



Segment Routing Traffic Engineering With OSPF

This chapter describes how Segment Routing traffic engineering can be implemented using OSPF.

- [Feature Information for Segment Routing Traffic Engineering With OSPF, on page 1](#)
- [Restrictions for Segment Routing Traffic Engineering With OSPF, on page 2](#)
- [Information About Segment Routing Traffic Engineering With OSPF, on page 2](#)
- [How to Configure Segment Routing Traffic Engineering With OSPF, on page 10](#)
- [Verifying Configuration of the SR-TE Tunnels, on page 18](#)

Feature Information for Segment Routing Traffic Engineering With OSPF

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 <https://cfng.cisco.com/>. An account on Cisco.com is not required.

Table 1: Feature Information for Segment Routing Traffic Engineering With OSPF

Feature Name	Releases	Feature Information
Segment Routing Traffic Engineering With OSPF	Cisco IOS XE Amsterdam 17.3.2	<p>A Traffic Engineered (TE) tunnel is a container of TE LSP(s) instantiated between the tunnel ingress and the tunnel destination. A TE tunnel may instantiate one or more SR-TE LSP(s) that are associated with the same tunnel.</p> <p>The following commands were added or modified:</p> <p>show mpls traffic-eng tunnels, tunnel mpls traffic-eng path-option 10 dynamic segment-routing, tunnel mpls traffic-eng path-option 10 segment-routing, tunnel mpls traffic-eng path-option 10 explicit name MIXED_PATH segment-routingtunnel mpls traffic-eng path-option 10 explicit name IP_PATH1 segment-routingtunnel mpls traffic-eng path-option 10 explicit name NODE_PATH segment-routing.</p>

Restrictions for Segment Routing Traffic Engineering With OSPF

- Segment Routing Traffic Engineering is supported only on OSPFv2.
- SR-TE is not supported on broadcast interfaces; it is supported only point-to-point interfaces.
- Only one instance of protocol should be enabled for TE at a given point of time.

Information About Segment Routing Traffic Engineering With OSPF

A Traffic Engineered (TE) tunnel is a container of TE LSP(s) instantiated between the tunnel ingress and the tunnel destination. A TE tunnel may instantiate one or more SR-TE LSP(s) that are associated with the same tunnel. The SR-TE LSP path may not necessarily follow the same IGP path to a destination node. In this case, the SR-TE path can be specified a set of prefix-SID(s) and/or adjacency-SID(s) of nodes and/or links to be traversed by the SR-TE LSP.

The head-end imposes the corresponding MPLS label stack on to outgoing packets to be carried over the tunnel. Each transit node along the SR-TE LSP path uses the incoming top label to select the next-hop, pop or swap the label, and forward the packet to the next node with the remainder of the label stack, until the packet reaches the ultimate destination. OSPF provides TE with the topology and SR related information. SR related information include SRGB/prefix/Adjacency SIDs of all nodes/links with SR enabled in the network.

Benefits of Using Segment Routing Traffic Engineering With OSPF

Segment routing traffic engineering offers a comprehensive support for all useful optimizations and constraints, for example:

- Latency
- Bandwidth
- Disjointness
- Resource avoidance

OSPFv2 provides the following functionalities for SR-TE:

- OSPFv2 provides SR information along with TE topology information to TE module.
- TE uses this information to construct SR TE path/tunnel comprising of one or more segments - with the combination of prefix and/or adjacency segments.
- For the prefixes TE is interested in, OSPF provides first hop resolution to setup the forwarding plane.
- SR TE tunnels are also advertised back into OSPF (like RSVP TE tunnels) for diverting traffic over the SR-TE tunnels.

OSPFv2 Segment Routing Traffic Engineering Functionalities

OSPFv2 perform the following functionalities for SR-TE:

- OSPFv2 provides SR information along with TE topology information to TE module.
- TE uses this information to construct SR TE path/tunnel comprising of one or more segments - with the combination of prefix and/or adjacency segments.
- For the prefixes TE is interested in, OSPF provides first hop resolution to setup the forwarding plane.
- SR TE tunnels are also advertised back into OSPF (like RSVP TE tunnels) for diverting traffic over the SR-TE tunnels.

Protected Adjacency SID

Segment routing creates protected adjacency SID for point to point to point interfaces and broadcast interfaces. It advertises them to the extended link-state advertisement (LSA) along with the unprotected adjacency SID. Protected adjacency SID can have a repair path, but it is not guaranteed to have a repair path.

Traffic Engineering Interfaces

In order to support SR-TE functionality, TE interfaces with various components, and with IGP (OSPF and ISIS) to distribute and receive information on TE topology. For SR-TE support, OSPF needs to additionally provide SR information to TE that it had received through various LSAs, for example,

- Router Information LSA
- Extended Prefix LSA
- Extended Link LSA

TE interfaces distribute information, such as bandwidth resources, constraints, capabilities, and other attributes, associated with the links that are configured for TE. The link information is distributed to other routers using opaque LSAs and is used by TE to create a local topology database. The topology database is a key element in allowing TE to compute a suitable constraint-based path for establishing an LSP. TE also interfaces with the IGP to notify when a TE headend interface can be considered for routing packets.

Unnumbered Support

IS-IS description of an unnumbered link does not contain remote interface ID information. The remote interface ID of an unnumbered link is required to include the unnumbered link as part of the SR-TE tunnel.

Segment Routing Traffic Engineering Support for Forwarding Adjacency

MPLS TE forwarding adjacency feature is supported by OSPF. In this, TE tunnel is considered as a link in the IGP network. TE tunnel interfaces are advertised in the IGP network like any other links. Routers can then use these links to compute the shortest path tree (SPT).



Note This feature is not supported with the SR-TE tunnels.

Segment Routing Traffic Engineering Support for Auto-route Announce

MPLS TE auto-route announce feature is supported by OSPF, that uses TE Tunnel as the first-hop, if the node is reachable via that tunnel. It allows the traffic to the nodes that are downstream to the tail-end of the TE tunnel flows through the tunnel. OSPF supports auto-route over the SR-TE tunnels similar to the MPLS TE tunnels setup using RSVP.

The TE tunnel that instantiates an SR-TE LSP can be Auto-route Announced (AA) into IGP (OSPF and ISIS) as an IGP shortcut. The IGP uses the TE tunnel as next hop and installs routes in RIB for all IP prefixes whose shortest path falls behind the TE tunnel destination. Auto-route announce for of TE tunnels is supported to carry IPV4 prefixes.

Auto-route Announce IP2MPLS

The auto-routeIP2MPLS feature is introduced for SR tunnels to avoid potential packet from looping indefinitely between the SR-TE tunnel headend/ingress and a node that is pointing/routing the packet back to the headend/ingress.

The solution consists in the headend programming in forwarding two sets of path(s) for the prefixes that are mapped over the SR-TE tunnel. The first is the pure IP route for the prefix(es) mapped on the and having the outgoing interface as the tunnel interface. This allows mapping IP traffic directly over the tunnel. The second is the MPLS path for the prefixes mapped on the tunnel. For this the prefix-SID label is programmed with the IGP shortest path outgoing interface(s), that is, non tunnel output interfaces.

SR-TE LSP Instantiation

A Traffic Engineered (TE) tunnel is a container of one or more instantiated TE LSPs. An SR-TE LSP is instantiated by configuring ‘segment-routing’ on the path-option of the TE tunnel. The traffic mapped to the tunnel is forwarded over the primary SR-TE instantiated LSP.

Multiple path-options can also be configured under the same tunnel. Each path-option is assigned a preference index or a path-option index that is used to determine the more favorable path-option for instantiating the primary LSP—the lower the path-option preference index, the more favorable the path-option. The other less favorable path-options under the same TE tunnel are considered secondary path-options and may be used once the currently used path-option is invalidated (for example, due to a failure on the path).



Note A forwarding state is maintained for the primary LSP only.

Tunnel Path Affinity Validation

The affinity of a tunnel path can be specified using the command **tunnel mpls traffic-eng affinity** under the tunnel interface.

The head-end validates that the specified SR path is compliant with the configured affinity. This necessitates that the paths of each segment of the SR path be validated against the specified constraint. The path is declared invalid against the configured affinity constraints if at least a single segment of the path does not satisfy the configured affinity.

SR-TE Traffic Load Balancing

SR-TE tunnels support the following load-balancing options:

Load Balancing on Port Channel TE Links

Port Channel interfaces carry the SR-TE LSP traffic. This traffic load balances over port channel member links as well as over bundle interfaces on the head or mid of an SR-TE LSP.

Load Balancing on Single Tunnel

While using the equal cost multi path protocol (ECMP), the path to a specific prefix-SID may point to multiple next-hops. And if the SR-TE LSP path traverses one or more prefix-SIDs that have ECMP, the SR-TE LSP traffic load-balances on the ECMP paths of each traversed prefix-SID from the head-end or any midpoint traversed node along the SR-TE LSP path.

Load Balancing on Multiple Tunnels

Multiple TE tunnels can be used as next-hop paths for routes to specific IP prefixes either by configuring static route on multiple tunnels, or auto-route announcing multiple parallel tunnels to the same destination. In such cases, the tunnels share the traffic load equally or load balance traffic on multiple parallel tunnels. It is also possible to allow Unequal Load Balance (UELB) with an explicit per tunnel configuration at the tunnel head-end. In this case, the tunnel load-share is passed from MPLS-TE to forwarding plane.

The tunnel load-share feature continues to work for TE tunnels that instantiate the SR-TE LSPs.

SR-TE Tunnel Reoptimization

TE tunnel reoptimization occurs when the head-end determines that there is a more optimal path available than the one currently used. For example, in case of a failure along the SR-TE LSP path, the head-end could detect and revert to a more optimal path by triggering reoptimization.

Tunnels that instantiate SR-TE LSP can re-optimize without affecting the traffic carried over the tunnel.

Re-optimization can occur because:

- The explicit path hops used by the primary SR-TE LSP explicit path are modified.
- The head-end determines the currently used path-option are invalid due to either a topology path disconnect, or a missing SID in the SID database that is specified in the explicit-path.
- A more favorable path-option (lower index) becomes available.

When the head-end detects a failure on a protected SR adjacency-SID that is traversed by an SR-TE LSP, it starts the invalidation timer. If the timer expires and the head-end is still using the failed path because it is unable to reroute on a different path, the tunnel state is brought 'down' to avoid a null route from being sent along with the traffic. Once the tunnel is down, services on the tunnel converge to take a different path.

The following is a sample output of a manual reoptimization example. In this example, the path-option is changed from **10** to **20**.

```
Router# mpls traffic-eng reoptimize tunnel 1 path-option 20
The targeted path-option is not in lock down mode. Continue? [no]: yes
Router# show mpls traffic-eng tunnels tunnel1
Name: R1_t1 (Tunnell) Destination: 10.6.6.6
Status:
  Admin: up      Oper: up      Path: valid      Signalling: connected
  path option 20, (SEGMENT-ROUTING) type explicit IP_PATH (Basis for Setup)
  path option 10, (SEGMENT-ROUTING) type dynamic
Config Parameters:
  Bandwidth: 0 kbps (Global) Priority: 6 6 Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
    Path-invalidation timeout: 45000 msec (default), Action: Tear
    AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: explicit path option 20 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours, 9 minutes
    Time since path change: 14 seconds
    Number of LSP IDs (Tun_Instances) used: 1819
    Current LSP: [ID: 1819]
    Uptime: 17 seconds
    Selection: reoptimization
    Prior LSP: [ID: 1818]
    ID: path option unknown
    Removal Trigger: reoptimization completed
  Tun_Instance: 1819
Segment-Routing Path Info (isis level-1)
  Segment0[Node]: 10.4.4.4, Label: 114
  Segment1[Node]: 10.5.5.5, Label: 115
  Segment2[Node]: 10.6.6.6, Label: 116
```

SR-TE with Lockdown Option

The **lockdown** option prevents SR-TE from re-optimizing to a better path. However, it does not prevent signaling the existence of a new path.

```
interface Tunnell
ip unnumbered Loopback1
tunnel mode mpls traffic-eng
tunnel destination 10.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 6 6
tunnel mpls traffic-eng path-option 10 segment-routing lockdown
tunnel mpls traffic-eng path-selection metric igp
tunnel mpls traffic-eng load-share 10 (Tunnell) Destination:
10.6.6.6
Status:
  Admin: up      Oper: up      Path: valid      Signalling: connected
  path option 10, (LOCKDOWN) type segment-routing (Basis for Setup)
Config Parameters:
  Bandwidth: 0 kbps (Global) Priority: 6 6 Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
```

```

Path Selection:
  Protection: any (default)
  Path-invalidation timeout: 45000 msec (default), Action: Tear
  AutoRoute: enabled LockDown: enabled Loadshare: 10 [200000000]
  auto-bw: disabled
  Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: segment-routing path option 10 is active
  BandwidthOverride: disabled LockDown: enabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours, 22 minutes
    Time since path change: 1 minutes, 26 seconds
    Number of LSP IDs (Tun_Instances) used: 1822
  Current LSP: [ID: 1822]
    Uptime: 1 minutes, 26 seconds
    Selection: reoptimization
  Prior LSP: [ID: 1821]
    ID: path option unknown
    Removal Trigger: configuration changed
  Tun_Instance: 1822
  Segment-Routing Path Info (isis level-1)
    Segment0[Node]: 10.6.6.6, Label: 116

```

SR-TE Tunnel Protection

Protection for SR TE tunnels can take any of the following alternatives:

IP-FRR Local Repair Protection

On an SR-TE LSP head-end or mid-point node, IP-FRR is used to compute and program the backup protection path for the prefix-SID or adjacency-SID label.

With IP-FRR, backup repair paths are pre-computed and pre-programmed by IGP's *before* a link or node failure. The failure of a link triggers its immediate withdrawal from the TE topology (link advertisement withdrawal). This allows the head-end to detect the failure of an SR-TE LSP traversing the failed adjacency-SID.

When a protected adjacency-SID fails, the failed adjacency-SID label and associated forwarding are kept functional for a specified period of time (5 to 15 minutes) to allow all SR TE tunnel head-ends to detect and react to the failure. Traffic using the adjacency-SID label continues to be FRR protected even if there are subsequent topology updates that change the backup repair path. In this case, the IGP's update the backup repair path while FRR is active to reroute traffic on the newly-computed backup path.

When the primary path of a protected prefix-SID fails, the PLR reroutes to the backup path. The head-end remains transparent to the failure and continues to use the SR-TE LSP as a valid path.

IP-FRR provides protection for adjacency and prefix-SIDs against link failures only.

Tunnel Path Protection

Path protection is the instantiation of one or more standby LSPs to protect against the failure of the primary LSP of a single TE tunnel.

Path protection protects against failures by pre-computing and pre-provisioning secondary paths that are failure diverse with the primary path-option under the same tunnel. This protection is achieved by computing a path that excludes prefix-SIDs and adjacency-SIDs traversed by the primary LSP or by computing a path that excludes SRLGs of the primary SR-TE LSP path.

In the event of a failure of the primary SR-TE LSP, at least one standby SR-TE LSP is used for the tunnel. Multiple secondary path-options can be configured to be used as standby SR-TE LSPs paths.

SR-TE LSP Path Verification

SR-TE tunnel functionality requires that the head-end perform initial verification of the tunnel path as well as the subsequent tracking of the reachability of the tunnel tail-end and traversed segments.

Path verification for SR-TE LSP paths is triggered whenever MPLS-TE is notified of any topology changes or SR SID updates.

The SR-TE LSP validation steps consist of the following checks:

Topology Path Validation

The head-end validates the path of an SR-TE LSP for connectivity against the TE topology. MPLS-TE head-end checks if links corresponding to the adjacency SIDs are connected in the TE topology.

For newly-instantiated SR-TE LSPs, if the head-end detects a discontinuity on any link of the SR-TE path, that path is considered invalid and is not used. If the tunnel has other path-options with valid paths, those paths are used to instantiate the tunnel LSP.

For TE tunnels with existing instantiated SR-TE LSP, if the head-end detects a discontinuity on any link, the head-end assumes a fault has occurred on that link. In this case, the local repair protection, such as the IP FRR, come in to effect. The IGP's continue to sustain the protected adjacency label and associated forwarding after the adjacency is lost for some time. This allows the head-ends enough time to reroute the tunnels onto different paths that are not affected by the same failure. The head-end starts a tunnel invalidation timer once it detects the link failure to attempt to reroute the tunnel onto other available path-options with valid paths.

If the TE tunnel is configured with other path-options that are not affected by the failure and are validated, the head-end uses one of those path-options to reroute (and re-optimize) the tunnel by instantiating a new primary LSP for the tunnel using the unaffected path.

If no other valid path-options exist under the same tunnel, or if the TE tunnel is configured with only one path-option that is affected by the failure, the head-end starts an invalidation timer after which it brings the tunnel state to 'down'. This action avoids a null route from being sent along with traffic flowing over the affected SR-TE LSP, and allows services riding over the tunnel to reroute over different available paths at the head-end. There is an invalidation drop configuration that keeps the tunnel 'up', but drops the traffic when the invalidation timer expires.

For intra-area SR-TE LSPs, the head-end has full visibility over the LSP path, and validates the path to the ultimate LSP destination. However, for inter-area LSPs, the head-end has partial visibility over the LSP path—only up to the first ABR. In this case, the head-end can only validate the path from the ingress to the first ABR. Any failure along the LSP beyond the first ABR node is invisible to the head-end, and other mechanisms to detect such failures, such as BFD over LSP are assumed.

SR SID Validation

SID hops of an SR-TE LSP are used to determine the outgoing MPLS label stack to be imposed on the outgoing packets carried over the SR-TE LSP of a TE tunnel. A database of global and local adjacency-SIDs is populated from the information received from IGP's and maintained in MPLS-TE. Using a SID that is not available in the MPLS TE database invalidates the path-option using the explicit-path. The path-option, in this case, is not used to instantiate the SR TE LSP. Also, withdrawing, adding, or modifying a SID in the MPLS-TE SID-database, results in the MPLS-TE head-end verifying all tunnels with SR path-options (in-use or secondary) and invokes proper handling.

LSP Egress Interface

When the SR-TE LSP uses an adjacency-SID for the first path hop, TE monitors the interface state and IGP adjacency state associated with the adjacency-SID and the node that the SR-TE LSP egresses on. If the interface or adjacency goes down, TE can assume a fault occurred on the SR-TE LSP path and take the same reactive actions described in the previous sections.



Note When the SR-TE LSP uses a prefix-SID for the first hop, TE cannot directly infer on which interface the tunnel egresses. TE relies on the IP reachability information of the prefix to determine if connectivity to the first hop is maintained.

IP Reachability Validation

MPLS-TE validates that the nodes corresponding to the prefix-SIDs are IP reachable before declaring the SR path valid. MPLS-TE detects path changes for the IP prefixes corresponding to the adjacency or prefix SIDs of the SR-TE LSP path. If the node announcing a specific SID loses IP reachability, due to a link or node failure, MPLS-TE is notified of the path change (no path). MPLS-TE reacts by invalidating the current SR-TE LSP path, and may use other path-options with a valid path, if any to instantiate a new SR-TE LSP.



Note Since IP-FRR does not offer protection against failure of a node that is being traversed by an SR-TE LSP (such as, a prefix-SID failure along the SR-TE LSP path), the head-end immediately reacts to IP route reachability loss for prefix-SID node by setting the tunnel state to 'down' and removes the tunnel forwarding entry if there are no other path-options with valid path for the affected tunnel.

Tunnel Path Resource Avoidance Validation

You can specify a set of addresses to be validated as excluded from being traversed by SR-TE tunnel packets. To achieve this, the head-end runs the per-segment verification checks and validates that the specified node, prefix or link addresses are indeed excluded from the tunnel in the SR path. The tunnel resource avoidance checks can be enabled per path using the commands below. The list of addresses to be excluded are defined and the name of the list is referenced in the path-option.

```
interface tunnel100
 tunnel mpls traffic-eng path-option 1 explicit name EXCLUDE segment-routing
 ip explicit-path name EXCLUDE enable
  exclude-address 192.168.0.2
  exclude-address 192.168.0.4
  exclude-address 192.168.0.3
!
```

SR-TE LSP Explicit Null

MPLS-TE tunnel head-end does not impose explicit-null at the bottom of the stack. When penultimate hop popping (PHP) is enabled for SR prefix SIDs or when an adjacency SID is the last hop of the SR-TE LSP, the packet may arrive at the tail-end without a transport label. However, in some cases, it is desirable that the packet arrive at the tail-end with explicit-null label, and in such case, the head-end will impose an explicit-null label at the top of the label stack.

Verbatim Path Support

MPLS TE LSPs usually require that all the nodes in the network are TE aware which means that they have IGP extensions to TE in place. However, some network administrators want the ability to build TE LSPs to traverse nodes that do not support IGP extensions to TE, but that do support RSVP extensions to TE. Verbatim LSPs are helpful when all or some of the intermediate nodes in a network do not support IGP extensions for TE.

When this feature is enabled, the IP explicit path is not checked against the TE topology database. Since the TE topology database is not verified, a Path message with IP explicit path information is routed using the shortest path first (SPF) algorithm for IP routing.

How to Configure Segment Routing Traffic Engineering With OSPF

Perform the following steps to configure Segment Routing Traffic Engineering With OSPF.

Enabling Segment Routing Traffic Engineering With OSPF

OSPF Segment Routing traffic engineering is enabled when the segment-routing is enabled along with mpls traffic engineering. SR-TE support is turned on in an area when you enable SR & MPLS TE in that area.

```
router ospf 10
  router-id 10.10.10.2
  segment-routing mpls
  mpls traffic-eng area 0
```

Configuring the Path Option for a TE Tunnel

When the path-option type for an operational SR tunnel is changed from SR to non-SR (for example, **dynamic**), the existing forwarding entry of the tunnel is deleted.

Segment Routing can be enabled or disabled on an existing secondary or an in-use path-option. If the tunnel uses a signaled RSVP-TE explicit path-option and segment routing is enabled on that tunnel, the RSVP-TE LSP is torn, and the SR-TE LSP is instantiated using the same path-option. Conversely, if segment routing is disabled on a path-option that is in use by the primary LSP, the tunnel goes down intermittently and a new RSVP-TE LSP will be signaled using the same explicit path.

If the segment-routing path-option is enabled on a secondary path-option (that is, not in-use by the tunnel's primary LSP), the tunnel is checked to evaluate if the newly specified SR-TE LSP path-option is valid and more favorable to use for the tunnel primary LSP.

```
Device(config)# interface tunnel 100
Device(config-if)# tunnel mpls traffic-eng path-option 1 explicit name foo segment-routing
Device(config-if)# tunnel mpls traffic-eng path-option 2 dynamic segment-routing
Device(config-if)# tunnel mpls traffic-eng path-option 3 segment-routing
```

Configuring SR Explicit Path Hops

The following explicit path hops are supported in SR-TE:

- IP addresses
- MPLS labels
- Mix of IP addresses and MPLS labels

For intra-area LSPs, the explicit path can be specified as a list of IP addresses:

```
Device(config)# ip explicit-path name foo
Device(config-ip-expl-path)# index 10 next-address 10.1.1.1 node address
Device(config-ip-expl-path)# index 20 next-address 10.12.12.2 link address
```



Note When using IP unnumbered interfaces, you cannot specify next hop address as an explicit path index. It should be node address or label.

The explicit path can also be specified as segment-routing SIDs:

```
Device(config)# ip explicit-path name foo
Device(config-ip-expl-path)# index 10 next-label 20
```

Configuring Tunnel Path Affinity Validation

The affinity of a tunnel path can be specified using the command **tunnel mpls traffic-eng affinity** under the tunnel interface.

The head-end validates that the specified SR path is compliant with the configured affinity. This necessitates that the paths of each segment of the SR path be validated against the specified constraint. The path is declared invalid against the configured affinity constraints if at least a single segment of the path does not satisfy the configured affinity.

```
interface Tunnel1
no ip address
tunnel mode mpls traffic-eng
tunnel destination 10.5.5.5
tunnel mpls traffic-eng priority 5 5
tunnel mpls traffic-eng bandwidth 100
tunnel mpls traffic-eng affinity 0x1 mask 0xFFFF
tunnel mpls traffic-eng path-option 10 dynamic segment-routing
Router# show tunnel ??
Name: R1_t1                               (Tunnel1) Destination: 10.5.5.5
Status:
  Admin: up           Oper: up           Path: valid           Signalling: connected
  path option 10, (SEGMENT-ROUTING) type dynamic (Basis for Setup, path weight 20)
Config Parameters:
  Bandwidth: 100      kbps (Global) Priority: 5 5  Affinity: 0x1/0xFFFF
  Metric Type: TE (default)
  Path Selection:
  Protection: any (default)
  Path-selection Tiebreaker:
  Global: not set Tunnel Specific: not set Effective: min-fill (default)
  Hop Limit: disabled
  Cost Limit: disabled
  Path-invalidation timeout: 10000 msec (default), Action: Tear
  AutoRoute: disabled LockDown: disabled Loadshare: 100 [0] bw-based
  auto-bw: disabled
```

```

Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: dynamic path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
Node Hop Count: 2
History:
  Tunnel:
    Time since created: 10 minutes, 54 seconds
    Time since path change: 34 seconds
    Number of LSP IDs (Tun_Instances) used: 55
  Current LSP: [ID: 55]
    Uptime: 34 seconds
  Prior LSP: [ID: 49]
    ID: path option unknown
    Removal Trigger: tunnel shutdown
  Tun_Instance: 55
Segment-Routing Path Info (isis level-1)
  Segment0[Link]: 192.168.2.1 - 192.168.2.2, Label: 46
  Segment1[Link]: 192.168.4.2 - 192.168.4.1, Label: 49

```

Configuring Affinity on an Interface

Perform the following steps to configure affinity on an interface:

```

interface GigabitEthernet2
ip address 192.168.2.1 255.255.255.0
ip router isis 1
negotiation auto
mpls traffic-eng tunnels
mpls traffic-eng attribute-flags 0x1
isis network point-to-point
ip rsvp bandwidth

```


Configuring Segment Routing Traffic Engineering With OSPF

Consider the following inter area and intra area use cases for configuring SR-TE with OSPF:

Configuring Intra Area Tunnel

Consider the following topology to configure intra area tunnel:

Figure 1: Intra Area Tunnel



All the routers are configured in the same area, Area 0.

Configuration at the head end router R1:

```
router ospf 10
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
segment-routing mpls
mpls traffic-eng router-id Loopback1
mpls traffic-eng area 0

interface GigabitEthernet2 //interface connecting to the router 2
ip address 10.101.1.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface GigabitEthernet4 //interface connecting to the router 4
ip address 10.101.1.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface loopback1
ip address 10.1.1.1/32
ip ospf 10 area 0
```

Configuration at the tail-end router R6:

```
router ospf 10
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
segment-routing mpls
mpls traffic-eng area 0
mpls traffic-eng router-id Loopback1
interface GigabitEthernet2 //interface connecting to the router 3
ip address 10.101.2.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels
```

Explicit Path SR-TE Tunnel 1

```

interface GigabitEthernet4 //interface connecting to the router 5
ip address 10.101.2.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface loopback1
ip address 10.6.6.6/32
ip ospf 10 area 0

```

Explicit Path SR-TE Tunnel 1

Consider tunnel 1 based only on IP addresses:

```

ip explicit-path name IP_PATH1
next-address 10.2.2.2
next-address 10.3.3.3
next-address 10.6.6.6
!
interface Tunnel1
ip unnumbered Loopback1
tunnel mode mpls traffic-eng
tunnel destination 10.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 6 6
tunnel mpls traffic-eng path-option 10 explicit name IP_PATH1 segment-routing
tunnel mpls traffic-eng path-selection metric igp
tunnel mpls traffic-eng load-share 10
end

```

Explicit Path SR-TE Tunnel 2

Consider tunnel 2 based on node SIDs

```

ip explicit-path name IA_PATH
next-label 114
next-label 115
next-label 116
!
interface Tunnel2
ip unnumbered Loopback1
tunnel mode mpls traffic-eng
tunnel destination 10.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 6 6
tunnel mpls traffic-eng bandwidth 10000 class-type 1
tunnel mpls traffic-eng path-option 10 explicit name IA_PATH segment-routing
tunnel mpls traffic-eng path-selection metric igp
tunnel mpls traffic-eng load-share 10
end

```

Explicit Path SR-TE Tunnel 3

Consider that tunnel 3 is based on a mix of IP addresses and label

```

ip explicit-path name MIXED_PATH enable
next-address 10.2.2.2
next-address 10.3.3.3
next-label 115
next-label 116

```

```

!
interface Tunnel3
 ip unnumbered Loopback1
 tunnel mode mpls traffic-eng
 tunnel destination 10.6.6.6
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 6 6
 tunnel mpls traffic-eng path-option 10 explicit name MIXED_PATH segment-routing
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng load-share 10

```



Note In the case of mixed path, IP next-hop cannot be used after using Node SIDs in the path. The following path will not be valid:

```

ip explicit-path name MIXED_PATH enable
next-label 115
next-label 116
next-address 10.2.2.2

```

Dynamic Path SR-TE Tunnel 4

Consider that tunnel 4 is based on adjacency SIDs

```

interface Tunnel4
 ip unnumbered Loopback1
 tunnel mode mpls traffic-eng
 tunnel destination 10.6.6.6
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 6 6
 tunnel mpls traffic-eng bandwidth 10000 class-type 1
 tunnel mpls traffic-eng path-option 10 dynamic segment-routing
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng load-share 10
end

```

Dynamic Path SR-TE Tunnel 5

Consider that tunnel 5 is based on Node SIDs

```


interface Tunnel5
 ip unnumbered Loopback1
 tunnel mode mpls traffic-eng
 tunnel destination 10.6.6.6
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 6 6
 tunnel mpls traffic-eng path-option 10 segment-routing
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng load-share 10

```

Configuring Inter Area Tunnel

Consider the following topology to configure inter area tunnel:

Figure 2: Inter Area Tunnel



All the routers are configured in the same area, area 0 except R6 which is configured in area 1.

Configuration at the head end router R1:

```
router ospf 10
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
segment-routing mpls
mpls traffic-eng router-id Loopback1
mpls traffic-eng area 0

interface GigabitEthernet2 //interface connecting to the router 2
ip address 10.101.1.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface GigabitEthernet4 //interface connecting to the router 4
ip address 10.101.1.1 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface loopback1
ip address 10.1.1.1/32
ip ospf 10 area 0
```

Configuration at the tail-end router R6:

```
router ospf 10
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
segment-routing mpls
mpls traffic-eng area 1
mpls traffic-eng router-id Loopback1
interface GigabitEthernet2 //interface connecting to the router 3
ip address 10.101.2.1 255.255.255.0
ip ospf 10 area 1
ip ospf network point-to-point
```



```

negotiation auto
mpls traffic-eng tunnels

interface GigabitEthernet4 //interface connecting to the router 5
ip address 10.101.2.1 255.255.255.0
ip ospf 10 area 1
ip ospf network point-to-point
negotiation auto
mpls traffic-eng tunnels

interface loopback1
ip address 10.6.6.6/32
ip ospf 10 area 1

```

Restrictions for Configuring Inter Area Tunnel

The following are the restrictions for configuring inter area tunnel:

- The dynamic option with node and adjacency SID are not supported.
- You can configure inter are tunnel using the explicit path containing only labels and/or IP address and labels.



Note The IP address can be used only be till the Area Border Router (ABR) and after that you need to specify only the labels.

Explicit Path SR-TE Tunnel 1

Consider tunnel 2 is based on node SIDs.

```

ip explicit-path name IA_PATH
next-label 114
next-label 115
next-label 116
!
interface Tunnel2
ip unnumbered Loopback1
tunnel mode mpls traffic-eng
tunnel destination 10.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 6 6
tunnel mpls traffic-eng bandwidth 10000 class-type 1
tunnel mpls traffic-eng path-option 10 explicit name NODE_PATH segment-routing
tunnel mpls traffic-eng path-selection metric igp
tunnel mpls traffic-eng load-share 10
end

```

Explicit Path SR-TE Tunnel 2

Consider that tunnel 3 is based on a mix of IP Addresses and label.

```

ip explicit-path name MIXED_PATH enable
next-address 10.2.2.2
next-address 10.3.3.3
next-label 115
next-label 116
!

```

```

interface Tunnel3
ip unnumbered Loopback1
tunnel mode mpls traffic-eng
tunnel destination 10.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 6 6
tunnel mpls traffic-eng path-option 10 explicit name MIXED_PATH segment-routing
tunnel mpls traffic-eng path-selection metric igp
tunnel mpls traffic-eng load-share 10

```

Verifying Configuration of the SR-TE Tunnels

Use the **show mpls traffic-eng tunnels *tunnel-number*** command to verify the configuration of the SR-TE tunnels.

Verifying Tunnel 1

```

Name: R1_t1 (Tunnel1) Destination: 10.6.6.6
Status:
  Admin: up Oper: up Path: valid Signalling: connected
  path option 10, (SEGMENT-ROUTING) type explicit IP_PATH (Basis for Setup)
Config Parameters:
  Bandwidth: 0 kbps (Global) Priority: 6 6 Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
    Path-invalidation timeout: 45000 msec (default), Action: Tear
    AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: explicit path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours
    Time since path change: 2 seconds
    Number of LSP IDs (Tun_Instances) used: 1814
    Current LSP: [ID: 1814]
    Uptime: 2 seconds
    Selection: reoptimization
    Prior LSP: [ID: 1813]
    ID: path option unknown
    Removal Trigger: configuration changed
  Tun_Instance: 1814
Segment-Routing Path Info (ospf 10 area 0)
  Segment0[Node]: 10.4.4.4, Label: 114
  Segment1[Node]: 10.5.5.5, Label: 115
  Segment2[Node]: 10.6.6.6, Label: 116

```

Verifying Tunnel 2

```

Name: R1_t2 (Tunnel1) Destination: 10.6.6.6

```

```

Status:
  Admin: up          Oper: up      Path: valid      Signalling: connected
  path option 10, (SEGMENT-ROUTING) type explicit IA_PATH (Basis for Setup)
Config Parameters:
  Bandwidth: 0      kbps (Global) Priority: 6 6  Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
  Path-invalidation timeout: 45000 msec (default), Action: Tear
  AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
  auto-bw: disabled
  Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: explicit path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours, 1 minutes
    Time since path change: 1 seconds
    Number of LSP IDs (Tun_Instances) used: 1815
  Current LSP: [ID: 1815]
    Uptime: 1 seconds
  Prior LSP: [ID: 1814]
    ID: path option unknown
    Removal Trigger: configuration changed
Tun_Instance: 1815
Segment-Routing Path Info (ospf 10 area 0)
  Segment0[ - ]: Label: 114
  Segment1[ - ]: Label: 115
  Segment2[ - ]: Label: 116

```

Verifying Tunnel 3

```

Name: R1_t3                                     (Tunnel1) Destination: 10.6.6.6
Status:
  Admin: up          Oper: up      Path: valid      Signalling: connected
  path option 10, (SEGMENT-ROUTING) type explicit MIXED_PATH (Basis for Setup)
Config Parameters:
  Bandwidth: 0      kbps (Global) Priority: 6 6  Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
  Path-invalidation timeout: 45000 msec (default), Action: Tear
  AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
  auto-bw: disabled
  Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: explicit path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours, 2 minutes
    Time since path change: 2 seconds
    Number of LSP IDs (Tun_Instances) used: 1816
  Current LSP: [ID: 1816]
    Uptime: 2 seconds
    Selection: reoptimization
  Prior LSP: [ID: 1815]
    ID: path option unknown
    Removal Trigger: configuration changed
Tun_Instance: 1816
Segment-Routing Path Info (ospf 10 area 0)

```

```

Segment0[Node]: 10.2.2.2, Label: 112
Segment1[Node]: 10.3.3.3, Label: 113
Segment2[ - ]: Label: 115
Segment3[ - ]: Label: 116

```

Verifying Tunnel 4

```

Name: R1_t4 (Tunnel1) Destination: 10.6.6.6
Status:
  Admin: up      Oper: up      Path: valid      Signalling: connected
  path option 10, (SEGMENT-ROUTING) type dynamic (Basis for Setup, path weight 30)
Config Parameters:
  Bandwidth: 0      kbps (Global) Priority: 6 6 Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
    Path-invalidation timeout: 45000 msec (default), Action: Tear
    AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: dynamic path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:
    Time since created: 6 days, 19 hours
    Time since path change: 2 seconds
    Number of LSP IDs (Tun_Instances) used: 1813
    Current LSP: [ID: 1813]
    Uptime: 2 seconds
    Prior LSP: [ID: 1806]
    ID: path option unknown
    Removal Trigger: configuration changed
  Tun_Instance: 1813
Segment-Routing Path Info (ospf 10 area 0)
  Segment0[Link]: 192.168.2.1 - 192.168.2.2, Label: 17
  Segment1[Link]: 192.168.4.2 - 192.168.4.1, Label: 25
  Segment2[Link]: 192.168.8.1 - 192.168.8.2, Label: 300

```

Verifying Tunnel 5

```

Name: R1_t5 (Tunnel1) Destination: 10.6.6.6
Status:
  Admin: up      Oper: up      Path: valid      Signalling: connected
  path option 10, type segment-routing (Basis for Setup)
Config Parameters:
  Bandwidth: 0      kbps (Global) Priority: 6 6 Affinity: 0x0/0xFFFF
  Metric Type: IGP (interface)
  Path Selection:
    Protection: any (default)
    Path-invalidation timeout: 45000 msec (default), Action: Tear
    AutoRoute: enabled LockDown: disabled Loadshare: 10 [200000000]
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Active Path Option Parameters:
  State: segment-routing path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
History:
  Tunnel:

```

```
Time since created: 6 days, 19 hours, 4 minutes
Time since path change: 14 seconds
Number of LSP IDs (Tun_Instances) used: 1817
Current LSP: [ID: 1817]
  Uptime: 14 seconds
  Selection: reoptimization
Prior LSP: [ID: 1816]
  ID: path option unknown
  Removal Trigger: configuration changed
Tun_Instance: 1817
Segment-Routing Path Info (ospf 10 area 0)
Segment0[Node]: 10.6.6.6, Label: 116
```

