

OSPFv2 Link-protection Topology Independent Loop Free Alternate Fast Reroute

This document describes OSPFv2 implementation of IP Fast Re-Route Feature (IP FRR) using TI -LFA (Topology Independent Loop Free Alternative).

- Feature Information for OSPFv2 Link-protection Topology Independent Loop Free Alternate Fast Reroute, on page 1
- Restrictions for Topology Independent Loop Free Alternate Fast Reroute, on page 2
- Information About OSPFv2 Link-Protection Topology Independent Loop Free Alternate Fast Reroute, on page 3
- How to Configure Topology Independent Loop Free Alternate Fast Reroute, on page 10
- Debugging Topology Independent Loop Free Alternate Fast Reroute, on page 15
- Examples: OSPFv2 Link-Protection Topology Independent Loop Free Alternate Fast Reroute, on page 15

Feature Information for OSPFv2 Link-protection Topology Independent Loop Free Alternate Fast Reroute

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://cfnng.cisco.com/. An account on Cisco.com is not required.

Table 1: Feature Information for OSPFv2 Link-protection Topology Independent Loop Free Alternate Fast Reroute

Feature Name	Releases	Feature Information
OSPFv2 Link-protection Topology Independent Loop Free Alternate Fast Reroute	Cisco IOS XE Amsterdam 17.3.2	Topology-Independent Loop-Free Alternate (TI-LFA) uses segment routing to provide link, node, and Shared Risk Link Groups (SRLG) protection in topologies where other fast reroute techniques cannot provide protection. The goal of TI-LFA is to reduce the packet loss that results while routers converge after a topology change due to a link failure. The following commands were introduced or modified: **fast-reroute per-prefix ti-lfa [area <area> [disable]]**, **fast-reroute per-prefix tie-break node-protecting index <index>**, **fast-reroute per-prefix tie-break node-protecting required index <index>**, **fast-reroute per-prefix tie-break srlg index <index>**, **fast-reroute per-prefix tie-break srlg required index <index>**, **ip ospf fast-reroute per-prefix protection disable**, **ip ospf fast-reroute per-prefix candidate disable**, **show ip ospf fast-reroute ti-lfa tunnels**.
IPv4 GRE-TP tunnel as protected link support for SR-MPLS TI-LFA with OSPFv2	Cisco IOS XE 17.18.2	From Cisco IOS XE 17.18.2 this feature extends IPv4 GRE-TP tunnel as protected link support for SR-MPLS TILFA with OSPFv2.

Restrictions for Topology Independent Loop Free Alternate Fast Reroute

- TI-LFA is supported only on OSPFv2.
- TI-LFA tunnels are created only if the router supports SR and it is configured with prefix SID. The prefix (or) node SID can be configured as connected SID (or) advertised using the SRMS (Segment Routing Mapping Server).
- TI-LFA is not supported on OSPF point to multi point interfaces.
- TI-LFA does not support Multi Topology Routing (MTR).
- TI-LFA does not create the repair path using virtual link, sham link (or) TE tunnels.
- TI-LFA tunnel is constructed and programmed by explicitly specifying the node (or) set of repair nodes through which the tunnel needs to traverse.
- 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 OSPFv2 Link-Protection Topology Independent Loop Free Alternate Fast Reroute

Topology-Independent Loop-Free Alternate (TI-LFA) uses segment routing to provide link, node, and Shared Risk Link Groups (SRLG) protection in topologies where other fast reroute techniques, such as RLFA (Remote Loop Free Alternative) cannot provide protection. The goal of TI-LFA is to reduce the packet loss that results while routers converge after a topology change due to a link failure.

These are the major benefits of using TI-LFA:

- Provides 100% coverage for all the prefixes.
- Prevents transient congestion and sub-optimal routing by leveraging on the post-convergence path.
- Protects Label Distribution Protocol (LDP) and IP traffic as well.

IP Fast Reroute and Remote Loop Free Alternate

IP Fast Reroute (FRR) is a set of techniques that allow rerouting the IP traffic around a failed link or failed node in the network.One of the techniques that is used is Loop Free Alternates (LFA), which is implemented using OSPF protocol. OSPF currently supports per-prefix directly connected LFA and remote LFA (RLFA). The problem with these LFA algorithms is the topology dependency; the LFA algorithms cannot find a loop-free alternate path through the network for all the topologies.

The per-prefix directly connected LFA (also known as DLFA) provides loop-free alternate path for most triangular topologies, but does not provide good coverage for rectangular or circular topologies. The Remote LFA implementation (RLFA) which uses MPLS forwarding with LDP signaling for tunneling the rerouted traffic to an intermediate node, extends the IPFRR coverage in ring or rectangular topologies. For each link, RLFA defines P-Space (set of nodes reachable from calculating node without crossing the protected link) and Q-Space (set of nodes that can reach the neighbor on the protected link without crossing the protected link itself). The nodes that belong to both P and Q-Spaces are called PQ nodes and can be used as the intermediate node for the protected traffic. RLFA forms targeted LDP session to the PQ node and form the RLFA tunnel. But for the topologies where P and Q-Spaces are disjoint, R-LFA does not provide protection for those prefixes.

Topology Independent Fast Reroute

Topology Independent Fast Reroute (TI-FRR) is a technique which uses segment routing to provide link protection in any topology assuming the metric on the links in the topology is symmetrical. TI-LFA does not guarantee a backup in the cases where bandwidth on a single link is asymmetrical. TI-LFA only considers loop-free repair paths that are on the post-convergence path. It helps to do better capacity planning of the network.

TI-LFA algorithm allows to create a full explicit path through the network. Using fully specified path may lead to issues in larger topologies due to the number of segments along the path. Specifying the whole path is however not necessary, only a subset of the path is needed to carry the traffic to an intermediate node (release

node) which does not loop the traffic back to the protecting node. The TI-LFA algorithm constructs a SR tunnel as the repair path. TI-LFA tunnel is constructed and programmed by explicitly specifying the node (or) set of repair nodes through which the tunnel needs to traverse. The traffic is carried on the tunnel (when the primary path fails) which is also on the post convergence path.

Topology-Independent Loop Free Alternate

When the local LFA and remote LFA are enabled, there is a good coverage of the prefixes to be protected. However, for some rare topologies that do not have a PQ intersect node, both local and remote LFA will fail to find a release node to protect the failed link. Furthermore, there is no way to prefer a post-convergence path, as the two algorithms have no knowledge of the post-convergence characteristics of the LFA.

To overcome the above limitation, topology-independent LFA (TI-LFA) is supported on an SR-enabled network and provides the following support:

- **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.
- **Local LFA for extended P space**—For nodes in the extended P space, local LFA is still the most economical method for the repair path. In this case, TI-LFA is not chosen.
- **Tunnel to PQ intersect node**—This is similar to remote LFA except that the repair path is guaranteed on the post convergence path using TI-LFA.
- **Tunnel to PQ disjoint node**—This capability is unique to the TI-LFA in the case when local and remote LFA cannot find a repair path.
- **Tunnel to traverse multiple intersect or disjoint PQ nodes**—TI-LFA provides complete coverage of all prefixes, up to the platform's maximum supported labels.
- **P2P and Broadcast interfaces for the protected link**—TI-LFA protects P2P and broadcast interfaces.
- **Asymmetrical links**—The OSPF metrics between the neighbors are not the same.
- **Multi-homed (anycast) prefix protection**—The same prefix may be originated by multiple nodes and TI-LFA protects the anycast prefixes also by providing post convergence repair path.
- **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.
- **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.

- **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-disjoint**—Prefers the SRLG disjoint repair path.

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.

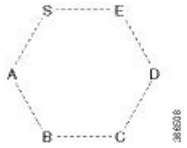
Effective with Cisco IOS-XE Release 3.18, node-protecting tie-breaker is disabled by default. Tie-breaker default and explicit tie-breaker on the same interface are mutually exclusive. The following tie-breakers are enabled by default on all LFAs:

- linecard-disjoint
- lowest-backup-metric
- SRLG-disjoint

P-Space

The set of routers that can be reached from S on the shortest path tree without traversing S-E is termed the P-space of S with respect to the link S-E.

Figure 1: A Simple Ring Topology



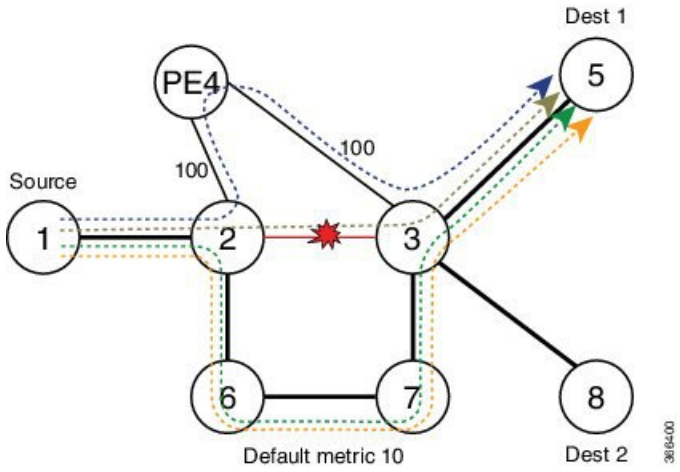
Q-Space

The set of routers from which the node E can be reached, by normal forwarding without traversing the link S-E, is termed the Q-space of E with respect to the link S-E.

Post-Convergence Path

Post convergence path is the path that OSPF uses after the link failure. TI-LFA always calculates the repair path which is the post convergence path. You can plan and dimension the post-convergence path to carry the traffic in the case of failure. TI-LFA enforces the post-convergence path by encoding it as a list of segments. The following figure shows an example of TI-LFA using post convergence path:

Figure 2: TI-LFA Using Post Convergence Path

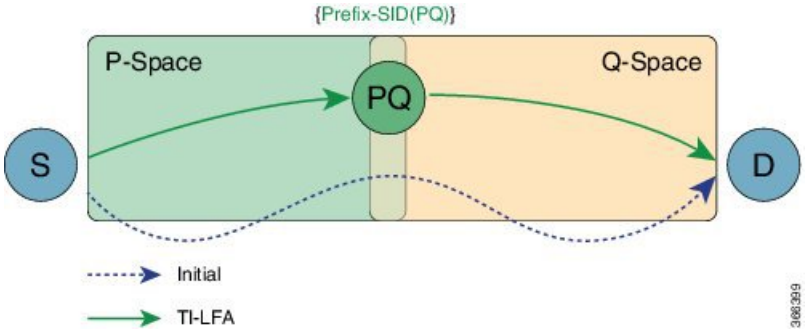


- It protects destination Node 5 on Node 2 against failure of link 2-3.
- Node 2 switches all the traffic destined to Node 5 via core links.

Per-Destination Link Protection

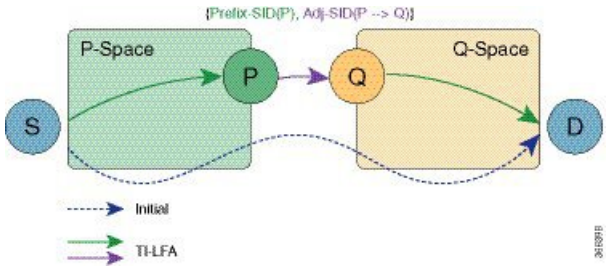
TI-LFA implementation provides per-destination link protection with the number of segments (labels)supported by the underlying hardware. The following figures show the implementation of TI-LFA:

Figure 3: TI-LFA: { Prefix-SID(PQ) }



If PQ is a direct neighbor of S, then no additional segment must be pushed.

Figure 4: TI-LFA: { Prefix-SID(P) , Adj -SID (P -> Q) }



Per Interface Loop Free Alternate Enablement

- TI-LFA can be enabled on an area basis.
- TI-LFA backup path is calculated only if TI-LFA protection is enabled on the primary interface which is to be protected. By default all the interfaces are enabled for protection.
- TI-LFA repair path is restricted by the number of labels supported by the hardware. If hardware supports only 2 labels then TI-LFA repair path can protect only those prefixes which can be protected by 2 or lesser segments. For those prefixes which need more than 2 segment remain unprotected.

Prefix Processing

Once TI-LFA path is calculated for the all the links, prefix processing starts. By default only intra and inter area prefixes are protected. For external prefixes to be protected, you need to enable segment routing globally under the OSPF level.

The primary and repair path should be of the same route type for the prefixes that are protected, that means, if the intra area needs to be protected then the TI-LFA repair path also calculates for the same intra area prefix whether the prefix is unique (or) anycast prefix.

Anycast Prefix Processing

OSPF TI-LFA also calculates the repair path for the anycast prefixes. Anycast prefixes (or) dual homed prefixes are the prefixes advertised by more than one routers. They could be intra, inter (or), external prefixes. The calculation of TI-LFA repair path for anycast prefixes is as below:

- Assume the prefix P1 is advertised by the routers R1 and R2. The prefix advertised by both the routers should be of the same route type, that is, both R1 and R2 should advertise the prefix as intra area prefix (or inter or external).
- Take the primary path is calculated towards R1 due to the lesser cost.
- When TI-LFA calculates the back up path, it calculates the post convergence path. So, post convergence path need not be towards R1. If the cost to reach R2 (in the post convergence) is shorter, then TI-LFA algorithm chooses the post convergence path towards R2. TI-LFA tunnel is formed towards R2.
- When R2 un-advertises the prefix, then the TI-LFA algorithm is re-calculated towards R1 for the repair path.

Per-Prefix Loop Free Alternate Tie-Break

IP FRR has the following tie break rules in the order given below. If you have more than one repair path available to choose the best path from, the following tie-break rules are applied. If more than one path matches all the tie break rules, then all the paths are used as repair paths.

- **Post Convergence**: Prefers backup path which is the post convergence path. This is enabled by default and user can not modify this.
- **Primary-path**: Prefers backup path from ECMP set.
- **Interface-disjoint**: Point-to-point interfaces have no alternate next hop for rerouting if the primary gateway fails. You can set the interface-disjoint attribute to prevent selection of such repair paths, thus protecting the interface.

- **Lowest-backup-metric**: Prefers backup path with lowest total metric. This is not applicable for TI-LFA since TI-LFA always chooses the back up path which is lowest cost.
- **LC-disjoint**: Prefers the back up path which is in different line card than that of the primary path.
- **Broadcast-interface-disjoint** : LFA repair paths protect links when a repair path and a protected primary path use different next-hop interfaces. However, on broadcast interfaces if the LFA repair path is computed via the same interface as the primary path and their next-hop gateways are different, in that case the node gets protected, but the link might not be. You can set the broadcast-interface-disjoint attribute to specify that the repair path never crosses the broadcast network the primary path points to, that means, it cannot use the interface and the broadcast network connected to it.
- **Load Sharing**: When more than one repair path matches the above rules, load share the backup paths. This rule also can be modified by the user.

Note The user can alter and define the tiebreak rules according to the requirement. In this way, the user can re-prioritize the sequence and/or remove some of the tie break indexes which are not needed.

Note The Lowest-backup-metric policy is not applicable for TI-LFA since TI-LFA always chooses the lowest back up path only.

You can see the above rules by using the following command:

```
R2#show ip ospf fast-reroute

            OSPF Router with ID (10.2.2.200) (Process ID 10)

Microloop avoidance is enabled for protected prefixes, delay 5000 msec

Loop-free Fast Reroute protected prefixes:

           Area        Topology name   Priority   Remote LFA Enabled     TI-LFA Enabled
              0                Base        Low                   No                Yes
    AS external                Base        Low                   No                Yes

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

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

Last SPF calculation started 17:25:51 ago and was running for 3 ms.
```

With the introduction of TI-LFA, the following two tie-break rules are enhanced.

- node-protection

- srlg-protection

The above two tie-break rules are not enabled by default. The user needs to configure the above mentioned tie-break policies.

Node Protection

TI-LFA node protection provides protection from node failures. Node protecting TI-LFA attempts to calculate the post conversion repair path that protects against the failure of a particular next-hop, not just the link to that particular next-hop.

Node protection is used as a tiebreaker in the implementation of the local LFA also. But when it is combined with TI-LFA, the back up path calculated post convergences with node protecting path. Per-Prefix TI-LFA node protection is disabled by default. The IPFRR TI-LFA node protection features is enabled when the corresponding tiebreak is enabled along with TI-LFA feature, that is,

```
router ospf 10
   [no] fast-reroute per-prefix ti-lfa [area <area> [disable]]
   [no] fast-reroute per-prefix tie-break node-protecting index <index>
   [no] fast-reroute per-prefix tie-break node-protecting required index <index>
```

When you enable node protection, all the other tie break rules also need to manually configured. The node protection is built over the link protection.

The difference between **node-protecting** and **node-protecting required** is in selecting the backup path. When you configure **node-protecting required**, then back up which is chosen has to be the path which does not go through the node (which is part of the link which we are protecting). If no such path is available, then no path is chosen as the backup path.

Shared Risk Link Groups Protection

A shared risk link group (SRLG) is a group of next-hop interfaces of repair and protected primary paths that have a high likelihood of failing simultaneously. The OSPFv2 Loop-Free Alternate Fast Reroute feature supports only SRLGs that are locally configured on the computing router. With the introduction of TI LFA, the post convergence path which does not share the SRLG group id with the primary path interface will be chosen. In that way, the user will be sure of the SRLG protection whenever the primary link fails.

The IPFRR TI-LFA SRLG protection features is enabled when the corresponding tiebreak is enabled along with Ti-LFA feature, that is,

```
router ospf 10
   [no] fast-reroute per-prefix ti-lfa [area <area> [disable]]
   [no] fast-reroute per-prefix tie-break srlg index <index>
   [no] fast-reroute per-prefix tie-break srlg required index <index>
```

When you enable SRLG protection, you need to manually configure all the other tie break rules. The difference between **srlg-protecting** and **srlg-protecting required** is in selecting the backup path. When you configure **srlg-protecting required**, then back up which is chosen has to be the path which does not share SRLG ID with the primary link which is protected. If no such path is available, then no path is chosen as the backup path.

Whereas, if you configure **srlg-protecting** alone then if the SRLG protection path is not available, the link protection path is chosen as the backup path. And when the SRLG protection path is available, the switchover happens to the SRLG protection path.

Node-Shared Risk Link Groups Protection

You can configure both node and SRLG protection tie breaks together. This means that the back up path needs to fulfil both the criteria of node protection as well as SRLG protection. 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 shortest path tree (SPT).

To enable node and SRLG protection tie breaks together, use the following command:

```
router ospf 10
   [no] fast-reroute per-prefix ti-lfa [area <area> [disable]]
   [no] fast-reroute per-prefix tie-break node-protecting index <index>
   [no] fast-reroute per-prefix tie-break srlg index <index>
```

The following show command is used to display the tie break policy:

```
R3#show ip ospf fast-reroute

            OSPF Router with ID (10.3.3.33) (Process ID 10)

Loop-free Fast Reroute protected prefixes:

           Area        Topology name   Priority   Remote LFA Enabled     TI-LFA Enabled
              0                Base        Low                   No                 No
              1                Base        Low                   No                 No
           1000                Base        Low                   No                 No
    AS external                Base        Low                   No                 No

  Repair path selection policy tiebreaks:
     0  post-convergence
    60  node-protecting
    70  srlg
   256  load-sharing

OSPF/RIB notifications:
 Topology Base: Notification Disabled, Callback Not Registered

Last SPF calculation started 00:00:06 ago and was running for 2 ms.
```

How to Configure Topology Independent Loop Free Alternate Fast Reroute

Enabling Topology Independent Loop Free Alternate Fast Reroute

By default, TI-LFA is disabled. You can use protocol enablement to enable TI-LFA.

Protocol enablement: Enables TI-LFA in router OSPF mode for all the OSPF areas. Perform the following steps to enable TI-LFA FRR.

```
[no] fast-reroute per-prefix ti-lfa [ area <area> disable]
```

```
  router ospf <process>
  fast-reroute per-prefix enable area <area> prefix-priority {low | high}
  fast-reroute per-prefix ti-lfa [ area <area> disable]
```

You can also use interface command to enable or disable IP FRR on specific interfaces.

```
interface <interface>
ip ospf fast-reroute per-prefix protection disable
ip ospf fast-reroute per-prefix candidate disable
ip ospf fast-reroute per-prefix protection ti-lfa [disable]
```



Note

- When TI-LFA is configured on the OSPF router and area wide, area specific configuration takes precedence.
- To protect external prefixes, TI-LFA should be enabled globally.

Configuring Topology Independent Loop Free Alternate Fast Reroute

This task describes how to enable per-prefix Topology Independent Loop-Free Alternate (TI-LFA) computation to converge traffic flows around link, node, and SRLG failures. TI-LFA can be configured on instance or area level inherited by lower levels. You can enable or disable per prefix FRR per interface level which is applicable for TI-LFA also.

Before you begin to configure, ensure that the following topology requirements are met:

- Router interfaces are configured as per the topology.
- Routers are configured with OSPF.
- Segment routing is enabled globally as well as under OSPF level.

1. Enables OSPF routing for the specified routing process and enters in router configuration mode.

   ```
   Device(config)# router ospf 10
   ```

2. Enables FRR.

   ```
   Device(config-router)# fast-reroute per-prefix enable prefix-priority low
   ```

3. Enables TI-LFA.

   ```
   Device(config-router)# fast-reroute per-prefix ti-lfa
   ```

4. Enables TI-LFA on the specific area.

   ```
   Device(config-router)# fast-reroute per-prefix ti-lfa area 0
   ```

5. Exits the TI-LFA mode.

   ```
   Device(config-router)# exit
   ```

6. Enters the interface mode.

   ```
   Device(config)#interface ethernet 0/0
   ```

7. If you do not wish to enable FRR on a specific inteface, use the protection disable command.

   ```
   Device(config-if)#ip ospf fast-reroute per-prefix protection disable
   ```

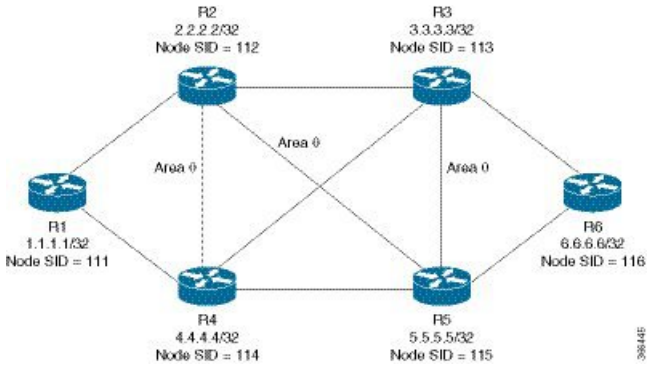
8. If you do not wish a specific interface to be enabled as a repair path, use the candidate disable command.

   ```
   Device(config-if)#ip ospf fast-reroute per-prefix candidate disable
   ```

Configuring Topology Independent Fast Reroute Tie-breaker

You need to enable segment routing on all the routers with prefix SIDs configured for all the nodes. Use the following topology as a reference to understand the configuration.

Figure 5: Configuration Example



Let us take the device R2 which is protecting the link between R2 and R3. The configuration at R2:

```
router ospf 10
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
segment-routing mpls
segment-routing area 0 mpls
fast-reroute per-prefix enable prefix-priority low
fast-reroute per-prefix ti-lfa
fast-reroute per-prefix ti-lfa area 0
fast-reroute per-prefix tie-break node-protecting index 60
fast-reroute per-prefix tie-break srlg index 70
mpls traffic-eng router-id Loopback1
mpls traffic-eng area 0

interface GigabitEthernet4   //interface connecting to the router 4
ip address 10.101.4.4 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
srlg gid 10
negotiation auto

interface GigabitEthernet3  //interface connecting to the router 3
ip address 10.101.3.3 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
srlg gid 10
negotiation auto

interface GigabitEthernet5   //interface connecting to the router 2
ip address 10.101.5.5 255.255.255.0
ip ospf 10 area 0
ip ospf network point-to-point
srlg gid 20
negotiation auto



interface loopback2
```

```
ip address 10.2.2.2/32
ip ospf 10 area 0
```



Note In all the other devices, configuration of segment routing and assignment of connected prefix SIDs need to be done.

How Node Protection Works: Using the same topology as an example, let us take the case where you are protecting the link between R2 and R3 and also the prefix which is leant from R6. In that case, let us assume that the primary path for the prefix is via R2-R3. So, our primary path is R2---R3---R6 and we are protecting the link R2---R3.

In this scenario, only link-protection is configured and enabled. When you enable TI-LFA under OSPF process, then you get the following paths provided the cost for all the paths are equal:

R2----R4----R5---R6

R2---R5----R3---R6

R2----R5---R6

If you have only link protection configured, then all the three paths will be chosen and they will share the load amongst them.

If you wish to configure node protection, then the backup would be calculated in such a way that the back up path does not contain the node that you are protecting. In this example, the node R3 in the back up is not required. As a result, only the following two paths would be chosen as the back up paths:

R2----R4----R5---R6

R2----R5---R6

It is possible that R2---R5---R3---R6 have the lesser cost than the above two paths. But since the node protection is configured, only the paths amongst the above two will be considered.

How SRLG Protection Works: SRLG protection further eliminates the back up paths in a such a way that the primary path and the backup does not share the same SRLG ID. Suppose the following back up paths are available:

R2----R4----R5---R6

R2----R5---R6

Then, the SRLG ID of (R2----R4) and (R2----R5) are compared against the primary interface (R2----R3) which is 10. It is noticed that only the interface R2----R5 has different SRLG ID which is 20. So, only the backup path R2---R5---R6 will be chosen.

Verifying Topology Independent Fast Reroute Tunnels

You can use the following command, to check the TI LFA tunnels:

```
Device#show ip ospf fast-reroute ti-lfa tunnels


            OSPF Router with ID (10.2.2.200) (Process ID 10)

                        Area with ID (0)
```

```
                    Base Topology (MTID 0)

Tunnel                Interface         Next Hop         Mid/End Point    Label
--------------------------------------------------------------------------------
MPLS-SR-Tunnel2       Et1/1             10.7.0.7         10.1.1.1         16020
MPLS-SR-Tunnel6       Et0/3             10.8.0.0         10.3.3.3         16003
MPLS-SR-Tunnel7       Et1/1             10.7.0.7         10.1.1.1         16020
                                                         10.5.5.5         16005
                                                         10.3.3.3         16003
MPLS-SR-Tunnel5       Et0/3             10.8.0.0         10.5.5.5         16005
MPLS-SR-Tunnel1       Et1/1             10.7.0.7         10.1.1.1         16020
                                                         10.5.5.5         16005
MPLS-SR-Tunnel3       Et1/1             10.7.0.7         10.6.6.6         16006
```

You can use the following command, to check the route in OSPF routing table with primary and repair path:

```
Device#show ip ospf rib 10.6.6.6

            OSPF Router with ID (10.2.2.200) (Process ID 10)

                Base Topology (MTID 0)

OSPF local RIB
Codes: * - Best, > - Installed in global RIB
LSA: type/LSID/originator

*>  10.6.6.6/32, Intra, cost 31, area 0
     SPF Instance 19, age 02:12:11
      contributing LSA: 10/10.0.0.0/10.6.6.6 (area 0)
     SID: 6
     CSTR Local label: 0
     Properties: Sid, LblRegd, SidIndex, N-Flag, TeAnn
     Flags: RIB, HiPrio
      via 10.7.0.7, Ethernet1/1 label 16006
       Flags: RIB
       LSA: 1/10.6.6.6/10.6.6.6
      PostConvrg repair path via 10.3.3.3, MPLS-SR-Tunnel6 label 16006, cost 81, Lbl cnt 1
       Flags: RIB, Repair, PostConvrg, IntfDj, LC Dj
       LSA: 1/10.6.6.6/10.6.6.6
```

You can use the following command, to display the route in the IP routing table:

```
Device#show ip route 10.6.6.6
Routing entry for 10.6.6.6/32
  Known via "ospf 10", distance 110, metric 31, type intra area
  Last update from 10.7.0.7 on Ethernet1/1, 00:25:14 ago
 SR Incoming Label: 16006
  Routing Descriptor Blocks:
  * 10.7.0.7, from 10.6.6.6, 00:25:14 ago, via Ethernet1/1, merge-labels
      Route metric is 31, traffic share count is 1
      MPLS label: 16006
      MPLS Flags: NSF
      Repair Path: 10.3.3.3, via MPLS-SR-Tunnel6
```

Debugging Topology Independent Loop Free Alternate Fast Reroute

You can use the following commands to debug TI-LFA FRR:

```
debug ip ospf fast-reroute spf
debug ip ospf fast-reroute spf detail
debug ip ospf fast-reroute rib
debug ip ospf fast-reroute rib [<access-list>]
```

Examples: OSPFv2 Link-Protection Topology Independent Loop Free Alternate Fast Reroute

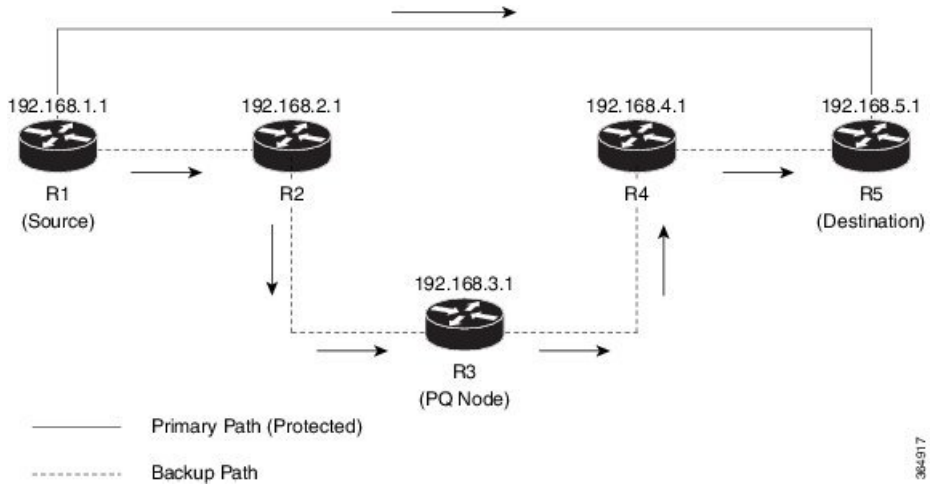
The following are the examples for the OSPFv2 Link-Protection TI-LFA FRR.

Example: Configuring Topology Independent Loop Free Alternate Fast Reroute

This example shows how to configure TI-LFA for segment routing TE tunnels using single or disjoint PQ nodes. The following are the two topologies used:

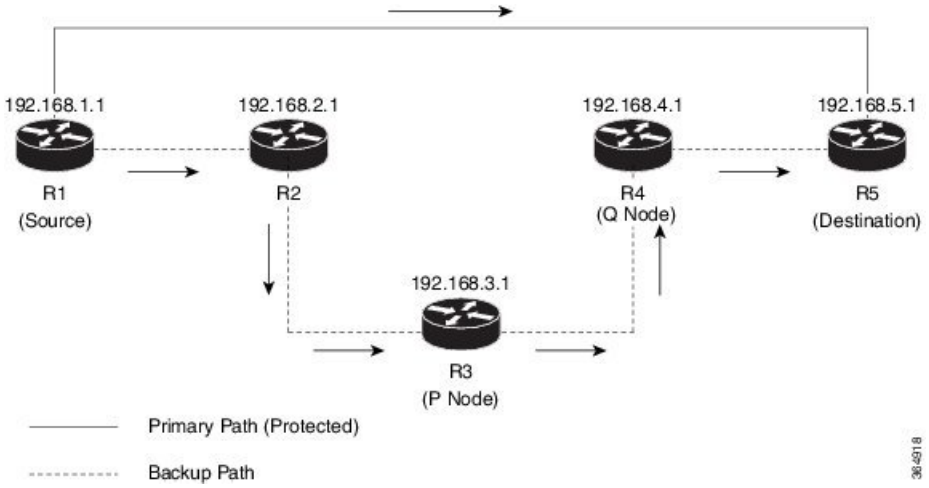
- Topology 1: A single PQ Node and therefore has two SIDs from the source router, R1 through the PQ Node to the destination router, R5.

Figure 6: Topology 1: Single PQ Node



- Topology 2: Disjoint PQ Nodes and therefore consists of three SIDs from the source router R1, through the P Node and the Q Node to the destination router, R5.

Figure 7: Topology 2: Disjoint PQ Nodes



Configure TI-LFA for OSPF on the source router (R1) interface connecting to the destination router (R5).

```
Device(config)# router ospf 10
Device(config-router)# fast-reroute per-prefix enable prefix-priority low
Device(config-router)# fast-reroute per-prefix ti-lfa
Device(config-router)# fast-reroute per-prefix ti-lfa area 0
Device(config-router)# exit
```