

Deploying Route Maps on Cisco Nexus Switches

White Paper



Contents

What You Will Learn	3
Introduction.....	3
Policy-Based Routing.....	3
Route Redistribution	4
Table Map	5
OSPF Table Map	6
IS-IS Table Map	7
BGP Table Map	7
EIGRP Table Map	8
Route Map for BGP Neighbor	8
Route Map for EIGRP Distribution List	9
Support Matrix	9

What You Will Learn

This document provides information about supported **route-map** statements and the resulting actions on Cisco Nexus® switches. It also describes supported **match** and **set** statements for routing protocols and policy-based routing (PBR), and it includes configuration examples. You can use this document as a quick reference for supported **route-map** statements on Cisco Nexus switches.

Introduction

A route map is a powerful tool used by network administrators to define customized routing policies based on the match and set criteria that can change packet or route information. This document provides details about the use of the route-map feature in Cisco® Nexus NX-OS Software.

Route maps are typically used to:

- Implement PBR
- Redistribute routes among routing protocols
- Filter routes or change route distance using table maps
- Modify attributes of a routing protocol (metrics, tags, distance, etc.)

Policy-Based Routing


PBR allows you to configure a defined policy for IPv4 and IPv6 traffic flows, thereby reducing reliance on routes derived from routing protocols. All packets received on an interface with PBR enabled are passed through enhanced packet filters or route maps. The route maps dictate the policy, determining where the packets are forwarded.

Route maps are composed of **match** and **set** statements that can be marked as **permit** or **deny**:

- 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.

Figure 1 shows a PBR configuration example.

Figure 1. PBR Configuration Example



Use the following steps to configure PBR:

1. Configure the PBR feature.

```
feature pbr
```

2. Configure a route map with match conditions to select packets for PBR. Configure the IPv4 and IPv6 next hop to specify the destination next hop. Precedence and next-hop availability using tracking can also be configured if required.

```
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
!
```

3. Apply the route map on the incoming interface.

```
interface ethernet 1/1
  ip policy route-map pbr-sample
```

4. Optionally, enable statistics for PBR.

```
route-map pbr-sample pbr-statistics
```

PBR implementations in Cisco NX-OS differ mainly as follows:

- The PBR feature is entirely implemented in hardware for both IPv4 and IPv6, resulting in little performance overhead on the supervisor CPU.
- PBR provides the capability to distribute traffic between virtual routing and forwarding (VRF) instances using the **set vrf** command.
- PBR statistics can be enabled using the **pbr-statistics** command for the route map.

Note that access control lists (ACLs) are supported for PBR, but not prefix lists.

Route Redistribution

Route redistribution allows routes from one routing protocol to be advertised to another routing protocol. NX-OS provides the capability to use a route map during redistribution for better control in importing the routes.

Route maps are applied to redistribution work with two routing protocols:

- The protocol that provides the original routing information
- The protocol to which the routing information is redistributed

Each routing protocol can support its own set of route attributes.

In redistribution route-map configuration:

- Route-map **match** commands verify the attributes of a route that are supported by the protocol that supplied the original route for redistribution.

- Route-map **set** commands modify the attributes of routes that are supported by the protocol to which the routes are redistributed.

Note that:

- Route maps are generic mechanisms. In some cases, certain set and match statements may not make sense. For example, the match length command in the redistribution route map has no effect on redistribution. However, it remains in the configuration of the route map and can be seen in the running configuration of the router. Such cases are marked “-“ in the support matrix.
- NX-OS enforces the use of the route map during redistribution. A route map without any **match** statement can be used to import all the routes from the source protocol.

The following configuration example shows route redistribution from Open Shortest Path First (OSPF) process 100 to the Border Gateway Protocol (BGP) that matches all OSPF routes with a tag of 100 and redistributes these routes to BGP with a metric (multi-exit discriminator [MED]) value of 100.

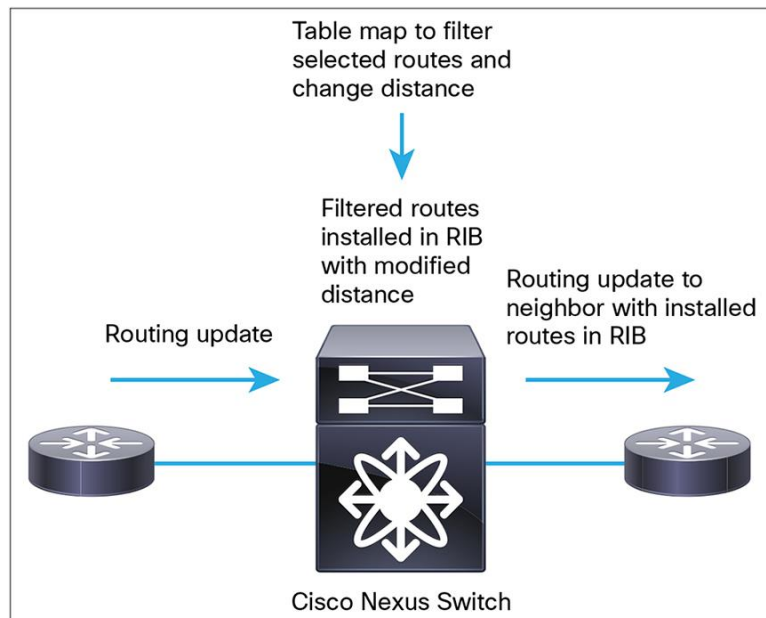
<pre>router bgp 200 address-family ipv4 unicast network 5.5.5.5/32 redistribute ospf 100 route-map tag</pre>	<pre>route-map test permit 10 match tag 100 set metric 100</pre>
--	--

Redistribution among other protocols can be configured similarly.

Table Map

A table map is a unique feature of NX-OS that allows the network administrator to filter routes or selectively modify the distance of the routes before the routes are sent to routing information base (RIB). The table map uses the route map to select routes based on a wide variety of parameters: metrics, level, type, next hop, outgoing interface, etc. (Figure 2).

Figure 2. Table Map



The table-map feature is available for:

- OSPFv2 and v3
- Enhanced Interior Gateway Routing Protocol (EIGRP) v4 and v6
- Intermediate System-to-Intermediate System (IS-IS) Protocol v4 and v6
- BGP

The table map allows more scalable solutions by preserving valuable RIB memory. The table map can be used for both filtering routes and changing the route distance.

- When the **table-map** command is used without the **filter** keyword, the route map referenced in the **table-map** command is used to set certain properties (such as the traffic index) of the routes for installation in the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.
- When the **table-map** command is used with the **filter** keyword, the route map referenced is also used to control whether or not a BGP route is to be downloaded to the RIB.

OSPF Table Map

OSPF supports a table map to filter and change the distances of IPv4 and IPv6 prefixes. This section provides a table-map configuration example for both OSPFv2 and OSPFv3.

The following example shows the configuration to change the OSPFv2 administrative distance for inter area routes to 150 and for external routes to 200.

```
route-map foo permit 10
  match route-type inter-area
  set distance 150

route-map foo permit 20
  match route-type external
  set distance 200

router ospf 201
  table-map foo
```

The following example shows the configuration of the **table-map** command with the **filter** keyword to use a route map (**Filter-OSPF**) that removes the next-hop path that is learned through VLAN 10, but not the next-hop path that is learned through VLAN 20.

```
route-map Filter-OSPF 10 deny
  match interface VLAN 10
  route-map Filter-OSPF 20 permit

route ospf p1
  table-map Filter-OSPF filter
```

The table-map configuration for OSPFv3 changes the distance and filter routes and is similar to the configuration for OSPFv2. The following example uses **route-map foo** to change the distance of inter area routes to 150 and distance of external routes to 200.

<pre>route-map foo permit 10 match route-type inter-area set distance 150 route-map foo permit 20 match route-type external set distance 200</pre>	<pre>router ospfv3 201 address-family ipv6 unicast table-map foo exit</pre>
--	---

IS-IS Table Map

The table map for IS-IS supports IPv4 and IPv6 route filtering and distance modification. The following example shows the table-map configuration of IS-IS on the basis of the outgoing interface value.

<pre>route-map isis-distance match interface eth1/1 set distance 10 route-map isis-filter match interface eth1/1</pre>	<pre>router isis 100 address-family ipv4 unicast table-map test address-family ipv6 unicast table-map isis-filter filter</pre>
---	--

BGP Table Map

Table maps are used on the BGP route reflector to save resources. The role of a dedicated route reflector is to propagate BGP updates without participating in the actual forwarding of transit traffic. Therefore, the route reflector does not need to have all BGP routes downloaded to its RIB or forwarding information base (FIB). It is beneficial for the route reflector to preserve its resources by not processing and storing those routes.

By default, BGP routes are downloaded to the RIB. To save resources on a dedicated route reflector, you can reduce or prevent such downloading by configuring a table map.

The following example allows only routes that match the community list to be downloaded to the RIB.

<pre>ip community-list 100 permit 65510:100 ! route-map bgp-to-rib permit 10 match community 100 !</pre>	<pre>router bgp 65510 address-family ipv4 unicast table-map bgp-to-rib filter</pre>
--	---

The following example shows the configuration to change the distance of a route on the basis of the community list.

<pre>ip community-list 100 permit 65510:100 ! route-map bgp-to-rib permit 10 match community 100 set distance 10 !</pre>	<pre>router bgp 65510 address-family ipv4 unicast table-map bgp-to-rib filter</pre>
--	---

EIGRP Table Map

EIGRP supports table maps to filter and change the distance of both IPv4 and IPv6 prefixes. The following example configures the table map for EIGRP on the basis of the next-hop value.

<pre>route-map eigrp-filter match ip next-hop prefix-list ipv4- nexthop route-map eigrp-distance match ipv6 next-hop prefix-list ipv6- nexthop set distance 10</pre>	<pre>router eigrp 100 address-family ipv4 unicast table-map eigrp-filter filter address-family ipv6 unicast table-map eigrp-distance</pre>
--	--

Route Map for BGP Neighbor

BGP neighbors use route maps to match incoming and outgoing IPv4 and IPv6 routes on the basis of their attributes to filter or set route properties. Route maps can be applied to both dynamic peers (prefix-based peers) and fixed peers.

The following example shows the configuration to specify the next-hop self using the outbound route map for the static neighbor.

<pre>route-map setrrnh permit 10 set ip next-hop peer-address route-map setrrnhv6 permit 10 set ipv6 next-hop peer-address</pre>	<pre>router bgp 200 neighbor 192.0.2.12 remote-as 200 address-family ipv4 unicast route-map setrrnh out exit address-family ipv6 unicast route-map setrrnhv6 out</pre>
--	--

The following example configures the list of accepted autonomous system (AS) numbers and the autonomous-system-path list for dynamic neighbors.

```
route-map BGPPeers
  match as-number 64496, 64501-64510
  match as-number as-path-list List1, List2

router bgp 64496
neighbor 192.0.2.0/8 remote-as
  route-map BGPPeers
```

Route Map for EIGRP Distribution List

The EIGRP distribution list allows you to filter incoming IPv4 and IPv6 routes on an interface using a route map. This capability provides flexibility in matching and setting route attributes. The following example shows the configuration of an EIGRP distribution list.

<pre>route-map metric-range match metric 50 set tag 5 router eigrp 1</pre>	<pre>interface ethernet 1/2 ip router eigrp 1 ip distribute-list eigrp 1 route-map metric-range in ipv6 distribute-list eigrp 1 route- map metric-range in</pre>
--	--

Support Matrix

The route-map support matrix for routing protocols is available at the following link:

[Route Map Support Matrix for Routing Protocols](#)

Note: Some statements, such as site of origin (**soo**) and extended community route targets, apply to Cisco Nexus 7000 Series Switches only, because these features are not supported on fixed-platform switches: Cisco Nexus 5000 and 6000 Series Switches.

The route-map support matrix for PBR is available at the following link:

[Route Map Support Matrix for Policy-Based Routing](#)




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)