

Policy Based Routing

This chapter describes how to configure the ASA to support policy based routing (PBR). The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- About Policy Based Routing, on page 1
- Guidelines for Policy Based Routing, on page 3
- Path Monitoring, on page 5
- Configure Policy Based Routing, on page 6
- History for Policy Based Routing, on page 9

About Policy Based Routing

Traditional routing is destination-based, meaning packets are routed based on destination IP address. However, it is difficult to change the routing of specific traffic in a destination-based routing system. With Policy Based Routing (PBR), you can define routing based on criteria other than destination network—PBR lets you route traffic based on source address, source port, destination address, destination port, protocol, or a combination of these.

Policy Based Routing:

- Lets you provide Quality of Service (QoS) to differentiated traffic.
- Lets you distribute interactive and batch traffic across low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths.
- Allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections.

Policy Based Routing can implement QoS by classifying and marking traffic at the network edge, and then using PBR throughout the network to route marked traffic along a specific path. This permits routing of packets originating from different sources to different networks, even when the destinations are the same, and it can be useful when interconnecting several private networks.

Why Use Policy Based Routing?

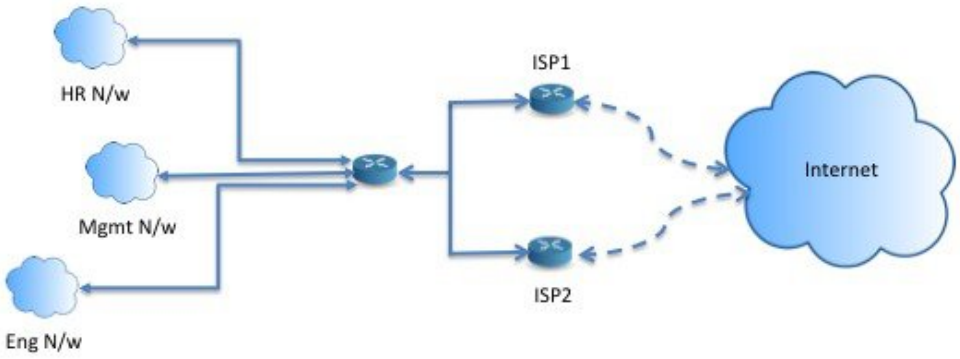
Consider a company that has two links between locations: one a high-bandwidth, low-delay expensive link, and the other a low-bandwidth, higher-delay, less-expensive link. While using traditional routing protocols, the higher-bandwidth link would get most, if not all, of the traffic sent across it based on the metric savings

obtained by the bandwidth and/or delay (using EIGRP or OSPF) characteristics of the link. PBR allows you to route higher priority traffic over the high-bandwidth/low-delay link, while sending all other traffic over the low-bandwidth/high-delay link.

Some applications of policy based routing are:

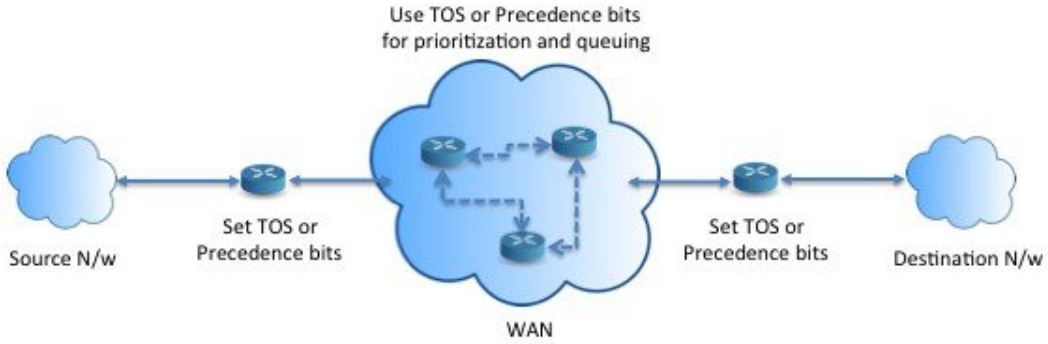
Equal-Access and Source-Sensitive Routing

In this topology, traffic from HR network & Mgmt network can be configured to go through ISP1 and traffic from Eng network can be configured to go through ISP2. Thus, policy based routing enables the network administrators to provide equal-access and source-sensitive routing, as shown here.



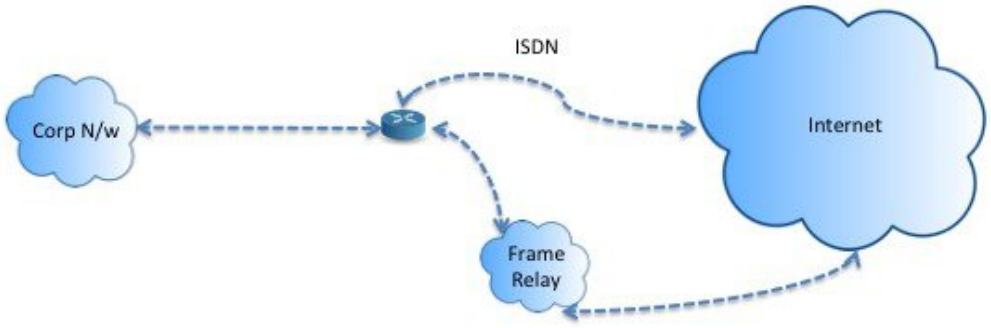
Quality of Service

By tagging packets with policy based routing, network administrators can classify the network traffic at the perimeter of the network for various classes of service and then implementing those classes of service in the core of the network using priority, custom or weighted fair queuing (as shown in the figure below). This setup improves network performance by eliminating the need to classify the traffic explicitly at each WAN interface in the core of backbone network.



Cost Saving

An organization can direct the bulk traffic associated with a specific activity to use a higher-bandwidth high-cost link for a short time and continues basic connectivity over a lower-bandwidth low-cost link for interactive traffic by defining the topology, as show here.



Load Sharing

In addition to the dynamic load-sharing capabilities offered by ECMP load balancing, network administrators can now implement policies to distribute traffic among multiple paths based on the traffic characteristics.

As an example, in the topology depicted in the Equal-Access Source Sensitive Routing scenario, an administrator can configure policy based routing to load share the traffic from HR network through ISP1 and traffic from Eng network through ISP2.

Implementation of PBR

The ASA uses ACLs to match traffic and then perform routing actions on the traffic. Specifically, you configure a route map that specifies an ACL for matching, and then you specify one or more actions for that traffic. Finally, you associate the route map with an interface where you want to apply PBR on all incoming traffic.

Note Before proceeding with configuration, ensure that the ingress and egress traffic of each session flows through the same ISP-facing interface to avoid unexpected behavior caused by asymetirc routing, specifically when NAT and VPN are in use.

Guidelines for Policy Based Routing

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Per-flow Routing

Since the ASA performs routing on a per-flow basis, policy routing is applied on the first packet and the resulting routing decision is stored in the flow created for the packet. All subsequent packets belonging to the same connection simply match this flow and are routed appropriately.

PBR Policies Not Applied for Output Route Look-up

Policy Based Routing is an ingress-only feature; that is, it is applied only to the first packet of a new incoming connection, at which time the egress interface for the forward leg of the connection is selected. Note that PBR

will not be triggered if the incoming packet belongs to an existing connection, or if NAT is applied and NAT chooses the egress interface.

PBR Policies Not Applied for Embryonic Traffic



Note An embryonic connection is where the necessary handshake between source and destination has not been made.

When a new internal interface is added and a new VPN policy is created using a unique address pool, PBR is applied to the outside interface matching the source of the new client pool. Thus, PBR sends traffic from the client to the next hop on the new interface. However, PBR is not involved in the return traffic from a host that has not yet established a connection with the new internal interface routes to the client. Thus, the return traffic from the host to the VPN client, specifically, the VPN client response is dropped as there is no valid route. You must configure a weighted static route with a higher metric on the internal interface.

Clustering

- Clustering is supported.
- In a cluster scenario, without static or dynamic routes, with ip-verify-reverse path enabled, asymmetric traffic may get dropped. So disabling ip-verify-reverse path is recommended.

IPv6 Support

IPv6 is supported

Path Monitoring Guidelines

Following are the guidelines for configuring the path monitoring on the interfaces:

- Interfaces must have an interface name.
- Management-only interfaces cannot be configured with the path monitoring. To configure the path monitoring, you must uncheck the **Dedicate this interface to management only** check box.
- Path monitoring is not supported on devices in Transparent or multicontext system mode.
- Auto monitoring types (auto, auto4, and auto6) are not supported for Tunnel interfaces.
- Path monitoring cannot be configured for the following interfaces:
  - BVI
  - Loopback
  - DVTI

Additional Guidelines

- All existing route map related configuration restrictions and limitations will be carried forward.

- Do not use route maps containing match policy lists for policy based routing. The match policy-list is only used for BGP.
- Unicast Reverse Path Forwarding (uRPF) validates the source IP address of packets received on an interface against the routing table and not against the PBR route map. When uRPF is enabled, packets received on an interface through PBR are dropped as they are without the specific route entry. Hence, when using PBR, ensure to disable uRPF.

Path Monitoring

Path monitoring, when configured on interfaces, derive metrics such as round trip time (RTT), jitter, mean opinion score (MOS), and packet loss per interface. These metrics are used to determine the best path for routing PBR traffic.

The metrics on the interfaces are collected dynamically using ICMP probe messages to the interface's default gateway or a specified remote peer.

Default Monitoring Timers

For metric collection and monitoring, the following timers are used:

- The interface monitor average interval is 30 seconds. This interval indicates the frequency to which the probes average.
- The interface monitor update interval is 30 seconds. This interval indicates the frequency at which the average of the collected values are calculated and made available for PBR to determine the best routing path.
- The interface monitor probe interval by ICMP is one second. This interval indicates the frequency at which an ICMP ping is sent.
- The application monitor probe interval by HTTP is 10 seconds. This interval indicates the frequency at which an HTTP ping is sent. Path monitoring uses the last 30 samples of HTTP ping for calculating the average metrics.



Note You cannot configure or modify the interval for any of these timers.

Typically, in PBR, traffic is forwarded through egress interfaces based on the priority value (interface cost) configured on them. From management center version 7.2, PBR uses IP-based path monitoring to collect the performance metrics (RTT, jitter, packet-lost, and MOS) of the egress interfaces. PBR uses the metrics to determine the best path (egress interface) for forwarding the traffic. Path monitoring periodically notifies PBR about the monitored interface whose metric got changed. PBR retrieves the latest metric values for the monitored interfaces from the path monitoring database and updates the data path.

Path monitoring functions only with dynamic metrics, and only if the RTT, jitter, packet-lost, or MOS variables are set on the interfaces. Path monitoring does not function with static metrics—interface cost (cost set in interface).

You must enable path monitoring for the interface and configure the monitoring type. The PBR policy page allows you to specify the desired metric for path determination. See Configure Policy Based Routing, on page 6.

Configure Path Monitoring

You can configure path monitoring to perform Policy Based Routing based on the network service groups. To use path monitoring without NSG, you can navigate to the **Interface** > **Edit** page and specify the path monitoring type. See Configure Policy Based Routing.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Interface Settings** > **Path Monitoring**.

Step 2 Select interface from **Interface** drop-down.

Step 3 Select the network service group (NSG) in the **Available Network Service Groups** box. To select multiple NSGs, use the control key and click on the required NSGs.

Step 4 Click **Add** to add the Network Service Groups.

Step 5 Click **Apply**.

Step 6 To remove the configuration, select the NSGs from the **Added Network Service Groups** box and click **Remove**, and then click **Apply**.

Configure Policy Based Routing

A route map is comprised of one or more route-map statements. Each statement has a sequence number, as well as a permit or deny clause. Each route-map statement contains match and set commands. The match command denotes the match criteria to be applied on the packet. The set command denotes the action to be taken on the packet.

- When a route map is configured with both IPv4 and IPv6 match/set clauses or when a unified ACL matching IPv4 and IPv6 traffic is used, the set actions will be applied based on destination IP version.
- When multiple next-hops or interfaces are configured as a set action, all options are evaluated one after the other until a valid usable option is found. No load balancing will be done among the configured multiple options.
- The verify-availability option is not supported in multiple context mode.

Procedure

Step 1 In ASDM, configure one or more standard or extended ACLs to identify traffic on which you want to perform Policy Based Routing. See **Configuration** > **Firewall** > **Advanced** > **ACL Manager** .

Step 2 Choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**, and click **Add**.

The **Add Route Map** dialog box appears.

Step 3 Enter the route map name and sequence number. You will use this same name for optional additional route map statements. The sequence number is the order in which the ASA assesses the route maps.

Step 4 Click **Deny** or **Permit**.

The ACL also includes its own permit and deny statements. For Permit/Permit matches between the route map and the ACL, the Policy Based Routing processing continues. For Permit/Deny matches, processing ends for this route map, and other route maps are checked. If the result is still Permit/Deny, then the regular routing table is used. For Deny/Deny matches, the Policy Based Routing processing continues.

Step 5 Click the **Match Clause** tab to identify the ACLs you created.

In the **IPv4** section, choose **Access List** from the drop-down menu, and then select one or more standard or extended ACLs from the dialog box.

Note
Ensure that the access list does not contain any inactive rules. You cannot set match ACL with inactive rules to a PBR.

If you use a standard ACL, matching is done on the destination address only. If you use an extended ACL, you can match on source, destination, or both.

Use the IPv4 section for both IPv4 and IPv6 ACLs. For the extended ACL, you can specify IPv4, IPv6, Identity Firewall, or Cisco TrustSec parameters. You can also include network-service objects. For complete syntax, see the ASA command reference.

Step 6 Click the **Policy Based Routing** tab to define policy for traffic flows.

Check one or more of the following set actions to perform for the matching traffic flows:

- **Set PBR next hop address**—For IPv4 and IPv6, you can configure multiple next-hop IP addresses in which case they are evaluated in the specified order until a valid routable next-hop IP address is found. The configured next-hops should be directly connected; otherwise the set action will not be applied.
- **Set default next-hop IP address**—For IPv4 and IPv6, if the normal route lookup fails for matching traffic, then the ASA forwards the traffic using this specified next-hop IP address.
- **Recursively find and set next-hop IP address—**Both the next-hop address and the default next-hop address require that the next-hop be found on a directly connected subnet. With this option, the next-hop address does not need to be directly connected. Instead a recursive lookup is performed on the next-hop address, and matching traffic is forwarded to the next-hop used by that route entry according to the routing path in use on the router.
- **Configure Next Hop Verifiability**—Verify if the next IPv4 hops of a route map are available. You can configure an SLA monitor tracking object to verify the reachability of the next-hop. Click **Add** to add next-hop IP address entries, and specify the following information.
  - **Sequence Number**—Entries are assessed in order using the sequence number.
  - **IP Address**—Enter the next hop IP address.
  - **Tracking Object ID**—Enter a valid ID.
- **Set interfaces**—This option configures the interface through which the matching traffic is forwarded. You can configure multiple interfaces, in which case they are evaluated in the specified order until a valid interface is found. When you specify **null0**, all traffic matching the route map will be dropped. There must be a route for the destination that can be routed through the specified interface (either static or dynamic).

- **Set Clause** > **Adaptive Interface Cost**—This option is on the Set Clause tab rather than the **Policy Based Routing** tab. This option sets the output interface based on the interface's cost. Click the **Available Interfaces** field and select the interfaces that should be considered. The egress interface is selected from the list of interfaces. If the costs of the interfaces are the same, it is an active-active configuration and packets are load-balanced (round-robin) on the egress interfaces. If the costs are different, the interface with the lowest cost is selected. Interfaces are considered only if they are up.
- **Set null0 interface as the default interface**—If a normal route lookup fails, the ASA forwards the traffic null0, and the traffic will be dropped.
- **Set do-not-fragment bit to either 1or 0**—Select the appropriate radio button.
- **Set differential service code point (DSCP) value in QoS bits**—Select a value from the IPv4 or IPv6 drop-down list.

Step 7 Click **OK**, and then click **Apply**.

Step 8 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and configure the ingress interfaces that should apply this route map to determine the egress interfaces.

a) Select an ingress interface and click **Edit**.
b) In **Route Map**, select the policy-based route map that should be applied.
c) If you used **Adaptive Interface Cost** to select the output interface in the route map, set the **Cost** value on the interface.

   The value can be 1-65535. The default is 0, which you can reset by deleting the value from this field. The lower the number, the higher the priority. For example, 1 has priority over 2.

d) For PBR to use flexible metrics in identifying the best path for routing packets, from the **Path Monitoring** drop-down list, select the relevant monitoring type:

   - **auto**—Sends ICMP probes to the IPv4 default gateway of the interface, if it exists (same as Auto IPv4). Else, sends to the IPv6 default gateway of the interface (same as Auto IPv6).
   - **ipv4**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, the adjacent field is enabled. Enter the IPv4 address in the field.
   - **ipv6**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, the adjacent field is enabled. Enter the IPv4 address in the field.
   - **auto4**—Sends ICMP probes to the IPv4 default gateway of the interface.
   - **auto6**—Send ICMP probes to the IPv6 default gateway of the interface.
   - **None**—To disable path monitoring for the interface.

e) Click **OK**, then **Apply**.

History for Policy Based Routing

Table 1: History for Route Maps

Feature Name	Platform Releases	Feature Information
Path monitoring through HTTP client	9.20(1)	PBR can now use the performance metrics (RTT, jitter, packet-lost, and MOS) collected by path monitoring through HTTP client on the application domain rather than the metrics on a specific destination IP. HTTP based path-monitoring can be configured on the interface using Network Service Group objects. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Path Monitoring**
Path monitoring metrics in PBR.	9.18(1)	PBR uses the metrics to determine the best path (egress interface) for forwarding the traffic. Path monitoring periodically notifies PBR with the monitored interface whose metric got changed. PBR retrieves the latest metric values for the monitored interfaces from the path monitoring database and updates the data path. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**
Policy based routing	9.4(1)	Policy Based Routing (PBR) is a mechanism by which traffic is routed through specific paths with a specified QoS using ACLs. ACLs let traffic be classified based on the content of the packet's Layer 3 and Layer 4 headers. This solution lets administrators provide QoS to differentiated traffic, distribute interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths, and allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections. We updated the following screens: Configuration > Device Setup > Routing > Route Maps > Policy Based Routing, Configuration > Device Setup > Routing > Interface Settings > Interfaces
IPv6 support for Policy Based Routing	9.5(1)	IPv6 addresses are now supported for Policy Based Routing. We modified the following screens: **Configuration > Device Setup > Routing > Route Maps > Add Route Map > Policy Based Routing** **Configuration > Device Setup > Routing > Route Maps > Add Route Maps > Match Clause**

Feature Name	Platform Releases	Feature Information
VXLAN support for Policy Based Routing	9.5(1)	You can now enable Policy Based Routing on a VNI interface. We modified the following screen: **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General**
Policy Based Routing support for Identity Firewall and Cisco Trustsec	9.5(1)	You can configure Identity Firewall and Cisco TrustSec and then use Identity Firewall and Cisco TrustSec ACLs in Policy Based Routing route maps. We modified the following screen: **Configuration > Device Setup > Routing > Route Maps > Add Route Maps > Match Clause**