Chapter 4
IP Access Control Lists (ACLs)

NOTE: This chapter applies only to flow-based ACLs (also called CPU-based ACLs). For information about Enhanced Performance hardware-based ACLs, see "EP Hardware-Based IP Access Control Lists (ACLs)" on page 5-1.

HP's IP Access Control Lists (ACLs) enable you to do the following:

- Permit or deny management access to an HP device
- Forward or drop packets
- Provide input to other features:
  - Route maps – used by BGP4 to filter route information; also used by Policy-Based Routing (PBR) to selectively modify and route IP packets based on their source IP address
  - Adaptive Rate Limiting
  - RIP offset lists – used to set the metric for routes redistributed into RIP
  - Network Address Translation (NAT) address pools

This chapter describes how to configure IP ACLs to filter forwarded traffic and to perform Policy-Based Routing (PBR). For information about other uses of IP ACLs, see the following:

- Permit or deny management access to an HP device – see the *Security Guide*
- Route maps (when used for BGP4) – see "Defining Route Maps" on page 12-67
- Adaptive Rate Limiting – see "Characterizing the Traffic" on page 6-15
- RIP offset lists – see "Configuring a RIP Offset List" on page 9-5
- NAT address pools – see "Configuring Dynamic NAT Parameters" on page 14-6

HP Routing Switches also enable you to configure other types of ACLs, which are used by BGP4:

- AS-path ACL – filters BGP4 routes based on Autonomous System (AS) path information in the routes
- Community ACL – filters BGP4 routes based on the community information in the routes
- IP prefix list – filters routes based on the prefix (network portion) of the destination address or on the next-hop gateway

In general, these other types of ACLs are used to provide input to other commands or processes, such as route aggregation, route redistribution, and so on. For information, see "Configuring BGP4" on page 12-1.

NOTE: For optimal performance, apply deny ACLs to inbound ports instead of outbound ports. This way, traffic is dropped as it tries to enter the HP device, instead of being dropped after it has been forwarded internally to the outbound port.

NOTE: Outbound ACLs do not filter broadcast traffic or any traffic (including ICMP replies) generated by the HP device itself.

NOTE: You cannot use ACLs to filter based on MAC information or Quality of Service (QoS) information.

To filter based on MAC information, see the "Defining MAC Address Filters" section in the "Configuring Basic Features" chapter of the *Installation and Basic Configuration Guide*.

To filter based on QoS information, see "Assigning IP and Layer 4 Sessions to Priority Queues" on page 2-16.

NOTE: On Routing Switches, ACLs do not apply to traffic that is switched between one port and another in the same VLAN.

Overview

The following sections describe IP ACLs. For configuration information, see the following sections:

NOTE: This chapter describes how to configure ACLs using the HP device's CLI. You also can create and modify ACLs using a text editor on a file server, then copy them to the device's running-config file. In fact, this method is a convenient way to reorder individual ACL entries within an ACL. See "Modifying ACLs" on page 4-25.

How Flow-Based ACLs Work

Flow-based ACLs (also called CPU-based ACLs) work as follows:

When the device receives an IP packet, the device checks the receiving port's ACL CAM entries for an entry with the same address information as the packet.

- If the CAM contains a matching entry, the device takes the action specified by the entry (permit or deny).

NOTE: CAM entries are not programmed when you apply an ACL to an interface. CAM entries are created by the CPU when a packet received by the device matches a CAM entry on the inbound interface, as described below. The Layer 4 CAM entries programmed by the CPU for ACL matches age out if unused for 70 seconds.

NOTE: The CAM can contain entries for ACLs with deny actions only if you enable this support by entering the **hw-drop-acl-denied-packet** command.

- If the CAM does not contain a matching entry, the device sends the packet to the CPU for ACL comparison.
  - If the packet matches an ACL applied to inbound traffic on the port and the ACL has the permit action, the CPU programs an ACL permit entry into the Layer 4 CAM for the port that received the packet. The CAM entry contains the packet's address information.
  - If the packet matches an ACL applied to inbound traffic on the port and the ACL has the deny action, the CPU drops the packet but does not program an entry into the Layer 4 CAM, unless you have enabled the CPU to do so by entering the **hw-drop-acl-denied-packet** command.
  - If the packet does not match any of the inbound ACLs on the interface (and therefore matches an implicit **deny ip any any**), the CPU drops the packet. The CPU does not program an entry into the Layer 4 CAM, unless you have enabled the CPU to do so by entering the **hw-drop-acl-denied-packet** command.
- If the packet's outbound interface has an ACL applied to the outbound traffic direction, the device sends the packet to the CPU for filtering and either drops the packet or forwards the packet on the outbound interface, depending on the results of the ACL comparison.

How Fragmented Packets are Processed

The descriptions above apply to non-fragmented packets. In 07.6.04 and later, the default processing of fragments by flow-based ACLs (and hardware-based ACLs also) is as follows:

- The first fragment of a packet is permitted or denied using the ACLs. The first fragment is handled the same way as non-fragmented packets, since the first fragment contains the Layer 4 source and destination application port numbers. The device uses the Layer 4 CAM entry if one is programmed, or applies the interface's ACL entries to the packet and permits or denies the packet according to the first matching ACL.
- For other fragments of the same packet, one of the following occurs:
  - If the device has a CAM entry for the packet (or for previous packets in the same flow), and has not been configured to send the fragments to the CPU, the device uses the CAM entry to forward the fragments in hardware.

    The fragments are forwarded even if the first fragment, which contains the Layer 4 information, was denied. Generally, denying the first fragment of a packet is sufficient, since a transaction cannot be completed without the entire packet. However, for stricter fragment control, you can send fragments to the CPU for filtering.

- If the device is configured to send fragments to the CPU for filtering, the device compares the source and destination IP addresses to the ACL entries that contain Layer 4 information.

  —If the fragment's source and destination addresses exactly match an ACL entry that has Layer 4 information, the device assumes that the ACL entry is applicable to the fragment and permits or denies the fragment according to the ACL entry. The device does not compare the fragment to ACL entries that do not contain Layer 4 information.

  —If both the fragment's source and destination addresses do not exactly match an ACL entry, the device skips the ACL entry and compares the packet to the next ACL entry. This is true even if either the source or destination address (but not both) does exactly match an ACL entry.

  —If the source and destination addresses do not exactly match any ACL entry on the applicable interface, the device drops the fragment.

NOTE: By default, 10 Gigabit Ethernet modules also forward the first fragment instead of using the ACLs to permit or deny the fragment.

You can modify the handling of denied fragments by flow-based ACLs or hardware-based ACLs. In addition, you can throttle the fragment rate on an interface that uses hardware-based ACLs. See "Dropping All Fragments That Exactly Match an ACL" on page 4-27 and "Enabling ACL Filtering of Fragmented Packets" on page 5-8.

Hardware Aging of Layer 4 CAM Entries for Flow-Based ACLs

Flow-based ACLs and hardware-based ACLs both use Layer 4 CAM entries. The device permanently programs hardware-based ACLs into the CAM. The entries never age out. In software release 07.5.04 and later, the device does age out Layer 4 CAM entries for flow-based ACLs. A Layer 4 CAM entry for a flow-based ACL ages out if the entry is unused for 70 seconds. The age time is not configurable.

After an entry ages out, its CAM space becomes available for other ACL entries or other features that use the Layer 4 CAM.

Types of IP ACLs

You can configure the following types of IP ACLs:

- Standard – Permits or denies packets based on source IP address. Valid standard ACL IDs are 1 – 99 or a string.
- Extended – Permits or denies packets based on source and destination IP address and also based on IP protocol information. Valid extended ACL IDs are a number from 100 – 199 or a string.

ACL IDs and Entries

ACLs consist of ACL IDs and ACL entries:

- ACL ID – An ACL ID is a number from 1 – 99 (for a standard ACL) or 100 – 199 (for an extended ACL) or a character string. The ACL ID identifies a collection of individual ACL entries. When you apply ACL entries to an interface, you do so by applying the ACL ID that contains the ACL entries to the interface, instead of applying the individual entries to the interface. This makes applying large groups of access filters (ACL entries) to interfaces simple.

  NOTE: This is different from IP access policies. If you use IP access policies, you apply the individual policies to interfaces.

- ACL entry – An ACL entry is a filter command associated with an ACL ID. The maximum number of ACL entries you can configure is a system-wide parameter and depends on the device you are configuring. You can configure up to the maximum number of entries in any combination in different ACLs. The total number of entries in all ACLs cannot exceed the system maximum.

NOTE: Up to 1024 entries are supported on Routing Switches using Management 1 and Management 2 modules. Management 4 modules can support up to 4096 ACL entries.

You configure ACLs on a global basis, then apply them to the incoming or outgoing traffic on specific ports. You can apply only one ACL to a port's inbound traffic and only one ACL to a port's outbound traffic. The software applies the entries within an ACL in the order they appear in the ACL's configuration. As soon as a match is found, the software takes the action specified in the ACL entry (permit or deny the packet) and stops further comparison for that packet.

Default ACL Action

The default action when no ACLs are configured on a device is to permit all traffic. However, once you configure an ACL and apply it to a port, the default action for that port is to deny all traffic that is not explicitly permitted on the port.

- If you want to tightly control access, configure ACLs consisting of permit entries for the access you want to permit. The ACLs implicitly deny all other access.
- If you want to secure access in environments with many users, you might want to configure ACLs that consist of explicit deny entries, then add an entry to permit all access to the end of each ACL. The software permits packets that are not denied by the deny entries.

NOTE: Do not apply an empty ACL (an ACL ID without any corresponding entries) to an interface. If you accidentally do this, the software applies the default ACL action, deny all, to the interface and thus denies all traffic.

Controlling Management Access to the Device

You can use standard ACLs to control Telnet/SSH, Web, and SNMP access to an HP device. See the *Security Guide*.

ACL Logging

ACL logging is disabled by default. However, when you configure an ACL entry, you can enable logging for that entry by adding the **log** parameter to the end of the CLI command for the entry.

When you enable logging for an ACL entry, statistics for packets that match the permit or deny conditions of the ACL entry are logged. For example, if you configure a standard ACL entry to deny all packets from source address 209.157.22.26, statistics for packets that are explicitly denied by the ACL entry are logged in the HP device's Syslog buffer and in SNMP traps sent by the device.

The first time an ACL entry permits or denies a packet, the software immediately generates a Syslog entry and SNMP trap. The software also starts a five-minute timer. The timer keeps track of all packets explicitly denied by the ACL entries. After five minutes, the software generates a single Syslog entry for each ACL entry that has denied a packet. The message indicates the number of packets denied by the ACL entry during the previous five minutes.

If no ACL entries explicitly permit or deny packets during an entire five-minute timer interval, the timer stops. The timer restarts when an ACL entry explicitly permits or denies a packet.

NOTE: The timer for logging packets denied by Layer 2 filters is separate.

NOTE: The software generates log entries only when packets are explicitly permitted or explicitly denied by ACLs. The software does not generate log entries for implicitly permitted or denied entries.

Usage Guidelines for Access Control Lists (ACLs)

HP IP ACLs have the following uses:

- Filtering forwarded traffic through the device
- Controlling management access to the device itself
- Providing input to other features, such as route maps

Using ACLs as Input to Other Features

You can use ACLs to provide input to other features such as route maps and distribution lists. When you use an ACL this way, use permit statements in the ACL to specify the traffic that you want to send to the other feature. If you use deny statements, the traffic specified by the deny statements is not supplied to the other feature.

If you use an ACL in a route map and you use a wildcard character as the source IP address, make sure you apply the route map to interfaces, not globally. Otherwise, a loop can occur.

Using ACLs and Network Address Translation (NAT) on the Same Interface

NOTE: These guidelines do not apply to devices that are using the T-Flow management module. You can configure ACLs and NAT on the same port without needing to follow these guidelines.

You can use ACLs and NAT on the same interface, so long as you follow these guidelines:

- Do not enable NAT on an interface until you have applied ACLs (as described below) to the interface. If NAT is already enabled, you must disable it, apply the ACLs, then re-enable NAT on the interface.
- Enable the strict TCP mode.
- On the inside NAT interface (the one connected to the private addresses), apply inbound ACLs that permit TCP, UDP, and ICMP traffic to enter the device from the private sub-net.

You can use a standard ACL to permit all traffic (including TCP, UDP, and ICMP traffic) or an extended ACL with separate entries to explicitly permit TCP, UDP, and ICMP traffic.

NOTE: You do not need to apply ACLs to permit TCP, UDP, and ICMP traffic unless you are applying other ACLs to the interface as well. If you do not plan to apply any ACLs to a NAT interface, then you do not need to apply the ACLs to permit TCP, UDP, and ICMP traffic.

Here is an example of how to configure device to use ACLs and NAT on the same interfaces. In this example, the inside NAT interface is port 1/1 and the outside NAT interface is port 2/2.

The following commands enable the strict TCP mode and configure an ACL to permit all traffic from the 10.10.200.x sub-net. A second ACL denies traffic from a specific host on the Internet.

```
HP9300(config)# ip strict-acl-tcp
HP9300(config)# access-list 1 permit 10.10.200.0 0.0.0.255
HP9300(config)# access-list 2 deny 209.157.2.184
```

The following commands configure global NAT parameters.

```
HP9300(config)# ip nat inside source list 1 pool outadds overload
HP9300(config)# ip nat pool outadds 204.168.2.1 204.168.2.254 netmask 255.255.255.0
```

The following commands configure the inside and outside NAT interfaces. Notice that the ACLs are applied to the inbound direction on the inside NAT interface, and are applied ***before*** NAT is enabled. In this example, ACL 1 permits all traffic to come into the inside interface from the private sub-net. ACL 2 denies traffic from a specific host from going out the interface to the private sub-net.

```
HP9300(config)# interface ethernet 1/1
HP9300(config-if-1/1)# ip address 10.10.200.1 255.255.255.0
HP9300(config-if-1/1)# ip access-group 1 in
HP9300(config-if-1/1)# ip access-group 2 out
HP9300(config-if-1/1)# ip nat inside
HP9300(config-if-1/1)# interface ethernet 2/2
HP9300(config-if-2/2)# ip address 204.168.2.78 255.255.255.0
```

```
HP9300(config-if-2/2)# ip nat outside
```

NOTE: If you are using software release 07.6.04 or later, enter the **ip rebind-acl** command at the global CONFIG level of the CLI to place the **ip strict-acl-tcp** command into effect.

Requirement for Applying ACL Configuration Changes

NOTE: This section applies to software release 07.6.04 and later.

For flow-based and hardware-based ACLs, if you make an ACL configuration change, you must reapply the ACLs to their interfaces to place the change into effect. An ACL configuration change includes any of the following:

- Adding, changing, or removing an ACL or an entry in an ACL
- Changing a PBR policy
- Enabling or disabling the TCP strict mode or UDP strict mode (flow-based ACLs only)
- Changing EP ToS-based QoS mappings (since EP QoS uses the Layer 4 CAM)

Reapplying Modified ACLs

To reapply ACLs following an ACL configuration change, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# ip rebind-acl all
```

Syntax: [no] ip rebind-acl <num> | <name> | all

Disabling or Re-Enabling Access Control Lists (ACLs)

A Routing Switch cannot actively use both IP access policies and ACLs for filtering IP traffic. When you boot a Routing Switch with software release 06.5.00 or higher, the software checks the device's startup-config file for **ip access-policy-group** commands, which associate IP access policies with ports. If the software finds an **ip access-policy-group** command in the file, the software disables all packet-forwarding ACLs (those associated with specific ports) and also prevents you from applying an ACL to a port.

The next time you save the startup-config file, the software adds the following command near the top of the file, underneath the **ver** (software version) statement:

ip dont-use-acl

This command disables all packet-forwarding ACLs (those associated with specific ports) and also prevents you from associating an ACL with a port. However, the command does not remove existing ACLs from the startup-config file. In addition, the command does not affect ACLs used for controlling management access to the device.

Enabling ACL Mode

If you try to apply an ACL to a port when the ACL mode is disabled (when the **ip dont-use-acl** command is in effect), a message is displayed, as shown in the following CLI example:

```
HP9300(config-if-e1000-1/1)# ip access-group 1 out
Must enable ACL mode first by using no ip dont-use-acl command and removing all ip
access-policy-group commands from interfaces, write memory and reload
```

As the message states, if you want to use ACLs, you must first enable the ACL mode. To do so, use either of the following methods.

USING THE CLI

To enable the ACL mode, enter the following commands:

```
HP9300(config-if-e1000-1/1)# exit
```

```
HP9300(config)# no ip dont-use-acl
HP9300(config)# write memory
HP9300(config)# end
HP9300# reload
```

The **write memory** command removes the **ip dont-use-acl** command from the startup-config file. The **reload** command reloads the software. When the software finishes loading, you can apply ACLs to ports.

The commands that configure the IP access policies and apply them to ports remain in the startup-config file in case you want to use them again, but they are disabled. If you later decide you want to use the IP access policies again instead of ACLs, you must disable the ACL mode again. See the following section.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the Routing Switch using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.
4. Click on the General link to display the IP configuration panel.
5. Select the Enable radio button next to Access Control List.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Disabling ACL Mode

If the ACL mode is enabled, a message is displayed when you try to apply an IP access policy to a port, as shown in the following CLI example:

```
HP9300(config-if-e1000-1/1)# ip access-policy-group 1 in
Must disable ACL mode first by using ip dont-use-acl command, write memory and
reload
```

To use the IP access policies, you first must disable the ACL mode using either of the following methods.

USING THE CLI

To disable the ACL mode, enter the following commands:

```
HP9300(config-if-e1000-1/1)# exit
HP9300(config)# ip dont-use-acl
HP9300(config)# write memory
HP9300(config)# end
HP9300# reload
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the Routing Switch using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.
4. Click on the General link to display the IP configuration panel.
5. Select the Disable radio button next to Access Control List.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring Standard ACLs

This section describes how to configure standard ACLs with numeric IDs.

- For configuration information on named ACLs, see “Configuring Named ACLs” on page 4-22.
- For configuration information on extended ACLs, see “Configuring Extended ACLs” on page 4-12.

Standard ACLs permit or deny packets based on source IP address. You can configure up to 99 standard ACLs. You can configure up to 1024 individual ACL entries on a device. There is no limit to the number of ACL entries an ACL can contain except for the system-wide limitation of 1024 total ACL entries.

USING THE CLI

To configure a standard ACL and apply it to outgoing traffic on port 1/1, enter the following commands.

```
HP9300(config)# access-list 1 deny host 209.157.22.26 log
HP9300(config)# access-list 1 deny 209.157.29.12 log
HP9300(config)# access-list 1 deny host IPHost1 log
HP9300(config)# access-list 1 permit any
HP9300(config)# int eth 1/1
HP9300(config-if-1/1)# ip access-group 1 out
HP9300(config)# write memory
```

The commands in this example configure an ACL to deny packets from three source IP addresses from being forwarded on port 1/1. The last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries.

Standard ACL Syntax

Syntax: [no] access-list <num> deny | permit <source-ip> | <hostname> <wildcard> [log]

or

Syntax: [no] access-list <num> deny | permit <source-ip>/<mask-bits> | <hostname> [log]

Syntax: [no] access-list <num> deny | permit host <source-ip> | <hostname> [log]

Syntax: [no] access-list <num> deny | permit any [log]

Syntax: [no] ip access-group <num> in | out

The <num> parameter is the access list number and can be from 1 – 99.

The **deny | permit** parameter indicates whether packets that match a policy in the access list are denied (dropped) or permitted (forwarded).

The <source-ip> parameter specifies the source IP address. Alternatively, you can specify the host name.

NOTE: To specify the host name instead of the IP address, the host name must be configured using the HP device’s DNS resolver. To configure the DNS resolver name, use the **ip dns server-address…** command at the global CONFIG level of the CLI.

The <wildcard> parameter specifies the mask value to compare against the host address specified by the <source-ip> parameter. The <wildcard> is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet’s source address must match the <source-ip>. Ones mean any value matches. For example, the <source-ip> and <wildcard> values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C sub-net 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in CIDR format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “209.157.22.26 0.0.0.255” as “209.157.22.26/24”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into ones. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of sub-net lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP sub-net masks in CIDR format, the mask is saved in the file in "/<mask-bits>" format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE: If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with sub-net mask in the display produced by the **show ip access-list** command.

The **host** <source-ip> | <hostname> parameter lets you specify a host IP address or name. When you use this parameter, you do not need to specify the mask. A mask of all zeros (0.0.0.0) is implied.

The **any** parameter configures the policy to match on all host addresses.

The **log** argument configures the device to generate Syslog entries and SNMP traps for packets that are permitted or denied by the access policy.

NOTE: You can enable logging on ACLs and filters that support logging even when the ACLs and filters are already in use. To do so, re-enter the ACL or filter command and add the **log** parameter to the end of the ACL or filter. The software replaces the ACL or filter command with the new one. The new ACL or filter, with logging enabled, takes effect immediately.

The **in | out** parameter specifies whether the ACL applies to incoming traffic or outgoing traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or virtual interface.

NOTE: If the ACL is for the inbound traffic direction on a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface. See "Configuring Named ACLs" on page 4-22.

USING THE WEB MANAGEMENT INTERFACE

To configure a standard ACL:

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to System or IP. You can access the ACL configuration panels from either location.
4. Select the Standard ACL link.

   - If the device does not already have some standard ACLs, the Standard ACL configuration panel is displayed, as shown in the following example.

   Otherwise, if the device already has some standard ACLs, the Standard ACL table is displayed. This table lists the configured ACLs. Select the Add Standard ACL link to display the Standard ACL configuration panel,

as shown in the following example.

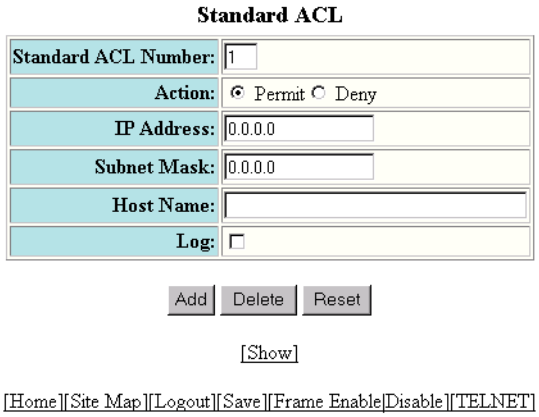

5. Change the ACL number in the Standard ACL Number field or use the ACL number displayed in the field.

NOTE: You cannot specify an ACL name.

6. Select the ACL action. You can select Permit or Deny:
   - Permit – Forwards traffic or allows management access for the specified IP source.
   - Deny – Drops traffic or denies management access for the specified IP source.

NOTE: If the ACL is a forwarding ACL, the action forwards or drops the traffic. If the ACL is a management access ACL, the action permits or denies management access.

7. Enter the source information. You can enter the source IP address and network mask or the host name.
   - If you enter the address, you also must enter the network mask. To specify “any”, enter “0.0.0.0”.
   - If you enter a host name instead of an IP address, when you click Add to add the ACL, the Web management interface sends a DNS query for the address. For the query to be successful, the device must have network access to a DNS server and the server must have an Address record for the host. In addition, the device must be configured with a DNS domain name and the IP address of the DNS server.
8. If you specified the Deny action, optionally enable logging by selecting the Log checkbox. If you enable logging for this ACL entry, the software generates Syslog entries for traffic that the ACL denies.
9. Select the IP Access Group link from the tree view.
   - If the device does not already have some ACLs applied to interfaces, the IP Access Group configuration panel is displayed, as shown in the following example.
   - Otherwise, if the device already has some ACLs applied to interfaces, the IP Access Group table is displayed. Select the Add link to display the IP Access Group configuration panel, as shown in the

following example.

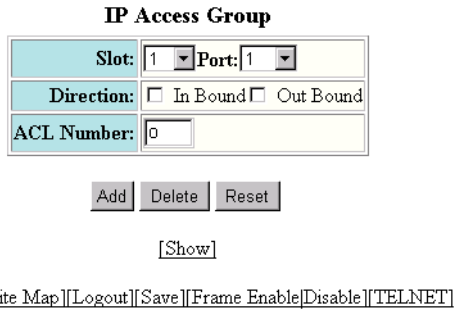

10. Select the Slot and port from the Slot and Port pulldown menus.
11. Specify the traffic direction to which the ACL applies. You can select one or both of the following:
    - In Bound – The ACL applies to traffic received on the port from other devices.
    - Out Bound – The ACL applies to traffic this HP device queues for transmission on the port.
12. Enter the ACL number in the ACL Number field.

NOTE: You cannot specify an ACL name.

13. Click the Add button to save the ACL and the association of the ACL with an interface to the device's running-config file.
14. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on Save to Flash.

Configuring Extended ACLs

This section describes how to configure extended ACLs.

- For configuration information on named ACLs, see "Configuring Named ACLs" on page 4-22.
- For configuration information on standard ACLs, see "Configuring Standard ACLs" on page 4-9.

Extended ACLs let you permit or deny packets based on the following information:

- IP protocol
- Source IP address or host name
- Destination IP address or host name
- Source TCP or UDP port (if the IP protocol is TCP or UDP)
- Destination TCP or UDP port (if the IP protocol is TCP or UDP)

The IP protocol can be one of the following well-known names or any IP protocol number from 0 – 255:

- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- Internet Gateway Routing Protocol (IGRP)

- Internet Protocol (IP)
- Open Shortest Path First (OSPF)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

For TCP and UDP, you also can specify a comparison operator and port name or number. For example, you can configure a policy to block web access to a specific website by denying all TCP port 80 (HTTP) packets from a specified source IP address to the website's IP address.

USING THE CLI

To configure an extended access list that blocks all Telnet traffic received on port 1/1 from IP host 209.157.22.26, enter the following commands.

```
HP9300(config)# access-list 101 deny tcp host 209.157.22.26 any eq telnet log
HP9300(config)# access-list 101 permit ip any any
HP9300(config)# int eth 1/1
HP9300(config-if-1/1)# ip access-group 101 in
HP9300(config)# write memory
```

Here is another example of commands for configuring an extended ACL and applying it to an interface. These examples show many of the syntax choices. Notice that some of the entries are configured to generate log entries while other entries are not thus configured.

```
HP9300(config)# access-list 102 perm icmp 209.157.22.0/24 209.157.21.0/24
HP9300(config)# access-list 102 deny igmp host rkwong 209.157.21.0/24 log
HP9300(config)# access-list 102 deny igrp 209.157.21.0/24 host rkwong log
HP9300(config)# access-list 102 deny ip host 209.157.21.100 host 209.157.22.1 log
HP9300(config)# access-list 102 deny ospf any any log
HP9300(config)# access-list 102 permit ip any any
```

The first entry permits ICMP traffic from hosts in the 209.157.22.*x* network to hosts in the 209.157.21.*x* network.

The second entry denies IGMP traffic from the host device named "rkwong" to the 209.157.21.*x* network.

The third entry denies IGRP traffic from the 209.157.21.*x* network to the host device named "rkwong".

The fourth entry denies all IP traffic from host 209.157.21.100to host 209.157.22.1 and generates Syslog entries for packets that are denied by this entry.

The fifth entry denies all OSPF traffic and generates Syslog entries for denied traffic.

The sixth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 102 to the incoming and outgoing traffic on port 1/2 and to the incoming traffic on port 4/3.

```
HP9300(config)# int eth 1/2
HP9300(config-if-1/2)# ip access-group 102 in
HP9300(config-if-1/2)# ip access-group 102 out
HP9300(config-if-1/2)# exit
HP9300(config)# int eth 4/3
HP9300(config-if-4/3)# ip access-group 102 in
HP9300(config)# write memory
```

Here is another example of an extended ACL.

```
HP9300(config)# access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24
HP9300(config)# access-list 103 deny tcp 209.157.21.0/24 eq ftp 209.157.22.0/24
HP9300(config)# access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24 lt telnet
neq 5
HP9300(config)# access-list 103 deny udp any range 5 6 209.157.22.0/24 range 7 8
HP9300(config)# access-list 103 permit ip any any
```

The first entry in this ACL denies TCP traffic from the 209.157.21.*x* network to the 209.157.22.x network.

The second entry denies all FTP traffic from the 209.157.21.*x* network to the 209.157.22.x network.

The third entry denies TCP traffic from the 209.157.21.*x* network to the 209.157.22.*x* network, if the TCP port number of the traffic is less than the well-known TCP port number for Telnet (23), and if the TCP port is not equal to 5. Thus, TCP packets whose TCP port numbers are 5 or are greater than 23 are allowed.

The fourth entry denies UDP packets from any source to the 209.157.22.*x* network, if the UDP port number from the source network is 5 or 6 and the destination UDP port is 7 or 8.

The fifth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 103 to the incoming and outgoing traffic on ports 2/1 and 2/2.

```
HP9300(config)# int eth 2/1
HP9300(config-if-2/1)# ip access-group 103 in
HP9300(config-if-2/1)# ip access-group 103 out
HP9300(config-if-2/1)# exit
HP9300(config)# int eth 2/2
HP9300(config-if-2/2)# ip access-group 103 in
HP9300(config-if-2/2)# ip access-group 103 out
HP9300(config)# write memory
```

Filtering on IP Precedence and ToS Values

To configure an extended IP ACL that matches based on IP precedence, enter commands such as the following:

```
HP9300(config)# access-list 103 deny tcp 209.157.21.0/24 209.157.22.0/24
precedence internet
HP9300(config)# access-list 103 deny tcp 209.157.21.0/24 eq ftp 209.157.22.0/24
precedence 6
HP9300(config)# access-list 103 permit ip any any
```

The first entry in this ACL denies TCP traffic from the 209.157.21.*x* network to the 209.157.22.x network, if the traffic has the IP precedence option “internet” (equivalent to “6”).

The second entry denies all FTP traffic from the 209.157.21.*x* network to the 209.157.22.x network, if the traffic has the IP precedence value “6” (equivalent to “internet”).

The third entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

To configure an IP ACL that matches based on ToS, enter commands such as the following:

```
HP9300(config)# access-list 104 deny tcp 209.157.21.0/24 209.157.22.0/24 tos normal
HP9300(config)# access-list 104 deny tcp 209.157.21.0/24 eq ftp 209.157.22.0/24 tos 13
HP9300(config)# access-list 104 permit ip any any
```

The first entry in this IP ACL denies TCP traffic from the 209.157.21.*x* network to the 209.157.22.x network, if the traffic has the IP ToS option “normal” (equivalent to “0”).

The second entry denies all FTP traffic from the 209.157.21.*x* network to the 209.157.22.x network, if the traffic has the IP precedence value “13” (equivalent to “max-throughput”, “min-delay”, and “min-monetary-cost”).

The third entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

Extended ACL Syntax

Syntax: access-list <num> deny | permit <ip-protocol> <source-ip> | <hostname> <wildcard> [<operator> <source-tcp/udp-port>] <destination-ip> | <hostname> [<icmp-type>] <wildcard> [<operator> <destination-tcp/udp-port>] [precedence <name> | <num>] [tos <name> | <num>] [log]

Syntax: [no] access-list <num> deny | permit host <ip-protocol> any any [log]

Syntax: [no] ip access-group <num> in | out

The <num> parameter indicates the ACL number and be from 100 – 199 for an extended ACL.

The **deny | permit** parameter indicates whether packets that match the policy are dropped or forwarded.

The <ip-protocol> parameter indicates the type of IP packet you are filtering. In release 07.6.04 and later, you can specify a well-known name for any protocol whose number is less than 255. For other protocols, you must enter the number. Enter “?” instead of a protocol to list the well-known names recognized by the CLI.

The <source-ip> | <hostname> parameter specifies the source IP host for the policy. If you want the policy to match on all source addresses, enter **any**.

The <wildcard> parameter specifies the portion of the source IP host address to match against. The <wildcard> is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet’s source address must match the <source-ip>. Ones mean any value matches. For example, the <source-ip> and <wildcard> values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C sub-net 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in Classless Interdomain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of “209.157.22.26 0.0.0.255” as “209.157.22.26/24”. The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of sub-net lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP sub-net masks in CIDR format, the mask is saved in the file in “/<mask-bits>” format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE: If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with sub-net mask in the display produced by the **show ip access-list** command.

The <destination-ip> | <hostname> parameter specifies the destination IP host for the policy. If you want the policy to match on all destination addresses, enter **any**.

The <icmp-type> parameter specifies the ICMP protocol type.

NOTE: This parameter applies only if you specified **icmp** as the <ip-protocol> value.

NOTE: If you do not specify a message type, the ACL applies to all types of ICMP messages. The <num> parameter can be a value from 0 – 255.

This parameter can have one of the following values:

- echo
- echo-reply
- information-request
- mask-reply
- mask-request
- parameter-problem
- redirect
- source-quench
- time-exceeded
- timestamp-reply
- timestamp-request
- unreachable
- <num>

The <operator> parameter specifies a comparison operator for the TCP or UDP port number. This parameter applies only when you specify **tcp** or **udp** as the IP protocol. For example, if you are configuring an entry 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 **neq**.
- **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.

  NOTE: This operator applies only to destination TCP ports, not source TCP ports.

The <tcp/udp-port> parameter specifies the TCP or UDP port number or well-known name. In release 07.6.04 and later, you can specify a well-known name for any application port whose number is less than 1024. For other application ports, you must enter the number. Enter “?” instead of a port to list the well-known names recognized by the CLI.

The **in | out** parameter specifies whether the ACL applies to incoming traffic or outgoing traffic on the interface to which you apply the ACL. You can apply the ACL to an Ethernet port or a virtual interface.

NOTE: If the ACL is for the inbound traffic direction on a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface. See "Configuring Named ACLs" on page 4-22.

The **precedence** <name> | <num> parameter of the **ip access-list** command specifies the IP precedence. The precedence option for of an IP packet is set in a three-bit field following the four-bit header-length field of the packet's header. You can specify one of the following:

- **critical** or **5** – The ACL matches packets that have the critical precedence. If you specify the option number instead of the name, specify number 5.
- **flash** or **3** – The ACL matches packets that have the flash precedence. If you specify the option number instead of the name, specify number 3.
- **flash-override** or **4** – The ACL matches packets that have the flash override precedence. If you specify the option number instead of the name, specify number 4.
- **immediate** or **2** – The ACL matches packets that have the immediate precedence. If you specify the option number instead of the name, specify number 2.
- **internet** or **6** – The ACL matches packets that have the internetwork control precedence. If you specify the option number instead of the name, specify number 6.
- **network** or **7** – The ACL matches packets that have the network control precedence. If you specify the option number instead of the name, specify number 7.
- **priority** or **1** – The ACL matches packets that have the priority precedence. If you specify the option number instead of the name, specify number 1.
- **routine** or **0** – The ACL matches packets that have the routine precedence. If you specify the option number instead of the name, specify number 0.

The **tos** <name> | <num> parameter of the **ip access-list** command specifies the IP ToS. You can specify one of the following:

- **max-reliability** or **2** – The ACL matches packets that have the maximum reliability ToS. The decimal value for this option is 2.
- **max-throughput** or **4** – The ACL matches packets that have the maximum throughput ToS. The decimal value for this option is 4.
- **min-delay** or **8** – The ACL matches packets that have the minimum delay ToS. The decimal value for this option is 8.
- **min-monetary-cost** or **1** – The ACL matches packets that have the minimum monetary cost ToS. The decimal value for this option is 1.
- **normal** or **0** – The ACL matches packets that have the normal ToS. The decimal value for this option is 0.
- <num> – A number from 0 – 15 that is the sum of the numeric values of the options you want. The ToS field is a four-bit field following the Precedence field in the IP header. You can specify one or more of the following. To select more than one option, enter the decimal value that is equivalent to the sum of the numeric values of all the ToS options you want to select. For example, to select the **max-reliability** and **min-delay** options, enter number 10. To select all options, select 15.

The **log** parameter enables SNMP traps and Syslog messages for packets denied by the ACL.

NOTE: You can enable logging on ACLs and filters that support logging even when the ACLs and filters are already in use. To do so, re-enter the ACL or filter command and add the **log** parameter to the end of the ACL or filter. The software replaces the ACL or filter command with the new one. The new ACL or filter, with logging enabled, takes effect immediately.

USING THE WEB MANAGEMENT INTERFACE

To configure an extended ACL:

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to System or IP. You can access the ACL configuration panels from either location.
4. Select the Extended ACL link.
    - If the device does not already have some extended ACLs, the Extended ACL configuration panel is displayed, as shown in the following example.

- Otherwise, if the device already has some extended ACLs, the Extended ACL table is displayed. This table lists the configured ACLs. Select the Add Extended ACL link to display the Extended ACL configuration panel, as shown in the following example.

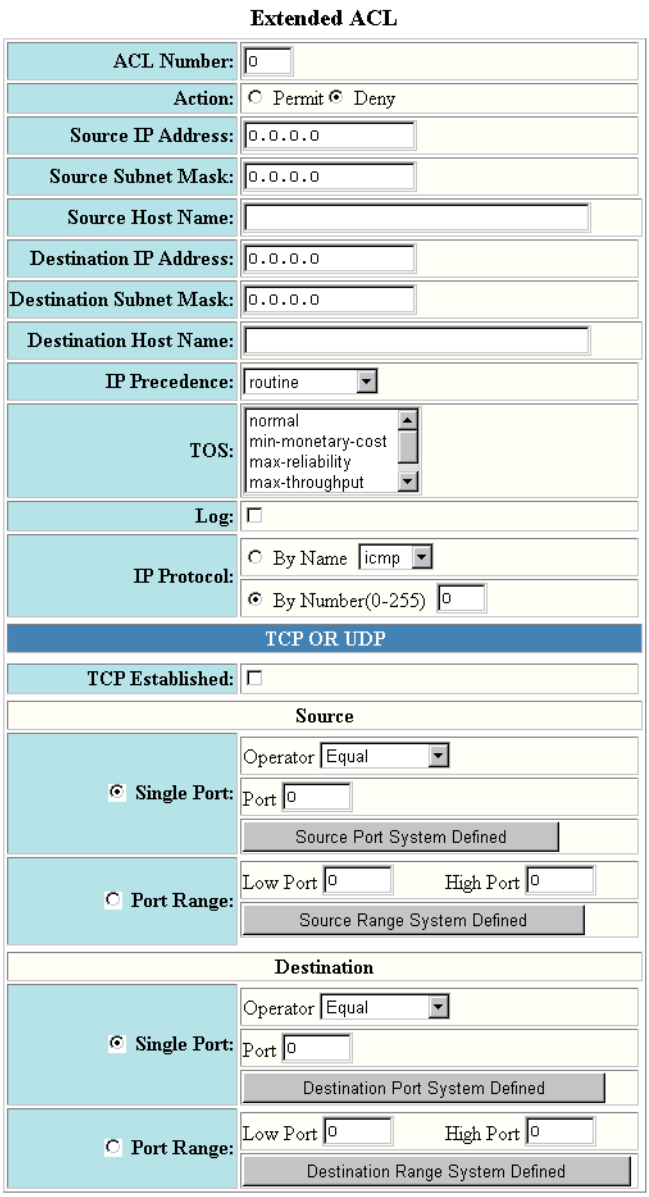

5. Change the ACL number in the ACL Number field or use the ACL number displayed in the field.

NOTE: You cannot specify an ACL name.

6. Select the ACL action. You can select Permit or Deny:
   - Permit – Forwards traffic that matches the ACL.
   - Deny – Drops traffic that matches the ACL.
7. Enter the source IP information. You can enter the source IP address and network mask or the host name.
   - If you enter the address, you also must enter the network mask. To specify “all”, enter “0.0.0.0”.

- If you enter a host name instead of an IP address, when you click Add to add the ACL, the Web management interface sends a DNS query for the address. For the query to be successful, the device must have network access to a DNS server and the server must have an Address record for the host. In addition, the device must be configured with a DNS domain name and the IP address of the DNS server.

8. Enter the destination IP information. The options and requirements are the same as those for entering the source IP information.
9. Select the IP precedence from the IP Precedence pulldown menu (optional). The precedence option for of an IP packet is set in a three-bit field following the four-bit header-length field of the packet's header. You can select one of the following:
   - routine – The ACL matches packets that have the routine precedence.
   - priority – The ACL matches packets that have the priority precedence.
   - immediate – The ACL matches packets that have the immediate precedence.
   - flash – The ACL matches packets that have the flash precedence.
   - flash-override – The ACL matches packets that have the flash override precedence.
   - critical – The ACL matches packets that have the critical precedence.
   - internet – The ACL matches packets that have the internetwork control precedence.
   - network – The ACL matches packets that have the network control precedence.
   - none – The ACL does not use the IP precedence as part of the comparison when filtering.
10. Select the Type of Service (ToS) from the ToS menu (optional). You can select one or more of the following:
    - normal – The ACL matches packets that have the normal ToS.
    - min-monetary-cost or – The ACL matches packets that have the minimum monetary cost ToS.
    - max-reliability – The ACL matches packets that have the maximum reliability ToS.
    - max-throughput – The ACL matches packets that have the maximum throughput ToS.
    - min-delay – The ACL matches packets that have the minimum delay ToS.

    NOTE: To select more than one ToS option, hold the CTRL key while selecting each option.

11. If you specified the Deny action, optionally enable logging by selecting the Log checkbox. If you enable logging for this ACL entry, the software generates Syslog entries for traffic that the ACL denies.
12. Specify the IP protocol. You can specify the protocol by name or by number.
    - To specify the IP protocol by name, select the By Name radio button, then select the protocol from the pulldown menu. You can select one of the following: icmp, igmp, igrp, ip, ospf, tcp, udp.
    - To specify the IP protocol by number, select the By Number radio button, then enter the decimal number of the protocol.
13. If you specified "tcp" or "udp" for the IP protocol, use the following steps to configure the source and destination TCP or UDP options. Otherwise, go to Step 18.
14. Select the Established checkbox if you selected the TCP protocol and you want the ACL to apply to established TCP sessions after you apply the ACL to an interface. Specifically, if you select this option, the ACL 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. If you do not select this option, the ACL applies only to sessions that begin after you apply the ACL to an interface.
15. Select the comparison operator for the source TCP or UDP port. You can select one of the following:
    - Equal – The ACL applies to the TCP or UDP port you specify in the next step.
    - NotEqual – The ACL applies to all TCP or UDP ports ***except*** the port you specify in the next step.

- LessThan – The ACL applies to TCP or UDP ports whose numbers are less than the number of the port you specify in the following step.
- GreaterThan – The ACL applies to TCP or UDP ports whose numbers are greater than the number of the port you specify in the following step.

NOTE: The comparison operators apply only when you are filtering on individual source and destination TCP or UDP ports. If you are filtering on a range of ports, the operators do not apply. Instead, the ACL matches on any TCP or UDP port that is equal to a port within the specified range.

---

16. Specify the source TCP or UDP port. You can specify a single port or a range of ports.
    - To specify a single port, select the radio button next to Single Port. Enter the port number in the Port field. Alternatively, you can select a well-known port name. To do so, select the Source Port System Defined button to change the port number entry field into a pulldown menu containing well-known port names. Select the port from the pulldown menu.
    - To specify a port range, select the radio button next to Port Range. Enter the low port number in the range in the Low Port field and the high port number in the HighPort field. Alternatively, select the Source Range System Defined button to change the entry fields into pulldown menus containing well-known names. Even if you specify the ports by name, you still must select the lower-numbered port first, then select the higher-numbered port.
17. Specify the destination TCP or UDP port. You can specify a single port or a range of ports. The procedures and requirements are the same as those for selecting the source TCP or UDP port. See the previous step.
18. Select the IP Access Group link from the tree view.
    - If the device does not already have some ACLs applied to interfaces, the IP Access Group configuration panel is displayed, as shown in the following example.
    - Otherwise, if the device already has some ACLs applied to interfaces, the IP Access Group table is displayed. Select the Add link to display the IP Access Group configuration panel, as shown in the following example.

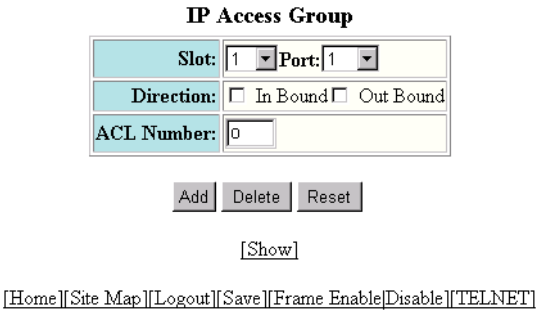

19. Select the Slot and port from the Slot and Port pulldown menus.
20. Specify the traffic direction to which the ACL applies. You can select one or both of the following:
    - In Bound – The ACL applies to traffic received on the port from other devices.
    - Out Bound – The ACL applies to traffic this HP device queues for transmission on the port.
21. Enter the ACL number in the ACL Number field.

NOTE: You cannot specify a named ACL.

---

22. Click the Add button to save the ACL and the association of the ACL with an interface to the device's running-config file.

23. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on Save to Flash.

Configuring Named ACLs

When you configure an IP ACL, you can refer to the ACL by a numeric ID or by a name.

- If you refer to the ACL by a numeric ID, you can use 1 – 99 for a standard ACL or 100 – 199 for an extended ACL.
- If you refer to the ACL by a name, you specify whether the ACL is a standard ACL or an extended ACL, then specify the name.

You can configure up to 100 named standard IP ACLs and 100 named extended IP ACLs. You also can configure up to 100 standard ACLs and 100 extended ACLs by number. Regardless of how many ACLs you have, the device can have a maximum of 1024 ACL entries, associated with the ACLs in any combination. (On HP 9300 series Chassis devices with Management 2 modules, the maximum is 2048.)

To configure a named IP ACL, use the following CLI method.

USING THE CLI

The commands for configuring named ACL entries are different from the commands for configuring numbered ACL entries. The command to configure a numbered ACL is **access-list**. The command for configuring a named ACL is **ip access-list**. In addition, when you configure a numbered ACL entry, you specify all the command parameters on the same command. When you configure a named ACL, you specify the ACL type (standard or extended) and the ACL number with one command, which places you in the configuration level for that ACL. Once you enter the configuration level for the ACL, the command syntax is the same as the syntax for numbered ACLs.

The following examples show how to configure a named standard ACL entry and a named extended ACL entry.

Configuration Example for Standard ACL

To configure a named standard ACL entry, enter commands such as the following.

```
HP9300(config)# ip access-list standard Net1
HP9300(config-std-nacl)# deny host 209.157.22.26 log
HP9300(config-std-nacl)# deny 209.157.29.12 log
HP9300(config-std-nacl)# deny host IPHost1 log
HP9300(config-std-nacl)# permit any
HP9300(config-std-nacl)# exit
HP9300(config)# int eth 1/1
HP9300(config-if-1/1)# ip access-group Net1 out
```

The commands in this example configure a standard ACL named "Net1". The entries in this ACL deny packets from three source IP addresses from being forwarded on port 1/1. Since the implicit action for an ACL is "deny", the last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries. For an example of how to configure the same entries in a numbered ACL, see "Configuring Standard ACLs" on page 4-9.

Notice that the command prompt changes after you enter the ACL type and name. The "std" in the command prompt indicates that you are configuring entries for a standard ACL. For an extended ACL, this part of the command prompt is "ext". The "nacl" indicates that are configuring a named ACL.

Syntax: ip access-list extended | standard <string> | <num>

The **extended | standard** parameter indicates the ACL type.

The <string> parameter is the ACL name. You can specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, “ACL for Net1”). The <num> parameter allows you to specify an ACL number if you prefer. If you specify a number, you can specify from 1 – 99 for standard ACLs or 100 – 199 for extended ACLs.

NOTE: For convenience, the software allows you to configure numbered ACLs using the syntax for named ACLs. The software also still supports the older syntax for numbered ACLs. Although the software allows both methods for configuring numbered ACLs, numbered ACLs are always formatted in the startup-config and running-config files in using the older syntax, as follows.

```
access-list 1 deny host 209.157.22.26 log
access-list 1 deny 209.157.22.0 0.0.0.255 log
access-list 1 permit any
access-list 101 deny tcp any any eq http log
```

The options at the ACL configuration level and the syntax for the **ip access-group** command are the same for numbered and named ACLs and are described in “Configuring Standard ACLs” on page 4-9.

Configuration Example for Extended ACL

To configure a named extended ACL entry, enter commands such as the following.

```
HP9300(config)# ip access-list extended “block Telnet”
HP9300(config-ext-nacl)# deny tcp host 209.157.22.26 any eq telnet log
HP9300(config-ext-nacl)# permit ip any any
HP9300(config-ext-nacl)# exit
HP9300(config)# int eth 1/1
HP9300(config-if-1/1)# ip access-group “block Telnet” in
```

The options at the ACL configuration level and the syntax for the **ip access-group** command are the same for numbered and named ACLs and are described in “Configuring Extended ACLs” on page 4-12.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure IP ACLs using the Web management interface.

Adding a Comment to an ACL Entry

You can optionally add comment text to describe entries in an ACL. The comment text appears in the output of **show** commands that display ACL information.

For example, the following commands add comments to entries to a numbered ACL, ACL 100:

```
HP9300(config)# access-list 100 remark The following line permits TCP packets
HP9300(config)# access-list 100 permit tcp 192.168.4.40/24 2.2.2.2/24

HP9300(config)# access-list 100 remark The following permits UDP packets
HP9300(config)# access-list 100 permit udp 192.168.2.52/24 2.2.2.2/24

HP9300(config)# access-list 100 deny ip any any
```

If the ACL is a named ACL, (for example, you entered TCP/UDP instead of 100), enter the following commands:

```
HP9300(config)# access-list TCP/UDP remark The following line permits TCP packets
HP9300(config)# access-list TCP/UDP permit tcp 192.168.4.40/24 2.2.2.2/24
HP9300(config)# access-list TCP/UDP remark The following permits UDP packets
HP9300(config)# access-list TCP/UDP permit udp 192.168.2.52/24 2.2.2.2/24
HP9300(config)# access-list TCP/UDP deny ip any any
```

Syntax: [no] access-list <acl-num> | <acl-name> remark <comment-text>

Enter the number of the ACL for <acl-num>. Beginning with software release 07.6.04, you can add a comment to a named ACL by entering the ACL's name for <acl-name>.

The <comment-text> can be up to 128 characters in length. The comment must be entered separately from the actual ACL entry; that is, you cannot enter the ACL entry and the ACL comment with the same **access-list** command. Also, in order for the remark to be displayed correctly in the output of **show** commands, the comment must be entered immediately before the ACL entry it describes.

You can use the **show running-config** or **show access-list** commands to display the ACL and comments

The following shows an example of a numbered ACL with a comment text in a show running-config display:

```
HP9300# show running-config
…
access-list 100 remark The following line permits TCP packets
access-list 100 permit tcp 192.168.4.40/24 2.2.2.2/24
access-list 100 remark The following line permits UDP packets
access-list 100 permit udp 192.168.2.52/24 2.2.2.2/24
access-list 100 deny ip any any
```

The following shows the comment text for the ACL named TCP/UDP in a show running-config display:

```
HP9300# show running-config ...

access-list TCP/UDP remark The following line permits TCP packets
access-list TCP/UDP permit tcp 192.168.4.40/24 2.2.2.2/24
access-list TCP/UDP remark The following line permits UDP packets
access-list TCP/UDP permit udp 192.168.2.52/24 2.2.2.2/24
access-list TCP/UDP deny ip any any
```

Syntax: show running-config

The following example show the comment text for a numbered ACL in a show access-list display:

```
HP9300# show access-list 100
IP access list rate-limit 100 aaaa.bbbb.cccc

Extended IP access list 100 (Total flows: N/A, Total packets: N/A)
  ACL Comments:  The following line permits TCP packets
  permit tcp 0.0.0.40 255.255.255.0 0.0.0.2 255.255.255.0 (Flows: N/A, Packets: N/A)
  ACL Comments:  The following line permits UDP packets
  permit udp 0.0.0.52 255.255.255.0 0.0.0.2 255.255.255.0 (Flows: N/A, Packets: N/A)
  deny ip any any (Flows: N/A, Packets: N/A)
```

The next example shows the comment text for a named ACL in a show access-list display:

```
HP9300# show access-list TCP/UDP

IP access list rate-limit 100 aaaa.bbbb.cccc

Extended IP access list TCP/UDP (Total flows: N/A, Total packets: N/A)
ACL Comments:  The following line permits TCP packets
permit tcp 0.0.0.40 255.255.255.0 0.0.0.2 255.255.255.0 (Flows: N/A, Packets: N/A)
ACL Comments:  The following line permits UDP packets
permit udp 0.0.0.52 255.255.255.0 0.0.0.2 255.255.255.0 (Flows: N/A, Packets: N/A)
deny ip any any (Flows: N/A, Packets: N/A)
```

Syntax: show access-list <acl-num> | <acl-name> | all

Reapplying ACLs to Interfaces

If you change an ACL, or enable or disable TCP or UDP strict mode, HP recommends that you enter the following command to rebind the ACL configuration:

```
HP9300(config)# ip rebind-acl all
```

Syntax: [no] ip rebind-acl <num> | <name> | all

This command reapplies all ACLs to their interfaces.

To reapply a specific ACL, enter a command such as the following:

```
HP9300(config)# ip rebind-acl 101
```

This command reapplies ACL 101 only.

Modifying ACLs

NOTE: This section applies to standard ACLs and to extended ACLs.

When you use the HP device's CLI or Web management interface to configure an ACL, the software places the ACL entries in the ACL in the order you enter them. For example, if you enter the following entries in the order shown below, the software always applies the entries to traffic in the same order.

```
HP9300(config)# access-list 1 deny 209.157.22.0/24

HP9300(config)# access-list 1 permit 209.157.22.26
```

Thus, if a packet matches the first ACL entry in this ACL and is therefore denied, the software does not compare the packet to the remaining ACL entries. In this example, packets from host 209.157.22.26 will always be dropped, even though packets from this host match the second entry.

You can use the CLI to reorder entries within an ACL by individually removing the ACL entries and then re-adding them. To use this method, enter "**no**" followed by the command for an ACL entry, and repeat this for each ACL entry in the ACL you want to edit. After removing all the ACL entries from the ACL, re-add them.

This method works well for small ACLs such as the example above, but can be impractical for ACLs containing many entries. Therefore, HP devices provide an alternative method. The alternative method lets you upload an ACL list from a TFTP server and replace the ACLs in the device's running-config file with the uploaded list. Thus, to change an ACL, you can edit the ACL on the file server, then upload the edited ACL to the device. You then can save the changed ACL to the device's startup-config file.

ACL lists contain only the ACL entries themselves, not the assignments of ACLs to interfaces. You must assign the ACLs on the device itself.

NOTE: The only valid commands that are valid in the ACL list are the **access-list** and **end** commands. The HP device ignores other commands in the file.

To modify an ACL by configuring an ACL list on a file server:

1. Use a text editor to create a new text file. When you name the file, use 8.3 format (up to eight characters in the name and up to three characters in the extension).

   NOTE: Make sure the HP device has network access to the TFTP server.

2. Optionally, clear the ACL entries from the ACLs you are changing by placing commands such as the following at the top of the file:

   ```
   no access-list 1
   no access-list 101
   ```

   When you load the ACL list into the device, the software adds the ACL entries in the file after any entries that already exist in the same ACLs. Thus, if you intend to entirely replace an ACL, you must use the **no access-list** <num> command to clear the entries from the ACL before the new ones are added.

3. Place the commands to create the ACL entries into the file. The order of the separate ACLs does not matter, but the order of the entries within each ACL is important. The software applies the entries in an ACL in the order they are listed within the ACL. Here is an example of some ACL entries:

   ```
   access-list 1 deny host 209.157.22.26 log
   access-list 1 deny 209.157.22.0 0.0.0.255 log
   access-list 1 permit any
   access-list 101 deny tcp any any eq http log
   ```

   The software will apply the entries in ACL 1 in the order shown and stop at the first match. Thus, if a packet is denied by one of the first three entries, the packet will not be permitted by the fourth entry, even if the packet matches the comparison values in this entry.

4. Enter the command "**end**" on a separate line at the end of the file. This command indicates to the software that the entire ACL list has been read from the file.

5. Save the text file.

6. On the HP device, enter the following command at the Privileged EXEC level of the CLI:

   copy tftp running-config <tftp-ip-addr> <filename>

   NOTE: This command will be unsuccessful if you place any commands other than **access-list** and **end** (at the end only) in the file. These are the only commands that are valid in a file you load using the **copy tftp running-config…** command.

7. To save the changes to the device's startup-config file, enter the following command at the Privileged EXEC level of the CLI:

   write memory

Here is a complete example of an ACL configuration file.

```
no access-list 1
no access-list 101
access-list 1 deny host 209.157.22.26 log
access-list 1 deny 209.157.22.0 0.0.0.255 log
access-list 1 permit any
access-list 101 deny tcp any any eq http log
end
```

NOTE: Do not place other commands in the file. The HP device reads only the ACL information in the file and ignores other commands, including **ip access-group** commands. To assign ACLs to interfaces, use the CLI.

Dropping All Fragments That Exactly Match an ACL

For a packet fragment that is sent to the CPU for processing, the device compares the fragment's source and destination IP addresses against the interface's ACL entries. By default, if the fragment's source and destination IP addresses exactly match an ACL entry that also has Layer 4 information (source and destination TCP or UDP application ports), the device permits or denies the fragment according to the ACL.

On an individual interface basis, you can configure a Standard (non-EP) device to automatically drop a fragment whose source and destination IP addresses exactly match an ACL entry that has Layer 4 information, even if that ACL entry's action is permit. To do so, enter the following command at the configuration level for an interface:

```
HP9300(config-if-1/1)# ip access-group frag deny
```

Syntax: [no] ip access-group frag deny

NOTE: EP devices also support the **ip access-group frag deny** command but the command performs a different service on EP devices. See "Enabling ACL Filtering of Fragmented Packets" on page 5-8.

Enabling Hardware Filtering for Packets Denied by ACLs

By default, packets denied by ACLs are filtered by the CPU. You can enable the device to create Content Addressable Memory (CAM) entries for packets denied by ACLs. This causes the filtering to occur in hardware instead of in the CPU.

When you enable hardware filtering of denied packets, the first time the device filters a packet denied by an ACL, the device sends the packet to the CPU for processing. The CPU also creates a CAM entry for the denied packet. Subsequent packets with the same address information are filtered using the CAM entry. The CAM entry ages out after two minutes if not used.

To enable hardware filtering of denied packets, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# hw-drop-acl-denied-packet
```

Syntax: [no] hw-drop-acl-denied-packet

Enabling ICMP Unreachable Messages for Traffic Denied by ACLs

By default, an HP device does not send a message to another device when an ACL on the HP device denies a packet from the other device. You can enable a Routing Switch to send an ICMP unreachable message to a device when an ACL denies a packet from the device.

To enable the ICMP unreachable messages, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# acl-denied-icmp-msg
```

Syntax: [no] acl-denied-icmp-msg

The command applies globally to all ACLs configured on the device.

NOTE: This command applies only to Routing Switches.

NOTE: This command does not take effect in the following cases:

—Hardware-based ACLs are enabled.

—The **hw-drop-acl-denied-packet** command is in effect.

In either case, all packets denied by the ACL are dropped by hardware without sending an ICMP message.

ACL Filtering for Traffic Switched Within a Virtual Routing Interface

By default, an HP device does not filter traffic that is switched from one port to another within the same virtual routing interface, even if an ACL is applied to the interface. You can enable the device to filter switched traffic within a virtual routing interface. When you enable the filtering, the device uses the ACLs applied to inbound traffic to filter traffic received by a port from another port in the same virtual routing interface. This feature does not apply to ACLs applied to outbound traffic.

To enable filtering of traffic switched within a virtual routing interface, enter the following command at the configuration level for the interface:

```
HP9300(config-vif-1)# ip access-group ve-traffic
```

Syntax: [no] ip access-group ve-traffic

Applying an ACL to a Subset of Ports on a Virtual Interface

You can apply an ACL to a virtual routing interface. The virtual interface is used for routing between VLANs and contains all the ports within the VLAN. If the ACL is for the inbound traffic direction, you also can specify a subset of ports within the VLAN containing a specified virtual interface when assigning an ACL to that virtual interface.

Use this feature when you do not want the ACLs to apply to all the ports in the virtual interface's VLAN or when you want to streamline ACL performance for the VLAN.

NOTE: This feature applies only to a virtual interface's inbound direction. You cannot use this feature to specify a subset of ports for a virtual interface's outbound direction.

To apply an ACL to a subset of ports within a virtual interface, enter commands such as the following:

```
HP9300(config)# vlan 10 name IP-subnet-vlan
HP9300(config-vlan-10)# untag ethernet 1/1 to 2/12
HP9300(config-vlan-10)# router-interface ve 1
HP9300(config-vlan-10)# exit
HP9300(config)# access-list 1 deny host 209.157.22.26 log
HP9300(config)# access-list 1 deny 209.157.29.12 log
HP9300(config)# access-list 1 deny host IPHost1 log
HP9300(config)# access-list 1 permit any
HP9300(config)# interface ve 1
HP9300(config-vif-1)# ip access-group 1 in ethernet 1/1 ethernet 1/3 ethernet 2/1 to
2/4
```

The commands in this example configure port-based VLAN 10, add ports 1/1 – 2/12 to the VLAN, and add virtual routing interface 1 to the VLAN. The commands following the VLAN configuration commands configure ACL 1. Finally, the last two commands apply ACL 1 to a subset of the ports associated with virtual interface 1.

Syntax: [no] ip access-group <num> in ethernet <portnum> [<portnum>...] to <portnum>

Enabling Strict TCP or UDP Mode

By default, when you use ACLs to filter TCP or UDP traffic, the HP device does not compare all TCP or UDP packets against the ACLs.

For TCP and UDP, the device first compares the source and destination information in a TCP control packet or a UDP packet against entries in the session table. The session table contains forwarding entries based on Layer 3 and Layer 4 information.

- If the session table contains a matching entry, the device forwards the packet, assuming that the first packet the device received with the same address information was permitted by the ACLs.
- If the session table does not contain a matching entry, the device sends the packet to the CPU, where the software compares the packet against the ACLs. If the ACLs permit the packet (explicitly by a permit ACL entry or implicitly by the absence of a deny ACL entry), the CPU creates a session table entry for the packet's forwarding information and forwards the packet.

For TCP, this behavior by default applies only to control packets, not to data packets. Control packets include packet types such as SYN (Synchronization) packets, FIN (Finish) packets, and RST (Reset) packets.

For tighter access or forwarding control, you can enable the device to perform strict TCP or UDP ACL processing. The following sections describe the strict modes in more detail.

Enabling Strict TCP Mode

By default, when you use ACLs to filter TCP traffic, the HP device does not compare all TCP packets against the ACLs. Instead, the device compares TCP control packets against the ACLs, but not data packets. Control packets include packet types such as SYN (Synchronization) packets, FIN (Finish) packets, and RST (Reset) packets.

In normal TCP operation, TCP data packets are present only if a TCP control session for the packets also is established. For example, data packets for a session never occur if the TCP SYN for that session is dropped. Therefore, by filtering the control packets, the HP device also implicitly filters the data packets associated with the control packets. This mode of filtering optimizes forwarding performance for TCP traffic by forwarding data packets without examining them. Since the data packets are present in normal TCP traffic only if a corresponding TCP control session is established, comparing the packets for the control session to the ACLs is sufficient for filtering the entire session including the data.

However, it is possible to generate TCP data packets without corresponding control packets, in test or research situations for example. In this case, the default ACL mode does not filter the data packets, since there is no corresponding control session to filter. To filter this type of TCP traffic, use the strict ACL TCP mode. This mode compares all TCP packets to the configured ACLs, regardless of whether the packets are control packets or data packets. If the ACLs permit the packet, the device creates a session entry for forwarding other TCP packets with the same Layer 3 and Layer 4 addresses.

NOTE: Regardless of whether the strict mode is enabled or disabled, the device always compares TCP control packets against the configured ACLs before creating a session entry for forwarding the traffic.

NOTE: If the device's configuration currently has ACLs associated with interfaces, remove the ACLs from the interfaces before changing the ACL mode.

To enable the strict ACL TCP mode, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# ip strict-acl-tcp
```

Syntax: [no] ip strict-acl-tcp

This command configures the device to compare all TCP packets against the configured ACLs before forwarding them.

To disable the strict ACL mode and return to the default ACL behavior, enter the following command:

```
HP9300(config)# no ip strict-acl-tcp
```

NOTE: If you are using software release 07.6.04 or later, enter the **ip rebind-acl** command at the global CONFIG level of the CLI to place the **ip strict-acl-tcp** or **no ip strict-acl-tcp** command into effect.

Enabling Strict UDP Mode

By default, when you use ACLs to filter UDP traffic, the HP device does not compare all UDP packets against the ACLs. Instead, the device compares the source and destination information against entries in the session table. The session table contains forwarding entries based on Layer 3 and Layer 4 information.

- If the session table contains a matching entry, the device forwards the packet, assuming that the first packet the device received that contains the same address information was permitted by the ACLs.
- If the session table does not contain a matching entry, the device sends the packet to the CPU, where the software compares the packet against the ACLs. If the ACLs permit the packet (explicitly by a permit ACL entry or implicitly by the absence of a deny ACL entry), the CPU creates a session table entry for the packet's forwarding information and forwards the packet.

For tighter control, the software provides the strict ACL UDP mode. When you enable strict UDP processing, the device sends every UDP packet to the CPU and compares the packet against the configured ACLs.

NOTE: If the device's configuration currently has ACLs associated with interfaces, remove the ACLs from the interfaces before changing the ACL mode.

To enable the strict ACL UDP mode, enter the following command at the global CONFIG level of the CLI:

```
HP9300(config)# ip strict-acl-udp
```

Syntax: [no] ip strict-acl-udp

This command configures the device to compare all UDP packets against the configured ACLs before forwarding them.

To disable the strict ACL mode and return to the default ACL behavior, enter the following command:

```
HP9300(config)# no ip strict-acl-udp
```

NOTE: If you are using software release 07.6.04 or later, enter the **ip rebind-acl** command at the global CONFIG level of the CLI to place the **ip strict-acl-udp** or **no ip strict-acl-udp** command into effect.

Configuring ACL Packet and Flow Counters

You can configure counters for packets and flows that match entries in an ACL. Using the CLI, you can display the contents of the counters and clear them.

- The ACL packet counter feature provides an accurate count of packets matching individual ACL entries.
- The ACL flow counter feature provides an approximate count of flows matching individual ACL entries. This feature can be used for troubleshooting purposes to provide an indication of flow activity against an ACL. Each time the HP device receives the first packet of a flow matching an entry in an ACL list, the flow counter for that ACL entry is incremented by one. If a flow lasts longer than two minutes, the flow counter for the ACL entry is incremented again.

  NOTE: The ACL flow counter feature is designed to monitor the general volume of flow activity for an ACL. It is not intended to be used for accounting purposes.

The ACL flow and packet counters are incremented differently depending on whether they are permit or deny flows.

- For permit flows, only flows are counted. If a permit flow lasts longer than two minutes, the flow counter is incremented again.
- For deny flows, only packets are counted.

By default the ACL packet and flow counters are disabled. To activate them, enter the following command:

```
HP9300(config)# enable-acl-counter
```

Syntax: [no] enable-acl-counter

Once the ACL packet and flow counters are enabled, you can disable them with the **no** form of the **enable-acl-counter** command. Disabling and then re-enabling the ACL packet and flow counters resets them to zero.

To display the packet and flow counters for ACL 100:

```
HP9300# show access-list 100
Extended IP access list 100 (Total flows: 432, Total packets: 42000)
  permit tcp 1.1.1.0 0.0.0.255 any (Flows: 80, Packets: 12900)
  deny udp 1.1.1.0 0.0.0.255 any (Flows: 121, Packets: 20100)
  permit ip 2.2.2.0 0.0.0.255 any (Flows: 231, Packets: 9000)
```

Syntax: show access-list <acl-num> | <acl-name> | all

To clear the flow counters for ACL 100:

```
HP9300# clear access-list 100
```

Syntax: clear access-list <acl-num> | <acl-name> | all

Using ACLs to Filter ARP Packets

Starting with software release 07.6.034, you can use ACLs to filter ARP packets. Without this feature, ACLs cannot be used to permit or deny incoming or outgoing ARP packets. (Although an ARP packet contains an IP address just as an IP packet does, it is not an IP packets and is not subject to the normal filtering provided by ACLs.)

When an HP device receives an ARP request, the source MAC and IP addresses are stored in the device's ARP table. A new record in the ARP table overwrites existing records that contain the same IP address. This behavior can cause a condition called "ARP hijacking", when two hosts with the same IP address try to send an ARP request to the HP device.

Normally ARP hijacking is not a problem because IP assignments are done dynamically; however, in some cases, such as when the **ip follow** command is used, ARP hijacking can occur.

The **ip follow** command allows a router interface to share the IP address of another router interface. **Ip follow** conserves IP addresses, while separating Layer 2 traffic from different sources by port-based VLAN. Since multiple VLANs and the router interfaces that are associated with each of the VLANs share the same IP segment, it is possible for two hosts in two different VLANs to fight for the same IP address in that segment. ARP filtering using ACLs protects an IP host's record in ARP table from being overwritten by a hijacking host.

Using ACLs to filter ARP request checks the source IP address in the received ARP packet. Only packets with the permitted IP address will be allowed to be to be written in the ARP table; others are dropped.

Configuration Considerations:

- This feature is available on all devices running Layer 3 code. On a T-Flow module, this filtering occurs on the management processor.
- The feature is available on physical interfaces and virtual routing interfaces. It is supported on the following physical interface types: Ethernet and trunks.
- ACLs used to filter ARP packets a virtual routing interface can be inherited from a previous interface if the virtual routing interface is defined as a follower virtual routing interface

Configuring ACLs for ARP Filtering

To implement the ACL ARP filtering feature, enter commands such as the following:

```
HP9300(config)# access-list 101 permit ip host 192.168.2.2 any
HP9300(config)# access-list 102 permit ip host 192.168.2.3 any
HP9300(config)# access-list 103 permit ip host 192.168.2.4 any
```

```
HP9300(config)# vlan 2
HP9300(config-vlan-2)# tag ethe 1/1 to 1/2
HP9300(config-vlan-2)# router-interface ve 2
HP9300(config-vlan-2)# vlan 3
HP9300(config-vlan-3)# tag ethe 1/1 to 1/2
HP9300(config-vlan-3)#router-int ve 3
HP9300(config-vlan-3)# vlan 4
HP9300(config-vlan-4)# tag ethe 1/1 to 1/2
HP9300(config-vlan-4)# router-int ve 4
HP9300(config-vlan-4)# interface ve 2
HP9300(config-ve-2)# ip access-group 101 in
HP9300(config-ve-2)# ip address 192.168.2.1/24
HP9300(config-ve-2)# ip use-acl-on-arp 103
HP9300(config-ve-2)# exit

HP9300(config)# interface ve 3
HP9300(config-ve-3)# ip access-group 102 in
HP9300(config-ve-3)# ip follow ve 2
HP9300(config-ve-3)# no ip follow acl
HP9300(config-ve-3)# ip use-acl-on-arp
HP9300(config-ve-3)# exit

HP9300(config-vlan-4)# interface ve 4
HP9300(config-ve-4)# ip follow ve 2
HP9300(config-ve-4)# ip use-acl-on-arp
HP9300(config-ve-4)# exit
```

Syntax: [no] ip use-acl-on-arp [<access-list-number>]

When the **use-acl-on-arp** command is configured, the ARP module checks the source IP address of the ARP request packets received on the interface. It then applies the specified ACL policies to the packet. Only the packet with the IP address that the ACL permits will be allowed to be to be written in the ARP table; those that are not permitted will be dropped.

The <access-list-number> parameter identifies the ID of the standard ACL that will be used to filter the packet. Only the source and destination IP addresses will be used to filter ARP packet. You can do one of the following for <access-list-number>:

- Enter an ACL ID to explicitly specify the ACL to be used for filtering. In the example above, the line `HP9300(config-ve-2)# ip use-acl-on-arp 103` specifies ACL 103 to be used as the filter.
- Allow the ACL ID to be inherited from the IP ACLs that have been defined for the device. In the example above, the line `HP9300(config-ve-3)# ip use-acl-on-arp` does not define an ACL, but allows the ACL to be inherited from the IP ACL 102. Also in the example, the line `HP9300(config-ve-4)# ip use-acl-on-arp` allows the ACL to be inherited from IP ACL 101 because of the ip follow relationship between virtual routing interface 2 and virtual routing interface 4. Virtual routing interface 2 is configured with IP IP ACL 101; thus virtual routing interface 4 inherits IP ACL 101.

ARP requests will not be filtered by ACLs if one of the following conditions occur:

- If the ACL is to be inherited from an IP ACL, but there is no IP ACL defined.
- An ACL ID is specified for the **use-acl-on-arp** command, but no IP address or “any any” filtering criteria have been defined under the ACL ID.

Displaying ACLs

To display the ACLs configured on a device, use the following method.

USING THE CLI

To display ACLs, enter the **show ip access-lists** command. Here is an example:

```
HP9300(config)# show ip access-lists
```

```
Extended IP access list 101
    deny tcp host 209.157.22.26 host 209.157.22.26 eq http log
```

Syntax: show ip access-lists [<num>]

Displaying ACL Log Entries

The first time an entry in an ACL permits or denies a packet and logging is enabled for that entry, the software generates a Syslog message and an SNMP trap. Messages for packets permitted or denied by ACLs are at the warning level of the Syslog.

When the first Syslog entry for a packet permitted or denied by an ACL is generated, the software starts an ACL timer. After this, the software sends Syslog messages every one to ten minutes, depending on the value of the timer interval. If an ACL entry does not permit or deny any packets during the timer interval, the software does not generate a Syslog entry for that ACL entry. For more information about the timer, see “Configuring the Layer 4 Session Log Timer” on page 4-34.

NOTE: For an ACL entry to be eligible to generate a Syslog entry for permitted or denied packets, logging must be enabled for the entry. The Syslog contains entries only for the ACL entries that deny packets and have logging enabled.

To display Syslog entries, use one of the following methods.

USING THE CLI

Enter the following command from any CLI prompt:

```
HP9300(config)# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 38 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Log Buffer (50 entries):

21d07h02m40s:warning:list 101 denied tcp 209.157.22.191(0)(Ethernet 4/18
0010.5a1f.77ed) -> 198.99.4.69(http), 1 event(s)
00d07h03m30s:warning:list 101 denied tcp 209.157.22.26(0)(Ethernet 4/18
0010.5a1f.77ed) -> 198.99.4.69(http), 1 event(s)
00d06h58m30s:warning:list 101 denied tcp 209.157.22.198(0)(Ethernet 4/18
0010.5a1f.77ed) -> 198.99.4.69(http), 1 event(s)
```

In this example, the two-line message at the bottom is the first entry, which the software immediately generates the first time an ACL entry permits or denies a packet. In this case, an entry in ACL 101 denied a packet. The packet was a TCP packet from host 209.157.22.198 and was destined for TCP port 80 (HTTP) on host 198.99.4.69.

When the software places the first entry in the log, the software also starts the five-minute timer for subsequent log entries. Thus, five minutes after the first log entry, the software generates another log entry and SNMP trap for denied packets.

In this example, the software generates the second log entry five minutes later.

The time stamp for the third entry is much later than the time stamps for the first two entries. In this case, no ACLs denied packets for a very long time. In fact, since no ACLs denied packets during the five-minute interval following the second entry, the software stopped the ACL log timer. The software generated the third entry as soon as the ACL denied a packet. The software restarted the five-minute ACL log timer at the same time. As long as at least one ACL entry permits or denies a packet, the timer continues to generate new log entries and SNMP traps every five minutes.

You can also configure the maximum number of ACL-related log entries that can be added to the system log over a one-minute period. For example, to limit the device to 100 ACL-related syslog entries per minute:

```
HP9300(config)# max-acl-log-num 100
```

Syntax: [no] max-acl-log-num <num>

You can specify a number between 0 – 4096. The default is 256. Specifying 0 disables all ACL logging.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the Routing Switch using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of configuration options.
3. Select the System Log link.

Configuring the Layer 4 Session Log Timer

In releases 07.6.04 and later, you can configure the Layer 4 session log timer, which is used for keeping track of packets explicitly denied by an ACL.

When you enable logging for an ACL entry, statistics for packets that match the permit or deny conditions of the ACL entry are logged in the HP device's Syslog buffer and in SNMP traps sent by the device. The first time an ACL entry permits or denies a packet, the software immediately generates a Syslog entry and SNMP trap. The software also starts the Layer 4 session log timer. The timer keeps track of all packets explicitly denied by the ACL entries. When the timer expires, the software generates a single Syslog entry for each ACL entry that has denied a packet. The message indicates the number of packets denied by the ACL entry from the time that the timer was started. If no ACL entries explicitly permit or deny packets during an entire timer interval, the timer stops. The timer restarts when an ACL entry explicitly permits or denies a packet.

To store information about denied packets during the timer interval, the device makes entries in its Layer 4 session table. If a large number of packets are denied by the ACL during the timer interval, it can consume a large portion of the device's Layer 4 resources. To prevent this from happening, starting in release 07.6.04, you can configure the timer interval to be a shorter length of time. In releases prior to 07.6.04, the timer interval was set to 5 minutes and was not configurable.

For example, to set the timer interval to 2 minutes, enter the following command:

```
HP9300(config)# ip access-list logging-age 2
```

Syntax: ip access-list logging-age <minutes>

You can set the timer to between 1 and 10 minutes. The default is 5 minutes.

Displaying and Clearing Flow-Based ACL Statistics

You can display statistics for packets permitted or denied by Standard flow-based (CPU-based) ACLs.

Displaying ACL Statistics for Flow-Based ACLs

To display ACL statistics for flow-based ACLs, enter the following command:

```
HP9300(config)# show ip acl-traffic

ICMP inbound packets received 400
ICMP inbound packets permitted 200
ICMP inbound packets denied 200
```

Syntax: show ip acl-traffic

The command lists a separate set of statistics for each of the following IP protocols:

- ICMP
- IGMP

- IGRP
- IP
- OSPF
- TCP
- UDP
- Protocol number, if an ACL is configured for a protocol not listed above

For TCP and UDP, a separate set of statistics is listed for each application port.

Clearing Flow-Based ACL Statistics

To clear the ACL statistics, enter the following command at the Privileged EXEC level of the CLI:

```
HP9300(config)# clear ip acl-traffic
```

Syntax: clear ip acl-traffic

Displaying and Clearing ACL Filters for ARP

Displaying ACL Filters for ARP

To determine what ACLs have been configured to filter ARP requests, enter a command such as the following:

```
HP9300(config)# show acl-on-arp
Port   ACL ID   Filter Count
2      103      10
3      102      23
4      101      12
```

Syntax: show acl-on-arp [ethernet [<portnum>] | loopback [<num>] | ve [<num>]]

If port number or the interface number is not specified, all ports on the device that use ACLs for ARP filtering will be included in the display.

The Filter Count column shows how many ARP packets have been dropped on the interface since the last time the count was cleared.

Clearing Filter Count

To clear the filter count for all interfaces on the device, enter a command such as the following:

```
HP9300(config)# clear acl-on-arp
```

Syntax: clear acl-on-arp

The command resets the filter count on all interfaces in a device back to zero

Policy-Based Routing (PBR)

NOTE: This section describes the PBR support on Standard devices. For information about PBR on EP devices, see “Hardware-Based Policy-Based Routing (PBR)” on page 5-14.

Policy-Based Routing (PBR) allows you to use ACLs and route maps to selectively modify and route IP packets based on their source IP address.

You can configure the Routing Switch to perform the following types of PBR based on a packet’s Layer 3 and Layer 4 information:

- Select the next-hop gateway. (See “Configuration Examples” on page 4-38 for a complete configuration example.)

- Specify the default next-hop IP address if there is no explicit next-hop selection for the packet.
- Send the packet to the null interface (null0).

HP's PBR routing is based on standard and extended ACLs and route-maps. The ACLs classify the traffic. Route maps that match on the ACLs set routing attributes for the traffic. HP's implementation of PBR uses high performance switching algorithms including route caches and route tables.

NOTE: Standard PBR is supported only on Chassis Routing Switches.

NOTE: Source routing occurs in the CPU, not in the ASICs.

Configuring PBR

To configure PBR:

- Configure ACLs that contain the source IP addresses for the IP traffic to which you want to apply PBR.
- Configure a route map that matches on the ACLs and sets route information.
- Apply the route map globally or to individual interfaces.

NOTE: If you are using software release 07.6.04 or later, enter the **ip rebind-acl** command at the global CONFIG level of the CLI to place ACL configuration changes into effect.

Configure the ACLs

PBR uses route maps to change the routing attributes in IP traffic. This section shows an example of how to configure a standard ACL to identify the source sub-net for IP traffic.

To configure a standard ACL to identify a source sub-net, enter a command such as the following:

```
HP9300(config)# access-list 101 permit 209.157.23.0 0.0.0.255
```

The command in this example configures a standard ACL that permits traffic from sub-net 209.157.23.0/24. After you configure a route map that matches based on this ACL, the software uses the route map to set route attributes for the traffic, thus enforcing PBR.

NOTE: Do not use an access group to apply the ACL to an interface. Instead, use a route map to apply the ACL globally or to individual interfaces for PBR, as shown in the following sections.

Syntax: [no] access-list <num> deny | permit <source-ip> | <hostname> <wildcard> [log]

or

Syntax: [no] access-list <num> deny | permit <source-ip>/<mask-bits> | <hostname> [log]

Syntax: [no] access-list <num> deny | permit host <source-ip> | <hostname> [log]

Syntax: [no] access-list <num> deny | permit any [log]

The <num> parameter is the access list number and can be from 1 – 99.

The **deny | permit** parameter indicates whether packets that match a policy in the access list are denied (dropped) or permitted (forwarded).

NOTE: If you are configuring the ACL for use in a route map, always specify **permit**. Otherwise, the Routing Switch drops the traffic instead of further processing the traffic using the route map.

The <source-ip> parameter specifies the source IP address. Alternatively, you can specify the host name.

NOTE: To specify the host name instead of the IP address, the host name must be configured using the HP device's DNS resolver. To configure the DNS resolver name, use the **ip dns server-address…** command at the global CONFIG level of the CLI.

The <wildcard> parameter specifies the mask value to compare against the host address specified by the <source-ip> parameter. The <wildcard> is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet's source address must match the <source-ip>. Ones mean any value matches. For example, the <source-ip> and <wildcard> values 209.157.22.26 0.0.0.255 mean that all hosts in the Class C sub-net 209.157.22.x match the policy.

If you prefer to specify the wildcard (mask value) in CIDR format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of "209.157.22.26 0.0.0.255" as "209.157.22.26/24". The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 209.157.22.26/24 or 209.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 209.157.22.0/24 (if you have enabled display of sub-net lengths) or 209.157.22.0 0.0.0.255 in the startup-config file.

If you enable the software to display IP sub-net masks in CIDR format, the mask is saved in the file in "/<mask-bits>" format. To enable the software to display the CIDR masks, enter the **ip show-subnet-length** command at the global CONFIG level of the CLI. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format.

NOTE: If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with sub-net mask in the display produced by the **show ip access-list** command.

The **host** <source-ip> | <hostname> parameter lets you specify a host IP address or name. When you use this parameter, you do not need to specify the mask. A mask of all zeros (0.0.0.0) is implied.

The **any** parameter configures the policy to match on all host addresses.

The **log** argument configures the device to generate Syslog entries and SNMP traps for packets that are permitted or denied by the access policy.

NOTE: You can enable logging on ACLs and filters that support logging even when the ACLs and filters are already in use. To do so, re-enter the ACL or filter command and add the **log** parameter to the end of the ACL or filter. The software replaces the ACL or filter command with the new one. The new ACL or filter, with logging enabled, takes effect immediately.

Configure the Route Map

After you configure the ACLs, you can configure a PBR route map that matches based on the ACLs and sets routing information in the IP traffic.

NOTE: The match and set statements described in this section are the only route-map statements supported for PBR. Other route-map statements described in the documentation apply only to the protocols with which they are described.

To configure a PBR route map, enter commands such as the following:

```
HP9300(config)# route-map test-route permit 101
HP9300(config-routemap test-route)# match ip address 1
HP9300(config-routemap test-route)# set ip next-hop 192.168.2.1
HP9300(config-routemap test-route)# exit
```

The commands in this example configure an entry in a route map named "test-route". The **match** statement matches on IP information in ACL 1. The **set** statement changes the next-hop IP address for packets that match to 192.168.2.1.

Syntax: route-map <map-name> permit | deny <num>

The <map-name> is a string of characters that names the map. Map names can be up to 32 characters in length. You can define up 50 route maps on the Routing Switch.

The **permit | deny** parameter specifies the action the Routing Switch will take if a route matches a match statement.

- If you specify **deny**, the Routing Switch does not advertise or learn the route.
- If you specify **permit**, the Routing Switch applies the match and set statements associated with this route map instance.

The <num> parameter specifies the instance of the route map you are defining. Each route map can have up to 50 instances. Routes are compared to the instances in ascending numerical order. For example, a route is compared to instance 1, then instance 2, and so on.

Syntax: match ip address <ACL-num-or-name>

The <ACL-num> parameter specifies a standard or extended ACL number or name.

Syntax: set ip [default] next hop <ip-addr>

This command sets the next-hop IP address for traffic that matches a match statement in the route map.

If you specify **default**, the route map sets the next-hop gateway only if the Routing Switch does not already have explicit routing information for the traffic.

Syntax: set [default] interface null0 | [<portnum>...]

This command redirects the traffic to the specified interface. You can send the traffic to the null0 interface, which is the same as dropping the traffic. You can specify more than one interface, in which case the Routing Switch uses the first available port. If the first port is unavailable, the Routing Switch sends the traffic to the next port in the list.

If you specify **default**, the route map redirects the traffic to the specified interface only if the Routing Switch does not already have explicit routing information for the traffic.

Enabling PBR

After you configure the ACLs and route map entries, you can enable PBR globally, on individual interfaces, or both as described in this section. To enable PBR, you apply a route map you have configured for PBR globally or locally.

Enabling PBR Globally

To enable PBR globally, enter a command such as the following at the global CONFIG level:

```
HP9300(config)# ip policy route-map test-route
```

This command applies a route map named “test-route” to all interfaces on the device for PBR.

Syntax: ip policy route-map <map-name>

Enabling PBR Locally

To enable PBR locally, enter commands such as the following:

```
HP9300(config)# interface ve 1
HP9300(config-vif-1)# ip policy route-map test-route
```

The commands in this example change the CLI to the Interface level for virtual interface 1, then apply the “test-route” route map to the interface. You can apply a PBR route map to Ethernet ports or virtual interfaces.

Syntax: ip policy route-map <map-name>

Configuration Examples

The following sections provide configuration examples for the following uses of PBRs:

- Setting the next hop
- Setting the next hop, if the Routing Switch does not have an explicit next hop configured for the traffic

- Discarding traffic by sending it to a null interface

Setting the Next Hop

The following commands configure the Routing Switch to apply PBR to traffic from IP sub-nets 209.157.23.x, 209.157.24.x, and 209.157.25.x. In this example, route maps specify the next-hop gateway for packets from each of these sub-nets.

- Packets from 209.157.23.x are sent to 192.168.2.1.
- Packets from 209.157.24.x are sent to 192.168.2.2.
- Packets from 209.157.25.x are sent to 192.168.2.3.

The following commands configure three standard ACLs. Each ACL contains one of the ACLs listed above. Make sure you specify **permit** instead of deny in the ACLs, so that the Routing Switch permits the traffic that matches the ACLs to be further evaluated by the route map. If you specify **deny**, the Routing Switch denies the traffic from further evaluation and instead drops the packets. Notice that these ACLs specify **any** for the destination address.

```
HP9300(config)# access-list 101 permit 209.157.23.0 0.0.0.255
HP9300(config)# access-list 102 permit 209.157.24.0 0.0.0.255
HP9300(config)# access-list 103 permit 209.157.25.0 0.0.0.255
```

The following commands configure three entries in a route map called “test-route”. The first entry (permit 1) matches on the IP address information in ACL 1 above. For IP traffic from sub-net 209.157.23.0/24, this route map entry sets the next-hop IP address to 192.168.2.1.

```
HP9300(config)# route-map test-route permit 101
HP9300(config-routemap test-route)# match ip address 101
HP9300(config-routemap test-route)# set ip next-hop 192.168.2.1
HP9300(config-routemap test-route)# exit
```

The following commands configure the second entry in the route map. This entry (permit 2) matches on the IP address information in ACL 2 above. For IP traffic from sub-net 209.157.24.0/24, this route map entry sets the next-hop IP address to 192.168.2.2.

```
HP9300(config)# route-map test-route permit 102
HP9300(config-routemap test-route)# match ip address 102
HP9300(config-routemap test-route)# set ip next-hop 192.168.2.2
HP9300(config-routemap test-route)# exit
```

The following commands configure the third entry in the test-route route map. This entry (permit 3) matches on the IP address information in ACL 3 above. For IP traffic from sub-net 209.157.25.0/24, this route map entry sets the next-hop IP address to 192.168.2.3.

```
HP9300(config)# route-map test-route permit 103
HP9300(config-routemap test-route)# match ip address 103
HP9300(config-routemap test-route)# set ip next-hop 192.168.2.3
HP9300(config-routemap test-route)# exit
```

The following command enables PBR by globally applying the test-route route map to all interfaces.

```
HP9300(config)# ip policy route-map test-route
```

Alternatively, you can enable PBR on specific interfaces, as shown in the following example. The commands in this example configure IP addresses in the three source sub-nets identified in ACLs 1, 2, and 3, then apply route map test-route the interface.

```
HP9300(config)# interface ve 1
HP9300(config-vif-1)# ip address 209.157.23.1/24
HP9300(config-vif-1)# ip address 209.157.24.1/24
HP9300(config-vif-1)# ip address 209.157.25.1/24
HP9300(config-vif-1)# ip policy route-map test-route
```

Setting the Next Hop When no Next Hop Is Explicitly Configured

The following commands configure a PBR to set the next-hop gateway for traffic, but only if the Routing Switch does not already have a next-hop gateway specified for the traffic. In this example, a route map specifies the next-hop gateway for packets from sub-net 192.168.1.x.

The following command configures a standard ACL for the sub-net.

```
HP9300(config)# access-list 104 permit 192.168.1.0 0.0.0.255 any
```

The following commands configure an entry in a route map called “test-route-if-no-gateway”. The first entry (permit 4) matches on the IP address information in ACL 4 above. For IP traffic from sub-net 192.168.1.0/24, this route map entry sets the next-hop IP address to 192.111.1.1, but only if the Routing Switch does not already have a gateway configured for the sub-net.

```
HP9300(config)# route-map test-route-if-no-gateway permit 104
HP9300(config-routemap test-route-if-no-gateway)# match ip address 104
HP9300(config-routemap test-route-if-no-gateway)# set ip default next-hop
192.111.1.1
HP9300(config-routemap test-route-if-no-gateway)# exit
```

The following command enables PBR by globally applying the route map to all interfaces.

```
HP9300(config)# ip policy route-map test-route-if-no-gateway
```

Alternatively, you can enable PBR on specific interfaces, as shown in the following example. The commands in this example configure IP addresses in the source sub-net identified in ACL 4, then apply route map test-route-if-no-gateway to the interface.

```
HP9300(config)# interface ve 2
HP9300(config-vif-1)# ip address 192.168.1.34/24
HP9300(config-vif-1)# ip policy route-map test-route-if-no-gateway
```

Setting the Output Interface to the Null Interface

The following commands configure a PBR to send all traffic from 192.168.1.204/32 to the null interface, thus dropping the traffic instead of forwarding it.

```
HP9300(config)# access-list 106 permit 209.168.1.204 0.0.0.0
```

The following commands configure an entry in a route map called “file-13”. The first entry (permit 6) matches on the IP address information in ACL 6 above. For IP traffic from the host 209.168.1.204/32, this route map entry sends the traffic to the null interface instead of forwarding it, thus sparing the rest of the network the unwanted traffic.

```
HP9300(config)# route-map file-13 permit 106
HP9300(config-routemap file-13)# match ip address 106
HP9300(config-routemap file-13)# set interface null0
HP9300(config-routemap file-13)# exit
```

The following command enables PBR by globally applying the route map to all interfaces.

```
HP9300(config)# ip policy route-map file-13
```

Alternatively, you can enable the PBR on specific interfaces, as shown in the following example. The commands in this example configure IP addresses in the source sub-net identified in ACL 6, then apply route map file-13 to the interface.

```
HP9300(config)# interface ethernet 3/11
HP9300(config-if-3/11)# ip address 192.168.1.204/32
HP9300(config-if-3/11)# ip policy route-map file-13
```