

Policy-Based Routing

This chapter describes how to configure Firewall Threat Defense to support policy-based routing (PBR) through Firewall Management Center's policy-based routing page. The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- About policy-based routing, on page 1
- Licenses for policy-based routing, on page 3
- Best practices for policy-based routing, on page 3
- Path monitoring metrics for best path route determination, on page 5
- Configure policy-based routing policy, on page 8
- Configure policy-based routing, on page 11
- Configure PBR with path monitoring, on page 16
- Useful CLIs for monitoring PBR, on page 18
- Troubleshoot PBR, on page 21
- History for policy-based routing, on page 23

About policy-based routing

Policy-based routing (PBR) extends traditional routing protocols to provide greater control over traffic flow. It enables routing decisions based on criteria beyond destination IP, such as source and destination ports, protocols, and applications. Furthermore, PBR facilitates secure, application-specific traffic breakout in large-scale network deployments.

Policy-based routing characteristics

In traditional routing, packets are routed based on the destination IP address. Changing the routing of specific traffic in a destination-based routing system is difficult. PBR extends and complements the mechanisms provided by routing protocols.

PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link. With PBR, you can define routing that is based on criteria other than destination network such as source port, destination address, destination port, protocol, applications, or a combination of these objects.

You can use PBR to classify the network traffic based on applications, username, group membership, and security group association. This routing method applies to scenarios involving numerous devices accessing applications and data in large network deployment. In large deployments, network traffic is typically backhauled to a hub as encrypted traffic in a route-based VPN. These topologies often result in issues such as packet

latency, reduced bandwidth, and packet drop. Resolving these issues requires costly, complex deployments and management.

PBR policy enables you to securely break out traffic for specified applications. You can configure PBR policy in the Secure Firewall Management Center user interface to allow the applications to be directly accessed.

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 gets most, if not all, of the traffic sent across it based on the metric savings obtained by the bandwidth, delay, or both (using EIGRP or OSPF) characteristics of the link. With PBR, you can route higher priority traffic over the high-bandwidth/low-delay link, while sending all other traffic over the low-bandwidth/high-delay link.

These scenarios demonstrate policy-based routing applications:

- **Direct internet access**: In this topology, application traffic from the branch office can be routed directly to the internet instead of through the VPN tunnel connecting to the headquarters. The branch Firewall Threat Defense is configured with an internet exit point. The PBR policy is applied on the ingress interface (*Inside 1*) to identify the traffic based on the applications, user identity (username and group membership), and Security Group Tag (security group association) defined in the ACL. Correspondingly, the traffic is forwarded through the egress interfaces directly to the internet or to the IPsec VPN tunnel.
- **Equal-access and source-sensitive routing**: In this topology, traffic from the HR and Mgmt networks 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, in this example.
- **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 route the traffic from HR network through ISP1 and traffic from Eng network through ISP2 and thus share the load.

Figure 1: Direct internet access scenario

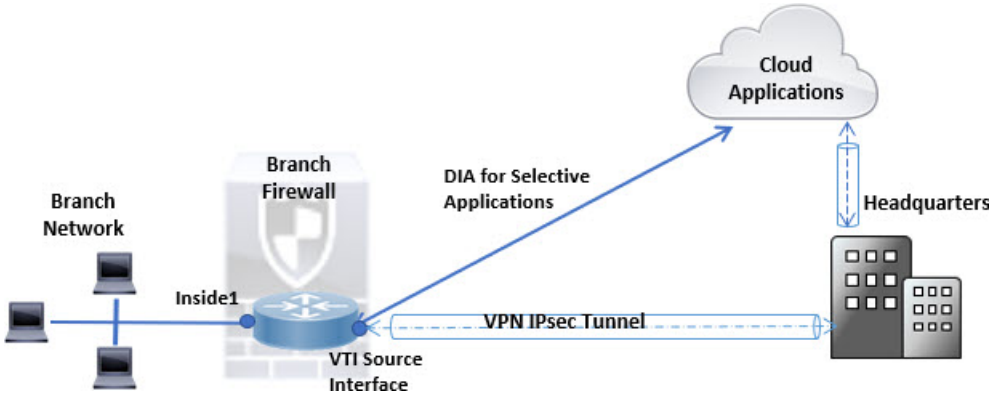
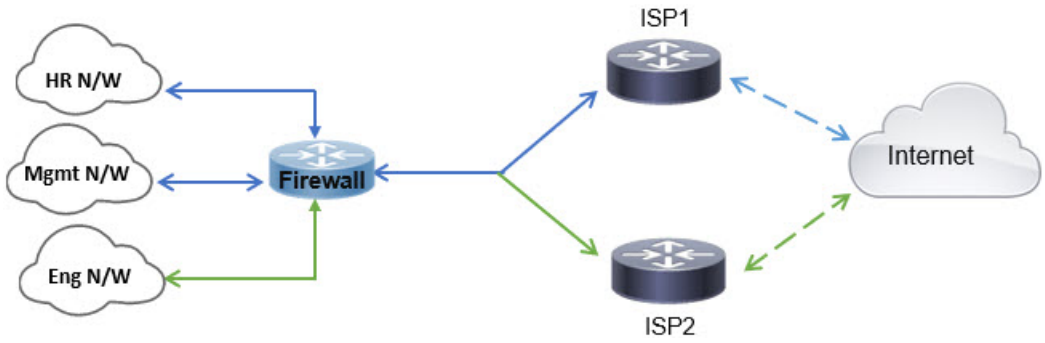


Figure 2: Equal access scenario



Licenses for policy-based routing

Policy-based routing (PBR) has specific licensing requirements depending on the features and services used in your configuration.

- All PBR policy features are supported with the Essentials Firewall Threat Defense smart license.
- Secure Client licensing is required if you are configuring PBR for routing traffic through RA VPN tunnel interface.
- Cisco ISE license is required for ISE usage in PBR policy.

Best practices for policy-based routing

Device guidelines

Ensure proper version compatibility and device-specific configurations when implementing PBR.

- PBR through Firewall Management Center's policy-based routing page is supported only from version 7.1 or later on both the Firewall Management Center and the Firewall Threat Defense device.
- When you upgrade Firewall Management Center or Firewall Threat Defense to version 7.1 or later, the PBR configuration in the device is removed. You must configure PBR again using the PBR page. If the managed device is earlier than version 7.1, you must configure PBR again using FlexConfig with deploy option set to "every time."
- On cluster devices, do not configure application, user identity, and Security Group Tag (SGT) based PBR policies, because they are not supported on cluster devices.

Interface guidelines

Configure appropriate interface types and settings for PBR policies.

- Configure only routed-mode interfaces and non management-only interfaces in the Global virtual router as ingress or egress interfaces for the PBR policy. You cannot configure user-defined virtual router interfaces for the policy.
- The interfaces that you want to define in the policy must have a logical name.

- Configure Static VTIs only as egress interfaces.
- Do not choose Dynamic VTIs for configuring the PBR egress interfaces.

You can apply PBR for managing both IPv4 and IPv6 traffic.

Application-based PBR and DNS configuration

Ensure proper DNS configuration for application-based PBR functionality.

- Application-based PBR uses DNS snooping for application detection. Application detection succeeds only if the DNS requests pass through Firewall Threat Defense in a clear-text format; the DNS traffic is not encrypted.
- You must configure trusted DNS servers for the application detection to succeed.

For more information on configuring DNS servers, refer to Configure DNS server settings.

PBR policies not applied for output route look-up

PBR is an ingress-only feature; The system applies PBR 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

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 matches the source of the new client pool. Thus, PBR sends traffic from the client to the next hop on the new interface. However, PBR does not process return traffic from a host until the new internal interface establishes a connection with 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. To prevent the response from being dropped, configure a weighted static route with a higher metric on the internal interface.

HTTP-based path monitoring guidelines

Configure HTTP-based path monitoring with appropriate interface types and version considerations.

- Configure HTTP-based path monitoring only on physical, port-channel, subinterfaces, and static tunnel interfaces. Do not configure it on cluster devices.
- HTTP uses only IPv4 to ping applications. IPv4 metrics are applied for routing and forwarding both IPv4 and IPv6 traffic.
- HTTP-based application monitoring is enabled by default in Secure Firewall Management Center version 7.4 and higher. However, when you upgrade from previous versions, this option is not enabled by default. You must manually enable it.

Additional guidelines

Follow these additional configuration considerations for effective PBR implementation.

- All existing configuration restrictions and route map limitations remain in effect.
- When you define the ACL for the policy match criteria, you can select multiple predefined applications from a list to form an Access Control Entry (ACE). In Firewall Threat Defense, the predefined applications

are stored as Network Service objects and the group of applications as Network Service Groups (NSG). You can create a maximum of 1024 such NSGs. The application or network service group is detected through first-packet classification. Currently, you cannot add to or modify the predefined applications list.

- 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. Therefore, when using PBR, ensure you disable uRPF.

Path monitoring metrics for best path route determination

Path monitoring metrics for best path route determination is a dynamic routing method that uses various monitoring approaches, such as ICMP and HTTP, to collect performance metrics and determine the best path.

Path monitoring methods and operation

PBR uses either static cost or path monitoring (dynamic metrics) to route its traffic.

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.

ICMP-based path monitoring method

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

HTTP-based path monitoring method

Path monitoring calculates dynamic metrics for each remote peer associated with an interface. HTTP is preferred over ICMP when monitoring multiple applications and determining the best path through a policy on a branch firewall, because:

- HTTP-ping can derive the performance metrics of the path up to the application layer of the server, where the application is hosted.
- Because the application domain is tracked instead of the IP address, you do not need to change the firewall configuration when the application server IP address changes.



Note You can configure both ICMP and HTTP on the same interface. If the destination in the policy matches to any of the domain IP, the corresponding metrics are used. If the destination does not match to any of the configured domains, the metrics from the ICMP are used by PBR to select the outgoing interface.

Default monitoring timers for metric collection

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.

PBR and path monitoring integration

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 policy, on page 8.

PBR and HTTP-based path monitoring

From management center version 7.4, PBR can be configured to use HTTP-based path monitoring to collect the performance metrics of the application domains and not just one destination IP address. Path monitoring begins only after a DNS entry for a domain is detected; it does not start immediately when HTTP-based application monitoring is configured. After obtaining the resolved IP address for the domain, path monitoring sends an HTTP request and receives a response. If DNS resolves multiple IP addresses for a domain, path monitoring probes and monitors the application using the first resolved IP address. Path monitoring continues until the IP address changes or until HTTP-based monitoring is disabled.

Based on the HTTP request and response durations, path monitoring computes the performance metrics for the application. Path monitoring sends collected metrics to PBR at regular intervals so that PBR can make routing and forwarding decisions for traffic from the configured ingress interface. If traffic arrives before path monitoring sends metrics to PBR, the routing table determines the traffic flow. Once metrics are available, PBR uses them to make routing decisions for subsequent traffic.



Note Based on the Network Service Groups in the match ACL of the policy, you can apply PBR for multiple domains having multiple IP addresses.

The management center associates applications and NSGs with the egress interface only when the PBR configuration meets these criteria:

- The match ACL contains the monitored applications.
- The PBR policy is configured with any one of the interface ordering values (metric type):
  - Minimal Jitter
  - Maximum Mean Opinion Score
  - Minimal Round-Trip Time
  - Minimal Packet Loss

Configure path monitoring settings

Configure path monitoring to collect metrics such as RTT, jitter, MOS, and packet loss on specified interfaces for effective traffic management and routing optimization.

The PBR policy relies on flexible metrics, such as round-trip time (RTT), jitter, mean opinion score (MOS), and packet loss of the interfaces to identify the best routing path for its traffic. Path monitoring collects these metrics on the specified interfaces. On the **Interfaces** page, you can configure path monitoring settings on interfaces to send ICMP probes or HTTP pings to collect metrics.

Procedure

Step 1 Choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Path Monitoring** tab.

Step 4 To configure ICMP based monitoring of the interface, click the **Enable IP based Monitoring** check box.

Step 5 From the **Monitoring Type** drop-down list, select the relevant option:

- **Auto**—Sends ICMP probes to the IPv4 default gateway of the interface. If the IPv4 gateway does not exist, path monitoring sends the probes to the IPv6 default gateway of the interface.
- **Peer IPv4**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, enter the IPv4 address in the **Peer IP To Monitor** field.
- **Peer IPv6**—Sends ICMP probes to the specified peer IPv6 address (next-hop IP) for monitoring. If you select this option, enter the IPv6 address in the **Peer IP To Monitor** field.
- **Auto IPv4**—Sends ICMP probes to the default IPv4 gateway of the interface.
- **Auto IPv6**—Send ICMP probes to the default IPv6 gateway of the interface.

Note

- The Auto options are not available for VTI interfaces. You must specify the peer address.
- Only one next-hop is monitored to a destination. That is, you cannot specify more than one peer address to monitor for an interface.

Step 6 By default, the **Enable HTTP based Application Monitoring** check box is selected. All applications selected for path monitoring in the match ACL of the PBR policy are listed, provided this interface is configured as the egress interface in the policy. To disable the HTTP-based monitoring of the interface, clear the check box.

Step 7 Click **Ok**.

Step 8 To save the settings, click **Save**.

Path monitoring is configured on the interface with the specified settings. The system will collect metrics using ICMP probes or HTTP pings based on your configuration to optimize routing paths.

Add path monitoring dashboard

Add the Path Monitoring Dashboard to enable viewing path monitoring metrics with customizable display options through predefined correlations and widget management.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Select the device, and click **Add new dashboard** (+).

Step 3 Enter a name for the custom dashboard.

Step 4 In the **Metrics** area, click the **Add from predefined correlations** button.

Step 5 From the list, click **Interface - Path metrics**.

By default, all four metrics and an additional metric field are selected to display as widgets in the dashboard.

To exclude any, click **Delete** (🗑).

Step 6 Click **Add dashboard**.

Configure policy-based routing policy

This task enables you to configure Policy-Based Routing (PBR) to control traffic routing based on specific criteria rather than relying solely on the routing table.

You can configure the PBR policy on the Policy Based Routing page by specifying the ingress interfaces, match criteria (Extended Access Control List), and egress interfaces.

Before you begin

To use the path monitoring metrics for configuring the traffic forwarding priority over egress interfaces, you must configure the path monitoring settings for the interfaces. See Configure path monitoring settings, on page 7.

Follow these steps to configure policy-based routing policy:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Policy Based Routing**.

On the Policy Based Routing page, you can view the configured policy. The grid shows ingress interfaces and the combination of the policy-based route access list and egress interfaces.

Step 4 To configure the policy, click **Add**.

Step 5 In the **Add Policy Based Route** dialog box, select the **Ingress Interface** from the drop-down list.

Note
You can only select interfaces with logical names that belong to a global virtual router from the drop-down list. You cannot configure a VLAN interface with a logical name as the source (ingress) interface.

Step 6 To specify the match criteria and the forward action in the policy, click **Add**.

Step 7 In the **Add Forwarding Actions** dialog box, do the following:

a) From the **Match ACL** drop-down, choose the extended access control list object. You can predefine the ACL object (see Configure an extended ACL object) or click the **Add** (+) icon to create the object. In the **New Extended Access List Object** box, enter a name, click **Add** to open the **Add Extended Access List Entry** dialog box, where you can define the network, port, user identity, SGT, or application match criteria for the PBR policy.

Note
You can either have destination address or application/user identity/SGT defined in an ACE.

To selectively apply PBR on the incoming interface, you can define *Block* criteria in the ACE. When the traffic matches the block rule of the ACE, the traffic is forwarded to the egress interface based on the routing table.

b) From the **Send To** drop-down list:

- To select the configured interfaces, choose **Egress Interfaces**.
- To specify the IPv4/IPv6 next hop addresses, choose **IP Address**. Proceed to Step 7.e, on page 10

c) If you have selected **Egress Interfaces**, from the **Interface Ordering** drop-down, choose the relevant option:

- By **Interface Priority**—The traffic is forwarded based on the priority of the interfaces. Traffic is routed to the interface with the least priority value first. When the interface is not available, the traffic is then forwarded to the interface with the next lowest priority value. For example, let us assume that interface1, interface2, and interface3 are configured with priority values *0*,*1*, and *2*, respectively. The traffic is forwarded to interface1. If interface1 becomes unavailable, the traffic is then forwarded to interface2.

Note
To configure the priority for the interfaces, click **Configure Interface Priority** on the Policy Based Routing page. In the dialog box, provide the priority number against the interfaces, and then click **Save**. You can also configure the priority for an interface in the Interface Settings.

When the priority value is the same for all the interfaces, the traffic is balanced among the interfaces. By default, the interface priority for PBR is set to 0. Interfaces with the lowest priority value are preferred for traffic forwarding.

- By **Order**—The traffic is forwarded based on the sequence of the interfaces specified here. For example, let us assume that interface1, interface2, and interface3 are selected in this order, interface2, interface3, interface1. The traffic is forwarded first to interface2, then to interface3, regardless of their priority values.
- By **Minimal Jitter**—The traffic is forwarded to the interface that has the lowest jitter value. You need to enable Path Monitoring on the interfaces for PBR to obtain the jitter values.
- By **Maximum Mean Opinion Score**—The traffic is forwarded to the interface that has the maximum mean opinion score (MOS). You need to enable Path Monitoring on the interfaces for PBR to obtain the MOS values.
- By **Minimal Round Trip Time**—The traffic is forwarded to the interface that has the minimal round trip time (RTT). You need to enable Path Monitoring on the interfaces for PBR to obtain the RTT values.
- By **Minimal Packet Loss**—The traffic is forwarded to the interface that has the minimal packet loss. You need to enable Path Monitoring on the interfaces for PBR to obtain the packet loss values.

d) In the **Available Interfaces** box, all the interfaces with their priority values are listed. From the list of interfaces, click the **Add** (+) button to add to the selected egress interfaces. Proceed to Step 8.

Note
The selected interface route must be present in the routing table.

e) If you have selected **IP Address**, enter the IP addresses separated by commas in the **IPv4 Addresses** or **IPv6 Addresses** fields. The traffic is forwarded as per the sequence of the specified IP addresses.

Note
When multiple next-hop IP addresses are provided, the traffic is forwarded as per the sequence of the specified IP addresses until a valid routable next-hop IP address is found. The configured next-hops should be directly connected.

f) From the **Don't Fragment** drop-down list, select Yes, No, or None. If the DF (Don't Fragment) flag is set to *Yes*, the intermediate routers never perform fragmentation of a packet.To specify the current interface as the default for forwarding, check the **Default Interface** check box.

g) The **IPv4 Settings** and **IPv6 Settings** tabs enable you to specify the recursive and default settings:

Note
For a route-map, you can only specify either IPv4 or IPv6 next-hop settings.

- **Recursive**—The route map configuration is applied only when the specified next-hop address and the default next-hop address are found on a directly connected subnet. However, you could use the recursive option, where the next-hop address need not be directly connected. Here, 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 current routing path of the router.
- **Default**—If the normal route lookup fails to match traffic, the traffic is forwarded to this specified next-hop IP address.

h) Check the **Peer Address** check box to use the next-hop address as the peer address.

Note
You cannot configure a route map with both default next-hop address and peer address.

i) For IPv4 settings, you can check whether the next IPv4 hops of a route map are available under **Verify Availability**—click the **Add** (+) button and add the next-hop IP address entries:

- **IP Address**—Enter the next hop IP address.
- **Sequence**—Entries are assessed in order using the sequence number. Ensure that no duplicate sequence numbers are entered. The valid range is 1 to 65535.
- **Track**—Enter a valid ID. The valid range is 1 to 255.

Step 8 To save the policy, click **Save** and **Deploy**.

The Firewall Threat Defense uses ACLs to match traffic and perform routing actions on the traffic. Typically, you configure a route map that specifies an ACL against which traffic is matched, and then you specify one or more actions for that traffic. With the use of path monitoring, PBR can now select the best egress interface for routing the traffic. Finally, you associate the route map with an interface on which you want to apply PBR on all incoming traffic.

Configure policy-based routing

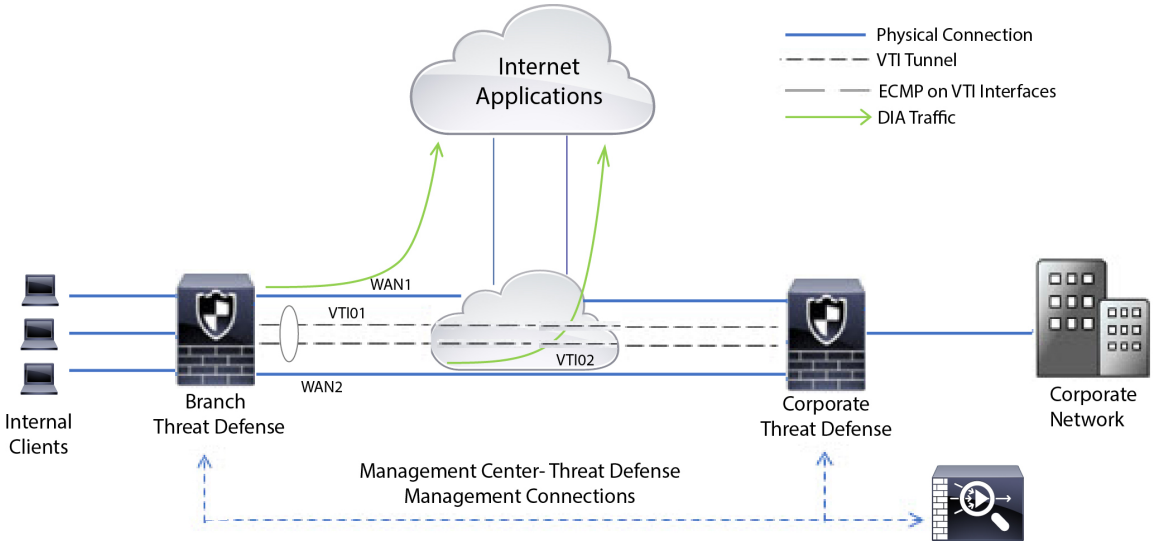
Configure PBR to enable direct internet access for specific web-based applications, optimize traffic flow, and achieve load balancing between WAN and VTI interfaces.

Consider a typical corporate network scenario where all the branch network traffic passes through a route-based VPN of the corporate network and diverges to the extranet, when required. Accessing web-based applications through the corporate network to support daily operations may increase network size and maintenance costs. This example illustrates how to configure PBR for direct internet access.

The following figure depicts the topology of a corporate network. The branch network is connected to the corporate network through a route-based VPN. Traditionally, the corporate Firewall Threat Defense is configured to handle both the internal and external traffic of the branch office. With the PBR policy, the branch Firewall Threat Defense is configured with a policy that routes specific traffic to the WAN network instead of the virtual tunnels. The rest of the traffic flows through the route-based VPN, as usual.

This example also illustrates how to configure the WAN and VTI interfaces with ECMP zones to achieve load balancing.

Figure 3: Configuring policy-based routing on Branch Firewall Threat Defense in Firewall Management Center



Before you begin

This example assumes that you have already configured WAN and VTI interfaces for the branch Firewall Threat Defense in Firewall Management Center.

Procedure

Step 1 Configure policy-based routing for the branch Firewall Threat Defense, select the ingress interfaces:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.

c) In the **Add Policy Based Route** dialog box, select the interfaces (say, *Inside 1*, and *Inside 2*) from the **Ingress Interface** drop-down list.

Step 2 Specify the match criteria:

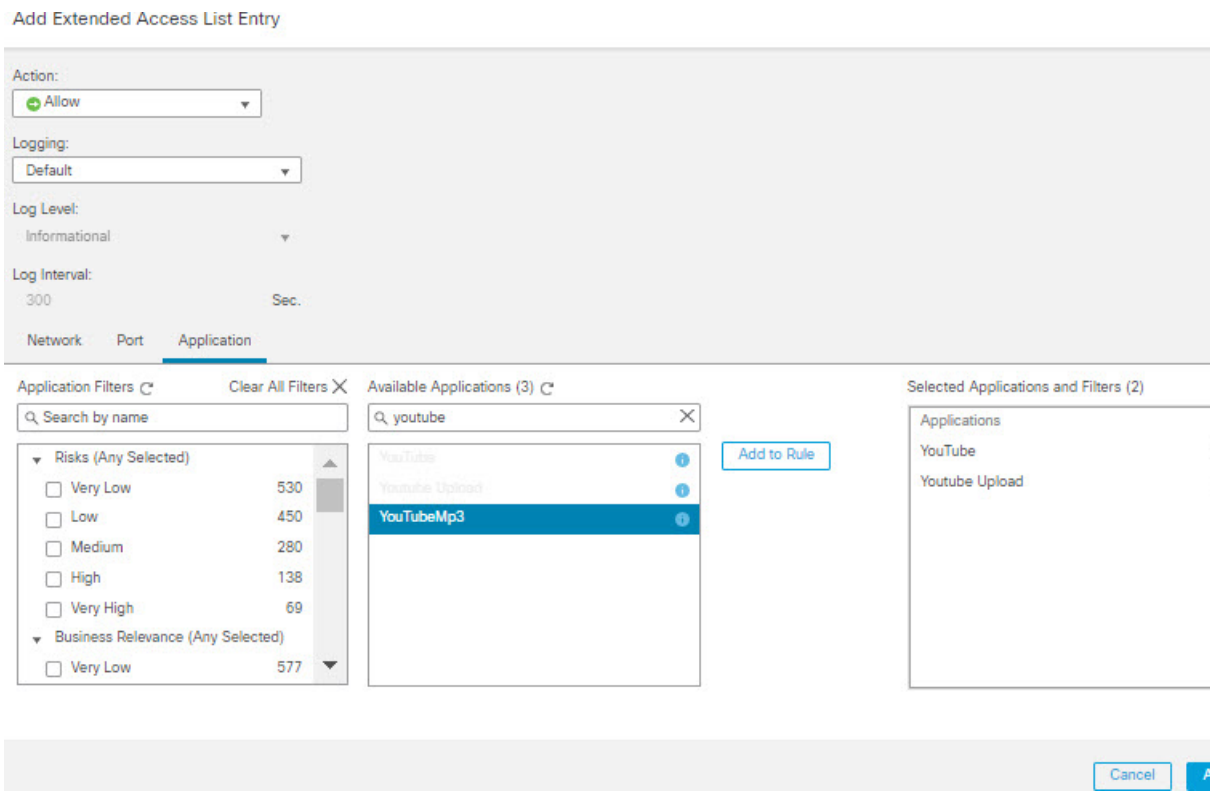
a) Click **Add**.

b) To define the match criteria, click the **Add** (+) button.

c) In **New Extended Access List Object**, enter the name for the ACL (say, *DIA-FTD-Branch*), and click **Add**.

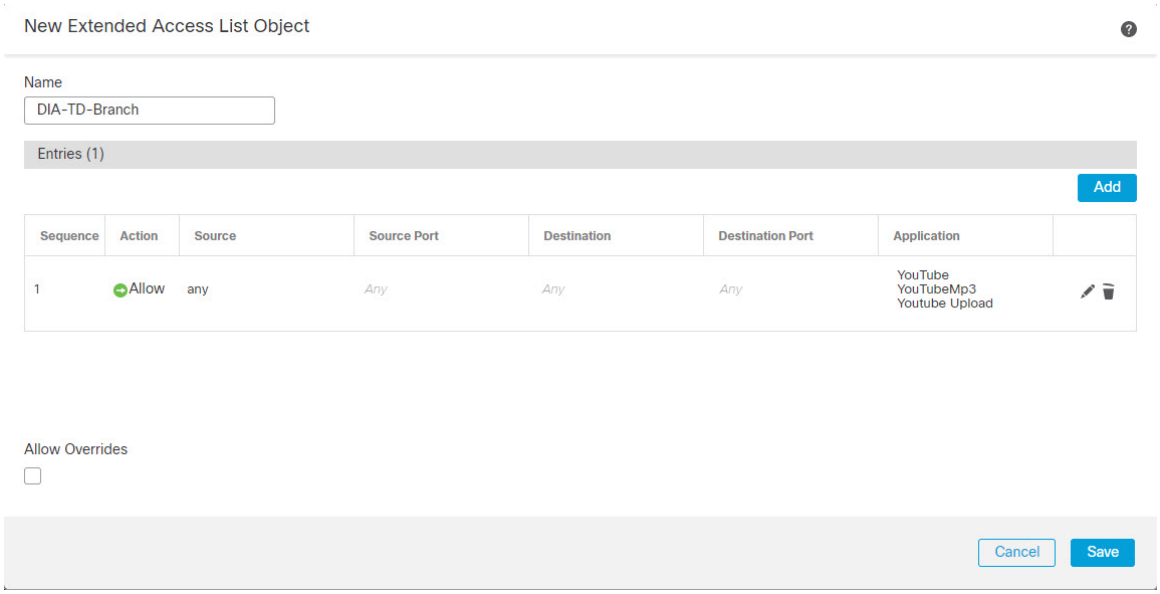
d) In the **Add Extended Access List Entry** dialog box, choose the required web-based applications from the **Application** tab:

Figure 4: Applications tab



On the Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

Figure 5: Extended ACL



e) Click **Save**.

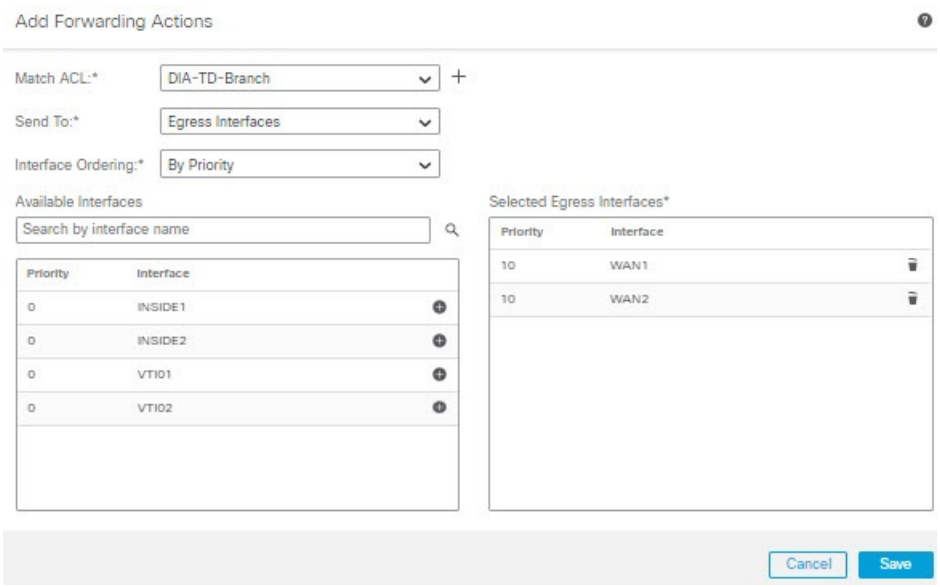
f) Select *DIA-FTD-Branch* from the **Match ACL** drop-down list.

Step 3 Specify the egress interfaces:

a) From the **Send To** and **Interface Ordering** drop-down lists, choose Egress Interfaces, and Interface Priority respectively.

b) Under **Available Interfaces**, click the ⊕ button against the respective interface names to add *WAN1* and *WAN2*:

Figure 6: Configuring policy-based routing



c) Click **Save**.

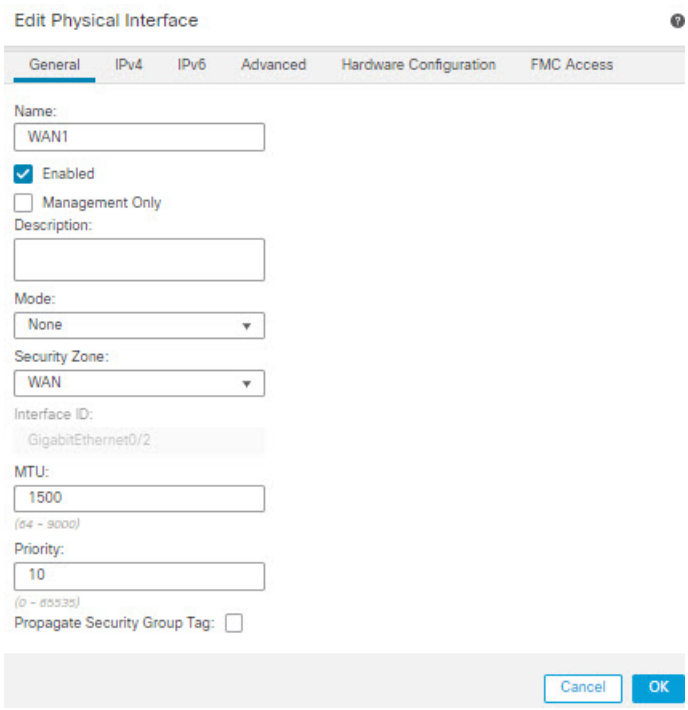
Step 4 Configure interface priority:

You can set the priority value for the interfaces either in the **Edit Physical Interface** page, or in the **Policy Based Routing** page (**Configure Interface Priority**). In this example, the Edit Physical Interface method is described.

a) Choose **Devices** > **Device Management**, and edit the branch Firewall Threat Defense.

b) Set the priority for the interfaces. Click **Edit** against the interface and enter the priority value:

Figure 7: Setting interface priority



c) Click **Ok** and **Save**.

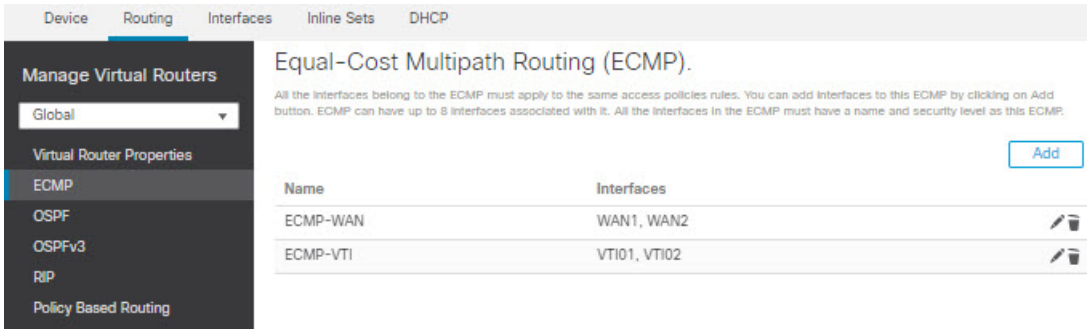
Step 5 Create ECMP zones for load balancing:

a) In the **Routing** page, click **ECMP**.

b) To associate interfaces to the ECMP zone, click **Add**.

c) Select *WAN1* and *WAN 2* and create an ECMP zone—*ECMP-WAN*. Similarly, add *VTI01* and *VTI02* and create an ECMP zone—*ECMP-VTI*:

Figure 8: Associating interfaces with ECMP zone

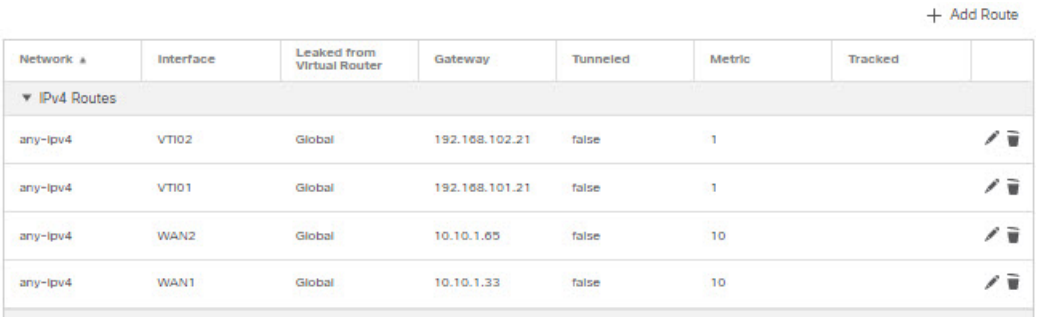


Step 6 Configure static routes for the zone interfaces for load balancing:

a) In the **Routing** page, click **Static Route**.

b) Click **Add** and specify the static routes for *WAN1*, *WAN2*, *VTI01*, and *VTI02*. Ensure that you specify the same metric value for the interfaces belonging to the same ECMP zones (Step 5):

Figure 9: Configuring static routes for ECMP zone interfaces



Note
Ensure that the zone interfaces have the same destination address and metric, but different gateway addresses.

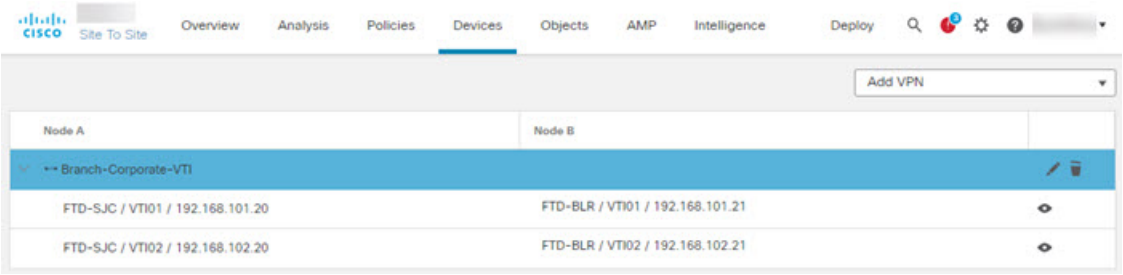
Step 7 Configure trusted DNS on the WAN objects of the branch Firewall Threat Defense to ensure secure flow of traffic to the internet:

a) Choose **Devices** > **Platform Settings**, and create a DNS policy on the branch Firewall Threat Defense.
b) To specify the trusted DNS, edit the policy and click **DNS**.
c) To specify the DNS servers for the DNS resolution to be used by WAN objects, in the **DNS Settings** tab, provide the DNS server group details and select WAN from the interface objects.
d) Use the **Trusted DNS Servers** tab to provide specific DNS servers that you trust for the DNS resolution.

Step 8 Click **Save**, and then click **Deploy**.

Any *YouTube* related access requests from the branch inside network *INSIDE1* or *INSIDE2* are routed to *WAN1* or *WAN2* as they would match the *DIA-FTD-Branch* ACL. Any other request, say *google.com*, are routed through *VTI01* or *VTI02* as configured in the Site to Site VPN Settings.

Figure 10: Site to site VPN settings



With the ECMP configured, the network traffic is seamlessly balanced.

Configure PBR with path monitoring

Configure PBR with path monitoring to optimize application traffic routing by dynamically selecting the best egress interface based on performance metrics such as jitter, round-trip time, and packet loss.

This example describes how to configure PBR with path monitoring for the following applications with flexible metrics:

- Audio or video sensitive applications (example, WebEx Meetings) with Jitter.
- Cloud-based application (example, Office365) with RTT.
- Network-based access control (with a specific source and destination) with packet loss.

Before you begin

1. This example assumes that you are aware of the basic configuration steps for PBR.
2. You have configured ingress and egress interfaces with logical names. In this example, the ingress interface is named *Inside1*, and egress interfaces are named *ISP01*, *ISP02*, and *ISP03*.

Follow these steps to configure PBR with path monitoring:

Procedure

Step 1 Path monitoring configuration on interfaces *ISP01*, *ISP02*, and *ISP03*:

For the metrics collection on the egress interfaces, you must enable and configure path monitoring on them.

a) Choose **Devices** > **Device Management**, and edit Firewall Threat Defense.
b) Under the **Interfaces** tab, edit the interface (in our example, *ISP01*)
c) Click the **Path Monitoring** tab, select the **Enable Path Monitoring** check box, and then specify the monitoring type (see Configure path monitoring settings, on page 7).
d) Click **Ok** and **Save**.
e) Repeat these steps to configure the path monitoring settings for *ISP02* and *ISP03*.

Step 2 Configure policy-based routing for a branch in an organization Firewall Threat Defense, select the ingress interfaces:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.
c) In the **Add Policy Based Route** dialog box, select *Inside 1* from the **Ingress Interface** drop-down list.

Step 3 Specify the match criteria:

a) Click **Add**.
b) To define the match criteria, click the **Add** (+) button.
c) In **New Extended Access List Object**, enter the name for the ACL (example, *PBR-WebEx*), and click **Add**.
d) In the **Add Extended Access List Entry** dialog box, choose the necessary web-based applications, such as WebEx Meetings) from the **Application** tab.

Remember
On Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

e) Click **Save**.
f) Select *PBR-WebEx* from the **Match ACL** drop-down list.

Step 4 Specify the egress interfaces:

a) From the **Send To** drop-down list, choose Egress Interfaces.
b) From the **Interface Ordering** drop-down list, choose By Minimal Jitter.
c) Under **Available Interfaces**, click the **Right Arrow** (>)button against the respective interface names to add *ISP01*, *ISP02*, and *ISP03*.
d) Click **Save**.

Step 5 Repeat Step 2 and Step 3 to create PBRs for the *Inside1* interface to route Office365 and network-based access control traffic:

a) Create a match criteria object, example *PBR-Office365*, and select the Office365 application from the **Application** tab.
b) From the **Interface Ordering** drop-down list, choose By Minimal Round Trip Time.
c) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.
d) Now, create a match criteria object, example *PBR-networks*, and specify the source and destination interface in the **Network** tab.
e) From the **Interface Ordering** drop-down list, choose By Minimal Packet Loss.
f) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.

Step 6 Click **Save**, and then click **Deploy**.

Step 7 To view path monitoring metrics, Choose **Devices** > **Device Management**, and from **More** (⋮) click **Health Monitor**. To view the metric details for the interfaces of the device, you must add the path metrics dashboard. For details, see Add path monitoring dashboard, on page 8.

The WebEx, Office365, and networks-based ACL traffic are forwarded through the best route derived from the metrics value collected on *ISP01*, *ISP02*, and *ISP03*.

Useful CLIs for monitoring PBR

Run the monitoring commands described in this topic from the Firewall Threat Defense device CLI.

Interface configurations

To view the interface configurations of the device, run the `show run interface` command:

```
> show run interface
!
interface Ethernetl/l
 description Outside ispl handoff
 nameif outside1
 security-level 0
 zone-member ECMP-WAN
 ip address dhcp setroute
 policy-route cost 10
 policy-route path-monitoring 8.8.8.8
 policy-route path-monitoring object-group network-service FMC_NSG_4295470581 policy-route
 path-monitoring object-group network-service FMC_NSG_4295470600
!
interface Ethernet1/2
 description Outside isp2 handoff
 nameif outside2
 security-level 0
 zone-member ECMP-WAN
```

```
 ip address 192.133.243.240 255.255.255.192
 policy-route cost 20
 policy-route path-monitoring 8.8.8.8
 policy-route path-monitoring object-group network-service FMC_NSG_4295470581 policy-route
 path-monitoring object-group network-service FMC_NSG_4295470600
!
```

DNS configurations

Application-based routing only uses trusted DNS servers to resolve domains. To view the DNS configurations of the device, run the `show run DNS` command:

```
> show run dns
DNS server-group DefaultDNS
dns trusted-source 10.100.0.5
dns trusted-source 10.200.0.5
```

Route map configurations

When you configure PBR on the device, the Management Center auto-generates the route-map and applies it to the specified ingress interface. To view the route maps of the device, run the `show run route-map` command:

```
> show run route-map
!
route-map FMC_VPN_CONNECTED_DIST_RMAP_1000 permit 10
 match interface inside-employee
 set community 1000
!
route-map FMC GENERATED PBR 1729024850865 permit 5
 match ip address Cloud-storage-apps-acl
  set adaptive-intertace cost outside1 outside2

!
route-map FMC_GENERATED PBR 1729024850865 permit 10
match ip address Social-media-apps-acl
set adaptive-interface rtt outsidel outside2

!
route-map FMC GENERATED PBR 1729024850865 permit 15
match ip address Conferencing-apps-acl
set adaptive-interface jitter outside1 outside2

!
route-map FMC_GENERATED_PBR_1729024850865 permit 20
match ip address Corp-internal-apps-acl
set adaptive-interface cost outsidel_static_vti_1 outside2_static_vti_4
```

Access lists and network service groups configurations

The route-map applied to the ingress interface can reference an extended access control list. To view the details of an access list for PBR, run the `show run access list <access list_name>` command:

```
> show run access-list Cloud-storage-apps-acl
access-list Cloud-storage-apps-acl extended permit ip any object-group-network-service
FMC_NSG_4295470562
```

The network-service objects and object-groups are configured in extended access control lists and referenced in policy-based routing route maps and access control groups. To view the NSG configurations, run the `show object-group network-service <network-service-groups-name>` command. The *network-service-groups-name* is derived from the above show command for an access list.

```
> show object-group network-service FMC_NSG_4295470562
 object-group network-servire FMC_NSG_4295470562 (id=@xfdff0000)
 network-service-member "Box" dynamic
 description File storage and transfer site.
 app-id 1326
 domain box.com (bid=436735707) ip (hitcnt=0)
 domain boxcloud.com (bid=436924171) ip (hitcnt=0)
 domain box.net (bid=437080553) ip (hitcnt=0)
 domain box.org (bid=437174273) ip (hitcnt=0)
 domain boxcdn.net (bid=437272231) ip (hitcnt=0)
 domain boxrelay.com (bid=437481703) ip (hitcnt=0)
 domain boxenterprise.net (bid=437626005) ip {hitcnt=0)
 domain boxinvestorrelations.com (bid=437672765) ip (hitcnt=0)
 domain segment-box.com (bid=437886771) ip (hitcnt=0)
 domain box-corp.com (bid=437924995) ip (hitcnt=0)
 domain boxcn.net (bid=438072833) ip (hitcnt=0)
network-service-member "Dropbox" dynamic
 description Cloud based tile storage.
 app-id 125
 domain dropbox.com (bid=24259639) ip (hitcnt=0)
 domain cfl.dropboxstatic.com (bid=24495525) ip (hitcnt=0)
 domain dl.dropboxusercontent.com (bid=24596237) ip (hitcnt=0)
 domain dropboxapi.com (bid=24694467) ip (hitcnt=0)
 domain dropboxbusiness.com (bid=24859859) ip (hitcnt=0)
 domain dropboxcaptcha.com (bid=25008145) ip {hitcnt=0)
 domain dropbox-dns.com (bid=25087753) ip (hitcnt=0)
 domain dropboxer.net (bid=25236751) ip (hitcnt=0)
 domain dropboxusercontent.com (bid=25324335) ip (hitcnt=0)
 domain getdropbox.com (bid=25437501) ip (hitcnt=0)
 domain cloudon.com (bid=25580229) ip (hitcnt=0)
```

Path monitoring configurations

To view the path monitoring metrics collected on the egress interface, run the `show path-monitor` command:

```
> show path-monitor
Interface: outside2 (Ethernetl/2)
Remote peer: 8.8.8.8
    Remote peer reachable: Yes
    RTT average: 9138 microsecondes) Jitter: 1093 microsecond(s)
    Packet loss: 0% MOS: 4.39
    Last updated: 12 second(s) ago
Interface: outside2 (Ethernetl/2)
Remote NSG: FMC_NSG_4295470581
    Network Service: Facebook Domain name: fbsbx.com Remote peer reachable: Yes
    RTT average: 17460 microsecond(s) Jitter: 911 microseconde)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago

    Network Service: Facebook
    Domain name: facebook.net
    Remote peer reachable: Yes
    RTT average: 17444 microsecondes)
    Jitter: 836 microsecondes)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago

    Network Service: Instagram
    Domain name: instagram.com Remote peer reachable: Yes
    RTT average: 17576 microsecondes)
    Jitter: 429 microsecondes)
    Packet loss: 0%
```

```
    MOS: 4.39
    Last updated: 12 secondes) ago

Interface: outside2 (Ethernetl/2)
Remote NSG: FMC_NSG_4295470600
    Network Service: WebEx
    Domain name: webex.com Remote peer reachable: Yes RTT average: 18537 microsecond(s)
Jitter: 318 microseconde)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago
    Network Service: Zoom Domain name: zoom.com Remote peer reachable: Yes
    RTT average: 98196 microsecond(s) Jitter: 4120 microseconde)
    Packet loss: 0%
    MOS: 4.34
    Last updated: 12 second(s) ago
```

Troubleshoot PBR

This task helps you debug PBR configuration when packets get dropped by verifying routes, monitoring metrics, and validating interface selection.

When PBR is not working as intended, you need to systematically verify the configuration components including route maps, path monitoring, and interface selection to identify the root cause of packet drops.

Before you begin

Follow these steps to troubleshoot PBR configuration when packets get dropped:

Procedure

Step 1 Check that all essential routes for recursive resolution are present by using the show route or show ipv6 route commands in the appropriate tables.

Unless the route map for the interfaces participating in PBR is updated with correct routes, PBR will not work as intended.

Step 2 Display the remote address that path-monitoring uses to monitor metrics by using the **show running-config interface** command.

Example:

```
interface GigabitEthernet0/0
nameif outside_0
security-level 0
zone-member ecmp-zone
ip address 20.0.0.3 255.255.255.0 -> This is egress interface "show running-config" output,
 the monitored address and cost metric value is determined in this output.
policy-route cost 1
policy-route path-monitoring
20.0.0.4
!
int GigabitEthernet 0/3
!
interface GigabitEthernet0/3
nameif outside_3
security-level 0
ip address 11.1.1.2 255.255.255.0 -> This is ingress interface "show running-config" output,
```

```
 the specific route-map will be used by PBR to determine the next route.
policy-route route-map rtt-tes
```

This command shows the monitored address and cost metric values for path monitoring by sending ICMP packets or HTTP pings.

Step 3 Examine the **show path-monitoring** and **show run route-map** outputs for packet loss.

Example:

```
ciscoasa(config)# show path-monitoring
Interface: outside_0  -> The remote address used for ICMP monitoring
Remote peer: 20.0.0.4
Version: 6223
Remote peer reachable: Yes -> If this value turns "No", then the ICMPv4/v6 packet is not
reaching the required remote address.
RTT average: 1920 microsecond(s)
Jitter: 394 microsecond(s)
Packet loss: 0%
MOS: 4.40
Last updated: 17 second(s) ago -> The data should be updated by path monitoring after every
 30 seconds. The 'show route-map' would show the
updated metric values.
Interface: outside_2
Remote peer: 40.0.0.4
Version: 6223
Remote peer reachable: Yes
RTT average: 1935 microsecond(s)
Jitter: 433 microsecond(s)
Packet loss: 0%
MOS: 4.40
Last updated: 17 second(s) ago
```

Example:

```
ciscoasa(config)# show route-map
route-map rtt-test, permit, sequence 10
Match clauses:
Set clauses:
adaptive-interface rtt outside_0 (1920) outside_2 (1935) outside_1 (1971) -> Displays the
metric type (rtt) that is used by the policy route to select the adequate interface to send
 the packet. The interface list where cost of each interface is given. As the metric type
is "rtt" and considering the minimum rtt value, the "outside_0" interface route will be
selected by PBR.
route-map mos-test, permit, sequence 10
Match clauses:
Set clauses:
adaptive-interface mos outside_0 (378) outside_1 (390) outside_2 (440) - > As the metric
type is "mos", considering the maximum value of mos, the "outside_2" interface will be
selected by PBR.
```

The metric types (lost, rtt, jitter, cost) should select the interface with the minimum metric value for routing.

The metric type "mos" should select the interface with the maximum metric value for routing.

Step 4 Use **packet-tracer** command to validate the interface selection based on the metric type defined in the PBR.

Example:

```
packet-tracer input <interface> icmp <src ip address> 8 0 <dst ip
address> detailed
Phase: 3
Type: PBR-LOOKUP
Subtype: policy-route
Result: ALLOW
Elapsed time: 60656 ns
```

```
Config:
route-map rtt-test permit 10
match ip address allow_101_1_1_2
set adaptive-interface rtt outside_0 outside_2
Additional Information:
Matched route-map rtt-test, sequence 10, permit
Found next-hop 40.0.0.4 using egress ifc outside_2 -> PBR selects the adequate interface
from adaptiveinterface list given in "rtt-test" route-map.
```

Step 5 Use the **packet-tracer** command similar to the **debug policy-route** command.

When PBR successfully selects the route, the packet tracer output looks like this:

```
pbr: policy based route lookup called for 101.1.1.1/0 to 101.1.1.2/0 proto 1 sub_proto 8
received on
interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map rtt-test, sequence 10, permit; proceed with policy routing
pbr: policy based routing applied; egress_ifc = outside_2 : next_hop = 20.0.0.4
```

When PBR is not able to find the adequate route, it falls back to normal route lookup, and the packet tracer output looks like this:

```
pbr: policy based route lookup called for 100.1.1.1/0 to 100.1.1.2/0 proto 1 sub_proto 8
received on interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map mos-test, sequence 10, permit; proceed with policy routing
pbr: no route to 100.1.1.2 on adaptive-interface outside_2
pbr: no route to 100.1.1.2 on adaptive-interface outside_1
pbr: no route to 100.1.1.2 on adaptive-interface outside_0
pbr: policy based routing could not be applied; proceeding with normal route lookup
```

When a monitored remote address goes down, and path monitoring marks **Remote peer reachable** as **No** for that address, the PBR displays the a log that excludes the interface from adaptive-interface list.

```
pbr: policy based route lookup called for 100.1.1.1/0 to 101.1.1.2/0 proto 1 sub_proto 8
received on
interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map rtt-test, sequence 10, permit; proceed with policy routing
pbr: Path Monitoring Ifc Down : adaptive-interface outside_1 Excluded from PBR routing
pbr: policy based routing applied; egress_ifc = outside_2 : next_hop = 40.0.0.4
```

Note
The interface becomes eligible for adaptive PBR routing when the interface is reported as reachable in the path monitoring module.

History for policy-based routing

This reference provides a comprehensive history table that summarizes the introduction and evolution of policy-based routing features in Cisco Secure Firewall Management Center across releases, enabling quick tracking of feature availability and changes.

Table 1: History table

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Identity and SGT based PBR Policy	7.4.0	7.4.0	You can now classify the network traffic based on users and user groups, and SGTs in PBR policies. You can select the identity and SGT objects while defining the extended ACLs for the PBR policies. New/modified screens: New tabs added in the extended access list object for configuring the policy-based routing policy: **Objects** > **Object Management** > **Access Control Lists** > **Add Extended** page, **Users** and **Security Group Tag**.
HTTP based Path Monitoring	7.4.0	7.2.0	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 application monitoring option is enabled by default for the interface. You can configure a PBR policy with match ACL having the monitored applications and desired metric type for path determination. New/modified screens: New option in Interfaces page for enabling path monitoring: **Devices** > **Device Management** > **Edit Interfaces** > **Path Monitoring** > **Enable HTTP based Application Monitoring** check box.
Dual WAN/ISP Threat Defense Management Support	7.3.0	7.3.0	On a dual WAN-enabled threat defense, a single data interface was configured to communicate with the management center. Now, support to configure a secondary data interface is provided so that the communication channel is sustained when the primary data interface fails. The management center automatically configures PBR to route the SF-Tunnel traffic from the *tapnlp* (internal) interface to one of the available data interfaces based on the priority and SLA metric.
Next-hop settings for PBR route map	7.3.0	7.1.0	You can configure the next-hops for the PBR route-map while enabling packet forwarding actions. New/modified screens: New fields in Add/Edit Forwarding Actions page for configuring egress interfaces: **Device Management** > **Routing** > **Policy Based Routing** > **Add Forwarding Actions** page.
PBR and Path Monitoring	7.2.0	7.2.0	PBR uses path monitoring to collect the performance metrics (RTT, jitter, packet-lost, and MOS) of the egress interfaces. You must enable path monitoring for the interface and configure the monitoring type. You can configure a PBR policy with the desired metric for path determination. New/modified screens: New tab in Interfaces page for enabling path monitoring: **Devices** > **Device Management** > **Edit Interfaces** > **Path Monitoring** tab.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Configure policy based routing from the FMC web interface.	7.1.0	7.1.0	**Upgrade impact. Redo FlexConfigs after upgrade.** You can now configure policy based routing (PBR) from the FMC web interface. This allows you to classify network traffic based on applications and to implement direct internet access (DIA) to send traffic to the internet from a branch deployment. You can define a PBR policy and configure it on ingress interfaces, specifying match criteria and egress interfaces. Network traffic that matches the access control policy is forwarded through the egress interface based on priority or the order as configured in the policy. This feature requires Version 7.1+ on both the FMC and the device. When you upgrade the FMC to Version 7.1+, existing policy based routing FlexConfigs are removed. After you upgrade your devices to Version 7.1+, redo your policy based routing configurations in the FMC web interface. For devices that you do not upgrade to Version 7.1+, redo the FlexConfigs and configure them to deploy "every time." New/modified screens: **Devices** > **Device Management** > **Routing** > **Policy Based Routing**