
# Configuring Policy-Based Routing

This chapter describes how to configure policy-based routing on the Cisco NX-OS device.

This chapter includes the following sections:



## About Policy-Based Routing

With policy-based routing, you can configure a defined policy for IPv4 and IPv6 traffic flows that lessens the reliance on routes derived from routing protocols. All packets received on an interface with policy-based routing enabled are passed through enhanced packet filters or route maps. The route maps dictate the policy that determines where to forward packets.

Policy-based routing includes the following features:

- Source-based routing—Routes traffic that originates from different sets of users through different connections across the policy routers.
- Quality of Service (QoS)—Differentiates traffic by setting the precedence or type of service (ToS) values in the IP packet headers at the periphery of the network and leveraging queuing mechanisms to prioritize traffic in the core or backbone of the network (see the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide*).
- Load sharing—Distributes traffic among multiple paths based on the traffic characteristics.

# Policy Route Maps

Each entry in a route map contains a combination of match and set statements. The match statements define the criteria for whether appropriate packets meet the particular policy (that is, the conditions to be met). The set clauses explain how the packets should be routed once they have met the match criteria.

You can mark the route-map statements as permit or deny. You can interpret the statements as follows:

- If the statement is marked as permit and the packets meet the match criteria, the set clause is applied. One of these actions involves choosing the next hop.
- If a statement is marked as deny, the packets that meet the match criteria are sent back through the normal forwarding channels, and destination-based routing is performed.
- If the statement is marked as permit and the packets do not match any route-map statements, the packets are sent back through the normal forwarding channels, and destination-based routing is performed.

**Note** Policy-based routing is specified on the interface that receives the packets, not on the interface from which the packets are sent.

# Set Criteria for Policy-Based Routing

The Cisco Nexus 9000 Series switches support the following **set** commands for route maps used in policy-based routing:

- **set** {**ip** | **ipv6**} **next-hop** *address1* [*address2...*] [**load-share**]
- **set interface null0**

These **set** commands are mutually exclusive within the route-map sequence.

In the first command, the IP address specifies the adjacent next-hop router in the path toward the destination to which the packets should be forwarded. The first IP address associated with a currently up connected interface is used to route the packets.

**Note** You can optionally configure this command for next-hop addresses to load balance traffic for up to 32 IP addresses. In this case, Cisco NX-OS sends all traffic for each IP flow to a particular IP next-hop address.

If the packets do not meet any of the defined match criteria, those packets are routed through the normal destination-based routing process.

# Route-Map Processing Logic

When a packet is received on an interface that is configured with a route map, the forwarding logic processes each route-map statement according to the sequence number.

If the route-map statement encountered is a **route-map...permit** statement, the packet is matched against the criteria in the **match** command. This command may refer to an ACL that has one or more access control entries (ACEs). If the packet matches the permit ACEs in the ACL, the policy-based routing logic executes the action specified by the **set** command on the packet.

If the route-map statement encountered is a **route-map... deny** statement, the packet is matched against the criteria in the **match** command. This command may refer to an ACL that has one or more ACEs. If the packet matches the permit ACEs in the ACL, policy-based routing processing terminates, and the packet is routed using the default IP routing table.

**Note** The **set** command has no effect inside a **route-map... deny** statement.

If the route-map configuration does not contain a match statement, the policy-based routing logic executes the action specified by the **set** command on the packet. All packets are routed using policy-based routing.

If the route-map configuration references a match statement but the match statement references a non-existing ACL or an existing ACL without any access control entries (ACEs), the packet is routed using the default routing table.

If the next-hop specified in the **set** {**ip** | **ipv6**} **next-hop** command is down, is not reachable, or is removed, the packet is routed using the default routing table.

## Policy-Based Routing Filtering Options

You can identify traffic by using additional options. The following list includes most but not all additional filtering options.

Policy-based routing ACLs support the following additional filtering options:

- Layer 3 source and/or destination address
- TCP and UDP ports
- Precedence level*
- Differentiated Services Code Point (DSCP) value*
- TCP packets with the ACK, FIN, PSH, RST, SYN, or URG bit set*
- Established TCP connections*
- Packet length*

*Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards do not support these filtering options.

# Licensing Requirements for Policy-Based Routing

The following table shows the licensing requirements for this feature:

| Product | License Requirement |
|---|---|
| Cisco NX-OS | Policy-based routing requires an Enterprise Services license. Any feature not included in a license package is bundled with the nx-os image and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the *Cisco NX-OS Licensing Guide*. |

# Prerequisites for Policy-Based Routing

Policy-based routing has the following prerequisites:

- Assign an IP address on the interface and bring the interface up before you apply a route map on the interface for policy-based routing.

# Guidelines and Limitations

Policy-based routing has the following configuration guidelines and limitations:

- A policy-based routing route map can have only one match statement per route-map statement.
- A policy-based routing route map can have only one set statement per route-map statement, unless you are using IP SLA policy-based routing. For information on IP SLA policy-based routing, see the *Cisco Nexus 9000 Series NX-OS IP SLAs Configuration Guide.*

  **Note** Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards do not support IP SLA.

- A **match** command cannot refer to more than one ACL in a route map used for policy-based routing.
- The same route map can be shared among different interfaces for policy-based routing as long as the interfaces belong to the same virtual routing and forwarding (VRF) instance.
- Using a prefix list as a match criteria is not supported. Do not use a prefix list in a policy-based routing route map.
- Policy-based routing supports only unicast traffic. Multicast traffic is not supported.
- Policy-based routing is not supported with inbound traffic on FEX ports.
- Policy-based routing is not supported on FEX ports for Cisco Nexus 9300-EX Series switches.
- Only Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards support policy-based routing with Layer 3 port-channel subinterfaces.
- An ACL used in a policy-based routing route map cannot include deny access control entries (ACEs).
- Policy-based routing is supported only in the default system routing mode.
- The Cisco Nexus 9000 Series switches do not support the **set vrf** and **set default next-hop** commands.
- Policy-based routing traffic cannot be balanced if the next hop is recursive over ECMP paths. Instead, use the **set** {**ip** | **ipv6**} **next-hop** *ip-address* **load-share** command to specify the adjacent next hops.
- When you configure multiple features on an interface (such as PBR and ingress ACL), the ACLs for those features are merged for TCAM optimization. As a result, statistics are not supported.
- For PBR with VXLAN, the **load-share** keyword is not required.

  **Note** Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards do not support policy-based routing with VXLAN.

- Beginning with Cisco NX-OS Release 6.1(2)I3(2), the Cisco Nexus 9000 Series switches support policy-based ACLs (PBACLs), also referred to as object-group ACLs. For more information, see the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.

**Note** Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards do not support PBACLs.

- Beginning with Cisco NX-OS 7.0(3)I5(1), Cisco Nexus 9200 and 9300-EX Series switches support IPv4 and IPv6 policy-based routing. Cisco Nexus 9500 Series switches with the X9732C-EX line card support only IPv4 policy-based routing.
- Cisco Nexus 9500 Series switches with the X97xx-EX line cards support IPv4 (and not IPv6) policy-based routing.
- Beginning with Cisco NX-OS Release 7.0(3)F3(3), Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards support IPv4 and IPv6 policy-based routing. For these line cards, PBR policy has a higher priority over attached and local routes. Explicit white listing might be required if protocol neighbors are directly attached.
- The following guidelines and limitations apply to PBR over VXLAN EVPN:
  - PBR over VXLAN EVPN is supported only for Cisco Nexus 9300-EX and 9300-FX platform switches.
  - PBR over VXLAN EVPN does not support the following features: IP SLAs, VTEP ECMP, and the **load-share** keyword in the **set** {**ip** | **ipv6**} **next-hop** *ip-address* command.

# Default Settings

Table 16-1 lists the default settings for policy-based routing.

*Table 16-1 Default Policy-Based Routing Parameters*

| Parameters | Default |
|---|---|
| Policy-based routing | Disabled |

# Configuring Policy-Based Routing

This section includes the following topics:



## Enabling the Policy-Based Routing Feature

You must enable the policy-based routing feature before you can configure a route policy.

### SUMMARY STEPS

1. **configure terminal**

2. [**no**] **feature pbr**
3. (Optional) **show feature**
4. (Optional) **copy running-config startup-config**

### DETAILED STEPS

| | Command | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br><br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | [**no**] **feature pbr**<br><br>**Example:**<br>`switch(config)# feature pbr` | Enables the policy-based routing feature.<br>Use the **no** form of this command to disable the policy-based routing feature.<br>**Note** The **no feature pbr** command removes the policies applied under the interfaces. It does not remove the ACL or route-map configuration nor does it create a system checkpoint. |
| **Step 3** | **show feature**<br><br>**Example:**<br>`switch(config)# show feature` | (Optional) Displays enabled and disabled features. |
| **Step 4** | **copy running-config startup-config**<br><br>**Example:**<br>`switch(config)# copy running-config startup-config` | (Optional) Saves this configuration change. |

# Configuring a Route Policy

You can use route maps in policy-based routing to assign routing policies to the inbound interface. Cisco NX-OS routes the packets as soon as it finds a next hop and an interface.

### BEFORE YOU BEGIN

For switches other than the Cisco Nexus 9508 with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards, you must configure the IPv6 RACL TCAM region (using TCAM carving) before you apply the policy-based routing policy for IPv6 traffic. For instructions, see the "Configuring ACL TCAM Region Sizes" and "Configuring TCAM Carving - For Cisco NX-OS Release 6.1(2)I2(1) and Later Releases" sections in the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.

**Note** The switch has an IPv4 RACL TCAM region by default for IPv4 traffic.

### SUMMARY STEPS

1. **configure terminal**

2. **interface** *type slot/port*
3. {**ip** | **ipv6**} **policy route-map** *map-name*
4. **route-map** *map-name* [**permit** | **deny**] [*seq*]
5. **match** {**ip** | **ipv6**} **address access-list-name** *name* [*name...*]
6. (Optional) **set ip next-hop** *address1* [*address2...*] [**load-share**]
7. (Optional) **set ipv6 next-hop** *address1* [*address2...*] [**load-share**]
8. (Optional) **set interface null0**
9. (Optional) **copy running-config startup-config**

## DETAILED STEPS

| | Command | Purpose |
|---|---|---|
| **Step 1** | **configure terminal**<br><br>**Example:**<br>`switch# configure terminal`<br>`switch(config)#` | Enters global configuration mode. |
| **Step 2** | **interface** *type slot/port*<br><br>**Example:**<br>`switch(config)# interface ethernet 1/2`<br>`switch(config-if)#` | Enters interface configuration mode. |
| **Step 3** | {**ip** \| **ipv6**} **policy route-map** *map-name*<br><br>**Example**:<br>`switch(config-if)# ip policy route-map Testmap` | Assigns a route map for IPv4 or IPv6 policy-based routing to the interface. |
| **Step 4** | **route-map** *map-name* [**permit** \| **deny**] [*seq*]<br><br>**Example:**<br>`switch(config-if)# route-map Testmap`<br>`switch(config-route-map)#` | Creates a route map or enters route-map configuration mode for an existing route map. Use *seq* to order the entries in a route map. |
| **Step 5** | **match** {**ip** \| **ipv6**} **address access-list-name** *name* [*name...*]<br><br>**Example**:<br>`switch(config-route-map)# match ip address access-list-name ACL1` | Matches an IPv4 or IPv6 address against one or more IP or IPv6 access control lists (ACLs). This command is used for policy-based routing and is ignored by route filtering or redistribution. |
| **Step 6** | **set ip next-hop** *address1* [*address2...*] [**load-share**]<br><br>**Example**:<br>`switch(config-route-map)# set ip next-hop 192.0.2.1` | (Optional) Sets the IPv4 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured.<br><br>Use the optional **load-share** keyword to load balance traffic across a maximum of 32 next-hop addresses. For PBR with VXLAN, the **load-share** keyword is not required. |

| | Command | Purpose |
|---|---|---|
| Step 7 | **set ipv6 next-hop** *address1* [*address2...*] [**load-share**]<br><br>**Example:**<br>`switch(config-route-map)# set ipv6 next-hop 2001:0DB8::1` | (Optional) Sets the IPv6 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured.<br><br>Use the optional **load-share** keyword to load balance traffic across a maximum of 32 next-hop addresses. For PBR with VXLAN, the **load-share** keyword is not required. |
| Step 8 | **set interface null0**<br><br>**Example:**<br>`switch(config-route-map)# set interface null0` | (Optional)<br><br>Sets the interface used for routing. Use the **null0** interface to drop packets. |
| Step 9 | **copy running-config startup-config**<br><br>**Example:**<br>`switch(config-route-map)# copy running-config startup-config` | (Optional) Saves this configuration change. |

# Verifying the Policy-Based Routing Configuration

To display the policy-based routing configuration information, perform one of the following tasks:

| Command | Purpose |
|---|---|
| **show** [**ip** \| **ipv6**] **policy** [*name*] | Displays information about an IPv4 or IPv6 policy. |
| **show route-map** [*name*] **pbr-statistics** | Displays policy statistics.<br><br>**Note** Cisco Nexus 9508 switches with N9K-X9636C-R, N9K-X9636C-RX, and N9K-X9636Q-R line cards do not support PBR statistics. |

Use the **route-map** *map-name* **pbr-statistics** command to enable policy statistics. Use the **clear route-map** *map-name* **pbr-statistics** command to clear these policy statistics.

# Configuration Examples for Policy-Based Routing

This example shows how to configure a simple route policy on an interface:

```
ip access-list pbr-sample
permit tcp host 10.1.1.1 host 192.168.2.1 eq 80
!
route-map pbr-sample
match ip address pbr-sample
set ip next-hop 192.168.1.1
!
route-map pbr-sample pbr-statistics
interface ethernet 1/2
ip policy route-map pbr-sample
```

The following output verifies this configuration:

```
switch# show route-map pbr-sample
route-map pbr-sample, permit, sequence 10
 Match clauses:
  ip address (access-lists): pbr-sample
 Set clauses:
  ip next-hop 192.168.1.1

switch# show route-map pbr-sample pbr-statistics
route-map pbr-sample, permit, sequence 10
 Policy routing matches: 84 packets
Default routing: 233 packets

switch# show ip policy
Interface   Route-map  Status  VRF-Name
Ethernet1/2 pbr-sample Active  --
```

# Related Documents

| Related Topic | Document Title |
|---|---|
| IP SLA PBR object tracking | *Cisco Nexus 9000 Series NX-OS IP SLAs Configuration Guide* |
| Troubleshooting information | *Cisco Nexus 9000 Series NX-OS Troubleshooting Guide* |