



Configure Segment Routing for OSPF Protocol

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
NC57 line cards in Native Mode: SR-MPLS - OSPFv2	Release 7.3.1	This feature is now supported on routers that have the Cisco NC57 line cards installed and operate in native mode. To enable the native mode, use the hw-module profile npu native-mode-enable command in the configuration mode. Ensure that you reload the router after configuring the native mode.

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 Cisco NCS 5500 Series Router, see the *Implementing OSPF* module in the *Routing Configuration Guide for Cisco NCS 5500 Series Routers*.

- [Enabling Segment Routing for OSPF Protocol, on page 2](#)
- [Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface, on page 4](#)
- [Configuring an Adjacency SID, on page 6](#)
- [Conditional Prefix Advertisement, on page 9](#)
- [Segment Routing ECMP-FEC Optimization, on page 11](#)
- [Display Neighbor Router Capabilities in OSPF Networks, on page 11](#)

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.

SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **segment-routing mpls**
4. **segment-routing sr-prefer**
5. **area** *area*
6. **segment-routing mpls**
7. **exit**
8. Use the **commit** or **end** command.

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# <code>configure</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	<p>router ospf <i>process-name</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config)# router ospf 1</pre>	Enables OSPF routing for the specified routing process and places the router in router configuration mode.
Step 3	<p>segment-routing mpls</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf)# segment-routing mpls</pre>	<p>Enables segment routing using the MPLS data plane on the routing process and all areas and interfaces in the routing process.</p> <p>Enables segment routing forwarding on all interfaces in the routing process and installs the SIDs received by OSPF in the forwarding table.</p>
Step 4	<p>segment-routing sr-prefer</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf)# segment-routing sr-prefer</pre>	Sets the preference of segment routing (SR) labels over label distribution protocol (LDP) labels.
Step 5	<p>area <i>area</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf)# area 0</pre>	Enters area configuration mode.
Step 6	<p>segment-routing mpls</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# segment-routing mpls</pre>	(Optional) Enables segment routing using the MPLS data plane on the area and all interfaces in the area. Enables segment routing forwarding on all interfaces in the area and installs the SIDs received by OSPF in the forwarding table.
Step 7	<p>exit</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# exit RP/0/RP0/CPU0:router(config-ospf)# exit</pre>	
Step 8	Use the commit or end command.	<p>commit —Saves the configuration changes and remains within the configuration session.</p> <p>end —Prompts user to take one of these actions:</p> <ul style="list-style-type: none"> • 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.

SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **area** *value*
4. **interface Loopback** *interface-instance*
5. **prefix-sid** [**algorithm** *algorithm-number*] {**index** *SID-index* | **absolute** *SID-value* } [**n-flag-clear**] [**explicit-null**]
6. Use the **commit** or **end** command.

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	router ospf <i>process-name</i> 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.

	Command or Action	Purpose
Step 3	<p>area <i>value</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf)# area 0</pre>	Enters area configuration mode.
Step 4	<p>interface <i>Loopback interface-instance</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# interface Loopback0 passive</pre>	Specifies the loopback interface and instance.
Step 5	<p>prefix-sid [<i>algorithm algorithm-number</i>] {<i>index SID-index</i> <i>absolute SID-value</i> } [<i>n-flag-clear</i>] [<i>explicit-null</i>]</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid index 1001</pre> <pre>RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid absolute 17001</pre>	<p>Configures the prefix-SID index or absolute value for the interface.</p> <p>Specify algorithm <i>algorithm-number</i> to configure SR Flexible Algorithm.</p> <p>Specify index <i>SID-index</i> for each node to create a prefix SID based on the lower boundary of the SRGB + the index.</p> <p>Specify absolute <i>SID-value</i> for each node to create a specific prefix SID within the SRGB.</p> <p>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 <i>n-flag-clear</i> keyword. OSPF does not set the N flag in the prefix-SID sub Type Length Value (TLV).</p> <p>To disable penultimate-hop-popping (PHP) and add an explicit-Null label, enter the <i>explicit-null</i> keyword. OSPF sets the E flag in the prefix-SID sub TLV.</p>
Step 6	Use the commit or end command.	<p>commit —Saves the configuration changes and remains within the configuration session.</p> <p>end —Prompts user to take one of these actions:</p> <ul style="list-style-type: none"> • 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 2: 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.

SUMMARY STEPS

1. **configure**
2. **router ospf** *process-name*
3. **area** *area*
4. **interface** *type interface-path-id*
5. **adjacency-sid** {**index** *adj-SID-index* | **absolute** *adj-SID-value*} [**protected**]
6. Use the **commit** or **end** command.

DETAILED STEPS**Procedure**

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	router ospf <i>process-name</i> 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 <i>area</i> Example: RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode.
Step 4	interface <i>type interface-path-id</i> Example: RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE0/0/0/1	Specifies the interface and enters interface configuration mode.
Step 5	adjacency-sid { index <i>adj-SID-index</i> absolute <i>adj-SID-value</i> } [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)#	Configures the Adj-SID index or absolute value for the interface. Specify index <i>adj-SID-index</i> for each link to create an Adj-SID based on the lower boundary of the SRLB + the index. Specify absolute <i>adj-SID-value</i> for each link to create a specific Adj-SID within the SRLB.

	Command or Action	Purpose
	<code>adjacency-sid absolute 15010</code>	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.	<p>commit —Saves the configuration changes and remains within the configuration session.</p> <p>end —Prompts user to take one of these actions:</p> <ul style="list-style-type: none"> • 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
```

!

!

!

Conditional Prefix Advertisement

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Segment Routing Conditional Prefix Advertisement for OSPF	Release 7.3.1	<p>In a typical Anycast scenario, if an advertising node becomes unavailable or unreachable while still advertising its Anycast SID, traffic could still be routed to the node and, as a result, get dropped.</p> <p>This feature allows a node to advertise its loopback address when it's connected to the domain, and to track the loopback addresses of the other nodes in the domain. If a node becomes unavailable or unreachable, it stops advertising its loopback address, allowing for a new path to be computed.</p>

In some situations, it's beneficial to make the OSPF prefix advertisement conditional. For example, an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) that has lost its connection to one of the areas or autonomous systems (AS) might keep advertising a prefix. If an ABR or ASBR advertises the Segment Routing (SR) SID with this prefix, the label stack of the traffic routed toward the disconnected area or AS might use this SID, which would result in dropped traffic at the ABR or ASBR.

ABRs or ASBRs are often deployed in pairs for redundancy and advertise a shared Anycast prefix SID. Conditional Prefix Advertisement allows an ABR or an ASBR to advertise its Anycast SID only when connected to a specific area or domain. If an ABR or ASBR becomes disconnected from the particular area or AS, it stops advertising the address for a specified interface (for example, Loopback).

Configure the conditional prefix advertisement under a specific interface. The prefix advertisement on this interface is associated with the route-policy that tracks the presence of a set of prefixes (prefix-set) in the Routing Information Base (RIB).

For faster convergence, the route-policy used for conditional prefix advertisement uses the new event-based **rib-has-route async** condition to notify OSPF of the following situations:

- When the last prefix from the prefix-set is removed from the RIB.
- When the first prefix from the prefix-set is added to the RIB.

Configuration

To use the conditional prefix advertisement in OSPF, create a prefix-set to be tracked. Then create a route policy that uses the prefix-set.

```

Router(config)# prefix-set prefix-set-name
Router(config-pfx)# prefix-address-1/length[, prefix-address-2/length,,,
prefix-address-16/length]
Router(config-pfx)# end-set

Router(config)# route-policy rpl-name
Router(config-rpl)# if rib-has-route async prefix-set-name then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy

```

To advertise the loopback address in OSPF conditionally, use the **advertise prefix route-policy** command under OSPF interface address-family configuration sub-mode.

```

Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface Loopback0
Router(config-ospf-ar-if)# advertise prefix route-policy rpl-name
Router(config-ospf-ar-if)# commit

```

Example

```

Router(config)# prefix-set domain_2
Router(config-pfx)# 2.3.3.3/32, 2.4.4.4/32
Router(config-pfx)# end-set
Router(config)# route-policy track_domain_2
Router(config-rpl)# if rib-has-route async domain_2 then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface Loopback0
Router(config-ospf-ar-if)# advertise prefix route-policy track_domain-2
Router(config-ospf-ar-if)# commit

```

Running Configuration

```

prefix-set domain_2
  2.3.3.3/32,
  2.4.4.4/32
end-set
!
route-policy track_domain_2
  if rib-has-route async domain_2 then
    pass
  endif
end-policy
!
router ospf 1
  area 0
  interface Loopback0
    advertise prefix route-policy track_domain_2
  !
!
!

```

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-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](#).

Display Neighbor Router Capabilities in OSPF Networks

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks on NCS 5700 Fixed Port Routers	Release 24.2.11	Introduced in this release on: NCS 5700 fixed port routers The feature support is now extended to NCS 5700 fixed port routers.

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks	Release 24.2.1	<p>Introduced in this release on: NCS 5500 fixed port routers; NCS 5500 modular routers (NCS 5500 line cards; NCS 5700 line cards [Mode: Compatibility; Native]).</p> <p>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.</p> <p>Previously, there was no straightforward method to display the segment routing and other capabilities of neighboring routers in an OSPF network.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> The capabilities keyword is introduced in the show ospf neighbor command. <p>YANG Data Models:</p> <ul style="list-style-type: none"> Cisco-IOS-XR-ipv4-ospf-oper.yang <p>See (GitHub, Yang Data Models Navigator)</p>

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 1 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
```