



Routing Information Base Support

The Routing Information Base (RIB) enhancement supports route redistribution and on-demand nexthop requirements.

- [Routing Information Base Support for Route Redistribution, on page 1](#)
- [OSPF Node SID Redistribution Support , on page 1](#)
- [Routing Information Base Support for On-Demand Next Hop, on page 4](#)
- [Additional References for Routing Information Base Support, on page 4](#)
- [Feature Information for Routing Information Base Support, on page 4](#)

Routing Information Base Support for Route Redistribution

Effective with Cisco IOS XE Everest 16.5.1, a requirement to redistribute labels associated with prefixes is introduced. To support redistribution requirements, the storage of local label per prefix is supported in RIB.

The local label is stored instead of the SID to ease use with different protocols which may use different SRGBs. The SID assigned by the destination protocol may not be the same as the SID associated with the source protocol.

The prefix reachability advertisement or an SRMS advertisement is the source of the SID. In SRMS advertisement, the destination protocols for redistribution do not advertise the SID in their prefix reachability advertisements, as this alters conflict resolution by indicating on other network nodes that the source of the advertisement was not from SRMS.

OSPF Node SID Redistribution Support

Effective Cisco IOS XE 16.7.1, when OSPF receives the redistributed prefixes from other IGPs and vice versa the prefix segment identifiers (SIDs) are also advertised which was not the case earlier. You needed to have the BGP LS (or) segment routing mapping server (SRMS) support to learn the SIDs across the IGP domains.

When the user enable redistribution under OSPF the prefix SID entries associated with the prefix entries are provided to OSPF. This gets advertised by OSPF to all its neighbor. The way OSPF advertises varies depending upon the role of OSPF in the network.

Information About OSPF Node SID Redistribution Support

NSSA ASBR

When you enable **redistribute ISIS instance ip** under OSPF which is Not-So-Stubby Area autonomous system boundary router (NSSA ASBR), it gets all the prefixes from IP routing information base (RIB) which are learnt by IS-IS along with the SID entries. OSPF generates Extended Prefix LSA (EPL) with the scope as area and the route type as RTYPE_NSSA1 or RTYPE_NSSA2 for the prefixes and advertises to all its neighbors. Similarly, when the redistribution is un-configured (or) when the prefixes become unavailable OSPF withdraws the EPL. When the redistributed route is a non-connected route then the OSPF sets the No-PHP flag but explicit NULL flag is not set. However, when the redistributed route is a connected route then OSPF sets the explicit NULL and No-PHP flag according to the configuration done in the SR policy.

When NSSA ABR receives the EPL, the ABR translates the LSA into opaque AS EPL and floods it to all its neighbors.

When a NSSA router which is neither ABR nor ASBR receives the EPL, it learns the prefix along with the SID entries and floods it to all its neighbors in the same area.


non-NSSA ASBR

When the user enabled **redistribute ISIS instance ip** under OSPF which is regular ASBR router, it gets all the prefixes from IP RIB which are learnt by IS-IS along with the SID entries. OSPF generates EPL with the scope as autonomous system (AS) and the route type as RTYPE_EXTERN1 or RTYPE_EXTERN2 for the prefixes and advertises to all its neighbors. Similarly when the redistribution is unconfigured (or) when the prefixes become unavailable, OSPF withdraws the EPL again with AS-Scope. When the redistributed route is a non-connected route then the OSPF sets the No-PHP flag but explicit NULL flag is not set. However, when the redistributed route is a connected route then OSPF sets the explicit NULL and No-PHP flag according to the configuration done in the SR policy. When a router receives the EPL with AS scope, it learns the prefix along with the SID entry and floods it to all its neighbors in all areas.

Redistributing Prefix

When IS-IS is enabled for redistribution of OSPF routes the prefixes are given along with the SID information so that the prefixes reach to other domain with the SID values. Refer to the below topology to understand the OSPF prefixes redistribution to the other domains:

Figure 1: OSPF Prefix Redistribution



R1 and R2 are enabled for OSPF. R2 and R3 are enabled for IS-IS. Both IS-IS and OSPF are enabled for Segment Routing. In R2, both IS-IS and OSPF are configured. Prefixes configured are:

1. 1.1.1/32 in R1 (enabled for OSPF with SID 1)
2. 2.2.2/32 in R2 (enabled for OSPF with SID 2)
3. 3.3.3/32 in R3 (enabled for ISIS SID 3)

When you enable SID redistribution in R2, then the prefix 3.3.3/32 is redistributed to R1. So, R1 knows the SID to reach the prefix R3.

```
conf t router isis 10 net 49.0001.0000.0000.0001.00 metric-style wide distribute link-state
segment-routing mpls router ospf 10 router-id 2.2.2.2 segment-routing mpls distribute
link-state
```

To enable redistribution of ISIS into OSPF routes:

```
conf t router ospf 10 redistribute isis 10 ip
```

Verifying OSPF Node SID Redistribution

Use the **show ip ospf rib redistribution detail** command to verify if OSPF is redistributing the prefixes from IS-IS.

```
Device# show ip ospf rib redistribution detail
OSPF Router with ID (2.2.2.2) (Process ID 10)

Base Topology (MTID 0)

OSPF Redistribution
3.3.3.3/32, type 2, metric 20, tag 0, from IS-IS Router
Attributes 0x1000000, event 1, PDB Index 4, PDB Mask 0x0
Source route metric 20, tag 0
SID 1003, SID Flags NP-bit, EPX Flags None
via 7.9.0.9, Ethernet0/0
```

Use the **show ip ospf segment-routing local-prefix** command to verify if the SID entries are advertised to its neighbor.

```
Device# show ip ospf segment-routing local-prefix

OSPF Router with ID (2.2.2.2) (Process ID 10)
Area 0:
Prefix:          Sid:  Index:          Type:          Source:
2.2.2.2/32      2    0.0.0.0      Intra          Loopback0
AS external:
Prefix:          Sid:  Index:          Type:          Source:
3.3.3.3/32      3    0.0.0.1      External       Redist
```

Use the **show ip ospf segment-routing sid-database** command to verify if the SIDs are received.

```
Device# show ip ospf segment-routing sid-database

OSPF Router with ID (1.1.1.1) (Process ID 10)
OSPF Segment Routing SIDs

Codes: L - local, N - label not programmed,
M - mapping-server

SID          Prefix          Adv-Rtr-Id      Area-Id  Type
-----
1            1.1.1.1/32     1.1.1.1         0        Intra
2            2.2.2.2/32     2.2.2.2         0        Intra
3            3.3.3.3/32     2.2.2.2         -        External
```

Use the **show ip route 3.3.3.3** command to verify if the IP routing entry is configured for the redistributed route.

```
Device# show ip route 3.3.3.3
Routing entry for 3.3.3.3/32
Known via "ospf 10", distance 110, metric 20, type extern 2, forward metric 20
```

```

Last update from 1.2.0.2 on Ethernet0/1, 00:00:01 ago
SR Incoming Label: 16003
Routing Descriptor Blocks:
* 1.3.1.3, from 2.2.2.2, 00:00:01 ago, via Ethernet1/1, merge-labels
  Route metric is 20, traffic share count is 1
  MPLS label: 16003
  MPLS Flags: NSF

```

Routing Information Base Support for On-Demand Next Hop

For On-Demand Next Hop (ODN) requirements, RIB supports a next hop called binding label which is provided by the supporting routing protocol (BGP). The binding label is used by the FIB to dynamically resolve the next hop.

The route producer installs a local binding label which identifies the ODN tunnel path associated with the next hop. The labeled traffic is sent via the tunnel and the label is distinct from the existing outlabel.

The following is the sample output of **show ip route** command where each next hop is updated to show the binding label.

```

Device# show ip route 10.10.10.2

Routing entry for 10.10.10.2/32
  Known via "isis", distance 115, metric 10, type level-1
  Redistributing via isis
  Last update from 200.200.200.2 on Ethernet0/0, 00:00:14 ago
  Incoming Label: 16100
  Routing Descriptor Blocks:
  * 200.200.200.2, from 10.10.10.2, 00:00:14 ago, via Ethernet0/0
    Route metric is 10, traffic share count is 1
    * Binding Label 4020, from 2.2.2.2, 00:00:14 ago,
      Route metric is 10, traffic share count is 1

```



Note The incoming labels are seen only after the SID redistribution is enabled.

Additional References for Routing Information Base Support

Related Documents

Related Topic	Document Title
Cisco IOS Commands	Cisco IOS Master Command List, All Releases

Feature Information for Routing Information Base Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Routing Information Base Support

Feature Name	Releases	Feature Information
Routing Information Base Support	Cisco IOS XE Everest 16.5.1b	The Routing Information Base (RIB) enhancement supports route redistribution and On-Demand Nextthop requirements. No new commands were added or modified.
OSPF Node SID Redistribution Support	Cisco IOS XE Fuji 16.7.1	Effective the Cisco IOS XE Fuji 16.7.1 release, when OSPF receives the redistributed prefixes from other IGPs and vice versa the prefix segment identifiers (SIDs) are also advertised which was not the case earlier. You need to have the BGP LS (or) segment routing mapping server (SRMS) support to learn the SIDs across the IGP domains. The following commands were added or modified for this feature: show ip ospf rib redistribution detail, show ip ospf segment-routing local-prefix, show ip ospf segment-routing sid-database, show ip route 3.3.3.3.

