Appendix D
Policies and Filters

The HP ProCurve 9304M, 9308M, and 6308M-SX routing switches and the 6208M-SX switch provide a robust array of policies and filters. You can configure policies and filters to do the following:

- Change Quality-of-Service priorities for individual ports, VLANs, Layer 4 flows, static MAC entries, and AppleTalk sockets.
- Configure protocol-based VLANs, IP sub-net VLANs, and IPX network VLANs within standard 802.1d port-based VLANs.
- Forward or drop IP packets based on source and destination IP addresses, Layer 4 information (such as TCP or UDP port), or both.
- Learn or drop IP/RIP routes on incoming traffic, based on network address or the IP/RIP neighbor's IP address.
- Control learning and advertisement of IP/RIP routes, based on network address or the IP/RIP neighbor's IP address.
- Forward or drop IPX packets based on source and destination network address and socket information.
- Control learning and advertisement of IPX RIP routes.
- Permit or deny access to IPX servers.
- Permit or deny AppleTalk zone and network information to reach other zones.
- Control learning and advertisement of routes learned from BGP4 neighbors. You can filter based on network address information, AS-path information, and community names.
- Redistribute routes among IP/RIP, OSPF, and BGP4.
- Filter on specific MAC addresses, on Layer 2 multicast packets, and on Layer 2 broadcast packets.

This appendix describes the various types of policies and filters. For each type of policy or filter, the CLI command syntax and the Web management links for configuring the policy or filter are provided. This appendix also refers you to specific configuration procedures.

Scope

Some policies and filters are configured and apply globally, while others are configured globally but apply to individual ports. The following table lists the scope for each type of policy and filter.

Table D.1: Scopes of Policies and Filters

Policy or Filter Type	Scope
QoS policy	Configured and applied to one of the following: • Ports • VLANs • Static MAC entries • Layer 4 sessions • AppleTalk sockets
Access policy (see forwarding filters)	See Forwarding filters
Forwarding filters • MAC forwarding filters • IP forwarding filters (same as IP access policy) • IPX forwarding filters • TCP/UDP forwarding filters	Configured globally, then applied locally to a port's inbound or outbound policy or filter group. You can use the same policy or filter in a port's inbound policy or filter group and outbound policy or filter group. You also can use the same policy or filter on multiple ports.
Address-lock filter	Configured and applied on individual ports.
Route filters • IP/RIP route filters • IPX RIP route filters • IPX SAP service filters	Configured globally and applied to individual ports
RIP neighbor filters	Configured and applied globally
AppleTalk zone and network filters	Configured and applied on individual ports.
BGP4 filters • BGP4 address • BGP4 AS-path • BGP4 community	Configured and applied globally and in route maps
Route redistribution filters • IP/RIP • OSPF • BGP4	Configured and applied globally

Default Filter Actions

By default, no policies or filters are defined on the routing switches and switch. The following table lists the default action when no policy or filter is configured and the default action after you configure a policy or filter. For some types of policies and filters, the default action changes once you configure a policy or filter, regardless of the policy or filter's contents.

Table D.2: Default Policy and Filter Actions

Policy or Filter Type	Default action when no policies or filters are configured	Default action after a policy or filter is configured
QoS policy	Queue all packets in normal or 0 priority queue	Queue all packets in normal or 0 priority queue unless explicitly configured for a higher queue
Access policy (see Forwarding filters)	See Forwarding filters	See Forwarding filters
Forwarding filters • MAC forwarding filters • IP forwarding filters (same as IP access policy) • IPX forwarding filters • TCP/UDP forwarding filters	Permit (forward) all packets	Deny (drop) all packets **Note**: The default action for AppleTalk zone and network filters is always permit. To deny all but specific zones, create permit filters for those zones, then create a deny filter and use the "additional zones" value with the filter.
Address-lock filter	Permit (forward) all packets	Permit only those packets whose source MAC addresses have been learned on the port; drop all others
Route filters • IP/RIP route filters • IP/RIP neighbor filters • IPX RIP route filters • IPX SAP service filters • AppleTalk zone and network filters • BGP4 address filters • BGP4 AS-path filters • BGP4 community filters	Permit (learn and advertise) all routes or services	Deny (do not learn or advertise) all routes or services
Route redistribution filter • IP/RIP • OSPF • BGP4	Do not redistribute routes	Do not redistribute routes unless explicitly redistributed by filter **Note**: For IP/RIP and OSPF, you must explicitly enable redistribution. Redistribution is enabled by default in BGP4.
Layer 2 broadcast and multicast filters	Allow outbound broadcasts and multicasts on the specified ports	Drop outbound broadcasts or multicasts on the specified ports

Policy and Filter Precedence

QoS

You can apply QoS policies to individual ports, VLANs, static MAC address, Layer 4 sessions, IP Multicast groups, and AppleTalk sockets. If a port is a member of two or more of these items and has different priorities, the priorities are merged. However, the resulting priority is never lower than the highest priority.

Precedence Among Filters on Different Layers

Generally, the device applies only the type of filter that applies to the traffic. For example, if a packet is a Layer 2 switched packet, then the device evaluates the packet against the port's MAC filters. If a packet is a routed IP packet, the device evaluates the packet against the port's IP access policies.

HP recommends that you do not use filters at different layers on the same port. For example, do not use MAC filters and IP access policies on the same port.

NOTE: You cannot use Layer 2 filters to filter for Layer 4 information. To filter for Layer 4 information, use IP access policies (filters).

NOTE: If you do choose to apply filters for multiple layers to the same port, note that Layer 2 MAC filters can affect the Layer 3 IP traffic that a port permits or denies on multinetted interfaces. A multinetted interface has multiple IP sub-net interfaces on the same port. MAC filters can filter on the Ethertype field. This field includes Layer 3 protocol information and identifies packets as IP packets, ARP packets, and so on.

If you configure a MAC filter, then leave the default action as "deny any", all packets from one of the IP sub-net addresses to another address on the same multinetted interface that do not match the filter are denied. This includes packet types such as IP and ARP. The result is that you have a Layer 2 filter but Layer 3 traffic is dropped. To avoid this, make sure you configure a filter to "permit any" traffic, thus changing the default action to permit for packets that are not denied by the other MAC filters.

Precedence Among Filters on the Same Layer

For most types of filters, a device applies filters based on the order in which you list them in a port's inbound or outbound filter list. For example, if you apply three filters, 3, 2, and 1024 to port 1/1's outbound filter list, the filters are applied in the following order: 3, 2, 1024.

You must configure the policies or filters before you can add them to a policy or filter group.

When you configure a policy or filter group, you must add all the policies or filters at the same time. You cannot edit policy or filter groups. To change a group, you must delete it, then add a new one.

NOTE: The devices apply Layer 2 broadcast and multicast filters in ascending numerical order, beginning with 1.

Policies

A policy is a set of rules that defines how the device handles packets. Table 3 lists the types of policies you can configure on the routing switches and the switch.

Table D.3: Policies

Policy Type	Supported on...		See page...
	Routing Switch	Switch	
Quality-of-Service (QoS) Policies	X	X	5
Layer 3 Policies			6
Protocol-based VLANs – either forward or drop Layer 3 traffic based on protocol (or, for IP sub-net VLANs and IPX network VLANs, sub-net or network address)	X	X	6
IP access policies – either forward or drop IP packets	X		8
Layer 4 Policies			28
TCP/UDP access policies – either forward or drop packets based on TCP or UDP port	X	X	9

Quality-of-Service Policies

The routing switches and switch support Quality-of-Service (QoS) through implementation of 802.1p/q prioritization. You can configure QoS policies for packets associated with the following items:

- Ports
- VLANs
- Static MAC entries
- Layer 4 sessions
- AppleTalk sockets.

The default queue for all packets is normal (or 0). You can change QoS policy by placing a port, VLAN, static MAC entry, Layer 4 session, IP Multicast group, or AppleTalk socket into a higher queue. See “Quality of Service Algorithm” on page C-1 for more information about the QoS algorithm.

Actions

QoS policies place packets in the specified queue for forwarding.

Scope

You can apply QoS policies to individual ports, VLANs, static MAC address, Layer 4 sessions, IP Multicast groups, and AppleTalk sockets. If a port is a member of two or more of these items and has different priorities, the priorities are merged. However, the resulting priority is never lower than the highest priority.

Syntax

Use the following CLI commands or Web management interface panels to configure QoS policies.

Table 2.4: QoS Policies

QoS Scope	CLI syntax	Web management links
Individual port	HP9300(config-if-1/1)# priority <0-7>	Port
VLAN	HP9300(config-vlan-8)# priority <0-7>	VLAN
Static MAC address[a]	HP9300(config)# static-mac-address <MAC-addr> ethernet <port-num> [priority <0-7>] [host-type|router-type]	Static Station
Layer 4 session	HP9300(config)# ip access-policy <num> priority <0-7> <ip-addr> <mask>|any <ip-addr> <mask>|any <tcp|udp> [<operator> [tcp/udp-port-num>]] HP9300(config-if-1/1)# ip access-policy-group <in|out> <policy-list> HP6308(config) ip policy <num> priority <0-7> tcp|udp <tcp/udp-port-num> global|local HP6308(config-if-1/1) ip-policy <num>	System->IP->IP Access Policy System->Layer 4 QoS
AppleTalk socket	HP9300(config)# appletalk qos socket <number> priority <0-7>	AppleTalk->Socket QoS

a.You can configure static MAC addresses on the switch but not on the routing switches.

Layer 3 Policies

Layer 3 policies are rules that control transmission and receipt of packets based on Layer 3 routing protocol information in the packets. You can configure the following types of Layer 3 policies:

- Protocol-based VLANs
- IP access policies (same as IP filters)

Protocol-Based VLANs

Within an 802.1d port-based VLAN, you can configure protocol-based VLANs that define Layer 3 broadcast domains for specific protocols. By configuring a port as a member of a protocol VLAN, you establish a forwarding policy for that port.

For example, if you have a port-based VLAN that contains ports 1 – 12, you can configure some or all of the ports in the VLAN as an AppleTalk protocol VLAN. AppleTalk broadcast traffic received on one of the ports in the AppleTalk VLAN is broadcast to the other ports in the AppleTalk VLAN, but not to ports outside the AppleTalk VLAN.

When a port in protocol-based VLAN receives a packet, the device examines the Layer 3 information in the packet to determine whether the packet type is the same as the protocol type of the VLAN.

- If the packet is the same type as the protocol of the VLAN, the device forwards the packet.
- If the packet is another protocol type, the device drops the packet.

For example, when a port in an AppleTalk VLAN receives an AppleTalk packet, the port forwards the packet. The same port drops IPX packets, unless the port also is a member of an IPX VLAN.

IP sub-net and IPX network VLANs are similar, except for these VLAN types the device examines the IP sub-net or IPX network address.

- If the IP sub-net or IPX network address matches the address of the IP sub-net VLAN or IPX network VLAN, the device forwards the packet.
- If the sub-net or network address does not match the VLAN, the device drops the packet.

See “Configuring VLANs” on page 17-1 for VLAN configuration rules and examples.

Actions

A device forwards a packet if its Layer 3 protocol information matches the protocol VLAN’s protocol type, IP sub-net, or IPX network; otherwise, the policy drops the packet.

Scope

The forwarding policy of a port-based VLAN applies only to that VLAN.

Syntax

Use the following CLI commands or Web management interface panels to configure VLAN policies.

Table 2.5: VLAN Policies

Scope	CLI syntax	Web management links
VLAN type	HP9300(config)# vlan <vlanID> by port HP9300(config-vlan-1)# [untagged] ethernet <port-num > [<to|ethernet> <port-num>]	System->VLAN-Protocol VLAN
Protocol-based VLANs	HP9300(config)# <ip-proto | ipx-proto | atalk-proto | decnet-proto | netbios-proto | other-proto> [name <string>] <ethernet> <number> [to <number>] HP9300(config-vlan-ip-proto)# dynamic <port-list> HP9300(config-vlan-ip-proto)# static <port-list> HP9300(config-vlan-ip-proto)# exclude <port-list>	System->VLAN-Protocol VLAN
IP sub-net VLANs	HP9300(config)# ip-subnet <ip address> <ip mask> [name <string>] HP9300(config-ip-subnet)# dynamic <port-list>	System->VLAN-Protocol VLAN
IPX network VLANs	HP9300(config-ip-subnet)# static <port-list> HP9300(config-ip-subnet)# exclude <port-list>	System->VLAN-Protocol VLAN

NOTE: The **untagged** command applies only if you are removing 802.1q tagging from the ports in the VLAN. 802/1q tagging allows a port to be a member of multiple port-based VLANs. Ports in a port-based VLAN are tagged by default. The default tag is 8100 and is a global parameter.

IP Access Policies

IP access policies are rules that determine whether the device forwards or drops IP packets. You create an IP access policy by defining an IP filter, then applying it to an interface. The filter consists of source and destination IP information and the action to take when a packet matches the values in the filter. You can configure an IP filter to permit (forward) or deny (drop) the packet.

You also can configure Layer 4 information in an IP filter. If you configure Layer 4 information, you are configuring a Layer 4 policy. See “TCP/UDP Access Policies” on page D-9.

You can apply an IP filter to inbound or outbound packets. When you apply the filter to an interface, you specify whether the filter applies to inbound packets or outbound packets. Thus, you can use the same filter on multiple interfaces and specify the filter direction independently on each interface.

Figure D.1 shows an example of an inbound IP access policy group applied to port 1 on slot 1 of a 9308M routing switch. In this example, packets enter the port from left to right. The first three packets have entered the port and have been permitted or denied. The two packets on the left have not yet entered the port. When they do, they will be permitted. Since the last policy in the group is a "permit any" policy, all packets that do not match another policy are permitted. The "permit any" policy changes the default action to permit.

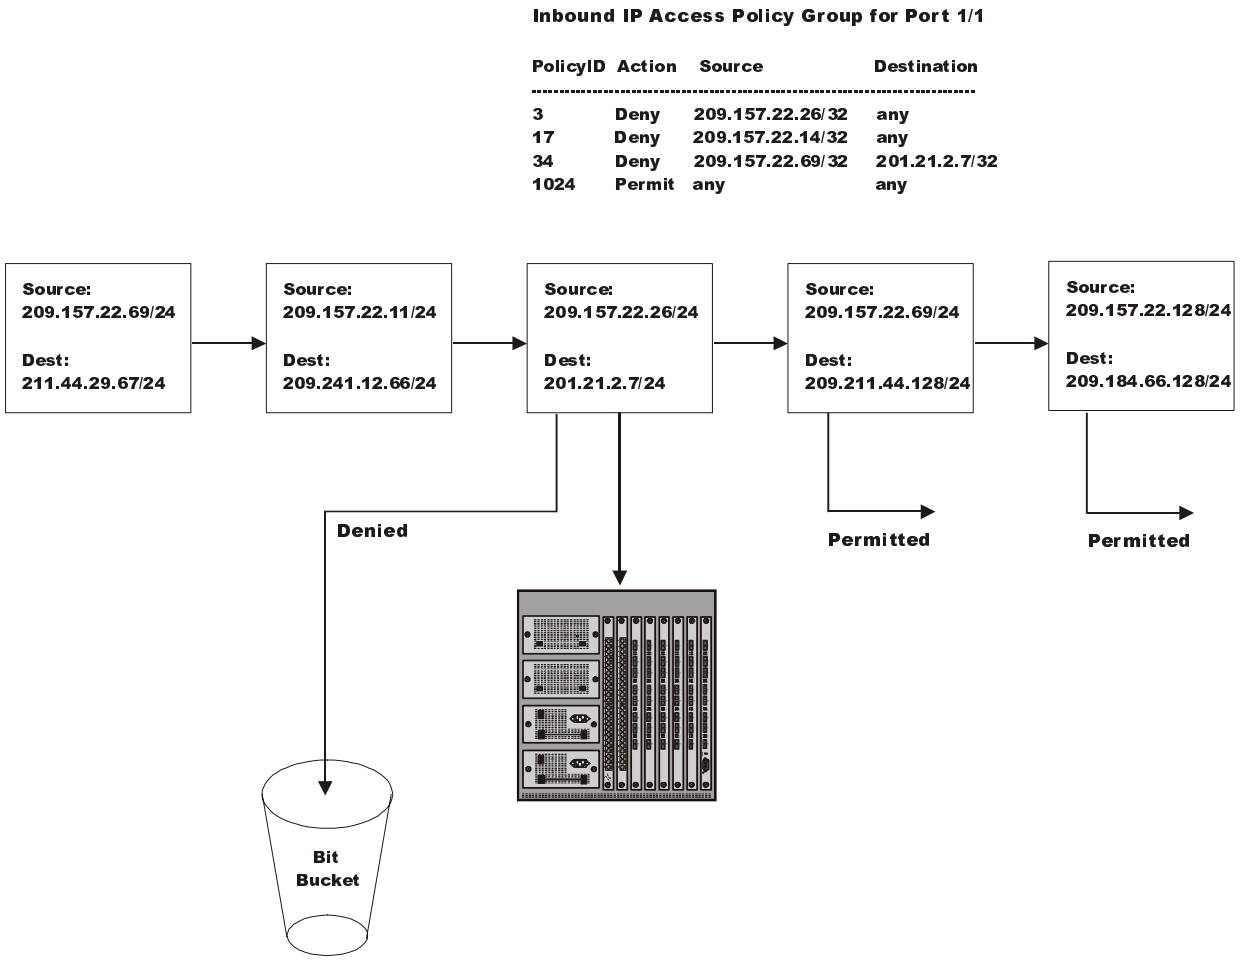

Figure D.1 IP access policies in inbound policy group for a port

Actions

IP access policies either forward or drop IP packets based on the IP source and IP destination addresses. You also can configure the policy to forward or drop a packet based on TCP/UDP port information. In this case, you are configuring a TCP/UDP access policy. See “TCP/UDP Access Policies” on page D-9.

Scope

You configure IP access policies globally, then apply them to individual ports. When you apply an IP policy to a port, you specify whether the policy applies to inbound or outbound packets. You can use the same policy in a port's inbound policy group and outbound policy group. When you configure a policy group, you must add all the policies to the group at one time. You cannot edit policy groups later. To change a policy group, you must delete the group and then add a new group.

Policies within the group are applied in positional order from left to right. Make sure you specify the filters in the order you want the device to apply them.

Syntax

Use the following CLI commands or Web management interface panels to configure IP access policies.

Table 2.6: IP Access Policies

CLI syntax	Web management links
HP9300(config)# ip access-policy <policy-num> <permit|deny> <ip-addr> <mask>|any <ip-addr> <mask>|any <tcp|udp> [<operator> [<tcp/udp-port-num>]] [log] HP9300(config-if-1/1)# ip access-policy-group <in|out> <policy-list>	System->IP Access Policy

Layer 4 Policies

Layer 4 policies are rules that control transmission and receipt of packets based on Layer 4 transport information. You can configure the following types of Layer 4 policies:

- TCP/UDP access policies (same as TCP/UDP filters)

TCP/UDP Access Policies

TCP/UDP access policies are IP filters that contain Layer 4 information. Layer 4 policies enable you to forward or drop packets for individual Layer 4 applications, giving you finer access control. You do not need to completely block an IP address to deny certain types of traffic from that address. You can selectively allow some types of traffic while dropping others. For example, you can configure a Layer 4 policy to drop web (HTTP) packets from a host but allow all other traffic from the host.

You can filter on the following Layer 4 application types:

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

For TCP and UDP, you also specify an operator and the port number or well-known name for the port. For example, if you want to filter on FTP traffic, you configure the filter to match on packets that contain the TCP application port number for FTP.

When you can configure a Layer 4 policy, you specify the source and destination IP address of the hosts or servers for which you are controlling access.

Figure D.2 shows an example of TCP/UDP access policies. Although this example does not explicitly identify these policies as inbound policies or outbound policies, when you apply the policies to individual ports you specify whether they are for inbound or outbound traffic.

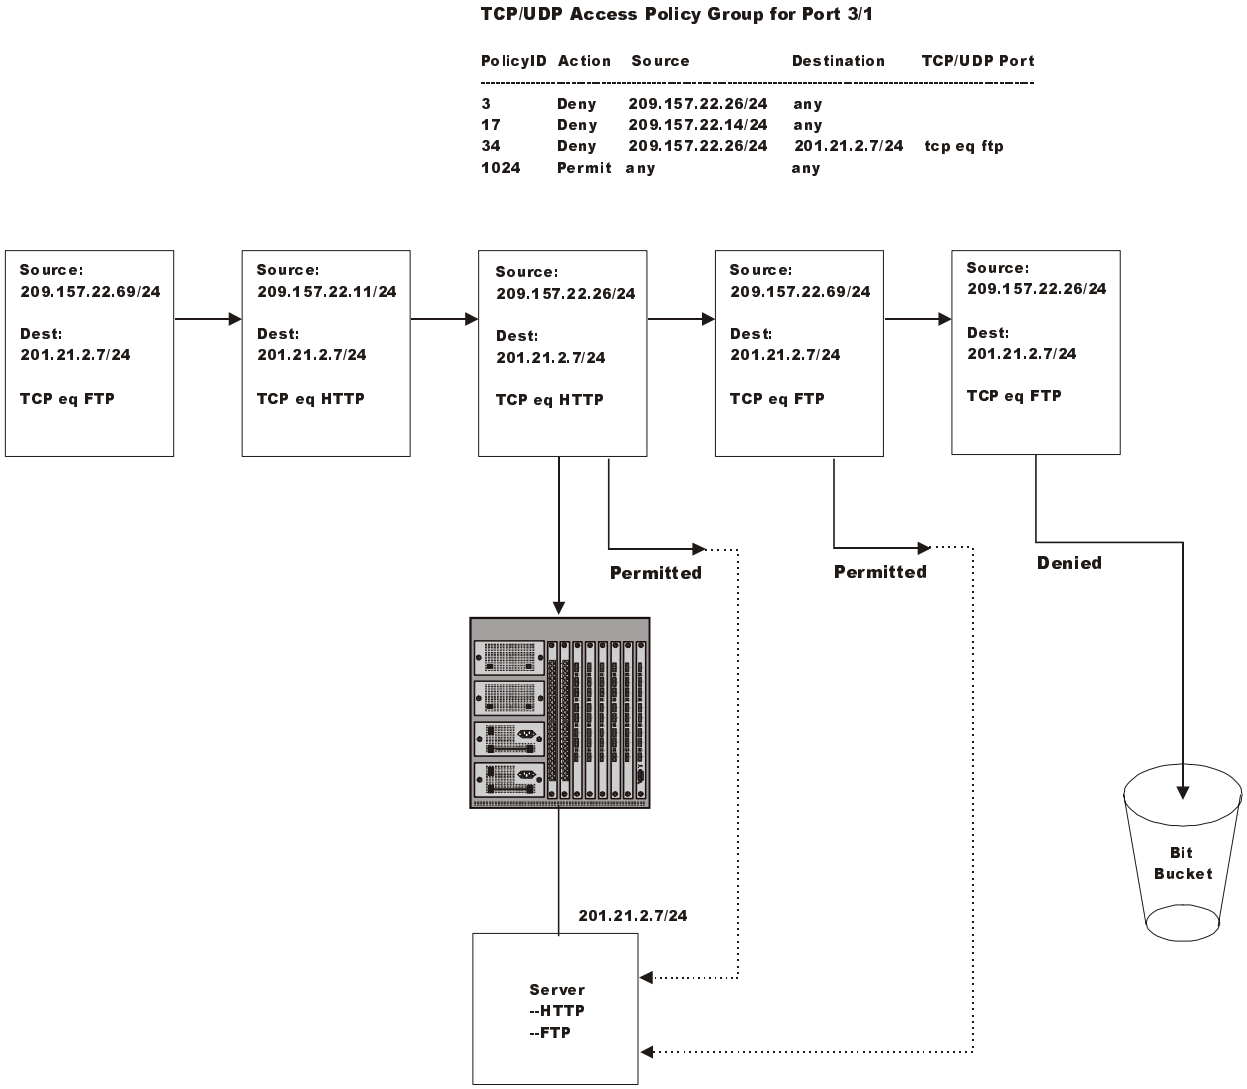

Figure D.2 TCP/UDP Access Policies

Actions

TCP/UDP access policies forward (permit) or drop (deny) IP packets based on the Layer 4 application information in the packets.

Scope

You configure TCP/UDP access policies globally, then apply them to individual ports. When you apply a TCP/UDP policy to a port, you specify whether the policy applies to inbound or outbound packets. You can use the same policy in a port's inbound policy group and outbound policy group. When you configure a policy group, you must add all the policies to the group at one time. You cannot edit policy groups later. To change a policy group, you must delete the group and then add a new group.

Policies within the group are applied in positional order from left to right. Make sure you specify the filters in the order you want the device to apply them.

Syntax

Use the following CLI commands or Web management interface panels to configure TCP/UDP access policies.

Table 2.7: TCP/UDP Access Policies

CLI syntax	Web management links
HP9300(config)# ip access-policy <policy-num> <permit|deny> <ip-addr> <mask>|any <ip-addr> <mask>|any <tcp|udp> [<operator> [<tcp/udp-port-num>]] [log] HP9300(config-if-1/1)# ip access-policy-group <in|out> <policy-list> HP6308(config) ip policy <num> priority <0-7> tcp|udp <tcp/udp-port-num> global|local HP6308(config-if-1/1) ip-policy <num>	System->IP Access Policy

Filters

A filter is a set of comparison values and an action. If a packet matches the set of values in the filter, the device takes the action specified in the filter. The routing switches and switch provide filters for Layer 2, Layer 3, and Layer 4. A filter looks at the appropriate fields in a packet to compare information related to one of the layers. For example, MAC filters look at the source and destination MAC address and, optionally, at the encapsulation information. IPX filters look at the source and destination network and socket information but do not look at the MAC information.

The following table lists the various types of filters you can configure on the routing switches and the switch.

Table 2.8: Filters

Filter Type	Supported on...		See page...
	Routing Switch	Switch	
Layer 2 Filters			12
MAC filters	X	X	12
Broadcast filters	X	X	13
Multicast filters	X	X	13
Address-lock filters	X	X	14
Layer 3 Filters			16
IP forwarding filters (same as IP access policies)	X		8
IP/RIP route filters	X		16
IP/RIP neighbor filters	X		18
IPX forwarding filters	X		19
IPX RIP filters	X		19
IPX SAP filters	X		20

Table 2.8: Filters (Continued)

Filter Type	Supported on...		See page...
	Routing Switch	Switch	
AppleTalk zone filters	X		21
AppleTalk network filters	X		22
BGP address filters	X		22
BGP AS-path filters	X		23
BGP community filters	X		24
IP/RIP redistribution filters	X		26
OSPF redistribution filters	X		27
BGP redistribution filters	X		27
Layer 4 Filters			28
TCP/UDP forwarding filters (same as TCP/UDP access policies)	X	X	9

Layer 2 Filters

Layer 2 filters control a device's transmission and receipt of packets based on MAC address information. The routing switches and switch provide the following types of Layer 2 filters:

- MAC address filters
- Address-lock filters

MAC Filters

MAC filters forward or drop packets based on the following information:

- Source MAC address
- Destination MAC address
- Encapsulation type and EtherType (optional)

A packet whose Layer 2 information matches the filter is either permitted (forwarded) or denied (dropped). You define a MAC filter on the global level, then apply it to an interface by adding it to that interface's inbound or outbound MAC filter group.

Action

MAC filters forward (permit) or drop (deny) packets.

Scope

You configure MAC filters globally, then apply them to individual ports. The filters do not take effect until applied to specific ports. You can apply them to a port's inbound MAC filter group, outbound MAC filter group, or both.

Syntax

Use the following CLI commands or Web management interface panels to configure MAC filters.

Table 2.9: MAC Filters

CLI syntax	Web management links
HP9300(config)# mac filter <filter-num> <permit|deny> <any|H.H.H> <any|H.H.H> <etype|IIc|snap> <operator> <frame-type> HP9300(config-if-1/1)# mac-filter-group <filter-list>	System->MAC Filter

Broadcast Filters

Broadcast filters are outbound filters that drop Layer 2 broadcast packets that match the filter criteria. You can filter on all broadcast traffic or on IP UDP broadcast traffic only. You also can specify a VLAN ID so that broadcasts are dropped only for the specified VLAN.

You can configure up to eight broadcast filters.

NOTE: Broadcast filters are applied in numerical order, beginning with filter 1.

Action

Broadcast filters forward (permit) or drop (deny) packets.

Scope

You configure broadcast filters globally, then apply them to individual ports. The filters do not take effect until applied to specific ports. The filters apply only to outbound traffic. The **exclude-ports** command specifies the ports on which you are excluding the filtered multicast packets.

Syntax

Use the following CLI commands or Web management interface panels to configure broadcast filters.

Table 2.10: Broadcast Filters

CLI syntax	Web management links
HP9300(config)# broadcast filter <filter-ID> any | ip udp [vlan <vlan-id>] exclude-ports ethernet <port-num> to <port-num> Or exclude-ports ethernet <port-num> ethernet <port-num>	Not available

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

Multicast Filters

Multicast filters are outbound filters that apply to packets that have a Layer 2 multicast address in the destination MAC address field. You can configure multicast filters to filter on all multicast addresses or a specific multicast address.

You can configure up to eight multicast filters.

NOTE: Multicast filters are applied in numerical order, beginning with filter 1.

Action

Multicast filters forward (permit) or drop (deny) packets.

Scope

You configure multicast filters globally, then apply them to individual ports. The filters do not take effect until applied to specific ports. The filters apply only to outbound traffic. The **exclude-ports** command specifies the ports on which you are excluding the filtered multicast packets.

Syntax

Use the following CLI commands or Web management interface panels to configure multicast filters.

Table 2.11: Multicast Filters

CLI syntax	Web management links
HP9300(config)# multicast filter <filter-ID> any|ip udp mac <multicast-address>|any [mask <mask>] [vlan <vlan-id>] exclude-ports ethernet <port-num> to <port-num> Or exclude-ports ethernet <port-num> ethernet <port-num>	Not available

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

Address-Lock Filters

Address-lock filters limit the number of MAC addresses that can be learned on a port. The port forwards only those packets that contain one of the source MAC addresses learned by the port. The port drops other packets. In addition, the device generates an SNMP trap for other packets received by the port.

Figure D.3 shows an example of an address-lock filter. In this example, the device is configured to learn only two MAC addresses on port 1/1. After the device learns two addresses, port 1/1 can forward only a packet whose source address is one of the two learned addresses. The port drops all other packets. This applies even to MAC broadcasts. If one of the packets learned on the port is not addressed to the MAC broadcast address, the port cannot forward MAC broadcasts.

The device learns MAC addresses from the source-MAC-address field of inbound packets received on the port.

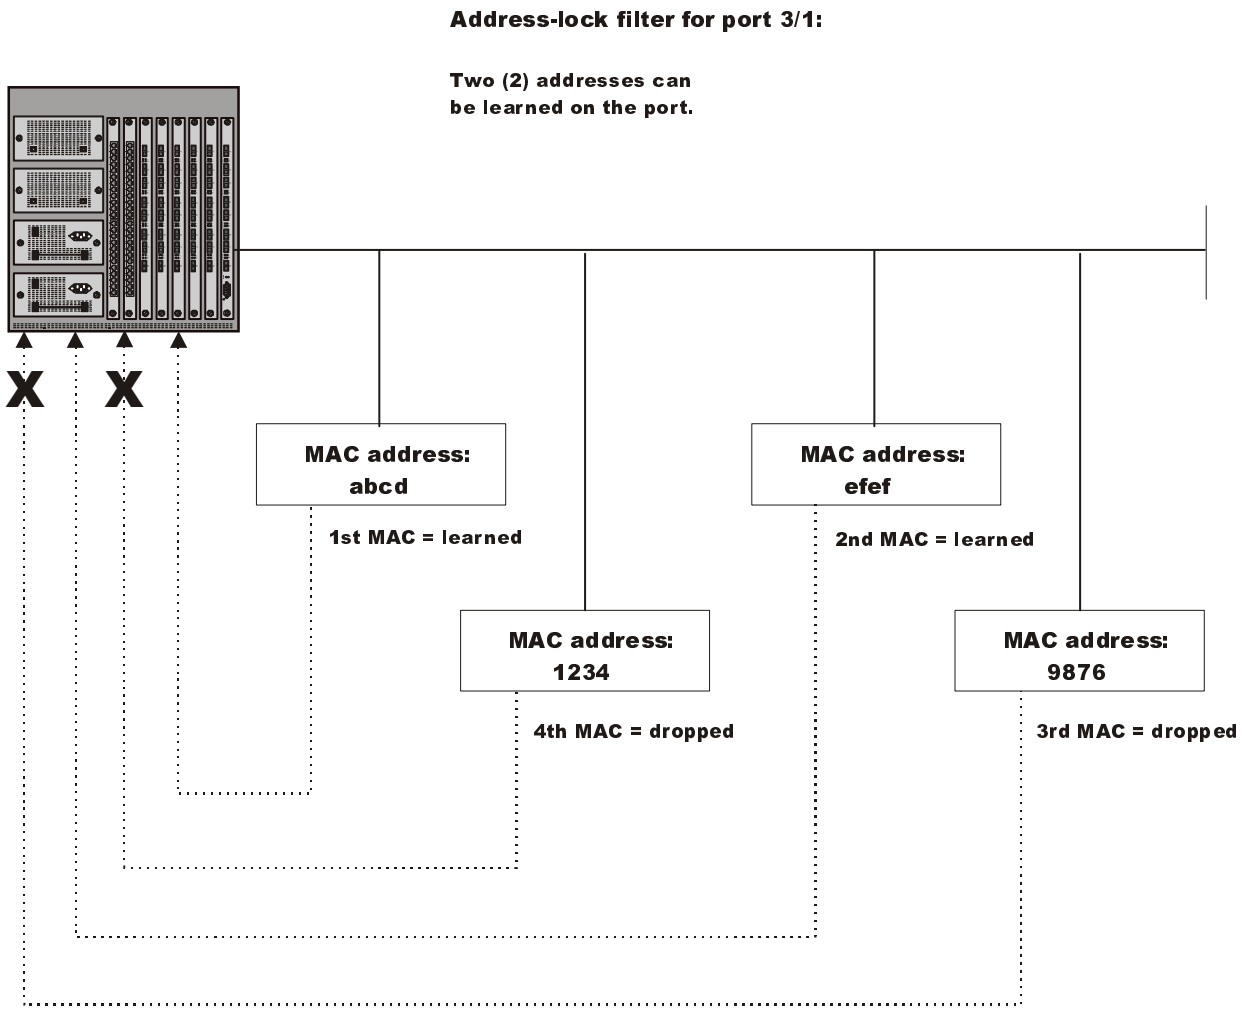

Figure D.3 Address-lock filter

Actions

Forward (permit) only those packets with a MAC address that the port has learned. Deny all other packets.

Scope

You configure a lock address filter globally, but you also specify the port as part of the filter.

Syntax

Use the following CLI commands or Web management interface panels to configure lock-address filters.

Table 2.12: Lock-Address Filters

CLI syntax	Web management links
HP6308(config)# lock-address ethernet <port> addr-count <num>	Port->Modify

Layer 3 Filters

Layer 3 filters control a device's transmission and receipt of packets based on routing protocol information in the packets. The routing switches and switch provide the following types of Layer 3 filters:

- IP forwarding filters (same as IP access policies, see "IP Access Policies" on page D-8)
- IP/RIP route filters
- IP/RIP neighbor filters
- IPX forwarding filters
- IPX RIP route and neighbor filters
- IPX SAP service filters
- AppleTalk zone filters
- AppleTalk network filters
- BGP route address filters
- BGP route AS-path filters
- BGP route community filters
- IP/RIP redistribution filters
- OSPF redistribution filters
- BGP redistribution filters

IP Filters

IP filters control the IP packets that the device sends and receives and the routes that the device learns or advertises. IP forwarding filters (IP Access policies) control transmission and receipt of IP packets, while IP/RIP route and neighbor filters control the routes that the device leans or advertises. Route filters filter on specific network addresses while neighbor filters filter on the IP addresses of the IP/RIP neighbors.

IP Forwarding Filters

IP forwarding filters determine whether to forward or drop an IP packet. IP forwarding filters on a switch or routing switch are called "IP access policies". See "IP Access Policies" on page D-8.

IP/RIP Route Filters

IP/RIP route filters control the routes that a device learns and advertises. Figure D.4 shows an example of a port with IP/RIP route filters. The port has filters for the inbound direction and the outbound direction. Notice that the same filter can be used for both directions. The inbound filters control the routes that the device learns; denied routes are not learned by the device. Outbound filters control the routes that the device advertises; denied routes are not advertised to RIP neighbors.

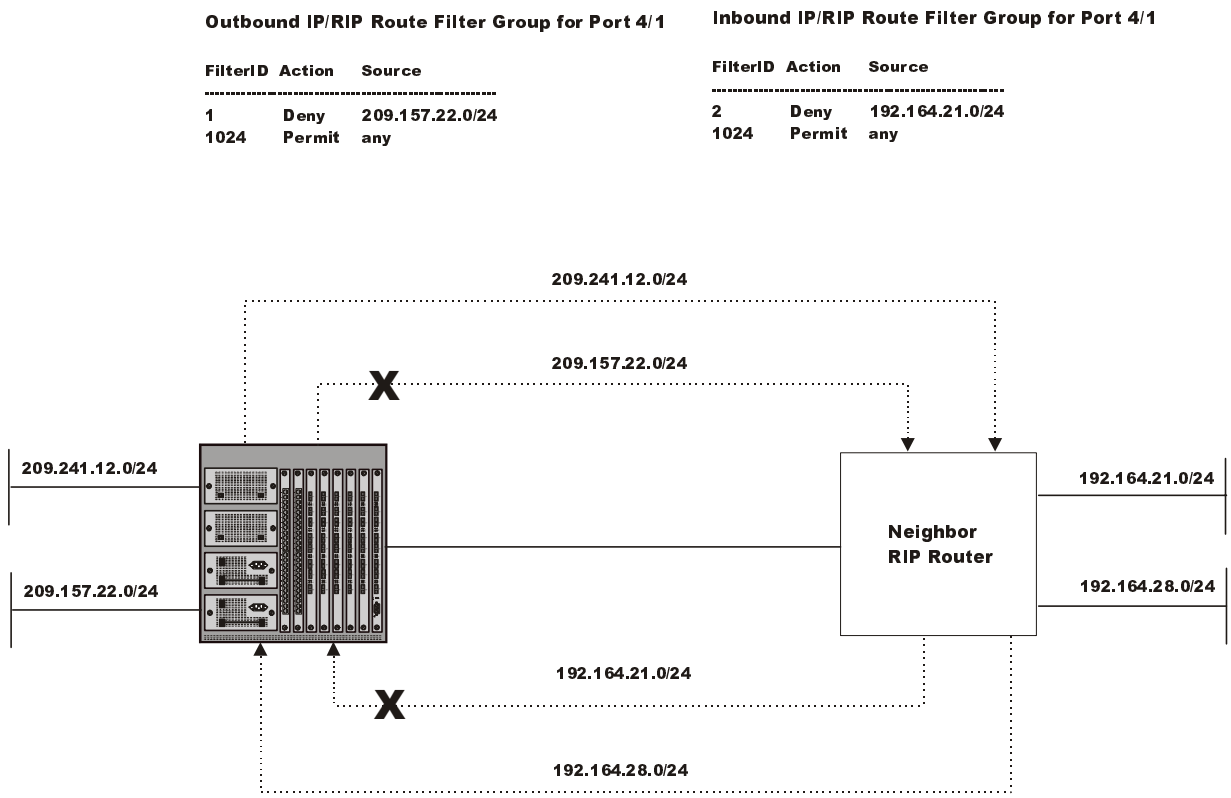

Figure D.4 IP/RIP route filters

Actions

- An IP/RIP route filter applied to outbound traffic on a port permits or denies advertisement of routes.
- An IP/RIP route filter applied to inbound traffic on a port permits or denies learning of the route. When the device learns an IP/RIP route, the route is added to the IP/RIP route table.

Scope

You configure IP/RIP route filters globally, then apply them to specific ports.

Syntax

Use the following CLI commands or Web management interface panels to configure IP/RIP route filters.

Table 2.13: IP/RIP Route Filters

CLI syntax	Web management links
HP9300(config-rip-router)# filter <filter-num> <permit|deny> <source-ip-address|any> <source-mask|any> HP9300(config-if-1/1)# ip rip filter-group in|out <filter-list>	RIP->RIP Route Filter

IP/RIP Neighbor Filters

IP/RIP neighbor filters specify the IP/RIP neighbors the device can receive updates from or send updates to. You identify the neighbor by specifying its IP address in the filter. Figure D.5 shows an example of an IP/RIP neighbor filter. In this example, the device is configured to drop all IP/RIP advertisements from the IP/RIP neighbor 192.99.26.1/24. Since this is an outbound filter, the filter does not affect advertisements received by the device from 192.99.26.1/24. The device can still learn IP/RIP routes from this neighbor.

Inbound IP/RIP Neighbor Filter for Port 4/3

FilterID	Action	Source
1	Deny	192.99.26.1/24
1024	Permit	any

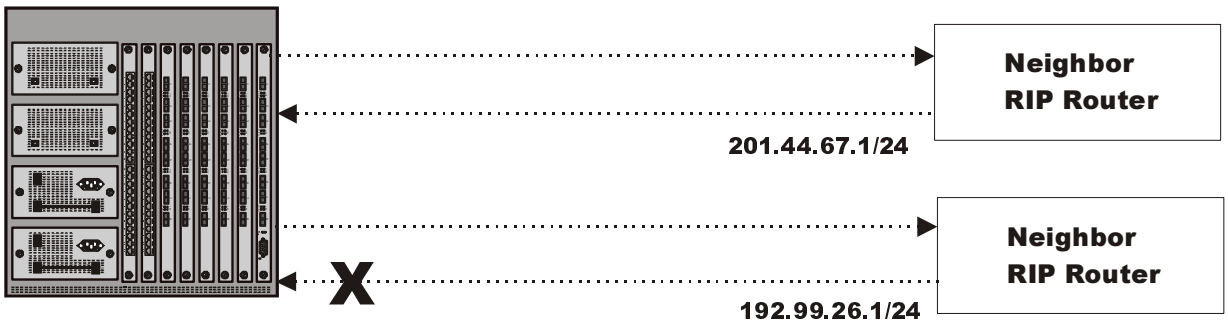

Figure D.5 IP/RIP neighbor filters

Actions

- An IP/RIP neighbor filter applied to outbound traffic on a port permits or denies advertisement of routes.
- An IP/RIP neighbor filter applied to inbound traffic on a port permits or denies learning of the routes advertised by the neighbor. When the device learns an IP/RIP route, the route is added to the IP/RIP route table.

Scope

You configure IP/RIP neighbor filters globally. They are automatically applied to all RIP ports as soon as you configure them.

Syntax

Use the following CLI commands or Web management interface panels to configure IP/RIP neighbor filters.

Table 2.14: IP/RIP Neighbor Filters

CLI syntax	Web management links
HP9300(config-rip-router)# neighbor <filter-num> permit|deny <source-IP-address|any>	RIP->RIP Neighbor Filter

IPX Filters

IPX filters control transmission and receipt of IPX packets, IPX RIP routes, and IPX Service Advertisement Protocol (SAP) messages. IPX forwarding filters filter on source and destination IPX address and socket information. IPX RIP filters filter based on a route's network address. IPX SAP filters filter based on server type and server name.

IPX Forwarding Filters

IPX forwarding filters control forwarding of IPX packets.

Action

- An IPX forward filter applied to inbound packets forwards or drops IPX packets received on the port.
- An IPX forward filter applied to outbound traffic forward or drops IPX packets sent to the port for forwarding.

Scope

You configure IPX forwarding filters globally, then apply them to specific ports.

Syntax

Use the following CLI commands or Web management interface panels to configure IPX forwarding filters.

Table 2.15: IPX Forwarding Filters

CLI syntax	Web management links
HP9300(config)# ipx forward-filter <filter-num> <permit|deny> <source-network-number|any> <source-node-number|any> <destination-network-number|any> <destination-node-number|any> <destination-socket-number|any> HP9300(config-if-1/1)# ipx forward-filter-group <in|out> <filter-list>	IPX->Forward Filter

IPX RIP Filters

IPX RIP filters control the IPX routes that the device learns or advertises.

Actions

- An IPX RIP filter applied to inbound packets learns or drops IPX routes received on the port.
- An IPX RIP filter applied to outbound packets advertises or does not advertise IPX routes.

Scope

You configure IPX RIP filters globally, then apply them to specific ports.

Syntax

Use the following CLI commands or Web management interface panels to configure IPX RIP filters.

Table 2.16: IPX RIP Filters

CLI syntax	Web management links
HP9300(config)# ipx rip-filter <filter-num> <permit|deny> <network-number|any> <network-mask|any> HP9300(config-if-1/1)# ipx rip-filter-group <in|out> <filter-list>	IPX->RIP Filter

IPX SAP Filters

IPX Service Advertisement Protocol (SAP) filters control client access to IPX servers.

Actions

- An IPX SAP filter applied to inbound packets learns or drops advertisements for the specific services.
- An IPX SAP filter applied to outbound traffic advertises or does not advertise services.

Scope

You configure IPX SAP filters globally, then apply them to specific ports.

Syntax

Use the following CLI commands or Web management interface panels to configure IPX SAP filters.

Table 2.17: IPX SAP Filters

CLI syntax	Web management links
HP9300(config)# ipx sap-filter <filter-num> <permit|deny> <server-type|any> <server-name|any> HP9300(config-if-1/1)# ipx sap-filter-group <in|out> <filter-list>	IPX->SAP Filter

Appletalk Filters

AppleTalk filters control access to AppleTalk zones and networks.

- AppleTalk zone filters permit or deny advertisement of zone names but allow network information to be learned and forwarded. Users cannot see the zone names in their Choosers but you can ping the networks. Zone filters are quite useful for reducing protocol overhead caused by "chatty" AppleTalk traffic. Use zone filtering to block information from a specific routing switch to Macintosh computers.
- AppleTalk network filters also can filter network information. When you configure an AppleTalk zone filter to deny zones, you can configure the filter to also deny the network information. To configure an AppleTalk filter to filter network information, use the RTMP filtering option with the filter.

Figure D.6 shows an example of an AppleTalk zone filter. In this example, Macintosh computers in the Marketing zone cannot see the Engineering zone. RTMP filtering is not used on this filter. Therefore, users in the Marketing zone can still ping individual devices in the Engineering zone. However, the overhead caused by unnecessary zone information exchanges between the two groups is eliminated.

To prevent users in the Marketing zone from even pinging individual devices in the Engineering zone, the RTMP filtering option can be used with the filter.

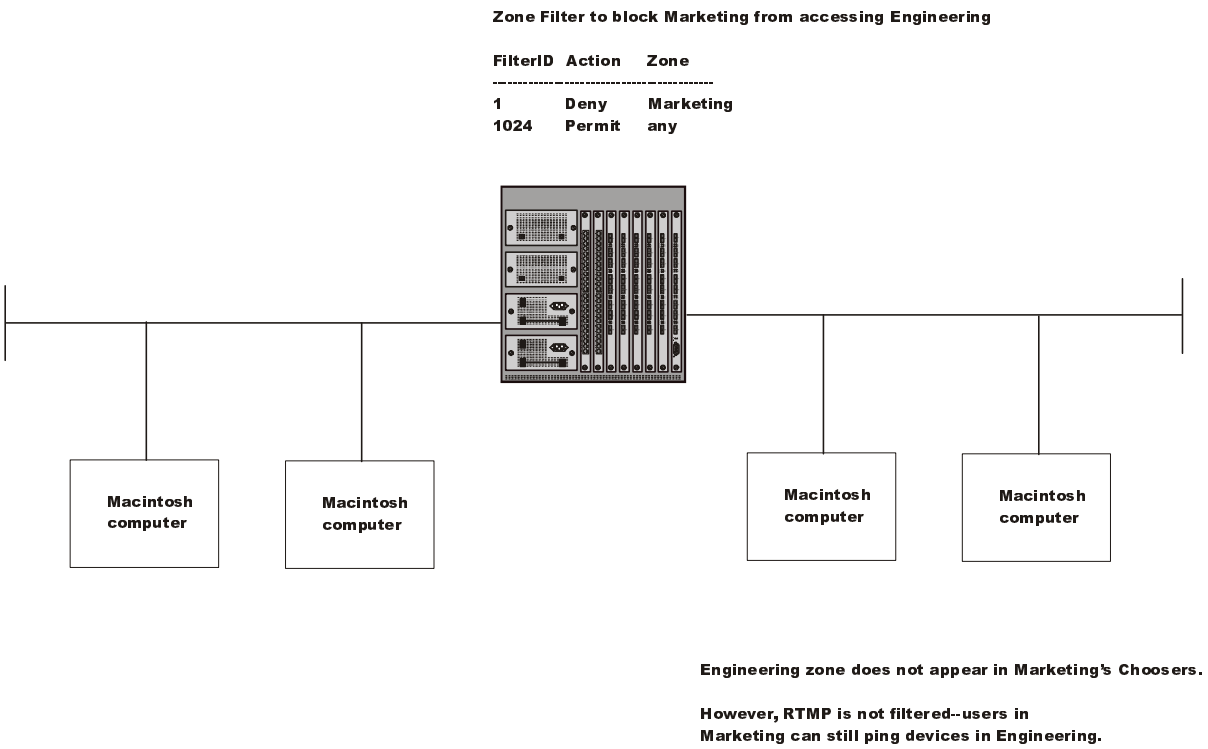

Figure D.6 AppleTalk zone filter

Appletalk Zone Filters

AppleTalk zone filters let you secure access to an AppleTalk zone. The filter controls whether the routing switch includes the zone in replies to a MAC chooser's ZIP GetZoneList request.

Actions

An AppleTalk zone filter permits (advertises) or denies (does not advertise) the specified zone. The zone does not appear in MAC user's choosers but you can still ping the networks that belong to the zone.

NOTE: Unlike other filters, the default action for AppleTalk filters does not change from permit to deny when you create a filter. To permit only specific zones and deny all others, create permit filters for the zones you want to permit, then use the following command to create a deny filter for all other zones:
appletalk deny zone additional-zones.

Scope

You configure and apply AppleTalk zone filters on individual ports.

Syntax

Use the following CLI commands or Web management interface panels to configure AppleTalk zone filters.

Table 2.18: AppleTalk Zone Filters

CLI syntax	Web management links
HP9300(config-if-1/1)# appletalk permit zone <string> HP9300(config-if-1/1)# appletalk deny zone <string>|additional-zones <rtmp-filtering|no-rtmp-filtering>	AppleTalk->Zone Filter

NOTE: If you use the **rtmp-filtering|no-rtmp-filtering** parameter, you are configuring an AppleTalk network filter. See the following section.

Appletalk Network Filters

Routing Table Maintenance Protocol (RTMP) filtering enhances a zone filter by hiding the cable ranges inside the zones used by other routing switches. The denied network numbers of the filtered zone will be removed from the RTMP packets.

The Macintosh chooser uses ZIP GetZoneList request to compile a list of zones available, so if the zone is not there the Macintosh computer cannot access it. RTMP filtering is useful for preventing downstream and adjacent routers from responding to GetZoneList requests that could give access to the zones you want to filter. All routing switches on the same segment should be configured with the same filters. You can prevent local Macintosh computers from accessing a zone but still allow the downstream routers with Macintosh computers attached to other networks to access those zones. To do so, do not use the RTMP filtering option with the zone filter.

When you configure an AppleTalk zone filter to also filter network information, the device removes route information for the networks in the specified zone before sending the RTMP packet out on the port.

Actions

AppleTalk network filters remove information about the networks in the denied zones before sending RTMP packets to Macintosh computers.

NOTE: AppleTalk network filters only deny information; they do not permit information.

Scope

You configure and apply AppleTalk network filters on individual ports.

Syntax

Use the following CLI commands or Web management interface panels to configure AppleTalk zone filters.

Table 2.19: AppleTalk Zone Filters

CLI syntax	Web management links
HP9300(config-if-1/1)# appletalk permit zone <string> HP9300(config-if-1/1)# appletalk deny zone <string>|additional-zones <rtmp-filtering|no-rtmp-filtering>	AppleTalk->Zone Filter

NOTE: If you do not use the **rtmp-filtering|no-rtmp-filtering** parameter, you are configuring an AppleTalk zone filter.

BGP4 Filters

Border Gateway Protocol version 4 (BGP4) filters control the routes that a device learns from BGP4 neighbors and advertises to BGP4 neighbors. You can configure filters to filter route information based on network address, AS-path, or community name.

BGP4 Address Filters

BGP4 address filters control whether the device learns or drops BGP4 route information based on the route's network address.

Actions

- A BGP4 address filter applied to inbound packets permits (learns) or denies (drops) the specified network address in BGP4 updates received from a BGP4 neighbor.
- A BGP4 address filter applied to outbound packets permits (advertises) or denies (drops) the specified network address in BGP4 updates the device sends to a BGP4 neighbor.

Scope

You define BGP4 address filters globally, then apply them as part of a BGP4 neighbor's distribute list or as part of a match statement in a route map.

Syntax

Use the following CLI commands or Web management interface panels to configure BGP4 address filters.

Table 2.20: BGP4 Address Filters

CLI syntax	Web management links
HP9300(config-bgp-router)# address-filter <num> permit|deny <IP-addr> <netmask>|any <IP-addr> <netmask>|any HP9300(config-bgp-router)# neighbor <router ID> remote-as <AS number> [advertisement-interval <num>] [distribute-list in|out <num,num,...>] [ebgp-multihop] [filter-list in|out <num,num,...>] [maximum-prefix <num>] [next-hop-self] [remote-as <AS number>] [route-map <map name>] [send-community] [weight <num>] HP9300(config-bgp-routemap RMAP_NAME)# match as-path-filters|community-filters|address-filters <num,num,...> [metric <num>] [next-hop <IP-addr>] [route-type <internal|external-type1|external-type2>] [tag <tag-value>]	BGP->Address Filter

NOTE: The **neighbor** command adds a BGP neighbor. The **distribute-list** parameter specifies a list of address filters and whether the list is applied to inbound or outbound BGP updates.

NOTE: The **match** command compares the information you configure for the command's parameters against BGP routes. You use this command when configuring a route map. If the comparison matches a route, set statements in the route map specify the action to take. See "Defining Route Maps" on page 12-36.

BGP4 AS-Path Filters

BGP4 AS-path filters control whether the device learns or drops BGP4 route information based on the route's AS-path. The ***AS-path*** is the list of BGP4 autonomous systems (ASs) through which the route information has traveled to reach the device.

Actions

- A BGP4 AS-path filter applied to inbound packets permits (learns) or denies (drops) routes for networks with the specified AS-path in BGP4 updates received from a BGP4 neighbor.
- A BGP4 AS-path filter applied to outbound packets permits (advertises) or denies (drops) routes for networks with the specified AS-path in BGP4 updates sent to a BGP4 neighbor.

Scope

You define BGP4 AS-path filters globally, then apply them as part of a BGP4 neighbor's distribute list or as part of a match statement in a route map.

Syntax

Use the following CLI commands or Web management interface panels to configure BGP4 AS-path filters.

Table 2.21: BGP4 AS-Path Filters

CLI syntax	Web management links
HP9300(config-bgp-router)# as-path-filter <num> permit|deny <AS-path> HP9300(config-bgp-router)# neighbor <router ID> remote-as <AS number> [advertisement-interval <num>] [distribute-list in|out <num,num,...>] [ebgp-multihop] [filter-list in|out <num,num,...>] [maximum-prefix <num>] [next-hop-self] [remote-as <AS number>] [route-map <map name>] [send-community] [weight <num>] HP9300(config-bgp-routemap RMAP_NAME)# match as-path-filters|community-filters|address-filters <num,num,...> [metric <num>] [next-hop <IP-addr>] [route-type <internal|external-type1|external-type2>] [tag <tag-value>]	BGP->AS Path Filter

NOTE: The <AS-path> value can be a regular expression. See "Using Regular Expressions" on page 12-33.

NOTE: The **neighbor** command adds a BGP neighbor. The **filter-list** parameter specifies a list of AS-path filters and whether the list is applied to inbound or outbound BGP updates.

NOTE: The **match** command compares the information you configure for the command's parameters against BGP routes. You use this command when configuring a route map. If the comparison matches a route, set statements in the route map specify the action to take. See "Defining Route Maps" on page 12-36.

BGP4 Community Filters

BGP4 community filters control whether the device learns or drops BGP4 route information based on the route's community membership.

Actions

- A BGP4 community filter applied to inbound packets permits (learns) or denies (drops) routes for networks with the specified community membership in BGP4 updates received from a BGP4 neighbor.
- A BGP4 AS-path filter applied to outbound packets permits (advertises) or denies (drops) routes for networks with the specified community membership in BGP4 updates sent to a BGP4 neighbor.

Scope

You define BGP4 community filters globally, then apply them as part of a BGP4 neighbor's distribute list or as part of a match statement in a route map.

Syntax

Use the following CLI commands or Web management interface panels to configure BGP4 community filters.

Table 2.22: BGP4 Community Filters

CLI syntax	Web management links
HP9300(config-bgp-router)# community-filter <filter-num> <permit|deny> <num>|internet|no-advertise|no-export HP9300(config-bgp-routemap RMAP_NAME)# match as-path-filters|community-filters|address-filters <num,num,...> [metric <num>] [next-hop <IP-addr>] [route-type <internal|external-type1|external-type2>] [tag <tag-value>]	BGP->Community Filter

NOTE: The **match** command compares the information you configure for the command's parameters against BGP routes. You use this command when configuring a route map. If the comparison matches a route, set statements in the route map specify the action to take. See "Defining Route Maps" on page 12-36.

Redistribution Filters

Redistribution filters control the exchange of routes between routing protocols. IP/RIP, OSPF, and BGP4 support redistribution of one another's routes. In addition, they all allow exchange of static routes.

You configure IP/RIP and OSPF redistribution filters to permit or deny routes for specific network addresses. Optionally, you can also filter on and modify the route metric. To configure redistribution, you configure redistribution filters in the protocol that will receive the routes. Redistribution is disabled by default in RIP and OSPF and enabled by default in BGP4.

BGP4 redistribution filters can filter based on a route's metric, weight, and also on the results of comparison of the route information with a route map. A ***route map*** is a named set of match conditions and parameter settings that a routing switch can use to modify route attributes and to control redistribution of routes. For more information, see "Defining Route Maps" on page 12-36.

BGP4 allows you to include the redistribution filters as part of a route map. A route map examines and modifies route information exchanged between BGP4 and IP/RIP or OSPF. See "Configuring BGP4" on page 12-1 for more information.

Figure D.7 shows an example of a redistribution filter. In this example, redistribution filters in OSPF are configured to redistribute two RIP routes into OSPF. Notice that unlike some other filter examples in this appendix, a filter for permitting all routes (to change the default action) is not configured. To maintain tight control over redistribution, the default action "deny any" is allowed to remain. Only routes that explicitly match the permit filters are permitted to be redistributed. Thus, in Figure D.7, the RIP route to 191.47.12.0/24 is not redistributed because there is no "permit any" filter that changes the default action from deny to permit.

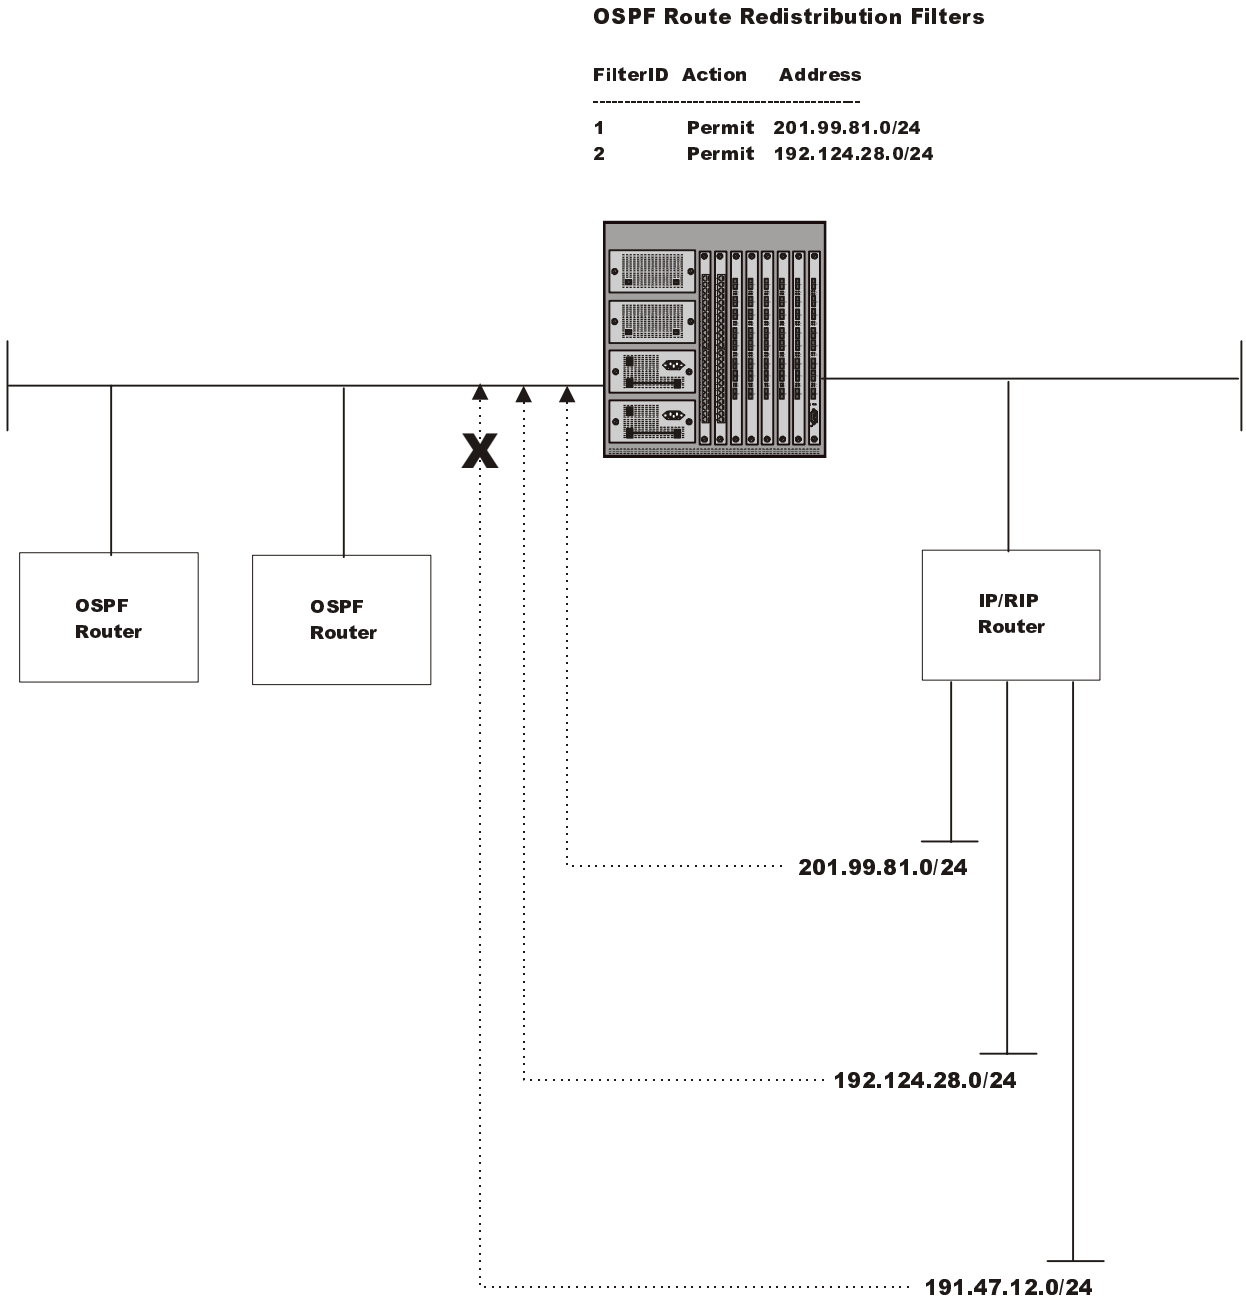

Figure D.7 OSPF redistribution filters

IP/RIP Redistribution Filters

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

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

Actions

IP/RIP redistribution filters permit (redistribute) or deny (do not redistribute) OSPF or BGP4 routes into IP/RIP.

Scope

You configure IP/RIP redistribution filters globally. They are automatically applied as soon as you configure them.

Syntax

Use the following CLI commands or Web management interface panels to configure IP/RIP redistribution filters.

Table 2.23: IP/RIP Redistribution Filters

CLI syntax	Web management links
HP9300(config-rip-router)# permit|deny redistribute <filter-num> all|bgp|ospf|static <ip-addr> <mask> [match-metric<value>|set-metric <value>]	RIP->Redistribution Filter

OSPF Redistribution Filters

OSPF redistribution filters control redistribution of routes from other protocols into OSPF. A device running OSPF can redistribute static routes, IP/RIP routes, and BG4P routes (if BGP4 is supported on the device) into OSPF.

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

Actions

OSPF redistribution filters permit (redistribute) or deny (don't redistribute) IP/RIP or BGP4 routes into OSPF.

Scope

You configure and apply OSPF redistribution filters globally.

Syntax

Use the following CLI commands or Web management interface panels to configure OSPF redistribution filters.

Table 2.24: OSPF Redistribution Filters

CLI syntax	Web management links
HP9300(config-ospf-router)# deny|permit redistribute <filter-num> all|bgp|rip|static address <ip address> [match-metric<value>|set-metric <value>]	OSPF->Redistribution Filter

BGP4 Redistribution Filters

BGP4 redistribution filters control redistribution of routes from other protocols into BGP4. A device running BGP4 can redistribute static routes, IP/RIP routes, and OSPF routes into BGP4.

Optionally, you can modify a route's metric and weight and use a route map to change additional attributes of the route.

Actions

BGP4 redistribution filters permit (redistribute) or deny (don't redistribute) IP/RIP or OSPF routes into IP/RIP.

Scope

You configure and apply BGP4 redistribution filters globally.

Syntax

Use the following CLI commands or Web management interface panels to configure BGP4 redistribution filters.

Table 2.25: BGP4 Redistribution Filters

CLI syntax	Web management links
HP9300(config-bgp-router)# redistribute <rip|ospf|static> [match <internal|external1|external2>] [metric <num>] [route-map <name>] [weight <num>]	BGP->Redistribute

NOTE: The optional **match internal|external1|external2** argument applies only to OSPF.

Layer 4 Filters

Layer 4 filters control IP traffic based on the Layer 3 and Layer 4 information in the packets. On switches and routing switches, Layer 4 filters are access policies that control access to Layer 4 applications based on TCP/UDP or other port number.

TCP/UDP Forwarding Filters

TCP/UDP forwarding filters are the same as TCP/UDP access policies. See “TCP/UDP Access Policies” on page D-9.