbers include the slot numbers. For example, you can select 1/5 – 1/8.
3. Select the number of ports you want to use in the trunk group. You can select 2 or 4.
4. Click in the checkbox next to Server to place a checkmark in the box if the other end of the trunk group is a server. If the other end of the connection is a switch or routing switch, do not click this checkbox.

5. Select the Add button and the Save To Flash link and then re-boot the system using the Reset link from the main menu. The configuration of trunk groups is complete for this system.

NOTE: The Reset button does not reset the switch or routing switch. The reset button is a browser feature that clears changes made to the screen before they are applied. To re-boot a system, select the Reset link on the main menu.

---

6. If the other end of the trunk group is a switch or routing switch, log in to the other device and follow steps 1, 2, 3 and 5 above.

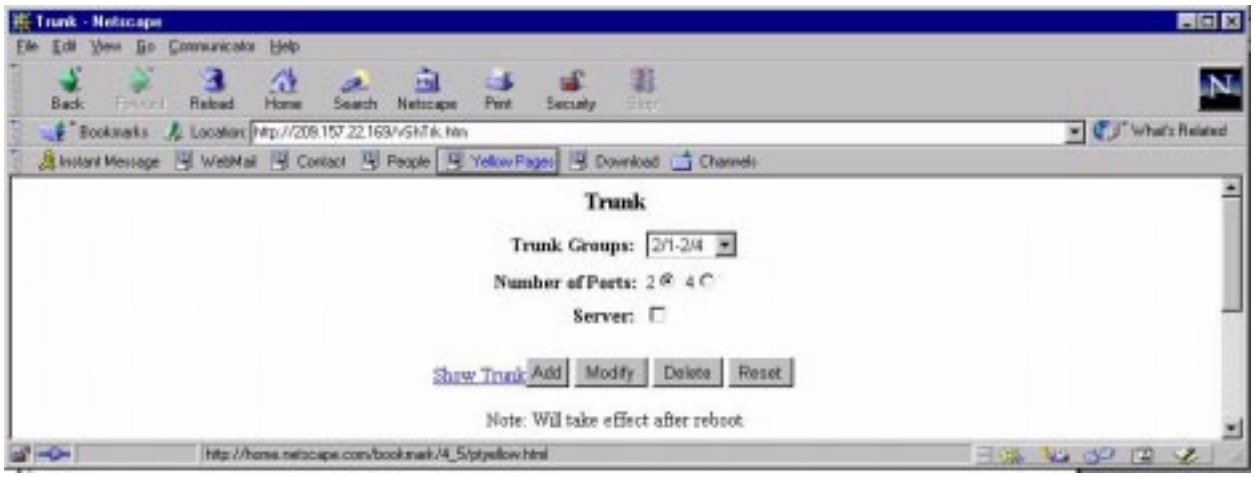

Figure 8.16 Trunk group configuration sheet

Example 2: Configuring a Trunk Group That Spans Multiple Gigabit Ethernet Modules in a 9304M or 9308M

To configure a trunk group that spans two modules in a 9304M or 9308M chassis, use one of the following methods.

USING THE CLI

To configure a trunk group consisting of two groups of ports, 1/1 – 1/4 on module 1 and 4/5 – 5/8 on module 4, enter the following commands:

```
HP9300(config)# trunk ethernet 1/1 to 1/4 ethernet 4/5 to 4/8

HP9300(config)# write mem

HP9300(config)# exit

HP9300# reload
```

NOTE: HP recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

Syntax: trunk [server|switch] ethernet <primary-portnum> to <portnum> ethernet <primary-portnum> to <portnum>

The **server|switch** parameter specifies whether the trunk ports will be connected to a server or to another switch or routing switch. This parameter affects the type of load balancing performed by the device. See “Trunk Group Load Sharing” on page 8-40. The default is **switch**.

Each **ethernet** parameter introduces a port group.

The **<primary-portnum> to <portnum>** parameters specify a port group. Notice that each port group must begin with a primary port. After you enter this command, the primary port of the first port group specified (which must be the group with the lower port numbers) becomes the primary port for the entire trunk group. For Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.

To configure a trunk group consisting of two groups of two ports each, enter commands such as the following:

```
HP9300(config)# trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4

HP9300(config)# write mem

HP9300(config)# exit

HP9300# reload
```

Notice that the groups of ports meet the criteria for a multi-slot trunk group. Each group contains the same number of ports (two) and begins on a primary port (1/1 and 3/3).

USING THE WEB MANAGEMENT INTERFACE

1. Select Trunk from the main menu. The Trunk Group configuration sheet is displayed on the screen as seen in Figure 8.16.

   NOTE: If trunk groups already exist, a table listing the configured trunk groups will appear instead. In this case, select Add Trunk.

2. Select Add Trunk.
3. Select 2 or 4 to indicate the number of ports in each group. Each group must have the same number of ports.
4. Select the port groups. Each group begins with the primary port number for that group. To select two groups, click on the first group, then hold down the CTRL key and click on the second group. Do not select more than two groups.
5. Select Server if you are connecting the trunk group ports to a server. Otherwise, the software assumes you are connecting the trunk groups ports to another switch or routing switch and uses the default value Switch.
6. Click the Add button to add the trunk group.
7. Select the Save To Flash link, then select Yes when prompted to save the trunk group configuration to flash memory.
8. Select Reload to reload the system and make the trunk group operational.
9. Connect the ports to the server, switch, or routing switch at the other end of the trunk.

NOTE: HP recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

Modifying Trunk Group Membership

You can change port membership by removing individual ports from the trunk group. To remove a port from a trunk group, use one of the following methods.

USING THE CLI

To remove ports 1/3 and 1/4 from the trunk group, enter the following command:

```
HP9300(config)# no trunk ethernet 1/3 to 1/4
```

Syntax: no trunk ethernet <port> [to <port>]

The <port> parameter indicates the port you are removing.

NOTE: Make sure you enter the lower port in the range before the "to" and the higher port in the range after the "to".

As a shortcut, you also can enter just the lower port in the range. The software automatically removes all higher ports in addition to the specified port. For example, to remove ports 1/3 and 1/4, you can enter the following command:

```
HP9300(config)# no trunk ethernet 1/3
```

The rules regarding trunk group membership are the same as in earlier software releases.

Therefore, for trunk group 1/1 – 1/4, the following commands are not valid:

```
HP9300(config)# no trunk ethernet 1/2
```

Or

```
HP9300(config)# no trunk ethernet 1/2 to 1/4
```

These commands are invalid because the trunk group cannot contain only a single port. These commands, if the software allowed them, would result in a trunk group consisting only of port 1/1.

Trunk groups can contain two ports or four ports but cannot contain only three ports. Therefore, the following command also is invalid for trunk group 1/1 – 1/4:

```
HP9300(config)# no trunk ethernet 1/4
```

This command is invalid because it would result in a trunk group containing three ports, 1/1 – 1/3.

USING THE WEB MANAGEMENT INTERFACE

1. Disconnect the ports to the server, switch, or routing switch at the other end of the trunk.
2. Select the Trunk link. A table listing the configured trunk groups is displayed.
3. Click the Modify button next to the trunk group you want to modify. The Trunk panel is displayed, as shown in Figure 8.16.
4. Select 2 or 4 to indicate the number of ports.
5. Select Server if you are connecting the trunk group ports to a server. Otherwise, the software assumes you are connecting the trunk groups ports to another switch or routing switch and uses the default value Switch.
6. Click the Modify button.
7. Select the Save To Flash link, then select Yes when prompted to save the trunk group configuration to flash memory.

NOTE: HP recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

NOTE: If you accidentally select a different port range by selecting a value in the Trunk Group field's pulldown menu, the software creates a new trunk group with the range and other values you select.

Deleting a Trunk Group

To delete a trunk group, use either of the following methods.

USING THE CLI

To delete a trunk group, use "**no**" in front of the command you used to create the trunk group. For example, to remove one of the trunk groups configured in the examples above, enter the following command:

```
HP9300(config)# no trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4
```

Syntax: no trunk ethernet <port> to <port>

USING THE WEB MANAGEMENT INTERFACE

To delete a trunk group:

1. Select the Trunk link. A table listing the configured trunk groups is displayed.
2. Click the Delete button next to the trunk group you want to delete.
3. Reboot the system by selecting the Reset link from the main menu. The selected trunk group is now inactive on this end of the path.

NOTE: If the other end of the trunk group is a switch or routing switch, log in to the other system and follow steps 1 – 3 above.

Displaying Trunk Group Configuration Information

To display configuration information for trunk groups, use of the following methods. Each method displays information for configured trunk groups and operational trunk groups. A configured trunk group is one that has been configured in the software but has not been placed into operation by a reset or reboot. An operational trunk group is one that has been placed into operation by a reset or reboot.

USING THE CLI

Enter the following command at any CLI level:

```
HP9300(config)# show trunk
Configured trunks:
Trunk Type   Ports
  1   Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4
Operational trunks:
Trunk Type   Ports                                  Duplex Speed Tag Priority
  1   Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4       None   None  No  level0
```

Syntax: show trunk

The following table describes the information displayed by the **show trunk** command.

Table 8.2: CLI Trunk Group Information

This Field...	Displays...
Trunk	The trunk group number. The software numbers the groups in the display to make the display easy to use.
Type	The type of trunk group, which can be one of the following: • Server – The trunk group is connected to server. • Switch – The trunk group is connected to another switch or routing switch.
Ports	The ports in the trunk group.
Duplex	The mode of the port, which can be one of the following: • None – The link on the primary trunk port is down. • Full – The primary port is running in full-duplex. • Half – The primary port is running in half-duplex. **Note**: This field and the following fields apply only to operational trunk groups.
Speed	The speed set for the port. The value can be one of the following: • None – The link on the primary trunk port is down. • 10 – The port speed is 10 Mbps. • 100 – The port speed is 100 Mbps. • IG – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.

USING THE WEB MANAGEMENT INTERFACE

Select the Trunk link. The configured trunk groups and operational trunk groups are listed.

This display shows the following information.

Table 8.3: Web Management Trunk Group Information

This Field...	Displays...
Connection Type	The type of trunk group, which can be one of the following: • Server – The trunk group is connected to server. • Switch – The trunk group is connected to another switch or routing switch.
Port Members	The ports in the trunk group.

Configuring DHCP Assist (switch only)

DHCP Assist allows the 9208M-SX switch to assist a routing switch that is performing multi-netting on its interfaces as part of its DHCP relay function.

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 switch does so by stamping each request with its IP gateway address in the DHCP discovery packet.

NOTE: The routing switches provide BootP/DHCP assistance by default on an individual port basis. See "Modifying the IP Address Used for Stamping BootP/DHCP Requests" on page 9-28.

By allowing multiple sub-net DHCP requests to be sent on the same wire, you can reduce the number of routing switch ports required to support secondary addressing as well as reduce the number of DHCP servers required, by allowing a server to manage multiple sub-net address assignments.

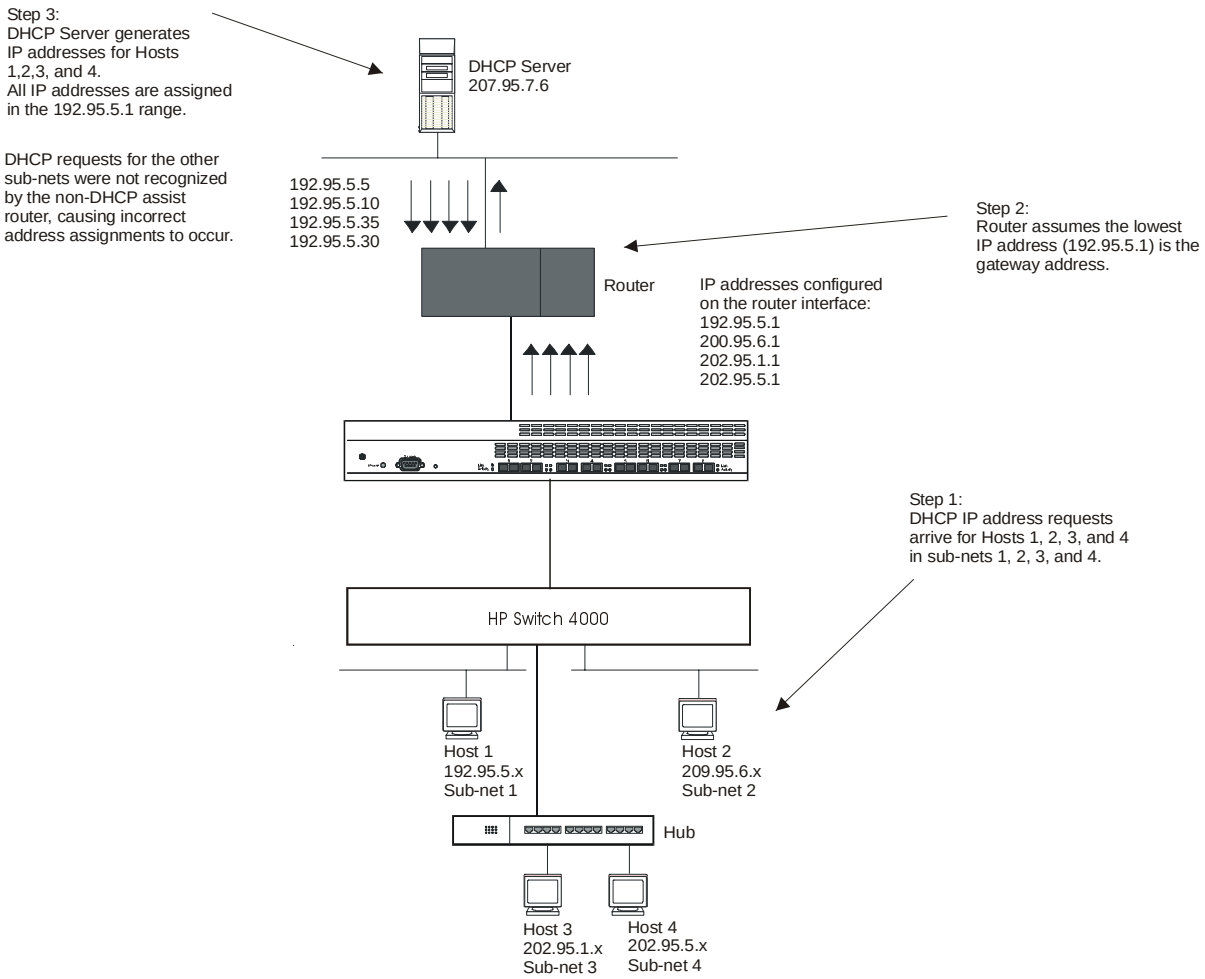

Figure 8.17 DHCP requests in a network without DHCP Assist on the switch

In a network operating without DHCP Assist, hosts can be assigned IP addresses from the wrong sub-net range because a router with multiple sub-nets configured on an interface cannot distinguish between DHCP discovery packets received from different sub-nets.

For example, in Figure 8.17 a host from each of the four sub-nets supported on a switch requests an IP address from the DHCP server. These requests are sent transparently to the router. Because the router is unable to

determine the origin of each packet by sub-net, it assumes the lowest IP address or the 'primary address' is the gateway for all ports on the switch and stamps the request with that address.

When the DHCP request is received at the server, it assigns all IP addresses within that range only.

With DHCP Assist enabled on the 6208M-SX switch, correct assignments are made because the switch provides the stamping service.

How DHCP Assist Works

Upon initiation of a DHCP session, the client sends out a DHCP discovery packet for an address from the DHCP server as seen in Figure 8.18. When the DHCP discovery packet is received at the 6208M-SX switch with the DHCP Assist feature enabled, the gateway address configured on the receiving interface is inserted into the packet. This address insertion is also referred to as stamping.

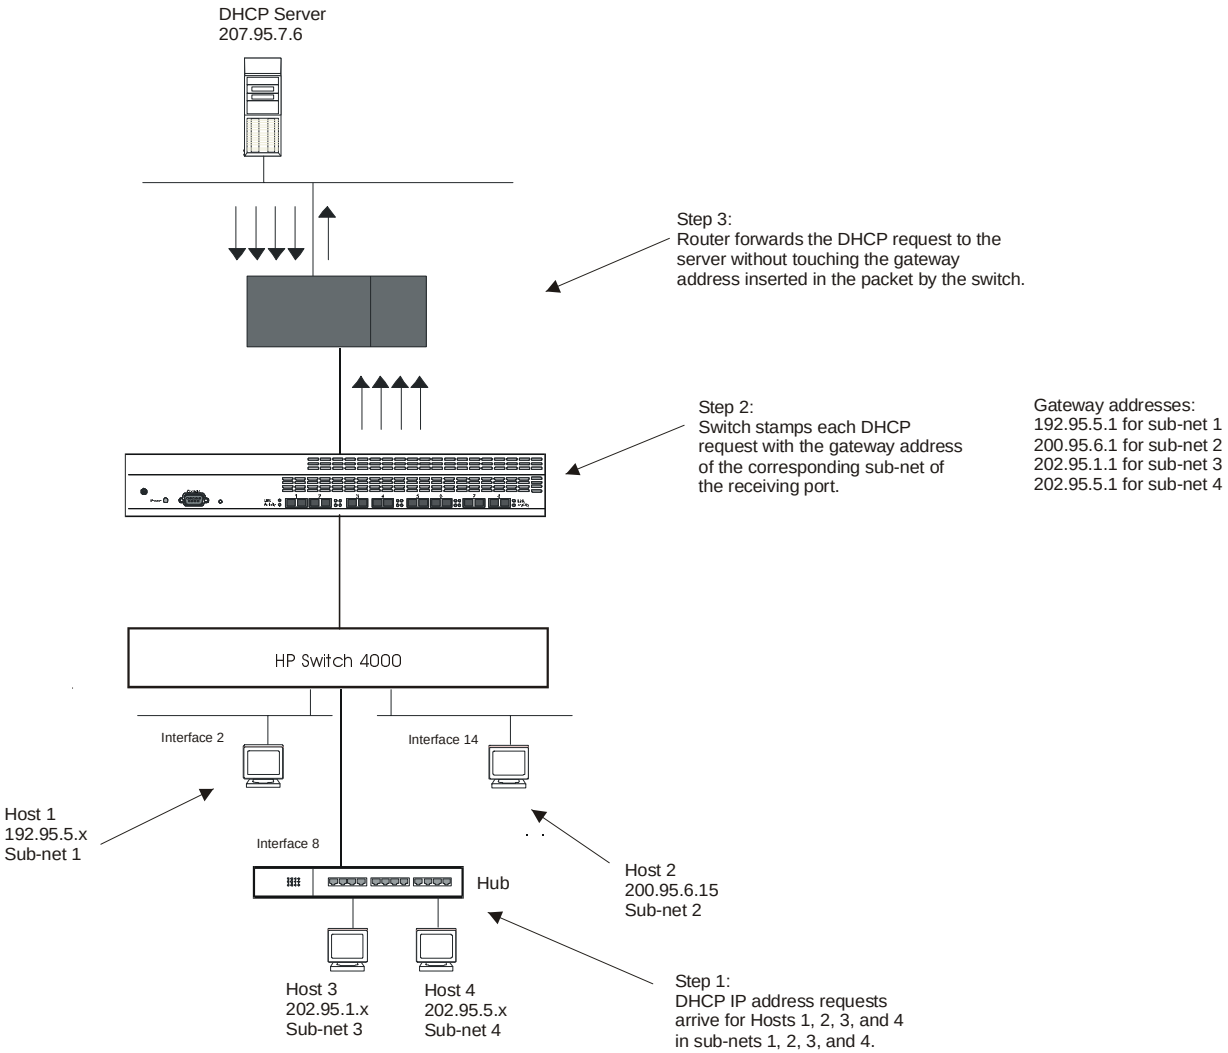

Figure 8.18 DHCP requests in a network with DHCP Assist operating on a 6208M-SX switch

When the stamped DHCP discovery packet is then received at the routing switch, it is forwarded to the DHCP server. The DHCP server then extracts the gateway address from each request and assigns an available IP address within the corresponding IP sub-net (Figure 8.19). The IP address is then forwarded back to the workstation that originated the request.

NOTE: The DHCP relay function of the connecting router needs to be turned on.

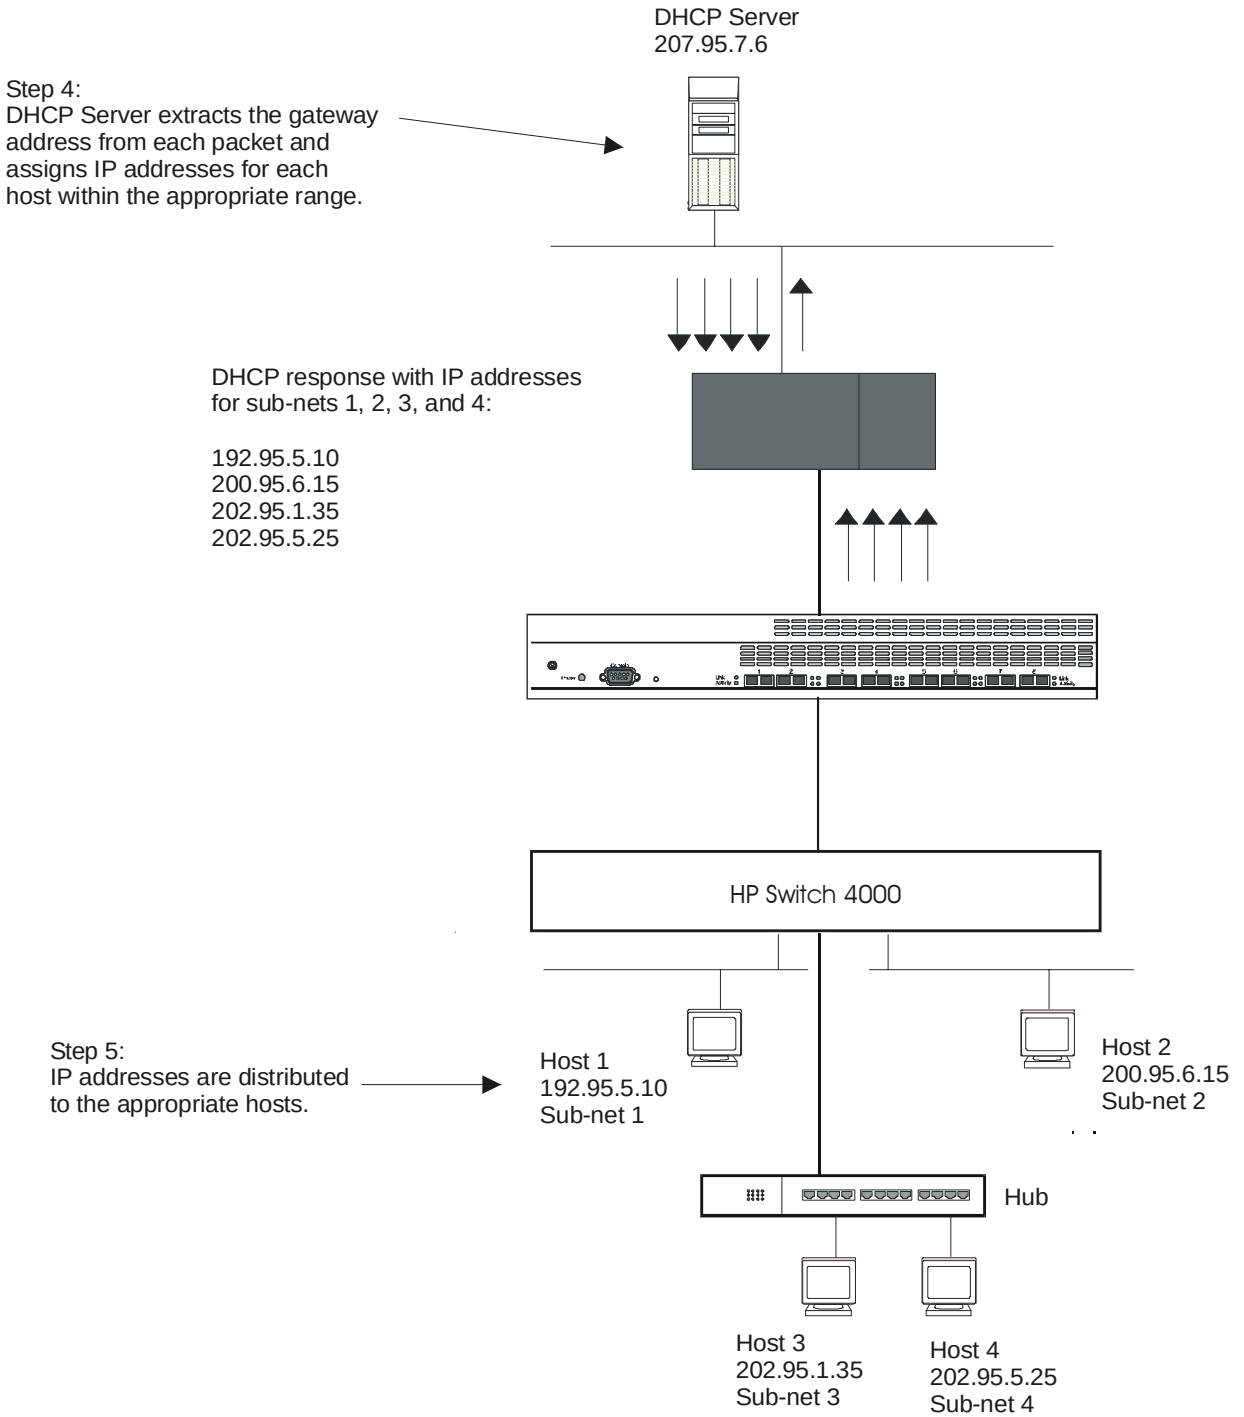

Figure 8.19 DHCP offers are forwarded back toward the requestors

Configuring DHCP Assist

A gateway address needs to be defined on the 6208M-SX switch for each sub-net. Once defined, a gateway list can be assigned to an interface on the switch. Each gateway address defined on the switch corresponds to an IP address of the router interface.

Up to eight addresses can be defined for each gateway list in support of ports that are multi-homed. When multiple IP addresses are configured for a gateway list, the switch inserts the addresses into the discovery packet in a round robin fashion.

Up to 32 gateway lists can be defined for each switch.

USING THE CLI

EXAMPLE: To create the configuration indicated in Figure 8.18 and Figure 8.19:

```
HP6208(config)# dhcp-gateway-list 1 192.95.5.1
HP6208(config)# dhcp-gateway-list 2 200.95.6.1
HP6208(config)# dhcp-gateway-list 3 202.95.1.1 202.95.5.1
HP6208(config)# int e 2
HP6208(config-if-2)# dhcp-gateway-list 1
HP6208(config-if-2)# int e8
HP6208(config-if-8)# dhcp-gateway-list 3
HP6208(config-if-8)# int e 14
HP6208(config-if-14)# dhcp-gateway-list 2
```

Syntax: dhcp-gateway-list <num> <ip-addr>

USING THE WEB MANAGMENT INTERFACE

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the DHCP Gateway link. The panel shown in Figure 8.20 will appear.
3. Enter the list ID in the List ID field.
4. Enter up to eight IP sub-nets for the list in the IP List fields.
5. Click the Add button to assign the change.
6. Repeat steps 1 – 5 for each gateway list.

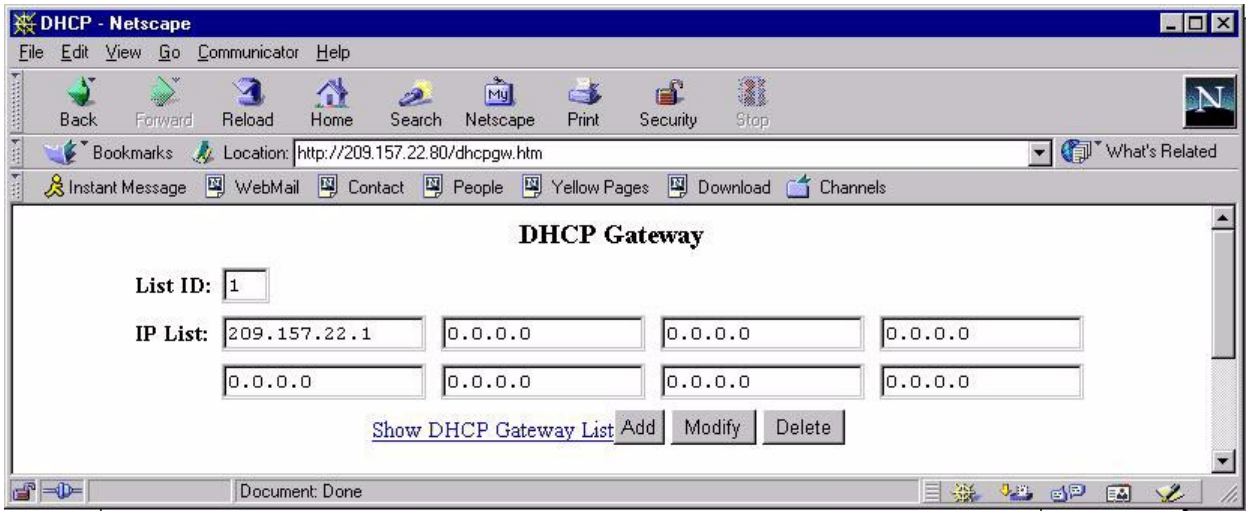

Figure 8.20 DHCP gateway list entry panel

Defining a DHCP Gateway List (switch only)

You can associate a gateway list with a port. You must configure a gateway list when DHCP Assist is enabled on the 6208M-SX switch. The gateway list contains a gateway address for each sub-net that will be requesting addresses from a DHCP server. The list allows the stamping process to occur. Each gateway address defined on the switch corresponds to an IP address of the router interface.

Up to eight addresses can be defined for each gateway list in support of ports that are multi-homed. When multiple IP addresses are configured for a gateway list, the switch inserts the addresses into the discovery packet in a round robin fashion.

Up to 32 gateway lists can be defined for each switch.

USING THE CLI

EXAMPLE: To define the sub-net address 192.95.5.1 as a gateway address and assign it to interface 2, enter the following:

```
HP6208(config)# dhcp-gateway-list 1 192.95.5.1

HP6208(config)# int e 2

HP6208(config-if-2)# dhcp-gateway-list 1
```

Syntax: dhcp-gateway-list <num> <ip-addr>

USING THE WEB MANAGEMENT INTERFACE

NOTE: To view the IP addresses associated with each index, select the DHCP Gateway link found on the System configuration sheet.

---

1. Select the Port link to display the port summary panel. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port that is to be modified to display the configuration panel shown in Figure 8.10.
3. Select the list number to assign to the port from the DHCP Gateway ID field's pulldown menu.
4. Click the Apply button to assign the new configuration.

NOTE: For details on configuring DHCP Assist and defining gateway lists, see “Configuring DHCP Assist” on page 8-51.

Enabling or Disabling IP Multicast Traffic Reduction (switch only)

This feature allows the HP 6208M-SX switch to limit the multicast of IGMP packets to only those ports on the switch that are identified as IP Multicast members.

- When configured to operate in the active mode, the switch will actively send out host queries to identify IP Multicast groups on the network and insert this information into the IGMP packet. Routers in the network generally handle this operation.
- In the passive mode, the switch will simply identify the packet as an IGMP packet and forward it accordingly. (The passive mode is sometimes called "IGMP snooping".)

By default, IP Multicast Traffic Reduction is disabled.

In most cases, the switch should be configured to operate in the passive mode. An exception is when the switch is in a stand-alone switched network with no external IP mulitcast router attachments.

NOTE: The change must be saved and the system reset to become active.

USING THE CLI

To enable IP Multicast Traffic Reduction on a switch to operate in the passive mode:

```
HP6208(config)# ip multicast passive
```

To enable IP Multicast Traffic Reduction on a switch to operate in the active mode:

```
HP6208(config)# ip multicast active
```

syntax: ip multicast <active | passive>

USING THE WEB MANAGEMENT INTERFACE

To enable IP Multicast Traffic Reduction on a switch:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select Enable next to IP Multicast.
3. Select either Active or Passive next to the IGMP option.
4. Click Apply to assign the changes.

Defining MAC Address Filters

MAC layer filtering enables you to build access lists based on MAC layer headers in the Ethernet/IEEE 802.3 frame. You can filter on the source and destination MAC addresses as well as other information such as the EtherType, LLC1 DSAP or SSAP numbers, and a SNAP EtherType. The filters apply to incoming traffic only.

NOTE: You cannot use Layer 2 filters to filter Layer 4 information. To filter Layer 4 information, use IP access policies. See “Defining IP Access Policies” on page 9-14.

You configure MAC filters globally, then apply them to individual interfaces. To apply MAC filters to an interface, you add the filters to that interface’s MAC filter group.

MAC filters provide one of two actions:

- permit – allows packets that meet the filter criteria to be forwarded.
- deny – prevents packets that match the filter criteria from being forwarded.

The device takes the action associated with the first matching filter. If the packet does not match any of the filters in the access list, the default action is to drop the packet. If you want the system to permit traffic by default, you must specifically indicate this by making the last entry in the access list a permit filter. Here is an example:
mac filter <last index number> permit any any

For routing switches, the MAC filter is applied only to those inbound packets that are to be switched. This includes those ports associated with a Virtual Ethernet (VE) interface. However, the filter is not applied to the VE; it is applied to the physical port.

NOTE: Use MAC Layer 2 filters only for switched traffic. If a routing protocol (for example, IP or IPX) is configured on an interface, a MAC filter defined on that interface is not applied to inbound packets. If you want to filter inbound route traffic, configure a route filter. For example, you can use the following command to filter IP route traffic: **ip filter-group <in | out> <filterID List>**

When you create a MAC filter, it takes effect immediately. You do not need to reset the system. However, you do need to save the configuration to flash memory to retain the filters across system resets.

For complete MAC filter examples, see “mac filter” on page B-108.

To define a MAC filter, use one of the following methods.

USING THE CLI

```
HP9300(config)# mac filter 1 deny 1543.6734.366e any snap eq 806

HP9300(config)# int e 1/1

HP9300(config-if-1/1)# mac filter-group 1
```

Syntax: mac filter <filter-num> <permit|deny> <any|H.H.H> <any|H.H.H> <etype|IIc|snap> <operator> <frame-type>

Syntax: mac-filter-group <filter-list>

NOTE: Remember that the filters must be applied as a group. For example, if you want to apply four filters to an interface, they must all appear on the same command line.

USING THE WEB MANAGEMENT INTERFACE

To define a MAC filter:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the MAC Filter link from the System configuration sheet to display the MAC Filter panel. See Figure 8.21.

   NOTE: If you have already configured MAC filters, a table listing the filters is displayed instead. Select Add MAC Filter.

   ---

3. Select the filter action by selecting Permit or Deny next to Action.
4. Enter the source MAC address in the Source Address field. Separate the bytes in the address with dashes. (See the example in Figure 8.21.)
5. Enter the comparison mask for the source address in the Source Mask field. The mask consists of "f"s and "0"s or the word "any".
   - An "f" indicates a significant bit. The software checks the indicated bit in each packet’s source MAC address.
   - A "0" indicates an insignificant bit. The software does not care what value is in the bit position.
   - "any" matches all bits and is equivalent to entering "ff-ff-ff-ff-ff-ff".

6. Enter the destination MAC address in the Destination Address field. Separate the bytes in the address with dashes.
7. Enter the comparison mask for the destination address in the Destination Mask field.
8. Select the frame type from the Frame Type field's pulldown menu.
9. Select an operator from the Operator field's pulldown menu to filter by protocol type.
10. Enter a protocol in the Protocol field.
11. Click the Add button to add the filter. The filter is now configured in the software but has not yet been applied to a port.
12. Select Filter Group to display the Filter Group panel as shown in Figure 8.22.

NOTE: If you have already configured MAC filter groups, a table listing the groups is displayed instead. Select Add MAC Filter Group.

---

13. Enter the filter numbers in the Filter Group field. Separate each filter number from the next one by a single space. The software applies the filters in the order you list them, from left to right. When a packet matches a filter, the software stops comparing the packet against the filter list and applies the action specified in the matching filter.
14. Click the Add button to add the filter group.

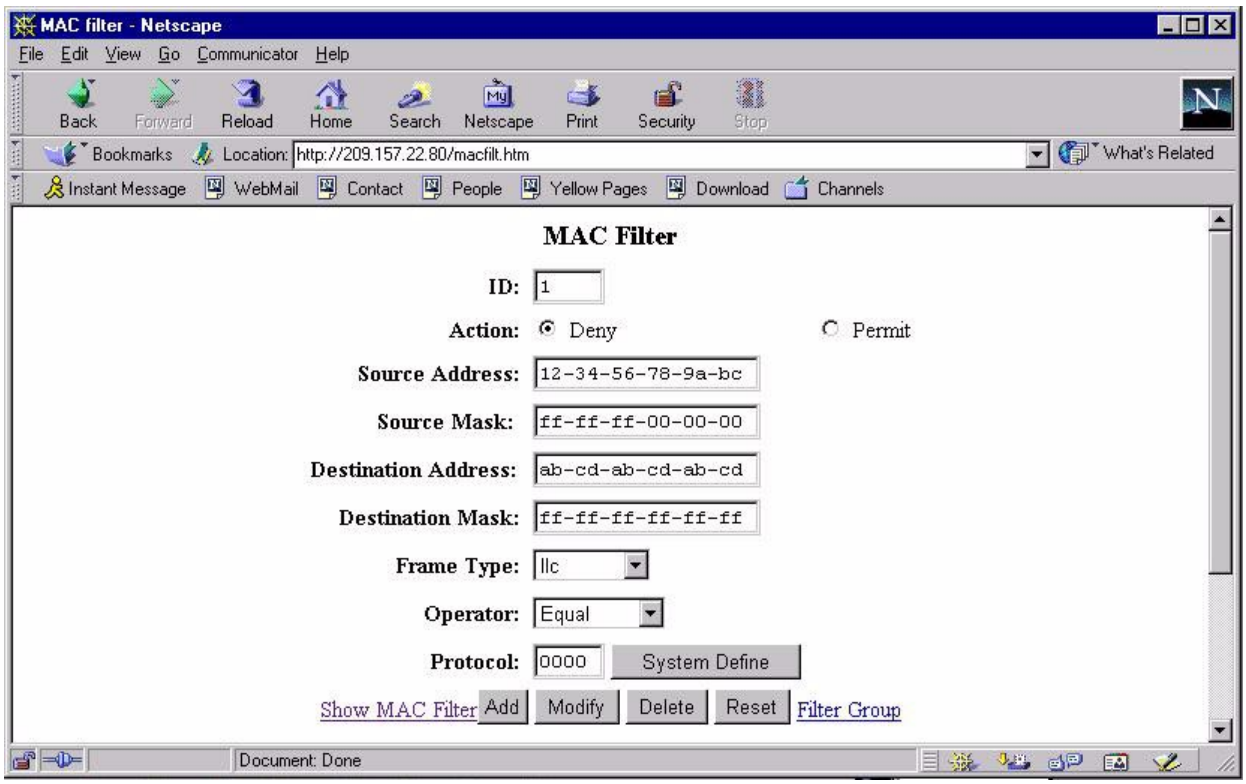

Figure 8.21 MAC Filter panel

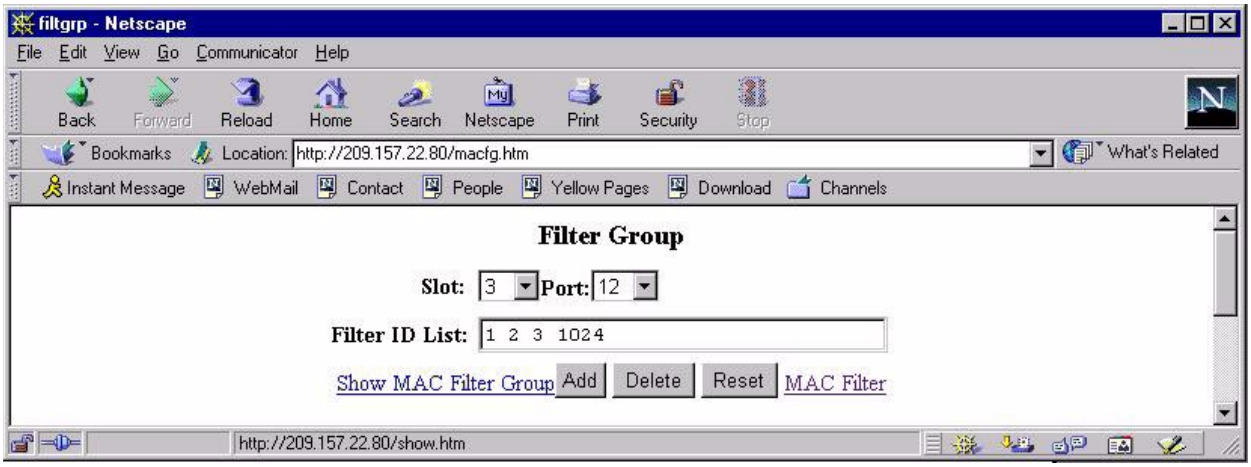

Figure 8.22 Filter Group panel

Enabling Logging of Packets Denied by MAC Filters

You can configure a device to generate Syslog entries and SNMP traps for packets that are denied by Layer 2 MAC filters. You can enable logging of denied packets on a global basis or an individual port basis.

See Example 4 in “show logging” on page B-242 for an example of the Syslog entries and a description of how the timer for the entries works.

USING THE CLI

To configure Layer 2 MAC filter logging globally, enter the following CLI commands at the global CONFIG level:

```
HP9300(config)# mac filter log_en
HP9300(config)# write mem
```

Syntax: [no] mac filter log_en

To configure Layer 2 MAC filter logging for MAC filters applied to ports 1/1 and 3/3, enter the following CLI commands:

```
HP9300(config)# int ethernet 1/1
HP9300(config-if-1/1)# mac filter-group log_en
HP9300(config-if-1/1)# int ethernet 3/3
HP9300(config-if-3/3)# mac filter-group log_en
HP9300(config0if-3/3)# write mem
```

Syntax: [no] mac filter-group log_en

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a Layer 2 MAC filter to generate Syslog entries and SNMP traps for denied packets using the Web management interface.

Changing the Maximum Number of Filters

By default, the devices support up to 64 filters per system and up to 32 filters per filter group on an interface. You can change these limits using the **system-max** command.

EXAMPLE:

```
HP9300(config)# system-max mac-filter-sys 128
HP9300(config)# system-max mac-filter-port 64
HP9300(config)# exit
HP9300# wr mem
HP9300# reload
```

Syntax: system-max mac-filter-port <value>

The <value> specifies the maximum number of MAC filters. The default and the valid range differ depending on the device. To display the default and range for your device, enter the show default values command. See "system-max" on page B-126 and "show default" on page B-218.

NOTE: The **system-max** commands take effect only after you reload the system.

Defining Broadcast and Multicast Filters

You can filter Layer 2 broadcast and multicast packets on specific ports.

- Layer 2 broadcast packets have the value "FFFFFFFFFFFF" (all ones) in the destination MAC address field. You can configure broadcast filters for all types of IP packets or for UDP packets.
- Layer 2 multicast packets have a multicast address in the destination MAC address field. You can configure multicast filters to filter on all MAC addresses or a specific multicast address.

You can configure up to eight of each type of filter.

To configure a Layer 2 broadcast or multicast filter, you define the filter globally to either filter out all types of broadcasts or to filter out only IP UDP broadcasts. After configuring a broadcast or multicast filter, you apply it to specific ports. Broadcast and multicast filters apply only to outbound traffic.

When defining the filter, you can specify a port-based VLAN ID. If a port is a member of more than one VLAN and is a tagged port, specifying a VLAN ID causes the filter to be applied only to traffic for the specified VLAN on the tagged ports to which you apply the filter. Otherwise, the filter applies to all the VLANs of which the port is a member.

The filters are applied in numerical order, beginning with filter number 1. As soon as the software finds a matching filter for a given packet, the filtering process stops for that packet. For example, if you configure filter 1 to filter all broadcast traffic and filter 2 to filter only IP UDP traffic, filter 1 will always be true for any broadcast packet, and thus the software will never consult filter 2 for ports that you configure to use filter 1.

Configuring a Layer 2 Broadcast Filter

To configure a broadcast filter, you must have access to the CONFIG level of the CLI. You can configure up to eight broadcast filters on a device.

Syntax: [no] broadcast filter <filter-ID> <any | ip udp> [vlan <vlan-id>]

Syntax: [no] exclude-ports ethernet <port-num> to <port-num>

Or

Syntax: [no] exclude-ports ethernet <port-num> ethernet <port-num>

The **exclude-ports** command specifies the ports to which the filter applies.

The **<filter-ID>** specifies the filter number and can a number from 1 – 8. The software applies the filters in ascending numerical order. As soon as a match is found, the software takes the action specified by the filter (block the broadcast) does not compare the packet against additional broadcast filters.

You can specify **any** or **ip udp** as the type of broadcast traffic to filter. The **any** parameter prevents all broadcast traffic from being sent on the specified ports. The **ip udp** parameter prevents all IP UDP broadcasts from being sent on the specified ports but allows other types of broadcast traffic.

If you specify a port-based VLAN ID, the filter applies only to the broadcast domain of the specified VLAN, not to all broadcast domains (VLANs) on the device.

As soon as you press Enter after entering the command, the CLI changes to the configuration level for the filter you are configuring. You specify the ports to which the filter applies at the filter's configuration level.

NOTE: This is the same command syntax as that used for configuring port-based VLANs. Use the first command for adding a range of ports. Use the second command for adding separate ports (not in a range). You also can combine the syntax. For example, you can enter **exclude-ports ethernet 1/4 ethernet 2/6 to 2/9**.

Configuration Examples

To configure a Layer 2 broadcast filter to filter all types of broadcasts, then apply the filter to ports 1/1, 1/2, and 1/3, enter the following commands:

```
HP9300(config)# broadcast filter 1 any
HP9300(config-bcast-filter-id-1)# exclude-ports ethernet 1/1 to 1/3
HP9300(config-bcast-filter-id-1)# write mem
```

To configure two filters, one to filter IP UDP traffic on ports 1/1 – 1/4, and the other to filter all broadcast traffic on port 4/6, enter the following commands:

```
HP9300(config)# broadcast filter 2 ip udp
HP9300(config-bcast-filter-id-2)# exclude-ports ethernet 1/1 to 1/4
HP9300(config-bcast-filter-id-3)# exit
HP9300(config)# broadcast filter 3 any
HP9300(config-bcast-filter-id-3)# exclude-ports ethernet 4/6
HP9300(config-bcast-filter-id-3)# write mem
```

To configure an IP UDP broadcast filter and apply that applies only to port-based VLAN 10, then apply the filter to two ports within the VLAN, enter the following commands:

```
HP9300(config)# broadcast filter 4 ip udp vlan 10
HP9300(config-bcast-filter-id-4)# exclude-ports eth 1/1 eth 1/3
HP9300(config-bcast-filter-id-1)# write mem
```

Configuring a Layer 2 Multicast Filter

To configure a multicast filter, you must have access to the CONFIG level of the CLI. You can configure up to eight multicast filters on a device.

Syntax: [no] multicast filter <filter-ID> any|ip udp mac <multicast-address>|any [mask <mask>]
[vlan <vlan-id>]

The parameter values are the same as the for the broadcast filter command. In addition, the multicast filter command requires the **mac <multicast-address> | any** parameter, which specifies the multicast address. Enter **mac any** to filter on all multicast addresses.

Enter **mac** followed by a specific multicast address to filter only on that multicast address. To filter on a range of multicast addresses, use the **mask <mask>** parameter. For example, to filter on multicast groups 0100.5e00.5200 – 0100.5e00.52ff, use **mask ffff.ffff.ff00**. The default mask matches all bits (is all Fs). You can leave the mask off if you want the filter to match on all bits in the multicast address.

Configuration Examples

To configure a Layer 2 multicast filter to filter all multicast groups, then apply the filter to ports 2/4, 2/5, and 2/8, enter the following commands:

```
HP9300(config)# multicast filter 1 any
HP9300(config-mcast-filter-id-1)# exclude-ports ethernet 2/4 to 2/5 ethernet 2/8
HP9300(config-mcast-filter-id-1)# write mem
```

To configure a multicast filter to block all multicast traffic destined for multicast addresses 0100.5e00.5200 – 0100.5e00.52ff on port 4/8, enter the following commands:

```
HP9300(config)# multicast filter 2 0100.5e00.5200 mask ffff.ffff.ff00
HP9300(config-mcast-filter-id-2)# exclude-ports ethernet 4/8
HP9300(config-mcast-filter-id-2)# write mem
```

The software calculates the range by combining the mask with the multicast address. In this example, all but the last eight bits in the mask are “significant bits” (ones). The last eight bits are zeros and thus match on any value. Each "f" or "0" is four bits.

Locking a Port To Restrict Addresses

Lock-address filters allow you to limit the number of devices that have access to a specific port. Access violations are reported as SNMP traps. By default this feature is disabled. A maximum of 2,048 entries can be specified for access. The default address count is eight.

USING THE CLI

EXAMPLE: To enable address locking for port 2 and place a limit of 15 entries:

```
HP6308(config)# lock e 2 addr 15
```

syntax: lock-address ethernet <port number> [addr-count <number>]

USING THE WEB MANAGEMENT INTERFACE

To enable address locking on a port:

1. Select the Port link to display the port summary panel. The Port link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left. A panel such as the one shown in Figure 8.9 is displayed.
2. Click the Modify button next to the port on which you want to enable address locking. A panel such as the one shown in Figure 8.9 is displayed.
3. Select Enable next to Lock Address.
4. Enter the maximum number of MAC addresses you want the device to learn on the port in the # of Address field.
5. Click the Apply button to assign the change.

NOTE: Before leaving the screen, you must click the Apply button for the port parameters to be assigned. Additionally, the configuration changes must be saved to flash (using the File menu option) for the changes to be preserved over a power cycle.

Configuring Basic Layer 3 Parameters

The procedures in this section describe how to enable the following Layer 3 protocols on the HP 9304M, 9308M, and 6308M-SX routing switches:

- IP
- IPX
- BGP4
- OSPF
- RIP
- DVMRP
- PIM
- AppleTalk
- VRRP
- SRP

By default, IP routing is enabled on routing switches. All other protocols are disabled, so you must enable them to configure and use them.

In some cases, you need to reset the system to activate a protocol. For specific details about this and configuration details on the protocols, refer to their individual chapters.

USING THE CLI

To enable a protocol, enter **router** at the global CONFIG level, followed by the protocol to be enabled. The following example shows how to enable OPSF:

```
HP9300(config)# router rip
HP9300(config)# end
HP9300# write mem
HP9300# reload
```

syntax: router appletalk|bgp|dvmrp|srp|ipx|ospf|pim|rip|vrrp

NOTE: The following protocols require a system reset before the protocol will be active on the system: PIM, DVMRP, RIP, and SRP. To reset a system, select the reset option under the File menu (Web) or enter the **reload** command at the privileged level of the CLI.

USING THE WEB MANAGEMENT INTERFACE

To enable protocols on a routing switch:

1. Select the System link to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Enable option next to the protocol(s) to be enabled.
3. Click the Apply button to assign the change.
4. Select the Save To Flash option from the File menu.
5. Select the Reset option from the File menu to reset the system.

   NOTE: The reset step is not required for the AppleTalk, OSPF, or BGP4 protocols to become active.

6. Click the Apply button to assign the change.

NOTE: Do not enable both SRP and VRRP. HP recommends that you use only one of these router redundancy protocols on a routing switch.

Configuring Layer 4 Quality of Service Parameters

This section describes how to assign Layer 4 sessions to Quality of Service (QoS) queues. QoS on the HP 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch is based on packet-based prioritization. When you assign a QoS priority to a port, you are selecting the queue in which the outbound packets for that port will be placed. Queues apply to outbound packets, not inbound packets. You can assign QoS priorities to the following system configuration areas:

- Ports – see "Modifying Port Priority (QoS)" on page 8-29.
- VLANs – see "Configuring VLANs" on page 17-1.
- Static MAC entries – see "Configuring Static MAC Entries" on page 8-35.
- Layer 4 sessions – see "Applying Layer 4 QoS Priority on the HP 6208M-SX Switch" on page 8-61 and "Applying Layer 4 QoS Priority on the HP 9304M, 9308M, or 6308M-SX Routing Switch" on page 8-63.
- AppleTalk sockets – see "AppleTalk QoS Socket" on page 16-25.

You can assign one of eight levels of priority, 0 – 7. Priority 0 is the lowest and 7 is the highest. The default priority is 0, the normal priority queue.

- 0 or 1 – These priorities correspond to queue 0, the normal priority queue.
- 2 or 3 – These priorities correspond to queue 1, a higher priority queue.
- 4 or 5 – These priorities correspond to queue 2, a higher priority queue.
- 6 or 7 – These priorities correspond to queue 3, the highest priority queue.

See "Quality of Service Algorithm" on page C-1 for additional information about the QoS algorithms.

Applying Layer 4 QoS Priority on the HP 6208M-SX Switch

The following sections show how to configure Layer 4 QoS policies on a 6208M-SX switch.

NOTE: The CLI syntax and Web management options for configuring Layer 4 QoS polices on the HP 9304M, 9308M, and 6308M-SX routing switches are different from those for the 6208M-SX switch. Make sure you use the appropriate set of instructions for the device you are configuring.

USING THE CLI

To assign a priority of 7 to FTP traffic on all ports on a 6208M-SX switch, enter the following commands:

```
HP6208(config)# ip policy 1 7 tcp ftp global
HP6208(config)# write mem
```

To assign a priority of 7 to HTTP traffic on ports 1 and 2 only, enter the following commands:

```
HP6208(config)# ip policy 2 7 tcp http local
HP6208(config)# int ethernet 1
HP6208(config-if-1)# ip-policy 2
HP6208(config-if-1)# int 2
HP6208(config-if-2)# ip-policy 2
HP6208(config)# write mem
```

Syntax: ip policy <num> priority <0-7> tcp|udp <tcp/udp-port-num> global|local

ip-policy <num>

The **<num>** parameter is the policy number.

The **priority <0-7>** parameter specifies the QoS priority level. The default is 0 (normal priority). The highest priority is 7.

The **tcp|udp <tcp/udp-port-num>** parameter specifies the TCP or UDP port to which you are applying the policy

The **global** and **local** parameters specify the scope of the policy:

- If you specify **global**, the policy applies to all ports.
- If you specify **local**, the policy will apply to the ports you specify. Use the following command on the Interface level of the CLI to apply the policy to a port: **ip-policy <num>**

Figure 8.23 shows the CLI syntax for configuring a QoS policy on the 6208M-SX switch. The value "<CR>" means "carriage return", also known as the Enter key.

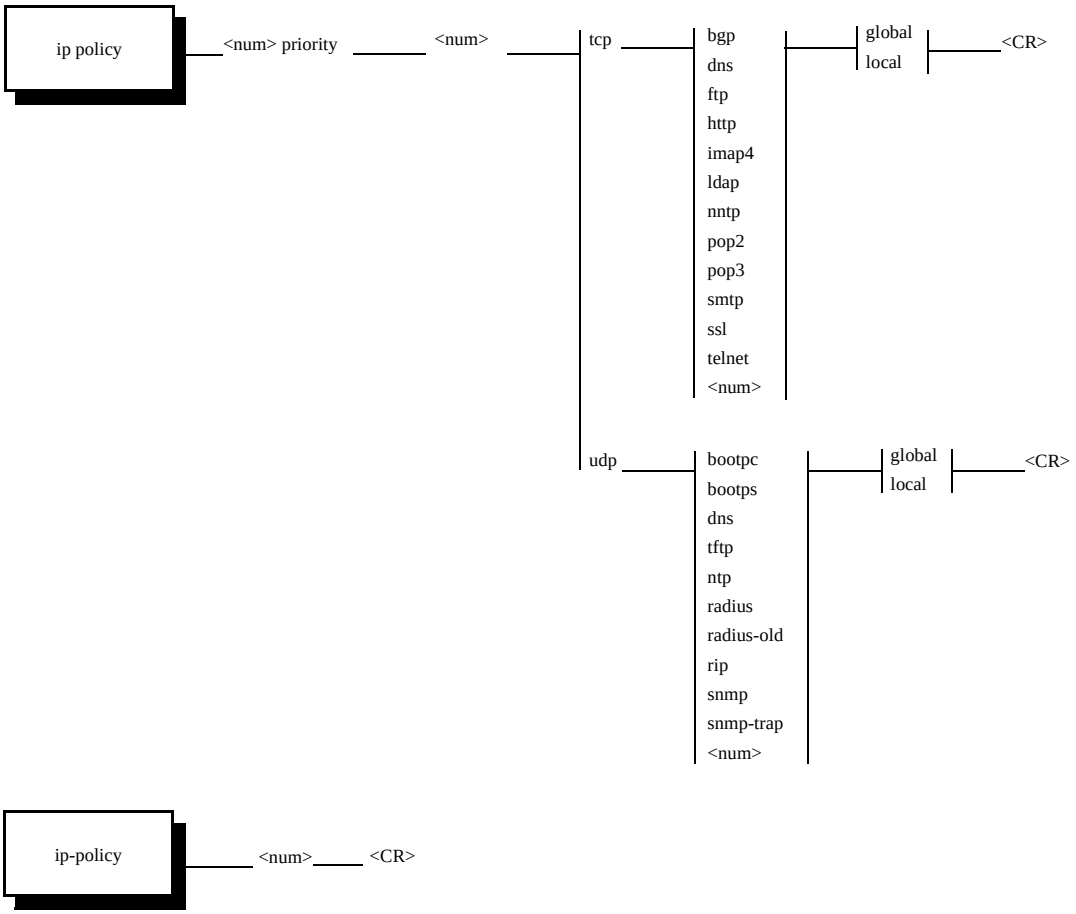

Figure 8.23 QoS IP policy syntax for the 6208M-SX switch

NOTE: The **ip policy** command allows you to configure global or local QoS policies. Use the **ip-policy** command (note the difference between "**ip policy**" and "**ip-policy**") at the Interface level of the CLI to apply a local policy to a specific interface.

USING THE WEB MANAGEMENT INTERFACE

To apply Layer 4 QoS priority on the 6208M-SX switch:

1. Select System to display the System configuration sheet.
2. Select the Layer 4 QoS link from the System configuration sheet to display a panel such as the one shown in Figure 8.24.
3. Enter the QoS policy ID in the ID field.
4. Select a priority scope of either Switch or Port. Selecting Switch applies the defined priority to all interfaces and port to a defined port. In this case, select Switch.

   NOTE: If you select Port as the scope, you also need to configure QoS on the port level by selecting the Port QoS link.

5. Select the priority to be assigned from the QoS field's pulldown menu.
6. Select the Layer 4 protocol type, which can be TCP or UDP. For this example, select TCP.
7. Select the TCP or UDP port from the TCP/UDP port field's pulldown menu. In this case, select FTP.
8. Click the Add button to assign the changes.

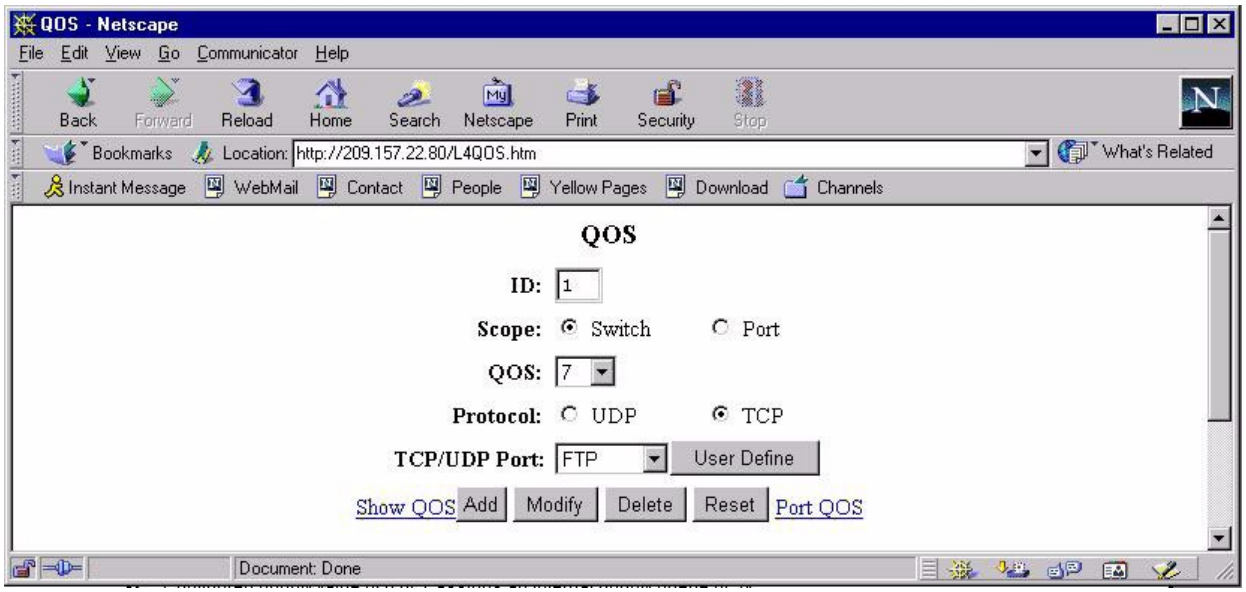

Figure 8.24 Layer 4 QoS entry panel on the 6208M-SX switch

Applying Layer 4 QoS Priority on the HP 9304M, 9308M, or 6308M-SX Routing Switch

To apply Layer 4 QoS on a routing switch, configure an IP access policy globally, then apply the policy to a port's outbound policy group.

NOTE: The CLI syntax and Web management options for configuring Layer 4 QoS polices on the HP 9304M, 9308M, and 6308M-SX routing switches are different from those for the 6208M-SX switch. For example, if you configure a global QoS policy on a switch, the policy is automatically applied. On a routing switch, you must configure the policy globally, then apply it to specific ports. Make sure you use the appropriate set of instructions for the device.

USING THE CLI

To assign a priority of 4 to all HTTP traffic on port 3/12 on a 9304M or 9308M, enter the following:

```
HP9300(config)# ip access-policy 1 priority 4 any any tcp eq http
HP9300(config)# int e 3/12
HP9300(config-if-3/12)# ip access-policy-group out 1
```

Syntax: ip access-policy <num> priority <0-7> <ip-addr> <mask>|any <ip-addr> <mask>|any icmp|igmp|igrp|ospf|tcp|udp|<num> [<operator> [tcp/udp-port-num>]]

ip access-policy-group in|out <policy-list>

NOTE: For backward compatibility, the routing switches also support the **ip filter** and **ip policy** commands. The parameters are the same as those for the **ip access-policy** command.

The **<num>** parameter is the policy number.

The **priority <0-7>** and **high|normal** parameters specify the QoS priority level. The default is 0 (normal priority). The highest priority is 7.

The **<ip-addr> <mask>|any <ip-addr> <mask>|any** parameters specify the source and destination IP addresses. If you specify a particular IP address, you also need to specify the mask for that address. If you specify **any** to apply the policy to all source or destination addresses, you do not need to specify **any** again for the mask. Make sure you specify a separate address and mask or **any** for the source and destination address.

The **icmp|igmp|igrp|ospf|tcp|udp|<num>** parameter specifies the Layer 4 port to which you are applying the policy. If you specify **tcp** or **udp**, you also can use the optional **<operator>** and **<tcp/udp-port-num>** parameters to fine-tune the policy to apply to specific TCP or UDP ports.

The **<operator>** parameter applies only if you use the **tcp** or **udp** parameter above. Use the **<operator>** parameter to specify the comparison condition for the specific TCP or UDP ports. For example, if you are configuring QoS for HTTP, specify **tcp eq http**. You can enter one of the following operators:

- **eq** – The policy applies to the TCP or UDP port name or number you enter after **eq**.
- **gt** – The policy applies to TCP or UDP port numbers greater than the port number or the numeric equivalent of the port name you enter after **gt**.
- **lt** – The policy applies to TCP or UDP port numbers that are less than the port number or the numeric equivalent of the port name you enter after **lt**.
- **neq** – The policy applies to all TCP or UDP port numbers except the port number or port name you enter after **lt**.
- **range** – The policy applies to all TCP or UDP port numbers that are between the first TCP or UDP port name or number and the second one you enter following the range parameter. The range includes the port names or numbers you enter. For example, to apply the policy to all ports between and including 23 (Telnet) and 53 (DNS), enter the following: **range 23 53**. The first port number in the range must be lower than the last number in the range.
- **established** – This operator applies only to TCP packets. If you use this operator, the QoS policy applies to TCP packets that have the ACK (Acknowledgment) or RST (Reset) bits set on (set to "1") in the Control Bits field of the TCP packet header. Thus, the policy applies only to established TCP sessions, not to new sessions. See Section 3.1, "Header Format", in RFC 793 for information about this field.

Figure 8.25 and Figure 8.26 show the CLI syntax for configuring a Layer 4 QoS policy on a routing switch.

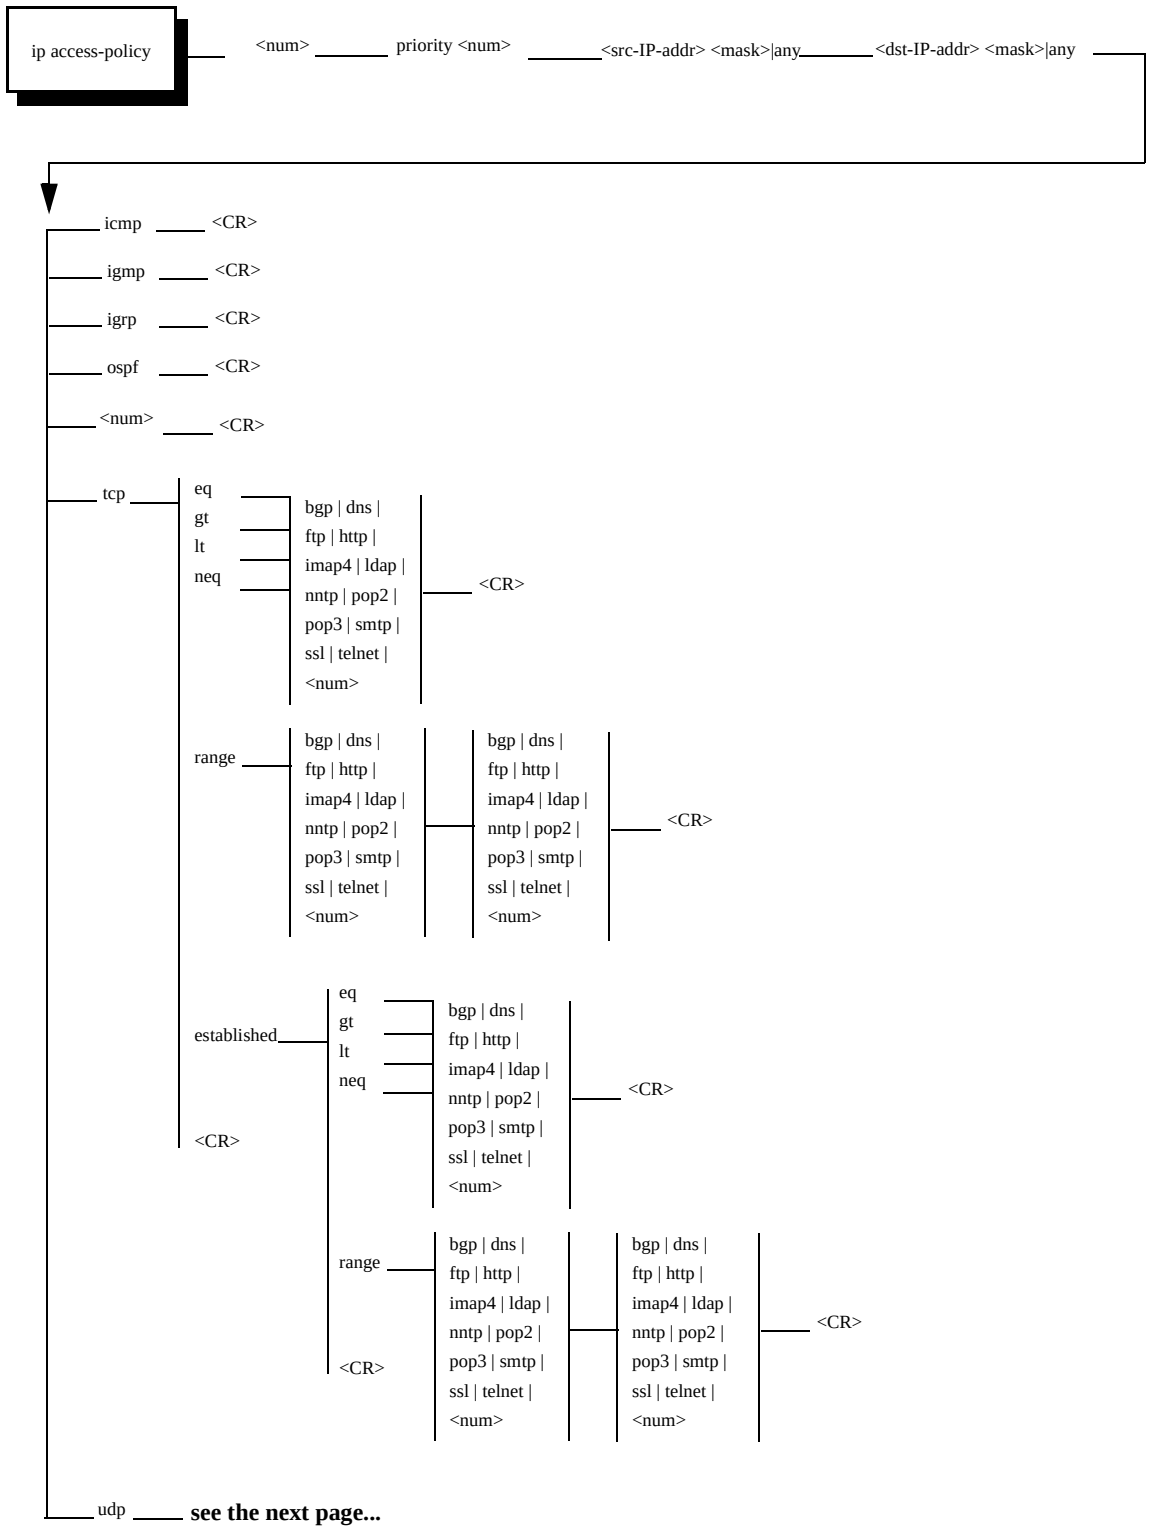

Figure 8.25 QoS IP policy syntax for a routing switch (1 of 2)

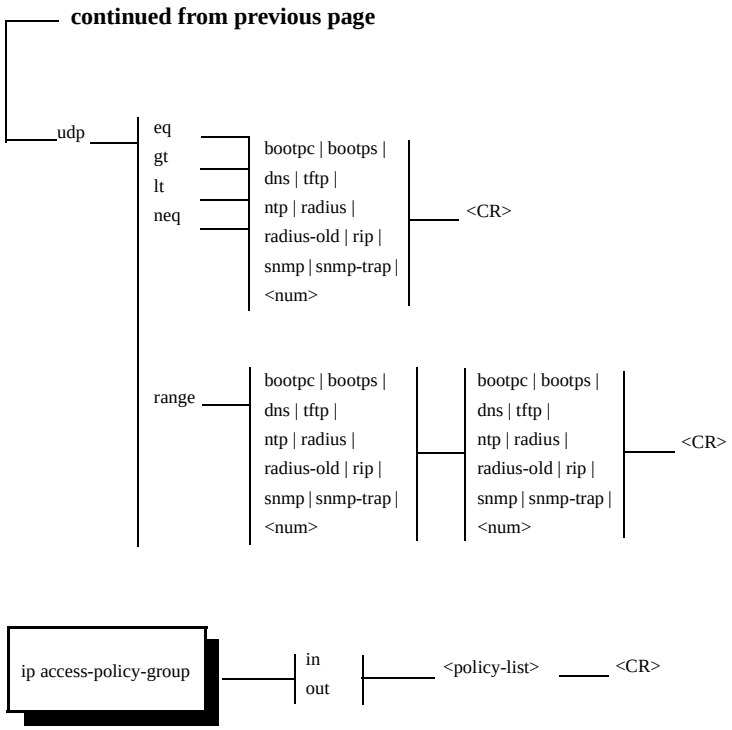

Figure 8.26 QoS IP policy syntax for a routing switch (2 of 2)

USING THE WEB MANAGEMENT INTERFACE

To apply Layer 4 QoS priority on the 9304M, 9308M, or 6308M-SX routing switch:

1. Select the IP link to display the IP panel.
2. Select IP Access Policy to display the panel shown in Figure 8.27.

NOTE: If you have already defined access policies, a table listing the policies is displayed instead. Select Add IP Access Policy.

---

3. Enter the policy ID in the ID field.
4. Select QoS next to Action.
5. Enter the source IP address and mask for the traffic flow you are prioritizing. A flow consists of Layer 3 source and destination IP addresses and the Layer 4 destination protocol port. Thus, you can set priority for a specific host. To set the Layer 4 QoS policy for all hosts, leave "0.0.0.0" in the Source Address, Source Mask, Destination Address, and Destination Mask fields.
6. Enter the Layer 4 protocol in the Protocol field. You can enter the protocol number or one of the following:
   - ICMP
   - IGMP
   - IGRP
   - OSPF
   - TCP
   - UDP

7. For the TCP and UDP protocols, select the comparison operator from the Operator field. You can select one of the following:
   - Greater – The policy applies to TCP or UDP port numbers greater than the port number or the numeric equivalent of the port name you specify.
   - Equal – The policy applies to the TCP or UDP port name or number you specify.
   - Less – The policy applies to TCP or UDP port numbers that are less than the port number or the numeric equivalent of the port name you specify.
   - Not Equal – The policy applies to all TCP or UDP port numbers except the port number or port name you specify.
   - Established (applies only to TCP) – This operator applies only to TCP packets. If you use this operator, the policy applies to TCP packets that have the ACK (Acknowledgment) or RST (Reset) bits set on (set to "1") in the Control Bits field of the TCP packet header. Thus, the policy applies only to established TCP sessions, not to new sessions. See Section 3.1, "Header Format", in RFC 793 for information about this field.
   - Range – The policy applies to all TCP or UDP port numbers that are between the first TCP or UDP port name or number and the second one you specify. The range includes the port names or numbers you enter. For example, to apply the policy to all ports between and including 23 (Telnet) and 53 (DNS), specify the following: 23 53. The first port number in the range must be lower than the last number in the range.
8. For the TCP and UDP protocols, enter the port number in the TCP/UDP port field.
9. Click the Add button to add the filter. The filter is configured but does not take effect until you add the filter to a port's IP access policy group.
10. Select Access Policy Group to display the Access Policy Group panel, shown in Figure 8.28.

NOTE: If you have already configured IP access policy groups, the groups are listed in a table. Select Add IP Access Policy Group.

---

11. Select the port (and slot, for chassis devices) to which you are applying the Layer 4 QoS policy.
12. Select Out Filter.

NOTE: QoS policies apply only to outbound traffic. The IP Access Policy Group panel contains an In Filter option because you also can configure IP Access Policies to permit or deny Layer 3 and Layer 4 traffic. See "Policies and Filters" on page D-1.

---

13. Enter the Layer 4 QoS policy number in the Filter ID List field. In this example, enter "1".
14. Select Add to add the filter to the port's IP access policy group.

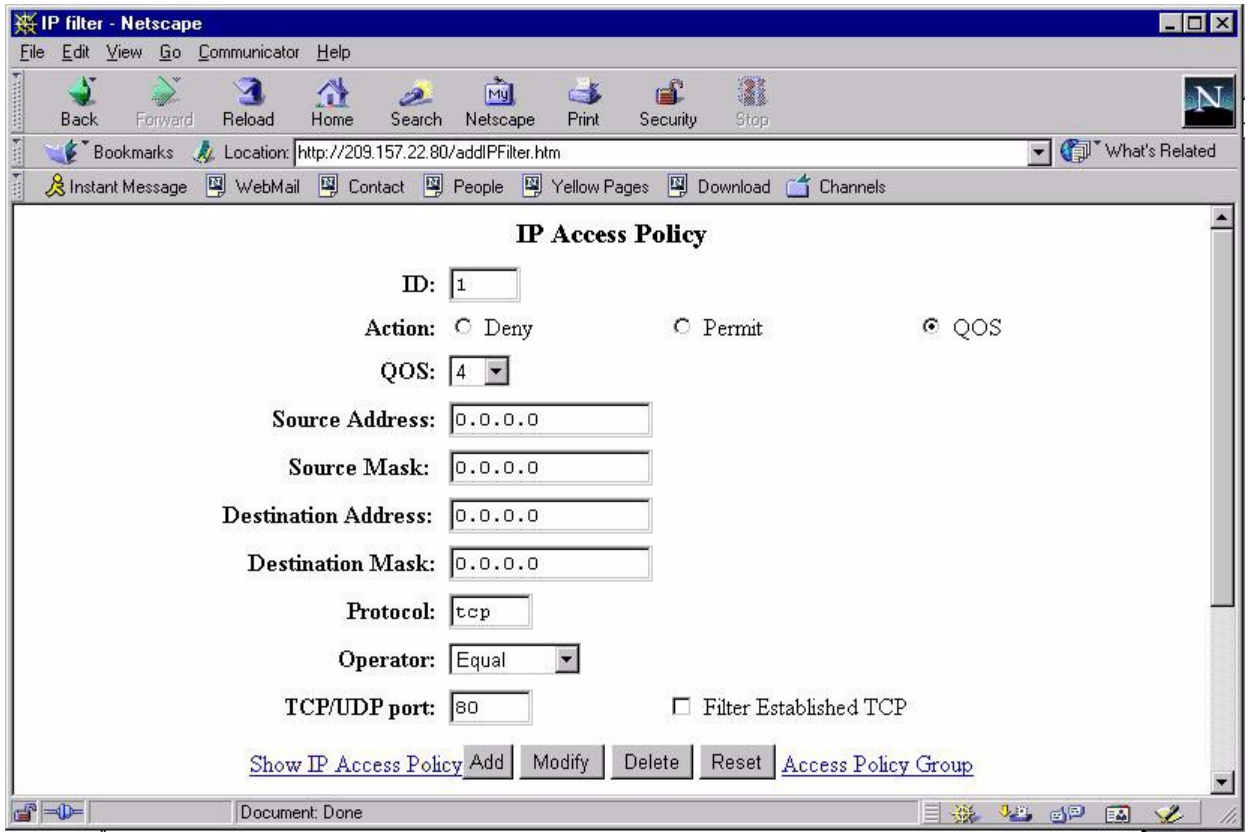

Figure 8.27 IP Access Policy panel

Figure 8.28 IP Access Policy Group panel

Modifying System Parameter Default Settings

The 9304M, 9308M, 6308M-SX routing switches and the 6208M-SX switch have default table sizes for the following parameters. The table sizes determine the maximum number of entries the tables can hold. You can adjust individual table sizes to accommodate your configuration needs.

- MAC address entries
- Layer 2 Port VLANs supported on a system
- Layer 3 Protocol VLANs supported on a system
- Layer 4 sessions supported
- IP cache size
- ARP entries
- IP routes
- IP route filters
- IP sub-nets per port and per device
- Static routes
- IGMP
- DVMRP routes
- IPX/SAP entries
- IPX/RIP entries
- IPX/SAP filters
- IPX/RIP filters
- IPX forwarding filters
- AppleTalk routes
- AppleTalk zones

The tables you can configure and the defaults and valid ranges for each table differ depending on the device you are configuring.

To display and configure the adjustable tables on a device, use one of the following methods.

NOTE: Changing the table size for a parameter reconfigures the device's memory. Whenever you reconfigure the memory on a device, you must save the change to the startup-config file, then reload the software to place the change into effect.

USING THE CLI

To display the configurable tables and their defaults and maximum values, enter the following command at any level of the CLI:

```
HP9300# show default values

sys log buffers:50         mac age time:300 sec       telnet sessions:5

ip arp age:20 min          bootp relay max hops:4     ip ttl:64 hops
ip addr per intf:24

when multicast enabled :
igmp group memb.:140 sec   igmp query:60 sec

when ospf enabled :

ospf dead:40 sec           ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec

when bgp enabled :
bgp local pref.:100        bgp keep alive:60 sec      bgp hold:180 sec
bgp metric:10              bgp local as:1             bgp cluster id:0
bgp ext. distance:20       bgp int. distance:200      bgp local distance:200

System Parameters    Default    Maximum
arp                  4000       16000
atalk-route          512        3072
atalk-zone-port      64         255
atalk-zone-sys       255        1024
dvmrp                2048       32000
igmp                 255        1024
ip-cache             16000      64000
ip-filter-port       32         256
ip-filter-sys        64         2048
ipx-forward-filter   32         256
ipx-rip-entry        2048       16384
ipx-rip-filter       32         256
ipx-sap-entry        4096       16384
ipx-sap-filter       32         256
l3-vlan              32         1024
ip-qos-session       128        32000
mac                  8000       64000
ip-route             10000      200000
ip-static-route      64         1024
vlan                 8          4096
mac-filter-port      16         256
mac-filter-sys       32         512
subnet-per-interface 24         64
subnet-per-system    256        512
```

Information for the configurable tables appears under the columns that are shown in bold type in this example. To simplify configuration, the command parameter you enter to configure the table is used for the table name. For example, to increase the capacity of the IP route table, enter the following commands:

```
HP9300(config)# system-max ip-route 120000

HP9300(config)# write mem

HP9300(config)# exit

HP9300# reload
```

NOTE: If you accidentally enter a value that is not within the valid range of values, the CLI will display the valid range for you.

USING THE WEB MANAGEMENT INTERFACE

To modify a table size using the Web management interface:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Parameter link on the System configuration sheet to display a panel such as the one shown in Figure 8.29. Depending on the device you are configuring, you might need to scroll down to see all the tables.
3. Click the Modify button next to the table you want to change.
4. Enter the new value for the table size.
5. Click the Add button to assign the changes.
6. Select the Save To Flash link, then click Yes when the interface asks you whether you want to save the configuration changes.
7. Select the Reload link to reload the software and place the change into effect.

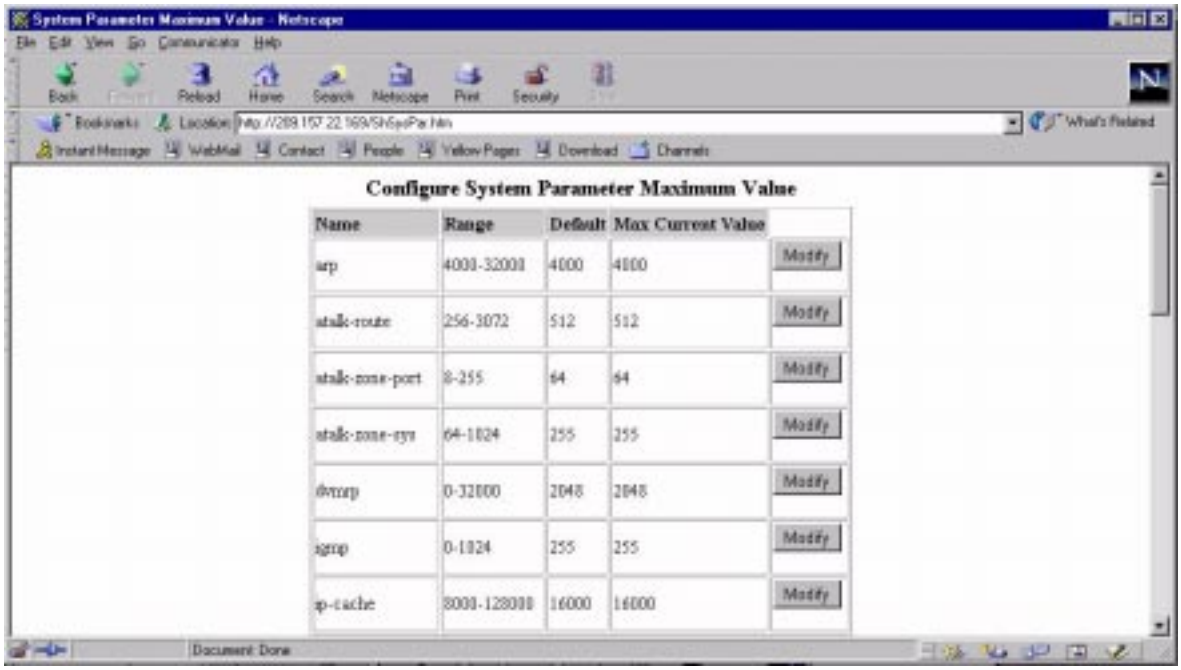

Figure 8.29 Table parameter configuration panel for a routing switch (scroll to display additional parameters)

Assigning a Mirror Port and a Monitor Port

This field assigns a port on the switch or routing switch to act as a diagnostic port. Once the mirror port is assigned, an external protocol analyzer can be attached to this port to monitor the activity of another port on the system.

Mirroring can be done on input or output traffic or on both input and output (full-duplex) traffic and any of the switch or routing switch ports, including the expansion ports, can operate as a mirror port. The default value for this field is none, meaning no mirror port is assigned.

Monitoring traffic on a port is a two step process:

- Enable a port to act as the mirror port.
- Identify the port on which the traffic is to be monitored (the monitor port).

USING THE CLI

EXAMPLE: Suppose you want to diagnose the input and output on traffic on port 3 on a module in slot 4 of a 9308M using port 1 in slot 4. To do so, enter the following:

```
HP9300(config)# mirror-port e 1/4

HP9300(config)# interface e 4/3

HP9300(config-if-4/3)# monitor both
```

syntax: mirror-port ethernet <port-num>

NOTE: To monitor just the input traffic, enter "in" instead of "both" in the above command. To monitor only the output traffic, enter "out" instead of "both" in the above command.

USING THE WEB MANAGEMENT INTERFACE

EXAMPLE: Suppose you want to diagnose the input and output on traffic on port 3 on a module in slot 4 of a 9308M using port 1 in slot 4. To do so:

1. Select System to display the System configuration sheet. The System link is located at the bottom of the main window. If frames are enabled in the Web management interface, the link also is available in the column of options on the left.
2. Select the Advance link at the bottom of the System configuration sheet.
3. Select the slot (if applicable) and port from the corresponding pulldown menus. In this example, select slot 4 and port 1.
4. Click the Apply button to assign the change.
5. Select the Port link to display the port summary panel.
6. Click the Modify button next to the port you want to monitor. In this example, select port 3 on the module in slot 4 (4/3).
7. Select Both from the Monitoring pulldown menu to initiate diagnosis on both input and output traffic on port 4/3.
8. Click the Apply button to assign the change.