

Configure Segment Routing for OSPF Protocol

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) developed by the OSPF working group of the Internet Engineering Task Force (IETF). Designed expressly for IP networks, OSPF supports IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets.

This module provides the configuration information to enable segment routing for OSPF.

Note For additional information on implementing OSPF on your , see the *Implementing OSPF* module in the .

- Enabling Segment Routing for OSPF Protocol, on page 1
- Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface, on page 3
- Configuring an Adjacency SID, on page 5
- Segment Routing ECMP-FEC Optimization, on page 8
- Multi-area Loopback Interface Support for OSPF, on page 8
- Display Neighbor Router Capabilities in OSPF Networks, on page 12

Enabling Segment Routing for OSPF Protocol

Segment routing on the OSPF control plane supports the following:

- OSPFv2 control plane
- Multi-area
- IPv4 prefix SIDs for host prefixes on loopback interfaces
- Adjacency SIDs for adjacencies
- MPLS penultimate hop popping (PHP) and explicit-null signaling

This section describes how to enable segment routing MPLS and MPLS forwarding in OSPF. Segment routing can be configured at the instance, area, or interface level.

Before you begin

Your network must support the MPLS Cisco IOS XR software feature before you enable segment routing for OSPF on your router.

Note You must enter the commands in the following task list on every OSPF router in the traffic-engineered portion of your network.

Procedure

	Command or Action	Purpose
Step 1	**configure** **Example:** `RP/0/RP0/CPU0:router# configure`	Enters mode.
Step 2	**router ospf** *process-name* **Example:** `RP/0/RP0/CPU0:router(config)# router ospf 1`	Enables OSPF routing for the specified routing process and places the router in router configuration mode.
Step 3	**segment-routing mpls** **Example:** `RP/0/RP0/CPU0:router(config-ospf)# segment-routing mpls`	Enables segment routing using the MPLS data plane on the routing process and all areas and interfaces in the routing process. Enables segment routing fowarding on all interfaces in the routing process and installs the SIDs received by OSPF in the forwarding table.
Step 4	**segment-routing sr-prefer** **Example:** `RP/0/RP0/CPU0:router(config-ospf)# segment-routing sr-prefer`	Sets the preference of segment routing (SR) labels over label distribution protocol (LDP) labels.
Step 5	**area** *area* **Example:** `RP/0/RP0/CPU0:router(config-ospf)# area 0`	Enters area configuration mode.
Step 6	**segment-routing mpls** **Example:** `RP/0/RP0/CPU0:router(config-ospf-ar)# segment-routing mpls`	(Optional) Enables segment routing using the MPLS data plane on the area and all interfaces in the area. Enables segment routing fowarding on all interfaces in the area and installs the SIDs received by OSPF in the forwarding table.
Step 7	**exit** **Example:**	

	Command or Action	Purpose
	`RP/0/RP0/CPU0:router(config-ospf-ar)# exit` `RP/0/RP0/CPU0:router(config-ospf)# exit`	
Step 8	Use the **commit** or **end** command.	**commit** —Saves the configuration changes and remains within the configuration session. **end** —Prompts user to take one of these actions: • **Yes** — Saves configuration changes and exits the configuration session. • **No** —Exits the configuration session without committing the configuration changes. • **Cancel** —Remains in the configuration session, without committing the configuration changes.

What to do next

Configure the prefix SID.

Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface

A prefix segment identifier (SID) is associated with an IP prefix. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels. A prefix SID is configured under the loopback interface with the loopback address of the node as the prefix. The prefix segment steers the traffic along the shortest path to its destination.

A prefix SID can be a node SID or an Anycast SID. A node SID is a type of prefix SID that identifies a specific node. An Anycast SID is a type of prefix SID that identifies a set of nodes, and is configured with n-flag clear. The set of nodes (Anycast group) is configured to advertise a shared prefix address and prefix SID. Anycast routing enables the steering of traffic toward multiple advertising nodes. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.

The prefix SID is globally unique within the segment routing domain.

This task describes how to configure prefix segment identifier (SID) index or absolute value on the OSPF-enabled Loopback interface.

Before you begin

Ensure that segment routing is enabled on an instance, area, or interface.

Procedure

	Command or Action	Purpose
Step 1	**configure** **Example:** `RP/0/RP0/CPU0:router# configure`	Enters mode.
Step 2	**router ospf** *process-name* **Example:** `RP/0/RP0/CPU0:router(config)# router ospf 1`	Enables OSPF routing for the specified routing process, and places the router in router configuration mode.
Step 3	**area** *value* **Example:** `RP/0/RP0/CPU0:router(config-ospf)# area 0`	Enters area configuration mode.
Step 4	**interface Loopback** *interface-instance* **Example:** `RP/0/RP0/CPU0:router(config-ospf-ar)# interface loopback 0`	Specifies the loopback interface and instance.

Step 5	**prefix-sid**{**index** *SID-index* | **absolute** *SID-value* } [**n-flag-clear**] [**explicit-null**]
	Example:
	`RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid index 1001`
	`RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid absolute 17001`
		Configures the prefix-SID index or absolute value for the interface.
		Specify **index** *SID-index* for each node to create a prefix SID based on the lower boundary of the SRGB + the index.
		Specify **absolute** *SID-value* for each node to create a specific prefix SID within the SRGB.
		By default, the n-flag is set on the prefix-SID, indicating that it is a node SID. For specific prefix-SID (for example, Anycast prefix-SID), enter the `n-flag-clear` keyword. OSPF does not set the `N` flag in the prefix-SID sub Type Length Value (TLV).
		To disable penultimate-hop-popping (PHP) and add an explicit-Null label, enter the `explicit-null` keyword. OSPF sets the `E` flag in the prefix-SID sub TLV.

Step 6	Use the **commit** or **end** command.	**commit** —Saves the configuration changes and remains within the configuration session. **end** —Prompts user to take one of these actions:

	Command or Action	Purpose
		• **Yes** — Saves configuration changes and exits the configuration session. • **No** —Exits the configuration session without committing the configuration changes. • **Cancel** —Remains in the configuration session, without committing the configuration changes.

Verify the prefix-SID configuration:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 7.0.0.1 self-originate
 OSPF Router with ID (10.0.0.1) (Process ID 1)
                Type-10 Opaque Link Area Link States (Area 0)
<...>
    Extended Prefix TLV: Length: 20
      Route-type: 1
      AF        : 0
      Flags     : 0x40
      Prefix    : 10.0.0.1/32

      SID sub-TLV: Length: 8
        Flags     : 0x0
        MTID      : 0
        Algo      : 0
        SID Index : 1001
```

Configuring an Adjacency SID

An adjacency SID (Adj-SID) is associated with an adjacency to a neighboring node. The adjacency SID steers the traffic to a specific adjacency. Adjacency SIDs have local significance and are only valid on the node that allocates them.

An adjacency SID can be allocated dynamically from the dynamic label range or configured manually from the segment routing local block (SRLB) range of labels.

Adjacency SIDs that are dynamically allocated do not require any special configuration, however there are some limitations:

- A dynamically allocated Adj-SID value is not known until it has been allocated, and a controller will not know the Adj-SID value until the information is flooded by the IGP.
- Dynamically allocated Adj-SIDs are not persistent and can be reallocated after a reload or a process restart.
- Each link is allocated a unique Adj-SID, so the same Adj-SID cannot be shared by multiple links.

Manually allocated Adj-SIDs are persistent over reloads and restarts. They can be provisioned for multiple adjacencies to the same neighbor or to different neighbors. You can specify that the Adj-SID is protected. If

the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.

Adjacency SIDs are advertised using the existing OSPF Adj-SID sub-TLV. The P-flag is defined for manually allocated Adj-SIDs.

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|B|V|L|G|P|     |
+-+-+-+-+-+-+-+-+
```

Table 1: Adjacency Segment Identifier (Adj-SID) Flags Sub-TLV Fields

Field	Description
P (Persistent)	This flag is set if the Adj-SID is persistent (manually allocated).

This task explains how to configure an Adj-SID on an interface.

Before you begin

Ensure that segment routing is enabled on the corresponding address family.

Use the **show mpls label table detail** command to verify the SRLB range.

Procedure

	Command or Action	Purpose
Step 1	**configure** **Example:** `RP/0/RP0/CPU0:router# configure`	Enters mode.
Step 2	**router ospf** *process-name* **Example:** `RP/0/RP0/CPU0:router(config)# router ospf 1`	Enables OSPF routing for the specified routing instance, and places the router in router configuration mode.
Step 3	**area** *area* **Example:** `RP/0/RP0/CPU0:router(config-ospf)# area 0`	Enters area configuration mode.
Step 4	**interface** *type interface-path-id* **Example:** `RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE0/0/0/1`	Specifies the interface and enters interface configuration mode.

	Command or Action	Purpose

Step 5	**adjacency-sid** {**index** *adj-SID-index* | **absolute** *adj-SID-value*} [**protected**]
	Example:
	`RP/0/RP0/CPU0:router(config-config-ospf-ar-if)# adjacency-sid index 10`
	`RP/0/RP0/CPU0:router(config-config-ospf-ar-if)# adjacency-sid absolute 15010`
		Configures the Adj-SID index or absolute value for the interface.
		Specify **index** *adj-SID-index* for each link to create an Ajd-SID based on the lower boundary of the SRLB + the index.
		Specify **absolute** *adj-SID-value* for each link to create a specific Ajd-SID within the SRLB.
		Specify if the Adj-SID is **protected**. For each primary path, if the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.

Step 6	Use the **commit** or **end** command.	**commit** —Saves the configuration changes and remains within the configuration session. **end** —Prompts user to take one of these actions: • **Yes** — Saves configuration changes and exits the configuration session. • **No** —Exits the configuration session without committing the configuration changes. • **Cancel** —Remains in the configuration session, without committing the configuration changes.

What to do next

Configure the SR-TE policy.

Protected Adjacency SID Backup Timer

OSPF advertises a protected adjacency SID for an adjacency when a backup path is available. Primary and backup paths are programmed into the label switching database (LSD) as rewrites.

When an adjacency goes down, OSPF stops advertising the protected adjacency SID immediately, and the backup path is promoted and installed as LSD rewrite. After a specified amount of time, the LSD rewrite is deleted. If the installed path fails again, the protection ends there and traffic through the original protected adjacency SID is permanently lost.

The Protected Adjacency SID Backup Timer provides a configurable maintenance time period. During this time period, OSPF updates the LSD rewrite with primary and backup (if available) paths to the neighbor upon topology changes.

Configuration

Use the **segment-routing protected-adjacency-sid-delay** command in OSPF configuration mode. The range is from 30 to 3600 seconds; the default is 900 seconds (15 min).

```
Router(config)# router ospf 1
Router(config-ospf)# segment-routing protected-adjacency-sid-delay 360
```

Running Configuration

```
router ospf 1
 segment-routing protected-adjacency-sid-delay 360
 area 1
  interface HundredGigE0/0/0/1
   fast-reroute per-prefix
   fast-reroute per-prefix ti-lfa enable
  !
 !
!
```

Segment Routing ECMP-FEC Optimization

ECMP-FECs are used for any ECMP programming on the system, such as MPLS LSP ECMP, VPN multipath, and EVPN multi-homingEVPN multi-homing.

The SR ECMP-FEC optimization solution minimizes ECMP-FEC resource consumption during underlay programming for an SR-MPLS network. This feature supports sharing the same ECMP-FEC, regular FEC, and Egress Encapsulation DB (EEDB) entries for all IPv4 and IPv6 Segment Routing prefixes with the same set of next hops. ECMP-FEC optimization is triggered when all the out_labels associated with the ECMP paths for a given prefix have the same value. If this rule is not met, then the prefix is programmed with a dedicated ECMP-FEC. Other prefixes that meet the rule are candidates for optimization.

Segment Routing Label Edge Router (LER) ECMP-FEC Optimization enables ECMP-FEC optimization originally developed for Label Switched Router (LSR) nodes (MPLS P) to be enabled on LER (Layer 3 MPLS PE) routers.

For usage guidelines, limitations, and configuration options, see Segment Routing ECMP-FEC Optimization.

Multi-area Loopback Interface Support for OSPF

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
Multi-area Loopback Interface for OSPF	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.

Feature Name	Release Information	Feature Description
Multi-area Loopback Interface for OSPF	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS You can save IP addresses and resources, prevent the use of multiple node SIDs for labels associated with loopback interfaces, and save time configuring multiple loopback interfaces for an Area Border Router (ABR) in a network. These improvements are possible as you can now configure a single loopback interface for multiple areas. With this feature, an ABR can use a single loopback interface for all areas it connects to, eliminating the need for separate loopback interfaces for each area. Previously, each loopback interface was linked to only one area. The feature introduces these changes: **CLI:** The **multi-area-interface** command is extended to support loopback interfaces.

Loopback Interface Optimization for ABRs

With this feature, you can reuse a single loopback interface for multiple areas, with no specific limit on the number of areas that can be configured for an ABR. This means that an ABR requires only one loopback interface for all areas it connects to instead of requiring a separate loopback interface for each area.

In earlier releases, configuring separate loopback interfaces for each area was necessary, resulting in configuring as many loopback interfaces as areas on the ABR.

Key Benefits of Multi-area Loopback Interface

The key benefits of the feature are:

- The feature offers an efficient networking solution, eliminating the need to configure multiple loopback interfaces.

- The feature improves scaling by enabling a 1:n ratio of loopback interfaces to multiple areas.

Usage Guidelines and Limitations for Multi-area Loopback

The following usage guidelines and limitations apply:

- The feature supports only IPv4 addresses.
- Associate the loopback interface to an area of an ABR, and then reuse it as a multi-area interface in other areas.
- The feature is limited to a single OSPF instance and does not extend to multiple OSPF instances.
- The feature is limited to a single VRF and does not extend to multiple VRFs.

Configurations

Configure the loopback interface

Perform the following task to configure the loopback interface IPv4 or IPv6 address, which is the primary loopback address that you want to reuse for all areas:

```
Router(config)#interface Loopback 0
Router(config-if)#ipv4 address 10.10.10.1/8
Router(config-if)#commit
```

Running Configuration

```
!
interface Loopback0
 ipv4 address 10.10.10.1 255.0.0.0
!
```

Associate the loopback interface to an area

Perform the following task to associate the primary loopback interface to area 1:

```
Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface Loopback 0
Router(config-ospf-ar-if)#commit
```

Running Configuration

```
!
router ospf 100
 router-id 10.10.10.1
 area 1
  interface Loopback 0
  !
 !
```

Reuse the loopback interface

Perform the following task to reuse the primary loopback interface in another OSPF area:

```
Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 0
Router(config-ospf-ar)#multi-area-interface Loopback 0
Router(config-ospf-ar-mif)#commit
```

Running Configuration

```
!
router ospf 100
 router-id 10.10.10.1
 area 0
  multi-area-interface Loopback 0
  !
 !
```

Verification

The following example shows how the loopback 0 interface is identified by area 1 in an ABR:

```
Router#show ospf routes 10.10.10.1/8 backup-path detail
Thu Mar 21 18:54:55.729 UTC

OSPF Route entry for 10.10.10.1/8
  Route type:  Intra-area
  Last updated: Mar 11 02:43:17.876
  Metric: 1
  SPF priority: 8,  SPF version: 38
  RIB version: 0,  Source: Unknown
  Label Info: Default 16004 SSPF 18004 Type SR
       10.10.10.1, directly connected, via Loopback0
        LSA: 1/10.10.10.1/10.10.10.1, Area: 1
```

The following example shows that summary LSA is not advertised by ABR in area 0 after you enable the loopback interface in area 1:

```
Router#show ospf database summary 10.10.10.1
Thu Mar 21 18:30:57.294 UTC


            OSPF Router with ID (10.10.10.1) (Process ID 1)

                Summary Net Link States (Area 0)

  LS age: 1760
  Options: (No TOS-capability, DC)
  LS Type: Summary Links (Network)
  Link State ID: 10.10.10.1 (Summary Network Number)
  Advertising Router: 10.0.0.1
  LS Seq Number: 8000006e
  Checksum: 0x5550
  Length: 28
  Network Mask: /32
        TOS: 0  Metric: 2
```

The following example indicates the router link-state advertisement (LSA) in area 0:

```
Router#show ospf database opaque-area ext-prefix 10.10.10.1/8
Thu Mar 21 18:57:40.914 UTC


            OSPF Router with ID (10.10.10.1) (Process ID 1)

                Type-10 Opaque Link Area Link States (Area 0)

  LS age: 345
```

```
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 7.0.0.3
Opaque Type: 7
Opaque ID: 3
Advertising Router: 10.10.10.1
LS Seq Number: 80000001
Checksum: 0xa3d6
Length: 56

  Extended Prefix TLV: Length: 8
    Route-type: 1
    AF        : 0
    Flags     : 0x40
    Prefix    : 10.10.10.1/32

    SID sub-TLV: Length: 8
      Flags     : 0x0
      MTID      : 0
      Algo      : 0
      SID Index : 4

    SID sub-TLV: Length: 8
      Flags     : 0x0
      MTID      : 0
      Algo      : 1
      SID Index : 2004
```

Display Neighbor Router Capabilities in OSPF Networks

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS You can now improve the user experience, understand the features and functionalities supported by neighboring routers, reduce operational work hours and human errors by gaining better visibility into the routing capabilities of directly connected neighboring routers in an OSPF network. These improvements are now possible when you enable Segment Routing on a neighboring connected router. Previously, there was no straightforward method to display the segment routing and other capabilities of neighboring routers in an OSPF network. The feature introduces these changes: **CLI:** • The **capabilities** keyword is introduced in the **show ospf neighbor** command. **YANG Data Models:** • `Cisco-IOS-XR-ipv4-ospf-oper.yang` See (GitHub, Yang Data Models Navigator)

The feature enables you to access detailed information of a neighboring router capability from the Router Information Link State Advertisement (RI LSA). This includes details such as segment routing feature, graceful restart and helper operation modes, hostname, and traffic engineering capabilities.

For more information about graceful restart and helper operation modes, see *Implementing OSPF* chapter in the *Routing Configuration Guide*.

Key Benefits of Displaying Neighbor Router Capabilities

- The feature ensures optimal path selection for faster packet forwarding.
- Reduces operational complexity, and enables you to make informed decisions for network upgrades.
- Perform compatibility checks with neighboring routers.

Usage Guidelines and Limitations to Display Neighbor Router Capabilities

The following usage guidelines and limitations apply:

- The RI LSA includes the capabilities of a router, only when Segment Routing capability is enabled for that router. This means that if Segment Routing is not enabled, the router does not advertise any capability in the RI LSA.
- The feature is area-oriented, meaning that it is specific to an OSPF area. This suggests that the capabilities advertised by a neighboring router are relevant and visible only within the same OSPF area.
- You must enable traffic engineering at the OSPF or at the area level.
- You must configure graceful restart and helper capabilities at the OSPF level.

Verify the Neighbor Router Capabilities

You can use the **show ospf neighbor capabilities** command to display the traffic engineering support, hostname, and graceful restart and helper capabilities of neighboring connected routers, which are enabled for Segment Routing.

```
Router#show ospf neighbor capabilities
Mon Apr 29 08:27:07.289 UTC

Capabilities advertised in RI LSA
  TE - Traffic Engineering
  GR - Graceful Restart
  GH - Graceful Restart Helper

  NA - Capabilities TLV Not Advertised

Neighbors for OSPF 1

Area 1
Neighbor ID            Address              Interface       SR-Enabled    Capabilities
192.168.199.137        209.165.200.225/27       Gi100/0/0/0     Yes           TE, GR, GH
    HostName: Device_3
192.168.48.1           209.165.200.230/27        Gi100/0/0/4     No            NA
192.168.48.200         209.165.200.235/27       Gi10/0/0/0      Yes           TE, GR, GH
    HostName: Device_0

Total neighbor count: 3
```

In this scenario, Area 1 consists of three connected routers. Routers 192.168.199.137 and 192.168.48.200 support Segment Routing, so their RI LSA includes the TE, GH, and GR capabilities. When the TE and GR are not configured, only the GH capability is advertised by default. Router 192.168.48.1, which does not support Segment Routing, does not include any capabilities in its RI LSA.

When you use the **show ospf neighbor** *neighbor-id* command to view the summary information of a specific neighboring router, the summary information includes the capabilities of that neighbor router if Segment Routing is enabled for it.

The following example shows the capabilities of neighbor 192.168.199.130 listed in the summary information:

```
Router#show ospf neighbor 192.168.199.130
Mon Nov 6 19:13:08.861 IST
* Indicates MADJ interface
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPF 1

 Neighbor 192.168.199.130, interface address 172.31.80.37
    In the area 0 via interface GigabitEthernet0/2/0/4
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:32
    Neighbor is up for 00:08:50
    Number of DBD retrans during last exchange 0
    Index 3/6, retransmission queue length 0, number of retransmission o
    First 0(0) /0(0) Next 0(0) /0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is l msec, maximum is 8 msec
    LS Ack list: NSR-sync pending 0, high water mark o
    Adjacency SIDs:
        Label: 24010,    Dynamic, Unprotected
        Label: 24011,    Dynamic,   Protected (Has backup)
    Neighbor Interface ID: 4
    Segment-Routing Enabled
      Strict SPF capable
      Start SID: 16000 Range Size: 8000
      Flex algo Enabled
    Capabilities advertised in RI LSA
      Graceful Restart capable
      Graceful Restart Helper capable
HostName: r5

Total neighbor count: 1
```