

IPv6 Loop-Free Alternate Fast Reroute

When a link or a router fails, distributed routing algorithms compute new routes that take into account the failure. The time taken for this computation is called routing transition. Until the transition is complete and all routers are converged on a common view of the network, the connectivity between the source and destination pairs is interrupted. You can use the IPv6 Loop-Free Alternate (LFA) Fast Reroute (FRR) feature to reduce the routing transition time to less than 50 milliseconds using a precomputed alternate next hop. When a router is notified of a link failure, the router immediately switches over to the repair path to reduce traffic loss.

IPv6 LFA FRR supports the precomputation of repair paths. The repair path computation is done by the Intermediate System-to-Intermediate System (IS-IS) routing protocol, and the resulting repair paths are sent to the IPv6 Routing Information Base (RIB). The repair path installation is done by Cisco Express Forwarding (formerly known as CEF).

- Prerequisites for IPv6 LFA FRR, on page 1
- Restrictions for IPv6 LFA FRR, on page 1
- Information About IPv6 LFA FRR, on page 2
- How to Configure IPv6 LFA FRR, on page 4
- Configuration Examples for IPv6 LFA FRR, on page 7
- Verifying IPv6 LFA FRR Configuration, on page 8
- Feature Information for Configuring IPv6 LFA FRR, on page 8

Prerequisites for IPv6 LFA FRR

There are no specific prerequisites for configuring IPv6 LFA FRR.

Restrictions for IPv6 LFA FRR

- Loop-Free Alternate (LFA) Fast Reroute (FRR) can protect paths that are reachable through an interface only if the interface is a point-to-point interface.
- Any type of tunnel interfaces cannot be used as a protected interface. However, tunnel can be a protecting (repair) tunnel.
- Loadbalance support is available for FRR-protected prefixes on per-prefix basis. If there are multiple equal backup paths, only one can be assigned to a prefix. Assignment is done based on hash function applied to IPv6 prefix. Different IPv6 prefixes have different result of hash function and therefore different backup paths are used.

- A maximum of eight FRR-protected interfaces can simultaneously undergo a cutover.
- Only Layer 3 VPN is supported.
- IPv6 multicast is not supported.
- Only physical and physical port-channel interfaces and subinterfaces are protected. Tunnels and virtual interfaces are not protected.
- The capability of LFA to find a backup path is limited by simplicity of the algorithm. The algorithm can find a backup path only if there is a direct IS-IS neighbor (other than primary one) which has primary path to a prefix, and that primary path does not point to the calculating router. If the network topology is such that LFA cannot cover significant percentage of primary paths with backup paths, it is recommended to use ISIS SRv6 TI-LFA algorithm to get good FRR coverage. For example, LFA algorithm provides good coverage in spine-leaf types of topologies, but not in ring topologies.

Information About IPv6 LFA FRR

IS-IS and IPv6 FRR

When a local link fails in a network, IS-IS recomputes new primary next-hop paths for all affected prefixes. These prefixes are updated in the RIB and the Forwarding Information Base (FIB). Until the primary path prefixes are updated in the forwarding plane, traffic directed towards the affected prefixes are discarded. This process can take hundreds of milliseconds.

In IPv6 FRR, IS-IS computes LFA next-hop routes for the forwarding plane to use in case of primary path failures. LFA is computed per prefix.

When there are multiple LFAs for a given primary path, IS-IS uses a tiebreaking rule to pick a single LFA for a primary path. In case of a primary path with multiple LFA paths, prefixes are distributed equally among LFA paths.

Repair Paths

Repair paths forward traffic during a routing transition. When a link or a router fails, due to the loss of a physical layer signal, initially, only the neighboring routers are aware of the failure. All other routers in the networkare unaware of the nature and location of this failure until information about this failure is propagated through a routing protocol, which may take several hundred milliseconds. It is, therefore, necessary to arrange for packets affected by the network failure to be steered to their destinations.

A router adjacent to the failed link employs a set of repair paths for packets that would have used the failed link. These repair paths are used from the time the router detects the failure until the routing transition is complete. By the time the routing transition is complete, all routers in the network revise their forwarding data and the failed link is eliminated from the routing computation.

Repair paths are precomputed in anticipation of failures so that they can be activated the moment a failure is detected.

The IPv6 LFA FRR feature uses the following repair paths:

- Equal Cost Multipath (ECMP) uses a link as a member of an equal cost path-split set for a destination. The other members of the set can be used as a repair path when the link fails.

- LFA is a next-hop that delivers a packet to its destination without looping back. Downstream paths are a subset of LFAs.

LFA Overview

LFA is a node other than the primary neighbor. Traffic is redirected to an LFA after a network failure. An LFA makes the forwarding decision without any knowledge of the failure.

An LFA must neither use a failed element nor use a protecting node to forward traffic. An LFA must not cause loops. By default, LFA is enabled on all supported interfaces as long as the interface can be used as a primary path.

Advantages of using per-prefix LFAs are as follows:

- The repair path forwards traffic during transition when the primary path link is down.
- All destinations having a per-prefix LFA are protected. This leaves only a subset (a node at the far side of the failure) unprotected.

LFA Calculation

The general algorithms to compute per-prefix LFAs can be found in RFC 5286. IS-IS implements RFC 5286 with a small change to reduce memory usage. Instead of performing a Shortest Path First (SPF) calculation for all neighbors before examining prefixes for protection, IS-IS examines prefixes after SPF calculation is performed for each neighbor. Because IS-IS examines prefixes after SPF calculation is performed, IS-IS retains the best repair path after SPF calculation is performed for each neighbor. IS-IS does not have to save SPF results for all neighbors.

Interaction Between RIB and Routing Protocols

A routing protocol computes repair paths for prefixes by implementing tiebreaking algorithms. The end result of the computation is a set of prefixes with primary paths, where some primary paths are associated with repair paths.

A tiebreaking algorithm considers LFAs that satisfy certain conditions or have certain attributes. When there is more than one LFA, configure the **fast-reroute per-prefix** command with the **tie-break** keyword. If a rule eliminates all candidate LFAs, then the rule is skipped.

A primary path can have multiple LFAs. A routing protocol is required to implement default tiebreaking rules and to allow you to modify these rules. The objective of the tiebreaking algorithm is to eliminate multiple candidate LFAs, select one LFA per primary path per prefix, and distribute the traffic over multiple candidate LFAs when the primary path fails.

Tiebreaking rules cannot eliminate all candidates.

The following attributes are used for tiebreaking:

- Downstream—Eliminates candidates whose metric to the protected destination is lower than the metric of the protecting node to the destination.
- Linecard-disjoint—Eliminates candidates sharing the same linecard with the protected path.
- Shared Risk Link Group (SRLG)—Eliminates candidates that belong to one of the protected path SRLGs.

- Load-sharing—Distributes remaining candidates among prefixes sharing the protected path.
- Lowest-repair-path-metric—Eliminates candidates whose metric to the protected prefix is higher.
- Node protecting—Eliminates candidates that are not node protected.
- Primary-path—Eliminates candidates that are not ECMPs.
- Secondary-path—Eliminates candidates that are ECMPs.

How to Configure IPv6 LFA FRR

Configuring FRR Support



Note LFA computations are enabled for all routes, and FRR is enabled on all supported interfaces.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ipv6 enable**
5. **ipv6 router isis** *area-tag*
6. **isis network point-to-point**
7. **exit**
8. **router isis** *area-tag*
9. **net** *net*
10. **metric-style wide**
11. **address-family ipv6**
12. **multi-topology**
13. **fast-reroute per-prefix** {**level-1 | level-2**} {**all | route-map** *route-map name*}
14. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Device(config)# interface GigabitEthernet0/0/0`	Configures an interface and enters interface configuration mode.
Step 4	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 on the interface. You can also enable IPv6 by configuring an IPv6 address.
Step 5	**ipv6 router isis** *area-tag* **Example:** `Device(config-if)# ipv6 router isis ipfrr`	Configures an IS-IS routing process for an IPv6 on an interface and attaches an area designator to the routing process.
Step 6	**isis network point-to-point** **Example:** `Device(config-if)# isis network point-to-point`	Enforces IS-IS point-to-point network type.
Step 7	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 8	**router isis** *area-tag* **Example:** `Device(config)# router isis ipfrr`	Enables the IS-IS routing protocol, specifies an IS-IS process, and enters router configuration mode.
Step 9	**net** *net* **Example:** `Device(config-router)# net 49.0001.0101.2800.0001.00`	Configures an IS-IS network entity (NET) for a routing process.
Step 10	**metric-style wide** **Example:** `Device(config-router)# metric-style wide`	Enables metric-style wide. **Note** It is recommended to run wide metric o
Step 11	**address-family ipv6** **Example:** `Device(config-router)# address-family ipv6`	Enters IPv6 configuration sub-mode.

	Command or Action	Purpose
Step 12	**multi-topology** **Example:** `Device(config-router-af)# multi-topology`	(Optional) Allows IS-IS to run in multi-topology mode in compliance with RFC 5120. Multi-topology allows for non-concurrent IPv4 and IPv6 topologies. **Note** IS-IS supports IPv6 LFA also in single-top
Step 13	**fast-reroute per-prefix** {**level-1** | **level-2**} {**all** | **route-map** *route-map name*} **Example:** `Device(config-router-af)# fast-reroute per-prefix level-2 all`	Enables per-prefix FRR in LFA mode. **Note** Configure the **all** keyword to protect all pr
Step 14	**end** **Example:** `Device(config-router-af)# end`	Exits router configuration mode and enters privileged EXEC mode.

Additional IS-IS IPv6 Commands

From Cisco IOS XE 17.15.1a, you can use the following optional commands to further fine-tune LFA FRR configurations:

Router IS-IS / Address-family IPv6 Mode Commands

fast-reroute tie-break {level-1 | level-2}

Configures the following tie-breakers that impact backup path calculation and selection:

```
downstream                      Prefer repair path via downstream node
linecard-disjoint               Prefer line card disjoint repair path
lowest-backup-path-metric       Prefer repair path with lowest total metric
node-protecting                 Prefer node protecting repair path
primary-path                    Prefer repair path from ECMP set
secondary-path                  Prefer non-ECMP repair path
srlg-disjoint                   Prefer SRLG disjoint repair path
```

fast-reroute interface disable <level>

Disables FRR protection on all interfaces by default. Interfaces where FRR is required can be configured explicitly using the interface level command.

fast-reroute load-sharing <level> disable

Disables load sharing between equal backup paths.

fast-reroute use-candidate-only <level>

Use as candidate interface only these allowed by the interface configuration.

Interface IS-IS IPv6 FRR Commands

isis ipv6 fast-reroute candidate <level> {disable}

Configures the interface for fast-reroute backup path.

isis ipv6 fast-reroute exclude <level> <interface>

Excludes another interface from being used for fast-reroute backup.

isis ipv6 fast-reroute protection <level> {disable}

Enables or disables fast-reroute protection on an interface.

isis ipv6 fast-reroute tie-break <level>

Creates the following set of tie-breakers specific for the interface:

```
default                          Use default tiebreakers set
downstream                       Prefer repair path via downstream node
linecard-disjoint                Prefer line card disjoint repair path
lowest-backup-path-metric        Prefer repair path with lowest total metric
node-protecting                  Prefer node protecting repair path
primary-path                     Prefer repair path from ECMP set
secondary-path                   Prefer non-ECMP repair path
srlg-disjoint                    Prefer SRLG disjoint repair path
```

Configuration Examples for IPv6 LFA FRR

Example: Configuring IPv6 LFA FRR

The following example shows basic configuration of IPv6 LFA FRR on the router interface and under router ISIS. IPv6 LFA FRR is enabled in level 2 for all ISIS IPv6 prefixes present in level 2.

```
interface Ethernet0/0
 ip unnumbered Loopback0
 ipv6 enable
  ipv6 router isis 1
   isis network point-to-point
 !
router isis 1
 net 49.0000.2222.2222.2222.00
 is-type level-2-only
 router-id Loopback0
  metric-style wide

 address-family ipv6
  multi-topology
  router-id Loopback0
   fast-reroute per-prefix level-2 all
exit-address-family
```

In the following example, only routes with tag 17 are protected.

```
router isis
 net 47.0004.004d.0001.0001.c11.1111.00
 address-family ipv6
  fast-reroute per-prefix level-2 route-map ipfrr-include
exit
route-map ipfrr-include
  match tag 17
```

Verifying IPv6 LFA FRR Configuration

Use the following show commands to verify IPv6 FRR and LFA configuration:

show isis ipv6 fast-reroute interfaces

```
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: Not Enabled, TI-LFA: Not Enabled
 Level-2 MT-2: FRR: Enabled, TI-LFA: Not Enabled
  Ethernet1/3: Protectable: Yes. Usable for repair: Yes
  Ethernet1/1: Protectable: Yes. Usable for repair: Yes
  Ethernet1/0: Protectable: Yes. Usable for repair: Yes
```

show isis ipv6 fast-reroute summary

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

Tag 1:
IPv6 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        0          0            0%
   Normal priority:      12         3            25%
   Total:                12         3            25%
```

show isis ipv6 rib

```
router# show isis ipv6 rib 604::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

* 604::1/128  prefix attr X:0 R:0 N:1
    via FE80::A8BB:CCFF:FE02:5E20/Ethernet0/2, type L2  metric 40 tag 0
    prefix attr: X:0 R:0 N:1
     (installed)
     repair path: via FE80::A8BB:CCFF:FE02:5A00/Ethernet0/0 metric: 40 (PP,DS,NP,SR)
      local LFA
      repair source: r604, metric to pfx: 40
    via FE80::A8BB:CCFF:FE02:5A00/Ethernet0/0, type L2  metric 40 tag 0
    prefix attr: X:0 R:0 N:1
     (installed)
     repair path: via FE80::A8BB:CCFF:FE02:5E20/Ethernet0/2 metric: 40 (PP,DS,NP,SR)
      local LFA
      repair source: r604, metric to pfx: 40
```

Feature Information for Configuring IPv6 LFA FRR

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 Configuring IPv6 LFA FRR

Feature Name	Releases	Feature Information
IPv6 Loop-Free Alternate Fast Reroute	Cisco IOS XE Release 17.15.1a	This feature was introduced. The following commands are introduced or modified as part of this feature: **fast-reroute tie-break {level-1 | level-2}** **isis ipv6 fast-reroute candidate <level> {disable}**