

Configure Route-Maps Applied with Redistribution Command

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Background Information](#)

[Route-Map Basics](#)

[Open Shortest Path First](#)

[Commands Used in Route-Maps Applied to Redistribution](#)

[Configure Unsupported Match and Set Commands in Route-Maps](#)

[Two-Protocol Nature of Route-Map Redistribution](#)

[Command Support Tables](#)

[Table 1 – Route-map Commands to Match Routes Installed into Routing Tables by Protocols](#)

[Table 2 – Commands that Modify Route Attributes During Redistribution into the Target Protocol](#)

[Conclusion](#)

[Related Information](#)

Introduction

This document describes how to configure route-maps that are applied with the **redistribute** command of dynamic routing protocols.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on the Cisco IOS® Software Release 12.3.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the [Cisco Technical Tips Conventions](#).

Background Information

This section provides an overview of route-maps used in Cisco IOS software.

Route-Map Basics

Route-maps have many features in common with widely known access control lists (ACLs). These are some of the traits common to both mechanisms:

- They are an ordered sequence of individual statements; each has a permit or deny result. Evaluation of ACL or route-maps consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found, and an action associated with the statement match is performed.
- They are generic mechanisms—criteria match, and match interpretation are dictated by the way they are applied. The same route-map applied to different tasks can have a different interpretation.

These are possible differences between route-maps and ACLs:

- Route-maps frequently use ACLs as the matched criteria.
- The main result from the evaluation of an access list is a yes or no answer. An ACL either permits or denies input data. Applied to redistribution, an ACL determines if a particular route can (route matches ACLs permit statement) or cannot (matches deny statement) be redistributed. Typical route-maps not only permit (some) redistributed routes but also modify information associated with the route when it is redistributed into another protocol.
- Route-maps are more flexible than ACLs and can verify routes based on criteria which ACLs cannot verify. For example, a route-map can verify if the type of route is internal or if it has a specific tag.
- Each ACL ends with an implicit deny statement, by design convention; there is no similar convention for route-maps. If the end of a route-map is reached when match attempts were made, the result depends on the specific application of the route-map. Fortunately, route-maps that are applied to redistribution behave the same way as ACLs: if the route does not match any clause in a route-map then the route redistribution is denied, as if the route-map contained deny statement at the end.

The dynamic protocol redistribute configuration command allows you to apply either an ACL or a route-map. The differences described in this section identify when to use a route-map in the redistribution process. Route-maps are preferred if you intend to either modify route information during redistribution or if you need more powerful capabilities to match capability than an ACL can provide. Conversely, if you simply need to selectively permit some routes based on their prefix or mask, Cisco recommends that you use an ACL (or equivalent prefix list) directly in the redistribute command. If you use a route-map to selectively permit some routes based on their prefix or mask, you typically use more configuration commands to achieve the same goal. Route maps are always applied for inbound traffic, and the route-map does not have any effect on the outbound traffic.

Open Shortest Path First

This is a typical Open Shortest Path First to Enhanced Interior Gateway Routing Protocol (OSPF-to-EIGRP) route-map, applied with a redistribute command:

```
!  
router eigrp 1  
  redistribute ospf 1 route-map ospf-to-eigrp  
  default-metric 20000 2000 255 1 1500  
!--- Output suppressed.
```

```

!
route-map ospf-to-eigrp deny 10
  match tag 6
  match route-type external type-2
!
route-map ospf-to-eigrp permit 20
  match ip address prefix-list pfx
  set metric 40000 1000 255 1 1500
!
route-map ospf-to-eigrp permit 30
  set tag 8
!

```

These are the important observations with this example:

- Route-map clauses are numbered. In this example, clauses have sequence numbers 10, 20, and 30. Sequence numbers allow you to do these actions:
 - Easily delete one specific clause but not affect other parts of the route-map.
 - Insert a new clause between two current clauses.

Cisco recommends that clauses are numbered in intervals of 10 to reserve number spaces to insert clauses in the future, if needed.

- Route-maps can have permit and deny clauses. In route-map ospf-to-eigrp, there is one deny clause (with sequence number 10) and two permit clauses. The deny clause rejects route matches from redistribution. Therefore, these rules apply:
 - If you use an ACL in a route-map permit clause, routes that are permitted by the ACL are redistributed.
 - If you use an ACL in a route-map deny clause, routes that are permitted by the ACL are not redistributed.
 - If you use an ACL in a route-map permit or deny clause, and the ACL denies a route, then the route-map clause match is not found, and the next route-map clause is evaluated.
- Each route-map clause has two types of commands:
 - **match** — Selects routes to which this clause must be applied.
 - **set** — Modifies information to redistributed into the target protocol.

For each redistributed route, the router first evaluates the match command of a clause in the route-map. If the matched criteria succeed, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes are modified by set commands. If the matched criteria fails then this clause is not applicable to the route, and Cisco IOS software proceeds to evaluate the route against the next clause in the route-map. Scan of the route-map continues until a clause is found whose match command(s) match the route or until the end of the route-map is reached.

- A match or set command in each clause can be missed or repeated several times, if one of these conditions exist:
 - If several match commands are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
 - If a match command refers to several objects in one command, either match (the logical OR algorithm is applied). For example, in the match ip address 101 121 commands, a route is permitted if it is permitted by access list 101 or access list 121.
 - If a match command is not present, all routes match the clause. In the previous example, all routes that reach clause 30 match; therefore, the end of the route-map is never reached.
 - If a set command is not present in a route-map permit clause, then the route is redistributed without modification of its current attributes.

Do not configure a set command in a deny route-map clause because the deny clause prohibits route redistribution. There is no information to modify.

A route-map clause without a matchset command performs an action. An empty permit clause allows a redistribution of the rest of the routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route-map is completely scanned but no explicit match is found).

Based on the information in this section, the previous OSPF-to-EIGRP route-map example does this:

- Prohibits the redistribution of all Type-2 external OSPF routes with tag 6 set.
- Redistributes into EIGRP all routes which match the pfx prefix list, with five metric values—40000, 1000, 255, 1, and 1500.
- Redistributes all other routes and sets their tags to 8 (the default).

Commands Used in Route-Maps Applied to Redistribution

This section contains these topics:

- [Configure Unsupported **match** and **set** Commands in Route-Maps](#)
- [Two-Protocol Nature of Route-Map Redistribution](#)
- [Command Support Tables](#)

Configure Unsupported Match and Set Commands in Route-Maps

Route-maps are generic mechanisms which you can use in many configurations, this includes the previously described **redistribute** command. For example, you can configure the **match length** command in a route-map for PBR to specify that a special action is performed when packets of a specific length are forwarded. However, you would not use the **match length** command in route-maps applied to redistribution.

You can configure **match** and **set** commands in a route-map that are not supported (or do not have effect) in a context where a route-map is applied (or intended to be applied at a later stage). An example of this situation can be the **match length** command used in a route-map applied to redistribution. In redistribution, a route-map is applied to each route installed into the routing table, by the protocol specified in the **redistribute** command. Therefore, when a router executes a route-map, the router interprets only the commands which make sense in the context of route-map application. In this example, the **match length** command mentioned in the redistribution route-map has no effect on redistribution. It remains in the configuration of the route-map, and it can be seen in the running configuration of the router. Route redistribution is not affected, however, whether this command is present in the route-map or not.

Therefore, the router allows the configuration of all kinds of **match** and **set** commands, but they must be applied logically to the situation. Otherwise, the configuration can be very confusing or can perform incorrect tasks.

Do not use commands with no effect in a route-map context — even if they seem harmless — because of these issues:

- Commands with no effect can obscure what you want to accomplish. They can cause confusion.
- Commands that are currently unsupported can become supported in future releases of Cisco IOS software. There can be unwanted changes in the route-map behavior after future software upgrades.

- Not all commands are completely harmless; for example, the **set metric +/-** command, which specifies the relative change of metric and which is used with BGP route advertisement. It can take the current metric of a route and increase or decrease it by a specified value before it propagates it.

The +/- form of this command is not supported by some protocols, see [EIGRP Route Map Support](#) and in other scenarios it is actually interpreted as the **set metric** command with the sign omitted. For example, consider this route-map:

```
!--- This redistribution route-map is very dangerous!  
route-map ospf-to-ospf permit 10  
  set metric +2  
!
```

This configuration appears to redistribute all routes from one OSPF process to another, while it increases the metric of all routes by two. Yet, it actually sets the metric of all routes to be the same, equal to 2. This is unexpected in the router configuration.

This route-map provides an even more counter-intuitive effect:

```
!--- This redistribution route-map is very dangerous!  
route-map ospf-to-ospf permit 10  
  set metric +2  
!
```

Rather than decrease the metric of redistributed routes, this configuration actually sets the metric to 367 (a positive value, because a negative metric is not possible when **set metric** is interpreted without the sign).

Two-Protocol Nature of Route-Map Redistribution

Route-maps 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 which 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.

The [Command Support Tables](#) section of this document lists the commands. They are categorized by **match** and **set** commands, to highlight the two-protocol nature of redistribution route-maps.

Command Support Tables

This section describes the commands that are supported in route-maps attached in the **redistribute** command. There are seven routing protocols from which routes can be redistributed; however, there are only five to which redistribution can take place. Connected and static routes are not dynamic routing protocols and can only provide information to be redistributed into other protocols.

This section does not include the **match** and **set** commands that are supported in Cisco IOS Software Release 12.3 route-maps but that are not applicable in the redistribution context.

Intermediate System-to-Intermediate System (IS-IS) and BGP can propagate information about Connectionless Network Service (CLNS) routes along with IP routes. To be thorough, the tables in this section also mention CLNS-related commands, which can be used in redistribution route-maps for these protocols.

You can use Routing Information Protocol (RIP), OSPF, IS-IS, and BGP to propagate IPv6 routes; redistribution route-maps for these protocols can contain IPv6-specific commands. The **match ip** and **set ip** commands are specific for redistribution of IPv4 prefixes. The **match ipv6** and **set ipv6** commands are specific for redistribution of IPv6 prefixes. You can use the **match clns** and **set clns** commands only if you use a route-map to redistribute CLNS routes to or from the routing protocol.

Table 1 and Table 2 use these conventions:

- Supported commands are marked with a **Yes**.
- Unsupported commands are marked with a dash (—).
- Unsupported commands known to perform an action (probably an unwanted action) are marked with a **No**.

Table 1 – Route-map Commands to Match Routes Installed into Routing Tables by Protocols

Command	Redistribution Support						
	connected	static	RIP	EIGRP	OSPF	IS-IS	BGP
match clns address	—	Yes	—	—	—	Yes	Yes
match clns next-hop	—	Yes	—	—	—	Yes	—
match interface	Yes	Yes	Yes	Yes	Yes	Yes	—
match ip address	Yes	Yes	Yes	Yes	Yes	Yes	Yes
match ip address prefix-list	Yes	Yes	Yes	Yes	Yes	Yes	Yes
match ip next-hop	—	Yes	Yes	Yes	Yes	Yes	Yes

match ip next-hop prefix-list	—	No	No	No	No	No	No
match ip route-source	—	—	Yes	Yes	Yes	—	Yes
match ip route-source prefix-list	—	—	No	No	No	—	No
match ipv6 address [prefix-list]	Yes	Yes	Yes	—	Yes	Yes	Yes
match ipv6 next-hop [prefix-list]	—	Yes	Yes	—	—	—	Yes
match ipv6 route-source [prefix-list]	—	—	Yes	—	—	—	Yes
match metric	—	—	Yes	Yes	Yes	Yes	Yes
match policy-list	Yes	Yes	Yes	Yes	Yes	Yes	Yes
match route-type external	—	—	—	Yes	Yes	Yes	Yes
match route-type internal	—	—	—	Yes	Yes	—	Yes
match route-type local	—	—	—	—	—	—	Yes
match route-type nssa-external	—	—	—	—	Yes	—	—
match route-type {level-1 level-2}	—	—	—	—	—	Yes	—
match tag	—	Yes	Yes	Yes	Yes	Yes	Yes

Table 2 – Commands that Modify Route Attributes During Redistribution into the Target Protocol

Command	Redistribution Support				
	RIP	EIGRP	OSPF	IS-IS	BGP
set as-path tag	—	—	—	—	Yes
set community	—	—	—	—	Yes

set ip next-hop	—	—	—	—	Yes
set ip next-hop peer-address	—	—	—	—	No
set ipv6 next-hop	—	—	—	—	Yes
set level {backbone stub-area}	—	—	No	—	—
set level {level-1 level-2 level-1-2}	—	—	—	Yes	—
set local-preference	—	—	—	—	Yes
set metric	Yes	—	Yes	Yes	Yes
set metric +/-	No	—	No	No	No
set metric eigrp-metric	—	Yes	—	—	—
set metric +/- eigrp-metric	—	No	—	—	—
set metric-type internal	—	—	—	Yes	—
set metric-type external	—	—	—	Yes	—
set metric-type {type-1 type-2}	—	—	Yes	—	—
set nlri	—	—	—	—	Yes
set origin	—	—	—	—	Yes
set tag	Yes	Yes	Yes	—	—
set weight	—	—	—	—	Yes

Conclusion

Route-maps are very powerful but complicated tools for the redistribution of routes. They allow for very

fine manipulation of routing information when it is redistributed between protocols. However, they can be dangerous and can create blackholes or suboptimal traffic flow in your network. You must design networks very carefully, if you plan to employ complex redistribution features between multiple routing protocols.

Related Information

- [Redistribution of Routing Protocols](#)
- [Protocol Support Page](#)
- [Cisco Technical Support & Downloads](#)