C H A P T E R **4**

Using Segment Routing with OSPF

✎
Note Using segment routing with OSPF is similar to using it with IS-IS. Before reading this chapter, see Using Segment Routing with IS-IS, page 3-1.

Segment Routing (SR) is a new paradigm of source routing driven by the IGPs and centered on inserting a list of instructions called segments into each packet. SR forwarding can be instantiated by MPLS or IPv6.

OSPF as an IGP propagates the segment information and may either use it internally (for example, to compute Fast Rerouting repair paths) or hand over the segment information to other clients within the router (for example, TE).

Using OSFP with Segment Routing

To provide a base SR functionality, OSPF interacts with several components:

- Segment Routing Application—It handles IGP and instance-independent global commands as well as manages global resources used by the SR such as, block of static MPLS labels available for SR Node SID MPLS instantiation).
- MPLS Forwarding Infrastructure (MFI)—Supports MPLS forwarding table.
- RIBv4—Supports the IPv4 routing table

Restrictions for Using Segment Routing with OSPF

- Segment routing must be configured at the router level before enabling it at the OSPF instance.
- OSPF instance must successfully register itself with both, Segment Routing Application and MFI and retrieve resources that are necessary for operation, such as SRGB range, dynamically-assigned MPLS labels for Adjacency SIDs, and so on.
- If more than one SR router in the network advertises SID for the same prefix, then the SID values and flags must match. If conflicting attributes are found when calculating the SID for the prefix, then all SIDs for the prefix are ignored.
- Network type point-to-point and broadcast are supported.

Enabling Segment Routing for OSPF

You can enable segment routing for IGPs under the router configuration sub mode, through commands. However, IGP segment routing are enabled only after the global SR is configured.

SR functionality configurable globally, that is outside the context of any particular IGP instance, include:

- Global enabling SR on the router.
- Specifying the range of MPLS label values to be used to instantiate SR SIDs into MPLS dataplane (SRGB block)
- Associating SID index with local prefix

Configuring SR on an OSPF Instance

After SR is enabled globally on the router, it must be enabled in those OSPF instances that need to run SR.

segment-routing [**area** *N*] {**mpls** | **disable**}

This command is accepted only if SR is already enabled globally. If the **area** keyword is specified then SR is enabled in that area only, otherwise it is enabled in all areas attached to this OSPF instance. The keyword **disable** can be specified only when the **area** keyword is also present.

Enabling this command causes OSPF to originate RI LSA, Extended Prefix and Extended Link LSAs. It enables MPLS on all interfaces in area(s) enabled for SR and programs SR MPLS labels for forwarding.

Enabling Advertisement of Mapping Server Prefix Ranges

Global SR configuration may contain prefix-to-SID mapping entries for prefixes that are not local to the router. Each of these entries specify the range of prefixes. Remote mapping entries can be used to find SIDs for prefixes connected to routers that do not supporting SR and hence, are not capable of advertising SIDs themselves. This is part of SR-LDP inter-working functionality.

OSPF learns the ranges configured in the global SR configuration through the SR application and advertises them in the Extended Prefix Range TLVs.

The following command allows an OSPF instance to advertise mapping entries configured in the global SR mode:

[no] segment-routing prefix-sid-map advertise-local

This command is configurable in the router mode. Its default state is *disabled*, That is, no mapping ranges are advertised by OSPF even if they are configured in the global SR mode.

Disabling Mapping Server Functionality

By default, OSPF considers prefix ranges received from mapping servers in the network when computing SIDs for prefixes. This is the core of SR-LDP inter-working label computation. However, if the SR-LDP feature needs to be disabled, use the following command:

[no] segment-routing prefix-sid-map receive

Default state of this command is *enabled*. That is, OSPF processes mapping ranges received from mapping servers in the network.

When this command is configured in its non-default form, or the **no** form, it prohibits OSPF from considering prefix-SID mappings from mapping servers.

> **Note** This command does not affect processing of 'native' SIDs, that is, those SIDs that are advertised in the Extended Prefix TLVs by routers to whom the prefix is locally connected.

Interface Mode Prefix Attributes Command

Extended Prefix TLV of the Extended Prefix LSA carries flags for the prefix and one of them is N-flag (Node). The N-flag indicates that any traffic sent along to the prefix is destined to the router originating the LSA. This flag typically marks the host routes of the router's loop-back.

By default, OSPF advertises the routes of its loopback with the N-flag. If the IP address of a loopback is not unique (for example, it is part of an anycast address), then OSPF must be configured to *not* advertise the N-flag with the prefix.

Use the following command to cause OSPF to clear N-flag on all prefixes configured on interface in the interface-mode:

ip ospf prefix-attributes n-flag-clear

Using OSFP Fast Reroute with Segment Routing

IP Fast Reroute is a set of techniques that allow rerouting IP traffic around a failed link or a failed node in the network within a very short time (< 50ms). One of the techniques to do this is Loop Free Alternates (LFA). Effective with Cisco IOS XE Release 3.18S, OSPF supports per-prefix directly connected LFA and remote LFA (RLFA).

The per-prefix directly connected LFA provides loop-free alternate path for most triangular topologies, but does not provide good coverage for rectangular or circular topologies. However, the RLFA, which uses MPLS forwarding with LDP signaling for tunneling the rerouted traffic to an intermediate node, extends the IPFRR coverage in ring or rectangular topologies. For each link, RLFA defines the P-Space (the set of nodes reachable from the calculating node without crossing the protected link) and Q-Space (the set of nodes that can reach the neighbor on the protected link without crossing the protected link itself). The nodes that belong to both P- and Q-Spaces are called PQ nodes and can be used as the intermediate nodes for the protected traffic. However, for topologies where the P- and Q-Spaces are completely disjoint, there is still no coverage by RLFA.

Topology Independent Fast st Reroute (TI-FRR) is a technique that uses Segment Routing to provide link protection in any topology, assuming the metric on the links in the topology is symmetrical. Even TI-LFA does not guarantee a backup in cases where the bandwidth on a single link is asymmetrical. TI-LFA only considers loop-free paths that are on the post-convergence path that helps you to better plan the capacity of the network.

Segment Routing allows creating a full explicit path through the network, but using such a fully specified path is not scalable in larger topologies due to the number of segments along the path. Specifying the whole path is, however, not necessary, and only a subset of the path is needed to carry the traffic to an intermediate node (release node) which does not loop the traffic back to the protecting node. When the release node is one of the neighbors of S (source or sender node), then we have a directly connected TI-LFA.

Effective with Cisco IOS XE Release 3.18S, an SR Tunnel is constructed to the release node (if remote) and used as the TI- LFA. The tunnel is constructed by explicitly forwarding through a set of one or more repair nodes. The tunnel is created using a SID stack consisting of one SID (Node SID or Adjacency SID) for every repair node and the corresponding label stack is pushed to the protected traffic.

Per Instance Enablement

- TI-LFA can be enabled on a per interface basis. One interface could be configured for LFA, another interface for LFA and RLFA, yet another interface for LFA and TI-LFA and another interface for LFA, RLFA, and TI-LFA.

- TI --LFA backup path is calculated only if TI-LFA protection is enabled on the output interface of the primary path.

- TI- LFA protection is not available for virtual links, sham links, and TE tunnels.

- Backup paths do not use virtual links, sham links, and TE tunnels.

Limitations of TI-LFA on OSPF

- TI-LFA is supported only on OSPFv2.

- TI-LFA only computes TI-LFA candidates that have been found based on a post-convergence path excluding the primary link, That is, TI-LFA only provides link protection. Node protection or SRLG protection is not supported.

- TI-LFA is calculated only if D and repair nodes (P, Q, PQ, or other) in the post-convergence path are segment routing capable. S does not need to be SR capable.

- TI-LFA is calculated only if Repair Nodes (P, Q, PQ, or other) in the post-convergence path have node SID. A node SID can be directed connected or one advertised by Mapping Server (SRMS).

- TI-LFA is restricted to a maximum of two or three segment tunnels, which may be lesser than the maximum labels that can be pushed by the router.

- When S is not SR capable, TI-LFA calculates a 0 or 1 segment LFA along the post-convergence path. The 1-segment LFA uses the LDP label to reach the PQ node in the post-convergence path.

- TI-LFA does not compute backup for virtual link, sham link, or TE tunnel path.

- TI-LFA does not use virtual link, sham link or TE tunnel as backup path.

- TI-LFA does not provide adjacency protection. It only protects prefixes. SR TE may provide protection for adjacency segments. Packets arriving with active adjacency segment do not have TI-LFA protection.

- TI-LFA is calculated only for prefixes having SR prefix SID or node SID.

- TI-LFA does not support Multi Topology Routing (MTR). LFA is calculated only for MTID zero.

Enabling and Disabling TI-FRR

TI-LFA can be enabled on interface, area, and process scopes. By default, TI-LFA is disabled globally.

TI-LFA can be enabled using following commands.

fast-reroute per-prefix ti-lfa

However, when TI-LFA is enabled globally, you have the option to disable it per area level, using following command:.

[no] fast-reroute per-prefix ti-lfa [area *area* **disable]**

This command is accepted even if Segment-Routing is not enabled on the OSF instance or area. In this case, a TI-LFA backup path is created if the post-convergence path has a PQ node with an LDP label to reach it. That is, RLFA through a post-convergence path.

At the process level:

- **fast-reroute per-prefix ti-lfa** [**area** *area* **disable**]
- **fast-reroute per-prefix enable** [**area** *area*] **prefix-priority** {**high** | **low**}
- **fast-reroute per-prefix remote-lfa** [**area** *area*] **maximum-cost** *cost*
- **fast-reroute per-prefix remote-lfa** [**area** *area*] **tunnel mpls-ldp**
- **fast-reroute per-prefix remote-lfa** [**area** *area*] **tunnel mpls-ldp**
- **fast-reroute per-prefix tie-break** *attribute* [**required**] **index** *index*

At the interface level:

- **ip ospf fast-reroute per-prefix protection disabled**
- **ip ospf fast-reroute per-prefix candidate disable**
- **ip ospf fast-reroute per-prefix protection ti-lfa** [**disable**]

Verifying TI-LFA on OSPF

The following show command displays the details about TI-LFA configuration:

show ip ospf fast-reroute ti-lfa

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Microloop avoidance is enabled for protected prefixes, delay 5000 msec

Loop-free Fast Reroute protected prefixes:

Area Topology name Priority Remote LFA Enabled TI-LFA Enabled
0 Base High Yes Yes

Repair path selection policy tiebreaks (built-in default policy):
0 post-convergence
10 primary-path
20 interface-disjoint
30 lowest-metric
40 linecard-disjoint
50 broadcast-interface-disjoint
256 load-sharing

OSPF/RIB notifications:
Topology Base: Notification Enabled, Callback Registered

Last SPF calculation started 00:02:42 ago and was running for 8 ms.
```

Verifying the Configuration of SR on OSPF

- The following **show** command displays information about the new LSAs, such as router information, Extended Prefix and Extended Link LSAs:

show ip ospf database opaque-area type ext-limit

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Type-10 Opaque Area Link States (Area 0)

LS age: 1378
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 8.0.0.7
Opaque Type: 8 (Extended Link)
Opaque ID: 7
Advertising Router: 4.4.4.4
LS Seq Number: 80000056
Checksum: 0xFEF8
Length: 76

TLV Type: Extended Link
Length: 52
Link connected to : another Router (point-to-point)
(Link ID) Neighboring Router ID: 12.12.12.12
(Link Data) Interface IP address: 0.0.0.7

Sub-TLV Type: Adj SID
Length : 7
Flags : L-Bit, V-bit
MTID : 0
Weight : 0
Label : 27

Sub-TLV Type: Adj SID
Length : 7
Flags : L-Bit, V-bit, B-bit
MTID : 0
Weight : 0
Label : 34

Sub-TLV Type: Local / Remote Intf ID
Local Interface ID : 7
Remote Interface ID : 0
```

show ip ospf database opaque-area type ext-prefix

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Type-10 Opaque Area Link States (Area 0)

LS age: 1392
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 7.0.0.0
Opaque Type: 7 (Extended Prefix)
Opaque ID: 0
Advertising Router: 4.4.4.4
LS Seq Number: 80000055
Checksum: 0xD156
Length: 44
```

```
TLV Type: Extended Prefix
Length: 20
Prefix : 4.4.4.4/32
AF : 0
Route-type: Intra
Flags : N-bit

Sub-TLV Type: Prefix SID
Length: 8
Flags : None
MTID : 0
Algo : IGP metric based SPT
SID : 4
```

show ip ospf database opaque-area type router-information

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Type-10 Opaque Area Link States (Area 0)

LS age: 1402
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 4.0.0.0
Opaque Type: 4 (Router Information)
Opaque ID: 0
Advertising Router: 4.4.4.4
LS Seq Number: 80000055
Checksum: 0x7B86
Length: 52

TLV Type: Router Information
Length: 4
Capabilities:
Graceful Restart Helper
Stub Router Support
Traffic Engineering Support

TLV Type: Segment Routing Algorithm
Length: 1
Algorithm: IGP metric based SPT

TLV Type: Segment Routing Range
Length: 12
Range Size: 8000

Sub-TLV Type: SID/Label
Length: 3
Label: 16000
```

- The following show command displays the interface output, but only if segment routing auto-enabled MPLS forwarding on the interface and SIDs were allocated to that interface

show ip ospf interface

```
Loopback0 is up, line protocol is up
Internet Address 13.13.13.13/32, Interface ID 29, Area 0
Attached via Interface Enable
Process ID 100, Router ID 13.13.13.13, Network Type LOOPBACK, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
0 1 no no Base
Enabled by interface config, including secondary ip addresses
```

```
Loopback interface is treated as a stub Host
BDI2104 is up, line protocol is up
Interface is unnumbered, Interface ID 40, Area 0
Using address of Loopback0 (13.13.13.13)
Attached via Interface Enable
Process ID 100, Router ID 13.13.13.13, Network Type POINT_TO_POINT, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
0 1 no no Base
Enabled by interface config, including secondary ip addresses
Transmit Delay is 1 sec, State POINT_TO_POINT, BFD enabled
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
oob-resync timeout 40
Hello due in 00:00:05
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Can be protected by per-prefix Loop-Free FastReroute
Can be used for per-prefix Loop-Free FastReroute repair paths
Not Protected by per-prefix TI-LFA
Segment Routing enabled for MPLS forwarding
Index 1/3/3, flood queue length 0
Next 0x0(0)/0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 8.8.8.8
Suppress hello for 0 neighbor(s)
```

- The following show command displays the information about segment routing
    - Per-instance state of the feature—whether enabled or disabled, whether operational or not
    - Registration status of OSPF with the SR application and MFI
    - Retrieved SRGB range

show ip ospf [*AS*] segment-routing

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Global segment-routing state: Enabled

Segment Routing enabled:
Area Topology name Forwarding
0 Base MPLS

SR Attributes
Prefer SR Labels
Do not advertise Explicit Null

Local MPLS label block (SRGB):
Range: 16000 - 23999
State: Created

Registered with SR App, client handle: 10
Connected map notifications active (handle 0x18), bitmask 0x1
Active policy map notifications active (handle 0x19), bitmask 0xC
Registered with MPLS, client-id: 100

Bind Retry timer not running
Adj Label Bind Retry timer not running
Adj Protected Label Bind Retry timer not running
```

- The following show command displays information about the SID database for one or for all SIDs. The output displays all prefixes that were advertised with a given SID and highlights which prefixes are local to the router.

✎

Note In the absence of a mis-configuration, there is only one prefix reported for each SID value.

show ip ospf [*AS*] segment-routing sid-database [*SID*]

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

OSPF Segment Routing SIDs

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

SID Prefix/Mask Adv-Rtr-Id Area-Id
-------- ------------------ ------------------ -------------
8   8.8.8.8/32 8.8.8.8 0
10  10.10.10.10/32 10.10.10.10 0
13 (L) 13.13.13.13/32 13.13.13.13 0
```

- The following show command displays the database of local prefixes along with their SID values and flags.

✎

Note This show command provides information only about locally-configured prefixes. Information about prefixes received from other routers is available either through LSDB (**show ip ospf database opaque-area**) or LRIB (**show ip ospf rib local**) show commands.

show ip ospf [*AS*] segment-routing local-prefix

```
OSPF Router with ID (13.13.13.13) (Process ID 100)

Area 0:

Prefix: Sid: Index: Interface:
13.13.13.13/32 13 0.0.0.0 Loopback0
```

Troubleshooting Segment Routing on OSPF

Use the following command to debug issues

- **debug ip ospf [*AS*] segment-routing**
- **debug ip ospf fast-reroute spf**
- **debug ip ospf fast-reroute spf detail**
- **debug ip ospf fast-reroute rib**
- **debug ip ospf fast-reroute rib [*access-list*]**

Additional References

Standards

Standard	Title
—	There are no new standards for this feature.

MIBs

MIB	MIBs Link
• —	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
—	There are no RFCs for this feature.