Chapter 9
Configuring IP and IP/RIP

This chapter describes how to configure the IP and IP/RIP protocols on the HP ProCurve 9304M, 9308M, and 6308M-SX routing switches using the CLI and Web management interface.

NOTE: IP routing and IP/RIP are supported only on the routing switches, not on the 6208M-SX switch.

To display IP and RIP configuration information and statistics, see “Configuring IP and IP/RIP” on page 9-3.

For complete syntax information for the CLI commands shown in this chapter, see “Command Line Interface Commands” on page B-1.

NOTE: 9304M and 9308M routing switches that use Redundant Management modules can contain a maximum of 80000 IP routes by default. The 6308M-SX and chassis devices that use other management modules can contain a maximum of 10000 IP routes by default. If you need to increase the capacity of the IP route table for BGP4, see “Modifying System Parameter Default Settings” on page 8-69.

Overview of IP/RIP

IP/RIP is a distance-vector protocol. IP/RIP routers transmit and receive RIP updates to and from neighboring routers. By default, the routing switches send RIP updates every 30 seconds. You can change the update interval and many other IP and IP/RIP parameters if needed.

The routing switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the ***IP route table*** as the route to the destination. Typically, the best path is the path with the fewest ***hops***. A hop is another router through which packets must travel to reach the destination. If the routing switch receives a RIP update from another router that contains a path with fewer hops than the path stored in the route table, the routing switch replaces the older route with the newer one. The routing switch then includes the new path in the updates it sends to other RIP routers.

Each entry in the IP/RIP routing table includes the destination address, the next hop address, and a ***metric***. The metric is equal to the number of hops required to reach a destination.

The IP/RIP protocol on the 9304M, 9308M, and 6308M-SX routing switches supports the following RIP types:

- Version 1
- V1 compatible with V2
- Version 2 (the default)

IP/RIP Features

RIP includes a number of features that help stabilize its performance in rapidly changing network conditions. These features include ***hop count limits***, ***hold downs***, ***split horizons***, and ***poison reverse updates***.

Hop Count Limit

A maximum of 15 hops is supported by IP/RIP. Any destination that is greater than 15 hops away is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

Hold Downs

A ***hold-down*** instructs routers to delay (hold down) action update messages received from routes that may be inactive. The period of time is generally longer than the time required to update the entire network with a routing change. This safeguard prevents an inactive route from being reinstated.

Split Horizons

Split horizons prevent routing loops from being generated by adjacent routers. This feature is useful when a router's path to a given router is through another router. Split horizons allow a routing broadcast to be modified so that routers with intermediate routers in their path to a destination router, are not seen as a path to the destination router by the intermediate router.

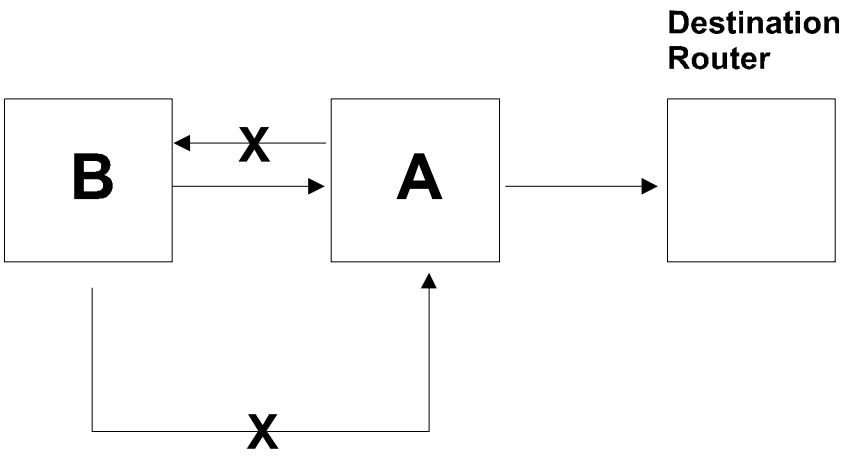

Figure 9.1 Split horizon in action

For example, in Figure 9.1, without split horizon operating, router A could see router B as a path to router X. However, if A were to route to B to reach router X, a loop would occur. A split horizon modifies a routing broadcast so that the intermediate router does not treat the source router as a path to the destination router. In Figure 9.1, the link with an "X" over it indicates a loop that is prevented by the split-horizon feature.

Poison Reverse Updates

Poison reverse updates are used to prevent larger loops within the network by setting the metric (cost) of neighboring routes to infinity. This will prevent two-hop loops.

IP/RIP Default Route Learning and Advertising

The 9304M, 9308M, and 6308M-SX routing switches can learn and advertise default IP/RIP routes. This feature can be enabled on a global or interface basis. By default, this feature is disabled.

Priority for learning of IP/RIP routes is in the following order:

1. Static IP/RIP routes.
2. IP/RIP routes learned from RIP.
3. IP/RIP routes learned from OSPF.

ICMP Host Unreachable Message for Undeliverable ARPs

If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

Configuring IP and IP/RIP

By default, the IP protocol is active on the 9304M, 9308M, and 6308M-SX routing switches at initial start-up, so there is no need to enable the protocol. However, you do need to assign IP addresses.

Static routes, IP access policies (sometimes called "IP filters"), and the UDP helper feature are components of the IP protocol. Additionally, the protocol comes with system (global) and interface level parameters that you can modify to better suit the needs of the network.

The following actions can be done at the IP and RIP levels of the CLI or from the IP and RIP configuration sheets of the Web management interface:

1. Enable IP/RIP.
2. Assign IP addresses to routing switch interfaces.
3. Modify global IP parameters (optional).
4. Modify interface IP parameters (optional).
5. Define static IP routes (optional).
6. Assign Static ARP and RARP entries (optional).
7. Define IP filters (optional).
8. Configure UDP helper (optional).
9. Define IP/RIP route filters (optional).
10. Define IP/RIP route filter groups (optional).
11. Modify the RIP global default parameters–metric value, update time parameters (optional).
12. Configure redistribution filters, if non-RIP routes are to be imported into RIP.
13. Modify or enable interface parameters—RIP type or poison reverse (optional).

Dynamic IP/RIP Configuration

This feature allows a routing switch to apply key IP/RIP configuration changes immediately without requiring a system reset. Here is a summary of those parameters:

- Enabling or disabling of RIP
- Adding a static route
- Enabling RARP or Proxy ARP
- Adding static ARP or RARP entries
- Setting the ARP cache aging value
- Enabling ICMP Router Discovery Protocol (IRDP)
- Adding a Relay BootP server address
- Setting RIP transmit intervals
- Assignment of RIP type—V1, V2 or V1/V2 compatible
- Activating RIP poison reverse

Enabling IP/RIP

The IP/RIP protocol is disabled by default. It must be enabled on the routing switch, and the system must be reset before you can use the protocol.

USING THE CLI

To enable RIP on a routing switch, enter the following commands:

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

syntax: router rip

NOTE: In the above example, the system is reset to enable the IP/RIP protocol. HP recommends that you configure all elements of the protocol before you reset the system.

USING THE WEB MANAGEMENT INTERFACE

1. Select the System link from the main menu.
2. Select the checkbox next to RIP.
3. Select the Save To Flash link from the main menu.
4. Select the Reload option from the main menu.

Assigning IP Addresses

Before attaching equipment to the routing switch, you must assign individual sub-net IP addresses and masks for each of the ports based on the desired and current network topology.

By default, you can configure up to 24 IP interfaces on each port, virtual interface, and loopback interface. The 9304M, 9308M, and 6308M-SX routing switches support both classical IP network masks (Class A, B, and C sub-net masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks.

- To enter a classical network mask, enter the mask in IP address format. For example, enter "209.157.22.99 255.255.255.0" for an IP address with a Class-C sub-net mask.
- To enter a prefix network mask, enter a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter "209.157.22.99/24" for an IP address that has a network mask with 24 significant bits (ones).

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the display to prefix format. See “Changing Network Mask Displays to Prefix Format” on page 9-10.

USING THE CLI

To assign an IP address for interface 1, enter the following commands:

```
HP9300(config)# interface ethernet 1/1
HP9300(config-if-1/1)# ip address 192.45.6.1 255.255.255.0
```

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

or

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

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

NOTE: You also can enter the IP address and mask in the following manner:
`HP9300(config-if-1/1)# ip address 192.45.6.1/24`

NOTE: Before exiting the Interface level of the CLI to configure IP interfaces on other routing switch ports, configure the remaining parameters for the IP interface. For details on configuring IP interface parameters, see “Modify IP and IP/RIP Interface Parameters (optional)” on page 9-26.

USING THE WEB MANAGEMENT INTERFACE

To assign an IP address:

1. Select the IP Address link from the IP configuration sheet. The panel shown in Figure 9.2 will appear.

   NOTE: If at least one IP address is already defined on the system, then a summary panel appears first. Select the Add IP Address link.

2. Select the port or slot/port combination that the address is to be assigned.
3. Enter the IP address of the sub-net.
4. Enter the sub-net mask.
5. Select the Secondary box if the IP address being defined is not the first address assigned to this interface.
6. Click the Add button to add the new IP address.

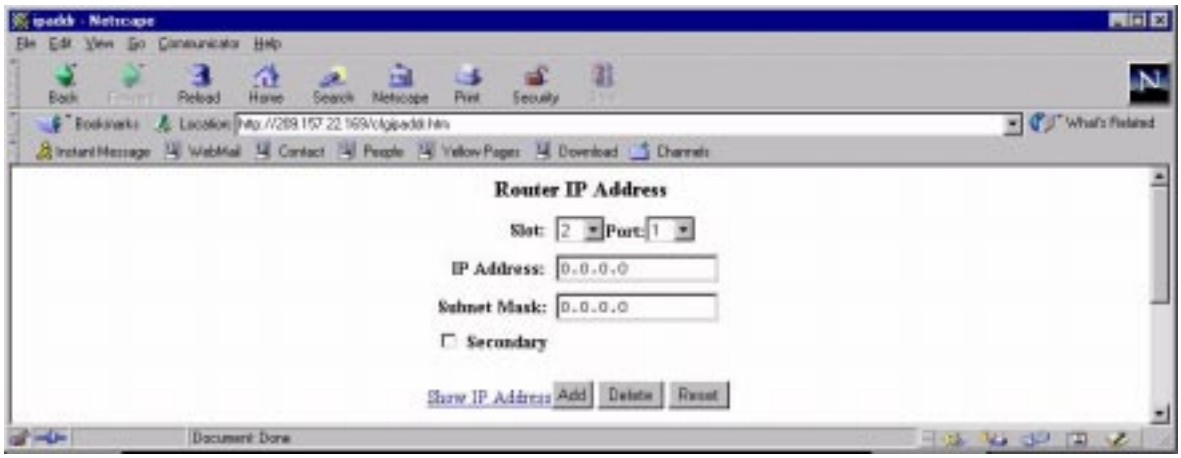

Figure 9.2 Assigning an IP address to an interface

Modifying Global IP and IP/RIP Parameters (optional)

Many IP/RIP parameters can be modified for the IP protocol on a global basis. Each of these parameters comes with a default setting and does not need to be modified unless your network configuration requires a change.

You can perform the following parameter configuration tasks:

- Modify the maximum number of hops for a BootP Relay server.
- Modify the ARP aging period.
- Modify the time-to-live (TTL) threshold.
- Enable or disable RDP.
- Enable or disable load sharing.
- Enable or disable proxy ARP.
- Enable or disable RARP.
- Configure global static ARP or RARP entries.

- Configure static IP routes.
- Configure IP access policies (IP forwarding filters).
- Enable or disable broadcast forwarding UDP Helper).
- Disable or re-enable directed broadcast forwarding.
- Change the display format for network masks to prefix format (CLI only).

Figure 9.3 shows the IP configuration sheet in the Web management interface. You can change many of the IP parameters using this display.

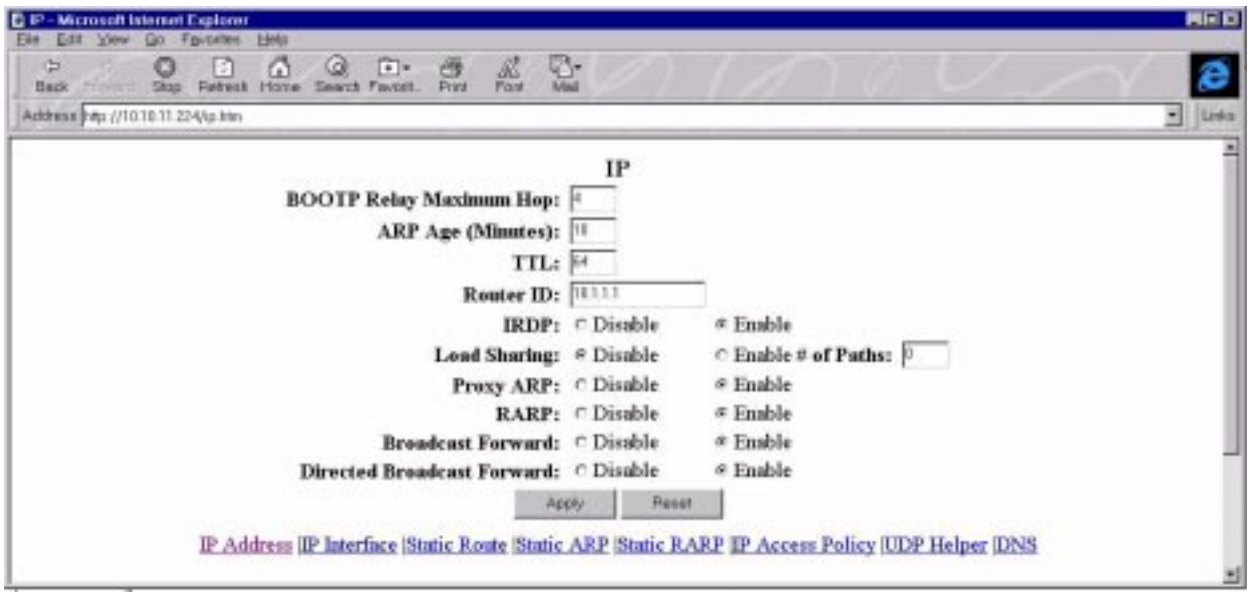

Figure 9.3 IP configuration sheet

Modifying the Maximum Number of Hops to a BootP Relay Server

The 9304M, 9308M, and 6308M-SX routing switches can support the relay of BootP requests to a BootP server outside of its network. You can modify the maximum number of hops that a request will traverse to a BootP server. The parameter value ranges from 1 – 15 hops. The default value is 4 hops.

USING THE CLI

To modify the maximum number of hops supported, enter the following command:

```
HP9300(config)# bootp-relay-max-hops 10
```

syntax: bootp-relay-max-hops <1-15>

USING THE WEB MANAGEMENT INTERFACE

To modify the maximum number of hops supported:

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Enter a value from 1 – 15 into the BootP Relay Maximum Hop field.
3. Select the Apply button to assign the changes.

Modifying the ARP Aging Period

The ARP aging period defines how long an inactive ARP entry remains in the ARP cache before the routing switch ages out the entry. The parameter value ranges from 0 – 240 minutes. If you enter 0, aging is disabled. The default value is 10 minutes.

USING THE CLI

To modify the ARP aging parameter to 20 minutes, enter the following command:

```
HP9300(config)# ip arp-age 20
```

syntax: ip arp-age <0-240>

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Enter a value from 0 – 240 into the ARP Age field.
3. Select the Apply button to assign the changes.

Modifying the tTTL Threshold

This parameter defines how long a packet will remain alive on the network. The range is from 1 – 255 hops. The default value for this parameter is 64 hops.

USING THE CLI

To modify the TTL threshold to 25, enter the following commands:

```
HP9300(config)# ip ttl 25
HP9300(config)# exit
```

syntax: ip ttl <1-255>

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Enter a value from 1 – 255 into the TTL field.
3. Select the Apply button to assign the changes.

Changing the Router ID

The OSPF and BGP4 protocols use router IDs to identify the routers that are running the protocols. A router ID is a valid, unique IP address and sometimes, is an IP address configured on the router. The router ID cannot be an IP address in use by another device. By default, the router ID is the lowest IP address configured on the routing switch. However, you can set the router ID to any valid IP address.

NOTE: The routing switches use the same router ID for both OSPF and BGP4. If the routing switch is already configured for OSPF, you may want to use the router ID that is already in use on the routing switch rather than set a new one. To display the router ID, enter the **show ip** CLI command at any CLI level or select the IP link in the Web management interface.

USING THE CLI

To set the router ID, enter a command such as the following:

```
HP9300(config)# ip router-id 209.157.22.26
```

Syntax: ip router-id <ip-addr>

The <ip-addr> can be any valid, unique IP address.

NOTE: You can specify an IP address used for an interface on the routing switch, but do not specify an IP address in use by another device.

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link to display the IP configuration sheet, which is shown in Figure 9.3.
2. Edit the value in the Router ID field to any valid IP address not in use on another router.
3. Click the Apply button to assign the change.

Enabling or Disabling IRDP

IIRDP allows routers to dynamically learn about routes on other networks. The routing switch advertises its IP addresses to other routers on the network and answer queries from those routers. The default value for this feature is enabled.

USING THE CLI

To disable IRDP on a routing switch, enter the following command:

```
HP9300(config)# no ip irdp
```

To re-enable IRDP on a routing switch, enter the following command:

```
HP9300(config)# ip irdp
```

syntax: [no] ip irdp

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Select IRDP.
3. Select the Apply button to assign the changes.

Enable or Disable Suppression of Directed Broadcasts

The H9304M, 9308M, and 6308M-SX routing switches allow directed IP broadcast forwarding by default, per section 5.3.5.2 in RFC 1812. However, if you want to suppress these directed broadcasts, you can do by entering the following command at the CONFIG level of the CLI:

```
HP9300(config)# no ip directed-broadcast
```

syntax: [no] ip directed-broadcast

The software makes the forwarding decision based on the routing switch's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last hop router.

Directly attached network or sub-network broadcast forwarding can be suppressed on the routing switches. Thus, you have the option to suppress directed broadcasts on directly attached networks or sub-networks on a global or per interface level.

To enable the suppression of directed broadcasts, enter the following command in the CONFIG mode:

```
HP9300(config)# ip directed-broadcast
```

Enable or Disable Load Sharing

Load sharing allows traffic to be sent across multiple paths of equal cost to a destination, resulting in a faster transmission. This feature is available when using the OSPF routing protocol. This feature is by default disabled.

NOTE: For information about configuring OSPF, see “Configuring OSPF” on page 10-1.

USING THE CLI

To enable load sharing for OSPF, enter the following command:

```
HP9300(config)# ip load-sharing [<num>]
```

syntax: [no] ip load-sharing [<num>]

You can specify from 2 – 8 paths. The default is 4.

See “Enable Load Sharing” on page 10-21 for more information about this feature.

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu to display the panel shown in Figure 9.3.
2. Enable the Load Sharing option.
3. Select the Apply button to assign the changes.

Disabling or Enabling Proxy ARP

Proxy ARP enables or disables a routing switch as proxy for devices on its sub-nets. As proxy, the routing switch responds to ARP requests from other devices on the network. By default, this feature is enabled on the routing switch.

USING THE CLI

To disable the proxy ARP function on the routing switch, enter the following command:

```
HP9300(config)# no ip proxy-arp
```

To re-enable the proxy ARP function on the routing switch, enter the following command:

```
HP9300(config)# ip proxy-arp
```

syntax: [no] ip proxy-arp

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu to display the panel shown in Figure 9.3.
2. Enable the Proxy ARP option.
3. Select the Apply button to assign the changes.

Enable or Disable RARP

You can enable or disable Reverse Address Resolution Protocol (RARP) on the routing switch. RARP allows retrieval of an IP address associated with a given MAC address. By default this feature is enabled.

USING THE CLI

To enable the RARP function on the routing switch, enter the following command:

```
HP9300(config)# ip rarp
```

syntax: [no] ip rarp

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Enable the RARP option.
3. Select the Apply button to assign the changes.

Enabling or Disabling Broadcast Forward

Broadcast forward allows the routing switch to make UDP helper assignments. Broadcast forward is used in conjunction with the UDP helper feature to define the type of application traffic (port number or socket) that is being forwarded to the server. By default this feature is enabled.

Additional configuration is required to configure the UDP helper feature. For more details on configuring UDP helper, see “Configuring UDP Helper (optional)” on page 9-32.

USING THE CLI

To enable the broadcast forwarding of snmp traps, enter the following command:

```
HP9300(config)# ip forward-protocol udp snmp-trap
```

Syntax: ip forward-protocol udp <UDP-application-name>|<UDP-application-num>

Possible values:

number	echo	snmp-trap
bootpc	mobile-ip	tacacs
bootps	netbios-dgm	talk
discard	netbios-ns	
dnsix	ntp	
tftp	snmp	

In addition, you can specify any UDP application by using the application's UDP port number.

NOTE: By default, when an IP helper address is configured on an interface, UDP broadcast forwarding is enabled for the following UDP packets: bootps, domain, tftp, time, netbios-dgm, netbios-ns, and tacacs.

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link from the main menu. The panel shown in Figure 9.3 will appear.
2. Enable the Broadcast Forward option.
3. Select the Apply button to assign the changes.

NOTE: To define the ports to be forwarded, select the UDP Helper link from the IP configuration sheet.

Changing Network Mask Displays to Prefix Format

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the displays to prefix format (example: /18) by entering the following command at the Privileged (Enable) or CONFIG level of the CLI:

```
HP9300(config)# ip show-subnet-length
```

Defining Static IP Routes

You can manually add static IP routes by entering a destination IP address and mask along with the IP address of the next hop router. You also can assign the default router as the destination by entering 0.0.0.0 0.0.0.0.

The routing switches support up to 16 static routes by default. You can increase this support to up to 64 routes if needed.

NOTE: In software release 05.0.00 and later, the software will replace a statically configured static default route with a learned default route if the learned route's administrative distance is lower than the statically configured default route's distance. However, the default administrative distance for static routes is changed to 1 in software release 05.2.00, so only directly-connected routes are preferred over static routes when the default administrative distances for the routes are used.

USING THE CLI

To enter static IP route 1 with a destination address of 192.0.0.0 255.0.0.0 and a next hop router IP address of 195.0.0.0 on interface 1/6, enter the following commands:

```
HP9300(config)# ip route 1 192.0.0.0 255.0.0.0 195.1.1.1
```

Syntax: ip route <routenum> <dest-ip-addr> <dest-mask> <next-hop-ip-addr> [<metric>] [distance <num>]

USING THE WEB MANAGEMENT INTERFACE

1. Select the Static Route option from the IP configuration sheet. The static route entry panel shown in Figure 9.4 will appear.

   NOTE: If static routes already exist on the routing switch, then the static route summary panel appears instead. In this case, select the Add Static Route link to reach the Static route entry panel.

   ---

2. Enter the IP address in the Network field.
3. Enter the IP mask.
4. Enter the address of the next hop router that provides access to that destination.
5. Enter a default metric for the route if a value other than the one configured at the interface level is desired. The default metric is 1.
6. Enter the administrative distance for the static route. Each type of route on the routing switch has a different default administrative distance. See "Changing Administrative Distances" on page 12-22.
7. Click the Add button to save the entry to the static route table.

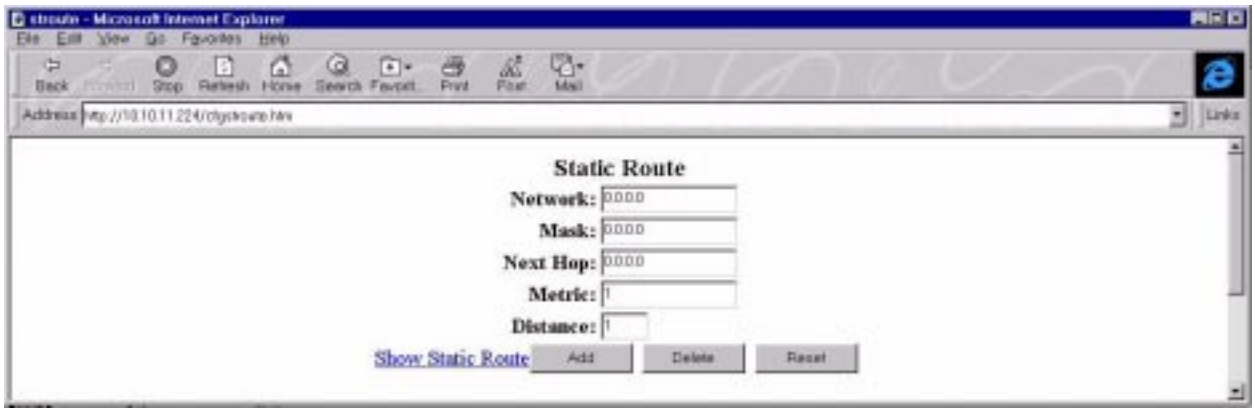

Figure 9.4 Defining an IP static route

Assigning Static ARP and RARP Entries (optional)

You can assign up to 16,000 static ARP and RARP entries.

USING THE CLI

To assign a static ARP entry on a chassis system, enter a command such as the following:

```
HP9300(config)# arp 1 192.53.4.2 1245.7654.2348 e 1/2
```

Syntax: arp <num> <ip-addr> <mac-addr> ethernet <portnum>

USING THE WEB MANAGEMENT INTERFACE

1. Select Static ARP from the IP configuration sheet. The panel shown in Figure 9.5 will appear.

NOTE: If any static ARP entries are defined on the routing switch, the static ARP summary panel appears first. In this case, select Add Static ARP.

2. Enter the IP address.
3. Enter the MAC address.
4. Select the port that the static ARP entry is to be assigned to from the pull down menu.
5. Click the Add button to save the entry to the static ARP table.

NOTE: You must be directly linked to an IP interface for which you are defining a static ARP.

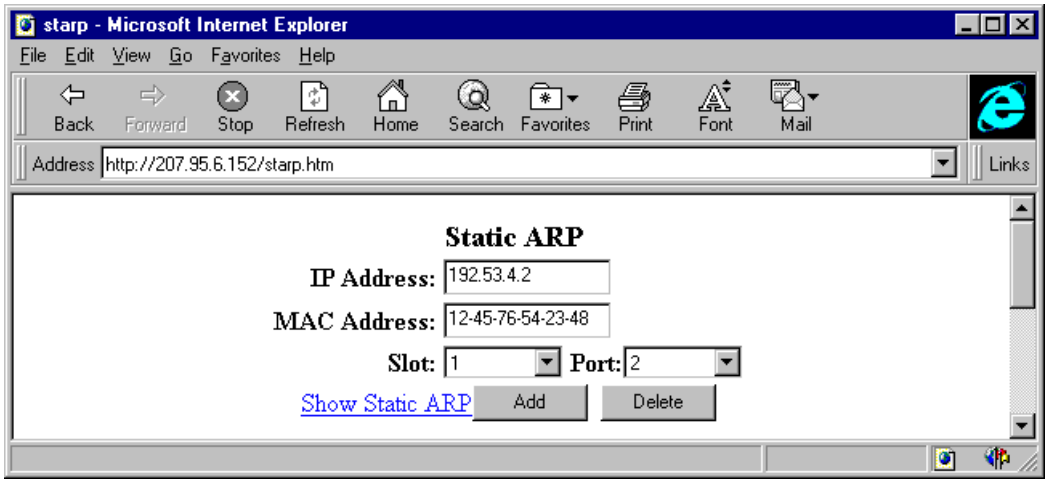

Figure 9.5 Static ARP entry panel

USING THE CLI

To assign a static IP RARP entry for static routes on a routing switch, enter the a command such as the following:

```
HP9300(config)# rarp 1 1245.7654.2348 192.53.4.2
```

syntax: rarp <number> <mac address>.<ip address> ethernet <port>

USING THE WEB MANAGEMENT INTERFACE

1. Select Static RARP from the IP configuration sheet. The panel shown in Figure 9.6 will appear.

NOTE: If any static RARP entries are defined on the routing switch, the static RARP summary panel appears first. In this case, select Add Static RARP.

2. Enter the MAC address.
3. Enter the IP address.
4. Click the Add button to save the entry to the static RARP table.

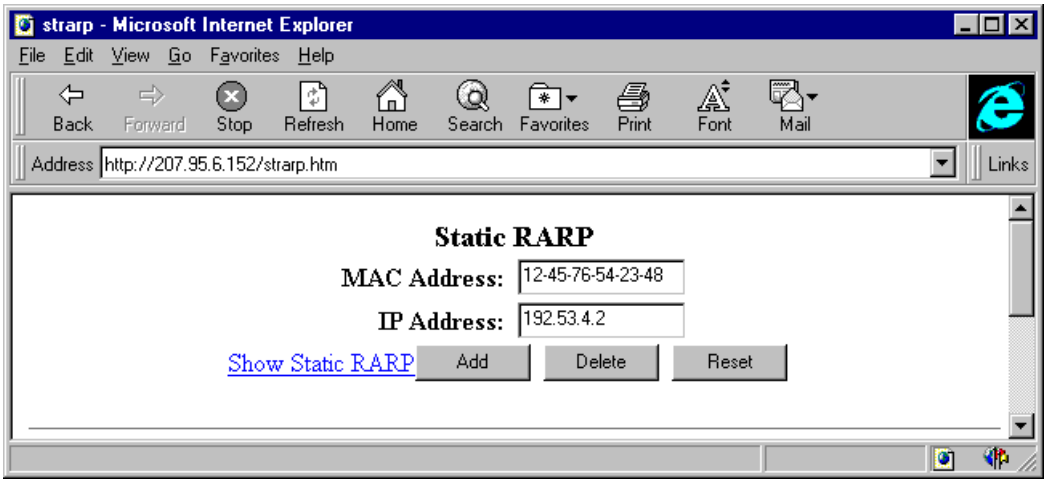

Figure 9.6 Static RARP entry panel

Assigning IP and IP/RIP Filters

You can define IP and IP/RIP filters on a global basis and assign filters on an interface basis. You also can define filters for redistributing routes among RIP and OSPF. This section describes how to perform the following filter tasks:

- Define IP access policies (permit and deny filters).
- Assign IP access policies to specific ports.
- Define IP/RIP filters.
- Assign IP/RIP filter groups to specific ports.
- Define IP/RIP neighbor filters.
- Define IP/RIP redistribution filters.

The following sections describe how to configure these access policies and filters. For more information, see “Policies and Filters” on page D-1.

Defining IP Access Policies

You can enhance network security by configuring IP access policies to explicitly permit or deny IP packets based on IP protocol, IP source and destination, IP protocol port, and even TCP or UDP application port.

NOTE: The routing switch permits all IP packets by default. However, once you configure an IP access policy, the routing switch denies all IP packets by default unless you explicitly permit them. Thus, if you want the routing switch to permit all IP packets except the ones you filter out, you must configure the last IP access policy to permit all IP packets. If a packet does not match other filters (and thus is not denied), the packet matches the last filter and is permitted.

You can filter on the following IP protocols:

- ICMP
- IGMP
- IGRP
- OSPF
- TCP
- UDP

In addition, if you filter on TCP or UDP, you also can specify a particular application port (such as "HTTP" or "80") or a logical expression consisting of an operator and port names or numbers. See the syntax descriptions below for details.

USING THE CLI

EXAMPLE 1: To configure an IP access policy that globally accepts all FTP traffic without regard to network orientation, use the wildcard value ‘any' in place of an IP address and enter the following command:

```
HP9300(config)# ip access-policy 1 permit any any tcp eq ftp
```

EXAMPLE 2: To configure an IP access policy that accepts only FTP traffic from a specific network, enter the following command:

```
HP9300(config)# ip access-policy 1 permit 192.38.5.54 255.255.255.0 195.38.5.53
255.255.255.0 tcp eq ftp
```

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

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

NOTE: For backward compatibility, the routing switch also supports 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 **deny|permit** parameter specifies the action the routing switch takes if a packet matches the policy.

- If you specify deny, the routing switch drops the packet.
- If you specify permit, the routing switch forwards the packet.

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 IP protocol 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 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.

The **log** parameter applies only to deny policies. This parameter generates a Syslog entry for packets that are denied by the policy. See “show logging” on page B-242.

Figure 9.7 and Figure 9.8 show the CLI syntax for configuring an IP access policy.

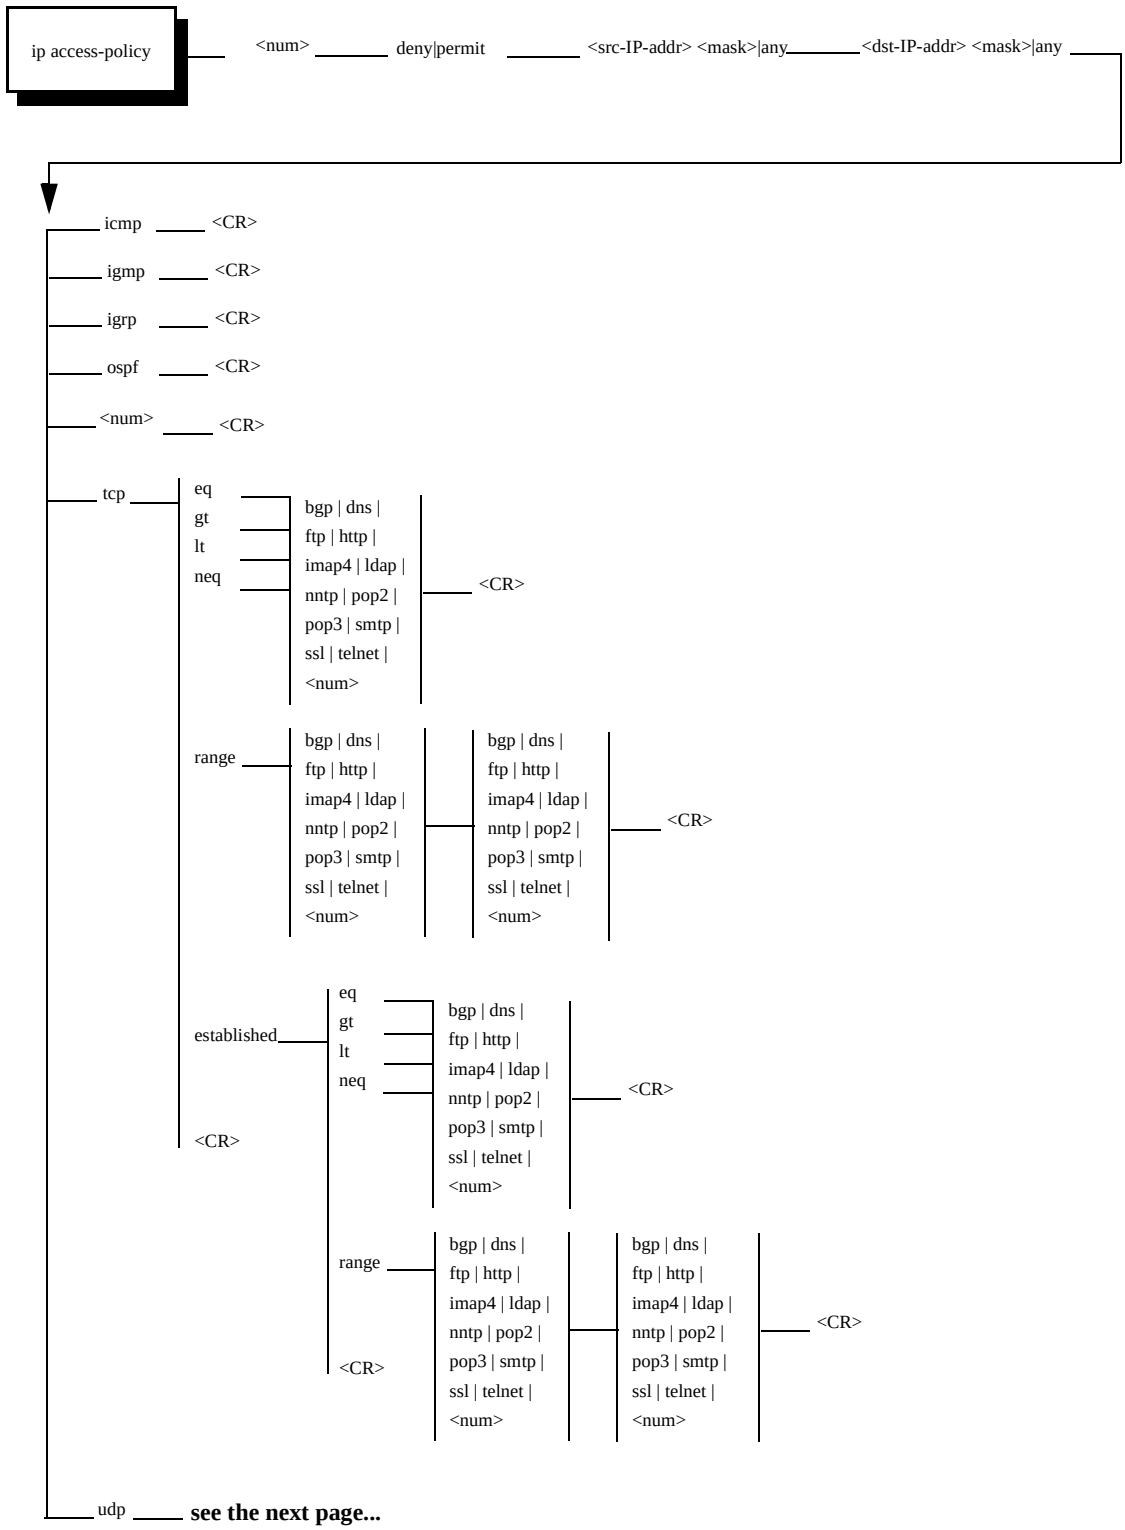

Figure 9.7 IP access policy syntax for an IP access policy (1 of 2)

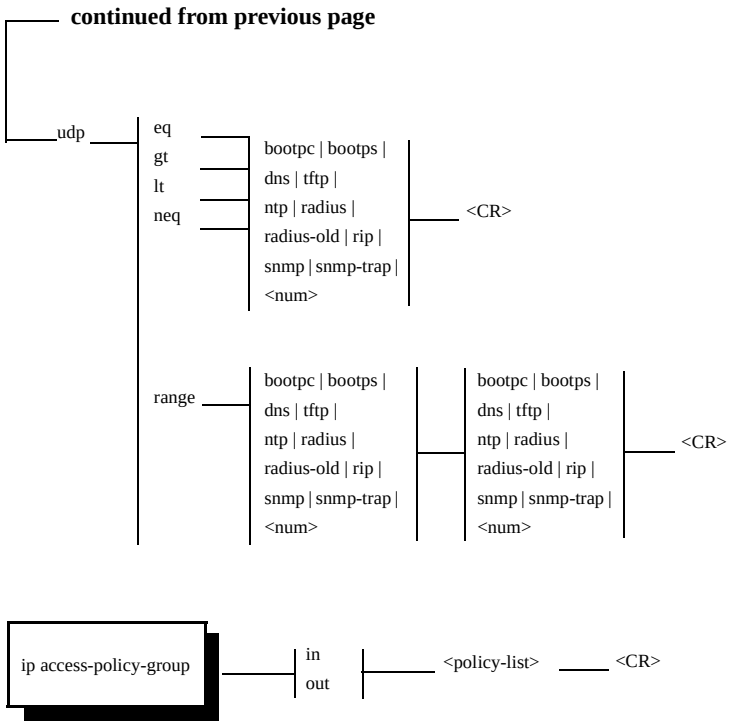

Figure 9.8 IP access policy syntax for an IP access policy (2 of 2)

USING THE WEB MANAGEMENT INTERFACE

To add an IP access policy:

1. Select the IP Access Policy link from the IP configuration sheet. The panel shown in Figure 9.9 will appear.

   NOTE: If IP filters are already defined on the routing switch, then the IP filter summary panel is displayed. In this case, select the Add IP Filter link.

2. Enter an ID for the access policy.
3. Select Permit, Deny, or QoS.

   NOTE: If you select QoS, you are configuring a Layer 4 Quality of Service (Qos) policy. See “Applying Layer 4 QoS Priority on the HP 9304M, 9308M, or 6308M-SX Routing Switch” on page 8-63.

4. Enter the source address and mask for the policy.

   NOTE: You can specify the wildcard value "any" in the source and destination IP address and mask fields to allow all traffic. Entering 0.0.0.0 represents "any". Likewise, to allow all protocols to be accepted by a filter, you can enter a single zero (0) in the protocol field.

5. Enter the destination address and mask for the policy.

6. If you want to filter on a specific IP protocol, select the protocol from the Protocol field's pulldown menu. For example, to filter on TCP packets, select TCP. You can enter the protocol number or select one of the following:
   - ICMP
   - IGMP
   - IGRP
   - OSPF
   - TCP
   - UDP
7. If you selected TCP or UDP, you can select a comparison operator. Select the operator from the Operator field's pulldown menu. 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. If you selected a comparison operator, enter the port number in the TCP/UDP port field. For example, if you selected TCP and Equal and you want to filter on HTTP traffic, enter the value 80 (the well-known port number for HTTP).

   NOTE: You must enter the port's number instead of the well-known name.

9. Click the Add button to assign the IP access policy.

Modify or Delete an IP Access Policy

1. Select IP Access Policy from the IP configuration sheet.
2. Click either the Modify or Delete button to the right of the IP policy you want to change or delete. If you click Modify, an entry panel for that interface appears. Make the desired changes and click Add to save the changes.

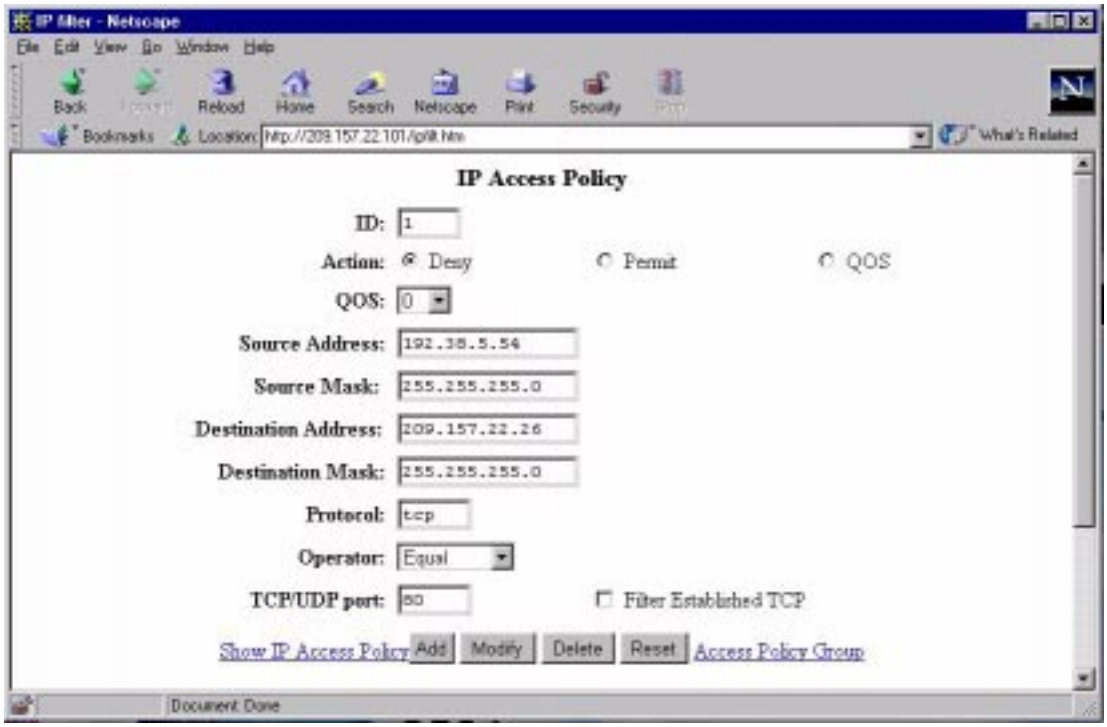

Figure 9.9 IP Access Policy entry panel

Applying IP Access Policies to Ports

Once you define an IP access policy, you can apply it to the inbound or outbound traffic on a port.

USING THE CLI

To assign IP access policies 2, 3, and 5 to port 1 on module 2 of a chassis, enter the following commands:

```
HP9300(config)# interface e 2/1
HP9300(config-if-2/1)# ip access-policy-group in 2 3 5
```

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

You also can specify policy ranges. For example, to apply policies 1 – 3, policy 9, and policies 11 – 25 to port 2/4's outbound policy group, enter the following command:

```
HP9300(config)# int ethernet 2/4
HP9300(config-if-2/4)# ip access-policy-group out 1 to 3 9 11 to 25
```

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

USING THE WEB MANAGEMENT INTERFACE

To assign IP filters 1, 2, and 5 to port 1 on module 2 of a chassis:

1. Select the Access Policy Group link from the IP filter configuration panel, shown in Figure 9.10.

NOTE: If at least one IP access policy group is already defined on the routing switch, then the IP access policy group summary panel is displayed first. In this case, select the Add IP Access Policy Group link.

2. Select the port or slot/port to which you are assigning the access policies.
3. Select either or both the In and Out options.
   - Selecting In applies the access policies to all incoming traffic on the port.
   - Selecting Out applies the access policies to all outgoing traffic on the port.
   - Selecting both options applies the access policies to both incoming and outgoing traffic.
4. Enter the access policy IDs in the Filter ID List field. To enter a range, enter the first policy number in the range, a space, a dash, another space, and then the second policy number. For example, enter "1 – 4" to specify the range 1 – 4.

NOTE: When specifying a range, you must use spaces on either side of the dash.

Figure 9.10 Assigning IP filters

Defining IP/RIP Route Filters

To define an IP/RIP filter, RIP must be enabled on the routing switch. A filter controls the routes that are stored in the IP routing table for inbound routes. For outbound routes, the filter defines the routes that are advertised through a given interface. You can define up to 64 route filters for a routing switch.

NOTE: A route is defined by its IP address and IP mask.

USING THE CLI

To enable RIP on the routing switch and then define IP/RIP filters, enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# filter 1 permit 192.53.4.1 255.255.255.0
HP9300(config-rip-router)# filter 2 permit 192.53.5.1 255.255.255.0
HP9300(config-rip-router)# filter 3 permit 192.53.6.1 255.255.255.0
HP9300(config-rip-router)# filter 4 deny 192.53.7.1 255.255.255.0
```

syntax: filter <filter-num> <permit|deny> <source-ip-address|any> <source-mask|any>

NOTE: Instead of specifying a specific route, you can specify all routes versus a specific sub-net by using the value **any**.

USING THE WEB MANAGEMENT INTERFACE

To define a RIP route filter:

1. Select RIP Route Filter from the RIP configuration sheet to display the entry panel shown in Figure 9.11.

   NOTE: If RIP route filters are already configured, a summary panel is displayed instead. In this case, select the Add RIP Route Filter link to reach the entry panel.

2. Enter the filter ID.
3. Select either Permit or Deny as the action.
4. Enter an IP address and mask or the wildcard value, 0.0.0.0, to allow all routes.
5. Click the Add button to save the filter.

To modify or delete a RIP route filter:

1. Select RIP Route Filter from the RIP configuration sheet to display a summary panel of all defined RIP route filters.
2. Click the Modify or Delete button next to the filter you want to change or delete. If the click Modify, enter the changes to either or both of the Action or IP Address fields and then click the Modify button to apply the changes. If you click Delete, the filter is removed immediately.

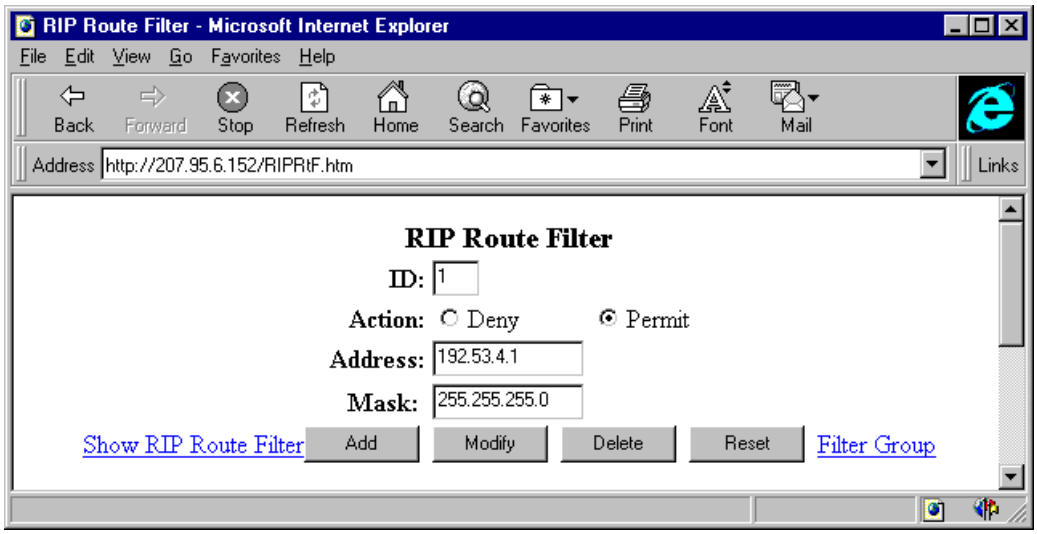

Figure 9.11 IP/RIP filter entry panel

Applying IP/RIP Route Filters to Ports

Once you define RIP route filters, you can assign them to individual ports. You also can specify whether the filters apply to advertisements sent by the routing switch or to updates received by the routing switch. Out filters apply to advertisements sent by the routing switch. In filters apply to updates received by the routing switch.

USING THE CLI

To assign route filters 2, 3, and 4 to all incoming routes on interface 2 of module 1, enter the following commands:

```
HP9300(config)# interface e 1/2
HP9300(config-if-1/2)# ip rip filter-group in 2 3 4
```

syntax: ip rip filter-group in|out <filter-list>

NOTE: If you specify **out** in the above example, filters 2, 3, and 4 are applied to all RIP routes being advertised. You also can assign filter groups on a global basis.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Filter Group link from the RIP filter configuration panel. The panel shown in Figure 9.12 will appear.
2. Select the port or slot/port to which the filter(s) will be assigned.
3. Select either or both of the In Filter and Out Filter options.
   - Selecting the In Filter option applies the filters to incoming traffic only.
   - Selecting the Out Filter option applies the filters to outgoing traffic only.
   - Selecting both options applies the filters to both incoming and outgoing traffic.
4. Enter the filters to be applied to the interface in the Filter ID List field.
5. Click the Add button to assign the changes.

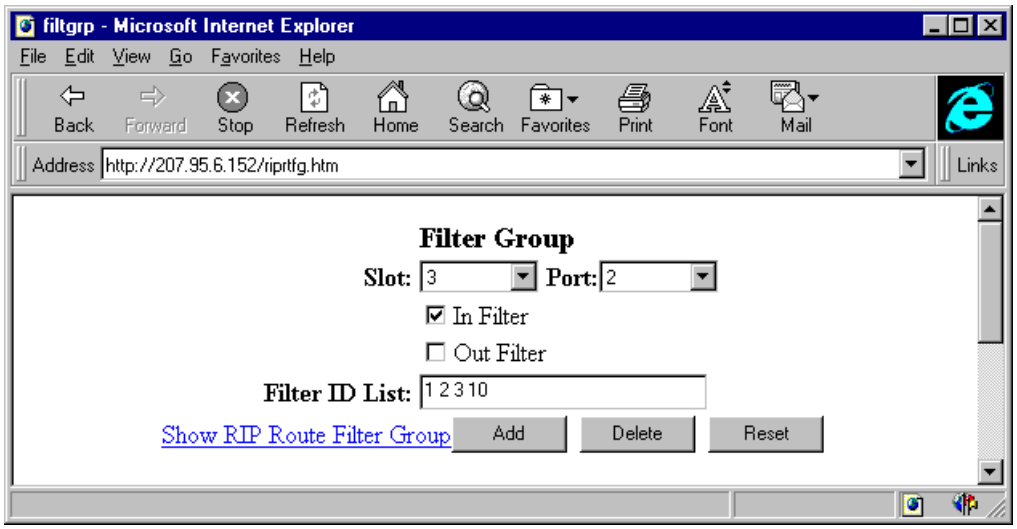

Figure 9.12 Assigning IP/RIP filters to an interface

Defining IP/RIP Neighbor Filters

By default, the routing switch learns RIP routes from all its RIP neighbors. Neighbor filters allow you to specify the neighbor routers from which the routing switch can receive RIP routes. You can define up to 64 neighbor filters.

Neighbor filters apply globally to all ports.

USING THE CLI

To configure a routing switch so that no RIP routes are learned from neighbor routers, enter the following command:

```
HP9300(config-rip-router)# neighbor 1 deny any
```

syntax: neighbor <filter-num> permit|deny <source-IP-address>|any

USING THE WEB MANAGEMENT INTERFACE

To define a RIP neighbor filter:

1. Select RIP Neighbor Filter from the RIP configuration sheet. The panel shown in Figure 9.13 will appear.
2. Enter the filter ID.
3. Select either the Permit or Deny action.
4. Enter the source IP address that will be filtered or 0.0.0.0 to filter on all neighboring routers.
5. Click the Add button to assign the filter.

To modify or delete a RIP neighbor filter:

1. Select RIP Neighbor Filter from the RIP configuration sheet. A summary panel of all defined RIP neighbor filters will appear.
2. Click the Modify or Delete button next to the filter that is to be changed or deleted. If you click Modify, enter the changes to the Action or IP Address fields and then click the Modify button apply the changes. If you click Delete, the filter is removed immediately.

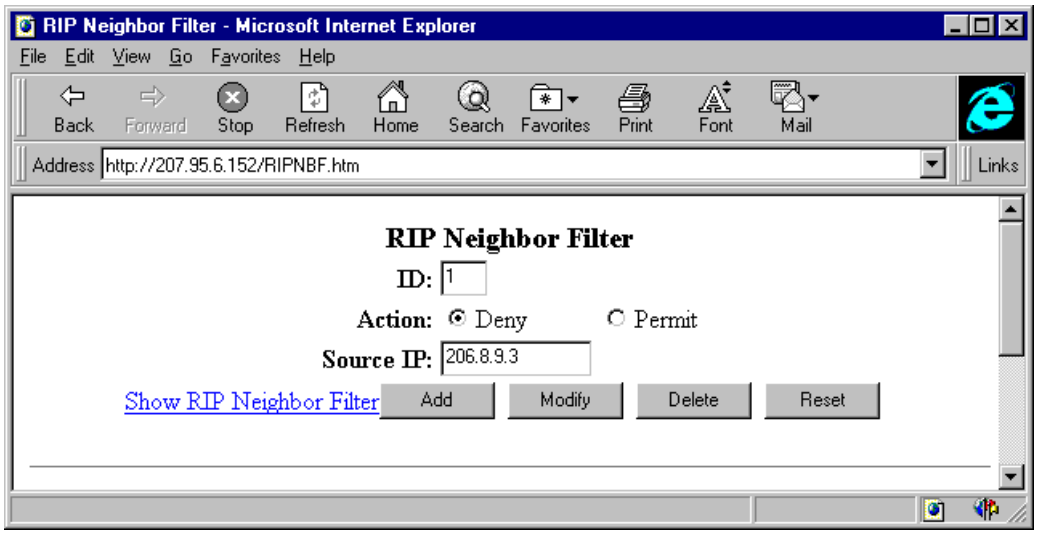

Figure 9.13 RIP neighbor filter entry panel

Defining Redistribution Filters

IP/RIP redistribution filters control redistribution of routes from other protocols into RIP. A routing switch running RIP can redistribute static routes, OSPF routes, and BGP4 routes (if BGP4 is supported on the device) into RIP.

Optionally, you can specify a metric that the route must match or you can set the metric on redistributed routes. By setting the metric, you can cause the routing switch to prefer IP/RIP routes or redistributed routes to the specified network.

USING THE CLI

EXAMPLE 1: To deny redistribution on all incoming routes received from the 207.92.0.0 network (by interface), enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# deny redis 2 all 207.92.0.0 255.255.0.0
```

EXAMPLE 2: To deny redistribution on OSPF routes only, enter the following command:

```
HP9300(config-rip-router)# deny redis 3 ospf 207.92.0.0 255.255.0.0
```

EXAMPLE 3: To deny redistribution by metric, enter the following command:

```
HP9300(config-rip-router)# deny redis 3 ospf 207.92.0.0 255.255.0.0 match-metric 10
```

Syntax: permit|deny redistribute <filter-num> all|bgp|ospf|static <ip-addr> <mask> [match-metric<value>|set-metric <value>]

The **all** parameter applies redistribution to all route types.

The **bgp** parameter applies redistribution to BGP4 routes only.

The **ospf** parameter applies redistribution to OSPF routes only.

The **static** parameter applies redistribution to the static route only.

The **<ipaddr> <mask>** parameters apply redistribution to the specified network and sub-net address.

The **match-metric <value>** parameter applies redistribution to those routes with a specific metric value; possible values are from 1 – 15.

The **set-metric <value>** parameter sets the RIP metric value that will be applied to those routes imported into RIP.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Redistribution Filter link from the RIP configuration sheet. The panel shown in Figure 9.14 will appear.
2. Enter an IP address and mask to filter on a specific network. You can use zeros (0.0.0.0) instead of a specific interface to allow all IP addresses or mask ranges.
3. Enter the filter ID.
4. Select either Permit or Deny as the action.
5. Select the types of routes you want to filter on—All, Static, OSPF, or BGP4.
6. Enable the Match Metric parameter to limit the import of routes to only those that match the metric specified in the Match Metric field.
7. Enable the Set Metric parameter to define and assign a specific metric to an imported route. If enabled, the specified value overrides the default metric defined on the RIP configuration sheet.
8. Click the Add button to assign the redistribution filter.

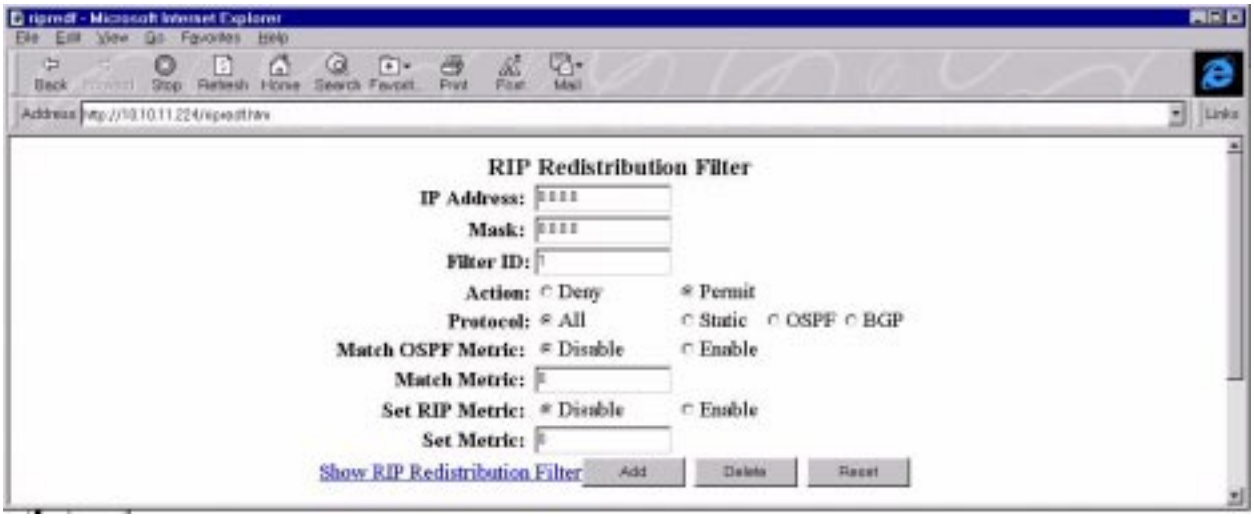

Figure 9.14 IP/RIP redistribution filter entry panel

Modify IP and IP/RIP Interface Parameters (optional)

IP and IP/RIP come with default settings for their interface parameters. You do not need to modify any of these parameters unless your network configuration requires a parameter change. You can configure the following interface parameters:

- IP interface parameters:
  - Encapsulation format
  - Maximum transmission unit (MTU)
  - Metric
  - IP address used for stamping BootP/DHCP requests
- RIP interface parameters:
  - RIP routing state on individual routing switch ports
  - RIP Version—version 1, version 2, or version 2 with version 1 compatibility
  - Poison reverse state
  - Filter groups

Modifying IP Interface Parameters

Use the procedures in this section to modify the following parameters:

- Encapsulation format
- Maximum transmission unit (MTU)
- Metric
- IP address used for stamping BootP/DHCP requests

Modifying Encapsulation Format

The encapsulation format parameter allows you to select the encapsulation format to be used on a port for MAC address encapsulation. This can vary by port. The options are Ethernet II or SNAP. The default format is Ethernet II.

USING THE CLI

To change the encapsulation type on interface 1/5 to Ethernet SNAP, enter the following commands:

```
HP9300(config)# int e 1/5
HP9300(config-if-1/5)#ip encapsulation ethernet_snap
```

syntax: ip encapsulation <ethernet_snap | ethernet_ii>

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP Interface link from the IP configuration sheet. The panel shown in Figure 9.15 will appear.
2. Select the port (and slot, if applicable).
3. Select the encapsulation type from the pulldown menu.
4. Select the Apply button to assign the changes.

Modifying the Size of the Maximum Transmission Unit (MTU)

The MTU field defines the maximum packet size to be accepted on a given port. The possible size for Ethernet II packets is 572 – 1500 bytes. Ethernet SNAP packets can be from 572 – 1492 bytes. The default value for Ethernet II packets is 1500. The default for SNAP packets is 1492.

USING THE CLI

To change the MTU for interface 1/5 to 1000, enter the following commands:

```
HP9300(config)# int e 1/5
HP9300(config-if-1/5)#ip mtu 1000
```

syntax: ip mtu <572-1500> (Ethernet SNAP); ip mtu <572-1492> (Ethernet II)

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP Interface link from the IP configuration sheet to display the panel shown in Figure 9.15.

NOTE: If at least one IP interface is defined on the routing switch, then a summary panel will appear first. In this case, select the Configure IP Interface link to reach the IP interface panel shown in Figure 9.15.

2. Enter an MTU value from 572 – 1500 if the interface is operating with Ethernet SNAP encapsulation. If the interface is operating with Ethernet II, enter a value from 572 – 1492.
3. Select the Apply button to assign the changes.

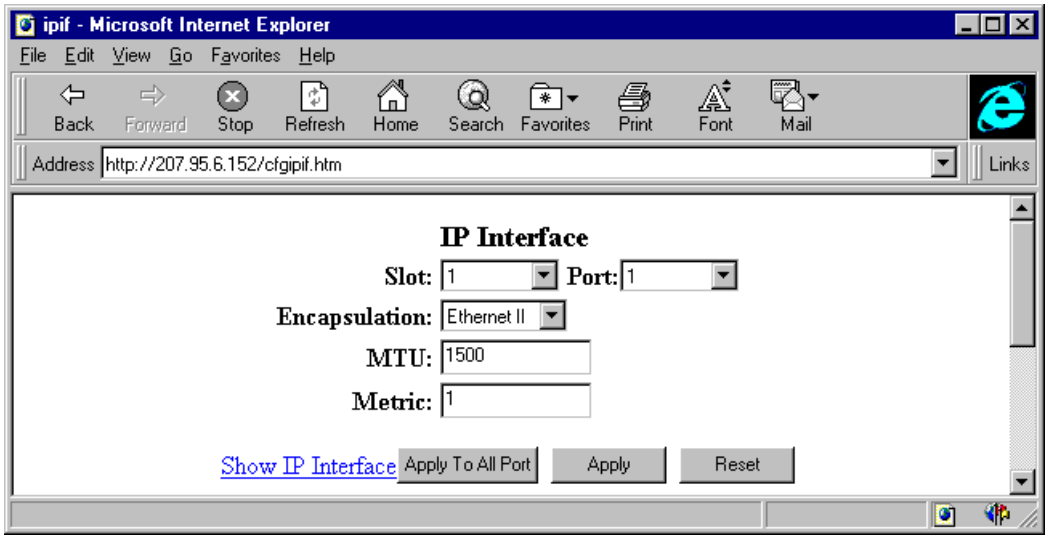

Figure 9.15 IP interface configuration panel

Modifying the Metric

Metric defines the cost that will be applied to all IP routes on an interface. A metric cost from 1 – 16 can be assigned. The default metric cost is 1.

USING THE CLI

To assign a route cost (metric) of 15 to interface 1/6:

```
HP9300(config)# int e 1/6
HP9300(config-if-1/6)# ip metric 15
```

syntax: ip metric <1-16>

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP Interface link from the main menu. The panel shown in Figure 9.15 will appear.
2. Enter a value from 1 – 16 for the metric.

NOTE: IP/RIP considers interfaces with a metric of 16 to be unreachable. Use this metric only if you do not want the interface to be used.

3. Select the Apply button to assign the changes.

Modifying the IP Address Used for Stamping BootP/DHCP Requests

The routing switch assists BootP/DHCP requests by stamping such requests with the IP address of the gateway that leads to the BootP/DHCP server. By default, the lowest numbered IP address on an interface is used as the address for stamping the requests. To change the address, use one of the following methods.

USING THE CLI

To change the IP address used for stamping BootP/DHCP requests on interface 1/1, enter the following commands:

```
HP9300(config)# int e 1/1
HP9300(config-if-1/1)# ip bootp-gateway 109.157.22.26
```

Syntax: ip bootp-gateway <ip-addr>

USING THE WEB MANAGEMENT INTERFACE

You cannot change the IP address used for stamping BootP/DHCP requests using the Web management interface.

IP/RIP Interface Parameters

Use the procedures in this section to modify the following parameters:

- RIP routing on individual routing switch ports
  - RIP Version—version 1, version 2, or version 2 with version 1 compatibility
  - Poison reverse
- Filter groups

NOTE: You also can define IP access policies, assign static IP routes and define static ARP and RARP entries for interfaces. For more details on these features, see the specific sections on their configuration within this chapter.

Enabling IP/RIP Routing on Interfaces and Modify Parameters (optional)

As autonomous systems, the 9304M, 9308M, and 6308M-SX routing switches can support multiple protocols on the same device. You can enable RIP on individual ports by selecting that port from the pulldown menu, assigning a version type, then either enabling or disabling the parameter poison reverse.

USING THE CLI

To enable RIP on an interface, define the type of RIP route and enable poison reverse for interface 1/1, enter the following commands:

```
HP9300(config)# int e1/1
HP9300(config-if-1/1)# ip rip v1-only
HP9300(config-if-1/1)# ip rip poison-reverse
HP9300(config-if-1/1)# end
```

```
HP9300# write memory
HP9300# reload
```

syntax: ip rip <v1-only|v1-compatible-v2|v2-only>; ***syntax:*** ip rip poison-reverse

USING THE WEB MANAGEMENT INTERFACE

To enable RIP routing on individual interfaces:

1. Select RIP Interface from the RIP configuration sheet. The panel shown in Figure 9.16 will appear.

   NOTE: If RIP is already defined on some interfaces, an interface configuration summary panel will appear. In this case, select Configure RIP Interface to add an interface.

2. Select the port or slot/port to be configured from the pulldown menu.
3. Assign the RIP type version from the pull down menu. Options are version 1, version 2, v1 compatible v2 or disabled. The default state is version 2.
4. Enable poison reverse, a loop prevention feature, if desired.
5. Select Apply to assign the changes.

NOTE: To assign the configured interface parameters to all other RIP interfaces on the routing switch, select the Apply All Port button.

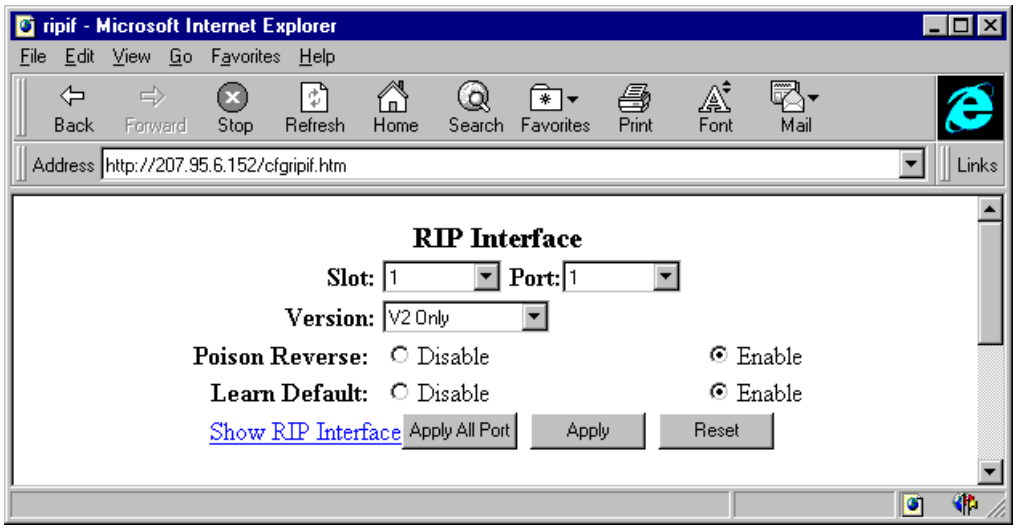

Figure 9.16 RIP interface display and entry panel

Modify Global IP/RIP Parameters

The IP/RIP protocol has some global parameters, which have default settings. You do not need to modify these parameters unless your network configuration requires a parameter change.

The following RIP parameters are modified at the RIP router level when using the CLI and at the RIP configuration sheet when using the Web management interface.

- Update time
- Enable or disable of redistribution
- Global default metric used for redistribution
- Enable IP/RIP Default Route Learning and Advertising

Modifying Update Time Value

The update time sets the time interval between the transmission of regular RIP response packets. Possible values are 1 – 1000 seconds. The default value is 30 seconds.

USING THE CLI

To modify the interval at which RIP response packets are transmitted to 120 seconds, enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# update 120
```

syntax: update-time <1-1000>

USING THE WEB MANAGEMENT INTERFACE

1. Select the RIP link from the main menu. The panel shown in Figure 9.17 will appear.
2. Enter a value from 1 – 1000 in the Update Time field.
3. Select the Apply button to assign the changes.

Figure 9.17 RIP configuration sheet

Enabling or Disabling Redistribution

When RIP is enabled, it imports external routes (OSPF routes, static routes, or BGP4 routes) into the RIP domain. Redistribution is disabled by default.

USING THE CLI

To enable redistribution for RIP, the user would enter the following:

```
HP9300(config)# router rip
HP9300(config-rip-router)# redistribution
```

syntax: redistribution

USING THE WEB MANAGEMENT INTERFACE

1. Select the RIP link from the main menu to display the panel shown in Figure 9.17.
2. Enable redistribution.
3. Select the Apply button to assign the changes.

Modifying the Redistribution Global Default Metric

The RIP redistribution metric allows you to define the global default metric (cost) assigned to all external routes imported into RIP for redistribution. Possible values are 1 – 15. The default value is 1.

USING THE CLI

To assign a global metric of 10 as the default cost, enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# default 10
```

syntax: default-metric <1-15>

USING THE WEB MANAGEMENT INTERFACE

1. Select the RIP link from the main menu. The panel shown in Figure 9.17 will appear.
2. Enter a value from 1 – 15 in the Redistribution Default Metric field.
3. Select the Apply button to assign the changes.

Modifying the Default Administrative Distance

The HP9304M, 9308M, and 6308M-SX routing switches can learn about networks from various protocols, including Border Gateway Protocol version 4 (BGP4), IP/RIP, and OSPF. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. The default administrative distance for IP/RIP routes is 120. See “Changing Administrative Distances” on page 12-22 for a list of the default distances for all route sources.

To select one route over another based on the source of the route information, the routing switch can use the administrative distances assigned to the sources. You can bias the routing switch’s decision by changing the default administrative distance for IP/RIP routes.

USING THE CLI

To change the administrative distance for IP/RIP routes to 140, enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# distance 140
```

syntax: distance <num>

USING THE WEB MANAGEMENT INTERFACE

1. Select the RIP link from the main menu. The panel shown in Figure 9.17 will appear.
2. Edit the value in the Distance field.
3. Select the Apply button to assign the changes.

Enabling IP/RIP Default Route Learning and Advertising

You can enable learning and advertising of IP/RIP routes on a global or interface basis.

USING THE CLI

To enable learning of default IP/RIP routes on a global basis, enter the following commands:

```
HP9300(config)# router rip
HP9300(config-rip-router)# learn-default
```

To enable learning of default IP/RIP routes on an interface basis, enter the following commands:

```
HP9300(config)# int e1
HP9300(config-if-1)# ip rip learn-default
```

syntax: learn-default

USING THE WEB MANAGEMENT INTERFACE

To enable learning of default IP/RIP routes:

1. Select the RIP interface link from the RIP configuration sheet. A summary panel of all RIP interfaces will appear.

NOTE: If RIP is already defined on some interfaces, an interface configuration summary panel will appear. In this case, select Configure RIP Interface. Select the Modify button next to the interface upon which learning of default routes is to be enabled. The RIP interface entry panel will appear.

2. Enable learn default.
3. Select the Apply button to assign the changes.

NOTE: To globally enable learning of default routes across all interfaces, select the Apply To All Ports button instead of the Apply button.

Configuring UDP Helper (optional)

The 9304M, 9308M, and 6308M-SX routing switches support relay of UDP/DHCP packets to their destinations for a specific application such as bootps, domain, tftp, and so on for cases where the destination server is not on the local LAN segment.

The following port sockets names are supported for the UDP helper feature:

number	echo	snmp
bootpc	mobile-ip	snmp-trap
bootps	netbios-dgm	tacacs
discard	netbios-ns	talk
dnsix	ntp	tftp

NOTE: You also can specify any UDP application by its number.

NOTE: By default, when an IP helper address is configured on an interface, UDP broadcast forwarding is enabled for the following UDP packets: bootps, domain, tftp, time, netbios-dgm, netbios-ns, and tacacs.

USING THE CLI

To configure the UDP/DHCP helper feature on interface 2 on chassis module 1, enter the following commands:

```
HP9300(config)# interface e 1/2
HP9300(config-if-1/2)# ip helper-address 1 207.95.7.6
```

syntax: ip helper-address <1-4> <ip address>

USING THE WEB MANAGEMENT INTERFACE

To configure the UDP/DHCP helper feature on an interface:

1. Select the UDP helper option on the IP configuration sheet to display the panel shown in Figure 9.18.
2. Select the port or slot/port to which the UDP helper packets will be forwarded from the pulldown menu(s).
3. Enter the IP address of the remote server for which the routing switch will be relaying the packets.
4. Select the Add button to apply the changes. You are now ready to assign applications to be forwarded, highlighted in the next section.

To select an application to be forwarded to the server by the routing switch:

1. Select System Broadcast Forward from the UDP helper entry panel. The panel shown in Figure 9.19 will appear.
2. From the pulldown menu, select the application(s) to be forwarded to the server. The chosen broadcast forwards are displayed in the display panel under the Selected Forward Ports heading. By default, the following applications are already be selected: bootps, domain, tftp, time, netbios-dgm, netbios-ns, and tacacs.
3. Select the Add button to apply the changes.

You can define you own protocol with the User Define panel. To do so:

1. Select the User Broadcast Forward link from the UDP helper configuration panel. The panel shown in Figure 9.20 will appear.
2. Enter a value from 1 – 65535.
3. Select the Add button to assign the change.

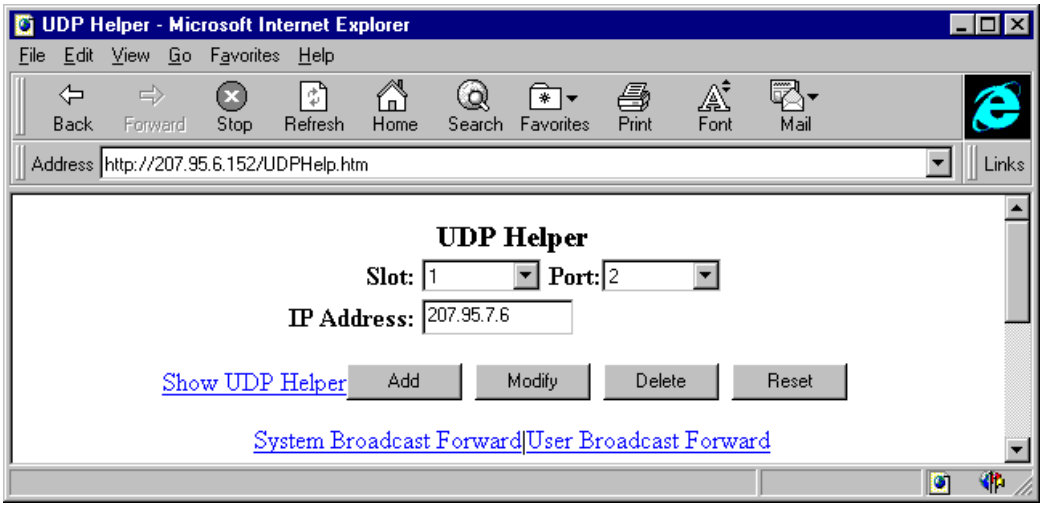

Figure 9.18 UDP helper configuration panel

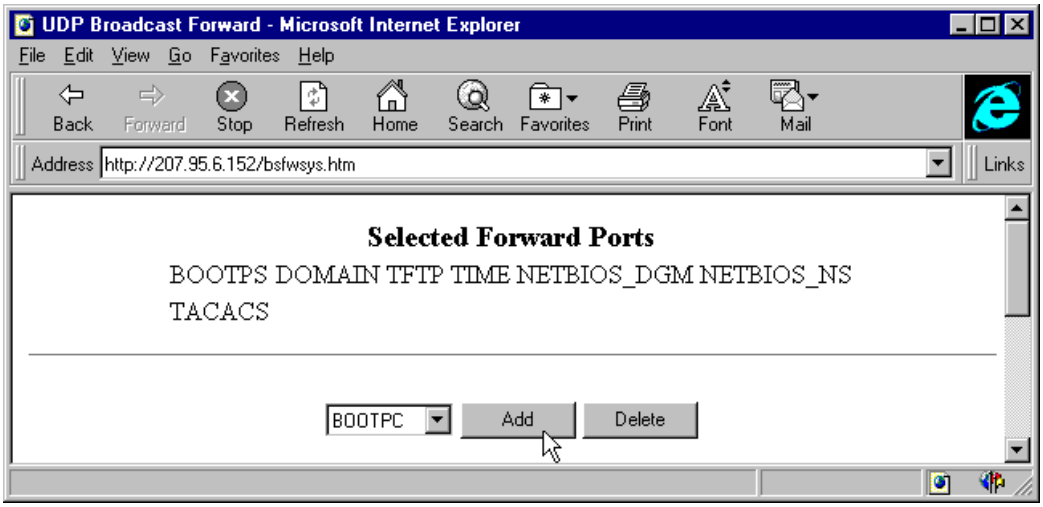

Figure 9.19 System broadcast forward entry panel

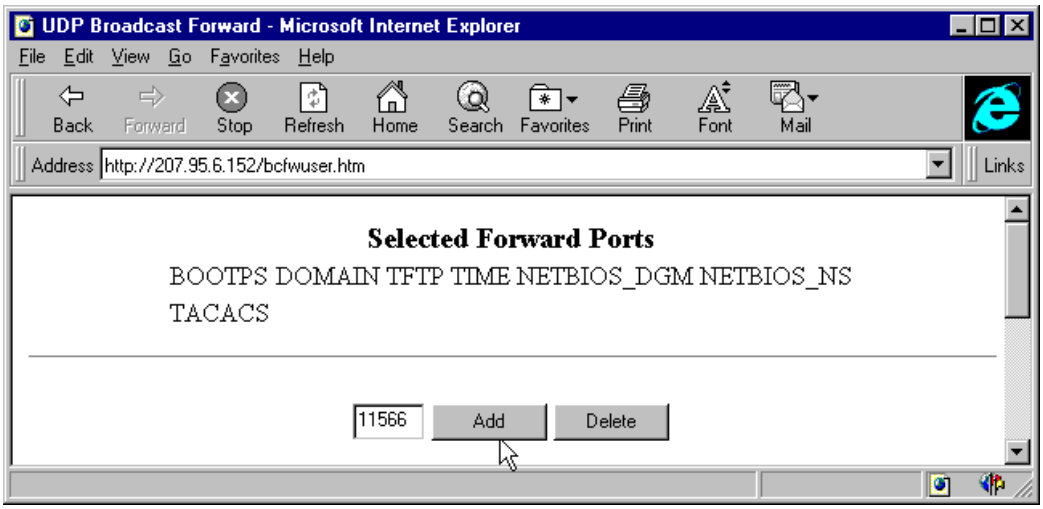

Figure 9.20 User-defined broadcast forward entry panel

Displaying IP and RIP Information

You can use CLI commands and Web management options to display the following IP information:

- Global IP parameter settings and IP access policies – see “Displaying Global IP Configuration Information” on page 9-35.
- IP interfaces – see “Displaying IP Interface Information” on page 9-40.
- ARP cache – see “Displaying the ARP Cache” on page 9-43.
- IP host cache – see “Displaying the IP Host Cache” on page 9-45.
- IP forwarding cache – see “Displaying the IP Forwarding Cache” on page 9-48.
- IP route table – see “Displaying the IP Route Table” on page 9-49.
- IP traffic statistics – see “Displaying IP Traffic Statistics” on page 9-50.
- RIP filters – see “Displaying RIP Filters” on page 9-59.

The sections below describe how to display this information.

In addition to the information described below, you can display the following IP information. This information is described in other parts of this guide.

- DVMRP information – “Commands – All Levels” on page B-63
- PIM information – “Commands – All Levels” on page B-63
- OSPF information – see “Displaying OSPF Information” on page 10-29.
- BGP4 information – see “Displaying BGP4 Information” on page 12-43.
- VRRP information – see “Displaying VRRP Configuration Information and Statistics” on page 13-16.
- SRP information – “Commands – All Levels” on page B-63

Displaying Global IP Configuration Information

To display global IP configuration information for the routing switch, use one of the following methods.

USING THE CLI

To display IP configuration information, enter the following command at any CLI level:

```
HP9300> show ip

Global Settings
  ttl: 64, arp-age: 10, bootp-relay-max-hops: 4
  router-id : 207.95.11.128
  enabled : UDP-Broadcast-Forwarding  IRDP  Proxy-ARP  RARP  OSPF
  disabled: BGP4 Load-Sharing  RIP  DVMRP  SRP  VRRP

Static Routes
  Index   IP Address        Subnet Mask       Next Hop Router   Metric Distance
  1       0.0.0.0           0.0.0.0           209.157.23.2      1      1

Policies
  Index   Action   Source           Destination      Protocol   Port  Operator
  1       deny     209.157.22.34    209.157.22.26    tcp        http  =
  64      permit   any              any
```

Syntax: show ip

NOTE: This command has additional options, which are explained in other sections in this guide, including the sections below this one.

This display shows the following information.

Table 2.4: CLI Display of Global IP Configuration Information

This Field...	Displays...
Global Settings	
ttl	The Time-To-Live (TTL) for IP packets. The TTL specifies the maximum number of router hops a packet can travel before reaching the routing switch. If the packet's TTL value is higher than the value specified in this field, the routing switch drops the packet. To change the maximum TTL, see "Modifying the tTTL Threshold" on page 9-7.
arp-age	The ARP aging period. This parameter specifies how many minutes an inactive ARP entry remains in the ARP cache before the routing switch ages out the entry. To change the ARP aging period, see "Modifying the ARP Aging Period" on page 9-7.
bootp-relay-max-hops	The maximum number of hops way a BootP server can be located from the routing switch and still be used by the routing switch's clients for network booting. To change this value, see "Modifying the Maximum Number of Hops to a BootP Relay Server" on page 9-6.

Table 2.4: CLI Display of Global IP Configuration Information (Continued)

This Field...	Displays...
router-id	The 32-bit number that uniquely identifies the routing switch. By default, the router ID is the numerically lowest IP interface configured on the routing switch. To change the router ID, see “Changing the Router ID” on page 9-7.
enabled	The IP-related protocols that are enabled on the routing switch.
disabled	The IP-related protocols that are disabled on the routing switch.
Static Routes	
Index	The row number of this entry in the IP route table.
IP Address	The IP address of the route’s destination.
Subnet Mask	The network mask for the IP address.
Next Hop Router	The IP address of the routing switch interface to which the routing switch sends packets for the route.
Metric	The cost of the route. Usually, the metric represents the number of hops to the destination.
Distance	The administrative distance of the route. The default administrative distance for static IP routes on the 9304M, 9308M, and 6308M-SX routing switches is 130. To list the default administrative distances for all types of routes or to change the administrative distance of a static route, see “Changing Administrative Distances” on page 12-22.
Policies	
Index	The policy number. This is the number you assigned the policy when you configured it.
Action	The action the routing switch takes if a packet matches the comparison values in the policy. The action can be one of the following: • deny – The routing switch drops that match this policy. • permit – The routing switch forward packets that match this policy.
Source	The source IP address the policy checks for in IP packets.
Destination	The destination IP address the policy checks for in IP packets.

Table 2.4: CLI Display of Global IP Configuration Information (Continued)

This Field...	Displays...
Protocol	The IP protocol the policy checks for in IP packets. The protocol can be one of the following: • ICMP • IGMP • IGRP • OSPF • TCP • UDP
Port	The Layer 4 TCP or UDP port the policy checks for in packets. The port can be displayed by its number or, for port types the routing switch recognizes, by the well-known name. For example, TCP port 80 can be displayed as HTTP. **Note**: This field applies only if the IP protocol is TCP or UDP.
Operator	The comparison operator for TCP or UDP port names or numbers. **Note**: This field applies only if the IP protocol is TCP or UDP.

USING THE WEB MANAGEMENT INTERFACE

To display global IP configuration information:

1. Select the Summary link from the list of links beneath the display panel to display the Summary panel. If you have frames enabled, the links also are listed in the frame on the left side of the window.
2. Select Configuration next to Type, if it is not already selected.
3. Select the checkbox next to IP to place a checkmark in the box, if it does not already contain one.
4. Optionally, deselect other display options by clicking to remove the checkmarks from the checkboxes next to the options. Deselecting other display options simplifies the display you will receive after you click the Apply button.
5. Click the Apply button to display the Configuration information. The IP information is listed in the IP General section. Scroll down until you see this section.

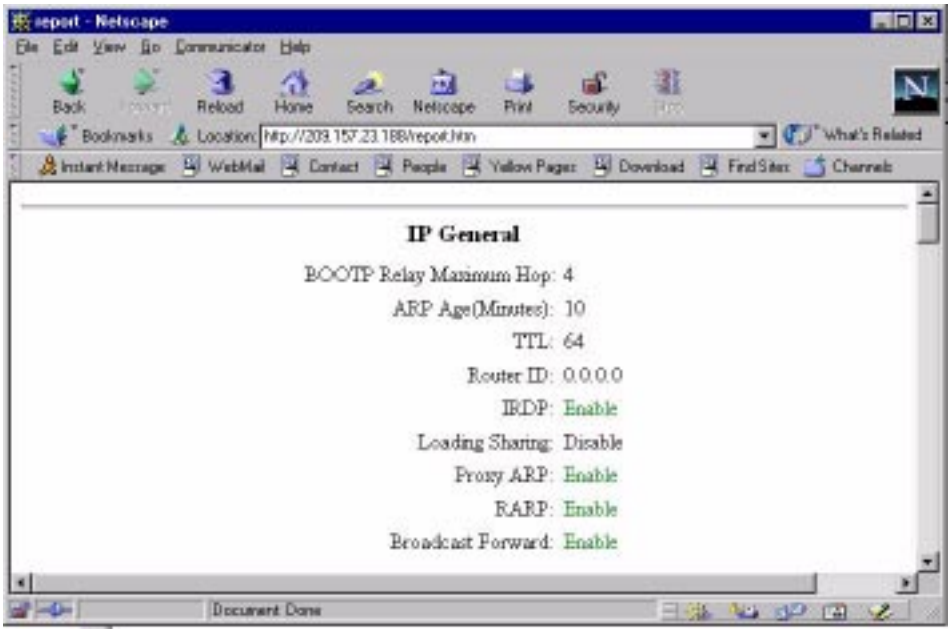

Figure 9.21 IP General information

The IP General display shows the following information.

Table 2.5: Web Display of Global IP Configuration Information

This Field...	Displays...
BOOTP Relay Maximum Hops	The maximum number of hops way a BootP server can be located from the routing switch and still be used by the routing switch's clients for network booting. To change this value, see "Modifying the Maximum Number of Hops to a BootP Relay Server" on page 9-6.
ARPAge	The ARP aging period. This parameter specifies how many minutes an inactive ARP entry remains in the ARP cache before the routing switch ages out the entry. To change the ARP aging period, see "Modifying the ARP Aging Period" on page 9-7.
TTL	The Time-To-Live (TTL) for IP packets. The TTL specifies the maximum number of router hops a packet can travel before reaching the routing switch. If the packet's TTL value is higher than the value specified in this field, the routing switch drops the packet. To change the maximum TTL, see "Modifying the tTTL Threshold" on page 9-7.
Router ID	The 32-bit number that uniquely identifies the routing switch. By default, the router ID is the numerically lowest IP interface configured on the routing switch. To change the router ID, see "Changing the Router ID" on page 9-7.

Table 2.5: Web Display of Global IP Configuration Information (Continued)

This Field...	Displays...
IRDP	The state of the IRDP protocol. The state can be one of the following: • Disable • Enable To change the state of the protocol, see “Enabling or Disabling IRDP” on page 9-8.
Load Sharing	The state of the IP load sharing feature. The state can be one of the following: • Disable • Enable To change the state of the feature, see “Enable or Disable Load Sharing” on page 9-8.
Proxy ARP	The state of the Proxy ARP feature. The state can be one of the following: • Disable • Enable To change the state of the feature, see “Disabling or Enabling Proxy ARP” on page 9-9.
RARP	The state of the RARP feature. The state can be one of the following: • Disable • Enable To change the state of the feature, see “Enable or Disable RARP” on page 9-9.
Broadcast Forward	The state of the IP broadcast forwarding feature. The state can be one of the following: • Disable • Enable To change the state of the feature, see “Enabling or Disabling Broadcast Forward” on page 9-9.

Displaying IP Interface Information

To display IP interface information, use one of the following methods.

USING THE CLI

To display IP interface information, enter the following command at any CLI level:

```
HP9300> show ip interface

Interface Ethernet 1/1
  port state: UP
  ip address: 99.1.1.1          subnet mask: 255.255.0.0
  encapsulation: ETHERNET, mtu: 1500, metric: 1
  RIP version: V2, Poison Reverse: on
  No Helper Addresses are configured.

Interface Ethernet 3/24
  port state: UP
  ip address: 100.100.1.6       subnet mask: 255.255.0.0
  encapsulation: ETHERNET, mtu: 1500, metric: 1
  No Helper Addresses are configured.

Interface Ve 1
  members: ethe 7/1 to 7/4
  active: none
  port state: DOWN
  ip address: 207.95.11.128     subnet mask: 255.255.255.0
  encapsulation: ETHERNET, mtu: 1500, metric: 1
  No Helper Addresses are configured.

Interface Ve 2
  members: ethe 7/5 to 7/8
  active: ethe 7/5 to 7/6
  port state: UP
  ip address: 209.157.23.188    subnet mask: 255.255.255.0
  encapsulation: ETHERNET, mtu: 1500, metric: 1
  No Helper Addresses are configured.
```

Syntax: show ip interface [ethernet <portnum>] | [ve <num>]

This display shows the following information.

Table 2.6: CLI Display of IP Interface Information

This Field...	Displays...
Interfaces configured on ports	
Interface Ethernet	The port number on which the interface is configured.
port state	The interface state, which can be one of the following: • DOWN • UP
ip address	The IP address of this interface.
subnet mask	The network mask for the IP address of this interface.
encapsulation	The frame type used to encapsulate packets on this interface. The frame type is always ETHERNET, which stands for Ethernet II.

Table 2.6: CLI Display of IP Interface Information (Continued)

This Field...	Displays...
mtu	The Maximum Transmission Unit (MTU), which specifies the maximum packet size for packets sent and received on this interface.
metric	The cost associated with this interface.
RIP version	The version of RIP running on the interface. The version can be one of the following: • V1 • V2 **Note**: This field appears only if RIP is enabled on the interface.
Poison Reverse	Indicates whether RIP poison reverse is enabled or disabled. This field can have one of the following values: • off • on **Note**: This field appears only if RIP is enabled on the interface.
UDP Helper Addresses	The UDP helper addresses that are configured on this interface are listed underneath the fields described above. If no helper addresses are configured, the following message is displayed instead: "No Helper Addresses are configured."
Virtual interfaces (VEs)	
Interface Ve	The virtual interface (VE) number, which you assigned when you configured the interface.
members	The ports on which this VE is configured.
active	The member ports that are currently active. If none of the port members are active, this field lists "none".
port state	The interface state, which can be one of the following: • DOWN • UP
ip address	The IP address of this interface.
subnet mask	The network mask for the IP address of this interface.
encapsulation	The frame type used to encapsulate packets on this interface. The frame type is always ETHERNET, which stands for Ethernet II.
mtu	The Maximum Transmission Unit (MTU), which specifies the maximum packet size for packets sent and received on this interface.
metric	The cost associated with this interface.
UDP Helper Addresses	The UDP helper addresses that are configured on this interface are listed underneath the fields described above. If no helper addresses are configured, the following message is displayed instead: "No Helper Addresses are configured."

USING THE WEB MANAGEMENT INTERFACE

To display IP interface information:

1. Select the IP link. The IP panel is displayed. This panel shows basic IP configuration information and also contains links to other information and configuration panels.
2. Select the IP Interface link from the bottom of the panel. The IP Interface panel is displayed, as shown in Figure 9.22.

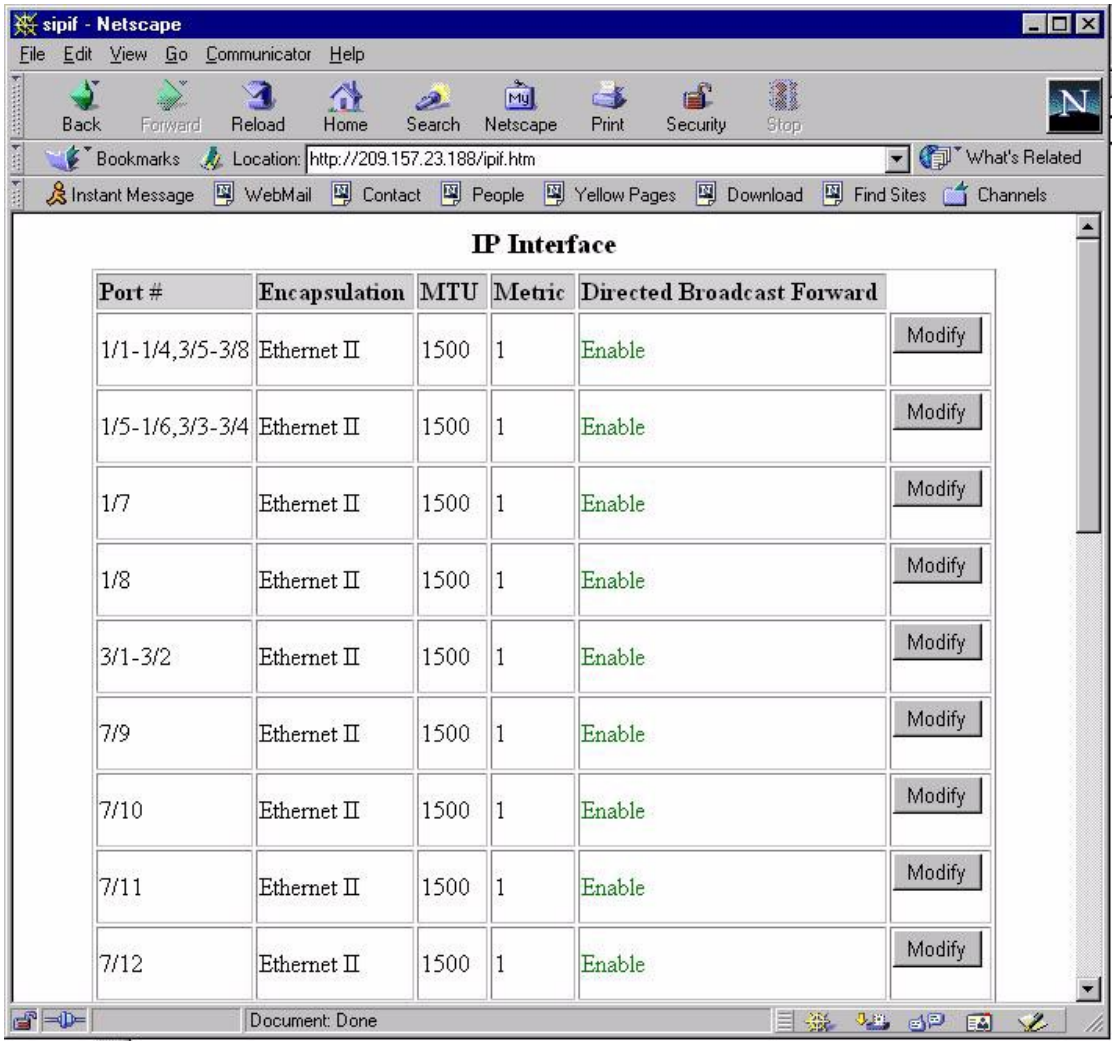

Figure 9.22 IP Interface table

This display shows the following information.

Table 2.7: Web Display of IP Interface Information

This Field...	Displays...
Port #	The physical port number or virtual interface (VE) number. VEs are shown as "v<num>", where <num> is the number you assigned to the VE when you configured it. For example, VE 1 is shown as "v1". If a range of ports is listed in this field, the interface is a trunk group. If two ranges of ports are listed, the interface is a trunk group that spans multiple chassis modules.
Encapsulation	The frame type used to encapsulate packets on this interface. The frame type is always Ethernet II.
MTU	The Maximum Transmission Unit (MTU), which specifies the maximum packet size for packets sent and received on this interface.
Metric	The cost associated with this interface.
Directed Broadcast Forward	The state of the directed broadcast forwarding feature. The state can be one of the following: • Disable • Enable To change the state of this feature, see “Enable or Disable Suppression of Directed Broadcasts” on page 9-8.

Displaying the ARP Cache

To display the ARP cache, use one of the following methods.

USING THE CLI

To display the ARP cache, enter the following command at any CLI level:

```
HP9300> show arp
```

Syntax: show arp [ethernet <num> | mac-address <xxxx.xxxx.xxxx> [<mask>] | <IP-addr> [<mask>]] [<num>]

Specify the MAC address mask as "f"s and "0"s, where "f"s are significant bits. Specify IP address masks in standard decimal mask format (for example, 255.255.0.0).

The optional <num> parameter lets you display the table beginning with a specific entry number.

This display shows the following information. The number in the left column of the CLI display is the row number of the entry in the ARP cache. This number is not related to the number you assign to static MAC entries.

Table 2.8: CLI Display of ARP Cache

This Field...	Displays...
IP Address	The IP address of the device.
MAC Address	The MAC address of the device.

Table 2.8: CLI Display of ARP Cache (Continued)

This Field...	Displays...
Type	The type, which can be one of the following: • Dynamic – The device learned the entry from an incoming packet. • Static – You added the entry manually.
Age	The number of minutes the entry has remained unused. If this value reaches the ARP aging period, the entry is removed from the cache. To display the ARP aging period, see "Displaying Global IP Configuration Information" on page 9-35. To change the ARP aging interval, see "Modifying the ARP Aging Period" on page 9-7. **Note**: Static entries do not age out.
Port	The port on which the entry was learned.

USING THE WEB MANAGEMENT INTERFACE

To display the IP host cache:

1. Select the Show link to display the Show Statistics panel. This panel lists statistics links for all the enabled features.
2. Select the ARP Cache link. The ARP Cache panel is displayed, as shown in Figure 9.23.

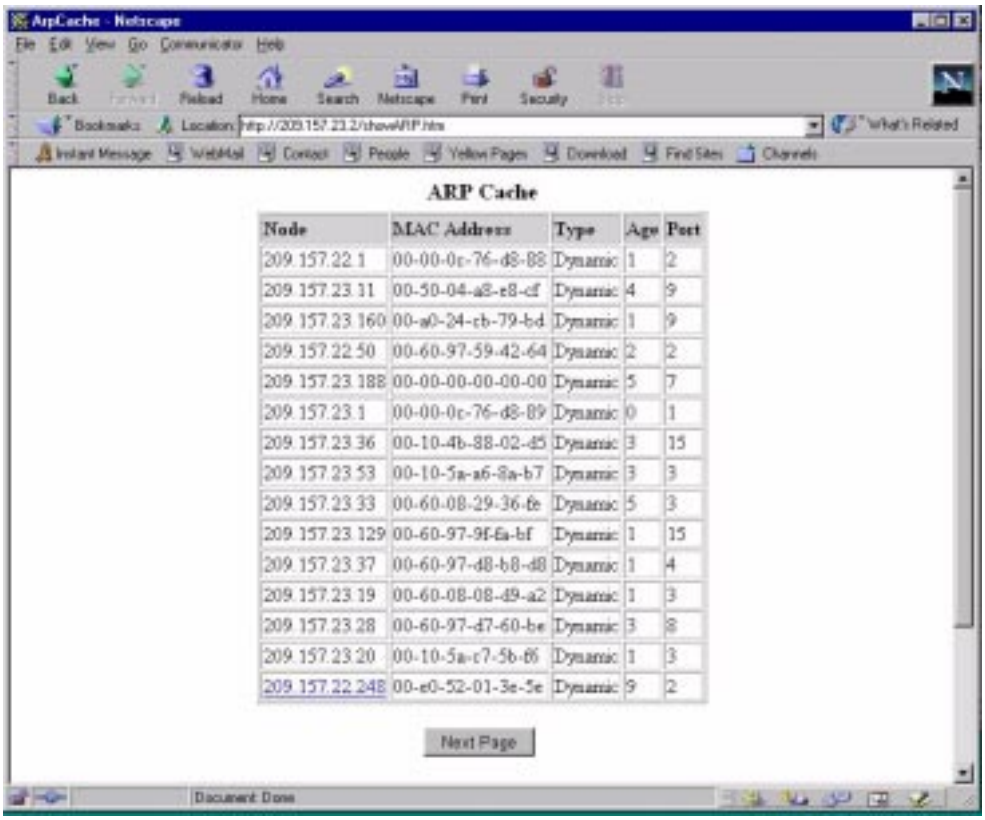

Figure 9.23 ARP Cache

This display shows the following information.

Table 2.9: Web Display of ARP Cache

This Field...	Displays...
Node	The IP address of the device.
MAC Address	The MAC address of the device.
Type	The type, which can be one of the following: • Dynamic – The device learned the entry from an incoming packet. • Static – You added the entry manually.
Age	The number of minutes the entry has remained unused. If this value reaches the ARP aging period, the entry is removed from the cache. To display the ARP aging period, see "Displaying Global IP Configuration Information" on page 9-35. To change the ARP aging interval, see "Modifying the ARP Aging Period" on page 9-7. **Note**: Static entries do not age out.
Port	The port on which the entry was learned.

Displaying the IP Host Cache

To display the IP host cache, use one of the following methods.

USING THE CLI

To display the IP host cache, enter the following command at any CLI level:

```
HP9300> show ip cache
```

Syntax: show ip cache [<ip-addr>] | [<num>]

The **<ip-addr>** parameter displays the cache entry for the specified IP address.

The **<num>** parameter displays the cache beginning with the row following the number you enter. For example, to begin displaying the cache at row 10, enter the following command: **show ip cache 9**.

The **show ip cache** command displays the following information.

Table 2.10: CLI Display of IP Host Cache

This Field...	Displays...
IP Address	The IP address of the destination.
Next Hop	The IP address of the next-hop router to the destination. This field contains either an IP address or the value DIRECT. DIRECT means the destination is either directly attached or the destination is an address on this routing switch. For example, the next hop for loopback addresses and broadcast addresses is shown as DIRECT.
MAC	The MAC address of the destination. **Note**: If the entry is type U (indicating that the destination is this routing switch), the address consists of zeroes.

Table 2.10: CLI Display of IP Host Cache (Continued)

This Field...	Displays...
Type	The type of host entry, which can be one or more of the following: • D – Dynamic • P – Permanent • F – Forward • U – Us • C – Complex Filter • W – Wait ARP • I – ICMP Deny • K – Drop • R – Fragment • S – Snap Encap
Port	The port through which this device reaches the destination. For destinations that are located on this device, the port number is shown as "n/a".
VLAN	Indicates the VLAN(s) the listed port is in.
Pri	The QoS priority of the port or VLAN.

USING THE WEB MANAGEMENT INTERFACE

To display the IP host cache:

1. Select the Show link to display the Show Statistics panel. This panel lists statistics links for all the enabled features.
2. Select the Cache link in the IP section of the panel. The IP Cache panel is displayed, as shown in Figure 9.24.

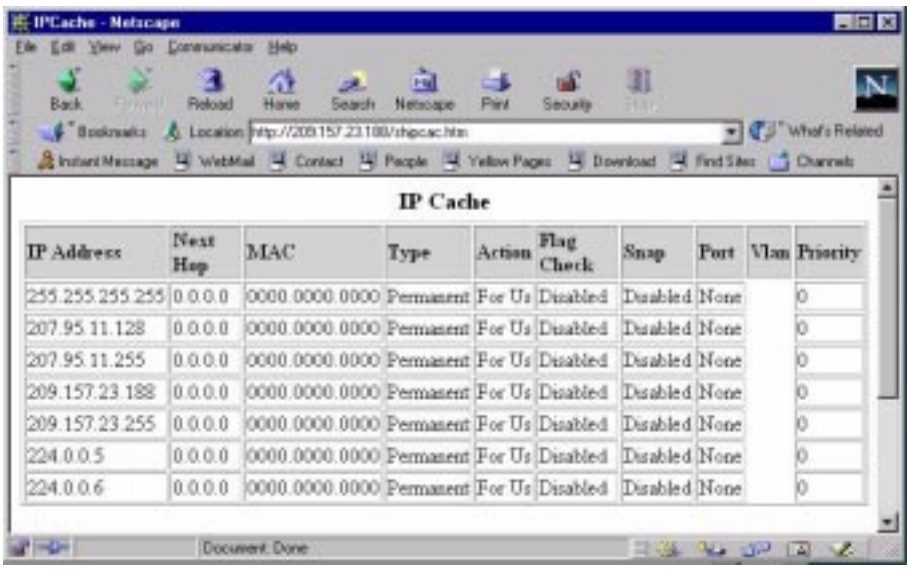

Figure 9.24 IP Host Cache

This display shows the following information.

Table 2.11: Web Display of IP Host Cache Information

This Field...	Displays...
IP Address	The IP address of the destination.
Next Hop	The IP address of the next-hop router to the destination. This field contains either an IP address or the value DIRECT. DIRECT means the destination is either directly attached or the destination is an address on this routing switch. For example, the next hop for loopback addresses and broadcast addresses is shown as DIRECT.
MAC	The MAC address of the destination. **Note**: If the entry is type U (indicating that the destination is this routing switch), the address consists of zeroes.
Type	The type of host entry, which can be one or more of the following: • D – Dynamic • P – Permanent • F – Forward • U – Us • C – Complex Filter • W – Wait ARP • I – ICMP Deny • K – Drop • R – Fragment • S – Snap Encap
Action	This information is used by HP customer support.
Flag Check	This information is used by HP customer support.
Snap	This information is used by HP customer support.
Port	The port through which this device reaches the destination. For destinations that are located on this device, the port number is shown as "n/a".
VLAN	Indicates the VLAN(s) the listed port is in.
Priority	The QoS priority of the port or VLAN.

Displaying the IP Forwarding Cache

To display the IP forwarding cache, use one of the following methods.

USING THE CLI

To display the IP forwarding cache, enter the following command at any CLI level:

```
HP9300> show ip flow-cache
```

Syntax: show ip flow-cache [<ip-addr>]

The **<ip-addr>** parameter displays the cache entry for the specified IP address.

The **show ip flow-cache** command displays the following information.

Table 2.12: CLI Display of IP Forwarding Cache

This Field...	Displays...
Source	The source IP address of the entry.
Dest.	The destination IP address of the entry.
Proto	This information is used by HP customer support.
Est	This information is used by HP customer support.
Port	This information is used by HP customer support.
Prio	This information is used by HP customer support.
Age	This information is used by HP customer support.
Total flow-cache entries used	The number of available flow-cache entries that are in use.
free	The number of flow-cache entries allocated in memory that are still available for use. On some devices, you can change the number of flow-cache entries available for use. See “Modifying System Parameter Default Settings” on page 8-69.

USING THE WEB MANAGEMENT INTERFACE

You cannot display the IP forwarding cache using the Web management interface

Displaying the IP Route Table

To display the IP route table, use one of the following methods.

USING THE CLI

To display the IP route table, enter the following command at any CLI level:

```
HP9300> show ip route

Total number of IP routes: 514
Starting index: 1  B:BGP D:Directly-Connected  R:RIP  S:Static  O:OSPF
IA:OSPF inter area  E1:OSPF external type 1  E2:OSPF external type 2

Destination       NetMask           Gateway           Port   Cost   Type

1.1.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.2.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.3.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.4.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.5.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.6.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.7.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.8.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.9.0.0           255.255.0.0       99.1.1.2          1/1    2      R
1.10.0.0          255.255.0.0       99.1.1.2          1/1    2      S
```

Syntax: show ip route [<IP-addr> | <num> | bgp | ospf | rip]

The **<ip-addr>** parameter displays the route to the specified IP address.

The **<num>** option display the route table entry whose row number corresponds to the number you specify. For example, if you want to display the tenth row in the table, enter “10”.

The **bgp** option displays the BGP4 routes.

The **ospf** option displays the OSPF routes.

The **rip** option displays the RIP routes.

The following table lists the information displayed by the **show ip route** command.

Table 2.13: CLI Display of IP Route Table

This Field...	Displays...
Destination	The destination network of the route.
NetMask	The network mask of the destination address.
Gateway	The next-hop router.
Port	The port through which this routing switch sends packets to reach the route's destination.
Cost	The route's cost.

Table 2.13: CLI Display of IP Route Table (Continued)

This Field...	Displays...

Type	The route type, which can be one of the following:
	• BGP – The route was learned from BGP.
	• D – The destination is directly connected to this routing switch.
	• R – The route was learned from RIP.
	• S – The route is a static route.
	• O – The route is an OSPF route. Unless you use the ospf option to display the route table, "O" is used for all OSPF routes. If you do use the ospf option, the following type codes are used:
	  • O – OSPF intra area route (within the same area).
	  • IA – The route is an OSPF inter area route (a route that passes from one area into another).
	  • E1 – The route is an OSPF external type 1 route.
	  • E2 – The route is an OSPF external type 2 route.

USING THE WEB MANAGEMENT INTERFACE

To display the IP route table:

1. Select the Show link to display the Show Statistics panel.
2. Select Routing Table in the IP section.

Displaying IP Traffic Statistics

To display IP traffic statistics, use one of the following methods.

USING THE CLI

To display IP traffic statistics, enter the following command at any CLI level:

```
HP9300> show ip traffic

  1439 received, 1180 sent, 0 forwarded
  0 filtered, 0 fragmented, 0 reassembled, 0 bad header
  0 no route, 0 unknown proto, 0 no buffer, 55 other errors


ICMP Statistics

Received:

  7 total, 0 errors, 0 unreachable, 0 time exceed
  0 parameter, 0 source quench, 0 redirect, 0 echo,
  0 echo reply, 0 timestamp, 0 timestamp reply, 0 addr mask
  0 addr mask reply, 0 irdp advertisment, 0 irdp solicitation

Sent:

  6 total, 0 errors, 0 unreachable, 0 time exceed
  0 parameter, 0 source quench, 0 redirect, 0 echo,
  0 echo reply, 0 timestamp, 0 timestamp reply, 0 addr mask
  0 addr mask reply, 6 irdp advertisment, 0 irdp solicitation
```

```
UDP Statistics

  206 received, 0 sent, 128 no port, 0 input errors


TCP Statistics

  0 active opens, 0 passive opens, 0 failed attempts
  0 active resets, 0 passive resets, 0 input errors
  1048 in segments, 1173 out segments, 0 retransmission

RIP Statistics

  0 requests sent, 0 requests received
  0 responses sent, 0 responses received
  0 unrecognized, 0 bad version, 0 bad addr family, 0 bad req format
  0 bad metrics, 0 bad resp format, 0 resp not from rip port
  0 resp from loopback, 0 packets rejected
```

Syntax: show ip traffic

The **show ip traffic** command displays the following information.

Table 2.14: CLI Display of IP Traffic Statistics

This Field...	Displays...
IP statistics	
received	The total number of IP packets received by the device.
sent	The total number of IP packets originated and sent by the device.
forwarded	The total number of IP packets received by the device and forwarded to other devices.
filtered	The total number of IP packets filtered by the device.
fragmented	The total number of IP packets fragmented by this device to accommodate the MTU of this device or of another device.
reassembled	The total number of fragmented IP packets that this device re-assembled.
bad header	The number of IP packets dropped by the device due to a bad packet header.
no route	The number of packets dropped by the device because there was no route.
unknown proto	The number of packets dropped by the device because the value in the Protocol field of the packet header is unrecognized by this device.
no buffer	This information is used by HP customer support.
other errors	The number of packets that this device dropped due to error types other than the types listed above.

Table 2.14: CLI Display of IP Traffic Statistics (Continued)

This Field...	Displays...
ICMP statistics	
The ICMP statistics are derived from RFC 792, "Internet Control Message Protocol", RFC 950, "Internet Standard Subnetting Procedure", and RFC 1256, "ICMP Router Discovery Messages". Statistics are organized into Sent and Received. The field descriptions below apply to each.	
total	The total number of ICMP messages sent or received by the device.
errors	This information is used by HP customer support.
unreachable	The number of Destination Unreachable messages sent or received by the device.
time exceed	The number of Time Exceeded messages sent or received by the device.
parameter	The number of Parameter Problem messages sent or received by the device.
source quench	The number of Source Quench messages sent or received by the device.
redirect	The number of Redirect messages sent or received by the device.
echo	The number of Echo messages sent or received by the device.
echo reply	The number of Echo Reply messages sent or received by the device.
timestamp	The number of Timestamp messages sent or received by the device.
timestamp reply	The number of Timestamp Reply messages sent or received by the device.
addr mask	The number of Address Mask Request messages sent or received by the device.
addr mask reply	The number of Address Mask Replies messages sent or received by the device.
irdp advertisement	The number of ICMP Router Discovery Protocol (IRDP) Advertisement messages sent or received by the device.
irdp solicitation	The number of IRDP Solicitation messages sent or received by the device.
UDP statistics	
received	The number of UDP packets received by the device.
sent	The number of UDP packets sent by the device.
no port	The number of UDP packets dropped because the packet did not contain a valid UDP port number.
input errors	This information is used by HP customer support.

Table 2.14: CLI Display of IP Traffic Statistics (Continued)

This Field...	Displays...
TCP statistics	
The TCP statistics are derived from RFC 793, "Transmission Control Protocol".	
active opens	The number of TCP connections opened by this device by sending a TCP SYN to another device.
passive opens	The number of TCP connections opened by this device in response to connection requests (TCP SYNs) received from other devices.
failed attempts	This information is used by 6208M-SX switch customer support.
active resets	The number of TCP connections this device reset by sending a TCP RESET message to the device at the other end of the connection.
passive resets	The number of TCP connections this device reset because the device at the other end of the connection sent a TCP RESET message.
input errors	This information is used by HP customer support.
in segments	The number of TCP segments received by the device.
out segments	The number of TCP segments sent by the device.
retransmission	The number of segments that this device restransmitted because the retransmission timer for the segment had expired before the device at the other end of the connection had acknowledged receipt of the segment.
RIP statistics	
The RIP statistics are derived from RFC 1058, "Routing Information Protocol".	
requests sent	The number of requests this device has sent to another RIP router for all or part of its RIP routing table.
requests received	The number of requests this device has received from another RIP router for all or part of this device's RIP routing table.
responses sent	The number of responses this device has sent to another RIP router's request for all or part of this device's RIP routing table.
responses received	The number of responses this device has received to requests for all or part of another RIP router's routing table.
unrecognized	This information is used by HP customer support.
bad version	The number of RIP packets dropped by the device because the RIP version was either invalid or is not supported by this device.
bad addr family	The number of RIP packets dropped because the value in the Address Family Identifier field of the packet's header was invalid.
bad req format	The number of RIP request packets this routing switch dropped because the format was bad.

Table 2.14: CLI Display of IP Traffic Statistics (Continued)

This Field...	Displays...
bad metrics	This information is used by HP customer support.
bad resp format	The number of responses to RIP request packets this routing switch dropped because the format was bad.
resp not from rip port	This information is used by HP customer support.
resp from loopback	The number of RIP responses received from loopback interfaces.
packets rejected	This information is used by HP customer support.

USING THE WEB MANAGEMENT INTERFACE

To display IP traffic statistics:

1. Select the Show link to display the Show Statistics panel. This panel lists statistics links for all the enabled features.
2. Select the Traffic link in the IP section of the panel. The IP Traffic panel is displayed, as shown in Figure 9.25 and Figure 9.26. You might need to scroll down to see all the fields.

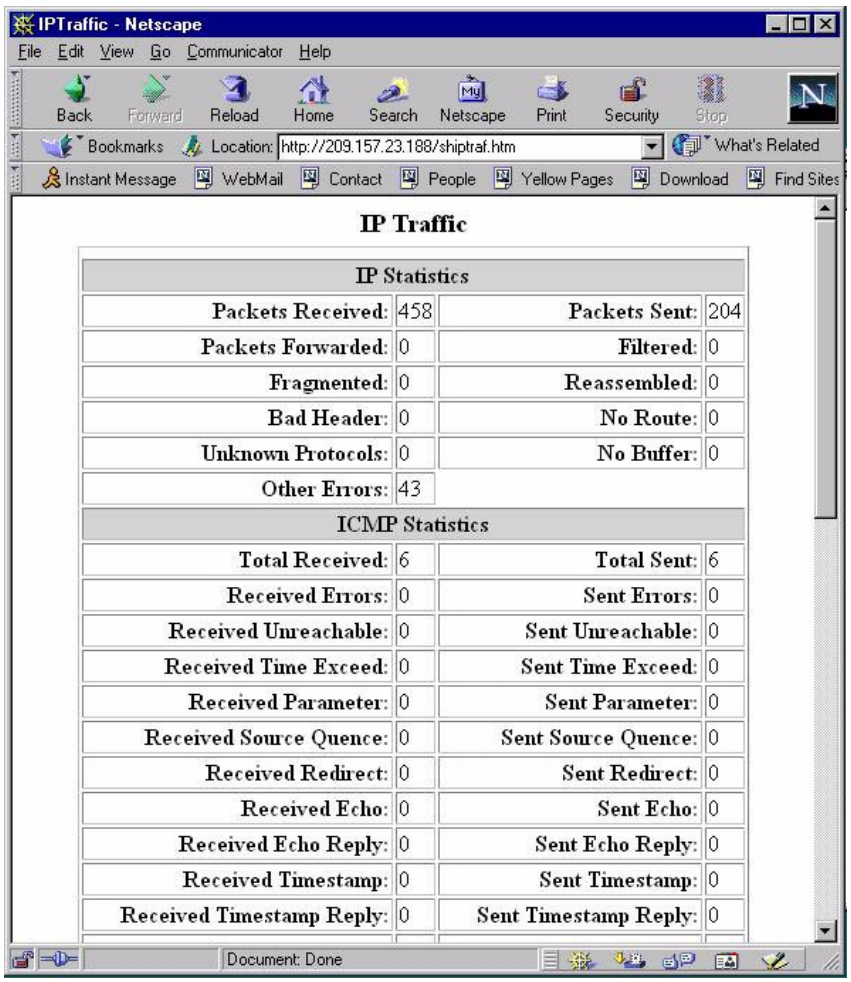

Figure 9.25 IP Traffic panel (1 of 2)

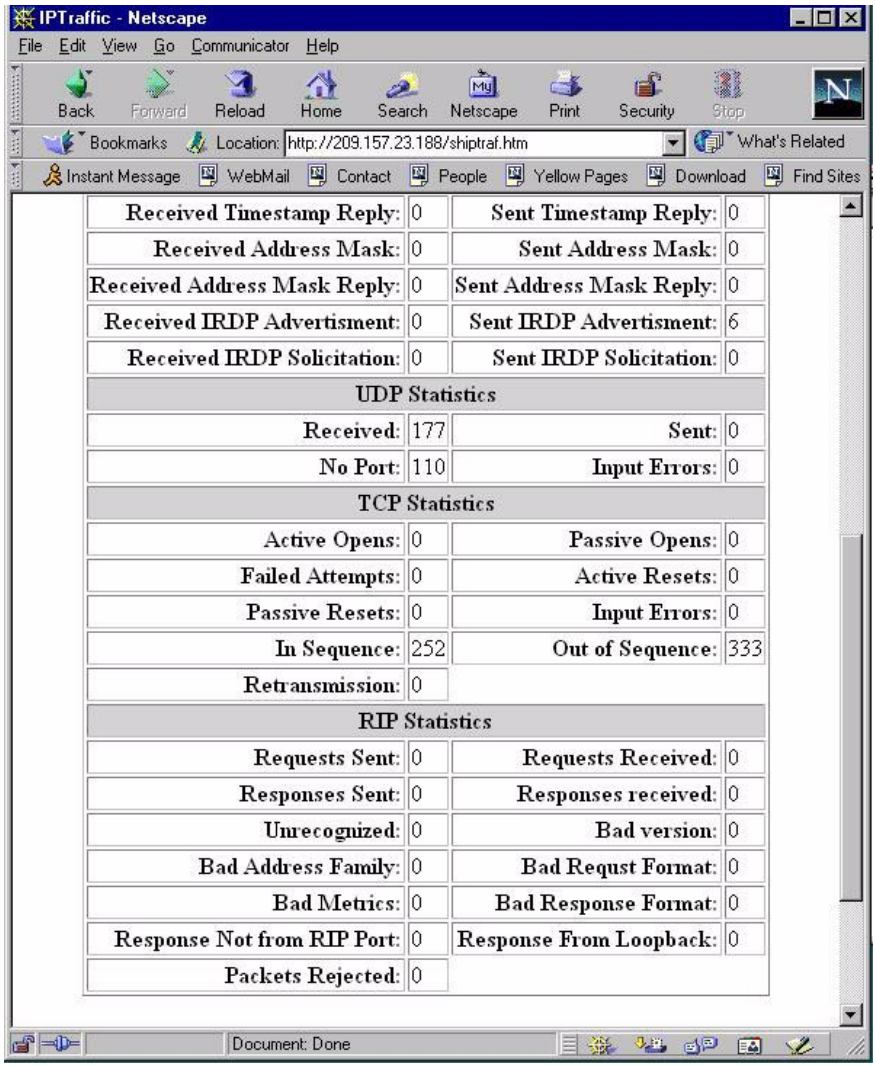

Received Timestamp Reply:	0	Sent Timestamp Reply:	0
Received Address Mask:	0	Sent Address Mask:	0
Received Address Mask Reply:	0	Sent Address Mask Reply:	0
Received IRDP Advertisment:	0	Sent IRDP Advertisment:	6
Received IRDP Solicitation:	0	Sent IRDP Solicitation:	0
UDP Statistics			
Received:	177	Sent:	0
No Port:	110	Input Errors:	0
TCP Statistics			
Active Opens:	0	Passive Opens:	0
Failed Attempts:	0	Active Resets:	0
Passive Resets:	0	Input Errors:	0
In Sequence:	252	Out of Sequence:	333
Retransmission:	0		
RIP Statistics			
Requests Sent:	0	Requests Received:	0
Responses Sent:	0	Responses received:	0
Unrecognized:	0	Bad version:	0
Bad Address Family:	0	Bad Requst Format:	0
Bad Metrics:	0	Bad Response Format:	0
Response Not from RIP Port:	0	Response From Loopback:	0
Packets Rejected:	0		

Figure 9.26 IP Traffic panel (2 of 2)

This display shows the following information.

Table 2.15: Web Display of IP Traffic Statistics

This Field...	Displays...
IP statistics	
Packets Received	The number of IP packets received by the device.
Packets Sent	The number of IP packets originated and sent by the device.
Packets Forwarded	The number of IP packets received from another device and forwarded by this device.
Filtered	The number of IP packets filtered by this device.
Fragmented	The number of IP packets fragmented by this device before sending or forwarding them.
Reassembled	The number of fragmented IP packets received and re-assembled by the device.
Bad Header	The number of packets dropped because they had a bad header.
No Route	The number of packets dropped because they had no route information.
Unknown Protocols	The number of packets dropped because they were using an unknown protocol.
No Buffer	The number of packets dropped because the device ran out of buffer space.
Other Errors	The number of packets dropped due to errors other than the ones listed above.
ICMP statistics	
Total Received	The number of ICMP packets received by the device.
Total Sent	The number of ICMP packets sent by the device.
Received Errors	This information is used by HP customer support.
Sent Errors	This information is used by HP customer support.
Received Unreachable	The number of Destination Unreachable messages received by the device.
Sent Unreachable	The number of Destination Unreachable messages sent by the device.
Received Time Exceed	The number of Time Exceeded messages received by the device.
Sent Time Exceed	The number of Time Exceeded messages sent by the device.
Received Parameter	The number of Parameter Problem messages received by the device.
Sent Parameter	The number of Parameter Problem messages sent by the device.
Received Source Quench	The number of Source Quench messages received by the device.
Sent Source Quench	The number of Source Quench messages sent by the device.

Table 2.15: Web Display of IP Traffic Statistics (Continued)

This Field...	Displays...
Received Redirect	The number of Redirect messages received by the device.
Sent Redirect	The number of Redirect messages sent by the device.
Received Echo	The number of Echo messages received by the device.
Sent Echo	The number of Echo messages sent by the device.
Received Echo Reply	The number of Echo messages received by the device.
Sent Echo Reply	The number of Echo messages sent by the device.
Received Timestamp	The number of Timestamp messages received by the device.
Sent Timestamp	The number of Timestamp messages sent by the device.
Received Timestamp Reply	The number of Timestamp Reply messages received by the device.
Sent Timestamp Reply	The number of Timestamp Reply messages sent by the device.
Received Address Mask	The number of Address Mask Request messages received by the device.
Sent Address Mask	The number of Address Mask Request messages sent by the device.
Received Address Mask Reply	The number of Address Mask Replies messages received by the device.
Sent Address Mask Reply	The number of Address Mask Replies messages sent by the device.
Received IRDP Advertisement	The number of ICMP Router Discovery Protocol (IRDP) Advertisement messages received by the device.
Sent IRDP Advertisement	The number of IRDP Advertisement messages sent by the device.
Received IRDP Solicitation	The number of IRDP Solicitation messages received by the device.
Sent IRDP Solicitation	The number of IRDP Solicitation messages sent by the device.
UDP statistics	
Received	The number of UDP packets received by the device.
Sent	The number of UDP packets sent by the device.
No Port	The number of UDP packets dropped because the packet did not contain a valid UDP port number.
Input Errors	This information is used by HP customer support.
TCP statistics The TCP statistics are derived from RFC 793, "Transmission Control Protocol".	
Active Opens	The number of TCP connections opened by this device by sending a TCP SYN to another device.
Passive Opens	The number of TCP connections opened by this device in response to connection requests (TCP SYNs) received from other devices.
Failed Attempts	This information is used by HP customer support.

Table 2.15: Web Display of IP Traffic Statistics (Continued)

This Field...	Displays...
Active Resets	The number of TCP connections this device reset by sending a TCP RESET message to the device at the other end of the connection.
Passive Resets	The number of TCP connections this device reset because the device at the other end of the connection sent a TCP RESET message.
Input Errors	This information is used by HP customer support.
In Segments	The number of TCP segments received by the device.
Out Segments	The number of TCP segments sent by the device.
Retransmission	The number of segments that this device restransmitted because the retransmission timer for the segment had expired before the device at the other end of the connection had acknowledged receipt of the segment.
RIP statistics The RIP statistics are derived from RFC 1058, "Routing Information Protocol".	
Requests Sent	The number of requests this device has sent to another RIP router for all or part of its RIP routing table.
Requests Received	The number of requests this device has received from another RIP router for all or part of this device's RIP routing table.
Responses Sent	The number of responses this device has sent to another RIP router's request for all or part of this device's RIP routing table.
Responses Received	The number of responses this device has received to requests for all or part of another RIP router's routing table.
Unrecognized	This information is used by HP customer support.
Bad Version	The number of RIP packets dropped by the device because the RIP version was either invalid or is not supported by this device.
Bad Address Family	The number of RIP packets dropped because the value in the Address Family Identifier field of the packet's header was invalid.
Bad Req Format	The number of RIP request packets this routing switch dropped because the format was bad.
Bad Metrics	This information is used by HP customer support.
Bad Response Format	The number of responses to RIP request packets this routing switch dropped because the format was bad.
Resp Not From RIP Port	This information is used by HP customer support.
Resp From Loopback	The number of RIP responses received from loopback interfaces.
Packets Rejected	This information is used by HP customer support.

Displaying RIP Filters

To display the RIP filters configured on the routing switch, use one of the following methods.

USING THE CLI

To display RIP filters, enter the following command at any CLI level:

```
HP9300> show ip rip

               RIP Route Filter Table
  Index   Action   Route IP Address    Subnet Mask
  1       deny     any                 any

               RIP Neighbor Filter Table
  Index   Action   Neighbor IP Address
  1       permit   any
```

Syntax: show ip rip

This display shows the following information.

Table 2.16: CLI Display of RIP Filter Information

This Field...	Displays...
Route filters The rows underneath "RIP Route Filter Table" list the RIP route filters. If no RIP route filters are configured on the device, the following message is displayed instead: "No Filters are configured in RIP Route Filter Table".	
Index	The filter number. You assign this number when you configure the filter.
Action	The action the routing switch takes if a RIP route packet matches the IP address and sub-net mask of the filter. The action can be one of the following: • deny – RIP route packets that match the address and network mask information in the filter are dropped. If applied to an interface's outbound filter group, the filter prevents the routing switch from advertising the route on that interface. If applied to an interface's inbound filter group, the filter prevents the routing switch from adding the route to its IP route table. • permit – RIP route packets that match the address and network mask information are accepted. If applied to an interface's outbound filter group, the filter allows the routing switch to advertise the route on that interface. If applied to an interface's inbound filter group, the filter allows the routing switch to add the route to its IP route table.
Route IP Address	The IP address of the route's destination network or host.
Subnet Mask	The network mask for the IP address.

Table 2.16: CLI Display of RIP Filter Information (Continued)

This Field...	Displays...
Neighbor filters The rows underneath "RIP Neighbor Filter Table" list the RIP neighbor filters. If no RIP neighbor filters are configured on the device, the following message is displayed instead: "No Filters are configured in RIP Neighbor Filter Table".	
Index	The filter number. You assign this number when you configure the filter.
Action	The action the routing switch takes for RIP route packets to or from the specified neighbor: • deny – If the filter is applied to an interface's outbound filter group, the filter prevents the routing switch from advertising RIP routes to the specified neighbor on that interface. If the filter is applied to an interface's inbound filter group, the filter prevents the routing switch from receiving RIP updates from the specified neighbor. • permit – If the filter is applied to an interface's outbound filter group, the filter allows the routing switch to advertise RIP routes to the specified neighbor on that interface. If the filter is applied to an interface's inbound filter group, the filter allows the routing switch to receive RIP updates from the specified neighbor.
Neighbor IP Address	The IP address of the RIP neighbor.

USING THE WEB MANAGEMENT INTERFACE

To display RIP filter information:

1. Select the RIP link.
2. Select one of the following links:
   - RIP Route Filter
   - RIP Neighbor Filter
   - Redistribution Filter