



Configuring OSPF Local RIB

- [Prerequisite for OSPF Local RIB, on page 1](#)
- [Restriction for OSPF Local RIB, on page 1](#)
- [Information About OSPF Local RIB, on page 1](#)
- [How to Configure OSPF Local RIB, on page 2](#)
- [Configuration Examples for OSPF Local RIB, on page 4](#)
- [Feature History for OSPF Local RIB, on page 8](#)

Prerequisite for OSPF Local RIB

Before configuring this feature, ensure the OSPF routing protocol is configured.

Restriction for OSPF Local RIB

This feature is available only for IP Version 4 networks.

Information About OSPF Local RIB

A router that is running OSPF maintains a local RIB in which it stores all routes to destinations that it has learned from its neighbors. At the end of each SPF, OSPF attempts to install the best (that is, the least-cost) routes to a destination present in the local RIB into the global IPv4 routing table. The global RIB will be updated only when routes are added, deleted, or changed. Routes in the local RIB and Forwarding Information Base (FIB) will not compute when intermediate results are computed during SPF, resulting in fewer dropped packets in some circumstances.

By default, the contents of the global RIB are used to compute inter-area summaries, NSSA translation, and forwarding addresses for type-5 and type-7 LSAs. Each of these functions can be configured to use the contents of the OSPF local RIB instead of the global RIB for their computation. Using the local RIB for the computation may be slightly faster in some circumstances, but because the local RIB has information for only a particular instance of OSPF, using it for the computation may yield incorrect results. Potential problems that may occur include routing loops and null routes. It is recommended that you not change the default values because they are conservative and preserve the current global RIB behavior.

By default, OSPF installs discard routes to null0 for any area range (internal) or summary-address (external) prefixes that it advertises to other routers. Installation of a discard route can prevent routing loops in cases

where portions of a summary do not have a more specific route in the RIB. Normally, internal discard routes are installed with an administrative distance of 110, while external discard routes have an administrative distance of 254.

There may be rare circumstances, however, when some other values are needed. For example, if one OSPF process installs a route that exactly matches an area range configured on another OSPF process, the internal discard routes for the second OSPF process could be given a higher (less desirable) administrative distance.

OSPF Local RIB Path Limit

The OSPF Local RIB Path Limit feature allows network administrators to control the number of paths OSPF installs in its Local RIB for a specific prefix.

How to Configure OSPF Local RIB

Although it is recommended to keep the default settings for the commands described in the following sections, it is optional to change the defaults settings.

Changing the Default Local RIB Criteria

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router ospf <i>process-id</i> [vrf <i>vpn-name</i>] Example: Device(config)# router ospf 23	Configures an OSPFv2 routing process and enters router configuration mode.
Step 4	local-rib-criteria [forwarding-address] [inter-area-summary] [nssa-translation] Example: Device(config-router)# local-rib-criteria forwarding-address	Specifies that the OSPF local RIB will be used for route validation.
Step 5	end Example:	Returns to privileged EXEC mode.

	Command or Action	Purpose
	<code>Device(config-router)# end</code>	
Step 6	show ip ospf <i>process-id</i> rib [redistribution] [<i>network-prefix</i>] <i>[network-mask]</i> [detail] Example: <code>Device# show ip ospf 23 rib</code>	Displays information for the OSPF local RIB or locally redistributed routes.

Changing the Administrative Distance for Discard Routes

It is recommended that you keep the default settings. However, you can follow the steps in this section to change the administrative distance for discard routes.

Procedure

	Command or Action	Purpose
Step 1	enable Example: <code>Device> enable</code>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <code>Device# configure terminal</code>	Enters global configuration mode.
Step 3	router ospf <i>process-id</i> [vrf <i>vpn-name</i>] Example: <code>Device(config)# router ospf 23</code>	Configures an OSPFv2 routing process and enters router configuration mode.
Step 4	discard-route [external [<i>distance</i>]] [internal [<i>distance</i>]] Example: <code>Device(config-router)# discard-route external 150</code>	Specifies the administrative distance to be used for internal and external discard routes. <p>Note You can now specify the administrative distance for internal and external discard routes.</p>
Step 5	end Example: <code>Device(config-router)# end</code>	Returns to privileged EXEC mode.
Step 6	show ip route [<i>ip-address</i> [<i>mask</i>]] [longer-prefixes] <i>protocol</i> [<i>process-id</i>] list	Displays the current state of the routing table.

	Command or Action	Purpose
	<p>[<i>access-list-number</i> <i>access-list-name</i>] static download]</p> <p>Example:</p> <pre>Device# show ip route ospf 23</pre>	<p>Note Entering the show ip route command will verify the changed administrative distance values for external and internal discard routes.</p>

Example

The sample output displayed for the **show ip route** command confirms that the administrative distance for the IP route 192.168.0.0/24 is 110.

```
Device# show ip route 192.168.0.0 255.255.255.0
```

```
Routing entry for 192.168.0.0/24
```

```
    Known via "ospf 1", distance 110, metric 0, type intra area
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Null0
```

```
    Route metric is 0, traffic share count is 1
```

Configuration Examples for OSPF Local RIB

Example: Configuring OSPF Local RIB

In the following example, when the maximum-paths parameter is set to four in the default configuration mode, RIB has successfully installed four paths.

```
Device# show ip route ospf
45.1.1.0 [110/30] via 13.1.20.2, 00:03:03, Ethernet0/2.20
                [110/30] via 13.1.19.2, 00:03:03, Ethernet0/2.19
                [110/30] via 13.1.18.2, 00:03:03, Ethernet0/2.18
                [110/30] via 13.1.17.2, 00:03:03, Ethernet0/2.17
```

The following is a sample output of the **show ip ospf rib** command. In the given example, the default value is set to four for maximum-paths. Within the OSPF Local RIB, it is visible that the first four paths have been successfully installed into RIB and the last two paths have been rejected because of exceeding maximum-path limit.

```
Device# show ip ospf rib 45.1.1.2
45.1.1.0/24, Intra, cost 30, area 0
    SPF Instance 26, age 00:00:21
```

```
Flags: RIB
  via 13.1.20.2, Ethernet0/2.20, label 1048578, strict label 1048578
    Flags: RIB
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 13.1.13.2, Ethernet0/2.13, label 1048578,
strict label 1048578, cost 30
    Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, Downstr,
LoadShare
    LSA: 1/4.4.4.4/4.4.4.4
  via 13.1.19.2, Ethernet0/2.19, label 1048578, strict label 1048578
    Flags: RIB
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 12.1.18.2, Ethernet0/0.18, label 1048578,
strict label 1048578, cost 30
    Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, NodeProt,
Downstr, LoadShare
    LSA: 1/4.4.4.4/4.4.4.4
  via 13.1.18.2, Ethernet0/2.18, label 1048578, strict label 1048578
    Flags: RIB
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 12.1.11.2, Ethernet0/0.11, label 1048578,
strict label 1048578, cost 30
    Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, NodeProt,
Downstr, LoadShare
    LSA: 1/4.4.4.4/4.4.4.4
  13.1.17.2, Ethernet0/2.17, label 1048578, strict label 1048578
    Flags: RIB
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 13.1.12.2, Ethernet0/2.12, label 1048578,
strict label 1048578, cost 30
    Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, Downstr,
LoadShare
    LSA: 1/4.4.4.4/4.4.4.4
  via 13.1.12.2, Ethernet0/2.12, label 1048578, strict label 1048578
    Flags: MaxPath
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 13.1.15.2, Ethernet0/2.15, label 1048578,
strict label 1048578, cost 30
    Flags: Repair, PostConvrg, IntfDj, BcastDj, PrimPath, Downstr,
LoadShare
    LSA: 1/4.4.4.4/4.4.4.4
  via 13.1.11.2, Ethernet0/2.11, label 1048578, strict label 1048578
    Flags: MaxPath
    LSA: 1/4.4.4.4/4.4.4.4
    Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 13.1.16.2, Ethernet0/2.16, label 1048578,
```

```

strict label 1048578, cost 30
  Flags: Repair, PostConvrg, IntfDj, BcastDj, PrimPath, Downstr,
LoadShare
  LSA: 1/4.4.4.4/4.4.4.4

```

The following is a sample output of the **show ip ospf rib** command. With OSPF LRIB path limit feature, when default value for maximum-paths is set to four, only the four paths installed in the RIB are seen in the OSPF Local RIB.

```

Device# show ip ospf rib 45.1.1.2
45.1.1.0/24, Intra, cost 30, area 0
  SPF Instance 27, age 00:01:01
  Flags: RIB
    via 13.1.16.2, Ethernet0/2.16, label 1048578, strict label 1048578
      Flags: RIB
      LSA: 1/4.4.4.4/4.4.4.4
      Source: 4.4.4.4 (area 0)
      PostConvrg repair path via 13.1.20.2, Ethernet0/2.20, label 1048578,
strict label 1048578, cost 30
      Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, Downstr,
LoadShare
      LSA: 1/4.4.4.4/4.4.4.4
    via 13.1.15.2, Ethernet0/2.15, label 1048578, strict label 1048578
      Flags: RIB
      LSA: 1/4.4.4.4/4.4.4.4
      Source: 4.4.4.4 (area 0)
      PostConvrg repair path via 12.1.16.2, Ethernet0/0.16, label 1048578,
strict label 1048578, cost 30
      Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, NodeProt,
Downstr, LoadShare
      LSA: 1/4.4.4.4/4.4.4.4
    via 13.1.14.2, Ethernet0/2.14, label 1048578, strict label 1048578
      Flags: RIB
      LSA: 1/4.4.4.4/4.4.4.4
      Source: 4.4.4.4 (area 0)
      PostConvrg repair path via 12.1.15.2, Ethernet0/0.15, label 1048578,
strict label 1048578, cost 30
      Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, NodeProt,
Downstr, LoadShare
      LSA: 1/4.4.4.4/4.4.4.4
via 13.1.13.2, Ethernet0/2.13, label 1048578, strict label 1048578
  Flags: RIB
  LSA: 1/4.4.4.4/4.4.4.4
  Source: 4.4.4.4 (area 0)
  PostConvrg repair path via 12.1.20.2, Ethernet0/0.20, label 1048578,
strict label 1048578, cost 30
  Flags: RIB, Repair, PostConvrg, IntfDj, BcastDj, PrimPath, NodeProt,
Downstr, LoadShare
  LSA: 1/4.4.4.4/4.4.4.4

```

Example: Configuring maximum-paths per-prefix-distribution

The following is a sample output of the **maximum-paths per-prefix distribution** command. The **maximum-paths per-prefix distribution** command is disabled by default. In this case, all prefixes utilize the same paths.

```
Device# show ip route
144.1.11.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.12.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.13.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.14.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.15.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.16.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
144.1.17.1 [110/21] via 13.1.16.2, 00:02:23, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:02:23, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:02:23, Ethernet0/2.14
                  [110/21] via 13.1.13.2, 00:02:23, Ethernet0/2.13
```

In the following example, when **maximum-paths per-prefix distribution** command is enabled, various prefixes utilize different sets of paths.

```
Device# show ip route
144.1.11.1 [110/21] via 13.1.16.2, 00:00:04, Ethernet0/2.16
                  [110/21] via 13.1.15.2, 00:00:04, Ethernet0/2.15
                  [110/21] via 13.1.14.2, 00:00:04, Ethernet0/2.14
                  [110/21] via 12.1.11.2, 00:00:04, Ethernet0/0.11
144.1.12.1 [110/21] via 12.1.14.2, 00:00:04, Ethernet0/0.14
                  [110/21] via 12.1.13.2, 00:00:04, Ethernet0/0.13
                  [110/21] via 12.1.12.2, 00:00:04, Ethernet0/0.12
                  [110/21] via 12.1.11.2, 00:00:04, Ethernet0/0.11
144.1.13.1 [110/21] via 12.1.17.2, 00:00:04, Ethernet0/0.17
                  [110/21] via 12.1.16.2, 00:00:04, Ethernet0/0.16
                  [110/21] via 12.1.15.2, 00:00:04, Ethernet0/0.15
                  [110/21] via 12.1.14.2, 00:00:04, Ethernet0/0.14
144.1.14.1 [110/21] via 12.1.20.2, 00:00:04, Ethernet0/0.20
                  [110/21] via 12.1.19.2, 00:00:04, Ethernet0/0.19
```

```

[110/21] via 12.1.18.2, 00:00:04, Ethernet0/0.18
[110/21] via 12.1.17.2, 00:00:04, Ethernet0/0.17
144.1.15.1 [110/21] via 13.1.17.2, 00:00:04, Ethernet0/2.17
[110/21] via 13.1.12.2, 00:00:04, Ethernet0/2.12
[110/21] via 13.1.11.2, 00:00:04, Ethernet0/2.11
[110/21] via 12.1.20.2, 00:00:04, Ethernet0/0.20
144.1.16.1 [110/21] via 13.1.16.2, 00:00:04, Ethernet0/2.16
[110/21] via 13.1.15.2, 00:00:04, Ethernet0/2.15
[110/21] via 13.1.14.2, 00:00:04, Ethernet0/2.14
[110/21] via 13.1.13.2, 00:00:04, Ethernet0/2.13
144.1.17.1 [110/21] via 13.1.16.2, 00:00:04, Ethernet0/2.16
[110/21] via 12.1.13.2, 00:00:04, Ethernet0/0.13
[110/21] via 12.1.12.2, 00:00:04, Ethernet0/0.12
[110/21] via 12.1.11.2, 00:00:04, Ethernet0/0.11

```

Feature History for OSPF Local RIB

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Table 1: Feature History for the OSPF Local RIB Feature

Release	Feature	Feature Information
Cisco IOS XE Fuji 16.9.2	OSPF Local RIB	With the OSPF Local RIB feature, each OSPF protocol instance has its own local Routing Information Base (RIB). The OSPF local RIB serves as the primary state for OSPF SPF route computation. The global RIB is not updated with intermediate results during the SPF. Instead, the global RIB is updated only when routes are added, deleted, or changed, thereby reducing global RIB computation. This reduced update activity may result in fewer dropped packets.
Cisco IOS XE 17.14.1	OSPF Local RIB Path Limit Enhancement	With the introduction of the OSPF Local RIB Path Limit, network administrators can now control the number of paths OSPF installs in the Local RIB for a specific prefix.

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.