

Configure Topology Independent - Loop Free Alternate

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[IP-Fast Reroute](#)

[TI-LFA](#)

[PQ node](#)

[Configure](#)

[ISIS](#)

[OSPF](#)

[Verify](#)

[Repair Node is a Direct Neighbor \(LFA\)](#)

[Repair Node is a PQ Node \(rLFA\)](#)

[Repair is a Q node, Neighbor of the Last P Node \(TI-LFA\)](#)

[Troubleshoot](#)

[Related Information](#)

Introduction

This document describes how to configure Topology Independent - Loop Free Alternate (TI-LFA) in Segment Routing. The focus will be on TI-LFA configuration and verification on Cisco IOS® XR.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Knowledge of Segment Routing and configuration
- Cisco IOS XR

Components Used

The information in this document is based on Cisco IOS XR.

The information in this document was created from the devices in a specific lab environment. All of

the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

IP-Fast Reroute

- Classic LFA cannot protect all destinations in most of the current networks. A limitation of LFA is that, even if one or more LFAs exist, the optimal LFA might not always be provided.
- Remote LFA (rLFA) extends the coverage to 90-95% of the destinations, but it also does not always provide the most desired repair path. RLFA also adds more operational complexity when you require a targeted LDP session to the RLFAs to protect LDP traffic.

TI-LFA provides a solution to these limitations while maintaining the simplicity of the IPFRR solutions.

TI-LFA

TI-LFA in Segment Routing addresses challenges in Fast Reroute (FRR) which were earlier not possible with classic LFA path calculations.

TI-LFA is a segment routing based LFA FRR, which brings:

- Simplicity
- Entirely automated
- No directed LDP session
- No RSVP-TE tunnels
- Incremental deployment
- An optimal backup path along the post-convergence path
- Prevents transient congestion and suboptimal routing

When FRR is not enabled, there is no backup route installed in the RIB.

```
RP/0/0/CPU0:R1#show route 192.0.2.6
Routing entry for 192.0.2.6/32
Known via "isis 111", distance 115, metric 30, labeled SR, type level-2
Routing Descriptor Blocks
198.51.100.2, from 192.0.2.6, via GigabitEthernet0/0/0/0
Route metric is 30
No advertising protos.
```

PQ node

rLFA and TI-LFA use terminology such as P space, Q space, or PQ node while computing repair paths. Further details can be found in the verify section,

P Space: Set of routers that R1 can reach without traversing over the failed path in the pre-convergence state as per Shortest Path Tree (SPT) algorithm.

Q Space: Set of routers that R6 can reach without traversing over the failed path in the pre-convergence state as per SPT algorithm.

PQ Space: Intersection of R1's P space and R6's Q space.

Configure

A simple CLI in order to enable TI-LFA under IGP (Intermediate System-to-Intermediate System (ISIS), Open Shortest Path First (OSPF)) interfaces configuration mode as shown here.

ISIS

```
router isis 111
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
```

OSPF

```
router ospf 111
area 0
interface GigabitEthernet0/0/0/0
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa enable
```

Verify

Use this section to confirm that your configuration works properly.

Repair Node is a Direct Neighbor (LFA)


R1 (192.0.2.1) usually calculates its least-cost path towards R6 (192.0.2.6) and installs in the RIB. Traffic will be forwarded from R1 to R6 via R1 --R2 -- R3 -- R6 (Primary Path).

Without LFA, if there is a link failure between R1 --x-- R2, traffic between **R1 > R6** will be discarded for a few milliseconds until R1 recalculates and finds another route via R4.

With LFA enabled, R1 pre-installs a route to R6 via R4 as a backup.

This criterion should satisfy for backup route installation in the case of LFA,

- R4 least-cost path to R6 should not be via R1
- total cost from R4 to R6 should be lower than the current primary path i.e. cost from **R1 > R6** as shown in the image.



```
RP/0/0/CPU0:R1#show route 192.0.2.6/32
```

```
Routing entry for 192.0.2.6/32
```

```
Known via "isis 111", distance 115, metric 30, labeled SR, type level-2
```

```
Routing Descriptor Blocks
```

```
198.51.100.2, from 192.0.2.6, via GigabitEthernet0/0/0/0, Protected !Primary Path
```

```
Route metric is 30
```

```
198.51.100.21, from 192.0.2.6, via GigabitEthernet0/0/0/1, Backup (Local-LFA) !Backup Path
```

```
Route metric is 32
```

```
No advertising protos.
```

```
RP/0/0/CPU0:R1#show isis fast-reroute 192.0.2.6/32 detail
```

```
L2 192.0.2.6/32 [30/115] medium priority
```

```
via 198.51.100.2, GigabitEthernet0/0/0/0, R2, SRGB Base: 48000, Weight: 0
```

```
FRR backup via 198.51.100.21, GigabitEthernet0/0/0/1, R4, SRGB Base: 48000, Weight: 0, Metric: 32
```

```
P: No, TM: 32, LC: No, NP: Yes, Yes, SRLG: Yes
```

```
src R6.00-00, 192.0.2.6, prefix-SID index 6, R:0 N:1 P:0 E:0 V:0 L:0
```

```
RP/0/0/CPU0:R1#show cef 192.0.2.6/32 detail
```

```
192.0.2.6/32, version 1056, labeled SR, internal 0x1000001 0x81 (ptr 0xa12dbd34) [1], 0x0 (0xa12c12fc), 0xa28 (0xa170e1dc)
```

```
local adjacency 198.51.100.2
```

```
Prefix Len 32, traffic index 0, precedence n/a, priority 1
```

```
via 198.51.100.2/32, GigabitEthernet0/0/0/0, 11 dependencies, weight 0, class 0, protected
[flags 0x400]
path-idx 0 bkup-idx 1 NHID 0x0 [0xa175c4b8 0x0]
next hop 198.51.100.2/32
local label 48006 labels imposed {48006}
```

```
via 198.51.100.21/32, GigabitEthernet0/0/0/1, 11 dependencies, weight 0, class 0, backup (Local-
LFA) [flags 0x300]
path-idx 1 NHID 0x0 [0xa166e338 0x0]
next hop 198.51.100.21/32
local adjacency
local label 48006 labels imposed {48006}
Load distribution: 0 (refcount 2)
Hash OK Interface Address
0 Y GigabitEthernet0/0/0/0 198.51.100.2
```

Repair Node is a PQ Node (rLFA)


In most cases, LFA repair paths are not available (as they are not able to fulfil the criterion previously mentioned).

rLFA addresses the limitations with LFA, by selecting a node to act as a repair node even if they are not direct neighbors to R1. rLFA works when you identify a PQ node.

R1's P-space includes all nodes that R1 can reach without traversing **R1 > R2** in the pre-convergence state (the state is achieved when all the prefixes are installed and there are no further updates left).

R6's Q-space includes all nodes that R6 can reach without traversing **R1 > R2** in a pre-convergence state.

As R5 is the only node that is part of P and Q-space, R5 is elected as a PQ node and will act as a repair node for the link **R1 > R2** as shown in the image.



Note: Metric has been changed from 12 to 30 (R4 -- R5) in order to demonstrate rLFA behaviour.

```
RP/0/0/CPU0:R1#show route 192.0.2.6/32
Routing entry for 192.0.2.6/32
Known via "isis 111", distance 115, metric 30, labeled SR, type level-2
Routing Descriptor Blocks
198.51.100.2, from 192.0.2.6, via GigabitEthernet0/0/0/0, Protected !Primary path
Route metric is 30
198.51.100.21, from 192.0.2.6, via GigabitEthernet0/0/0/1, Backup (TI-LFA) !Backup path
Repair Node(s): 192.0.2.5
Route metric is 50
No advertising protos.
```

```
RP/0/0/CPU0:R1#show isis fast-reroute 192.0.2.6/32 detail
L2 192.0.2.6/32 [30/115] medium priority
via 198.51.100.2, GigabitEthernet0/0/0/0, R2, SRGB Base: 48000, Weight: 0
Backup path: TI-LFA (link), via 198.51.100.21, GigabitEthernet0/0/0/1 R4, SRGB Base: 48000,
Weight: 0
P node: R5.00 [192.0.2.5], Label: 48005
Prefix label: 48006
Backup-src: R6.00
```

P: No, TM: 50, LC: No, NP: No, No, SRLG: Yes
src R6.00-00, 192.0.2.6, prefix-SID index 6, R:0 N:1 P:0 E:0 V:0 L:0

RP/0/0/CPU0:R1#show cef 192.0.2.6/32 detail

192.0.2.6/32, version 1166, labeled SR, internal 0x1000001 0x81 (ptr 0xa12dc41c) [1], 0x0
(0xa12c19e0), 0xa28 (0xa170e1b0)
local adjacency 198.51.100.2

via 198.51.100.2/32, GigabitEthernet0/0/0/0, 10 dependencies, weight 0, class 0, protected
[flags 0x400]

path-idx 0 bkup-idx 1 NHID 0x0 [0xa175c4b8 0x0]

next hop 198.51.100.2/32

local label 48006 labels imposed {48006}

via 198.51.100.21/32, GigabitEthernet0/0/0/1, 10 dependencies, weight 0, class 0, backup (TI-
LFA) [flags 0xb00]

path-idx 1 NHID 0x0 [0xa166e338 0x0]

next hop 198.51.100.21/32, Repair Node(s): 192.0.2.5

local adjacency

local label 48006 labels imposed {48005 48006}

Load distribution: 0 (refcount 3)


Hash OK Interface Address

0 Y GigabitEthernet0/0/0/0 198.51.100.2

Repair is a Q node, Neighbor of the Last P Node (TI-LFA)

rLFA will not be able to provide protection if the metric between **R4 > R5** is raised to 50. R1's P-space for Link **R1 > R2** includes only R3. R6's Q-space for Link **R1 > R2** includes R3, R4, and R5. No repair nodes are available.

So, even rLFA does not guarantee a backup path in all network topologies. TI-LFA address the limitations of rLFA and provides a backup path for the link **R1 > R2** in such scenarios as shown in the image.



Note: Metric has been changed from 30 to 50 (R4 -- R5) in order to demonstrate Segment Routing TI-LFA behaviour (which was not possible with rLFA).

```
RP/0/0/CPU0:R1#show route 192.0.2.6/32
Routing entry for 192.0.2.6/32
Known via "isis 111", distance 115, metric 30, labeled SR, type level-2
Routing Descriptor Blocks
198.51.100.2, from 192.0.2.6, via GigabitEthernet0/0/0/0, Protected !Primary Path
Route metric is 30
198.51.100.21, from 192.0.2.6, via GigabitEthernet0/0/0/1, Backup (TI-LFA) !Backup Path
Repair Node(s): 192.0.2.4, 192.0.2.5
Route metric is 70
No advertising protos.
```

```
RP/0/0/CPU0:R1#show isis fast-reroute 192.0.2.6/32 detail
L2 192.0.2.6/32 [30/115] medium priority
via 198.51.100.2, GigabitEthernet0/0/0/0, R2, SRGB Base: 48000, Weight: 0
Backup path: TI-LFA (link), via 198.51.100.21, GigabitEthernet0/0/0/1 R4, SRGB Base: 48000,
Weight: 0
P node: R4.00 [192.0.2.4], Label: ImpNull
Q node: R5.00 [192.0.2.5], Label: 24003
Prefix label: 48006
```



```
Backup-src: R6.00
P: No, TM: 70, LC: No, NP: No, No, SRLG: Yes
src R6.00-00, 192.0.2.6, prefix-SID index 6, R:0 N:1 P:0 E:0 V:0 L:0
```

```
RP/0/0/CPU0:R1#show cef 192.0.2.6/32 detail
```

```
192.0.2.6/32, version 1192, labeled SR, internal 0x1000001 0x81 (ptr 0xa12dc41c) [1], 0x0
(0xa12c165c), 0xa28 (0xa170e310)
local adjacency 198.51.100.2
```

```
via 198.51.100.2/32, GigabitEthernet0/0/0/0, 12 dependencies, weight 0, class 0, protected
[flags 0x400]
```

```
path-idx 0 bkup-idx 1 NHID 0x0 [0xa175c170 0xa175c4b8]
```

```
next hop 198.51.100.2/32
```

```
local label 48006 labels imposed {48006}
```

```
via 198.51.100.21/32, GigabitEthernet0/0/0/1, 12 dependencies, weight 0, class 0, backup (TI-
LFA) [flags 0xb00]
```

```
path-idx 1 NHID 0x0 [0xa166e16c 0xa166e338]
```

```
next hop 198.51.100.21/32, Repair Node(s): 192.0.2.4, 192.0.2.5
```

```
local adjacency
```

```
local label 48006 labels imposed {ImplNull 24003 48006}
```

```
Load distribution: 0 (refcount 7)
```

```
Hash OK Interface Address
```

```
0 Y GigabitEthernet0/0/0/0 198.51.100.2
```

```
RP/0/0/CPU0:R1#
```

Troubleshoot

This section provides information you can use in order to troubleshoot your configuration.

Refer to; [Troubleshooting Segment Routing - Cisco Systems](#)

Related Information

- [Introduction to Segment Routing - Cisco Systems](#)
- [Segment Routing Technology deep-dive and advanced use cases - Cisco Systems](#)
- [Technical Support & Documentation - Cisco Systems](#)