



Policy Based Routing

This chapter describes how to configure Threat Defense to support policy based routing (PBR) through 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](#)
- [Guidelines and Limitations for Policy Based Routing, on page 3](#)
- [Path Monitoring, on page 4](#)
- [Configure Policy-Based Routing Policy, on page 7](#)
- [Configuration Example for Policy Based Routing, on page 10](#)
- [Configuration Example for PBR with Path Monitoring, on page 15](#)
- [History for Policy Based Routing, on page 17](#)

About Policy Based Routing

In traditional routing, packets are routed based on the destination IP address. However, it is difficult to change the routing of specific traffic in a destination-based routing system. Policy Based Routing (PBR) gives you more control over routing by extending and complementing the existing 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 is applicable in scenarios where, numerous devices access applications and data in a large network deployment. Traditionally, large deployments have topologies that backhaul all the network traffic 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. Overcoming these issues involves high-cost complex deployments and management.

PBR policy enables you to securely breakout 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.

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 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.

Following are a few scenarios where you can use Policy Based Routing:


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 threat defense is configured with an internet exit point and 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, as shown 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 route the traffic from HR network through ISP1 and traffic from Eng network through ISP2 and thus share the load.

Guidelines and Limitations for Policy Based Routing

Firewall Mode Guidelines

PBR is supported only on routed firewall mode.

Device Guidelines

- PBR through management center's Policy Based Routing page is supported only from Version 7.1+ on both the management center and the device.
- When you upgrade management center or threat defense to version 7.1 and higher, the PBR configuration in the device is removed. You must configure PBR again using the Policy Based Routing page. If the managed device is lower than version 7.1, you must configure PBR again using FlexConfig with deploy option set to "every time."
- Configuring application, user identity, and Security Group Tag (SGT) based PBR policy on cluster devices is not supported.

Interface Guidelines

- Only routed interfaces and non management-only interfaces belonging to the Global virtual router can be configured as ingress or egress interface.
- PBR is not supported on user-defined virtual routers.
- Only interfaces that have a logical name can be defined in the policy.
- Static VTIs can be configured only as egress interfaces.
- 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 asymmetric routing, specifically when NAT and VPN are in use.

IPv6 Support

PBR supports IPv6.

Application-Based PBR and DNS Configuration

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

For more information on configuring DNS servers, see [DNS](#).

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.

HTTP-based Path Monitoring Guidelines

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

Additional Guidelines

- All existing configuration restrictions and limitations of route map will be carried forward.
- While defining the ACL for the policy match criteria, you can select multiple applications from a list of predefined applications to form an Access Control Entry (ACE). In 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. Hence, when using PBR, ensure to disable uRPF.

Path Monitoring

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

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

Path monitoring computes flexible metrics for multiple remote peers per interface. To monitor and determine best path for multiple applications through a policy on a branch firewall, HTTP is preferred over ICMP for the following reasons:

- HTTP-ping can derive the performance metrics of the path up to the application layer of the server, where the application is hosted.
- The need to change the firewall configuration whenever the application server IP address is changed is removed as the application domain is tracked instead of the IP address.



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 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

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 7](#).

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 does not commence monitoring immediately after HTTP-based application monitoring is configured. It starts monitoring only when a DNS entry is snooped for a domain. With the information on the resolved IP for the domain, it sends and receives the HTTP request and response respectively. When DNS resolves multiple IP addresses for a single domain, the first resolved IP address will be used for probing and monitoring the application. It continues to monitor till the IP address changes or the HTTP-based path monitoring is disabled.

Based on the HTTP request and response durations, path monitoring computes the performance metrics for the application. The collected metrics are forwarded to PBR periodically, for making the routing and forwarding decision for the traffic arising from the configured ingress interface. If traffic arrives before the path monitoring could send its metrics to PBR, the traffic flow follows path chosen by the routing table. For the subsequent traffic flows that arrive after path monitoring metrics are available, PBR applies its routing decision based on the metrics and forwards the 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.

In HTTP-based path monitoring of applications, the management center associates the applications/NSGs to the egress interface only when the PBR configuration meets the following criteria:

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

Configure Path Monitoring Settings

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 interfaces with settings for path monitoring to send the ICMP probes or HTTP pings for metrics collection.

Procedure

-
- Step 1** Select **Devices > Device Management** and click **Edit** (✎) for your 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**, and to save the settings, click **Save**.
-

Configure Policy-Based Routing Policy

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 6](#).

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Click **Routing**.

Step 3 Click **Policy Based Routing**.

The Policy Based Routing page displays the configured policy. The grid displays the list of ingress interfaces and a 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 Only interfaces that have logical names and that belong to a global virtual router are listed in the drop-down.

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 Extended ACL Objects](#)) 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 9](#)

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 *Gig0/1*, *Gig0/2*, and *Gig0/3* are configured with priority values *0*, *1*, and *2* respectively. The traffic is forwarded to *Gig0/1*. If *Gig0/1* becomes unavailable, the traffic is then forwarded to *Gig0/2*.

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 **Order**—The traffic is forwarded based on the sequence of the interfaces specified here. For example, let us assume that *Gig0/1*, *Gig0/2*, and *Gig0/3* are selected in the following order, *Gig0/2*, *Gig0/3*, *Gig0/1*. The traffic is forwarded to *Gig0/2* first, then to *Gig0/3*, irrespective 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 7.k, on page 10](#)
- 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.
- g) To specify the current interface as the default for forwarding, check the **Default Interface** check box.
- h) The **IPv4 Settings** and **IPv6 Settings** tab enables 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.
- i) 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.
- j) 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.

k) Click **Save**.

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


The 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.

Add Path Monitoring Dashboard

To view the path monitoring metrics, you must add the path monitoring dashboard to the Health Monitoring page of the device.

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 the four metrics are selected for displaying as portlets in the dashboard with an additional metric field. You can exclude any of them by clicking **Delete** ().

- Step 6** Click **Add Dashboard**.
-

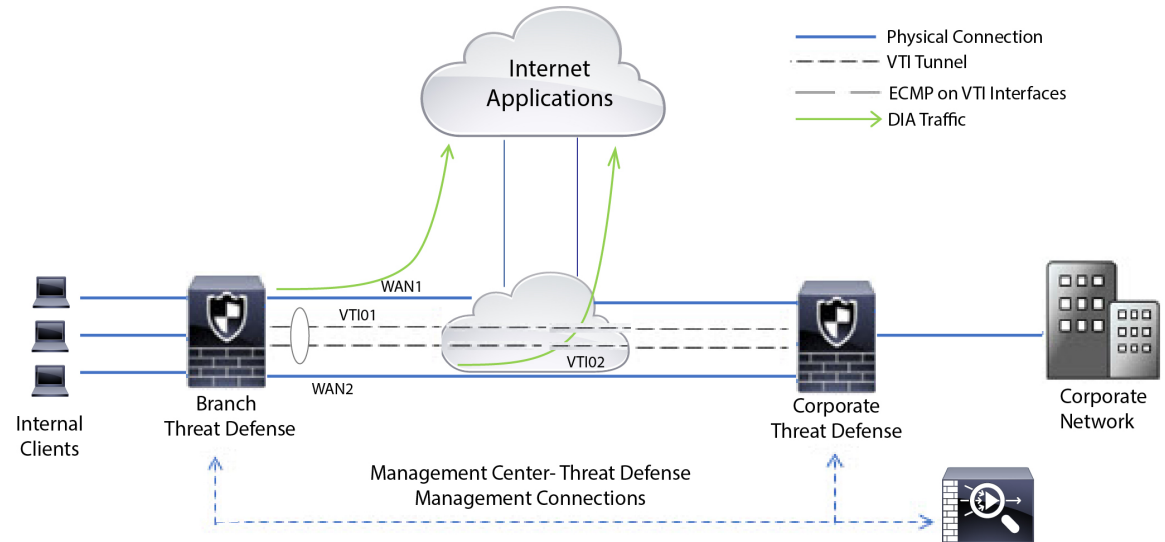
Configuration Example for Policy Based Routing

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 the web-based applications that address day-to-day operations through the corporate network results in huge network expansion and maintenance costs. This example illustrates the PBR configuration procedure 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 threat defense is configured to handle both the internal and external traffic of the branch office. With the PBR policy, the branch 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 the configuring of the WAN and the VTI interfaces with ECMP zones to achieve load balancing.

Figure 1: Configuring Policy Based Routing on Branch Threat Defense in Management Center



Before you begin

This example assumes that you have already configured WAN and VTI interfaces for the branch threat defense in management center.

Procedure

-
- Step 1** Configure policy based routing for the branch threat defense, select the ingress interfaces:
- Choose **Devices > Device Management**, and edit the threat defense device.
 - Choose **Routing > Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.
 - 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:
- Click **Add**.
 - To define the match criteria, click the **Add (+)** button.
 - In **New Extended Access List Object**, enter the name for the ACL (say, *DIA-FTD-Branch*), and click **Add**.
 - In the **Add Extended Access List Entry** dialog box, choose the required web-based applications from the **Application** tab:

Figure 2: Applications Tab

Add Extended Access List Entry

Action:

Logging:

Log Level:

Log Interval: 300 Sec.

Network Port **Application**

Application Filters Clear All Filters

Available Applications (3)

Selected Applications and Filters (2)

Applications

YouTube

Youtube Upload

YouTubeMp3

Cancel Add

On the 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 3: Extended ACL

New Extended Access List Object

Name

Entries (1)

Sequence	Action	Source	Source Port	Destination	Destination Port	Application
1	<input type="text" value="Allow"/>	any	Any	Any	Any	YouTube YouTubeMp3 Youtube Upload

Allow Overrides

Cancel Save

e) Click **Save**.

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

Step 3

Specify the egress interfaces:

- From the **Send To** and **Interface Ordering** drop-down lists, choose Egress Interfaces, and By Priority respectively.
- Under **Available Interfaces**, click the **+** button against the respective interface names to add *WAN1* and *WAN2*:

Figure 4: Configuring Policy Based Routing

Add Forwarding Actions

Match ACL:* +

Send To:*

Interface Ordering:*

Available Interfaces

Search by interface name

Priority	Interface
0	INSIDE1 +
0	INSIDE2 +
0	VTI01 +
0	VTI02 +

Selected Egress Interfaces*

Priority	Interface
10	WAN1 -
10	WAN2 -

- c) Click **Save**.

Step 4

Interface priority configuration:

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.

- Choose **Devices** > **Device Management**, and edit the branch threat defense.
- Set the priority for the interfaces. Click **Edit** against the interface and enter the priority value:

Figure 5: Setting Interface Priority

Edit Physical Interface

General IPv4 IPv6 Advanced Hardware Configuration FMC Access

Name: WAN1

Enabled
 Management Only

Description:

Mode: None

Security Zone: WAN

Interface ID: GigabitEthernet0/2

MTU: 1500
(64 - 9000)

Priority: 10
(0 - 65535)

Propagate Security Group Tag:

Cancel OK

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 6: Associating Interfaces with ECMP Zone

Device Routing Interfaces Inline Sets DHCP

Manage Virtual Routers

Global

Virtual Router Properties

- ECMP
- OSPF
- OSPFv3
- RIP
- Policy Based Routing

Equal-Cost Multipath Routing (ECMP).

All the interfaces belong to the ECMP must apply to the same access policies rules. You can add interfaces to this ECMP by clicking on Add button. ECMP can have up to 8 interfaces associated with it. All the interfaces in the ECMP must have a name and security level as this ECMP.

Add









Name	Interfaces
ECMP-WAN	WAN1, WAN2
ECMP-VTI	VTI01, VTI02

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 7: Configuring Static Routes for ECMP Zone Interfaces





Network	Interface	Leaked from Virtual Router	Gateway	Tunneled	Metric	Tracked	
+ Add Route							
▼ IPv4 Routes							
any-ipv4	VTI02	Global	192.168.102.21	false	1		 
any-ipv4	VTI01	Global	192.168.101.21	false	1		 
any-ipv4	WAN2	Global	10.10.1.65	false	10		 
any-ipv4	WAN1	Global	10.10.1.33	false	10		 

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 threat defense to ensure secured flow of traffic to the internet:
- Choose **Devices > Platform Settings**, and create a DNS policy on the branch threat defense.
 - To specify the trusted DNS, **Edit** the policy and click **DNS**.
 - 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.
 - Use the **Trusted DNS Servers** tab to provide specific DNS servers that you trust for the DNS resolution.
- Step 8** Save and 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 8: Site to Site VPN Settings

Node A	Node B	
+ Branch-Corporate-VTI  		
FTD-SJC / VTI01 / 192.168.101.20	FTD-BLR / VTI01 / 192.168.101.21	
FTD-SJC / VTI02 / 192.168.102.20	FTD-BLR / VTI02 / 192.168.102.21	

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

Configuration Example for PBR with Path Monitoring

This example details the configuration of 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*.

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 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 6](#)).
 - d) Click **Ok** and **Save**.
 - e) Repeat the same steps and configure the path monitoring settings for *ISP02* and *ISP03*.
- Step 2** Configure policy-based routing for a branch in an organization threat defense, select the ingress interfaces:
- a) Choose **Devices** > **Device Management**, and edit the 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 required web-based applications (example, WebEx Meetings) from the **Application** tab.
- Remember** On 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 same interface (*Inside1*) to route Office365 and network-based access control traffic:
- Create a match criteria object, example *PBR-Office365*, and select the Office365 application from the **Application** tab.
 - From the **Interface Ordering** drop-down list, choose By Minimal Round Trip Time.
 - Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.
 - Now, create a match criteria object, example *PBR-networks*, and specify the source and destination interface in the **Network** tab.
 - From the **Interface Ordering** drop-down list, choose By Minimal Packet Loss.
 - Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.
- Step 6** **Save and 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 10.

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*.

History for Policy Based Routing

Table 1:

Feature	Minimum Management Center	Minimum Threat Defense	Details
Identity and SGT based PBR Policy	7.4.0	7.4.0	<p>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.</p> <p>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.</p>
HTTP based Path Monitoring	7.4.0	7.2.0	<p>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.</p> <p>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.</p>

Feature	Minimum Management Center	Minimum Threat Defense	Details
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 autoconfigures PBR to route the SF-Tunnel traffic from the <i>tapnlp</i> (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.
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