



IPv4 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 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 IPv4 Loop-Free Alternate Fast Reroute 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.

IPv4 Loop-Free Alternate Fast Reroute 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 Routing Information Base (RIB). The repair path installation is done by Cisco Express Forwarding (formerly known as CEF) and Open Shortest Path First (OSPF).

- [Finding Feature Information, on page 1](#)
- [Prerequisites for IPv4 Loop-Free Alternate Fast Reroute, on page 1](#)
- [Restrictions for IPv4 Loop-Free Alternate Fast Reroute, on page 2](#)
- [Information About IPv4 Loop-Free Alternate Fast Reroute, on page 2](#)
- [How to Configure IPv4 Loop-Free Alternate Fast Reroute, on page 4](#)
- [Configuration Examples for IPv4 Loop-Free Alternate Fast Reroute, on page 7](#)
- [Feature Information for Configuring IPv4 Loop-Free Alternate Fast Reroute, on page 8](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for IPv4 Loop-Free Alternate Fast Reroute

- 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.

- When a LAN interface is physically connected to a single neighbor, you should configure the LAN interface as a point-to-point interface so that it can be protected through LFA FRR.

Restrictions for IPv4 Loop-Free Alternate Fast Reroute

- A Multiprotocol Label Switching (MPLS) traffic engineering (TE) tunnel cannot be used as a protected interface. However, an MPLS TE tunnel can be a protecting (repair) interface as long as the TE tunnel is used as a primary path.
- Loadbalance support is available for FRR-protected prefixes, but the 50 ms cutover time is not guaranteed.
- A maximum of eight FRR-protected interfaces can simultaneously undergo a cutover.
- Only Layer 3 VPN is supported.
- IPv4 multicast is not supported.
- IPv6 is not supported.
- IS-IS will not calculate LFA for prefixes whose primary interface is a tunnel.
- LFA calculations are restricted to interfaces or links belonging to the same level or area. Hence, excluding all neighbors on the same LAN when computing the backup LFA can result in repairs being unavailable in a subset of topologies.
- Only physical and physical port-channel interfaces are protected. Subinterfaces, tunnels, and virtual interfaces are not protected.
- A TE label switched path (LSP) can be used as a backup path. However, the primary path has to be a physical interface, which can be used to achieve FRR in ring topologies.
- Border Gateway Protocol (BGP) Prefix-Independent Convergence (PIC) and IP FRR can be configured on the same interface as long as they are not used for the same prefix.

Information About IPv4 Loop-Free Alternate Fast Reroute

IS-IS and IP FRR

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

In IP 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 network are 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 IPv4 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 provide an alternative path when the link fails.
- LFA is a next-hop route 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 IPv4 Loop-Free Alternate Fast Reroute

Configuring Fast Reroute 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. **ip address** *ip-address mask*
5. **ip router isis** *area-tag*
6. **isis tag** *tag-number*

7. **exit**
8. **interface** *type number*
9. **ip address** *ip-address mask*
10. **ip router isis** *area-tag*
11. **isis tag** *tag-number*
12. **exit**
13. **router isis** *area-tag*
14. **net** *net*
15. **fast-reroute per-prefix** {*level-1* | *level-2*} {*all* | **route-map** *route-map-name*}
16. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Device(config)# interface GigabitEthernet0/0/0	Configures an interface and enters interface configuration mode.
Step 4	ip address <i>ip-address mask</i> Example: Device(config-if)# ip address 10.1.1.1 255.255.255.0	Sets a primary or secondary IP address for an interface.
Step 5	ip router isis <i>area-tag</i> Example: Device(config-if)# ip router isis ipfrr	Configures an IS-IS routing process for an IP on an interface and attaches an area designator to the routing process.
Step 6	isis tag <i>tag-number</i> Example: Device(config-if)# isis tag 17	Sets a tag on the IP address configured for an interface when the IP prefix is added to an IS-IS link-state packet (LSP).

	Command or Action	Purpose
Step 7	exit Example: Device(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 8	interface <i>type number</i> Example: Device(config)# interface GigabitEthernet0/0/1	Configures an interface and enters interface configuration mode.
Step 9	ip address <i>ip-address mask</i> Example: Device(config-if)# ip address 192.168.255.2 255.255.255.0	Sets a primary or secondary IP address for an interface.
Step 10	ip router isis <i>area-tag</i> Example: Device(config-if)# ip router isis ipfrr	Configures an IS-IS routing process for an IP on an interface and attaches an area designator to the routing process.
Step 11	isis tag <i>tag-number</i> Example: Device(config-if)# isis tag 17	Sets a tag on the IP address configured for an interface when the IP prefix is added to an IS-IS LSP.
Step 12	exit Example: Device(config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 13	router isis <i>area-tag</i> Example: Device(config)# router isis ipfrr	Enables the IS-IS routing protocol, specifies an IS-IS process, and enters router configuration mode.
Step 14	net <i>net</i> Example: Device(config-router)# net 49.0001.0101.2800.0001.00	Configures an IS-IS network entity (NET) for a routing process.
Step 15	fast-reroute per-prefix { level-1 level-2 } { all route-map <i>route-map-name</i> } Example: Device(config-router)# fast-reroute per-prefix level-2 all	Enables per-prefix FRR. <ul style="list-style-type: none"> • Configure the all keyword to protect all prefixes.


	Command or Action	Purpose
Step 16	end Example: Device(config-router)# end	Exits router configuration mode and enters privileged EXEC mode.

Configuration Examples for IPv4 Loop-Free Alternate Fast Reroute

Example: Configuring IPv4 Loop-Free Alternate Fast Reroute Support

The figure below shows IPv4 LFA FRR protecting BGP next hops by using interface tags.

Figure 1: Sample IPv4 LFA FRR Configuration



The following example shows how to configure IPv4 LFA FRR on Router A as shown in the above figure. Router A will advertise prefixes 10.0.0.0/24 and 192.168.255.0/24 along with the tag 17.

```
Device# configure terminal
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.1.1 255.255.255.0
Device(config-if)# ip router isis ipfrr
Device(config-if)# isis tag 17
Device(config-if)# exit
Device(config)# interface GigabitEthernet0/0/1
Device(config-if)# ip address 192.168.255.2 255.255.255.0
Device(config-if)# ip router isis ipfrr
Device(config-if)# isis tag 17
Device(config-if)# exit
Device(config)# router isis ipfrr
Device(config-router)# net 49.0001.0001.0001.0001.00
Device(config-router)# fast-reroute per-prefix level-2
```

The following example shows how to configure IPv4 LFA FRR on other routers as shown in the above figure. Other routers can use tag 17 to calculate repair paths for the two prefixes configured in Router A.

```
Device(config)# router isis
Device(config-router)# net 47.0004.004d.0001.0001.c11.1111.00
Device(config-router)# fast-reroute per-prefix level-2 route-map ipfrr-include
Device(config-router)# exit
Device(config)# route-map ipfrr-include
Device(config-router)# match tag 17
```

Feature Information for Configuring IPv4 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 www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Configuring IPv4 Loop-Free Alternate Fast Reroute