Chapter 6
Policy-Based Routing

Policy-Based Routing (PBR) provides a flexible mechanism for forwarding data packets based on polices configured by a network administrator.

This chapter contains the following sections:

Policy-Based Routing Overview

PBR enables you to implement policies that selectively cause packets to take different paths. PBR provides a routing mechanism for networks that rely on Application Layer support, such as antivirus (AV), deep inspection (DI), or anti-spam, web filtering, and/or that require an automatic way to specific applications.

When a packet enters the security device, ScreenOS checks for PBR as the first part of the route-lookup process, and the PBR check is transparent to all non-PBR traffic. PBR is enabled at the interface level and configured within a virtual router context; but you can choose to bind PBR policies to an interface, a zone, a virtual router (VR), or a combination of interface, zone, or VRs.

You use the following three building blocks to create a PBR policy:

- Extended access lists
- Match groups
- Action groups

Extended Access-Lists

Extended access-lists list the match criteria you define for PBR policies. PBR match criteria determine the path of a particular data traffic flow. Match criteria include the following:

- Source IP address
- Destination IP address
- Source port
- Destination port
- Protocol, such as HTTP
- Quality of Service (QoS) priority (optional)

Match Groups

Match groups provide a way to organize (by group, name and priority) extended access lists. Match groups associate an extended access-list ID number with a unique match group name and a match-group ID number. This match-group ID number defines the order in which you want the security device to process the extended ACL lists. You can assign multiple extended access-lists to the same match-group.

Action Groups

Action groups specify the route that you want a packet to take. You specify the “action” for the route by defining the next interface, the next-hop, or both.

Each configured action entry is monitored for reachability as follows:

NOTE: Monitoring reachability does not refer to Layer 3 tracking or Layer 2 Address Resolution Protocol (ARP) lookups.

- **Next-Interface Only Reachability**

  If you associate the action entry with only a next-interface, link state determines reachability.

  If the next-interface is up, the action entry is reachable. Any interface including all the logical interfaces, such as tunnel, aggregate, or redundant, that are visible in the VR in which the policy resides are candidates for next-interface.

  For example, if you configure the action entry with a NULL interface, the action entry is reachable all the time. With a NULL interface as the next interface, PBR lookup always succeeds; so, ScreenOS stops the route lookup and discards the packet(s).

- **Next-Hop Only Reachability**

  If you associate the action group with a next-hop only, that next-hop must be reachable through a route entry in the destination routes routing table. The configured next-hop is reachable as long as a valid route exists in the destination routes routing table to resolve the next-hop.

- **Next-Interface** and **Next-Hop Reachability**

  If you configure both next-interface and next-hop reachability, the configured next-hop must be reachable through the configured next-interface.

  If the next-hop is reachable through the next-interface, the action entry is reachable. Any interface including all the logical interfaces, such as tunnel, aggregate, or redundant, that are visible in the VR in which the policy resides are candidates to be a next-interface.

  If the next hop is reachable but the next interface is a NULL interface, ScreenOS drops the packet. If you configure the action entry with a NULL interface as the next interface and the next hop as a static route, ScreenOS passes the packet(s) to the static route.

At the time of configuration, you also assign a sequence number to specify the order in which you want the action group entry processed.

Route Lookup with Policy-Based Routing

When you enable policy-based routing on an interface, ScreenOS checks all traffic sent to that interface for policy-based routing. When a packet enters the security device, ScreenOS checks the in-interface for a PBR policy configuration. If PBR is enabled on that in-interface, the following actions are applied to the packet:

1. ScreenOS applies the PBR policy bound to the in-interface to the packet.

2. If no interface-level PBR policy exists, then ScreenOS applies the PBR policy bound to the zone associated with the in-interface to the packet.

3. If no zone-level PBR policy exists, then ScreenOS applies the PBR policy bound to the VR associated with the in-interface to the packet.

ScreenOS locates the match group and then processes the action group entries. The first reachable action entry from the action-group with a valid route is used to forward the packet. If no reachable route exists among the action entries, then a regular route lookup is performed.

If the action entry is reachable, ScreenOS performs a route lookup with the preferred interface as the next-interface (if specified) and the next-hop as the IP address (if specified) instead of using the destination IP. If a route matches the indicated next-interface and next-hop, ScreenOS forwards the packet. Otherwise, ScreenOS uses the destination IP address.

NOTE: For more information about route lookup, see *Volume 2: Fundamentals*.

Configuring Policy-Based Routing

Figure 19 shows one way PBR differentiates service-traffic paths by sending HTTP traffic along one path and HTTPS traffic along another. Figure 19 shows two nodes, one at 172.18.1.10 and another at 172.18.2.10. When the security device receives HTTP traffic, ScreenOS routes the traffic through the 172.24.76.1 router; and when the security device receives HTTPS traffic, ScreenOS routes the traffic through the 172.24.76.2 router.

The opposite is true for the 172.18.2.10 node. HTTP traffic from the 172.18.2.10 node flows to the 172.24.76.2 router, and HTTPS traffic flows to the 172.24.76.1 router.

Figure 19: Routing HTTP and HTTPS Traffic with Policy-Based Routing

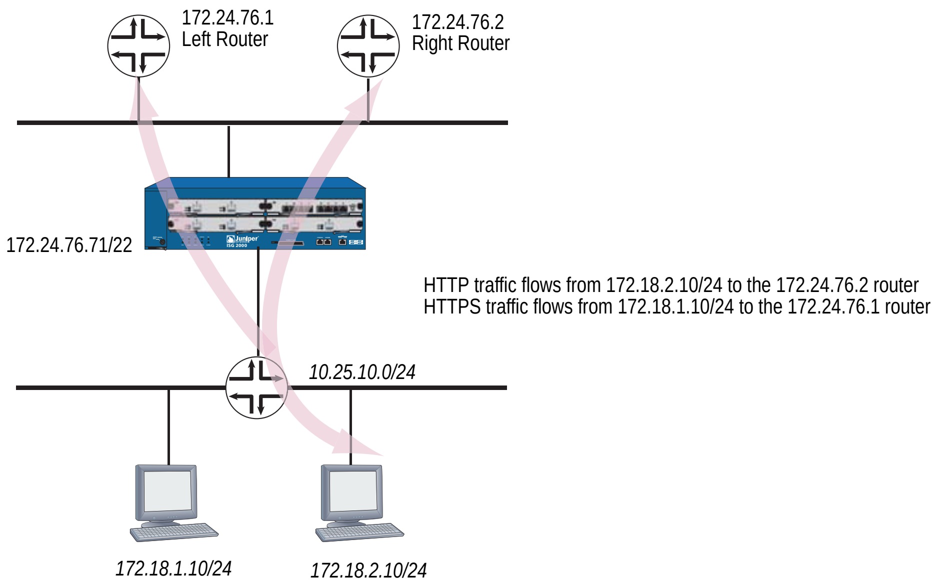
172.24.76.1
Left Router
172.24.76.2
Right Router
172.24.76.71/22
HTTP traffic flows from 172.18.2.10/24 to the 172.24.76.2 router
HTTPS traffic flows from 172.18.1.10/24 to the 172.24.76.1 router
10.25.10.0/24
172.18.1.10/24
172.18.2.10/24

Configuring an Extended Access List

You can configure an extended access list with the web user interface (WebUI) or the command line interface (CLI) from within a virtual router context. First, you configure the extended access list on the ingress virtual router (VR).

In this example on page 133, the ingress VR is the trust-vr. If you are using the CLI, you need to enter the virtual router context. This example requires two access lists: 10 and 20. The access sequence number is a number from 1 to 99. Entries 1 and 2 are required for each extended access list.

NOTE: Optionally, you can also add a type of service (TOS) number, which is a number from 1 to 255. A TOS number is not required in this example.

Access list 10 defines the source IP address as 172.18.1.10, the destination port as 80, and the protocol as TCP. The destination point for access list 10 defines the destination IP address as 172.18.2.10, the destination port as 443, and the protocol as TCP.

Access list 20 defines the source IP address as 172.18.2.10, the destination port as 443, and the protocol as TCP. The destination point for access list 10 defines the destination IP address as 172.18.1.10, the destination port as 80, and the protocol as TCP.

In the CLI after configuring the extended access list, you exit the virtual router context. The WebUI example only shows the creation of access list 10.

WebUI

Network > Routing > PBR > Extended ACL List: Select the virtual router from the dropdown list, then click **New** to view the Configuration page.

Enter the following information to create access list 10 entries:

Creating Access List 10

Extended ACL ID: 10
Sequence No.: 1
Source IP Address/Netmask: 172.18.1.10/32
Destination Port: 80-80
Protocol: TCP

Click **OK**. ScreenOS returns you to a list of access lists.

Click **New** to configure a second entry for access list 10 and enter the following information:

Creating Access List 10

Extended ACL ID : 10
Sequence No.: 2
Source IP Address/Netmask: 172.18.2.10/32
Destination Port: 443-443
Protocol: TCP

Click **OK**. ScreenOS returns you to a list of access lists.

CLI

```
set vrouter trust-vr
set access-list extended 10 src-ip 172.18.1.10/32 dest-port 80-80 protocol tcp
    entry 1
set access-list extended 10 src-ip 172.18.2.10/32 dest-port 443-443 protocol
    tcp entry 2
set access-list extended 20 src-ip 172.18.2.10/32 dest-port 80-80 protocol tcp
    entry 1
set access-list extended 20 src-ip 172.18.1.10/32 dest-port 443-443 protocol
    tcp entry 2
exit
```

Configuring a Match Group

You can configure a match group with the WebUI or the CLI from within a virtual router context.

In the example on page 133, you need to configure two match-groups: Left Router and Right Router. You bind extended access list 10 to Left Router and extended access list 20 to Right Router. A match group name is a unique identifier of no more than 31 alphanumeric characters.

The ingress VR is the trust-vr. If you are using the CLI, you need to enter the virtual router context. In the CLI after configuring the extended access list, you exit the virtual router context.

The WebUI example only shows the creation of a match group for Left Router.

WebUI

Network > Routing > PBR > Match Group > Select the correct virtual router from the dropdown list, then click **New** to view the Match Group Configuration page. Enter the following information to configure Left Router:

Match Group Name: left_router
Sequence No.: 1
Extended ACL: Select 10 from the dropdown list.

CLI

```
set vrouter trust-vr
set match-group name left_router
set match-group left ext-acl 10 match-entry 1
set match-group name right_router
set match-group right ext-acl 20 match-entry 1
exit
```

Configuring an Action Group

You can configure an action group with the WebUI or the CLI within a virtual routing context.

In the example on page 133 two different action groups are possible: the security device can forward to traffic to the left router or the right router. For this reason, you need to configure two different action groups.

To configure these two action-groups, you perform the following tasks:

1. Enter the virtual routing context. In this example, the virtual router is the trust-vr.

2. Name the action-group with a meaningful, unique name. In this example, the names **action-right** and **action-left** are descriptive of the possible traffic flows.

3. Configure the action-group details. In this example, you set the next-hop address for each action-group and then assign a number to indicate the processing priority. In this example, the priority of each action-group is 1.

WebUI

Network > Routing > PBR > Action Group > Click **New** to view the Configuration page

CLI

```
set vrouter trust-vr
set action-group name action-right
set action-group action-right next-hop 172.24.76.2 action-entry 1
set action-group name action-left
set action-group action-left next-hop 172.24.76.1 action-entry 1
exit
```

Configuring a PBR Policy

You can configure a PBR policy with the WebUI or the CLI from within a virtual router context.

Each PBR policy needs to have a unique name. In this example, the policy is named **redirect-policy**.

A PBR policy can contain a match group name and action group name. In this example, traffic can flow two different ways, so two different statements are required: **action-left** with sequence number 1 and **action-right** with sequence number 2. The policy statement with sequence number 1 is processed first.

WebUI

Network > Routing > PBR > Policy > Click **New** to view the Configuration page

CLI

```
set vrouter trust-vr
set pbr policy name redirect-policy
set pbr policy redirect-policy match-group left action-group action-left 1
set pbr policy redirect-policy match-group right action-group action-right 2
exit
```

Binding a Policy-Based Routing Policy

You can bind a PBR policy to an interface, a zone, or a virtual router with the WebUI or the CLI from within a virtual router context.

Binding a Policy-Based Routing Policy to an Interface

You can bind the PBR policy **redirect-policy** to the ingress interface. In this example, the interface is the **trust** interface.

WebUI

Network > Routing > PBR > Policy Binding

CLI

```
set interface trust pbr redirect-policy
```

Binding a Policy-Based Routing Policy to a Zone

You can bind the PBR policy **redirect-policy** to a zone. In this example, the zone is the **Trust** zone.

WebUI

Network > Routing > PBR > Policy Binding

CLI

```
set zone trust pbr redirect-policy
```

Binding a Policy-Based Routing Policy to a Virtual Router

You can bind the PBR policy **redirect-policy** to a virtual router. In this example, the virtual router is the **trust-vr**.

WebUI

Network > Routing > PBR > Policy Binding

CLI

set vrouter trust-vr pbr redirect-policy

Viewing Policy-Based Routing Output

You can view policy-based routing-related information with the WebUI or the CLI.

Viewing an Extended Access List

You can view the entire list of extended access lists from the WebUI or the CLI.

In the CLI you can specify to view one particular extended access list. In the second CLI example, the sample output shows that two extended access lists exist in the trust-vr, but the user indicated extended access list 2. As specified, ScreenOS returned two access-list entries, 10 and 20, for the second extended access list only.

WebUI

Network > Routing > PBR > Access List Ext

CLI 1

get vrouter trust-vr pbr access-list configuration

Sample output:

```
set access-list extended 1 src-ip 172.16.10.10/32 dest-ip 192.169.10.10/32
dest-port 80-80 protocol tcp entry 1
set access-list extended 1 src-port 200-300 entry 2
set access-list extended 2 dest-port 500-600 protocol udp entry 10
set access-list extended 2 dest-ip 50.50.50.0/24 protocol udp entry 20
```

CLI 2

get vrouter trust-vr pbr access-list 2

Sample output:

```
PBR access-list: 2 in vr: trust-vr, number of entries: 2
------------------------------------------------
PBR access-list entry: 10
------------------------
dest port range 500-600
protocols: udp
PBR access-list entry: 20
------------------------
dest ip-address 50.50.50.0/24
protocols: udp
```

Viewing a Match Group

You can view match group details from the WebUI or the CLI.

WebUI

Network > Routing > PBR > Match Group

CLI

get vrouter trust-vr pbr match-group config

Sample output:

```
set match-group name pbr1_mg
set match-group pbr1_mg ext-acl 1 match-entry 1
set match-group name pbr1_mg2
set match-group pbr1_mg2 ext-acl 2 match-entry 10
```

Viewing an Action Group

You can view action group details from the WebUI or the CLI.

WebUI

Network > Routing > PBR > Action Group

CLI 1

get vrouter trust-vr pbr action-group configuration

Sample output:

```
set action-group name pbr1_ag
set action-group pbr1_ag next-interface ethernet2 next-hop 10.10.10.2
    action-entry 1
set action-group name pbr1_ag2
set action-group pbr1_ag2 next-hop 30.30.30.30 action-entry 10
set action-group pbr1_ag2 next-interface ethernet3 action-entry 20
set action-group pbr1_ag2 next-interface ethernet3 next-hop 60.60.60.60
    action-entry 30
```

CLI 2

get vrouter trust-vr pbr match-group name pbr1_ag2

Sample output:

```
device-> get vr tr pbr action-group name pbr1_ag2
PBR action-group: pbr1_ag2 in vr: trust-vr number of entries: 3
------------------------------------------------
PBR action-group entry: 10
next-interface: N/A, next-hop: 30.30.30.30
------------------------
PBR action-group entry: 20
next-interface: ethernet3, next-hop: 0.0.0.0
------------------------
PBR action-group entry: 30
next-interface: ethernet3, next-hop: 60.60.60.60
------------------------
```

Viewing a Policy-Based Routing Policy Configuration

You can view policy-based routing policy configuration details from the WebUI or the CLI. In the CLI you can choose to view the configuration or you can enter the policy name to view a single policy configuration.

WebUI

Network > Routing > PBR > Policy

CLI

get vrouter trust-vr pbr policy config

Sample output:

```
set pbr policy name pbr1_policy
set pbr policy pbr1_policy match-group pbr1_mg2 action-group pbr1_ag2 50
set pbr policy pbr1_policy match-group pbr1_mg action-group pbr1_ag 256
```

CLI

get vrouter trust-vr pbr policy name pbr1_policy

Sample output:

```
PBR policy: pbr1_policy in vr: trust-vr number of entries: 2
------------------------------------------------
PBR policy entry: 50
match-group: pbr1_mg2, action-group: pbr1_ag2
------------------------
PBR policy entry: 256
match-group: pbr1_mg, action-group: pbr1_ag
------------------------
```

Viewing a Complete Policy-Based Routing Configuration

You can view a policy-based routing configuration from the WebUI or the CLI.

WebUI

Network > Routing > PBR > Access List Ext
Network > Routing > PBR > Match Group
Network > Routing > PBR > Action Group
Network > Routing > PBR > Policy

CLI

get vrouter trust-vr pbr configuration

Sample output:

```
set access-list extended 1 src-ip 172.16.10.10/32 dest-ip 192.169.10.10/32
dest-port 80-80 protocol tcp entry 1
set access-list extended 1 src-port 200-300 entry 2
set access-list extended 2 dest-port 500-600 protocol udp entry 10
set access-list extended 2 dest-ip 50.50.50.0/24 protocol udp entry 20
set match-group name pbr1_mg
set match-group pbr1_mg ext-acl 1 match-entry 1
set match-group name pbr1_mg2
set match-group pbr1_mg2 ext-acl 2 match-entry 10
set action-group name pbr1_ag
set action-group pbr1_ag next-interface ethernet2 next-hop 10.10.10.2
action-entry 1
set action-group name pbr1_ag2
set action-group pbr1_ag2 next-hop 30.30.30.30 action-entry 10
set action-group pbr1_ag2 next-interface ethernet3 action-entry 20
set action-group pbr1_ag2 next-interface ethernet3 next-hop 60.60.60.60
action-entry 30
set pbr policy name pbr1_policy
set pbr policy pbr1_policy match-group pbr1_mg2 action-group pbr1_ag2 50
set pbr policy pbr1_policy match-group pbr1_mg action-group pbr1_ag 256
```

Advanced PBR Example

PBR allows you to define and offload only the types of traffic that ScreenOS needs to process. In processing specific types of traffic, such as traffic requiring antivirus (AV) scanning, the network does not get bottlenecked by scanning packet types that do not need to be scanned for viruses.

NOTE: You could also configure PBR to send traffic specific for anti-spam, deep inspection (DI), intrusion detection and prevention (IDP), web filtering, or caching.

You can combine several types of Juniper Networks security devices to work together to provide services while keeping network processing speed fast and AV scanning manageable. Figure 20 shows a security device running PBR to segregate AV traffic from all other traffic (right).

Figure 20: Selective Routing by Traffic Type

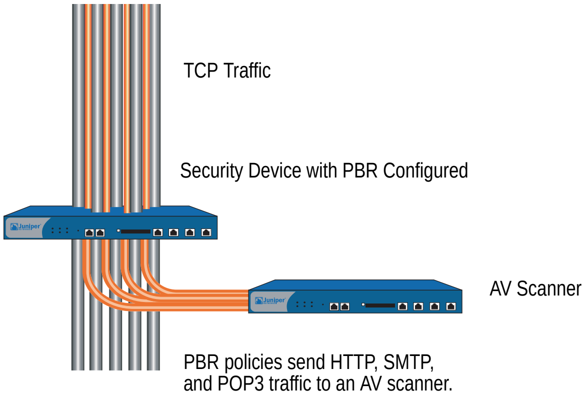
TCP Traffic
Security Device with PBR Configured
AV Scanner
PBR policies send HTTP, SMTP,
and POP3 traffic to an AV scanner.

For example, if you want use PBR to offload only HTTP, SMTP, and POP3 traffic for AV processing, at a minimum you need to use at least one security device with four available 10/100 interfaces to provide routing and one security device to provide the application (AV) support.

NOTE: If you have only three 10/100 interfaces available, you can place a switch between the two security devices and use VLAN tagging (802.1q) to set up the same paths for the traffic.

In the following example, you perform the following steps to set up the security device that provides the routing paths:

1. Configure routing.
2. Configure PBR.
3. Bind the PBR policies to the appropriate interfaces.

The next sections explain each of these steps. The examples show only CLI commands and output.

For information about configuring AV, see *Volume 4, Attack Detection and Defense Mechanisms*.

Routing

In this example, you need to create two custom zones:

- **av-dmz-1** for the trust-vr
- **av-dmz-2** for the untrust-vr

To set up the zones, enter the following commands:

```
set zone name av-dmz-1
set zone name av-dmz-2
```

Using the information shown in Table 15, you set up four 10/100 Ethernet interfaces.

Table 15: Interface Configuration for Routing

Interface Name	Zone	Virtual Router	IPv4 Address
E1	trust	trust-vr	10.251.10.0/24
E2	av-dmz-1	trust-vr	192.168.100.1/24
E3	av-dmz-2	untrust-vr	192.168.101.1/24
E4	untrust	untrust-vr	172.24.76.127/24

To set up the interfaces, enter the following commands:

```
set interface e1 zone trust vrouter trust-vr ip 10.251.10.0/24
set interface e2 zone av-dmz-1 vrouter trust-vr ip 192.168.100.1/24
set interface e3 zone av-dmz-2 vrouter untrust-vr ip 192.168.101.1/24
set interface e4 zone untrust vrouter untrust-vr ip 172.24.76.127/24
```

After setting up the zones, interfaces and routes, you need to perform the following two tasks:

1. Configure a static route from the untrust-vr to the trust-vr. Assign a gateway IP address of 10.251.10.0/24 and a preference value of 20 to the entry:

   set vrouter "untrust-vr"
   set route 10.251.10.0/24 vrouter "trust-vr" preference 20
   exit

2. Configure the NULL interface with a preference value greater than zero (0) from the Trust interface to the Untrust interface:

   set vrouter "trust-vr"
   set route 0.0.0.0/0 vrouter "untrust-vr" preference 20
   exit

You can verify the changes with the **get route** command:

```
Routing Table:
IPv4 Dest-Routes for <untrust-vr> (6 entries)
----------------------------------------------------------------------------
   ID         IP-Prefix      Interface         Gateway   P Pref    Mtr     Vsys
----------------------------------------------------------------------------
*   6         0.0.0.0/0           eth4     172.24.76.1   C    0      1     Root
*   3     10.251.10.0/24           n/a        trust-vr   S   20      0     Root
*   4     172.24.76.0/22          eth4         0.0.0.0   C    0      0     Root
*   2   192.168.101.1/32          eth3         0.0.0.0   H    0      0     Root
*   5   172.24.76.127/32          eth4         0.0.0.0   H    0      0     Root
*   1   192.168.101.0/24          eth3         0.0.0.0   C    0      0     Root

IPv4 Dest-Routes for <trust-vr> (5 entries)
------------------------------------------------------------
   ID         IP-Prefix      Interface         Gateway   P Pref    Mtr     Vsys
------------------------------------------------------------
*   5         0.0.0.0/0            n/a      untrust-vr   S   20      0     Root
*   1     10.251.10.0/24          eth1         0.0.0.0   C    0      0     Root
*   4   192.168.100.1/32          eth2         0.0.0.0   H    0      0     Root
*   3   192.168.100.0/24          eth2         0.0.0.0   C    0      0     Root
*   2     10.251.10.1/32          eth1         0.0.0.0   H    0      0     Root
```

You are now ready to configure PBR.

PBR Elements

After you configure the interfaces and routes, you configure PBR. For PBR to work correctly, you must configure the following items for the trust-vr:

- Extended access list
- Match group
- Action group
- PBR policy

Extended Access Lists

For this example, you determine that you want to send HTTP (port 80), SMTP (port 110), and POP3 (port 25) traffic for AV processing. To send these three types of packets to a security device, you set up an extended access list in the trust-vr.

NOTE: You do not need to set up an extended access list for the return traffic because the security device performs a session lookup before a route lookup and then applies a PBR policy as necessary. Return traffic has an existing session.

When any client in the 10.251.10.0/24 subnet initiates traffic that uses TCP to port 80, 110, or 25, you want ScreenOS to match that traffic to extended access list criteria and to perform the action associated with the access list. The action forces ScreenOS to route the traffic as you indicate and not like other traffic. Each access list needs three entries, one for each kind of TCP traffic that you are targeting.

To configure the extended access list for the trust-vr, enter the following commands:

```
set vrouter "trust-vr"
set access-list extended 10 src-ip 10.251.10.0/24 dest-port 80-80 protocol tcp
    entry 1
set access-list extended 10 src-ip 10.251.10.0/24 dest-port 110-110 protocol
    tcp entry 2
set access-list extended 10 src-ip 10.251.10.0/24 dest-port 25-25 protocol tcp
    entry 3
exit
```

Match Groups

A match group associates an extended access list with a meaningful name that gets referenced in the PBR policy. You first enter a virtual router context, then create a match group, and finally add an entry that associates the newly created match group name with an access list and entry number.

To create match groups in the trust-vr, enter the following commands:

```
set vrouter trust-vr
set match-group name av-match-trust-vr
set match-group av-match-trust-vr ext-acl 10 match-entry 1
exit
```

Action Group

Next, you create an action-group, which indicates where to send the packet. For this example, you create an action group for the trust-vr with the action set to send the traffic to the next hop.

CAUTION: If the action is to send traffic to the next interface, the link-state change will activate/deactivate the routing policy.

With next hop, the action resolves with Address Resolution Protocol (ARP).

For the trust-vr, you redirect traffic with the next hop statement through 192.168.100.254 by entering the following commands:

```
set vrouter trust-vr
set action-group name av-action-redirect-trust-vr set action-group
    av-action-redirect-trust-vr next-hop 192.168.100.254 action-entry 1
exit
```

PBR Policies

Next, you define the PBR policy, which requires the following elements:

- PBR policy name
- Match group name
- Action group name

To configure the PBR policy, enter the following commands:

```
set vrouter trust-vr
set pbr policy name av-redirect-policy-trust-vr
set pbr policy av-redirect-policy-trust-vr match-group av-match-trust-vr
    action-group av-action-redirect-trust-vr 1
exit
```

Interface Binding

Finally, you bind the PBR policy to the ingress interface, e1.

To bind the PBR policy to its ingress interface, enter the following commands:

```
set interface e1 pbr av-redirect-policy-trust-vr
```

Advanced PBR with High Availability and Scalability

Using the previous PBR example as a foundation, you can add resilience to your network with high availability (HA) and/or scalability.

Resilient PBR Solution

A robust PBR solution might include the following device configurations:

- Two security devices that provide networking
- Two other security devices that provide AV scanning

Each pair of devices runs NetScreen Redundancy Protocol (NSRP) in an active/passive configuration to provide failover protection. For the two security devices that are performing routing, one device takes over the routing function if a hardware failure occurs. In the case of the pair that is providing the AV scanning, if a failure occurs in one of the devices, the other device takes over the scanning function.

NOTE: For more information, see *Volume 11: High Availability*.

Scalable PBR Solution

PBR solutions scale well. If you need more capacity, you can add more security devices. By dividing the /24 subnet into two /25 subnets, you can configure one extended access list for the lower /25 subnet and another extended access list for the higher /25 subnet, then add two security devices to provide scanning services in the DMZ.

You can also implement load balancing if you create an active/active NSRP configuration. One device could process traffic from the lower /25 subnet, and the other device could process traffic from the higher /25 subnet. Each device backs up the other.

Chapter 7
Multicast Routing

This chapter introduces basic multicast routing concepts. It contains the following sections:

Overview

Enterprises use multicast routing to transmit traffic, such as data or video streams, from one source to a group of receivers simultaneously. Any host can be a source, and the receivers can be anywhere on the Internet.

IP multicast routing provides an efficient method for forwarding traffic to multiple hosts because multicast-enabled routers transmit multicast traffic only to hosts that want to receive the traffic. Hosts must signal their interest in receiving multicast data, and they must join a multicast group in order to receive the data. Multicast-enabled routers forward multicast traffic only to receivers interested in receiving the traffic.