

IS-IS SRv6 Link-protection Topology Independent Loop Free Alternate Fast Reroute

This document describes the functionalities and IS-IS implementation of IPv6 Fast Re-Route feature (IPv6FRR) using Segment Routing over IPv6 (SRv6) Topology Independent Loop Free Alternative (TI-LFA) link protection.

- Feature Information for IS-IS SRv6 Link-protection TI-LFA FRR, on page 1
- Prerequisites for IS-IS SRv6 Link-protection TI-LFA FRR , on page 1
- Restrictions for IS-IS SRv6 Link-protection TI-LFA FRR, on page 2
- Information About IS-IS SRv6 Link-protection TI-LFA FRR, on page 2
- How to Configure IS-IS SRv6 Link-protection TI-LFA FRR , on page 4

Feature Information for IS-IS SRv6 Link-protection TI-LFA FRR

Table 1: Feature Information for IS-IS SRv6 Link-protection TI-LFA FRR

Feature Name	Releases	Feature Information
IS-IS SRv6 Link-protection Topology Independent Loop Free Alternate Fast Reroute	Cisco IOS XE 17.15.1a	This feature was introduced.
IPv6 GRE-TP tunnel as protected link support for SRv6 TI-LFA with ISIS	Cisco IOS XE 17.18.2	From Cisco IOS XE 17.18.2, this feature extends IPv6 GRE-TP tunnel as protected link support for SRv6 TILFA with ISIS.

Prerequisites for IS-IS SRv6 Link-protection TI-LFA FRR

- SRv6 must be enabled on all the nodes, before configuring SRv6 TI-LFA. To enable SRv6, see chapter Segment Routing over IPv6.

```
segment-routing srv6
 locators
  locator A
   prefix CAFE:0:601::/48
   format usid-f3216
```

```
  !
router isis 1
net 49.0000.1111.1111.1111.00
 is-type level-2-only
 router-id Loopback0
 metric-style wide
address-family ipv6
  multi-topology
  router-id Loopback0
  segment-routing srv6
   locator A
    level-2
exit-address-family
!
```

Restrictions for IS-IS SRv6 Link-protection TI-LFA FRR

- The IGP cannot guarantee loop-free paths when tunnels are used for protection or repair, as it is unaware of the tunnel's underlying physical path.
- Forwarding failures on tunnel interfaces can occur in cases where local link failures do not occur. For this reason, in order to quickly detect forwarding failures on tunnel interfaces it is advisable to use the BFD protocol.

  Note that aggressive BFD timers must be used with care on tunnel interfaces as they can lead to false indication of failures
- Due to limitations on this platform, IGP TI-LFA cannot guarantee 50ms FRR failover.

Information About IS-IS SRv6 Link-protection TI-LFA FRR

When the local LFA is enabled, there is not always a good coverage of the prefixes to be protected.

To overcome the above limitation, effective Cisco IOS-XE Release 17.15.1a, topology-independent LFA (TI-LFA) is supported on an SRv6-enabled network.

Topology-Independent Loop Free Alternate

TI-LFA provides supports for the following:

- Link Protection—The LFA provides repair path for failure of the link.
- Local LFA—Whenever a local LFA on the post convergence path is available, it is preferred over TI-LFA because local LFA does not require additional SID for the repair path. That is, the label for the PQ node is not needed for the release node.
- TI-LFA path to traverse multiple intersect or disjoint PQ nodes, up to the platform’s maximum supported labels—TI-LFA provides complete coverage of all prefixes.
- P2P interfaces for the protected link—TI-LFA protects P2P interfaces.
- Asymmetrical links—The ISIS metrics between the neighbors are not the same.
- Multi-homed (anycast) prefix protection—The same prefix may be originated by multiple nodes.

- Protected prefix filtering—The route-map includes or excludes a list of prefixes to be protected and the option to limit the maximum repair distance to the release node.
- Tiebreakers—A subset of existing tiebreakers, applicable to TI-LFA, is supported.

Topology Independent Loop Free Alternate Tie-break

Local and remote LFA use default or user-configured heuristics to break the tie when there is more than one path to protect the prefix. The attributes are used to trim down the number of repair paths at the end of the TI-LFA link protection computation before the load balancing. Local LFA and remote LFA support the following tiebreakers:

- Linecard-disjoint—Prefers the line card disjoint repair path
- Lowest-backup-path-metric—Prefers the repair path with lowest total metric
- Node-protecting—Prefers node protecting repair path
- SRLG-disjoint—Prefers SRLG disjoint repair path
- Load-sharing—Distributes repair paths equally among links and prefixes

When there are two repair paths for a particular prefix, the path that the output port on different line card than that of the primary port is chosen as the repair path. For TI-LFA link protection, the following tiebreakers are supported:

- Linecard-disjoint—Prefers the line card disjoint repair path.
- LC disjoint index—If both the repair paths are on the same line card as that of the primary path, then, both paths are considered as candidates. If one of the path is on a different line card, then that path is chosen as the repair path.
- SRLG index—If both the repair paths have the same SRLG ID as that of the primary path, then, both the paths are considered as candidates. If one of the path has a different srlg id, then path is chosen as the repair path.
- Node-protecting—For TI-LFA node protection, the protected node is removed when computing the post-convergence shortest path. The repair path must direct traffic around the protected node.

The SRLG ID can be configured for each interface. When there are two repair paths for a prefix, the configured SRLG ID for the repair path is compared with that of the primary path SRLG ID. If the SRLG IDs for the secondary path is different than that of the primary, that path is chosen as the repair path. This policy comes into effect only when the primary path is configured with an SRLG ID. It is possible to configure both node and SRLG protection modes for the same interface or the same protocol instance. In that case, an additional TI-LFA node-SRLG combination protection algorithm is run. The TI-LFA node-SRLG combination algorithm removes the protected node and all members of the interface with the same SRLG group when computing the post-convergence SPT.

Interface Fast Reroute Tiebreakers

Interface fast reroute (FRR) tiebreakers are also needed for TI-LFA node and SRLG protection. When interface and protocol instance FRR tiebreakers both are configured, the interface FRR tiebreakers take precedence over the protocol instance. When interface FRR tiebreakers are not configured, the interface inherits the protocol instance FRR tiebreakers.

The following interface FRR tiebreaker commands apply only to the particular interface.

```
isis fast-reroute tie-break
[level-1 | level-2] linecard-disjoint
priority
isis fast-reroute tie-break
[level-1 | level-2] lowest-backup-metric
priority
isis fast-reroute tie-break
[level-1 | level-2] node-protecting
priority
isis fast-reroute tie-break
[level-1 | level-2] srlg-disjoint
priority
isis fast-reroute tie-break
[level-1 | level-2] default
```

Tie-breaker default and explicit tie-breaker on the same interface are mutually exclusive.

The following tie-breakers are enabled by default on all LFAs:

- linecard-disjoint
- lowest-backup-metric
- srlg-disjoint

How to Configure IS-IS SRv6 Link-protection TI-LFA FRR

Perform the following steps to configure SRv6 Link-protection Topology Independent Loop Free Alternate Fast Reroute.

Configuring Topology Independent Loop Free Alternate Fast Reroute

You can enable TI-LFA using any of the following two methods:

1. **Protocol enablement**—Enables TI-LFA in router isis mode for all IS-IS interfaces. Optionally, use the interface command to exclude the interfaces on which TI-LFA should be disabled.

   For example, to enable TI-LFA for all IS-IS interfaces:

   ```
   router isis 1
    address-family ipv6
      fast-reroute per-prefix {level-1 | level-2} {all | route-map route-map-name}
      fast-reroute ti-lfa {level-1 | level-2} [maximum-metric value]
   ```



Note The **isis fast-reroute per-prefix level-x** command enables local LFA and is required to enable TI-LFA.

2. **Interface enablement**—Enable TI-LFA selectively on each interface.

   ```
   interface interface-name
     isis fast-reroute protection {level-1 | level-2}
     isis fast-reroute ti-lfa protection {level-1 | level-2} [maximum-metric value]
   ```

The **maximum-metric** option specifies the maximum repair distance which a node is still considered eligible as a release node.

When both interface and protocol are TI-LFA enabled, the interface configuration takes precedence over the protocol configuration. TI-LFA is disabled by default.

To disable TI-LFA on a particular interface, use the following command:

```
interface interface-name
  isis fast-reroute ti-lfa protection {level-1 | level-2} disable
```

Examples: Configuring IS-IS SRv6 Link-protection TI-LFA FRR

Example 1: Base IS-IS SRv6 TILFA configuration

```
segment-routing srv6
 locators
  locator A
   prefix CAFE:0:601::/48
   format usid-f3216
  !
router isis 1
net 49.0000.1111.1111.1111.00
 is-type level-2-only
 router-id Loopback0
 metric-style wide
address-family ipv6
  router-id Loopback0
  multi-topology
  segment-routing srv6
   locator A
  level-2
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
exit-address-family
!
```

Example 2: EnableTI-LFA node-protecting tie-breaker on all IS-IS level-2 interfaces with priority 100. All other tie-breakers are disabled.

```
router isis
 address-family ipv6

  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
  fast-reroute tie-break level-2 node-protecting 100
```

Example 3: Enable TI-LFA node-protecting tie-breaker with priority 100 and TI-LFA SRLG protection with priority 200 on all IS-IS level-2 interfaces. All other tie-breakers are disabled because the node-protecting tie-breaker is configured.

```
router isis
 address-family ipv6

   fast-reroute per-prefix level-2 all
   fast-reroute ti-lfa level-2
   fast-reroute tie-break level-2 node-protecting 100
   fast-reroute tie-break level-2 srlg-disjoint 200
```

Example 4: Enable TI-LFA node-protecting tie-breaker with priority 100 on all IS-IS level-2 interfaces except on Ethernet0/0. For those IS-IS interfaces, all other tie-breakers are disabled. Ethernet0/0 overwrites the inheritance and uses the default set of tie-breakers with linecard-disjoint, lowest-backup-path-metric, srlg-disjoint enabled.

```
router isis
address-family ipv6
 fast-reroute per-prefix level-2 all
 fast-reroute ti-lfa level-2
 fast-reroute tie-break level-2 node-protecting 100
!
 interface ethernet0/0
  ipv6 router isis
  isis ipv6 fast-reroute tie-break level-2 default
```

Example 5: Enable TI-LFA using the default tie-breaker on all IS-IS interfaces except on Ethernet0/0. On Ethernet0/0, enable TI-LFA node-protecting with priority 100 and disable all other tie-breakers.

```
router isis
address-family ipv6
 fast-reroute per-prefix level-2 all
 fast-reroute ti-lfa level-2
!
 interface ethernet0/0
  ipv6 router isis
  isis ipv6 fast-reroute tie-break level-2 node-protecting 100
```

Example 6: Enable TI-LFA node-protecting tie-breaker with priority 200 and linecard-disjoint tie-breaker with priority 100 on all IS-IS level-2 interfaces. All other tie-breakers are disabled.

```
router isis
address-family ipv6

 fast-reroute per-prefix level-2 all
 fast-reroute ti-lfa level-2
 fast-reroute tie-break level-2 linecard-disjoint 100
 fast-reroute tie-break level-2 node-protecting 200
```

Verifying the Tiebreaker

To view tiebreakers enabled on the interface, use the following command:

show running all | section interface *interface-name*

To view tiebreakers enabled on the router mode, use the following command:

show running all | section router isis

Verifying the Primary and Repair Paths

Use the following show commands to verify primary and repair paths:

```
router#show isis ipv6 rib 605::1/128

IS-IS IPv6 process 1, local RIB

Repair path attributes:
```

```
    DS - Downstream, LC - Linecard-Disjoint, NP - Node-Protecting
    PP - Primary-Path, SR - SRLG-Disjoint

* 605::1/128  prefix attr X:0 R:0 N:1
    via FE80::A8BB:CCFF:FE02:5E20/Ethernet0/2, type L2  metric 30 tag 0
    prefix attr: X:0 R:0 N:1
     (installed)
     repair path: via FE80::A8BB:CCFF:FE02:5A00/Ethernet0/0 metric: 30 (DS,NP)
      TI-LFA node-protecting, link-protecting
        SRv6-Fwd-Id 25165857
        P node: r604 SID CAFE:0:604:: uN (PSP/USD)
      repair source: r605, metric to pfx: 60
```

```
router#show ipv6 route 605::1/128
Routing entry for 605::1/128
  Known via "isis 1", distance 115, metric 30, type level-2
  Route count is 1/1, share count 0
  Routing paths:
    FE80::A8BB:CCFF:FE02:5E20, Ethernet0/2 [Primary, repair via " Bind Label: 25165857"]
      Route metric is 30, traffic share count is 1
      From FE80::A8BB:CCFF:FE02:5E20
      Last updated 00:59:08 ago
     Bind Label: 25165857 [Repair]
      Route metric is 30, traffic share count is 1
      Last updated 00:59:08 ago
```

```
router#show ipv6 cef 605::1/128
605::1/128
  nexthop FE80::A8BB:CCFF:FE02:5E20 Ethernet0/2
    repair: recursive 25165857
```

```
router#show ipv6 cef 605::1/128 internal
605::1/128, epoch 0, RIB[I], refcnt 4, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00028000
  ifnums:
    Ethernet0/2(4): FE80::A8BB:CCFF:FE02:5E20
  path list 7F514ADDD0E0, 43 locks, per-destination, flags 0x16D [shble, hvsh, rif, rcrsv,
 hwcn, bldmp]
    path 7F514ABD30B8, share 1/1, type attached nexthop, for IPv6, flags [has-rpr]
      nexthop FE80::A8BB:CCFF:FE02:5E20 Ethernet0/2, IPV6 adj out of Ethernet0/2, addr
FE80::A8BB:CCFF:FE02:5E20 7F514B05AA48
        repair: recursive 25165857[Binding-Label:Default] (7F514ABD3188)
    path 7F514ABD3188, share 1/1, type recursive, for IPv6, flags [rpr, rpr-only]
      recursive via 25165857[Binding-Label:Default], repair, fib 7F514B624830, 1 terminal
fib, blbl:Default:25165857
      path list 7F514ADDD2F0, 3 locks, per-destination, flags 0x28049 [shble, rif, hwcn,
sb-oce, spec-order]
          path 7F514ABD34C8, share 1/1, type attached nexthop, for IPv6
            nexthop FE80::A8BB:CCFF:FE02:5A00 Ethernet0/0, IPV6 adj out of Ethernet0/0,
addr FE80::A8BB:CCFF:FE02:5A00 7F5145297A28
  output chain:
    loadinfo 7F514B15E908, per-session, 1 choice, flags 0185, 25 locks
      flags [Per-session, for-rx-IPv6, 2buckets, indirection]
      1 hash bucket
        < 0 > FRR Primary (0x7F514B15D9B8)
                 <primary: IPV6 adj out of Ethernet0/2, addr FE80::A8BB:CCFF:FE02:5E20
7F514B05AA48>
                <repair:  SRv6 SID List OCE 0x7F514B05D4D8 (7) 1 Segments, Flags 0x0 [none]

                            Segment List (1) mode:insert
                              CAFE:0:604:: [SL-MSD:16 END-POP-MSD:16 SRH-Inst:1]
                           PushCounter(SRv6 Encap) 7F514489F2E0
```

```
                              SRv6 Encap OCE 0x7F514B05FC68 (4) fwd-id:25165857 CAFE:0:604::
                                Encap Flags: 00000000 [none]
                                IPv6 TC: 0   Flow Label: 0       Hop Limit: 64
                                  Src: 601::1
                                  Dst: CAFE:0:604::
                              IPV6 adj out of Ethernet0/0, addr FE80::A8BB:CCFF:FE02:5A00
7F5145297A28>
      Subblocks:  None
```

Verifying SRv6 Configuration

Use the **show segment-routing srv6 locator** command to view SRv6 locators:

```
router# show segment-routing srv6 locator
  Name                  ID    Algo    Prefix             Status    Flags
  myLoc1                 3      0      2001:0:8::/48        Up        U
  myLocBestEffort        5      0      2001:0:1::/48        Up        U
```

Use the **show isis srv6 locators** command to view SID locators:

```
router# show isis srv6 locators
ISIS SRv6 Locators:
Tag sr:
Name             Prefix                       Level
----             ------                       -----
loc1             FC01:101:2::/48              2
```

```
router# show isis srv6 locators detail
ISIS SRv6 Locators:
Tag sr:
Name             Prefix                        Level
----             ------                        -----
loc1             FC01:101:2::/48               2

Level-1 metric: 0
 Level-2 metric: 0
 End-SIDs:
   FC01:101:2::
```

Verifying the IS-IS Topology Independent Loop Free Alternate Paths

Use the following show commands to verify ISIS TI-LFA configuration:

```
router#show isis ipv6 fast-reroute ti-lfa fwd-ids

Tag 1:

SRv6-Fwd-Id: 25165858 via FE80::A8BB:CCFF:FE02:5A00 Ethernet0/0 (2)
  Uncompressed SID List, SID count: 1
    P node: r603.00-00 SID CAFE:0:603:: uN (PSP/USD)
  Compressed SID List, SID count: 1
    SID[1]: CAFE:0:603::, MSD SL:16 End-POP:16

SRv6-Fwd-Id: 25165857 via FE80::A8BB:CCFF:FE02:5A00 Ethernet0/0 (2)
  Uncompressed SID List, SID count: 1
    P node: r604.00-00 SID CAFE:0:604:: uN (PSP/USD)
  Compressed SID List, SID count: 1
    SID[1]: CAFE:0:604::, MSD SL:16 End-POP:16

SRv6-Fwd-Id: 25165856 via FE80::A8BB:CCFF:FE02:5E20 Ethernet0/2 (4)
  Uncompressed SID List, SID count: 1
    P node: r604.00-00 SID CAFE:0:604:: uN (PSP/USD)
```

```
  Compressed SID List, SID count: 1
    SID[1]: CAFE:0:604::, MSD SL:16 End-POP:16

router#
```

```
router# show isis ipv6 fast-reroute summary

Tag 1:
IPv6 Fast-Reroute Protection Summary:

Prefix Counts: Total Protected Coverage
High priority: 0 0 0%
Normal priority: 25 22 88%
Total: 25 22 88%
```

```
router# show isis ipv6 fast-reroute interfaces


Tag 1 - Fast-Reroute Platform Support Information:

SRv6 TI-LFA: Supported by platform
Level-1 MT-2: FRR: Enabled, TI-LFA: Enabled
Level-2 MT-2: FRR: Not Enabled, TI-LFA: Not Enabled

TenGigabitEthernet2/3/0.3: Protectable: Yes. Usable for repair: Yes
GigabitEthernet2/1/1.2: Protectable: Yes. Usable for repair: Yes
GigabitEthernet2/1/1.1: Protectable: Yes. Usable for repair: Yes
Tunnel122: Protectable: Yes. Usable for repair: Yes
Tunnel121: Protectable: Yes. Usable for repair: Yes
Tunnel16: Protectable: Yes. Usable for repair: Yes
Loopback0: Protectable: Yes. Usable for repair: Yes
```