Configuring Policy-Based Routing

This chapter contains the following sections:

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information"chapter or the Feature History table in this chapter.

Information About Policy Based Routing

Policy-based routing allows you to configure a defined policy for IPv4 and IPv6 traffic flows, lessening 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, determining where to forward packets.

Route maps are composed of match and set statements that you can mark as permit or deny. You can interpret the statements as follows:

- If the packets match any route map statements, all the set statements are 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.

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 7000 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. If the statement is marked as a deny, the packets that meet the match criteria are sent back through the normal forwarding channels (destination-based routing is performed). If the statement is marked as permit and the packets meet the match criteria, all the set clauses are applied. If the statement is marked as permit and the packets do not meet the match criteria, those packets are also forwarded through the normal routing channel.

Note Policy 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 set criteria in a route map is evaluated in the order listed in the route map. Set criteria specific to route maps used for policy-based routing are as follows:

1. List of interfaces through which the packets can be routed—If more than one interface is specified, the first interface that is found to be up is used for forwarding the packets.
2. List of specified IP addresses—The IP address can specify 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 the set criteria for next-hop addresses to load balance traffic across up to 16 IP addresses. In this case, Cisco NX-OS sends all traffic for each IP flow to a particular IP next-hop address.

3. List of default interfaces—If there is no explicit route available to the destination address of the packet being considered for policy routing, the route map routes it to the first up interface in the list of specified default interfaces.
4. List of default next-hop IP addresses—Route to the interface or the next-hop address specified by this set statement only if there is no explicit route for the destination address of the packet in the routing table.

Note You can optionally configure the set criteria for the default next-hop addresses to load balance traffic across a maximum of 16 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.

Local Policy Routing

Local policy routing allows you to apply a route map to local (device-generated) traffic. All packets originating on the device that are not normally policy routed are subject to local policy routing.

Route Map Support Matrix for Policy-Based Routing

The following tables include the configurable match and set statements for policy-based routing on Cisco Nexus 70xx and 77xx Series switches running the latest shipping release. For specific release information, see the Feature History for Policy-Based Routing, on page 12.

The following legend applies to the tables:

- Yes—The statement is supported for policy-based routing.
- No—The statement is not supported for policy-based routing.
- If a statement does not apply for policy-based routing, there is an em dash (—) in the column next to the statement.
- Where clarification is required, information is added in the appropriate row/column.

Table 1: SET Route Map Statements for Policy-Based Routing

SET Route Map Statement	Policy-Based Routing (PBR)
IPv4 Next Hop	Yes
IPv6 Next Hop	Yes
Default IPv4 Next Hop	Yes
Default IPv6 Next Hop	Yes
IPv4 Next Hop Verify Availability	Yes
IPv6 Next Hop Verify Availability	Yes

SET Route Map Statement	Policy-Based Routing (PBR)
Default IPv4 Next Hop Verify Availability	Yes
Default IPv6 Next Hop Verify Availability	Yes
Interface null0	Yes
VRF	Yes
IPv4 Precedence	Yes
IPv6 Precedence	Yes
Interface, GRE Ethernet	No

Table 2: MATCH Route Map Statements for Policy-Based Routing

MATCH Route Map Statement	Policy-Based Routing (PBR)
Tag	Yes
Packet Length	Yes
VLAN ID	Yes
MAC ACL	Yes
IPv4 Prefix List	No
IPv6 Prefix List	No
IP ACL	Yes

Prerequisites for Policy-Based Routing

Policy-based routing has the following prerequisites:

- Install the correct license.
- You must enable policy-based routing.
- 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.
- If you configure VDCs, install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*).

Guidelines and Limitations for Policy-Based Routing

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

- Cisco NX-OS uses recursive next hops. You do not need to enter any commands for recursive next hops like you do for Cisco IOS.
- A policy-based routing route map can have only one match or set statement per route-map statement.
- 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.
- Prior to Cisco NX-OS Release 8.0(1) setting a tunnel interface or an IP address via a tunnel interface as a next hop in a policy-based routing policy is not supported. Applying policy-based routing or **ip policy route-map** on tunnel interfaces is also not supported. From Cisco NX-OS Release 8.0(1) onwards GRE next hop is supported on policy-based routing.
- Policy-based routing is not supported with inbound traffic on FEX ports.
- Using a prefix-list as a match criteria is not supported. Do not use a prefix-list in a policy-based routing route-map.
- Beginning with Cisco NX-OS Release 5.2(4), policy-based routing and WCCPv2 are supported on the same interface. However, policy-based routing with statistics and WCCPv2 is supported on the same interface only if bank chaining is disabled.
- Beginning with Cisco NX-OS Release 6.1(3), you can configure the device to support deny access control entries (ACEs) in a sequence for the following sequence-based features: VACLs and QoS. For more information, see the "Configuring VLAN ACLs" chapter in the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.
- PBR marks the next-hop as down even when the next hop and the corresponding tracks are up. This issue is due to the RPM that does not effectively process the tracks.

  Currently the object tracking manager (OTM) does not support forward referencing for track objects. Track objects must be created in the OTM before they are used in any configuration.

  Perform the following steps to configure the track objects with RPM or PBR so that the PBR next-hop issue does not occur:

  1. Create the track object in OTM using the **track** *<object id>* command.
  2. Use the configured track object in a route map using the **set ip next-hop verify-availability** *<ip1>* **track** *<object id>* command.
  3. Apply the route map to an interface using the **ip policy route-map** *<map-name>* command.

Default Settings for Policy-Based Routing

Table 3: Default Policy-Based Routing Parameters

Parameters	Default
Policy-based routing	Disabled

Configuring Policy-Based Routing

Enabling the Policy-Based Routing

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

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# [**no**] **feature pbr**	Enables the policy-based routing feature. Use the **no feature pbr** command to disable the policy-based routing feature and remove all associated configuration.
Step 3	(Optional) switch(config)# **show feature**	Displays enabled and disabled features.
Step 4	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

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 packet as soon as it finds a next hop and an interface.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot/port*	Enters interface configuration mode.
Step 3	switch(config-if)# **ip policy route-map** *map-name*	Assigns a route map for IPv4 policy-based routing to the interface.
Step 4	switch(config-if)# **ipv6 policy route-map** *map-name*	Assigns a route map for IPv6 policy-based routing to the interface.
Step 5	(Optional) switch(config-route-map)# **end**	Exits route-map configuration mode and enters the privileged executive mode.

	Command or Action	Purpose
Step 6	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to add a route map to an interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip policy route-map Testmap
switch(config)# exit
switch(config)# copy running-config startup-config
```

You can configure the following optional match parameters for route maps in route-map configuration mode:

Command	Purpose
match ip address access-list-name *name [name...]* Example: `switch(config-route-map)# match ip address access-list-name ACL1`	Matches an IPv4 address against one or more IP access control lists (ACLs). This command is used for policy-based routing and is ignored by route filtering or redistribution.
match ipv6 address access-list-name *name [name...]* Example: `switch(config-route-map)# match ipv6 address access-list-name ACLv6`	Matches an IPv6 address against one or more IPv6 ACLs. This command is used for policy-based routing and is ignored by route filtering or redistribution.
match length *min max* Example: `switch(config-route-map)# match length 64 1500`	Matches against the length of the packet. This command is used for policy-based routing.
match mac-list *maclist [...maclist]*	Matches against a list of MAC addresses. This command is used for policy-based routing.
match metric *metric-value [+- deviation-number] [...metric-value [+- deviation-number]]* [*+- deviation-number*] [*... metric-value +- deviation-number*]] Example: `switch(config-route-map)# match metric 10`	Matches against the routing protocol metric. This command is used for policy-based routing.
match vlan *vlan-range* `switch(config-route-map)# match vlan 64`	Matches against the VLAN ID of the packet. This command is used for policy-based routing.

You can configure the following optional set parameters for route maps in route-map configuration mode:

Command	Purpose

set ip next-hop *address2* [*address2...*] [**load-share** | **peer-address** | **unchanged** | **verify-availability**]
Example:
`switch(config-route-map)#`
`set ip next-hop 192.0.2.1`
	Sets the IPv4 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only.
	• Use the optional **load-share** keyword to load balance traffic across a maximum of 16 next-hop addresses.
	• Use the optional **peer-address** keyword to the next hop to be the Border Gateway Protocol (BGP) peering address.
	• Use the optional **unchanged** keyword to specifiy that the next-hop attribute in the BGP update to the eBGP peer is unmodified.
	• Use the optional **verify-availability** keyword to verify the reachability of the tracked object.

set ip default next-hop *address2* [*address2...*] [**load-share** | **verify-availability**]
Example:
`switch(config-route-map)#`
`set ip default next-hop 192.0.2.2`
	Sets the IPv4 next-hop address for policy-based routing when there is no explicit route to a destination. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only.
	• Use the optional **load-share** keyword to load balance traffic across a maximum of 16 next-hop addresses.
	• Use the optional **verify-availability** keyword to verify the reachability of the tracked object.
	Note For software-forwarded traffic, the route that is present in the unicast routing table (of the VRF in which packet was received) for packet-specified destination takes preference over what is specified in **set ip default next-hop** command, when there is condition match. Even if there is a default route present in the VRF, that default route overrides what is set in the command. This applies to software-forwarded traffic only.

Command	Purpose
set ipv6 next-hop *address2* [*address2...*] [**load-share** | **peer-address** | **unchanged** | **verify-availability**] Example: `switch(config-route-map)#` **`set ipv6 next-hop 2001:0DB8::1`**	Sets the IPv6 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only. • Use the optional **load-share** keyword to load balance traffic across a maximum of 16 next-hop addresses. • Use the optional **peer-address** keyword to the next hop to be the Border Gateway Protocol (BGP) peering address. • Use the optional **unchanged** keyword to specifiy that the next-hop attribute in the BGP update to the eBGP peer is unmodified. • Use the optional **verify-availability** keyword to verify the reachability of the tracked object.
set ipv6 default next-hop *address2* [*address2...*] [**load-share** | **verify-availability**] Example: `switch(config-route-map)#` **`set ipv6 default next-hop 2001:0DB8::2`**	Sets the IPv6 next-hop address for policy-based routing when there is no explicit route to a destination. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only. • Use the optional **load-share** keyword to load balance traffic across a maximum of 16 next-hop addresses. • Use the optional **verify-availability** keyword to verify the reachability of the tracked object.
set ip precedence *precedence-value* Example: `switch(config-route-map)#` **`set ip precedence highv4`**	Sets the precedence value in the IPv4 packet header.
set ipv6 precedence *precedence-value* Example: `switch(config-route-map)#` **`set ipv6 precedence highv6`**	Sets the precedence value in the IPv6 packet header.
set ipv6 precedence address prefix-list *prefix-list-name* Example: `switch(config-route-map)#` **`set ipv6 precedence address prefix-list acl1`**	Sets the IPv6 map routes to be injected.

Command	Purpose
set interface {*null10* | *tunnel-te*} Example: `switch(config-route-map)#` **`set interface null0`**	Sets the interface used for routing. Use the **null0** interface to drop packets. Use the **tunnel-te** interface to forward packets on the MPLS TE tunnel.
set vrf *vrf-name* Example: `switch(config-route-map)#` **`set vrf MainVRF`**	Sets the VRF for next-hop resolution.

Configuring Local Policy Routing

You can enable local policy routing for packets generated by the device and specify which route map the device should use.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# {**ip** | **ipv6**} **local policy route-map** *map-name*	Configures IPv4 or IPv6 local policy route maps for packets generated by the device.
Step 3	(Optional) **show** {**ip** | **ipv6**} **local policy**	Displays the route map used for IPv4 or IPv6 local policy routing.
Step 4	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring a Deny ACE

Beginning with Cisco NX-OS Release 6.1(3), you can configure the device to support deny access control entries (ACEs) in a sequence for the following sequence-based features: VACL and Quality of service (QoS).

When deny ACEs are enabled, the traffic that matches a deny ACE (an ACL rule with the **deny** keyword) in a class-map-acl is recursively matched against subsequent class-map-acls until it hits a permit ACE.

Note In earlier releases, an ACL used in a policy-based routing route map cannot include a deny statement.

Before you begin

Ensure that you are in the default or admin VDC.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# [**no**] **hardware access-list allow deny ace**	Enables deny ACEs in a sequence. The **no** form of the command disables deny ACEs.
Step 3	(Optional) **show running-config aclmgr**	Displays the ACL configuration.
Step 4	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Verifying the Policy-Based Routing Configuration

To display 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 {**ip** | **ipv6**} **local policy** [**vrf** *name*]	Displays the route map used for IPv4 or IPv6 local policy routing.
show route-map [*name*] **pbr-statistics**	Displays policy statistics.

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

Configuration Examples for Policy Based-Routing

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

```
feature pbr
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:

```
n7000# 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

n7000# show route-map pbr-sample pbr-statistics

route-map pbr-sample, permit, sequence 10
 Policy routing matches: 84 packets

Default routing: 233 packets
```

Configuration Example for Local Policy Routing

The following example sends packets with a destination IP address matching that allowed by extended access list 131 to the router at IP address 172.30.3.20:

```
ip local policy route-map xyz
!
route-map xyz
match ip address 131
set ip next-hop 172.30.3.20
```

Related Documents for Policy-Based Routing

Related Topic	Document Title
Policy-based routing CLI commands	*Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*
VDCs and VRFs	*Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference*

Standards for Policy-Based Routing

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for Policy-Based Routing

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 4: Feature History for Policy-Based Routing

Feature Name	Release	Feature Information
Route map support matrix	6.2(2)	Added the route map support matrix for policy-based routing.

Feature Name	Release	Feature Information
Policy-based routing	6.1(3)	Added support for deny access control entries (ACEs) in a sequence for the following sequence-based features: VACLs, policy-based routing, and QoS.
Policy-based routing	5.2(4)	Added support for policy-based routing and WCCPv2 on the same interface if bank chaining is disabled.
Interfaces	5.2(1)	Added support for **set interface route-map** command.
IPv6 policies	4.2(1)	Added support for IPv6 policies.
Policy-based routing	4.0(1)	This feature was introduced.