



IS-IS Remote Loop-Free Alternate Fast Reroute

The Intermediate System-to-Intermediate System (IS-IS) remote loop-free alternate (LFA) fast reroute (FRR) uses a backup route, precomputed using a dynamic routing protocol, to avoid traffic loss whenever a network fails. The backup routes (repair paths) are precomputed and installed in the router as the backup for the primary paths. Once the router detects a link or adjacent node failure, it switches to the backup path to avoid traffic loss.

IS-IS remote LFA FRR allows the backup path to be more than one hop away. This feature is particularly useful in some topologies, such as the commonly used ring topology, where an LFA does not have to be directly connected to the protecting router.

- [Finding Feature Information, on page 1](#)
- [Prerequisites for IS-IS Remote Loop-Free Alternate Fast Reroute, on page 1](#)
- [Restrictions for IS-IS Remote Loop-Free Alternate Fast Reroute, on page 2](#)
- [Information About IS-IS Remote Loop-Free Alternate Fast Reroute, on page 2](#)
- [How to Configure IS-IS Remote Loop-Free Alternate Fast Reroute, on page 5](#)
- [Configuration Examples for IS-IS Remote Loop-Free Alternate Fast Reroute, on page 6](#)
- [Additional References, on page 6](#)
- [Feature Information for IS-IS Remote Loop-Free Alternate Fast Reroute, on page 7](#)

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 IS-IS Remote Loop-Free Alternate Fast Reroute

Before performing the tasks in this module, you should be familiar with the concepts described in the “Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute” module.

IS-IS supports LFA FRR only when Multiprotocol Label Switching (MPLS) traffic engineering (TE) is also supported.

Restrictions for IS-IS Remote Loop-Free Alternate Fast Reroute

- An MPLS TE tunnel interface is treated as a point-to-point interface and hence the 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.
- IPv4 multicast is not supported.
- IPv6 is not supported.
- IS-IS will not calculate a loop-free alternate (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 might result in repairs being unavailable in a subset of topologies.
- Only physical interfaces and physical port-channel interfaces are protected. Subinterfaces, tunnels, and virtual interfaces are not protected.

Information About IS-IS Remote 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 Routing Information Base (RIB) and the Forwarding Information Base (FIB). Until the primary prefixes are updated in the forwarding plane, traffic directed toward the affected prefixes is 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 or the failure of a Bidirectional Forwarding Detection (BFD) session, 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. Therefore, packets affected by the network failure need to be steered to their destination.

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.

When a protected element fails, a repair node carries traffic around it toward the destination. When the protecting node detects this failure, it directs traffic around the protected element towards the repair node. In general, a repair node may be directly connected to the protecting node.

The IS-IS remote 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.

Loop-Free Alternate

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.

Remote LFA FRR

Some topologies (for example the commonly used ring-based topology) require protection that is not afforded by LFA FRR alone. Consider the topology shown in the figure below:

Figure 1: Remote LFA FRR with Ring Topology



The red looping arrow represents traffic that is looping immediately after a failure between node A and C (before network reconvergence). Device A tries to send traffic destined to F to next-hop B. Device B cannot be used as an LFA for prefixes advertised by nodes C and F. The actual LFA is node D. However, node D is not directly connected to the protecting node A. To protect prefixes advertised by C, node A must tunnel the packet around the failed link A-C to node D, provided that the tunnel does not traverse the failing link.

Remote LFA FRR enables you to tunnel a packet around a failed link to a remote loop-free alternate that is more than one hop away. In the figure above, the green arrow between A and D shows the tunnel that is automatically created by the remote LFA feature to bypass looping.

Loop-Free Alternate 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 Sender Policy Framework (SPF) for all neighbors before examining prefixes for protection, IS-IS examines prefixes after SPF is performed for each neighbor. Because IS-IS examines prefixes after SPF is performed, IS-IS retains the best repair path after each neighbor SPF is performed. 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 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 omitted.

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 IS-IS Remote Loop-Free Alternate Fast Reroute

Configuring IS-IS Remote Loop-Free Alternate Fast Reroute Tunnel

Perform this task to configure an IS-IS LFA FRR path that redirects traffic to a remote LFA tunnel.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router isis** [*area-tag*]
4. **fast-reroute remote-lfa** {*level-1* | *level-2*} **mpls-ldp** [**maximum-metric** *metric-value*]
5. **end**

DETAILED STEPS

	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	router isis [<i>area-tag</i>] Example:	Enables the IS-IS routing protocol and specifies an IS-IS process. <ul style="list-style-type: none"> • Enters router configuration mode.

	Command or Action	Purpose
	Device(config)# router isis ipfrr	
Step 4	fast-reroute remote-lfa {level-1 level-2} mpls-ldp [maximum-metric metric-value] Example: Device(config-router)# fast-reroute remote-lfa level-1 mpls-ldp	Configures an FRR path that redirects traffic to a remote LFA tunnel for either level 1 or level 2 packets. <ul style="list-style-type: none"> Use the maximum-metric metric-value keyword-argument pair to specify the maximum metric value required to reach the release node.
Step 5	end Example: Device(config-router)# end	Exits router configuration mode and enters privileged EXEC mode.

Configuration Examples for IS-IS Remote Loop-Free Alternate Fast Reroute

Example: Configuring IS-IS Remote Loop-Free Alternate Fast Reroute

The following example shows how to enable LFA FRR for all level 2 packets:

```
Router(config)# router isis
Router(router-config)# fast-reroute remote-lfa level-2 mpls-ldp
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
IS-IS commands: complete command syntax, command mode, defaults, command history, usage guidelines, and examples	<i>Cisco IOS IP Routing: IS-IS Command Reference</i>
Overview of Cisco IS-IS conceptual information with links to all the individual IS-IS modules	“Integrated IS-IS Routing Protocol Overview”
Understanding IS-IS IPv4 LFA FRR	“Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute”

Standards and RFCs

Standard/RFC	Title
RFC 5286	<i>Basic Specification for IP Fast Reroute: Loop-Free Alternates</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for IS-IS Remote 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 IS-IS Remote Loop-Free Alternate Fast Reroute

Feature Name	Releases	Feature Information
IS-IS Remote Loop-Free Alternate Fast Reroute	15.2(2)S 15.2(1)SY	The IS-IS Remote Loop-Free Alternate Fast Reroute feature enables a backup repair path in the event of node failure, even if the path is multiple hops away. The following commands were introduced or modified: fast-reroute remote-lfa and show isis fast-reroute .

