Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute

First Published: March 30, 2011

Last Updated: March 30, 2011

When a link or router fails, distributed routing algorithms compute new routes that take into account the change. The time taken for computation is called the 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 IS-IS IPv4 Loop-Free Alternate (LFA) Fast Reroute (FRR) feature to reduce the routing transition time to less than 50 milliseconds by using a precomputed alternate next hop. When a router is notified of a link failure, it immediately switches over to the repair path to reduce traffic loss.

The IS-IS IPv4 Loop-Free Alternate Fast Reroute feature 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 results (the repair paths) are sent to the routing information base (RIB). The repair path installation is done by Cisco Express Forwarding.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see 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 for Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute, on page 9.

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

Prerequisites for Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute

- LFA Fast Reroute (FRR) can protect paths 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 Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute

- A Multiprotocol Label Switching (MPLS) traffic engineering (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 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 of 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 Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute

IS-IS and IP FRR

In a normal network, when a local link fails, 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, the traffic directed towards the affected prefixes are discarded. This process can take hundreds of milliseconds.

In IP FRR, IS-IS computes the LFA next-hop routes for the forwarding plane to use in the event 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 one single LFA for a primary path. In the case of one primary path with multiple LFA paths, prefixes are distributed equally among the 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 the failure until information about the failure is propagated through a routing protocol. This process takes a minimum of several hundred of milliseconds. It is therefore necessary to arrange for the 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 the packets that would have used the failed link. These repair paths are used starting 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 have revised 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 when a failure is detected.

A repair node carries the traffic around the protected element towards its destination when the protected element fails. A protecting node detects the failure of the protected element and directs the traffic around the protected element when the traffic fails. The failed traffic is forwarded towards the repair node when the protected element fails. In general, a repair node may or may not be directly connected to the repairing node.

The IS-IS IPv4 Loop-Free Alternate Fast Reroute feature uses the following types of repair paths:

- Equal Cost Multi Path (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

Loop Free Alternate is a node other than the primary neighbor. Traffic is redirected to the 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 the traffic. An LFA must not cause loops. By default, the 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:

- Provide load sharing.

- Protect all destinations that have a per-prefix LFA, leaving only a subset (a node at the far side of the failure) unprotected.

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 the 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 done for each neighbor. Because IS-IS examines prefixes after SPF is done, 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 end result of the computation is a set of prefixes with primary paths, where some of the 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, the **fast-reroute per-prefix** command with the **tie-break** keyword should be configured. If a rule eliminates all candidate LFAs, then the rule is skipped.

A given primary path can have multiple LFAs. A routing protocol is required to implement the 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.

- Load-sharing—Distributes the 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 Equal Cost Multipaths (ECMPs).

- Secondary-path—Eliminates candidates that are ECMPs.

- Shared Risk Link Group (SRLG)—Eliminates candidates that belong to one of the protected path SRLGs.

How to Configure IS-IS 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

|	**Command or Action**	**Purpose**
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Router(config)# interface ethernet 1/1`	Configures an interface and enters interface configuration mode.
Step 4	**ip address** *ip-address mask* **Example:** `Router(config-if)# ip address 10.1.1.1 255.255.255.0`	Sets a primary or secondary IP address for an interface.

	Command or Action	Purpose
Step 5	**ip router isis** *area-tag* **Example:** `Router(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** *tag-number* **Example:** `Router(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).
Step 7	**exit** **Example:** `Router(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 8	**interface** *type number* **Example:** `Router(config)# interface ethernet 1/0`	Configures an interface and enters interface configuration mode.

Step 9	**ip address** *ip-address mask*
	
	Example:
	
	`Router(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** *area-tag* **Example:** `Router(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** *tag-number* **Example:** `Router(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:** `Router(config-if)# exit`	Exits interface configuration mode and enters global configuration mode.
Step 13	**router isis** *area-tag* **Example:** `Router(config)# router isis ipfrr`	Enables the IS-IS routing protocol and specifies an IS-IS process. • Enters router configuration mode.

	Command or Action	Purpose
Step 14	**net** *net* **Example:** `Router(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** *route-map-name*} **Example:** `Router(config-router)# fast-reroute per-prefix level-2 all`	Enables per-prefix FRR. • Configure the **all** keyword to protect all prefixes.
Step 16	**end** **Example:** `Router(config-router)# end`	(Optional) Exits router configuration mode and enters privileged EXEC mode.

Configuration Examples for IS-IS IPv4 Loop-Free Alternate Fast Reroute

Example: Configuring IP FRR Support

Figure 1 shows IP FRR protecting Border Gateway Protocol (BGP) next-hops by using interface tags.
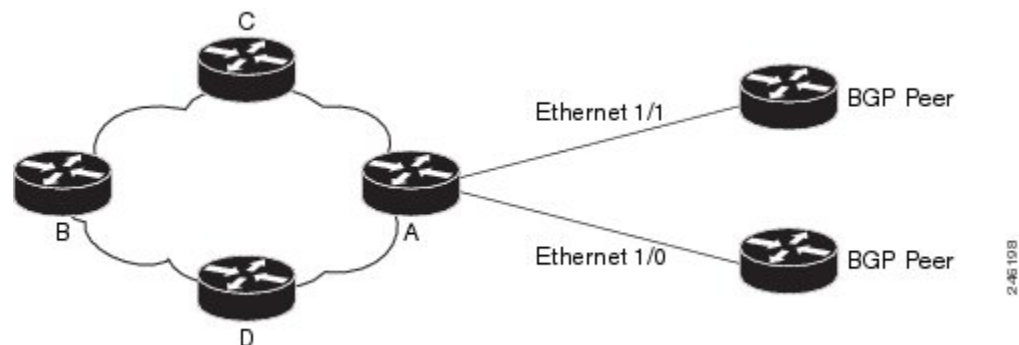
The following example shows how to configure IP FRR on Router A as shown in Figure 1. Router A will advertise the prefixes, 10.0.0.0/24 and 192.168.255.0/24 along with the tag 17.

```
Router# configure terminal
Router(config)# interface ethernet 1/1
```

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

The following example shows how to configure IP FRR on other routers as shown in Figure 1 . Other routers can use tag 17 to calculate repair paths for the two prefixes configured in Router A.

```
Router(config)# router isis
Router(config-router)# net 47.0004.004d.0001.0001.c11.1111.00

Router(config-router)# fast-reroute per-prefix level-2 route-map ipfrr-include
Router(config-router)# exit
Router(config)# route-map ipfrr-include
Router(config-router)# match tag 17
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
MPLS commands	Cisco IOS Multiprotocol Label Switching Command Reference
IS-IS commands	Cisco IOS IP Routing: ISIS Command Reference

Standards

Standard	Title
None	—

MIBs

MIB	MIBs Link
None	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 5286	*Basic Specification for IP Fast Reroute: Loop-Free Alternates*

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 Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute

Table 1: Feature Information for Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute, on page 10 lists the release history for this feature.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn . An account on Cisco.com is not required.

Note   Table 1: Feature Information for Configuring IS-IS IPv4 Loop-Free Alternate Fast Reroute, on page 10 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.

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

Feature Name	Releases	Feature Information
IS-IS IPv4 Loop-Free Alternate Fast Reroute	15.1(2)S 15.2(1)SY1	When a link or router fails, distributed routing algorithms compute new routes that take into account the change. The time taken for computation is called the routing transition. Until the transition is complete and all routers are converged on a common view of the network, connectivity between the source and destination pairs is interrupted. You can use this feature to reduce the routing transition time to less than 50 milliseconds by using a precomputed alternate next hop. When a router is notified of a link failure, it immediately switches over to the repair path to reduce traffic loss. The IS-IS IPv4 Loop-Free Alternate Fast Reroute feature focuses on the precomputation of repair paths. The repair path computation is done by the IS-IS routing protocol and the results (the repair paths) are sent to the RIB. The repair path installation is done by Cisco Express Forwarding. The following commands were introduced or modified: **debug isis fast-reroute**, **fast-reroute load-sharing**, **fast-reroute per-prefix**, **fast-reroute tie-break**, **show isis fast-reroute**.