



Segment Routing Traffic Engineering

Segment Routing (SR) enables any network node, such as a Server, Provider Edge (PE), Aggregator, or Provider (P) to engineer an explicit path for each of its traffic classes.



Note

This explicit path does not depend on a hop-by-hop signaling technique, such as Label Distribution Protocol (LDP) or Resource Reservation Protocol (RSVP); it only depends on a set of "segments" that are preprogrammed and advertised by the link-state routing protocol.

These segments act as topological sub-paths that can be combined together to form the desired path. In Segment Routing, the path is encoded in each packet itself in the form of SR Segment Identifiers (SIDs).

There are two types of segments—prefix and adjacency.

- A prefix segment represents the shortest path (as computed by IGP) to reach a specific prefix; a node segment is a special prefix segment that is bound to the loopback address of a node.
- An adjacency segment represents a specific adjacency to a neighbor node.

A node segment can be a multi-hop path while an adjacency segment is a one-hop path.

A segment is represented by a 32-bit entity called Segment ID (SID). A prefix-SID is globally unique, and the operator ensures such uniqueness. An adjacency SID is locally unique to the node, and is automatically generated by the node attached to the adjacency.

The Segment Routing control-plane can be applied to the MPLS data-plane. In this case, the prefix-SID in the MPLS data-plane is represented as an LSP whose path flows along the shortest-path to the prefix, whereas an adjacency-SID is represented as cross-connect entry pointing to a specific egress data-link.

A Traffic Engineered (TE) tunnel is a container of TE LSPs instantiated between the tunnel ingress and the tunnel destination. A TE tunnel can instantiate one or more SR-TE LSPs 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 through either a set of prefix-SIDs, or adjacency-SIDs of nodes, or both, and links to be traversed by the SR-TE LSP.

The headend imposes the corresponding MPLS label stack on 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.

The set of hops or segments that define an SR-TE LSP path are provisioned by the operator.

- [Restrictions for SR-TE, on page 2](#)

- [SR-TE LSP Instantiation, on page 9](#)
- [SR-TE and TI-LFA, on page 16](#)

Restrictions for SR-TE

- SR-TE statistics counters not supported.
- The Cisco ASR 920 routers support five label stacks of which four are transport labels.



Note

All five labels cannot be part of the SR-TE Tunnel label stack. One label must be a service label.

- The routers do not support unequal load balancing when using the load-share option.
- ECMP at single SR-TE tunnel level is not supported.
- SR-TE FRR with PoCH as the primary path is not supported. However, you can provision SR-TE tunnel without FRR over PoCH and backup tunnel over PoCH.
- The SR-TE dynamic tunnels do not support node protection. Therefore, node protection cannot achieve less than 50ms convergence with dynamic SR-TE tunnels.

Configuring a Path Option for a TE Tunnel

The **segment-routing** keyword indicates that the specified path is programmed as an SR path.

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

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 is 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.

Configuring SR Explicit Path Hops

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

```
Router(config)# ip explicit-path name foo
Router(config-ip-expl-path)# index 10 next-address 1.1.1.1 --> node address
```

```
Router(config-ip-expl-path)# index 20 next-address 12.12.12.2 --> link address
```

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

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

The following SR-TE explicit path hops are supported:


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



Note IP addresses cannot be used after using the label in MIXED_PATH.

Use Case: Segment Routing Traffic Engineering Basic Configuration

Consider the following topology:



Configuration at the headend router, R1:

```
interface GigabitEthernet0/0/2
ip address 100.101.1.1 255.255.255.0
ip router isis 1
isis network point-to-point
negotiation auto
mpls traffic-eng tunnels
router isis 1
net 49.0001.0010.0100.1001.00
```

```

is-type level-1
ispf level-1
metric-style wide
log-adjacency-changes
segment-routing mpls
segment-routing prefix-sid-map advertise-local
mpls traffic-eng router-id Loopback1
mpls traffic-eng level-1

```

Configuration at the tailend router, R6

```

interface GigabitEthernet0/02
ip address 100.101.1.1 255.255.255.0
ip router isis 1
isis network point-to-point
negotiation auto
mpls traffic-eng tunnels
router isis 1
 net 49.0001.0060.0600.6006.00
 ispf level-1
 metric-style wide
 log-adjacency-changes
 segment-routing mpls
segment-routing prefix-sid-map advertise-local
mpls traffic-eng router-id Loopback1
mpls traffic-eng level-1

```

Explicit Path SR-TE Tunnel 1

Consider tunnel 1 based only on IP addresses:

```

ip explicit-path name IP_PATH1
 next-address 2.2.2.2
 next-address 3.3.3.3
 next-address 6.6.6.6
!
interface Tunnel1
 ip unnumbered Loopback1 poll point-to-point
 tunnel mode mpls traffic-eng
 tunnel destination 6.6.6.6
 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 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 poll point-to-point
 tunnel mode mpls traffic-eng
 tunnel destination 6.6.6.6

```

```

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 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 2.2.2.2
next-address 3.3.3.3
next-label 115
next-label 116
!
interface Tunnel3
ip unnumbered Loopback1 poll point-to-point
tunnel mode mpls traffic-eng
tunnel destination 6.6.6.6
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 10

```

Dynamic Path SR-TE Tunnel 4

Consider that tunnel 4 is based on adjacency SIDs

```

interface Tunnel4
ip unnumbered Loopback1 poll point-to-point
tunnel mode mpls traffic-eng
tunnel destination 6.6.6.6
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 10
end

```

Dynamic Path SR-TE Tunnel 5

Consider that tunnel 5 is based on Node SIDs

```

interface Tunnel5
ip unnumbered Loopback1 poll point-to-point
tunnel mode mpls traffic-eng
tunnel destination 6.6.6.6
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 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: 6.6.6.6
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 (isis level-1)
  Segment0[Node]: 4.4.4.4, Label: 114
  Segment1[Node]: 5.5.5.5, Label: 115
  Segment2[Node]: 6.6.6.6, Label: 116
```

Verifying Tunnel 2

```
Name: R1_t2                                     (Tunnel1) Destination: 6.6.6.6
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 (isis level-1)
    Segment0[ - ]: Label: 114
    Segment1[ - ]: Label: 115
    Segment2[ - ]: Label: 116

```

Verifying Tunnel 3

```

Name: R1_t3                                     (Tunnell) Destination: 6.6.6.6
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 (isis level-1)
    Segment0[Node]: 2.2.2.2, Label: 112
    Segment1[Node]: 3.3.3.3, Label: 113
    Segment2[ - ]: Label: 115
    Segment3[ - ]: Label: 116

```

Verifying Tunnel 4

```

Name: R1_t4                                     (Tunnell) Destination: 6.6.6.6
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 (isis level-1)
  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: 6.6.6.6
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 (isis level-1)
  Segment0[Node]: 6.6.6.6, Label: 116

```


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.

SR-TE LSP Explicit Null

MPLS-TE tunnel headend 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 tailend without a transport label. However, sometimes, it is desirable that the packet arrive at the tailend with explicit-null label, and in such cases, the headend imposes an explicit-null label at the top of the label stack.

SR-TE LSP Path Verification

SR-TE tunnel functionality requires that the headend perform initial verification of the tunnel path as well as the subsequent tracking of the reachability of the tunnel tailend 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 headend validates the path of an SR-TE LSP for connectivity against the TE topology. MPLS-TE headend checks if links corresponding to the adjacency SIDs are connected in the TE topology.

For newly instantiated SR-TE LSPs, if the headend 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 headend detects a discontinuity on any link, the headend assumes that a fault has occurred on that link. In this case, the local repair protection, such as the IP FRR, comes in to effect. The IGPs 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 headend 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 headend 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 headend starts an invalidation timer after which it brings the tunnel state to 'down'. This action avoids black-holing the traffic flowing over the affected SR-TE LSP, and allows services riding over the tunnel to reroute over different available paths at the headend. 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 headend has full visibility over the LSP path, and validates the path to the ultimate LSP destination. However, for interarea LSPs, the headend has partial visibility over the LSP path—only up to the first ABR. In this case, the headend 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 headend, 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 IGPs 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 headend 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 that 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 headend 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 Affinity Validation

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

The headend validates that the specified SR path is compliant with the configured affinity. This requires 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 Tunnell
no ip address
tunnel mode mpls traffic-eng
tunnel destination 5.5.5.5
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 (Tunnell) Destination: 5.5.5.5
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

```
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
```

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 headend 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 following commands. 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
!
```

Tunnel Path Loop Validation

The SR path is a concatenation of SR segments (combination of prefix and adjacency SIDs). It is possible that any of the traversed segment's underlying paths may traverse through the ingress of the tunnel. In this case, packets that are mapped on the SR tunnel may loop back again to the headend. To avoid this sub-optimal path, the headend detects and invalidates a looping SR path through the ingress node.

Loop path validation is implicitly enabled on SR path. However, it is possible to disable this validation by using the **verbatim** path-option keyword associated with the tunnel path-option.

The following is an example of the **verbatim** path-option keyword when IP address 6.6.6.6 is in a different area:

```
interface Tunnell
 tunnel mode mpls traffic-eng
 tunnel destination 6.6.6.6
 tunnel mpls traffic-eng priority 5 5
 tunnel mpls traffic-eng bandwidth 100
 tunnel mpls traffic-eng path-option 10 explicit name NODE_PATH segment-routing verbatim
Name: R1_t1 (Tunnell) Destination: 6.6.6.6
Status:
  Admin: up      Oper: up      Path: valid      Signalling: connected
  path option 10, (SEGMENT-ROUTING) type explicit (verbatim) NODE_PATH (Basis for Setup)
Config Parameters:
  Bandwidth: 100      kbps (Global)  Priority: 5 5      Affinity: 0x0/0xFFFF
  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 [ignore: Verbatim Path Option]
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: explicit path option 10 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: enabled
History:
  Tunnel:
    Time since created: 7 minutes, 43 seconds
    Time since path change: 0 seconds
    Number of LSP IDs (Tun_Instances) used: 49
    Current LSP: [ID: 49]
    Uptime: 0 seconds
    Prior LSP: [ID: 48]
    ID: path option unknown
    Removal Trigger: signalling shutdown
  Tun_Instance: 49
Segment-Routing Path Info (isis level-1)
  Segment0[Node]: 2.2.2.2, Label: 20012
  Segment1[Node]: 3.3.3.3, Label: 20013
  Segment2[ - ]: Label: 20016

```

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 ECMPs

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 any midpoint traversed node along the SR-TE LSP path.



Note ECMP within a single SR-TE tunnel is not supported.

Load Balancing on Multiple Tunnels

ECMP across multiple SR-TE tunnels is not supported.

SR-TE Tunnel Re-optimization

TE tunnel re-optimization occurs when the headend determines that there is a more optimal path available than the one currently used. For example, if there is a failure along the SR-TE LSP path, the headend could detect and revert to a more optimal path by triggering re-optimization.

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 headend 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 headend 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 headend is still using the failed path because it is unable to reroute on a different path, the tunnel state is brought 'down' to avoid black-holing 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 re-optimization 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 tunnelt1
Name: R1_t1 (Tunnelt1) Destination: 6.6.6.6
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]: 4.4.4.4, Label: 114
  Segment1[Node]: 5.5.5.5, Label: 115
  Segment2[Node]: 6.6.6.6, Label: 116
```



Note SR-TE does not support lossless re-optimization with multiple path options.



Note When FRR is configured and the primary path is brought back up, re-optimization time is in the order of seconds due to microloop.

SR-TE With lockdown Option

The **lockdown** option only 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 poll point-to-point
 tunnel mode mpls traffic-eng
 tunnel destination 6.6.6.6
 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-option 20 segment-routing
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng 10
Router# show mpls traffic-eng tunnels tunnell
Name: csr551_t1 (Tunnell) Destination: 6.6.6.6
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]: 6.6.6.6, Label: 116
```

SR-TE Tunnel Protection

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



Note

50-millisecond traffic protection is not guaranteed for path protection scenarios.

IP-FRR Local Repair Protection

On an SR-TE LSP headend 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 *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 headend 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 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 headend 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.

If the primary SR-TE LSP fails, 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 and TI-LFA

Restrictions for Using SR-TE and TI-LFA

- In case of primary and secondary path switchover, a microloop is created between routers. Use the **microloop avoidance rib-update-delay** command to bring down the convergence time

Consider the following topology:

-----ixia-2

|

------(R4)-----

||

||

(R3) (R1) ----ixia-1

||

||

------(R2)-----

```
router isis ipfrr
net 49.0001.0120.1201.2012.00
is-type level-2-only
ispf level-2
metric-style wide
log-adjacency-changes
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix level-2 all
fast-reroute ti-lfa level-2
microloop avoidance rib-update-delay 10000
```

- Before you configure SR-TE for TI-LFA, you must enable TI-LFA is enabled on all nodes. For more information, see *Using Segment Routing with OSPF* .

```
mpls traffic-eng tunnels
!
segment-routing mpls
connected-prefix-sid-map
address-family ipv4
1.1.1.1/32 index 11 range 1
exit-address-family
!
interface Loopback1
ip address 1.1.1.1 255.255.255.255
ip router isis 1
!
interface Tunnel1
ip unnumbered Loopback1 poll point-to-point
tunnel mode mpls traffic-eng
tunnel destination 6.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng path-option 10 explicit name IP_PATH segment-routing
!
interface GigabitEthernet2
ip address 192.168.1.1 255.255.255.0
ip router isis 1
negotiation auto
mpls traffic-eng tunnels
isis network point-to-point
!
interface GigabitEthernet3
ip address 192.168.2.1 255.255.255.0
ip router isis 1
negotiation auto
mpls traffic-eng tunnels
isis network point-to-point
!
router isis 1
```

```

net 49.0001.0010.0100.1001.00
is-type level-1
ispf level-1
metric-style wide
log-adjacency-changes
segment-routing mpls
fast-reroute per-prefix level-1 all
fast-reroute ti-lfa level-1
mpls traffic-eng router-id Loopback1
mpls traffic-eng level-1
!
ip explicit-path name IP_PATH enable
next-address 4.4.4.4
next-address 5.5.5.5
next-address 6.6.6.6

```

- To reduce or minimize traffic loss after a high availability (HA) switchover, MPLS TE NSR and IS-IS NSF must be enabled.

Use the **mpls traffic-eng nsr** command in global EXEC mode.

```
mpls traffic-eng nsr
```

Use the **nsf** command under IS-IS or OSPF.

```

router isis
nsf cisco
nsf interval 0

```

- The Cisco ASR routers support 500 SR-TE tunnels with two transport labels, two TI-LFA protection labels and one service label.
- SSO is not supported with SR-TE on the Cisco RSP2 Module.
- For TI-LFA restrictions, see [Restrictions for the TI-LFA](#).

Verifying the SR-TE With TI_LFA Configuration

```

Router# show mpls traffic-eng tunnels tunnell1
Name: PE1 (Tunnell1) Destination: 6.6.6.6
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: 7 7 Affinity: 0x0/0xFFFF
  Metric Type: TE (default)
  Path Selection:
    Protection: any (default)
  Path-invalidation timeout: 45000 msec (default), Action: Tear
  AutoRoute: enabled LockDown: disabled Loadshare: 0 [0] bw-based
  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: 4 hours, 25 minutes
    Time since path change: 4 hours, 21 minutes
    Number of LSP IDs (Tun_Instances) used: 37
    Current LSP: [ID: 37]
    Uptime: 4 hours, 21 minutes

```

```

Tun_Instance: 37
Segment-Routing Path Info (isis level-1)
  Segment0[Node]: 4.4.4.4, Label: 16014
  Segment1[Node]: 5.5.5.5, Label: 16015
  Segment2[Node]: 6.6.6.6, Label: 16016
Router# show isis fast-reroute ti-lfa tunnel

Tag 1:
Fast-Reroute TI-LFA Tunnels:
Tunnel  Interface  Next Hop      End Point      Label      End Point Host
MP1      Gi2             192.168.1.2   6.6.6.6        16016      SR_R6
MP2      Gi3             192.168.2.2   6.6.6.6        16016      SR_R6
Router# show frr-manager client
client-name
  ISIS interfaces detail
TunnelI/F : MP1
  Type : SR
  Next-hop : 192.168.1.2
  End-point : 6.6.6.6
  OutI/F : Gi2
  Adjacency State : 1
  Prefix0 : 6.6.6.6(Label : 16016)
TunnelI/F : MP2
  Type : SR
  Next-hop : 192.168.2.2
  End-point : 6.6.6.6
  OutI/F : Gi3
  Adjacency State : 1
  Prefix0 : 6.6.6.6(Label : 16016)
Router# show ip cef 6.6.6.6 internal

6.6.6.6/32, epoch 2, RIB[I], refcnt 6, per-destination sharing
sources: RIB, LTE
feature space:
  IPRM: 0x00028000
  Broker: linked, distributed at 1st priority
  LFD: 6.6.6.6/32 1 local label
  sr local label info: global/16016 [0x1A]
    contains path extension list
  sr disposition chain 0x7FC6B0BF2AF0
    label implicit-null
    IP midchain out of Tunnel1
    label 16016
    FRR Primary
    <primary: label 16015
      TAG adj out of GigabitEthernet3, addr 192.168.2.2>
  sr label switch chain 0x7FC6B0BF2B88
    label implicit-null
    TAG midchain out of Tunnel1
    label 16016
    FRR Primary
    <primary: label 16015
      TAG adj out of GigabitEthernet3, addr 192.168.2.2>
ifnums:
  Tunnel1(13)
path list 7FC6B0BDDDE0, 3 locks, per-destination, flags 0x49 [shble, rif, hwc]
  path 7FC7144D4300, share 1/1, type attached nexthop, for IPv4
  MPLS short path extensions: [rib | prfmfi | lblmrg | srlbl] MOI flags = 0x3 label
implicit-null
  nexthop 6.6.6.6 Tunnel1, IP midchain out of Tunnel1 7FC6B0BBB440
output chain:
  IP midchain out of Tunnel1 7FC6B0BBB440
  label [16016|16016]
  FRR Primary (0x7FC714515460)

```

```

<primary: label 16015
          TAG adj out of GigabitEthernet3, addr 192.168.2.2 7FC6B0BBB630>
<repair:  label 16015
          label 16014
          TAG midchain out of MPLS-SR-Tunnel1 7FC6B0BBAA90
          label 16016
          TAG adj out of GigabitEthernet2, addr 192.168.1.2 7FC6B0BBBA10>

```

**Note**

To ensure a less than 50-msec traffic protection with TI-LFA, SR-TE with dynamic path option must use the backup adjacency SID.

To create an SR-TE with dynamic path option, use the following configuration on every router in the topology:

```

router isis 1
fast-reroute per-prefix level-1 all

```

At the tunnel headend router:

```

interface Tunnel1
ip unnumbered Loopback1 poll point-to-point
tunnel mode mpls traffic-eng
tunnel destination 6.6.6.6
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng path-option 1 dynamic segment-routing
tunnel mpls traffic-eng path-selection segment-routing adjacency protected

```

Configuring TI-LFA With Mapping Server

Consider the following topology:

- IXIA-2 injects IS-IS prefixes, and IXIA-1 sends one-way traffic to IXIA-2
- In R1 10,000 prefixes are configured in the segment-routing mapping-server

The configuration on R1 is:

```

conf t
segment-routing mpls
global-block 16 20016
!
connected-prefix-sid-map
address-family ipv4
11.11.11.11/32 index 11 range 1
exit-address-family
!
!
mapping-server
!
prefix-sid-map
address-family ipv4
120.0.0.0/24 index 2 range 1 attach
200.0.0.0/24 index 1 range 1 attach
192.168.0.0/24 index 100 range 10000 attach
exit-address-family
!
!
!
!
!

```

```

interface Loopback0
ip address 11.11.11.11 255.255.255.255
ip router isis ipfrr
!
interface GigabitEthernet0/1/0
ip address 14.0.0.1 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/1/2
ip address 11.0.0.1 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/1/4
ip address 200.0.0.1 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
router isis ipfrr
net 49.0001.0110.1101.1011.00
is-type level-2-only
metric-style wide
log-adjacency-changes
nsf cisco
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix level-2 all
fast-reroute ti-lfa level-2
microloop avoidance rib-update-delay 10000

```

On R2 the configuration is

```

conf t
!
!
segment-routing mpls
!
connected-prefix-sid-map
address-family ipv4
12.12.12.12/32 index 12 range 1
exit-address-family
!
!
interface Loopback0
ip address 12.12.12.12 255.255.255.255
ip router isis ipfrr
!
interface GigabitEthernet0/1/0
ip address 12.0.0.1 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/1/1
ip address 11.0.0.2 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
router isis ipfrr

```

```

net 49.0001.0120.1201.2012.00
is-type level-2-only
metric-style wide
log-adjacency-changes
nsf cisco
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix level-2 all
fast-reroute ti-lfa level-2
microloop avoidance rib-update-delay 10000
!

```

On R3 the configuration is

```

conf t
!
mpls traffic-eng tunnels
!
segment-routing mpls
!
connected-prefix-sid-map
address-family ipv4
13.13.13.13/32 index 13 range 1
exit-address-family
!
!
interface Loopback0
ip address 13.13.13.13 255.255.255.255
ip router isis ipfrr
!
interface GigabitEthernet0/0/4
ip address 13.0.0.1 255.255.255.0
ip router isis ipfrr
load-interval 30
speed 1000
no negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/0/5
ip address 12.0.0.2 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
router isis ipfrr
net 49.0001.0130.1301.3013.00
is-type level-2-only
metric-style wide
log-adjacency-changes
nsf cisco
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix level-2 all
fast-reroute ti-lfa level-2
microloop avoidance rib-update-delay 10000
!

```

On R4 the configuration is:

```

conf t
!
mpls traffic-eng tunnels
!

```

```
segment-routing mpls
!
connected-prefix-sid-map
address-family ipv4
14.14.14.14/32 index 14 range 1
exit-address-family
!
!
interface Loopback0
ip address 14.14.14.14 255.255.255.255
ip router isis ipfrr
!
interface GigabitEthernet0/0/0
ip address 14.0.0.2 255.255.255.0
ip router isis ipfrr
negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/0/3
ip address 13.0.0.2 255.255.255.0
ip router isis ipfrr
speed 1000
no negotiation auto
isis network point-to-point
!
interface GigabitEthernet0/0/5
ip address 120.0.0.1 255.255.255.0
ip router isis ipfrr
speed 1000
no negotiation auto
isis network point-to-point
!
router isis ipfrr
net 49.0001.0140.1401.4014.00
is-type level-2-only
metric-style wide
log-adjacency-changes
nsf cisco
segment-routing mpls
segment-routing prefix-sid-map advertise-local
fast-reroute per-prefix level-2 all
fast-reroute ti-lfa level-2
microloop avoidance rib-update-delay 10000
!
```

