IS-IS Configuration Guide for Cisco 8000 Series Routers, Cisco IOS XR Releases

Cisco 8000 Series Routers
Updated September 8, 2026

Topics included

1 Getting Started....6

2 YANG Data Models for IS-IS Features....8
Using YANG Data Models....9

3 Introduction to IS-IS....10
IS-IS protocol....11
IS-IS protocol features in Cisco IOS XR....11
Standards and RFCs supported....11
IS-IS use cases....12
Multi-instance IS-IS instances....12
Restrictions on interface usage with multiple IS-IS instances....12

4 Basic IS-IS Configuration....14
IS-IS configuration grouping....15
Configuration modes for IS-IS routing in Cisco IOS XR....15
Enable IS-IS and configure level 1 or level 2 routing....16

5 IS-IS Topology Modes....18
Multi-topology IPv6 for IS-IS....19
Multitopology route restrictions....19
Configure multitopology IPv6 for IS-IS....19
Single-topology IPv6 for IS-IS....20
Configure single topology mode for IS-IS....20
Configure a global multicast topology and associate it with network interfaces....22
Enable multicast IS-IS topology....23
Assign an interface to an IS-IS topology....24

6 Route Management and Distribution....26
Default route....27
Conditional default route originations in IS-IS....27
Configure conditional default route origination in IS-IS based on BGP neighbor status....28
Configure route customization for IS-IS....32
Redistributing IS-IS routes between multiple instances....34
Configure a routing policy....34
Set priority for adding prefixes to RIB....35
Route tags in IS-IS....36

Configure IS-IS interface route tags....36
Static IS-IS BGP tracking....38
Best practices for static IS-IS neighbor advertisement and connectivity....39
Configure IS-IS static neighbor tracking....39

7 Metrics and SPF Calculation....44
IS-IS auto-cost reference bandwidth....45
How auto-cost metric calculation works....45
IS-IS auto-cost metric information....46
How IS-IS tracks interface bandwidth changes....47
Guidelines for auto-cost reference bandwidth....48
Configure auto-cost reference bandwidth....48
Configure SPF interval for an IS-IS single-topology address family....51
IS-IS SPF interval delay....52
Configure IETF-compliant SPF calculation intervals....53

8 Interface Management and Flooding Control....56
IS-IS interface types....57
IS-IS adjacency stagger....57
Restrictions for IS-IS adjacency stagger....57
Configure IS-IS adjacency stagger....57
Configure IS-IS adjacencies....58
IS-IS static neighbors....60
Configure an IS-IS static neighbor....61
ISIS link group....62
Configure link group profile....63
Associate an IS-IS interface address family with a link group....64
Interface flood blocks....65
Mesh groups....65
LSP flood limits....66
LSP fast flood control mechanisms for IS-IS networks....66
Configure LSP flooding parameters for IS-IS 69
Layer 2 ACLs on physical and subinterfaces....72
Restrictions for configuring Layer 2 ACLs on physical and subinterfaces....73
Configure ISIS PDU rejection on Layer 2 interfaces....73

9 LSP Lifetimes....78
Maximum LSP lifetime and refresh interval....79
Minimum LSP lifetime thresholds....79

10 IS-IS Protection and Resiliency Enhancements....82
IS-IS non-stop forwarding 83
Configure nonstop forwarding for IS-IS....84

ISIS Non-stop Routing....85
Configure ISIS-NSR....85
IS-IS OOR protection mechanisms....86
Overload bits in OOR conditions....89
Configure overload bit suppression during OOR conditions....89
Global weighted SRLG protections....90
Configure global weighted SRLG protection....90
IS-IS penalties for link delay anomalies....93
Configure IS-IS penalties for link delay anomalies....93
IS-IS penalties for link loss anomalies....94
How IS-IS penalty for link loss anomaly works....95
Configure IS-IS penalty for link loss anomaly....96
Inter-area route propagation limits in IS-IS....97
Deployment models for inter-area route propagation limits in IS-IS....99
Restrictions for inter-area route propagation limits in IS-IS....99
Defaults and ranges for inter-area route propagation limits in IS-IS....99
Configure a hard limit threshold on propagated routes....99

11 IS-IS Authentication and Key Management....104
IS-IS authentication methods....105
Configure authentication for IS-IS....106
Configure keychains for IS-IS....107

12 MPLS Traffic Engineering and IS-IS Integration....110
MPLS Traffic Engineering....111
Configure MPLS TE for IS-IS....112
MPLS TE adjacencies....113
MPLS label distribution protocol IGP synchronization....113
Configure MPLS LDP IS-IS synchronization....115
MPLS TE interarea tunnels....116
Label distribution protocol IGP auto-configurations....117
MPLS LDP-IGP synchronization scenarios with LDP graceful restart....117
MPLS LDP-IGP synchronization compatibilities with IGP NSF....117

13 Fast Reroute and Load Balancing....120
IP fast reroute....121
Configure IP or LDP fast reroute loop-free alternate for an IS-IS interface....122
Unequal-cost multipaths for IS-IS....123

14 IS-IS Overload Bits, Avoidance, and Shutdown Modes....126
Overload bits on routers....127
Overload bit behavior in multitopology mode....127
IS-IS overload bit avoidance....127

Configure IS-IS overload bit avoidance....128
Configure IS-IS router overload handling....129
IID-TLV behavior in IS-IS protocol instances....132
IS-IS protocol shutdown mode....133
Configure IS-IS protocol shutdown mode....135

15 Multicast Features....136
Multicast-intact....137
Enable multicast-intact for IS-IS routes....137
IS-IS multicast topologies....138

16 Management Information Base for IS-IS....140
Management information base (MIB)....141

1 Getting Started

Outlines the release history and update summary for the IS-IS Configuration Guide. It enables users to access the latest IOS XR features and release information through a centralized, continuously updated resource.

This cumulative guide provides a single, continuously updated version that includes all the latest IOS XR features and release updates. It simplifies your experience by letting you bookmark one link and access the complete guide, instead of navigating through multiple release-specific versions.

Specific changes or updates tied to individual releases are clearly called out within the relevant sections. For a list of features introduced in a specific release, refer to the Release Notes.

The table lists the release numbers for which this document has been updated since its initial publication.

Table 1: Changes to this document

Date	Summary
September 2026	Republished for Release 26.3.1
June 2026	First published for Release 26.2.1

2 YANG Data Models for IS-IS Features

Topics:

- Using YANG Data Models

Introduces YANG data models for IS-IS features, highlighting programmatic configuration and data collection in Cisco IOS XR to streamline network automation and management.

You will learn about the YANG data models for IS-IS features.

Using YANG Data Models

Explains how to leverage YANG data models for automating the configuration and monitoring of IS-IS features, outlining model usage and integration methods for streamlined network management.

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the Github repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPaths. To view a comprehensive list of the data models supported in a release, navigate to the *Available-Content.md* file in the repository.

You can also view the data model definitions using the YANG Data Models Navigator tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.

3 Introduction to IS-IS

Topics:

- IS-IS protocol
- IS-IS use cases
- Multi-instance IS-IS instances

Introduces the IS-IS protocol, outlining core features, standards support, implementation restrictions, common use cases, and multi-instance deployment options for configuring and managing IS-IS in network environments.

IS-IS protocol

Outlines IS-IS protocol fundamentals, provides details on protocol features in Cisco IOS XR, and lists supported standards and RFCs for comprehensive understanding and compliance.

An Intermediate System-to-Intermediate System (IS-IS) protocol is a link-state interior gateway protocol that

- supports hierarchical routing with Level 1 and Level 2 areas,
- synchronizes network topology information using Link State Packets (LSPs), and
- calculates optimal routing paths using the Shortest Path First (SPF) algorithm.

Background and evolution

Originally developed for the OSI protocol suite, IS-IS has evolved to support both IPv4 and IPv6 routing. It is widely used in enterprise and service provider networks because it provides scalability, fast convergence, and operational stability in large and complex topologies.

Key characteristics of IS-IS

IS-IS provides these key capabilities:

- Support for single-topology and multitopology routing for IPv4, IPv6, and multicast
- Hierarchical routing through Level 1 and Level 2 areas
- Integration with advanced networking features such as traffic engineering, fast reroute, and routing security
- Efficient operation in large-scale network environments with minimal manual route management

IS-IS protocol features in Cisco IOS XR

Lists the key IS-IS protocol features, configuration options, and integrations supported in Cisco IOS XR.

Cisco IOS XR implements IS-IS as a standards-based routing protocol for IPv4 and IPv6 networks. The implementation supports a broad set of features for flexible deployment and interoperability in modern network environments.

Cisco IOS XR supports these IS-IS capabilities:

- Single-topology and multitopology IS-IS for IPv4 and IPv6
- Flexible configuration of Level 1 and Level 2 routing
- Support for multiple areas and multiple IS-IS instances
- Control over SPF calculations and LSP flooding behavior
- Route-policy integration for route control and redistribution
- Security features, including HMAC-MD5 authentication and keychain management
- Integration with MPLS, fast reroute, traffic engineering, and multicast routing

Standards and RFCs supported

Lists the industry standards and RFCs that Cisco IOS XR's IS-IS implementation supports to ensure protocol interoperability and consistent operation in multi-vendor networks.

Cisco's IS-IS implementation adheres to these key standards and RFCs, ensuring interoperability and reliable operation across different network devices:

- ISO/IEC 10589: Specifies the IS-IS routing protocol for OSI environments, defining the original packet structure and protocol operation.
- RFC 1195: Extends IS-IS to support routing of IP (IPv4) traffic, enabling IS-IS for use in IP-based networks.
- IETF IS-IS Working Group drafts: Include ongoing protocol enhancements such as IPv6 support, traffic engineering, and multitopology routing. Cisco IOS XR supports these advanced features for modern networking.
- RFC 8405: Standardizes algorithms for shortest-path-first (SPF) interval tuning, improving efficiency and resiliency during network convergence.
- RFC 4444: Defines Management Information Base (MIB) objects for IS-IS, enabling effective SNMP-based network monitoring and management.

Compliance with these standards allows Cisco IOS XR's IS-IS to interoperate seamlessly across heterogeneous network environments.

IS-IS use cases

Describes practical scenarios and applications where the IS-IS protocol is implemented within network infrastructures.

IS-IS is widely deployed in these scenarios:

- Service provider core and aggregation networks: Enables efficient, scalable routing and high availability.
- Large enterprise campus networks: Supports rapid network convergence and robust protocol stability.
- Multicast routing domains: Facilitates multicast forwarding and group membership management.
- MPLS-enabled traffic engineering backbones: Provides optimal path selection and traffic optimization for MPLS networks.
- Dual-stack IPv4/IPv6 network migration: Assists organizations in migrating from IPv4 to IPv6 while maintaining seamless operations.

Multi-instance IS-IS instances

Explains multi-instance IS-IS operations, including configuration approaches and restrictions encountered when implementing multiple IS-IS instances on Cisco platforms.

Multi-instance IS-IS instance is a routing protocol feature that

- allows you to configure up to 16 separate IS-IS instances on a device
- enables MPLS to operate on multiple IS-IS processes when they run on distinct sets of interfaces, and
- prevents double-booking by enforcing that each interface associates with only one IS-IS instance, generating configuration errors if violated.

Each IS-IS instance is treated equally by the Routing Information Base (RIB). When redistributing routes between instances, the RIB does not automatically prefer Level 1 over Level 2 routes. You must manually enforce route preference by setting different administrative distances for each instance type.

Restrictions on interface usage with multiple IS-IS instances

Provides operational limitations and interface assignment rules when configuring multiple IS-IS instances.

These restrictions apply when configuring multiple IS-IS instances:

- An interface can be associated with only one IS-IS instance (process).
- Instances may not share an interface; each interface assignment is exclusive.

4 Basic IS-IS Configuration

Topics:

- IS-IS configuration grouping
- Configuration modes for IS-IS routing in Cisco IOS XR
- Enable IS-IS and configure level 1 or level 2 routing

Outlines fundamental IS-IS configuration concepts, routing modes, and procedural steps for enabling and deploying IS-IS level 1 or level 2 routing within Cisco IOS XR environments.

IS-IS configuration grouping

Describes IS-IS configuration groupings, presenting structured references to configuration elements and the organization of IS-IS parameters for streamlined deployment in Cisco IOS XR.

Cisco IOS XR organizes all IS-IS configurations within router IS-IS configuration mode. This includes both the global IS-IS settings and the portions of interface configurations associated with IS-IS.

To display the entire IS-IS configuration, including all configured IS-IS instances, interface assignments, and interface attributes, use the **show running router isis** command. This command presents the running configuration for all IS-IS instances in a consolidated format.

Configuration modes for IS-IS routing in Cisco IOS XR

Explains available configuration modes for IS-IS routing, detailing different methods for managing and applying IS-IS settings within Cisco IOS XR.

IS-IS routing in Cisco IOS XR utilizes several hierarchical configuration modes that organize routing setup and protocol parameters. Understanding the function and attributes of each mode enables efficient IS-IS configuration and troubleshooting.

IS-IS configuration modes overview

Configuration mode	Purpose	Entry command example	Key Attributes
Router Configuration Mode	Start IS-IS process and set global options	`configuration → router isis <name>`	Global IS-IS settings, process name, global enablement
Router Address Family Configuration Mode	Set protocol parameters for address families	`router isis <name> → address-family ipv4 unicast`	IPv4/IPv6 selection, multicast/unicast options, address-family features
Interface Configuration Mode	Apply IS-IS to router interfaces	`router isis <name> → interface GigabitEthernet 0/3/0/0`	Interface selection, IS-IS activation, interface-specific options
Interface Address Family Configuration Mode	Fine-tune address-family options per interface	`interface ... → address-family ipv4 unicast`	Address-family, per-interface features, parameter customization

Example configuration commands

These examples illustrate how to enter each IS-IS configuration mode.

- Router configuration mode

```
Router# configure
Router(config)# router isis isp
Router(config-isis)#
```

- Router address family configuration mode

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)#
```

- Interface configuration mode

```
Router(config)# router isis isp
Router(config-isis)# interface GigabitEthernet 0/3/0/0
Router(config-isis-if)#
```

- Interface address family configuration mode

```
Router(config)# router isis isp
Router(config-isis)# interface GigabitEthernet 0/3/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)#
```

Configuration hierarchy overview

These configuration modes are nested and hierarchical, enabling modular configuration from global IS-IS settings down to per-interface and per-address-family details. This approach ensures protocol flexibility and tailored routing optimization in Cisco IOS XR environments.

Enable IS-IS and configure level 1 or level 2 routing

Provides instructions for enabling IS-IS on network devices and configuring level 1 or level 2 routing, guiding users through the essential steps for IS-IS deployment in Cisco IOS XR.

Set up IS-IS protocol and define the router's participation in intra-area, backbone, or both routing domains.

IS-IS is a dynamic routing protocol used to exchange routing information in large networks. You can configure IS-IS before assigning IP addresses, but actual routing will not occur until an IP address is configured on at least one interface. The routing level determines how the router forms adjacencies and participates within the network.

Before you begin

- Ensure the router has at least one configured IP address.
- Identify the network entity title (NET) and desired IS-IS routing level for your topology.

Procedure

1. Enable IS-IS for the desired routing instance.

Example

```
Router# configure
Router(config)# router isis isp
```

By default, all IS-IS instances are automatically Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.

2. Configure a network entity title (NET) for the IS-IS instance.

   Example

   ```
   Router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
   ```

   - Specify a NET for each routing instance if you are configuring multi-instance IS-IS.
   - This example configures a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110.00.
   - To specify more than one area address, specify additional NETs. Although the area address portion of the NET differs, the systemID portion of the NET must match exactly for all of the configured items.

3. (Optional) Choose the IS-IS routing level for this instance.

   Example

   ```
   Router(config-isis)# is-type level-2-only
   ```

   By default, every IS-IS instance acts as a **level-1-2** router.

   Use these options as needed:

   - - **level-1**: Only intra-area routing (Level 1 adjacencies).
   - - **level-2-only**: Only backbone routing (Level 2 adjacencies).
   - - **level-1-2**: Default; both intra-area and backbone routing.

4. (Optional) Verify the IS-IS configuration.

   Example

   ```
   Router# show isis protocol
   ```

   Displays summary information about the IS-IS instance.

5 IS-IS Topology Modes

Topics:

- Multi-topology IPv6 for IS-IS
- Single-topology IPv6 for IS-IS
- Configure a global multicast topology and associate it with network interfaces
- Enable multicast IS-IS topology

Outlines IS-IS topology modes, including multitopology and single-topology IPv6 support, topology configuration procedures, interface association, and restrictions to optimize routing in complex network environments.

Multi-topology IPv6 for IS-IS

Introduces multitopology IPv6 support in IS-IS, detailing key restrictions and providing step-by-step instructions for configuring multitopology routing to enable flexible IPv6 deployment in diverse network architectures.

Multi-topology IPv6 for IS-IS is a routing protocol feature that

- enables simultaneous support for both IPv4 and IPv6 address families within IS-IS
- uses multitopology as default behavior when interfaces are configured for both IPv6 and IPv4, and
- requires explicit configuration to enable single-topology IPv6 using the **single-topology** command.
- By default, multi-topology support is activated in software as soon as interfaces for both IPv4 and IPv6 are detected in the IS-IS configuration stanza.
- To enable single-topology IPv6, configure the **single-topology** command in the IPv6 router address family configuration submode of the IS-IS router stanza.

Multitopology route restrictions

Outlines multitopology route restrictions, including supported VRF, protocols (PIM and IS-IS), topology selection limitations, default policies for route sources, and RPL configuration options for multicast and unicast deployments.

- Only the default VRF is currently supported in multitopology solutions.
- Currently, only protocol-independent multicast (PIM) and intermediate system-intermediate system (IS-IS) routing protocols are supported.
- Topology selection is restricted to (S, G) route sources for both SM and SSM. Static and IS-IS are the only interior gateway protocols (IGPs) that support multitopology deployment.
- For non-(S, G) route sources such as a rendezvous point or bootstrap router, or when no route policy is configured, the default policy is used; either a unicast-default or multicast-default table is selected based on OSPF, IS-IS, or MBGP configuration.
- When configuring routing policy language (RPL), both multicast and unicast keywords are available with the **address-family** (**ipv4** | **ipv6**) command, but only multicast SAFI topologies can be configured globally.

Configure multitopology IPv6 for IS-IS

Configure multitopology IPv6 for IS-IS when IPv6 must operate in a separate topology from IPv4.

Enable IS-IS to support IPv6 networks in their own independent topology, separate from IPv4, for greater routing flexibility.

Use this task when you require IPv6 routing to run in a dedicated topology within IS-IS, rather than sharing the IPv4 topology. This separation is essential for scenarios where the operational requirements for IPv4 and IPv6 differ.

Before you begin

- Confirm IS-IS is enabled and supported on your device.
- Ensure IPv6 address configuration is permitted on your interfaces.

Procedure

Configure the IS-IS instance and enable IPv6 on the interface.

Example

```
Router(config)# router isis isp
Router(config-isis)# net 49.0000.0000.0001.00
Router(config-isis)# interface POS0/3/0/0
Router(config-isis-if)# address-family ipv6 unicast
Router(config-isis-af)# metric-style wide level 1
Router(config-isis-af)# exit
Router(config-isis)# interface POS0/3/0/0
Router(config-isis-if)# ipv6 address 2001::1/64
Router(config-isis)# exit
```

Single-topology IPv6 for IS-IS

Explains single-topology IPv6 operation in IS-IS and guides users through configuration steps to implement and manage single-topology IPv6 within network environments.

Single-topology IPv6 mode is an IS-IS routing approach that

- enables IS-IS for IPv6 to be configured on interfaces alongside IPv4 network protocols
- requires all interfaces and routers in an IS-IS area or domain to support identical network layer protocols, and
- uses a shared link topology allowing a single shortest path first (SPF) computation for both IPv4 and IPv6 routes.

In single-topology mode, IPv6 topologies work with both narrow and wide metric styles, as used in IPv4 unicast topology. During single-topology operation, one shortest path first (SPF) computation for each level is used to compute both IPv4 and IPv6 routes. This is possible because both IPv4 IS-IS and IPv6 IS-IS routing protocols share a common link topology.

Configure single topology mode for IS-IS

Configure IS-IS in single topology mode to route both IPv4 and IPv6 traffic on a shared topology.

Enable IS-IS single topology mode so IPv6 uses the existing IPv4 topology.

Use single topology mode when you want IPv6 to follow the IPv4 topology instead of using a separate topology. This simplifies routing configuration when both IP families share the same network design.

Before you begin

- Confirm all required address families are enabled on IS-IS interfaces.
- The **single-topology** command must be entered within the IPv6 address-family submode.
- If you use **adjacency-check disable**, you do not need to enable all address families per interface.
- The default metric style is narrow; configure metrics in the IPv4 or IPv6 address-family stanza as needed.

Procedure

1. Configure the interface and IP addresses.

Example

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# interface HundredGigE 0/1/0/3
```

Use one these commands:

- For IPv4:

  ipv4 address *address mask*

- For IPv6:

  ipv6 address *ipv6-prefix | prefix-length*> **eui-64**

  ipv6 address *ipv6-address* **link-local**

  ipv6 enable

  Notes:

  - If IS-IS uses IPv4, configure IPv4 addressing on all required interfaces.
  - Use the **eui-64** keyword to derive the IPv6 address from the interface ID.
  - Use the **link-local** keyword for a specific link-local address.

2. Configure IS-IS routing and enable single topology mode.

   Example

   ```
   Router# configure
   Router(config)# router isis isp
   Router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
   Router(config-isis)# address-family ipv6 unicast
   Router(config-isis-af)# single-topology
   Router(config-isis-af)# exit
   ```

   - Add multiple NETs if advertising more than one area address.
   - Use the same system ID in all NETs for the same instance.

3. Configure IS-IS interface parameters and address families.

   Example

   ```
   Router(config-isis)# interface HundredGigE 0/1/0/3
   Router(config-isis-if)# circuit-type level-1-2
   Router(config-isis-if)# address-family ipv4 unicast
   Router(config-isis-if)# commit
   ```

4. (Optional) Verify the IS-IS interface state.

   Example

   ```
   Router# show isis interface HundredGigE 0/1/0/3
   ```

5. Verify the IS-IS topology view.

Example

```
Router# show isis topology
```

Single topology IS-IS for IPv4 and IPv6

```
router isis isp
 net 49.0000.0000.0001.00
 address-family ipv6 unicast
  single-topology
 interface POS0/3/0/0
  address-family ipv4 unicast
  address-family ipv6 unicast
!
interface POS0/3/0/0
 ipv4 address 10.0.1.3 255.255.255.0
 ipv6 address 2001::1/64
```

Configure a global multicast topology and associate it with network interfaces

Provides procedural instructions for configuring a global IS-IS topology and associating it with network interfaces to enable effective topology management and integration.

Set up a global topology in the default VRF, associate it with network interfaces, and control routing table attributes.

Use this procedure to configure multicast topologies as part of the IS-IS workflow on Cisco IOS XR routers. Associating a topology with interfaces ensures routes are correctly added to the routing table.

Before you begin

- Ensure you have access to global configuration mode on your router.
- Identify the topology name to use.
- Know which interfaces you intend to assign to the topology.

Procedure

1. Define the multicast topology and (optionally) set a prefix limit.

   Example

   ```
   Router(config)# address-family ipv4 multicast topology green
   Router(config-af)# maximum prefix 100
   ```

   This sets up a topology named "green" in the default VRF and limits it to 100 prefixes.

2. Associate one or more interfaces with the topology and enable it.

   For each interface you want to assign, enter interface configuration mode and enable the topology.

Example

```
Router(config-af)# interface GigabitEthernet 0/3/0/0
   Router(config-if)# address-family ipv4 multicast topology green
   Router(config-af)# interface GigabitEthernet 0/3/2/0
   Router(config-if)# address-family ipv4 multicast topology purple
```

This associates "GigabitEthernet 0/3/0/0" with the "green" topology and "GigabitEthernet 0/3/2/0" with the "purple" topology. Repeat these two commands for each interface and topology combination you require.

Enable multicast IS-IS topology

Provides instructions for enabling and managing IS-IS topologies by assigning unique topology names and IDs, supporting advanced routing for IPv4 and IPv6 multicast environments within multiple routing domains.

Establish and distinguish IS-IS topologies in a routing domain by associating a topology ID with a named topology.

Use this task to create and manage multiple routing domains within IS-IS by assigning unique topology names and IDs. This enables advanced routing features and supports both IPv4 and IPv6 multicast topologies.

Before you begin

- Ensure you have access to the device in configuration mode.
- Prepare the instance ID, topology name, and numeric topology ID (value between 6 and 4095).

Follow these steps to enable an IS-IS topology:

Procedure

1. Enter IS-IS router configuration mode:

   Example

   ```
   Router(config)# router isis purple
   ```

2. Associate a named topology with an address family:

   Example

   ```
   Router(config-isis)# address-family ipv4 multicast topology green
   ```

3. Assign the numeric topology ID to the named topology:

   Example

   ```
   Router(config-isis-af)# toplogy-id 122
   ```

   ID must be in the range 6 to 4095.

Assign an interface to an IS-IS topology

Provides instructions for assigning interfaces to IS-IS topologies, including prerequisites, step-by-step configuration, and verification to enable flexible network segmentation and routing instances such as Layer 3 VPNs or multicast domains.

Enable IS-IS multi-topology routing by associating interfaces with distinct logical topologies for improved network segmentation and control.

Use this task when you want to segregate network routing into separate topologies, such as for Layer 3 VPNs or multicast domains. Assigning interfaces to individual IS-IS topologies allows separate route computation and advertisement for each topology.

Before you begin

- Verify you have privileged EXEC access on the router.
- Ensure IS-IS routing is enabled and properly initialized on the device.
- Prepare the IS-IS instance ID, network entity title (NET), interface details, and desired topology names.

Procedure

1. Enter IS-IS router configuration mode:

   Example

   ```
   Routing(config)# router isis purple
   ```

2. Configure the network entity title for the IS-IS instance:

   Example

   ```
   Routing(config-isis)# net netname
   ```

   This sets the ISO NET used for IS-IS routing.

3. Specify the interface to associate with the IS-IS instance:

   Example

   ```
   Routing(config-isis)# interface gigabitethernet 0/3/0/0
   ```

   This places the interface into IS-IS interface configuration mode.

4. Assign the interface to a specific address family and topology:

   Example

   ```
   Routing(config-isis-if)# address-family ipv4 multicast topology green
   ```

   This places the interface into the desired IS-IS topology and address family.

5. Repeat steps 3 and 4 for each additional interface and topology as needed.

The interfaces are now associated with the specified IS-IS topologies. Verify the configuration using commands such as **show isis interface** or **show isis topology** to confirm correct placement.

6 Route Management and Distribution

Topics:

- Default route
- Conditional default route originations in IS-IS
- Configure route customization for IS-IS
- Configure a routing policy
- Set priority for adding prefixes to RIB
- Route tags in IS-IS
- Static IS-IS BGP tracking

Outlines key concepts and procedures for managing and distributing network routes, including default and conditional route origination, IS-IS customization, routing policy configuration, route tagging, and BGP status tracking for optimized routing operations.

Default route

Introduces default route concepts and their role in route management, helping users understand how default routes impact network forwarding decisions.

A default route is a routing configuration that

- directs packets destined for networks not explicitly listed in the routing table to a specified next hop,
- can be automatically or manually installed into a routing domain such as IS-IS, and
- is controlled and filtered using route policies to meet diverse network requirements.

Default route behavior in IS-IS routing domains

You can force a default route into an IS-IS routing domain. When you configure redistribution of routes into IS-IS, the software does not, by default, redistribute the default route. The **default-information originate** command generates a default route within IS-IS, which can be managed using route policies. Route policies allow you to identify the announcement level, specify other filtering options, and conditionally advertise the default route depending on the existence of another route in the router's routing table.

Conditional default route originations in IS-IS

Details conditional default route origination in IS-IS influenced by BGP neighbor status and provides instructions for configuring this functionality to enhance routing flexibility.

Conditional default route origination is an IS-IS routing feature that

- enables IS-IS to originate a default route only when specific BGP neighbor or route conditions are satisfied,
- improves network reliability and responsiveness by linking the advertisement of default routes to dynamic routing states, and
- prevents incorrect default advertisements when BGP neighbors or watched routes are not active.

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
Conditional Default Route Originating in IS-IS	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
Conditional Default Route Originating in IS-IS	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.

Feature Name	Release Information	Feature Description
Conditional Default Route Originating in IS-IS	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM
Conditional Default Route Originating in IS-IS	Release 7.3.2	The Conditional Default Route Originating in IS-IS feature allows you to enhance the granularity of the default route the IS-IS originates based on a condition. It enables IS-IS to originate the default route based on the presence of a specific route in the RIB originated by a particular BGP speaker. This feature improves the reaction time of the watched route in the RIB by avoiding periodical queries of the routing policy. This feature enables you to respond to the client in a timely fashion when the watched route changes in the RIB.

Conditional default route origination in IS-IS enables administrators to control whether IS-IS advertises a default route into the network based on real-time conditions such as the state of a specific BGP neighbor or the presence of certain BGP-learned routes in the Routing Information Base (RIB). This dynamic behavior ensures the default route is only distributed when upstream connectivity is truly available.

Configure conditional default route origination in IS-IS based on BGP neighbor status

Guides you through configuring conditional default route origination in IS-IS based on BGP neighbor or watched-route state, including prerequisite setup, route-policy creation, neighbor tracking, address family configuration, and verification procedures for route advertisement and tracking status.

Control IS-IS default route advertisement based on BGP neighbor status or watched-route state.

Use this configuration to prevent advertising a default route unless specific BGP neighbor or route conditions are met, enhancing routing decision accuracy and stability.

Before you begin

- IS-IS must be configured on the router.
- BGP neighbor tracking should be defined.

- Required route-policy design must be in place.

Procedure

1. Configure the IS-IS instance and apply the route-policy under the IPv4 and IPv6 address families.

   Example

```
Router(config)#router isis 1
Router(config-isis)#is-type level-2-only
Router(config-isis)#net 47.0000.0000.0005.00
Router(config-isis)#address-family ipv4 unicast
Router(config-isis-af)#metric-style wide
Router(config-isis-af)#mpls traffic-eng level-2-only
Router(config-isis-af)#mpls traffic-eng router-id 10.5.5.5
Router(config-isis-af)#default-information originate route-policy policy_1
Router(config-isis-af)#segment-routing mpls sr-prefer
Router(config-isis-af)#exit
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#metric-style wide
Router(config-isis-af)#default-information originate route-policy policy_1
Router(config-isis-af)#segment-routing mpls sr-prefer
Router(config-isis-af)#exit
Router(config-isis)#exit
```

2. Configure a route policy that originates the default route based on the BGP neighbor tracking state.

   Example

```
Router(config)#route-policy track_bgp_neighbor
Router(config-rpl)#if track track-bgp-neighbors is up then
Router(config-rpl-if)#pass
Router(config-rpl-if)#endif
Router(config-rpl)#end-policy
```

3. Configure the watched-route policy, neighbor tracking, and IPv6 prefix-set objects used by the route policy.

   Example

```
Router(config)#route-policy track-bgp-neighbors
Router(config-rpl)#if rib-has-route async (192.1.1.0/24, 192.1.2.0/24) and
source in (10.2.35.1) and track track-bgp-neighbors is up then
Router(config-rpl-if)#pass
Router(config-rpl-if)#endif
Router(config-rpl)#end-policy
Router(config)#track track-bgp-neighbors
Router(config-track)#type bgp neighbor address-family state
Router(config-track)#exit
Router(config)#address-family ipv4 unicast
Router(config)#neighbor 10.2.35.1
Router(config)#prefix-set bgp_ipv6_neighbor_id
Router(config-pfx)#10:2:35::1
Router(config-pfx)#end-set
Router(config)#prefix-set bgp_ipv6_watched_routes
Router(config-pfx)#192:1:1::/112
Router(config-pfx)#192:1:2::/112
```

```
Router(config-pfx)#end-set
Router(config)#route-policy default_route_policy_ipv6
Router(config-rpl)#if rib-has-route async bgp_ipv6_watched_routes and protocol
 is bgp 100 and source in bgp_ipv6_neighbor_id then
Router(config-rpl-if)#pass
Router(config-rpl-if)#else
Router(config-rpl-if)#drop
Router(config-rpl-if)#endif
Router(config-rpl)#end-policy
```

4. Verify the BGP neighbor state, watched routes, tracking state, and the default route advertisement.

Example
Verify the BGP neighbor status.

```
Router(config)# show bgp neighbor brief
Mon Mar 8 13:30:27.312 IST
Neighbor    Spk AS Description      Up/Down NBRState
10.2.35.1    0  100 02:18:39         Established
10:2:35::1   0  100 02:18:40            Established
```

Example
Verify the IPv4 RIB route.

```
Router# show route ipv4 192.1.1.0/24
Mon Mar 8 13:33:14.726 IST
Routing entry for 192.1.1.0/24
   Known via "bgp 100", distance 200, metric 0, type internal
   Installed Mar 8 11:11:52.738 for 02:21:22
   Routing Descriptor Blocks
     10.2.35.1, from 10.2.35.1
       Route metric is 0
  No advertising protos.
```

Example
Verify the IPv6 RIB route.

```
Router# show route ipv6 192:1:1::/112
Mon Mar 8 13:33:31.340 IST
Routing entry for 192:1:1::/112
    Known via "bgp 100", distance 200, metric 0, type internal
    Installed Mar 8 11:11:52.738 for 02:21:38
    Routing Descriptor Blocks
      10:2:35::1, from 10:2:35::1
         Route metric is 0
      No advertising protos.
```

Example
Verify BGP neighbor status tracking.

```
Router# show track track-bgp-neighbors
Mon Mar 8 13:52:16.746 IST
Track track-bgp-neighbors
    BGP Neighbor AF IPv4 Unicast NBR 10.2.35.1 vrf default
```

```
    Reachability is UP
       Neighbor Address Reachablity is Up
       BGP Neighbor Address-family state is Up
    12 changes, last change IST Mon Mar 08 2021 11:11:52.741
    Delay up 0 secs(default), down 0 secs(default)
```

Example
Verify the default route status in IS-IS address family.

```
Router# show isis
Mon Mar 8 13:34:39.412
IST IS-IS Router: 1
    System Id: 0000.0000.0005
    Instance Id: 0
    IS Levels: level-2-only
    Manual area address(es): 47
    Routing for area address(es): 47
!! . .
Topologies supported by IS-IS:
   IPv4 Unicast
    .
    .
      Originating default route active since Mar 08 2021 11:12:05.914 IST
   IPv6 Unicast
   .
   .
      Originating default route active since Mar 08 2021 11:12:05.917 IST
!!
```

Example
Verify the IS-IS database.

```
Router# show isis database detail verbose r5 | i 0.0.0.0/0
Mon Mar 8 13:47:10.624 IST Metric: 0 IP-Extended 0.0.0.0/0

Router# show isis database detail verbose r5 | i ::/0
Mon Mar 8 13:47:10.727 IST Metric: 0 MT (IPv6 Unicast) IPv6 ::/0
```

Example
Verify the IPv4 IS-IS routes.

```
Router# show isis ipv4 route 0.0.0.0/0
Mon Mar 8 13:44:58.226 IST
L2 0.0.0.0/0 [10/115]
    via 10.1.35.2, TenGigE0/0/0/31, r5, SRGB Base: 16000, Weight: 0
```

Example
Verify the IPv6 IS-IS routes.

```
Router# show isis ipv6 route 0::0/0
Mon Mar 8 13:45:02.699 IST
L2 ::/0 [10/115]
```

```
 via fe80::28a:96ff:fee7:f418, TenGigE0/0/0/31, r5, SRGB Base: 16000, Weight:
0
```

Configure route customization for IS-IS

Explains how to customize routes for IS-IS, including example procedures for redistributing IS-IS routes between multiple instances to support complex routing environments.

Enable advanced routing functions in IS-IS, including injecting default routes and redistributing routes between IS-IS instances.

Use this task to perform route customization in IS-IS. You can inject default routes, redistribute routes from other IS-IS instances, and adjust routing preferences. Customization is optional and depends on your network design.

Procedure

1. Enable IS-IS routing for a specific instance.

   Example

   ```
   Router(config)# router isis isp
   ```

   Enables IS-IS routing for the specified routing process, and places the router in router configuration mode.

   By default, all IS-IS instances are automatically Level 1 and Level 2. You can change the level of routing to be performed by a particular routing instance by using the is-type command.

2. (Optional) Set the overload bit if needed.

   Example

   ```
   Router(config-isis)# set-overload-bit
   ```

   > **Note**
   >
   > Overload bit behavior does not apply during NSF restarts.

3. Specify the address family.

   Example

   ```
   Router(config-isis)# address-family ipv4 unicast
   ```

4. (Optional) Inject a default route into the routing domain.

Example

```
Router(config-isis-af)# default-information originate
```

- If you specify a route policy, it controls when the route is advertised.
- Without a route policy, the default route is advertised unconditionally at Level 2.

5. (Optional) Redistribute routes from one IS-IS instance into another.

   Example

   ```
   Router(config-isis-af)# redistribute isis 2 level-1
   ```

   In this example, an IS-IS instance redistributes Level 1 routes from another IS-IS instance.

6. (Optional) Summarize IPv4 or IPv6 prefixes.

   Example

   ```
   Router(config-isis-af)# summary-prefix 10.1.0.0/16 level 1
   ```

   or

   ```
   Router(config-isis-af)# summary-prefix 3003:xxxx::/24 level 1
   ```

   - IPv6 prefixes must be configured in IPv6 address family mode, IPv4 prefixes in IPv4 address family mode.

7. (Optional) Configure the maximum number of parallel paths allowed.

   Example

   ```
   Router(config-isis-af)# maximum-paths 16
   ```

   Sets the maximum parallel paths in the routing table.

8. (Optional) Define administrative distance for IS-IS routes. You can also configure the attached bit in Level 1 LSPs.

   Example

   ```
   Router(config-isis-af)# distance 90
   ```

   Example

   ```
   Router(config-isis-af)# set-attached-bit
   ```

Redistributing IS-IS routes between multiple instances

Provides configuration and operational reference for redistributing IS-IS routes between multiple IS-IS instances, including example topologies, sample configurations, and explanations of key commands for preferred route selection and backbone communication.

Redistributing IS-IS routes between multiple instances allows network administrators to propagate routing information across IS-IS instances configured with different levels (Level 1, Level 2). This process supports seamless area-to-backbone communication and ensures preferred route selection. Key commands include setting the attached bit and explicit redistribution.

Sample configuration

The example below shows two IS-IS instances:

- Instance 1: Restricted to Level 2 (backbone)
- Instance 2: Restricted to Level 1 (area)

Objective: The Level 1 instance propagates routes to the Level 2 instance using redistribution. The administrative distance is set higher for Level 2 so Level 1 routes are preferred. The attached bit is set for Level 1 to support backbone routing.

```
! Level 2 instance configuration
router isis 1
  is-type level-2-only
  net 49.0001.0001.0001.0001.00
  address-family ipv4 unicast
    distance 116
    redistribute isis 2 level 2
!
interface GigabitEthernet 0/3/0/0
  address-family ipv4 unicast
!
! Level 1 instance configuration
router isis 2
  is-type level-1
  net 49.0002.0001.0001.0002.00
  address-family ipv4 unicast
    set-attached-bit send always-set
!
interface GigabitEthernet 0/1/0/0
  address-family ipv4 unicast
```

- **set-attached-bit send always-set**: Sets the attached bit for the Level 1 instance, allowing it to advertise backbone capability when redistributing routes.
- **redistribute isis 2 level 2**: Enables the propagation of routes from instance 2 (Level 1) to instance 1 (Level 2).
- **distance 116**: Sets the administrative distance higher for Level 2, making Level 1 routes preferred.

Configure a routing policy

Guides you through defining and applying a routing policy to control IS-IS routing decisions for multicast and unicast traffic, including matching source and destination addresses, and referencing RPF topology tables.

Define and apply a routing policy to control packet handling based on source and destination IP addresses.

Routing policies guide routers in forwarding packets according to network requirements. This process includes matching traffic, applying policies, and using RPF topology tables.

Before you begin

- Decide the policy name and match criteria.
- Know the RPF topology tables you will reference.

Procedure

1. Create a routing policy by specifying a policy name.

   Example

   ```
   Router(config)# route-policy mt1
   ```

2. Add a conditional statement to match specific destination addresses.

   Example

   ```
   Router(config-rpl)# if destination in 10.0.0.1, 10.0.0.11 then
   ```

3. Add a conditional statement to match source addresses and set the RPF topology table accordingly.

   Example

   ```
   Router(config-rpl-if)# if source in (10.10.10.10) then
   Router(config-rpl-if-2)# set rpf-topology ipv4 multicast topology greentable
   Router(config-rpl-if-2)# else
   Router(config-rpl-if-else-2)# set rpf-topology ipv4 multicast topology
   bluetable
   Router(config-rpl-if-2)# endif
   Router(config-rpl-if)# endif
   Router(config-rpl)# end-policy
   ```

Set priority for adding prefixes to RIB

Provides instructions for configuring IS-IS to prioritize prefix addition to the RIB using access lists, prefix lists, or tag values, optimizing routing convergence and enabling policy-based route selection.

Ensure that certain prefixes are added to the RIB before others, helping to optimize network convergence and meet specific routing policies.

Use this task when you need IS-IS to prioritize which prefixes are added to the RIB based on an access list, prefix list, or tag value.

Before you begin

- Verify if IS-IS routing is enabled on your device.
- Prepare the access list, prefix list, or tag values you want to use for prioritization.

Procedure

1. Enter IS-IS router configuration mode for the desired routing instance.

   Example

   ```
   Router(config)# router isis isp
   ```

2. Specify the address family (IPv4 or IPv6) and enter address family configuration mode.

   Example

   ```
   Router(config-isis)# address-family ipv4 unicast
   ```

3. Configure the router to generate and accept only wide-link metrics for the specified level.

   Example

   ```
   Router(config-isis-af)# metric-style wide level 1
   ```

4. Set the prefix priority for SPF processing using an access list, prefix list, or tag value.

   Example

   ```
   Router(config-isis-af)# spf prefix-priority high tag 3
   ```

Route tags in IS-IS

Describes IS-IS support for route tags and gives procedures to tag IS-IS interface routes, facilitating advanced route identification and management.

Route tag is a route attribute that

- associates a numeric tag with individual route prefixes in IS-IS
- enables prioritization of route installation in the routing information base (RIB) based on tag values, and
- allows route policies to match route prefixes by tag, such as selecting certain prefixes for redistribution.

Route tags provide network operators with a mechanism for marking routes for administrative purposes, such as signaling special handling or classification within routing and redistribution policies.

- You can assign a route tag to internal route prefixes so that these routes are installed ahead of external prefixes in the RIB.
- Route tags can be used in redistribution policies to select which prefixes are propagated to other routing protocols.

Configure IS-IS interface route tags

Provides instructions for configuring IS-IS interface route tags, applying and verifying tag values for route identification and policy enforcement, and demonstrates using route-policy for advanced tagging.

Assign tags to IS-IS interface routes to enable route identification, filtering, and enforcement of routing policies.

Tagging IS-IS interface routes is useful when you need to apply policy, prioritize, or filter certain directly connected routes. Tag values are advertised in IS-IS and can be used for advanced routing policy control.

Procedure

1. Enter IS-IS router configuration mode for your instance.

   Example

   ```
   Router(config)# router isis isp
   ```

2. Enter IS-IS address-family configuration mode.

   Example

   ```
   Router(config-isis)# address-family ipv4 unicast
   ```

   Use **ipv6** keyword as needed for IPv6 routing.

3. (Optional) Configure metric-style for the IS-IS process.

   Example

   ```
   Router(config-isis-af)# metric-style wide level 1
   ```

   This sets the metric style to wide for Level 1 IS-IS.

4. Exit address-family configuration mode.

   Example

   ```
   Router(config-isis-af)# exit
   ```

5. Enter interface configuration mode for the desired interface.

   Example

   ```
   Router(config-isis)# interface HundredGigE 0/1/0/3
   ```

   Replace with your actual interface type and number.

6. Enter IS-IS address-family configuration under the interface, and configure the tag for the interface's advertised routes.

Example

```
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# tag 3
```

Use **ipv6** as needed for IPv6 routing.

This associates the tag value 3 with routes advertised via this interface. Choose a value that matches your network policy requirements.

7. Verify the tag assignment.

Example

```
Router# show isis ipv4 route detail
```

Confirm that the expected tags are present in the output.

Tagging routes with route-policy: Example

```
route-policy isis-tag-555
  if destination in (5.5.5.0/24 eq 24) then
    set tag 555
    pass
  else
    drop
  endif
end-policy

router isis uut
  net 00.0000.0000.12a5.00
  address-family ipv4 unicast
    metric-style wide
    redistribute static level-1 route-policy isis-tag-555
    spf prefix-priority critical tag 13
    spf prefix-priority high tag 444
    spf prefix-priority medium tag 777
```

This example shows how route tags can be set and managed using route-policy in IS-IS.

Static IS-IS BGP tracking

Explains static IS-IS BGP status concepts, supporting users in monitoring and managing the relationship between IS-IS and BGP protocols.

The static IS-IS BGP tracking feature is a routing capability that

- integrates static IS-IS neighbor advertisement with BGP peer state monitoring
- enables conditional advertisement of static IS-IS links based on BGP neighbor status, and
- improves topology accuracy and operational reliability by dynamically controlling static neighbor advertisement.

Feature Name	Release Information	Description
Static IS-IS BGP status	Release 26.1.1	Introduced in this release on: Fixed Systems (8200, 8700); Centralized Systems (8600); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]) You can now boost network accuracy and reliability with the static IS-IS BGP tracking feature. It only advertises static IS-IS neighbors when the matching BGP peer is up, using a simple route-policy to track their status. Two-way connectivity checks are turned off for static neighbors to avoid false failures, and IS-IS extended TLVs make sure these neighbors are clearly shown in IS-IS LSPs for accurate network mapping.

Two-way connectivity checks

Two-way connectivity check mechanisms ensure link reliability by checking connectivity in both directions. For static neighbors, since the remote device does not run IS-IS, Two-way connectivity check mechanism is disabled to avoid false negatives. This suppression allows the network controller to accurately represent the topology without expecting a two-way handshake on links where IS-IS is not active on the remote end.

IS-IS extended TLVs and their role in static neighbor representation

IS-IS extended TLVs are iS-IS protocol elements that carry link state information for adjacencies. You rely on IS-IS Extended TLVs to represent static neighbors seamlessly within iS-IS LSPs. This ensures that controllers receive consistent and accurate topology information without needing to understand new or different link types, maintaining interoperability and simplifying network management.

Best practices for static IS-IS neighbor advertisement and connectivity

Outlines best practices for reliable static IS-IS neighbor advertisement and connectivity, including configuring route policies based on BGP neighbor state, suppressing TWCC for non-IS-IS peers, and using extended TLVs for topology consistency.

Follow these best practices for static IS-IS neighbor advertisement and connectivity to maintain accurate and reliable network topology information.

- Always configure a route policy that tracks the state of the BGP neighbor associated with a static IS-IS link. This ensures you advertise the static neighbor only when the BGP peer is operational.
- Suppress two-way connectivity check (TWCC) for static neighbors where the remote device does not run IS-IS. This avoids false negatives and maintains accurate topology representation.
- Use IS-IS extended TLVs to advertise static neighbors identically to real IS-IS adjacencies. This enables seamless distribution of topology information without requiring protocol changes.

Configure IS-IS static neighbor tracking

Guides you through configuring static IS-IS neighbors, enabling conditional advertisement based on BGP session state, setting up tracking and route policies, and verifying connectivity to ensure accurate IS-IS network topology.

Configure static IS-IS neighbors and enable conditional advertisement for enhanced topology accuracy and reliable link status.

Use this task when setting up static neighbors in an IS-IS network to maintain precise topology representation and dependable connectivity.

Before you begin

- Ensure you have access to the network controller and the necessary permissions to configure IS-IS settings.
- Ensure of these information:
  - IS-IS System-ID of the static neighbor
  - Either or both remote IPv4 and IPv6 addresses
  - Affinity link to advertise
- Enable logging static neighbor advertisement changes

Follow these steps to perform static IS-IS neighbor advertisement and connectivity:

Procedure

1. Configure IS-IS static neighbor by using the **static-neighbor** command.

   Example

   ```
   Router(config)# router isis 1
   Router(config-isis)# interface GigabitEthernet0/2/0/7
   Router(config-isis-if)# static-neighbor
   Router(config-isis-static-nbr)# system-id 0001.0002.0007
   Router(config-isis-static-nbr)# remote ipv4 address 10.2.2.2
   Router(config-isis-static-nbr)# affinity RED
   Router(config-isis-static-nbr)#log-advertisement-changes
   ```

2. View the running configuration using the **show running-config** command.

   Example

   ```
   router isis 1
    interface GigabitEthernet0/2/0/7
     static-neighbor
      system-id 0001.0002.0007
      remote ipv4 address 10.2.2.2
      affinity RED
      log-advertisement-changes
     !
    !
   !
   ```

3. Configure a static neighbor explicitly to define a peer for the routing protocol on an interface, thereby bypass the dynamic adjacency discovery process. Despite this, the routing protocol remains active on the interface, and the local IP subnet will still be advertised to the statically configured neighbor. Enable conditional static neighbor advertisement based on the BGP session state.

Example

```
Router(config)# router isis 1
Router(config-isis)# affinity-map RED bit-position 7
Router(config-isis)# interface GigabitEthernet0/2/0/7
Router(config-isis-if)# point-to-point
Router(config-isis-if)# static-neighbor
Router(config-isis-static-nbr)# system-id 0001.0002.0007
Router(config-isis-static-nbr)# remote ipv4 address 10.1.1.2
Router(config-isis-static-nbr)# Remote ipv6 address: 2001:DB8::1
Router(config-isis-static-nbr)# advertise route-policy isis-bgp
Router(config-isis-static-nbr)# exit
Router(config-isis-static)# address-family ipv4 unicast
Router(config-if)# exit
Router(config-if)# address-family ipv6 unicast
Router(config-if)# exit
```

4. Configure a track object to monitor the state of a BGP neighbor for the IPv4 unicast address family by entering global configuration mode and specifying the neighbor's IP address.

Example

```
Router(config)# track track-bgp-neighbors
Router(config-track)# type bgp neighbor address-family state
Router(config-track-bgp-nbr-af)# address-family ipv4 unicast
Router(config-track-bgp-neighbor)# neighbor 10.2.35.1
Router(config-track-af)# exit
Router(config-track)# exit
```

5. Configure a route policy named *isis-bgp* that permits advertisement of a static IS-IS neighbor only when the associated BGP neighbor is up.

Example

```
Router(config)# route-policy isis-bgp
Router(config-rpl)# if bgp-nbr is up then
Router(config-rpl-if)# pass
Router(config-rpl-if)# else
Router(config-rpl-if)# drop
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
```

6. Use **show isis static-neighbors** command to verify the status of the static neighbors.

Example

```
Router# show isis static-neighbors
Fri Sep  6 05:32:18.889 PDT

IS-IS 1 static neighbors:

System Id          Interface        State
0001.0002.0007     Gi0/2/0/7        Up

  Remote IPv4 Address: 10.1.1.2
  Remote IPv6 Address: 2001:DB8::1
  Affinity: RED
  Route-policy: isis-bgp
```

```
  Level-1:
    Advertised in MT (Standard (IPv4 Unicast)): Yes
    Last Transition in MT (Standard (IPv4 Unicast)): 07:55:35.177
    Advertised in MT (IPv6 Unicast): Yes
    Last Transition in MT (IPv6 Unicast): 09:06:41.918
  Level-2:
    Advertised in MT (Standard (IPv4 Unicast)): Yes
    Last Transition in MT (Standard (IPv4 Unicast)): 07:55:35.183
    Advertised in MT (IPv6 Unicast): Yes
    Last Transition in MT (IPv6 Unicast): 09:06:41.917
```

The show output shows that the static IS-IS neighbor advertisement is controlled by the configured route policy named *isis-bgp*

The show output also shows that the static IS-IS neighbor is currently being advertised in the IS-IS Multi-Topology (MT) instance for the standard IPv4 unicast address family. This confirms that the route policy condition was met, that is the BGP neighbor is up, and the static neighbor is actively included in the IS-IS topology for IPv4 unicast routing.

7 Metrics and SPF Calculation

Topics:

- IS-IS auto-cost reference bandwidth
- Configure SPF interval for an IS-IS single-topology address family
- IS-IS SPF interval delay

Outlines IS-IS metric calculation strategies, auto-cost reference bandwidth implementation, interface bandwidth tracking, and SPF interval configuration procedures to optimize protocol performance and maintain accurate network topology computations.

IS-IS auto-cost reference bandwidth

Introduces IS-IS auto-cost reference bandwidth, detailing how metric calculation works, providing reference information, explaining interface bandwidth tracking, recommending auto-cost usage, and guiding configuration procedures.

The IS-IS auto-cost reference bandwidth feature is a routing metric mechanism that

- provisions IS-IS metrics based on the physical bandwidth of a link
- adjusts IS-IS metrics for bundle interfaces when bandwidth changes, and
- optimizes route selection by lowering the cost for high-speed links.

This feature enhances network adaptability and ensures that IS-IS routing decisions accurately reflect the current link capacities.

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
IS-IS auto-cost reference bandwidth	Release 25.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Centralized Systems (8600 [ASIC: Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]) The IS-IS auto-cost reference bandwidth feature automates IS-IS metric provisioning based on physical link bandwidth, optimizing path selection and reducing operational overhead. This feature allows you to configure a reference bandwidth, which IS-IS then uses to automatically calculate interface metrics. It also dynamically adjusts metrics for bundle interfaces when member links change, ensuring accurate and efficient routing without manual intervention. CLI: • **auto-cost (IS-IS)** • **show isis interface** : The Bandwidth field was added in the show output.

How auto-cost metric calculation works

Outlines how IS-IS auto-cost metric calculation dynamically updates interface costs by evaluating bandwidth, applying reference values, generating LSPs, running SPF algorithms, and updating routing tables to optimize path selection based on real-time link conditions.

Summary

The auto-cost metric calculation process dynamically determines and updates interface costs within IS-IS, ensuring routing decisions reflect current link bandwidths.

The key components involved in the process are:

- IS-IS: The routing protocol that performs the calculation and updates.
- Interface bandwidth: The actual capacity of a link, which is a primary input for the metric calculation.
- Configured reference bandwidth: A user-defined value that scales interface bandwidths into IS-IS metrics.

- Link State PDU (LSP) generation: The creation of data packets to propagate updated link-state information.
- Shortest Path First (SPF) run: The algorithm execution to re-calculate optimal paths based on new metrics.

The auto-cost metric calculation process involves IS-IS using interface bandwidth and a configured reference bandwidth to dynamically determine and update interface costs. This process leverages LSP generation and SPF runs to propagate updates and recalculate optimal routing paths.

Workflow

These stages describe how auto-cost metric calculation works:

1. Condition check: IS-IS verifies that auto-cost reference bandwidth is configured on an active interface or under the IS-IS router process, and that a default metric is not explicitly configured.
2. Initial metric determination: IS-IS calculates the interface cost based on its current bandwidth and the configured reference bandwidth.
3. Bandwidth change detection: IS-IS receives a notification indicating a change in the interface's bandwidth.
4. Metric recalculation: IS-IS updates the auto-cost reference bandwidth metric using the new interface bandwidth.
5. Routing update initiation: The updated metric triggers an LSP PDU generation.
6. Path re-evaluation: An SPF run is initiated to re-calculate routing paths.
7. Routing table update: The new metric is reflected in the routing table, influencing routing decisions.

Result

The IS-IS routing table dynamically reflects current link capacities, optimizing route selection based on real-time bandwidth conditions.

IS-IS auto-cost metric information

Details the automatic calculation of IS-IS auto-cost metrics using reference bandwidth, outlines metric calculation logic and key parameters, and provides additional notes on integer rounding and optimal path selection for effective route optimization.

The metric calculates automatically using reference bandwidth to support optimal, real-time path selection.

- The IS-IS auto-cost metric dynamically assigns routing metrics based on an interface's physical bandwidth.
- This feature optimizes route selection by lowering the cost for high-speed links and ensures routing decisions accurately reflect current link capacities.
- The auto-cost metric is automatically calculated using the configured reference bandwidth, supporting optimal path selection by reflecting real-time link capacities.

Metric calculation logic

When **`granularity_bandwidth`** is less than or equal to **`link_bandwidth`**:

- The effective bandwidth is determined by rounding down the link_bandwidth to the nearest multiple of granularity_bandwidth.
- The metric is calculated by dividing the reference_bandwidth by this effective bandwidth.
- Formula:

```
metric = reference_bandwidth / (link_bandwidth - (link_bandwidth %
granularity_bandwidth))
```

- When **granularity_bandwidth** is greater than link_bandwidth or not configured:
- The metric is calculated by dividing the reference_bandwidth directly by the actual link_bandwidth.
- Formula:

```
metric = reference_bandwidth / link_bandwidth
```

Key components

- **reference_bandwidth:** Defined in kilobits per second (kbps), serves as the baseline for scaling interface bandwidths into IS-IS metrics.
- **link_bandwidth**: The actual physical capacity of the interface, measured in kbps.
- **granularity_bandwidth**: Configured in kbps, determines the increment for rounding down the link_bandwidth before metric calculation.
- **% (Modulus Operator)** : Returns the remainder of the division; used to find how much link_bandwidth exceeds the nearest lower multiple of granularity_bandwidth.

Additional notes

- The final calculated metric is always an integer.
- Fractional results from the division are typically rounded up to the next whole number to ensure conservative metric assignment, preventing high-bandwidth links from receiving an undeservedly low metric.

How IS-IS tracks interface bandwidth changes

Outlines how IS-IS dynamically tracks interface bandwidth changes, including notification processing, auto-cost and metric evaluation, routing update generation, SPF runs, routing table updates, and TE component notification to optimize path selection and traffic engineering.

Summary

IS-IS dynamically monitors interface bandwidth fluctuations to ensure that routing metrics are current, optimizing path selection and informing Traffic Engineering components of changes.

The key mechanisms involved in the process are:

- IS-IS: The routing protocol responsible for monitoring, checking configurations, and initiating updates.
- Interface bandwidth change notification: An event signaling a change in the physical capacity of a link.
- Auto-cost reference bandwidth configuration: The setting that enables dynamic metric calculation based on bandwidth.
- Higher-priority metric configurations: Other metric settings, such as `metric-fallback` and explicit default metrics, that can override auto-cost.
- LSP PDU generation: The process of creating and sending routing updates.
- SPF run: The algorithm that re-calculates optimal paths in the routing table.
- Traffic Engineering (TE) component: A system that utilizes routing information for traffic management.

Workflow

These stages describe how IS-IS tracks interface bandwidth changes:

1. Receive notification: IS-IS receives an alert indicating a change in an interface's bandwidth.

2. Check auto-cost status: IS-IS determines if auto-cost reference bandwidth is enabled for the affected interface.
3. Evaluate metric priority: If auto-cost is enabled, IS-IS verifies that no other higher-priority metric configurations, such as `metric-fallback` or explicit default metrics, are active on the interface.
4. Recalculate metric: If auto-cost is enabled and has the highest priority, IS-IS recalculates the metric for the interface using the new bandwidth information.
5. Initiate routing updates: IS-IS triggers an LSP PDU generation to broadcast the updated metric.
6. Re-evaluate paths: IS-IS initiates an SPF run to re-calculate optimal routing paths based on the new metric.
7. Update routing table: The routing table is updated to reflect the new metric, influencing future routing decisions.
8. Notify TE component: IS-IS informs the Traffic Engineering (TE) component about any changes to TE tunnel metrics.

Result

Interface metrics are dynamically adjusted in response to bandwidth changes, leading to optimized routing decisions and accurate information for traffic engineering.

Guidelines for auto-cost reference bandwidth

Provides guidance on deploying and managing the auto-cost reference bandwidth feature, including configuration best practices, interface selection, metric prioritization, monitoring Bundle-Ether interfaces, operational considerations, and performance impact on routers.

Consider these guidelines when deploying and managing the auto-cost reference bandwidth feature:

- Configure auto-cost reference bandwidth only on active interfaces that have bandwidth information available.
- Avoid configuring auto-cost reference bandwidth on passive interfaces, as this feature does not apply to passive interfaces.
- Avoid configuring both an explicit metric under router address family or interface address family and auto-cost reference bandwidth on the same interface. If an explicit metric is configured, the auto-cost reference bandwidth setting is ignored.
- Avoid relying on the auto-cost reference bandwidth metric if you have configured other metric types, such as metric-fallback or explicit default metrics, on the interface. The system will always prioritize these explicitly configured metrics over auto-cost.
- Monitor Bundle-Ether interfaces for changes in link status. The effective bandwidth for these interfaces is determined by the total bandwidth of all active member links. Any change in link status, such as a link coming up or going down, automatically triggers a recalculation of the auto-cost metric.
- Operate the auto-cost reference bandwidth feature entirely within the router. This feature has no impact on interoperability with other IS-IS implementations or RFC conformance.
- Expect no significant increase in CPU usage or impact on router performance when using this feature.

Configure auto-cost reference bandwidth

Outlines how to configure auto-cost reference bandwidth for IS-IS globally or per interface, with steps to enable automatic metric provisioning based on link bandwidth and verification using show commands.

Enable automatic IS-IS metric provisioning based on link bandwidth.

This task involves configuring the auto-cost reference bandwidth globally for all interfaces or specifically for an individual interface. The interface-specific configuration overrides the global setting.

Before you begin

Ensure you have access to the router's command-line interface and are in privileged EXEC mode.

Follow these steps to configure the auto-cost reference bandwidth:

Procedure

1. Apply a global auto-cost reference bandwidth for all interfaces.

   Example

   ```
   Router(config)# router isis 1
   Router(config-isis)# address-family ipv4 unicast
   Router(config-isis-af)# auto-cost reference-bandwidth 30000000 granularity
   100000
   ```

   The range for **reference-bandwidth** keyword is from 1 through 4294967295 kilobit per second.

   The range for **granularity** keyword is from 1 through 4294967295 kilobit per second.

2. Configure auto-cost reference bandwidth for a specific interface.

   Example

   ```
   Router(config-isis-af)# interface Bundle-Ether1
   Router(config-isis)# address-family ipv4 unicast
   Router(config-isis-af)# auto-cost reference-bandwidth 20000000 granularity
   100000
   ```

3. Use the **show running-config router isis** to verify the configured auto cost metric.

   Example

   ```
   Router# show running-config router isis 1

   Sun Aug 31 19:10:22.072 UTC
   router isis 1
    is-type level-2-only
    net 49.0001.0000.0000.00.0600.00
    nsr
    distribute link-state
    nsf ietf
    log adjacency changes
    trace mode enhanced
    lsp-refresh-interval 65000
    max-lsp-lifetime 65535
    affinity-map RED bit-position 3
    address-family ipv4 unicast
     metric-style wide
    auto-cost reference-bandwidth 32000000000 granularity 100000000
   segment-routing mpls sr-prefer ! address-family ipv6 unicast metric-style
   wide auto-cost reference-bandwidth 32000000000 granularity 100000000 microloop
    avoidance segment-routing microloop avoidance rib-update-delay 65535
   segment-routing mpls sr-prefer
   !
   ```

4. Use the **show isis database detail** command to display detailed IS-IS database information. You can verify the calculated autocost metric from this show output. In this show output, the calculated auto-cost metric is 80.

Example

```
Router# show isis database detail

Sun Aug 31 19:10:24.442 UTC

IS-IS 1 (Level-1) Link State Database

LSPID        LSP Seq Num     LSP Checksum     LSP Holdtime/Rcvd     ATT/P/OL
r1.00-00     * 0x00000010    0xf36d           65507/*                0/0/0

Area Address: 49.0001
LSP MTU: 1492
NLPID: 0xcc
NLPID: 0x8e
MT: Standard (IPv4 Unicast)
MT: IPv6 Unicast 0/0/0
IP Address: 192.0.2.1
IPv6 Address: 2001:DB8::1
Hostname: r1
Router Cap: 192.0.2.1 D 5:0
Metric: 80 IS-Extended r2.00 Metric: 80 IS-Extended r2.00
Metric: 80 IS-Extended r3.00
Metric: 80 MT (IPv6 Unicast) IS-Extended r2.00
Metric: 80 MT (IPv6 Unicast) IS-Extended r2.00
Metric: 80 MT (IPv6 Unicast) IS-Extended r3.00
Metric: 80 IP-Extended 198.51.100.0/24
Metric: 80 IP-Extended 198.51.100.1/24
Metric: 80 IP-Extended 198.51.100.2/24
Metric: 80 IP-Extended 198.51.100.3/24
Metric: 10 IP-Extended 192.0.2.1/32
Metric: 80 MT (IPv6 Unicast) IPv6 2001:DB8:1::/64
Metric: 80 MT (IPv6 Unicast) IPv6 2001:DB8:2::/64
Metric: 80 MT (IPv6 Unicast) IPv6 2001:DB8:3::/64
Metric: 80 MT (IPv6 Unicast) IPv6 2001:DB8:4::/64
Metric: 10 MT (IPv6 Unicast) IPv6 2001:DB8::1/128
```

5. Use the **show isis interface** command to verify the interface bandwidth and the calculated auto-cost metric. In this output, the bandwidth is 400000000 and the calculated auto-cost metric is 80.

Example

```
Router# show isis interface

FourHundredGigE0/0/0/0 Enabled
Adjacency Formation: Enabled
Prefix Advertisement: Enabled
Bandwidth: 400000000 Circuit Type: level-1 (Configured: level-1-2)
Media Type: P2P
Circuit Number: 0
Extended Circuit Number: 35
Last IIH Received: 19:10:25 (5.36 sec ago), 1497 octets
Last IIH Sent: 19:10:28 (2.07 sec ago), 1497 octets
Sending next P2P IIH in: 7 s IPv4 BFD: Disabled
IPv6 BFD: Disabled Level-1
Adjacency Count: 1
Adjacency Flaps: 0
LSP Pacing Interval: 33 ms
PSNP Entry Queue Size: 0
Hello Interval: 10 s
```

```
Hello Multiplier: 3 CLNS I/O
Protocol State: Up
MTU: 1497
SNPA: 7837.b7ad.0400
Layer-2 Multicast: Listening
All ISSs: IPv4 Unicast Topology: Enabled
Adjacency Formation: Running
Prefix Advertisement: Running
Policy (L1/L2): -/-
Metric (L1/L2): 80/80
Metric fallback:
Bandwidth (L1/L2): Inactive/Inactive
Anomaly (L1/L2): Inactive/Inactive
```

Configure SPF interval for an IS-IS single-topology address family

Provides instructions for configuring the SPF interval within a single-topology IS-IS environment to control the timing of shortest path calculations.

Adjust SPF calculation timing to tune router performance in IS-IS single-topology configurations.

IS-IS uses Shortest Path First (SPF) calculations to compute routes for each topology. In single-topology mode, only IPv4 requires SPF; IPv6 borrows the IPv4 topology and does not require a separate SPF calculation. Tuning SPF intervals can improve router performance and prevent resource issues. Incremental SPF (ISPF) can be enabled separately, allowing more efficient recalculation after topology changes. The SPF interval determines the minimum time between SPF runs.

Before you begin

- Ensure that IS-IS routing is enabled on your device.
- Understand the impact of SPF interval values on network stability and CPU utilization.

Follow these steps to configure the SPF interval for an IS-IS single-topology address family:

Procedure

1. Enter IS-IS routing configuration mode for the specified instance, and define the address family for single-topology mode:

   Example

   ```
   Router# configure
   Router(config)# router isis isp
   Router(config-isis)# address-family ipv4 unicast
   ```

   You can change the level of routing to be performed by a particular routing instance by using the **is-type** router configuration command.

2. Set the desired SPF interval values as needed.

   Example

   ```
   Router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 30
   ```

- These values control the minimum time between consecutive SPF calculations.
- If set too low, excessive SPF runs may waste CPU resources in unstable networks.
- If set too high, topology change reactions may be delayed, increasing the risk of lost packets.
- SPF interval does not apply to incremental SPF (ISPF).

3. (Optional) Enable incremental SPF calculation.

   Example

   ```
   Router(config-isis-af)# ispf
   ```

   This enables ISPF, which recalculates only portions of the topology affected by a change.

4. (Optional) Verify the SPF settings and calculation frequency.

   Example

   ```
   Router# show isis instance 1 spf-log ipv4
   ```

   Displays how often and why the router has run a full SPF calculation.

IS-IS SPF interval delay

Explains IS-IS SPF interval delay concepts, outlining the effects of SPF delay on IS-IS operations and supplying procedures for configuring IETF-compliant SPF calculation intervals to manage route computation frequency.

IS-IS SPF interval delay is a protocol timing mechanism that

- defines a standard algorithm to postpone SPF computations in IS-IS
- reduces computational load and churn on IGP nodes when multiple, temporally close network events trigger several SPF calculations, and
- lowers the probability and the duration of transient forwarding loops during native IS-IS convergence.

Table 4: Feature History Table

Feature Name	Release	Description
Setting SPF interval in IS-IS to postpone the IS-IS SPF computations	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8712-MOD-M • 8011-4G24Y4H-I

Feature Name	Release	Description
Setting SPF interval in IS-IS to postpone the IS-IS SPF computations	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM
Setting SPF interval in IS-IS to postpone the IS-IS SPF computations	Release 7.7.1	You can now define a standard algorithm to postpone the IS-IS SPF computations by setting an SPF interval. This reduces the computational load and churn on IGP nodes when multiple temporally close network events trigger multiple SPF computations. This algorithm also reduces the probability and the duration of transient forwarding loops during native IS-IS convergence when the protocol reacts to multiple temporally close events. This feature complies with RFC 8405 . This feature introduces the **spf-interval ietf** command.

Additional details

The algorithm specified in RFC 8405 provides a standardized approach to temporarily postpone IS-IS SPF computation. Applying this feature is optional but can be beneficial in environments with high protocol churn.

Configure IETF-compliant SPF calculation intervals

Configures SPF calculation intervals in IS-IS routing to postpone route computations and improve network stability. This process helps prevent excessive CPU utilization in environments experiencing frequent network changes.

Postponing SPF calculations by configuring intervals helps prevent excessive CPU utilization and stabilizes routing updates in environments experiencing frequent network changes. The IETF algorithm is recommended for most deployments, as per RFC 8405.

Procedure

1. Enable IS-IS routing for the target instance and enter router configuration mode.

Example

```
Router(config)# router isis isp
```

2. Specify the address family (IPv4 or IPv6) for unicast routing and enter address family configuration mode.

Example

```
Router(config-isis)# address-family ipv4 unicast
```

3. Set the SPF calculation interval type to IETF.

Example

```
Router(config-isis-af)# spf-interval ietf
```

4. Adjust delay values if needed for your environment using parameters like short-wait, initial-wait, and so on.

Example

```
Router(config-isis-af)# spf-interval ietf short-wait 500
```

5. Verify the new SPF interval configuration.

Example

```
Router# show run router isis
router isis 1
 net 49.0001.0000.0000.0100.00
 log adjacency changes
 address-family ipv4 unicast
  metric-style wide
  spf-interval ietf
 !
 address-family ipv6 unicast
  metric-style wide
  spf-interval ietf
 !
```

6. Monitor SPF calculation logs for expected delays and triggers.

Example

```
Router# show isis ipv4 spf-log last 5 detail
IS-IS 1 Level 2 IPv4 Unicast Route Calculation Log
                    Time Total Trig.
Timestamp    Type    (ms) Nodes Count First Trigger LSP    Triggers
------------ ----- ----- ----- ----- --------------------
----------------------
--- Wed Mar 16 2022 ---
15:31:49.763  FSPF     1     6     3         tb5-r4.00-00 LINKBAD PREFIXBAD
```

```
Delay:                  101ms (since first trigger)
                        261177ms (since end of last calculation)
Trigger Link:           tb5-r2.00
Trigger Prefix:         34.1.24.0/24
New LSP Arrivals:       0
SR uloop:               No
Next Wait Interval:     200ms
RIB Batches:            1 (0 critical, 0 high, 0 medium, 1 low)
Timings (ms):           +--Total--+
                         Real   CPU
  SPT Calculation:          1     1
  Route Update:             0     0
                        ----- -----
```

It is recommended to use the default delay values, which are listed in Syntax description . These default parameters are suggetsed by RFC 8405 . These should be appropriate for most networks.

However, you can configure different values if required.

8 Interface Management and Flooding Control

Topics:

- IS-IS interface types
- IS-IS adjacency stagger
- IS-IS static neighbors
- ISIS link group
- Interface flood blocks
- Mesh groups
- LSP flood limits
- Layer 2 ACLs on physical and subinterfaces

Outlines IS-IS interface types, adjacency configurations, neighbor setup, link grouping, flood control mechanisms, mesh groups, LSP flooding controls, and Layer 2 ACL implementation to enhance interface management and mitigate network flooding.

IS-IS interface types

Introduces IS-IS interface types, explaining their characteristics and roles within IS-IS routing environments to facilitate effective interface selection and management.

IS-IS interfaces can be configured in four distinct types, each controlling how route information is advertised and adjacencies are formed:

- Active – Advertises connected prefixes and forms adjacencies. This is the default mode for interfaces.
- Passive – Advertises connected prefixes but does not form adjacencies. Configured with the **passive** command. Used for important prefixes such as loopback addresses. For many prefixes, prefer route redistribution with an appropriate policy.
- Suppressed – Does not advertise connected prefixes but forms adjacencies. Configured with the **suppress** command.
- Shutdown – Neither advertises connected prefixes nor forms adjacencies. Configured with the **shutdown** command to disable the interface without removing IS-IS settings.

IS-IS adjacency stagger

Introduces IS-IS adjacency stagger, highlighting operational restrictions and providing step-by-step instructions for configuring both adjacency stagger and general IS-IS adjacencies to optimize protocol convergence and stability.

IS-IS adjacency stagger is a routing protocol feature that

- limits the number of adjacencies established at the same time after a process restart or reload
- enables you to control the initial and subsequent number of neighbors that can form adjacencies, and
- reduces instability caused by excessive packet exchanges and database synchronization during network recovery.

Additional details

When a router restarts or reloads, it must re-establish adjacencies and synchronize its routing database with each neighbor. These operations generate significant processing and bandwidth demands. If too many adjacencies form simultaneously, this can cause packet congestion, dropped packets, and network instability. IS-IS adjacency stagger helps prevent these issues by limiting the rate at which adjacencies are established.

Restrictions for IS-IS adjacency stagger

Lists the limitations of the IS-IS adjacency stagger feature, including unsupported interface types and protocols.

These restrictions apply to the IS-IS adjacency stagger feature:

- IS-IS adjacency stagger is only supported on point-to-point interfaces. It is not supported on LAN interfaces.
- IS-IS adjacency stagger does not support non-stop forwarding (NSF) mechanisms.

Configure IS-IS adjacency stagger

Configure IS-IS adjacency stagger to limit how many adjacencies form at the same time after restart or reload events.

Limit the number of IS-IS adjacencies that can form at the same time after device restart or reload, helping control network convergence and resource usage.

Before you begin

- Ensure that the interface is a point-to-point IS-IS interface.
- Confirm that NSF is not required for your workflow.

Procedure

Enter IS-IS router configuration mode and configure the adjacency stagger values.

Example

```
Router#config
Router(config)#router isis 1
Router(config-isis)#adjacency stagger 2 3
Router(config-isis)#commit
```

Configure IS-IS adjacencies

Guides you through configuring, tuning, and verifying IS-IS adjacencies, including enabling state change logging, adjusting hello parameters, securing adjacencies with authentication, and using show commands for interface and neighbor verification.

Configure, tune, and verify IS-IS adjacencies, including logging state changes, adjusting hello interval and multiplier, and securing adjacencies with authentication.

Adjusting IS-IS adjacency parameters can improve network stability, especially on congested links. Enabling logging helps you monitor state changes. Authentication prevents unwanted neighbor formation. Parameters generally apply to the interface for both Level 1 and Level 2 adjacencies, but Level modifiers are not used on point-to-point interfaces.

Before you begin

- Ensure that IS-IS routing is enabled on your device.
- Identify interfaces where adjacency tuning is needed.

Procedure

1. Enable logging of adjacency state changes.

   Example

   ```
   Router# configure
   Router(config)# router isis isp
   Router(config-isis)# log adjacency changes
   ```

   This setting generates log messages when IS-IS adjacencies change state.

2. Enter interface configuration mode for the targeted interface, and configure hello padding for IS-IS hello PDUs.

   Example

   ```
   Router(config-isis)# interface HundredGigE 0/1/0/3
   Router(config-isis-if)# hello-padding sometimes
   ```

Hello padding applies to only this interface and not to all interfaces.

3. Set the hello interval time in seconds, and the hello multiplier.

   Example

   ```
   Router(config-isis-if)# hello-interval 6
   Router(config-isis-if)# hello-multiplier 10
   ```

   The hello multiplier specifies how many IS-IS hello packets a neighbor must miss before the router declares the adjacency down.

4. Set hello authentication to prevent unauthorized adjacency formation.

   Example

   ```
   Router(config-isis-if)# hello-password text clear mypassword
   ```

   This setting ensures that the system includes authentication in hello packets and requires successful authentication of hello packets from the neighbor before establishing an adjacency.

5. (Optional) Use the **show isis instance isp adjacency** command to display IS-IS adjacencies.

   Example

   ```
   Router# show isis instance isp adjacency
   ```

6. (Optional) Use the **show isis adjacency-log** command to view a log of the most recent adjacency state changes.

   Example

   ```
   Router# show isis adjacency-log
   ```

7. (Optional) Use the **show isis interface** command to view details about the IS-IS interface.

   Example

   ```
   Router# show isis interface HundredGigE 0/1/0/1 brief
   ```

8. (Optional) Use the **show isis neighbors** command to view IS-IS neighbor information.

   Example

   ```
   Router# show isis neighbors summary
   ```

IS-IS static neighbors

Explains IS-IS static neighbor functionality and guides users through configuration steps to establish static neighbor relationships for improved topology management.

A IS-IS static neighbor is a network management feature that

- allows the advertisement of an IS-IS link without forming an actual IS-IS adjacency
- provides necessary link topology for controllers where IS-IS is not active, and
- ensures the advertised link is identical to regular IS-IS links.

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
IS-IS static neighbor	Release 25.4.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on the Cisco 8011-32Y8L2H2FH routers.
IS-IS static neighbor	Release 25.2.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]);Centralized Systems (8600 [ASIC: Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]) IS-IS static neighbor allows the advertisement of an IS-IS link without forming an actual IS-IS adjacency. This feature is useful when a link is required in the topology for the controller, but IS-IS is not actively running on the link. This feature introduces these changes: CLI: **router isis static-neighbor**

Additional notes

IS-IS static neighbor enables network controllers to recognize a link in the topology even if IS-IS is not actively running on that link. This feature is especially useful when a controller requires awareness of the link, but no regular IS-IS adjacency is formed.

When IS-IS static neighbor is configured, protocol operation on the interface stops, which prevents the formation of dynamic adjacencies. However, local IP subnets continue to be advertised alongside static neighbors, depending on the configuration. Use the parameters (system-id, remote address, affinity, and log advertisement changes) to define the neighbors accurately.

Usage Guidelines

Static IS-IS links are advertised when these conditions are met:

- IS-IS interface is up, including IP/IPv6 protocols and Connectionless Network Service (CLNS) routing IO
- static IS-IS is configured, and
- point-to-point (p2p) IS-IS is configured.

Configure an IS-IS static neighbor

Guides you through configuring an IS-IS static neighbor, including defining neighbor properties, adjusting adjacency and advertisement settings, enabling logging, verifying operational status, and monitoring interface state.

Establish an IS-IS static neighbor relationship to control adjacency parameters and advertisement policies on a router.

Before you begin

Collect these details:

- IS-IS System-id of the static neighbor
- Remote IPv4 or IPv6 addresses
- Affinity link to advertise
- Determine if you need to enable logging for static neighbor advertisement changes

Procedure

1. Configure the IS-IS routing process, and select the interface where you want to configure the static neighbor.

   Example

   ```
   Router# configure
   Router(config)# router isis 1
   Router(config-isis)# interface GigabitEthernet0/2/0/7
   ```

2. Enter the static-neighbor sub-mode, and specify the static neighbor properties.

   Example

   ```
   Router(config-isis-if)# static-neighbor
   ```

   - Set the system-id of the static neighbor.

     ```
     Router(config-isis-static-nbr)# system-id 0001.0002.0007
     ```

   - Specify the remote IPv4 address.

     ```
     Router(config-isis-static-nbr)# remote ipv4 address 10.2.2.2
     ```

   - (Optional) Specify the affinity link.

     ```
     Router(config-isis-static-nbr)# affinity RED
     ```

   - (Optional) Enable logging for advertisement changes.

     ```
     Router(config-isis-static-nbr)# log-advertisement-changes
     ```

3. Use the **show isis static-neighbor** command to verify the IS-IS static neighbor's operational status.

Example

```
Router# show isis static-neighbor IS-IS 1

IS-IS 1 static neighbors:
System Id          Interface        State
0001.0002.0007     Gi0/2/0/3        Up
Remote IPv4 Address: 2.2.2.2
Remote IPv6 Address: 7::2
Affinity: RED
[additional status details and syslog messages]
```

4. Use the **show isis interface** command to display IS-IS interface information, including adjacency formation and protocol state.

Example

```
Router# show isis interface Gi0/2/0/3

Gi0/2/0/3             Enabled
Adjacency Formation:      Disabled (Passive due to static neighbor
configuration)
Prefix Advertisement: Enabled
Protocol State: Up
Forwarding Address(es): Unknown (Intf passive due to static neighbor
configuration)
[other interface and topology details]
// ...snip...
```

ISIS link group

Describes ISIS link group features, offering instructions for configuring link group profiles and associating interfaces with link groups to improve network resilience and control.

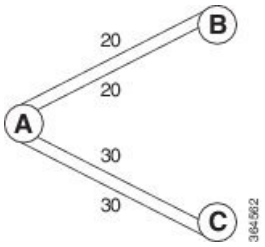
A ISIS link group is a routing feature that

- enables you to define a group of links for metric management
- adjusts the ISIS metric for remaining active links when the number of active links drops below a configured minimum, and
- restores the original metric when the link group is reestablished.

Additional details about ISIS link groups

The ISIS link group feature lets you optimize network routing by associating a set of links as a group. When the number of active IS-IS adjacencies within the group falls below the minimum threshold, a metric offset is applied to the remaining links. This dynamic adjustment influences route selection, helping traffic avoid degraded paths. Once links are restored to the threshold, ISIS removes the offset and returns metrics to their original values.

ISIS link group metric adjustment



Router A connects to both Router B and Router C. Between A and B, there are two layer 3 links with an ISIS metric of 20 each; between A and C, there are links with a metric of 30. Normally, traffic prefers the lower metric path (A→B). If one link between A and B fails and ISIS link groups are not configured, traffic still goes through B. With ISIS link group settings—using an offset of 20 and minimum-members of 2—when a link fails the metric rises to 40 (original 20 + offset 20), and traffic reroutes through C. Similar logic can apply for the path between A and C.

Configure link group profile

Guides you through defining and activating an IS-IS link group profile, setting metric offsets and membership requirements, optionally configuring revert-members, and verifying link group status using show commands.

Define and activate an IS-IS link group profile that specifies metric offsets and member requirements, supporting robust routing configurations.

A link-group profile is only active when both `minimum-members` and `metric-offset` values are set. If revert-members is not configured, the router defaults to using the minimum-members value for reverting the link group state.

Procedure

1. Enter IS-IS configuration mode, define the link group and set the metric offset.

   Example

   ```
   Router# configure
   Router(config)# router isis purple
   Router(config-isis)# link-group foo metric-offset 100
   ```

   The supported range for metric-offset is 1-16777214. Practical default for narrow metrics is 1-63).

2. Set the minimum number of active members required for the link group.

   Example

   ```
   Router(config-isis)# link-group foo minimum-members 2
   ```

   The supported range is 2 to 64 members.

3. (Optional) Configure the revert-members value if you want the router to revert the link group only after a specific member threshold is reached.

   Example

   ```
   Router(config-isis)# link-group foo revert-members 4
   ```

   The supported range is 2 to 64 members.

4. Use the **show isis interface** command to review the link group state and metric offset.

   Example

   ```
   Router# show isis interface gig 0/0/0/1

   Link-Group (L1/L2): Configured/Configured
   Metric-Offset (L1/L2): 100/100
   ```

5. Use the **show isis lsp** command to verify the updated metric value in the advertised LSP data.

Example: Link group profile configuration

This example shows a link-group profile configuration.

```
router isis 1
 is-type level-2-only
 net 49.1111.0000.0000.0006.00
 link-group foo
  metric-offset 100
  revert-members 4
  minimum-members 2
 address-family ipv4 unicast
  metric-style wide
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   link-group foo
```

Associate an IS-IS interface address family with a link group

Guides you through associating an IS-IS interface address family with a preconfigured link group, including configuration steps, level restrictions, prerequisites, and verification procedures.

Associate an IS-IS interface address family with a preconfigured link group to facilitate IS-IS protocol enhancements.

Each IS-IS interface and address family can specify only one link group association. By default, the association applies to both IS-IS levels unless restricted.

Before you begin

IS-IS routing must be enabled and the required link group must exist.

Procedure

---

1. Associate the interface address family with the link group:

   Example

   ```
   Router# configure
   Router(config)# router isis purple
   Router(config-isis)# interface GigabitEthernet 0/1/0/3
   Router(config-isis-if)# address-family ipv4 unicast
   Router(config-isis-if-af)# link-group access level 1
   ```

   Specify "level" only if you want to restrict the association to a single IS-IS level.

2. (Optional) Use the **show isis interface** command to verify the link group association.

Example

```
Router# show isis interface
```

Interface flood blocks

Details interface flood block mechanisms, explaining how flood blocking on specific interfaces helps manage network flooding and maintain stable routing environments.

A flood block is a network configuration feature that

- allows you to block specific interfaces from flooding IS-IS LSPs
- lets the remaining interfaces operate normally for flooding, and
- enables fine-tuning of the IS-IS flooding topology for reliability and scalability.

Implementation, operation, and maintenance considerations

Interface flood blocks prevent designated interfaces from participating in LSP flooding, while others continue normal operation. This approach is straightforward to implement but may be harder to maintain and more prone to mistakes than mesh groups over time.

Balancing IS-IS flooding topology for stability and scalability

The IS-IS flooding topology is adjusted rather than severely limited. Blocking too many interfaces can create network instability during failures, while blocking too few may not reach the desired scalability.

Improving robustness with CSNP

To improve robustness when all non-blocked interfaces fail, use the **csnp-interval** command in interface configuration mode to force periodic complete sequence number PDUs (CSNPs) on blocked point-to-point links. Periodic CSNPs enable the network to resynchronize.

Mesh groups

Describes mesh group concepts used to optimize protocol flooding and enhance scalability in complex network topologies.

A mesh group is a router interface set that

- limits flooding by restricting Link State Packet propagation,
- ensures network resilience by assuming dense connectivity between routers, and
- optimizes IS-IS behavior for stable operations.

Routers reachable via mesh group interfaces are assumed to be closely interconnected, with at least one redundant link connecting each router to every other router. This arrangement allows multiple links to fail without isolating any routers, maintaining network integrity.

In normal flooding, a new LSP received on an interface is flooded out all other interfaces. With mesh groups, when a new LSP is received over a mesh group interface, it is not flooded over other mesh group interfaces, reducing unnecessary traffic.

LSP flood limits

Explains LSP flood limitation strategies for IS-IS, including configuration procedures for LSP flooding parameters and the use of LSP fast flood control mechanisms to prevent excessive network load.

A flood limit is an IS-IS network design technique that

- restricts the flooding of link-state packets (LSPs) on specific interfaces or mesh groups
- improves network scalability in highly redundant or meshy topologies, and
- may reduce network reliability in the event of multiple failures.

Use cases for limiting LSP flooding

Limiting link-state packets (LSPs) may be desirable in certain “meshy” network topologies, such as highly redundant networks with a fully meshed set of point-to-point links over a nonbroadcast multiaccess (NBMA) transport. In these networks, full LSP flooding can limit network scalability.

Techniques for restricting LSP flooding

One way to restrict the size of the flooding domain is to introduce hierarchy with multiple Level 1 areas and a Level 2 area. Two other techniques can also be used, instead of or in addition to hierarchy:

- Blocking flooding on specific interfaces
- Configuring mesh groups

Consequences of flood limiting

Both blocking and mesh group techniques operate by restricting the flooding of LSPs. While scalability of the network is improved, network reliability may be reduced. A series of failures can prevent LSPs from being flooded throughout the network, even if links exist that would allow flooding if blocking or mesh groups had not restricted their use. Under these conditions, link-state databases of routers may become unsynchronized, and consequences such as persistent forwarding loops can result.

Recommendation for network design

Blocking or mesh groups should be used only if specifically required, and then only after careful network design.

LSP fast flood control mechanisms for IS-IS networks

Introduces LSP fast flood control mechanisms for IS-IS networks, outlining configurable parameters, dynamic flooding rate optimization, supported platforms, and enhancements for rapid and stable topology distribution prior to configuration or verification.

A fast flood control mechanism is an IS-IS protocol enhancement that

- increases the transmission rate of Link State Packets (LSPs) across IS-IS domains,
- dynamically adjusts flooding rates based on neighbor acknowledgment processing, and
- enables configurable parameters for flexible activation and rate control.

Table 6: Feature History Table

Feature Name	Release Name	Description
LSP Fast-Flooding on IS-IS Networks	Release 26.1.1	Introduced in this release on: Centralized Systems (8400 [ASIC: K100])(select variants only*) * This feature is now supported on Cisco 8404-SYS-D routers.
LSP Fast-Flooding on IS-IS Networks	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
LSP Fast-Flooding on IS-IS Networks	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8712-MOD-M • 8011-4G24Y4H-I

Feature Name	Release Name	Description
LSP Fast-Flooding on IS-IS Networks	Release 24.3.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) You can now accelerate the rate at which Link State Packets (LSPs) are distributed across an IS-IS network. Faster LSP distribution means faster network convergence. This faster convergence ensures that the most accurate topology information is quickly available across all routers on the network, reducing the chances of routing loops or misrouting. The feature introduces these changes: CLI: • **lsp-fast-flooding** • **max-lsp-tx** • **psnp-interval** • **remote-psnp-delay** **YANG Data Model:** Cisco-IOS-XR-um-router-isis-cfg (refer to GitHub and YANG Data Models Navigator) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Flexibility and activation

Unlike traditional IS-IS flooding, which operates with a fixed timer, fast-flooding mechanisms provide greater adaptability and control for network administrators. Adjustable parameters such as the local Partial Sequence Number PDU (PSNP) interval and maximum flooding rate can be set to meet specific network requirements. Fast-flooding is inactive by default, allowing administrators to enable it selectively where improved LSP flooding benefits the network.

Dynamic flooding rate optimization

Fast-flooding incorporates real-time adjustment of LSP transmission rates, using persistent monitoring of PSNP acknowledgments from neighbors. If acknowledgments are delayed, the flooding rate automatically decreases to prevent

overloading neighbor routers. If acknowledgments are received promptly, the mechanism increases the flooding rate up to the configured maximum, enhancing speed of topology distribution. This adaptive approach ensures optimal balance between rapid network convergence and stability.

For instance, configuring fast-flooding in a large IS-IS network enables LSP bursts at a default rate of 1000 LSPs per second, far surpassing traditional rates. The feature dynamically adapts to real-time conditions, preventing congestion while accelerating topology updates.

Configure LSP flooding parameters for IS-IS

Guides you through configuring IS-IS LSP flooding parameters globally and per interface to optimize scalability and performance, including tuning refresh intervals, lifetimes, flooding thresholds, mesh-group settings, and verifying operational status using show commands.

Tune IS-IS LSP flooding parameters to limit excessive flooding, reduce CPU load, and maintain stable IS-IS operations.

Flooding of IS-IS LSPs, if not controlled, can affect router scalability and lead to unnecessary resource consumption. Adjusting parameters lets you optimize IS-IS operation globally or per interface.

Before you begin

- Ensure IS-IS is already configured on your router.
- Decide whether you want to apply flooding parameters globally or on specific interfaces.
- Review network requirements for refresh intervals, LSP lifetimes, and flooding behavior.

Procedure

1. Enter IS-IS router configuration mode, and configure LSP flooding parameters globally.
   a) Set the LSP refresh interval (between LSP regenerations).

      Example

      ```
      Router# configure
      Router(config)# router isis isp
      Router(config-isis)# lsp-refresh-interval 10800 level 2
      ```

      The refresh interval should always be set lower than the **max-lsp-lifetime**.

   b) Configure the LSP check interval (database checksum validation).

      Example

      ```
      Router(config-isis)# lsp-check-interval 240
      ```

      This operation uses significant CPU resources and should be configured to run infrequently.

   c) Configure LSP generation interval to reduce flooding during instability:

      Example

      ```
      Router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5
      ```

During extended periods of network instability, repeated LSP recalculations can significantly increase CPU load on the local router. Additionally, flooding these recalculated LSPs to other intermediate systems increases network traffic and causes other routers to spend more time recalculating routes.

d) Set LSP MTU (maximum transmission unit).

Example

```
Router(config-isis)# lsp-mtu 1300
```

e) Set maximum LSP lifetime.

Example

```
Router(config-isis)# max-lsp-lifetime 11000
```

This is the length of time an LSP remains in a neighbor's database unless it is regenerated or refreshed.

f) Configure purge of LSPs with checksum errors.

Example

```
Router(config-isis)# ignore-lsp-errors disable
```

2. For interface-specific settings, enter interface configuration mode.

Example

```
Router(config-isis)# interface HundredGigE 0/1/0/3
```

3. Configure flooding thresholds and intervals per interface:

a) Set LSP interval for interface.

Example

```
Router(config-isis-if)# lsp-interval 100
```

b) Configure CSNP interval (broadcast CSNP packet frequency).

Example

```
Router(config-isis-if)# csnp-interval 30 level 1
```

> **Tip**
>
> - Frequent CSNPs increase router processing load.
> - Less frequent CSNPs may delay detection of differences.

c) Set retransmit interval (wait time for acknowledgment).

Example

```
Router(config-isis-if)# retransmit-interval 60
```

d) Set retransmit throttle interval (time between retransmissions).

Example

```
Router(config-isis-if)# retransmit-throttle-interval 1000
```

This time is usually greater than or equal to the **lsp-interval** command time because the reason for lost LSPs may be that a neighboring router is busy. A longer interval gives the neighbor more time to receive transmissions.

e) Optimize flooding in highly meshed networks.

Example

```
Router(config-isis-if)# mesh-group blocked
```

4. (Optional) Use the **show isis interface** command to check the IS-IS interface state.

Example

```
Router# show isis interface HundredGigE 0/1/0/1 brief
```

5. (Optional) Use the **show isis database** command to check the IS-IS LSP database.

Example

```
Router# show isis database level 1
```

6. (Optional) Use the **show isis lsp-log** command to check the LSP log information.

Example

```
Router# show isis lsp-log
```

7. (Optional) Use the **show isis database-log** command to check the IS-IS database log information.

Example

```
Router# show isis database-log level 1
```

Layer 2 ACLs on physical and subinterfaces

Describes implementing Layer 2 ACLs on physical and subinterfaces, noting configuration restrictions and providing instructions for enabling ISIS PDU rejection to enhance access security and traffic filtering.

Layer 2 ACL is a security configuration that

- binds to Layer 2 main interfaces, subinterfaces, or bundle interfaces,
- enables traffic filtering of ISIS packets on the ethernet port shared by these interfaces, and
- ensures ISIS traffic is dropped regardless of whether the ACL is applied to the main or subinterface, since all share the port.

Table 7: Feature History Table

Feature Name	Release Information	Feature Description
Support for a Configurable Knob to Reject ISIS PDU on Layer 2 Interfaces	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on the Cisco 8712-MOD-M routers.
Support for a Configurable Knob to Reject ISIS PDU on Layer 2 Interfaces	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Feature Name	Release Information	Feature Description
Support for a Configurable Knob to Reject ISIS PDU on Layer 2 Interfaces	Release 7.3.1	This feature enables you to use Layer 2 ACL to drop ISIS packets from certain ISIS destination MAC addresses. Dropping ISIS packets allows you to isolate a particular node from ISIS domain. This feature enables you to utilize the network bandwidth efficiently. This feature introduces the **ethernet-services access-list isis-drop-all-l2-pdus** command

Configuring a Layer 2 ACL on either the main interface or subinterface results in ISIS packet drops for both, because the ACL operates on their shared ethernet port.

Applying a Layer 2 ACL to a Layer 3-only interface does not affect ISIS traffic on Layer 2 interfaces.

Restrictions for configuring Layer 2 ACLs on physical and subinterfaces

Outlines the limitations of configuring Layer 2 ACLs on physical and subinterfaces, including ingress-only support, lack of per-interface statistics and ACL modification, restrictions on ACE changes, and specific ISIS destination MAC address matching requirements.

- This feature is supported only in the ingress direction.
- Per-interface statistics are not supported.
- Layer 2 ACL modification is not supported.
- Only remarks can be added, updated, or modified.
- Any insertion or modification of Layer ACL access control entries (ACE) is rejected. However, deletion of an ACE is accepted because it cannot be blocked.
- If you delete an ACE from an attached Layer 2 ACL, detach the Layer 2 ACL from all interfaces, modify the Layer 2 ACL, and reattach it to the interfaces to recover the deleted ACE.
- Layer 2 ACL supports matching only on ISIS destination MAC addresses. It does not support any other Layer 2 fields, such as source MAC (srcMAC), PCP, etc. Configure only one of the ISIS destination MAC addresses to drop ISIS packets. Non-ISIS destination MAC configuration is rejected.
- Hardware drops ISIS packets when you configure one of these destination MAC addresses:
  - 01:80:c2:00:00:14
  - 01:80:c2:00:00:15
  - 09:00:2b:00:00:04
  - 09:00:2b:00:00:05
  - 01:00:5e:90:00:02
  - 01:00:5e:90:00:03

Configure ISIS PDU rejection on Layer 2 interfaces

Configures Layer 2 ACLs to reject ISIS PDUs on Layer 2 interfaces for node isolation and bandwidth optimization. This process involves defining the ACL and applying it to specific interface types.

Set up Layer 2 ACLs and apply them to interfaces to drop ISIS packets based on destination MAC addresses.

This task enables you to isolate specific nodes from an ISIS domain and optimize usage of network bandwidth by filtering ISIS PDUs at Layer 2 interfaces.

Procedure

1. Create a Layer 2 ACL to deny ISIS PDUs and permit other traffic.

   Example

   ```
   Router# configure
   Router(config)# ethernet-services access-list isis-drop-all-l2-pdus
   Router(config)# 5 remark Drain ISIS between two routers
   Router(config)# 20 deny any host 0180.c200.0015
   Router(config)# 200 permit any any
   ```

2. Apply the ACL to the relevant interfaces:
   a) For a bundle interface:

      Example

      ```
      Router# configure
      Router(config)# interface Bundle-Ether 100 l2transport
      Router(config-if)# mtu 2000
      Router(config-if)# ethernet-services access-group isis-drop-all-l2-pdus
      ingress
      ```

   b) For a bundle subinterface:

      Example

      ```
      Router# configure
      Router(config)# interface Bundle-Ether101.101 l2transport
      Router(config-if)# encapsulation dot1q 101
      Router(config-if)# rewrite ingress tag pop 1 symmetric
      Router(config-if)# mtu 2000
      Router(config-if)# ethernet-services access-group isis-drop-all-l2-pdus
      ingress
      ```

   c) For a physical interface:

      Example

      ```
      Router# configure
      Router(config)# interface hundredGigE 0/0/0/0 l2transport
      Router(config-if)# mtu 2000
      Router(config-if)# ethernet-services access-group isis-drop-all-l2-pdus
      ingress
      ```

   d) For a physical subinterface:

      Example

      ```
      Router# configure
      Router(config)# interface hundredGigE 0/3/0/1.100 l2transport
      Router(config-if)# encapsulation dot1q 101
      ```

```
Router(config-if)# rewrite ingress tag pop 1 symmetric
Router(config-if)# mtu 2000
Router(config-if)# ethernet-services access-group isis-drop-all-l2-pdus
ingress
```

3. Use the **show access-lists ethernet-services** command to verify the ACL operation.

Example

```
Router# show access-lists ethernet-services l2 hardware ingress location
Thu Jan 21 04:22:12.667 UTC
ethernet-services access-list l2
20 deny any host 0180.c200.0014 (1243345)
200 permit any any
```

Example

```
Router# show access-lists ethernet-services
Sun Feb 14 12:52:09.539 PST
ethernet-services access-list isis-drop-all-l2-pdus
5 remark Drain ISIS between two routers.
20 deny any host 0180.c200.0015
200 permit any any
```

Example

```
Router# show access-lists ethernet-services isis-drop-all-l2-pdus
hardware ingress location 0/0/CPU0
Sun Feb 14 12:52:39.620 PST
ethernet-services access-list isis-drop-all-l2-pdus
20 deny any host 0180.c200.0015
200 permit any any
```

Example

```
Router# show access-lists ethernet-services isis-drop-all-l2-pdus
hardware ingress detail location 0/0/CPU0
Sun Feb 14 12:52:47.962 PST

isis-drop-all-l2-pdus Details:
Sequence Number: 20
NPU ID: 1
Number of DPA Entries: 1
ACL ID: 1
ACE Action: DENY
ACE Logging: DISABLED
Set TTL value: 0
Hit Packet Count: 0
Source MAC: 0000:0000:0000
Source MAC Mask: 0000:0000:0000
Destination MAC: 0180:C200:0015
Destination MAC Mask: FFFF:FFFF:FFFF
DPA Entry: 1
        Entry Index: 0
        DPA Handle: 0x93C84100
Sequence Number: 200
NPU ID: 1
```

```
Number of DPA Entries: 1
ACL ID: 1
ACE Action: PERMIT
ACE Logging: DISABLED
Set TTL value: 0
Source MAC: 0000:0000:0000
Source MAC Mask: 0000:0000:0000
Destination MAC: 0000:0000:0000
Destination MAC Mask: 0000:0000:0000
DPA Entry: 1
        Entry Index: 0
        DPA Handle: 0x93C84278
```

9 LSP Lifetimes

Topics:

- Maximum LSP lifetime and refresh interval
- Minimum LSP lifetime thresholds

Explains LSP lifetime parameters, including maximum and minimum thresholds, as well as refresh intervals, to guide configuration and management of LSPs in network environments.

Maximum LSP lifetime and refresh interval

Describes maximum LSP lifetime values and refresh interval settings, clarifying how to configure and manage these parameters to ensure proper LSP operation and stability.

A maximum LSP lifetime and refresh interval are IS-IS protocol settings that

- determine how frequently the router refreshes link-state packets (LSPs) in the IS-IS database
- define how long LSPs remain valid in a device's database before removal, and
- ensure routing stability by periodically refreshing LSPs and deleting outdated entries.

Default values

- By default, the router sends a periodic LSP refresh every 15 minutes.
- LSPs remain in the database for 20 minutes by default.
- If LSPs are not refreshed within their lifetime, they are deleted.

Configuration options

You can configure both the LSP refresh interval and the maximum lifetime.

Setting relationships

The refresh interval should always be less than the maximum lifetime; otherwise, LSPs might be deleted before they are refreshed.

Automatic adjustment

If no refresh interval is configured, the software automatically adjusts the interval as necessary to prevent premature LSP expiration.

Minimum LSP lifetime thresholds

Explains minimum LSP lifetime thresholds in IS-IS, including their role in preventing premature purging, configuring maximum lifetimes, handling Zero Age Lifetime behavior, and identifying corrupted lifetime events.

A minimum LSP lifetime threshold is an IS-IS protocol safeguard that

- prevents premature purging and unnecessary flooding of Link State Protocol Data Units (LSPs)
- ensures accurate handling of potentially corrupted Remaining Lifetime fields during LSP flooding events, and
- allows IS-IS to reset the Remaining Lifetime value of received LSPs to the maximum LSP lifetime for reliable network operation.

Maximum LSP lifetime configuration

By default, the maximum LSP lifetime is set to 1200 seconds. You can configure a different value using the **max-lsp-lifetime** *seconds* command. This rule ensures that, regardless of the Remaining Lifetime received, only the LSP originator can purge an LSP from the database before the maximum LSP lifetime elapses.

Zero Age Lifetime behavior

If the remaining lifetime for an LSP reaches 0, the LSP remains in the link state database for an additional 60 seconds, known as the Zero Age Lifetime. If the originating router does not refresh the LSP within this timeframe, the LSP is deleted from the database.

If the received LSP lifetime value is less than the Zero Age Lifetime (60 seconds), IS-IS generates an error indicating a corrupted lifetime event. For example:

```
Dec 14 15:36:45.663 : isis[1011]: RECV L2 LSP 1111.1111.1112.03-00 from
1111.1111.1112.03:
 possible corrupted lifetime 59 secs for L2 lsp 1111.1111.1112.03-00 from SNPA
 02e9.4522.5326 detected.
```

IS-IS saves the received remaining lifetime value in the LSP database. The value appears in the **show isis database** command output under the **Rcvd** field.

Further reference

For more information about the **show isis database** command, refer to the *IS-IS Commands* chapter in the *Routing Command Reference for Cisco NCS 5500 Series Routers*.

10 IS-IS Protection and Resiliency Enhancements

Topics:

- IS-IS non-stop forwarding
- ISIS Non-stop Routing
- IS-IS OOR protection mechanisms
- Overload bits in OOR conditions
- Global weighted SRLG protections
- IS-IS penalties for link delay anomalies
- IS-IS penalties for link loss anomalies
- Inter-area route propagation limits in IS-IS

Outlines IS-IS protection and resiliency features, including nonstop forwarding, NSR, enhancements during OOR conditions, overload bit suppression, global weighted SRLG protection, and penalty mechanisms for link anomalies, with configuration and operational guidelines for robust network stability.

IS-IS non-stop forwarding

Introduces IS-IS nonstop forwarding (NSF) support and provides instructions for configuring nonstop forwarding to ensure uninterrupted protocol operations during control plane restart.

IS-IS non-stop forwarding (NSF) is a network resiliency function that

- minimizes the amount of time a network is unavailable during IS-IS process restarts
- suppresses routing flaps, maintaining routing stability and performance, and
- allows data forwarding to continue while routing protocol state is restored.

IETF NSF: An NSF implementation based on RFC 5306 standards, requiring neighbor routers with NSF-aware software to assist in restoring routing information after a restart.

Cisco NSF: A proprietary implementation that stores the state needed for restart and recovers it using standard IS-IS protocol features, so it works even if neighboring devices do not support RFC 5306.

Impact of restarts on IS-IS adjacencies

When a router running IS-IS restarts, it typically loses adjacency with its peers, causing routing flaps that destabilize the network. NSF suppresses these flaps, keeps routes stable, and ensures minimal disruption. During an IS-IS restart, the router must relearn its IS-IS neighbors without resetting relationships and reacquire the network link-state database.

NSF deployment options

IS-IS NSF can be configured using two options:

- IETF NSF: Relies on cooperation from NSF-aware neighbor routers to rebuild routing information.
- Cisco NSF: Functions independently using standard IS-IS features and does not need neighbor support.

Behavior when neighbor routers lack IETF NSF support

If IETF NSF is enabled and a neighboring router does not support it, adjacency may flap, but nonstop forwarding continues for compatible neighbors. If no neighbors support IETF NSF, a restart falls back to a full cold start.

Hello intervals and BFD considerations for NSF

IS-IS protocol supports shorter hello intervals for neighbor management, but nonstop forwarding works only with the default hello interval. Fast connectivity loss detection is handled by Bidirectional Forwarding Detection (BFD), so hello intervals do not need adjustment for nonstop forwarding.

IS-IS timer guidelines for RP failover

A route processor (RP) failover can interrupt traffic for approximately 50 ms. Configure BFD with a total detection time greater than 50 ms to avoid BFD flap during RP failover. Calculate the total detection time by multiplying the BFD minimum interval by the multiplier.

Example:

- A BFD minimum interval of 15 ms with multiplier 3 produces a 45-ms detection time. Because this value is less than the observed interruption, BFD may flap.
- A BFD minimum interval of 20 ms with multiplier 3 produces a 60-ms detection time. Because this value exceeds the observed interruption, the BFD session is expected to remain up during the RP failover.

Examples

When configuring IS-IS NSF, if neighbors support RFC 5306, they will help restore routing information after a restart. Otherwise, Cisco NSF will work without neighbor support, maintaining traffic forwarding using protocol-defined mechanisms.

In a mixed-vendor environment, Cisco NSF enables resilience without relying on the IETF standard implementation.

Configure nonstop forwarding for IS-IS

Provides instructions for configuring nonstop forwarding in IS-IS, including enabling NSF, choosing protocol options, tuning restart parameters, and verifying NSF settings to ensure continuous traffic forwarding and minimize network disruptions during process restarts.

Ensure IS-IS continues routing traffic and synchronizing link-state information with neighbors across process restarts, such as RP failovers or IS-IS software upgrades, minimizing disruptions and user session loss.

Nonstop forwarding (NSF) enables the router to quickly recover from certain process restarts without causing IS-IS routing interruptions. This is important for service continuity in environments where link flaps or session loss is unacceptable.

Procedure

1. Enable IS-IS routing for the required routing instance.

   Example

   ```
   Router# configure
   Router(config)# router isis isp
   ```

2. Enable NSF and select the protocol based on your network environment.
   - For heterogeneous networks

     ```
     Router(config-isis)# nsf cisco
     ```

   - For homogeneous (IETF draft-based) networks

     ```
     Router(config-isis)# nsf ietf
     ```

   Note

   Use **cisco** if not all adjacent devices support NSF. Use **ietf** if all adjacent devices support IETF draft-based restartability.

3. Configure the number of permitted NSF restart acknowledgment resends.

   Example

   ```
   Router(config-isis)# nsf interface-expires 1
   ```

   If the resend limit is reached, NSF falls back to a cold restart.

4. Configure the number of seconds to wait for each restart acknowledgment.

Example

```
Router(config-isis)# nsf interface-timer 15
```

5. Configure the maximum route lifetime following an NSF restart.

Example

```
Router(config-isis)# nsf lifetime 20
```

> **Note**
>
> Set this value to match the time required for a full NSF restart. If set too high, stale routes may persist; if too low, valid routes may be purged too soon.

6. Use the **show running-config router isis isp** command to check if NSF settings appear under the relevant IS-IS instance.

ISIS Non-stop Routing

Explains ISIS Non-Stop Routing (NSR) functionality and guides users through configuration procedures to maintain routing continuity during process restarts.

Non-stop Routing (NSR) is a high availability routing protocol feature that

- suppresses IS-IS routing changes when devices with redundant route processors experience a switchover event
- reduces network instability and downtime by enabling continuous routing during such events, and
- transfers all necessary peering state information from the active to the standby processor to allow immediate protocol resumption upon switchover.

To preserve routing across process restarts, you must configure Non Stop Forwarding (NSF) in addition to NSR.

When NSR is enabled on a router with redundant route processors, switching from the active to the standby processor does not impact other IS-IS routers in the network, as all protocol state is preserved and routing continues seamlessly.

Configure ISIS-NSR

Guides you through configuring ISIS-NSR to ensure uninterrupted IS-IS routing, including enabling the NSR feature and verifying adjacency, status, and statistics to maintain high availability during process failures.

Enable ISIS Non-Stop Routing (NSR) so your router preserves IS-IS routing state and maintains adjacencies without interruption in the event of process failures.

Procedure

1. Enable the NSR feature.

Example

```
Router(config)# router isis 1
Router(config-isis)# nsr
```

2. Use the **show isis nsr adjacency** command to verify NSR adjacency status:

   Example

   ```
   Router# show isis nsr adjacency
   ```

3. Use the **show isis nsr status** command to verify NSR status:

   Example

   ```
   Router# show isis nsr status
   ```

   This command output displays HA readiness and sync status.

4. Use the **show isis nsr statistics** command to review NSR statistics:

   Example

   ```
   Router# show isis nsr statistics
   ```

   NSR statistics display numbers of adjacencies, LSPs, routes, and interface states for active and standby routers.

IS-IS OOR protection mechanisms

Describes IS-IS protection enhancements specifically designed for out-of-resource (OOR) conditions, emphasizing mechanisms that strengthen protocol resiliency when resources are limited.

IS-IS protection mechanism is a network feature set that

- continuously tracks IS-IS process memory usage to detect and respond to out-of-resource (OOR) conditions
- applies corrective actions, such as process restarts and LSP flow control, to mitigate memory leaks and excessive route installation, and
- ensures ongoing packet forwarding by using Nonstop Forwarding (NSF) during recovery actions to minimize network disruption.

OOR Condition: An OOR condition occurs when the IS-IS process uses more memory than its configured resource limit, potentially triggering process restarts or other protections.

Memory Leak: A memory leak describes a situation where the IS-IS process does not free memory that is no longer needed, gradually causing resource exhaustion.

Nonstop Forwarding (NSF): NSF is a mechanism allowing routers to keep forwarding packets without interruption, even while the IS-IS process restarts.

Table 8: Feature History Table

Feature Name	Release Name	Description
Protect IS-IS Processes in OOR Conditions	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
Protect IS-IS Processes in OOR Conditions	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
Protect IS-IS Processes in OOR Conditions	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Feature Name	Release Name	Description
Protect IS-IS Processes in OOR Conditions	Release 24.2.11	This feature enables prompt alerts for out-of-resource conditions in IS-IS processes that could otherwise cause network instability and disruption due to memory leaks and excessive link-state packets (LSPs). Additionally, you can disable the overload bit status flag in the router's LSP to prevent setting the overload-bit. We recommend consulting with Cisco for optimal results before making this change. Previously, during OOR conditions, IS-IS processes restarted themselves, but the OOR conditions could persist. This ability to protect IS-IS processes in OOR conditions is enabled by default and you can't disable it. The feature introduces these changes: CLI: • The feature introduces fields that indicate the memory state of the IS-IS protocol in the **show isis protocol** command. • **oor-set-overload-bit disable** command. **YANG Data Model** • New XPaths for `Cisco-IOS-XR-clns-isis-cfg` • `Cisco-IOS-XR-um-router-isis-cfg` (refer to GitHub and YANG Data Models Navigator)

Managing excessive link-state packets (LSPs)

Errors such as redistributing large numbers of prefixes from other protocols can result in excessive LSPs being sent to IS-IS. This flood of information can overwhelm router memory and CPU, potentially causing OOR conditions. IS-IS OOR protection mechanisms monitor and control LSP rates and route acceptance to maintain system stability.

Automatic resource limit monitoring

IS-IS regularly checks its memory usage against the maximum permitted value (RLIMIT_DATA). If memory consumption nears the threshold, alerts are generated and protective measures are triggered.

Suppression of FRR backup path computation

When IS-IS is not in the “Normal” memory state, Fast Reroute (FRR) backup path calculations and installations are temporarily disabled. When memory usage returns to “Normal” and remains so for at least 180 seconds, standard FRR operations resume.

- When a memory leak is detected, the IS-IS process restarts to reclaim memory, with NSF ensuring that traffic continues to flow during recovery.
- If IS-IS receives a sudden surge of routes, OOR mechanisms protect the router by regulating how many new routes and LSPs are processed.

Table 9: Memory usage states and associated monitoring intervals

Memory Usage (%)	Memory State	Periodic Monitoring Interval in Seconds
Less than 80	Normal	60
Greater than or equal to 80 and less than 90	Minor	30
Greater than or equal to 90 and less than 95	Severe	20
Greater than or equal to 95	Critical	10

Overload bits in OOR conditions

Explains how overload bit suppression operates in OOR conditions, describes its role in IS-IS to maintain network stability, and outlines critical considerations before configuring or verifying IS-IS overload behavior.

A overload bit is a protocol mechanism that

- signals a router is experiencing resource issues like memory shortages
- causes other routers to disregard its link-state packets (LSPs) when calculating shortest paths, and
- allows the affected router’s directly connected IP prefixes to remain reachable, even when its other paths are excluded from the area’s route calculations.

When a router enters OOR conditions, setting the overload bit in its LSPs prevents other routers from using it in their shortest path calculations. This protects the network from routing loops and forwarding failures due to incomplete or inaccurate routing information. If overload bit suppression is disabled, IS-IS continues forwarding based on possibly incomplete database information, risking network instability.

Best practice

Do not disable overload bit setting during OOR conditions—this mechanism is critical for stable network operation.

Configure overload bit suppression during OOR conditions

Guides you through configuring IS-IS to suppress overload-bit setting during out-of-resource conditions, evaluating operational risks, and verifying the protocol’s memory state.

Prevent IS-IS from automatically setting the overload bit during OOR conditions, enabling continued forwarding with incomplete link-state information.

Before you begin

Evaluate the operational risk of forwarding with an incomplete link-state database before applying this configuration.

Follow these steps to configure overload bit suppression during OOR conditions:

Procedure

1. Disable automatic overload-bit setting during out-of-resource conditions.

   Example

   ```
   Router#router isis 1
   Router(config-isis)#oor-set-overload-bit disable
   ```

2. Use the **show isis protocol** command to verify the IS-IS protocol state and review the memory-state fields in the output.

   Example

   ```
   Router#show isis protocol
   ```

Global weighted SRLG protections

Explains global weighted SRLG protection, details its impact on backup path selection and network resiliency with IS-IS, and outlines prerequisites, configuration types, and distribution methods for SRLG data.

A global weighted SRLG protection is a path selection mechanism that

- associates a weight with the SRLG value
- uses the SRLG weights to compute backup paths, and
- improves network resiliency by optimizing backup path selection in IS-IS routing.

A shared risk link group (SRLG) is a set of links that share a common resource and thus the same risk of failure. Existing loop-free alternate (LFA) implementations in interior gateway protocols (IGPs) support SRLG protection by considering only directly connected links for backup paths. This approach may fail if non-directly connected links sharing the same SRLG are part of the backup path calculation.

Additional details

To support global weighted SRLG protections, you need information about SRLGs on all links within the area topology. SRLG data for remote links can be distributed through IS-IS flooding or by manually configuring remote SRLGs.

Configure global weighted SRLG protection

Provides instructions for configuring global weighted SRLG protection to enhance IGP backup path computation, covering workflows for local SRLG values, remote SRLG flooding, and static remote SRLG provisioning to improve network resilience.

Enable global weighted SRLG protection in your network to enhance failure recovery by considering SRLG weights during the computation of backup paths in IGPs.

Use global weighted SRLG protection to enhance network resilience by assigning weights to SRLGs during backup-path computation. Depending on your network design, you can configure:

- Local SRLG values
- Remote SRLG flooding (to advertise values dynamically), or
- Static provisioning of remote SRLG values (when flooding is not possible).

Before you begin

Determine whether you want to use local SRLG protection, remote SRLG flooding, or static provisioning for remote links.

Follow these steps to configure global weighted SRLG protection (choose the workflow that matches your network requirements):

Procedure

1. Configure local SRLG values and enable global weighted SRLG protection on the local router.

   Example

```
Router(config)# srlg
Router(config-srlg)# interface TenGigE0/0/0/0
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# interface TenGigE0/0/0/1
Router(config-srlg-if)# name group1
Router(config-srlg)# name group value 100
Router(config)# router isis 1
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix srlg-protection
weighted-global
Router(config-isis-if-af)# fast-reroute per-prefix tiebreaker srlg-disjoint
index 1
Router(config-isis)# interface TenGigE0/0/0/0
Router(config-isis-if)# point-to-point
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix
Router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
Router(config-isis)# srlg
Router(config-isis-srlg)# name group1
Router(config-isis-srlg-name)# admin-weight 5000
```

2. Configure remote SRLG flooding on the local router.

   Example

```
Router(config)# router isis 1
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix srlg-protection
weighted-global
Router(config-isis-if-af)# fast-reroute per-prefix tiebreaker srlg-disjoint
index 1
Router(config-isis-if-af)# exit
Router(config-isis)# interface TenGigE0/0/0/0
Router(config-isis-if)# point-to-point
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix
Router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
Router(config-isis-if-af)# exit
Router(config-isis)# srlg
```

```
Router(config-isis-srlg)# name group1
Router(config-isis-srlg-name)# admin-weight 5000
```

3. Configure remote SRLG flooding when you want remote links to advertise their SRLG values dynamically.

   Example

```
Router(config)# srlg
Router(config-srlg)# interface TenGigE0/0/0/0
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# interface TenGigE0/0/0/1
Router(config-srlg-if)# name group1
Router(config-srlg)# name group value 100
Router(config-srlg)# exit
Router(config)# router isis 1
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# advertise application lfa link-attributes srlg
```

4. Configure static provisioning for remote SRLG values when flooding is not used. Perform these configurations on the local router.

   Example

```
Router(config)# srlg
Router(config-srlg)# interface TenGigE0/0/0/0
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# interface TenGigE0/0/0/1
Router(config-srlg-if)# name group1
Router(config-srlg)# name group value 100
Router(config-srlg)# exit
Router(config)# router isis 1
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix srlg-protection
weighted-global
Router(config-isis-if-af)# fast-reroute per-prefix tiebreaker srlg-disjoint
index 1
Router(config-isis)# interface TenGigE0/0/0/0
Router(config-isis-if)# point-to-point
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# fast-reroute per-prefix
Router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
Router(config-isis)# srlg
Router(config-isis-srlg)# name group1
Router(config-isis-srlg-name)# admin-weight 5000
Router(config-isis-srlg-name)# static ipv4 address 10.0.4.1 next-hop ipv4
address 10.0.4.2
Router(config-isis-srlg-name)# static ipv4 address 10.0.4.2 next-hop ipv4
address 10.0.4.1
```

IS-IS penalties for link delay anomalies

Explains IS-IS penalties for link delay anomalies, covering detection with SR-PM, automatic metric adjustments, advertisement and restoration processes, node reload behavior, and configuration guidance for managing IS-IS responses to link delay events.

A penalty for link delay anomaly is a network routing mechanism that

- automatically increases the IS-IS metric when a delay anomaly is detected
- advertises the updated metric in the network to steer traffic away from the affected link, and
- restores the original metric when the link recovers and the anomaly bit is reset.

Link anomaly detection overview

When Link Anomaly Detection is enabled in Segment Routing Performance Measurement (SR-PM), the Performance Measurement (PM) module monitors links for delay anomalies. If a delay anomaly is detected, PM sets an anomaly bit (A-bit) for the affected link.

IS-IS metric adjustment and advertisement

Upon detecting an anomaly (A-bit set), the IS-IS protocol automatically increases the link's metric by a user-defined penalty amount. This higher metric is then advertised to the network, making the link less attractive for routing decisions and steering traffic away from the problematic link.

Recovery and metric restoration

When the link recovers and delay metrics return to normal, the PM module clears the anomaly bit. IS-IS then restores the metric to its default or previously configured value, allowing the link to be used normally in routing again.

Node reload behavior

If a network node is reloaded, the IS-IS protocol advertises the default or configured metric (without penalty) for the link until a new performance measurement is made and any anomalies are reassessed.

Configuration guidance

For information about configuring link delay anomaly threshold values and more implementation details, see the *"Link Anomaly Detection with IGP Penalty"* topic in the *Segment Routing Configuration Guide.*

Configure IS-IS penalties for link delay anomalies

Guides you through configuring IS-IS to apply fallback metrics for interfaces experiencing delay anomalies, ensuring dynamic metric adjustment for optimized routing.

Ensure IS-IS dynamically increases link metrics for interfaces affected by delay anomalies, allowing optimized routing.

Before you begin

Ensure that link anomaly detection is configured in Segment Routing performance measurement before you apply the IGP penalty values.

Procedure

1. Configure the delay-anomaly fallback metric on the affected IS-IS interfaces.

Example

```
Router(config)#router isis 100
Router(config-isis)#interface GigabitEthernet 0/1/0/1
Router(config-isis-if)#address-family ipv4 unicast
Router(config-isis-if-af)#metric fallback anomaly delay increment 25
Router(config-isis-if-af)#exit
Router(config-isis-if)#exit
Router(config-isis)#interface GigabitEthernet 0/1/0/2
Router(config-isis-if)#address-family ipv4 unicast
Router(config-isis-if-af)#metric fallback anomaly delay multiplier 2
```

2. Review the running configuration to confirm the configured increment and multiplier values.

Example

```
router isis 100
 interface GigabitEthernet0/1/0/1
  address-family ipv4 unicast
   metric fallback anomaly delay increment 25
  !
 !
 interface GigabitEthernet0/1/0/2
  address-family ipv4 unicast
   metric fallback anomaly delay multiplier 2
  !
 !
!
```

IS-IS penalties for link loss anomalies

Explains iS-IS penalty for link loss anomaly and the technical context you need before you configure or verify IS-IS behavior.

IS-IS penalty for link loss anomaly is a routing protocol mechanism that

- increases the routing metric of a network interface when abnormal packet loss is detected
- causes IS-IS to prefer alternative paths to improve network stability, and
- facilitates automated traffic engineering to ensure rerouting away from problematic links.

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
IS-IS penalty for link loss anomaly	Release 25.1.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100]); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]) You can enhance network stability and routing efficiency by enabling the IS-IS routing protocol to dynamically adjust an interface's metric when the Performance Monitoring (PM) system detects a packet loss anomaly. This feature introduces these changes: CLI: The *loss* keyword is introduced in the **metric fallback anomaly** command.

This dynamic adjustment lets IS-IS automatically reroute traffic away from links experiencing substantial packet loss, helping maintain overall network stability and efficiency.

For information on configuring the synthetic loss measurement, refer to *Synthetic Loss Measurement* topic in the *Segment Routing Configuration Guide*.

How IS-IS penalty for link loss anomaly works

Describes how IS-IS penalties for link loss anomalies are applied in Cisco IOS XR, outlining the involved components, detection and notification workflow, metric adjustment process, and dynamic traffic rerouting to maintain optimal network reliability.

Summary

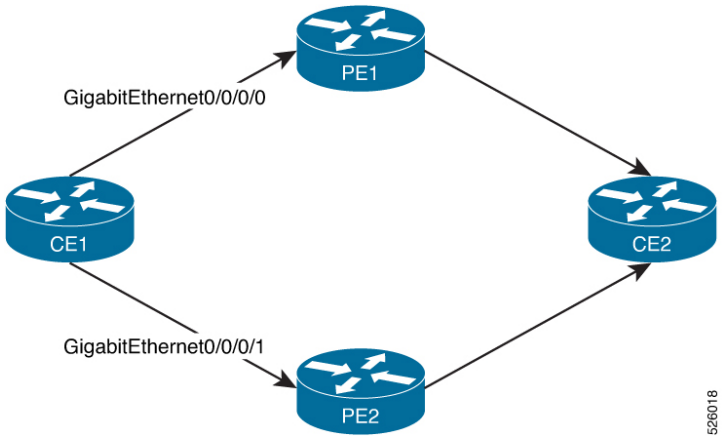
IS-IS penalty for link loss anomaly enables Cisco IOS XR software to dynamically adjust routing metrics whenever synthetic packet loss measurements indicate anomalies on particular network interfaces. This process ensures traffic is rerouted automatically to maintain optimal network reliability.

The key components involved in the process are:

- Customer edge devices (CE1 and CE2): Serve as the source and destination routers in the network topology.
- Provider edge routers (PE1 and PE2): Connect CE devices and provide alternate routing paths (Path 1 and Path 2).
- Cisco IOS XR software metric fallback anomaly loss command: Configured on interface(s) to modify metrics based on detected anomalies.
- Synthetic Loss Measurement: Continuously monitors packet loss and triggers anomaly detection as needed.

Workflow

Figure 1: Traffic flow under metric fallback anomaly loss



The process involves these stages:

1. Detection of anomaly: When the measured packet loss on Path 1 exceeds a pre-defined upper-bound threshold, the Cisco IOS XR software flags a loss anomaly.
2. Notification: Cisco IOS XR communicates the loss anomaly to the IS-IS routing process.
3. Metric adjustment: The metric fallback anomaly loss command on the affected interface instructs IS-IS to modify the interface metric.
4. Traffic rerouting: IS-IS recalculates its shortest path, and Cisco IOS XR reroutes traffic from Path 1 to Path 2.
5. Anomaly recovery: When packet loss recedes below the lower-bound threshold:
   - The synthetic loss measurement system clears the anomaly
   - IS-IS reverts the interface metric to its original value
   - Cisco IOS XR software reroutes traffic back to Path 1

Configure IS-IS penalty for link loss anomaly

Guides you through configuring IS-IS penalty settings for link loss anomalies by adjusting interface metrics for IPv4 and IPv6, and verifying the configuration to enhance routing integrity and mitigate unstable links.

Apply IS-IS penalties for link loss anomalies to improve routing calculation integrity.

Use this task when you need to adjust IS-IS metrics in response to link loss events, ensuring accurate path selection and mitigating the impact of unstable links.

Procedure

1. Enter the IS-IS interface configuration mode, and configure point-to-point service under the IS-IS interface.

   Example

   ```
   Router# configure terminal
   Router(config)# router isis 1
   Router(config-isis)# interface GigabitEthernet0/2/0/1
   Router(config-isis-if)# point-to-point
   ```

2. Configure the metric fallback anomaly loss increment for both IPv4 and IPv6 address families.

Example

```
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly loss increment 500
Router(config-isis-if-af)# exit
Router(config-isis-if)# address-family ipv6 unicast
Router(config-isis-if-af)# metric fallback anomaly loss increment 500
Router(config-isis-if-af)# exit
Router(config-isis-if)# exit
Router(config-isis)# commit
```

3. Use the **show isis database verbose** command to verify the metric fallback anomaly loss increment configuration.

Example

```
Router# show isis database verbose r100
Fri Oct 11 10:49:17.836 PDT
IS-IS 1 (Level-2) Link State Database
LSPID LSP Seq Num LSP Checksum LSP Holdtime/Rcvd ATT/P/OL
r100.00-00 * 0x0000000b 0xc725 1188 /* 0/0/0 Metric: 10 IS-Extended r101.00
Local Interface ID: 10, Remote Interface ID: 7 Interface IP Address: 10.1.1.100
 Neighbor IP Address: 10.1.1.101 Physical BW: 1000000 kbits/sec, 125000000
bytes/sec Link Average Delay: 704 us Link Min/Max Delay: 483/871 us Link Delay
 Variation: 127 us Metric: 10         MT (IPv6 Unicast) IS-Extended r101.00
Local Interface ID: 10, Remote Interface ID: 7 Interface IPv6 Address:
2001:1000::100 Neighbor IPv6 Address: 2001:1000::101 Physical BW: 1000000
kbits/sec, 125000000 bytes/sec Link Average Delay: 704 us Link Min/Max Delay:
 483/871 us Link Delay Variation: 127 us Total Level-2 LSP count: 1 Local
Level-2 LSP count: 1
```

Inter-area route propagation limits in IS-IS

Describes how to control the number of prefixes that IS-IS propagates between levels and how to monitor propagated-prefix limits.

Inter-area route propagation limits in IS-IS are configuration parameters that

- allow operators to specify the maximum number of route prefixes that can be propagated between IS-IS levels
- provide an option to generate a warning when the configured limit is exceeded
- when enforcement is enabled, stop advertising additional prefixes once the limit is reached
- provide a warning-only mode in which all eligible prefixes are advertised, even when the limit is exceeded, and
- help protect the network from excessive routing table growth caused by misconfiguration or unexpected routing protocol events.

Table 11: Feature History Table

Feature Name	Release Name	Description
Inter-area route propagation limits in IS-IS	Release 26.3.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Centralized Systems (8600 [ASIC: Q200]) ; Modular Systems (8800 [LC ASIC: Q200, P100]) This feature helps protect the network from excessive inter-area route propagation caused by misconfiguration or unexpected route growth by limiting the number of prefixes propagated between IS-IS levels. This feature, using the new **maximum-propagated-prefixes** command, allows operators to define propagation thresholds, receive early warnings, and optionally stop advertising additional prefixes when limits are exceeded. This feature introduces these changes: CLI: A new command, **maximum-propagated-prefixes,** has been introduced.

Inter-area route propagation limits in IS-IS

In IS-IS, routes learned in one level may be propagated into another level to support inter-area reachability. Prior to this enhancement, operators could control which routes were propagated by using policy, but they could not directly limit the number of propagated prefixes.

Starting with Cisco IOS XR 26.3.1, this feature adds a new command, **maximum-propagated-prefixes,** which allows operators to define a propagation threshold for inter-area routes. When the configured threshold is approached or exceeded, the router can generate warnings and, depending on configuration, either continue propagating routes or stop advertising additional prefixes. This feature helps protect the network from unexpected route growth between IS-IS levels.

When the propagated-prefix count exceeds the configured threshold, the router can either:

- generate a warning and continue advertising routes, or
- stop advertising additional propagated routes.

Benefits of inter-area route propagation limits in IS-IS

- Improved operational protection against accidental large-scale route propagation.
- Better control of IS-IS inter-area advertisements.
- Configurable warning thresholds for earlier visibility into route growth.
- Operational consistency with the existing IS-IS redistributed-prefix limiting model.

Deployment models for inter-area route propagation limits in IS-IS

Lists the two available deployment modes for inter-area route propagation limits in IS-IS and describes their effects on routing behavior.

The inter-area route propagation limits in IS-IS feature supports two deployment models:

- Warning-only mode (soft limit): The router generates a warning when the configured number of prefixes exceeds the soft limit, but continues advertising routes.
- Enforced mode (hard limit): The router generates a warning when a certain percentage of the hard limit is reached. IS-IS never advertises more propagated prefixes than the configured hard limit. If the number of eligible prefixes exceeds the hard limit, excess prefixes are omitted from the local system's LSPs.

These deployment models enable operators to monitor or strictly control the propagation of inter-area routes, helping prevent excessive route advertisement and maintain stable network operations. Use warning-only mode for monitoring purposes, and enforced mode when strict limits must be upheld.

Restrictions for inter-area route propagation limits in IS-IS

Outlines the mandatory restrictions and controls for limiting inter-area route propagation in IS-IS.

When you configure propagated-prefix limits in IS-IS, consider:

- Use route policy to control which prefixes are eligible for propagation.
- Use **maximum-propagated-prefixes** to control how many eligible prefixes IS-IS advertises.
- This feature applies to inter-area propagated routes. It does not limit redistributed routes.
- To limit redistributed routes, use **maximum-redistributed-prefixes**.
- To control propagation from level 2 into level 1, first enable level 2 into level 1 in IS-IS.
- Level 1 limits prefixes propagated into level 1.
- Level 2 limits prefixes propagated into level 2.
- Omit the level parameter so that the configuration applies to both levels.
- Limits are evaluated independently for the applicable area and topology.

Defaults and ranges for inter-area route propagation limits in IS-IS

Lists the default values, warning thresholds, and configurable ranges for IS-IS inter-area route propagation limits.

This reference details the configurable ranges and default behaviors for inter-area route propagation limits in IS-IS:

- Prefix limit range: 1–100000
- Warning-percentage range: 20–99
- Default warning percentage in enforced mode: 70%
- Default behavior without explicit configuration: warning-only at 20000 prefixes

Use these values to configure and monitor prefix propagation for IS-IS inter-area routes efficiently.

Configure a hard limit threshold on propagated routes

Configure either an enforced hard limit on the number of propagated routes in IS-IS, optionally with a customized warning percentage, or a warning-only threshold that does not limit propagation.

Uncontrolled IS-IS route propagation can result in excessive LSP growth and increased downstream routing state. By setting a hard limit and warning thresholds, you prevent the system from advertising more routes than intended, maintaining network stability and visibility.

Before you begin

- Ensure IS-IS is enabled on the router.
- Confirm the target propagation level is part of your IS-IS design.
- Have any required route policy configured.

Follow these steps to configure a hard limit and warning-only threshold on propagated routes:

Procedure

1. Enter IS-IS address-family configuration mode for the target instance.

   Example

   ```
   Router# config
   Router(config)# router isis core
   Router(config-isis)# address-family ipv4 unicast
   Router(config-isis-af)# propagate level 2 into level 1 route-policy pass-all
   Router(config-isis-af)# maximum-propagated-prefixes 1000 warning-percentage
   50
   ```

2. Configure the maximum propagated prefix limit for the desired level in one of the following ways.

   - This example configures the hard limit for level 2.

     ```
     Router(config-isis-af)# maximum-propagated-prefixes 100 level 2
     ```

   - This example changes the default early-warning percentage from 70% to the specified value.

     ```
     Router(config-isis-af)# maximum-propagated-prefixes 100 warning-percentage
      50 level 2
     ```

     The router tracks the number of propagated prefixes for the configured level. When the count exceeds the configured warning threshold, the router generates a warning. When the count exceeds the configured hard limit, the router stops advertising additional propagated routes.

     A warning message is displayed when the propagated prefix count exceeds the configured hard limit.

     ```
     Router(config-isis-af)# maximum-propagated-prefixes 100 warning-percentage
      50 level 2
     RP/0/1/CPU0:Aug 24 12:52:31.457 CEST: isis[1009]:
     %ROUTING-ISIS-4-MAX_PROPAGATED_PREFIXES : ISIS (eu): Topology IPv4 Unicast
      is exceeding the maximum number of propagated prefixes (100) into level
     2
     ```

   - This example configures the warning-only threshold that does not enforce a hard limit.

     ```
     Router(config-isis-af)# maximum-propagated-prefixes 90 warning-only level
      2
     ```

     The router generates warnings when the propagated-prefix count exceeds the configured threshold, but it continues to advertise propagated routes.

A warning message is displayed when the propagated prefix count exceeds the soft limit.

```
Router(config-isis-af)# maximum-propagated-prefixes 150 warning-percentage
 50 level 2
RP/0/1/CPU0:Aug 24 12:50:03.227 CEST: isis[1009]:
%ROUTING-ISIS-4-MAX_PROPAGATED_PREFIXES : ISIS (eu): Topology IPv4 Unicast
 is exceeding the soft limit of propagated prefixes (150) into level 2
```

3. Run the **show isis protocol** to verify propagated-prefix limits and review the advertised-prefix section for the relevant level.

Example

```
Router# show isis protocol
IS-IS Router:                       eu
  System Id:                        0000.0000.0005
  Hostname:                         ios
  IS Levels:                        level-1-2
  Manual area address:              49.0002
  Routing for area address:         49.0002
  Multi-Instance Id:                0
  Running on:                       RP/0/1/CPU0
  Job Id:                           1009
  PID:                              58446

  Started:                          Tue 2026 Sep 08 14:42:26.442 CEST
    Respawn count:                  3
    Null0 ready:                    Tue 2026 Sep 08 14:42:26.680 CEST
    First adjacency up:             Tue 2026 Sep 08 14:42:28.029 CEST
    Protocol shutdown:              No
    Non-stop forwarding:            Disabled
    Most recent startup mode:       Cold Restart

  Topologies supported by IS-IS
    IPv4 Unicast
      Distance:                     115
      Rib:                          connected
      Advertising local prefixes:   Passive Interfaces Only
      Redistributing:               no protocols
      Level-1
        Metric style:               Wide/Wide (generate/accept)
        Default Interface Metric:   10
        Advertising prefixes:       26
          Connected:                1
          Propagated:               25
            Warning when over:      40
          Redistributed:            0
            Limit:                  10000
            Warning when over:      7000 (70%)
      Level-2
        Metric style:               Wide/Wide (generate/accept)
        Default Interface Metric:   10
        Advertising prefixes:       61
          Connected:                1
          Propagated:               60
            Limit:                  60 (exceeded)
            Warning when over:      30 (50%)
          Redistributed:            0
            Limit:                  10000
            Warning when over:      7000 (70%)
    IPv6 Unicast
```

```
    Distance:                       115
    Rib:                            connected
    Advertising local prefixes:     Passive Interfaces Only
    Redistributing:                 no protocols
    Level-1
      Default Interface Metric:     10
      Advertising prefixes:         1
        Connected:                  1
        Propagated:                 0
          Warning when over:        10
        Redistributed:              0
          Limit:                    10000
          Warning when over:        7000 (70%)
    Level-2
      Default Interface Metric:     10
      Advertising prefixes:         18
        Connected:                  1
        Propagated:                 17
          Warning when over:        10 (exceeded)
        Redistributed:              0
          Limit:                    10000
          Warning when over:        7000 (70%)
Overload Bit:                       not configured
Maximum Metrics:                    not configured

LSP MTU:                            1492
LSP Full:                           level-1: No, level-2: No
LSP fast flooding:                  Disabled
  Max LSP transmit speed:           1000 LSPs/sec
  PSNP delay:                       500 ms
  PSNP interval:                    50 ms

SR-MPLS:
  SRLB allocated:                   15000 - 15999
  SRGB (index 0) notified:          never
  SRGB (index 0):                   allocated, 16000 - 23999

MFI LSD registered:                 Yes
  MFI LSD last registered:          Tue 2026 Sep 08 14:42:32.808 CEST
TE connection status:               Up
  TE connection up:                 Tue 2026 Sep 08 14:42:38.829 CEST
  TE connection flaps:              0
XTC connection status:              Up
  XTC connection up:                Tue 2026 Sep 08 14:42:38.827 CEST
  XTC connection flaps:             0
BFD connection status:              Up
  BFD connection up:                Tue 2026 Sep 08 14:42:27.600 CEST
  BFD connection flaps:             0
  BFD HA role change:               never
    BFD client HA role:             none
    BFD server HA role:             none

System memory state:                Normal
  Last state change:                never
IS-IS memory state:                 Normal
  Current IS-IS memory usage:       22073984 bytes (1%)
  Limit (libmemdbg):                1610612400
  Last state change:                never
  Usage during process start:       733936 bytes (0%)
  Monitoring interval:              60 sec
IS-IS memory usage per item:        items    bytes
  Nodes:                            7        2968
```

```
    LSPs:                           10       4454
    Links:                          28       34048
    Prefixes:                       380      370912
    Nets:                           139      108976
    Paths:                          334      32064

  Interfaces supported by IS-IS eu:
    Loopback0 is running passively (passive in configuration)
    GigabitEthernet0/0/0/35 is running actively (active in configuration)
    GigabitEthernet0/0/0/45 is running actively (active in configuration)
    GigabitEthernet0/0/0/56 is running actively (active in configuration)
    GigabitEthernet0/0/0/65 is running actively (active in configuration)
```

Review the advertised prefix counts and configured limits for each level.

4. Execute the **show isis route detail** command to display a route that is not advertised because the propagation limit was exceeded.

   Example

```
Router# show isis route 2.4.1.197/32 detail
Mon Aug 24 12:54:55.932 CEST

L1 2.4.1.197/32 [110/115] Label: None, medium priority
   Installed Aug 24 12:48:46.944 for 00:06:10
     via 1.1.45.4, GigabitEthernet0/0/0/45, brat, SRGB Base: 16000, Weight:
0
     src brat.00-02, 1.1.100.4
   L2 adv [110] native, propagated, over limit: not advertised
Router#
```

IS-IS generates warnings when the propagated prefix count exceeds the configured threshold, and stops advertising additional propagated routes once the hard limit is reached. If only a warning threshold is set, IS-IS continues advertising routes but generates notifications.

What to do next

Monitor route propagation levels using router logs and the **show isis protocol** command to ensure limits are effective.

11 IS-IS Authentication and Key Management

Topics:

- IS-IS authentication methods
- Configure keychains for IS-IS

Outlines IS-IS authentication methods and key management practices, covering the concepts and configuration steps required to secure IS-IS protocol operations with authentication and keychain mechanisms.

IS-IS authentication methods

Introduces IS-IS authentication principles and provides configuration procedures for enabling authentication on IS-IS, outlining step-by-step instructions to secure IS-IS message exchanges.

IS-IS authentication is a protocol feature that

- limits the establishment of adjacencies and the exchange of Link State Packets (LSPs) using password commands
- provides options for both plain-text and HMAC-MD5 authentication to protect against unauthorized access and ensure data integrity, and
- supports keychain configuration with hitless key rollover and adjustable tolerance for clock skew.

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
HMAC-MD5 Authentication	Release 25.4.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100], 8700 [ASIC: K100])(select variants only*) *This feature is supported on: • 8711-48Z-M • 8011-12G12X4Y-A • 8011-12G12X4Y-D
HMAC-MD5 Authentication	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
HMAC-MD5 Authentication	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) We have enhanced security for IS-IS protocol through authentication, preventing unauthorized devices from forming adjacencies with the router. HMAC-MD5 passwords ensure data integrity by using cryptographic checksums instead of sending passwords over the network. *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Types of IS-IS authentication

IS-IS supports two main authentication types:

- Plain-text authentication: Provides minimal protection. The password is configured on the router and is transmitted in clear text within IS-IS packets, making it visible to anyone able to capture these packets.
- HMAC-MD5 authentication: Offers enhanced security. The password itself is not sent, but is used to generate a cryptographic checksum, which protects the integrity of the exchanged data.

IS-IS authentication commands

IS-IS authentication uses specific commands to secure routing communications.

- Use the **hello-password** command to restrict adjacency formation between routers.
- Use the **lsp-password** command to limit the exchange of Link State Packets (LSPs).
- The area password is set at Level 1; the domain password is set at Level 2 using the **lsp-password** command.

Password configuration methods

IS-IS authentication passwords can be entered as

- Clear text (plain): Directly visible and used in packets.
- Encrypted form: Stored securely, but the plain-text password is still exchanged in certain packets (LSPs, SNPs, hello).

Global keychain configuration

The IS-IS keychain feature allows the protocol to reference globally configured keychains.

- Keychains are used for hello and LSP authentication, configurable at both the router and interface levels.
- Commands referencing the global keychain configuration direct IS-IS to obtain security parameters from the configured set.

Hitless key rollover and tolerance

IS-IS can implement hitless key rollover using keychains.

- Key rollover is time-based; in the event of clock skew between peers, the rollover process may be affected.
- The configurable tolerance setting allows you to extend the accept window (before and after) to accommodate time differences, ensuring seamless authentication transitions.

Configure authentication for IS-IS

Outlines procedures for configuring IS-IS authentication to secure routing exchanges, including choosing methods and passwords, applying authentication to LSPs and Hello packets, and enhancing network integrity on specified interfaces.

Configure IS-IS authentication to protect routing protocol exchanges and enhance security.

IS-IS authentication is used to secure routing messages by requiring a password, ensuring the integrity of route announcements, and preventing unauthorized changes to routing tables. This task is optional but recommended for maintaining network security.

Before you begin

- Decide which authentication method and password to use (for example, HMAC-MD5 or cleartext).

- Identify the interfaces requiring IS-IS authentication.

Procedure

1. Configure the IS-IS LSP authentication password.

   Example

   ```
   Router(config)# router isis isp
   Router(config-isis)# lsp-password hmac-md5 clear password1 level 1
   ```

   - Use **hmac-md5** for HMAC-MD5 authentication or **text** for cleartext authentication.
   - Use **clear** for unencrypted or **encrypted** for encrypted password entry.
   - **level 1** applies authentication to area LSPs/SNPs; **level 2** applies it to backbone LSPs/SNPs.
   - **send-only** authenticates only sent messages; **snp send-only** disables SNP password checking.

2. Enter interface configuration mode for each interface where you want IS-IS authentication.

   Example

   ```
   Router(config-isis)# interface GigabitEthernet 0/1/0/3
   ```

3. Configure the IS-IS Hello authentication password on the interface.

   Example

   ```
   Router(config-isis-if)# hello-password text clear mypassword
   ```

   This enables authentication for IS-IS Hello packets on this interface.

Configure keychains for IS-IS

Guides you through configuring keychains for IS-IS to secure protocol communications, including authentication setup for router and interface levels using keychains for LSPs, Sequence Number PDUs, and HELLO PDUs.

Secure IS-IS protocol messages by configuring keychains at both router and interface levels.

Keychains are used in IS-IS to authenticate Link State PDUs (LSPs) and HELLO PDUs. Configuring keychains enhances protocol security and enables verification of network communication integrity.

Before you begin

- Identify the keychain names and authentication keys for your IS-IS deployment.
- Ensure you are in global configuration mode on the device.

Procedure

1. Enter IS-IS router configuration mode, and configure the router-level keychain for IS-IS LSPs and Sequence Number PDUs.

   Example

   ```
   Router(config)# router isis isp
   Router(config-isis)# lsp-password keychain isis_a level 1
   ```

2. Enter interface configuration mode for the IS-IS-enabled interface, and configure the interface-level keychain for HELLO PDUs.

   Example

   ```
   Router(config-isis)# interface HundredGigE 0/1/0/3
   Router(config-isis-if)# hello-password keychain isis_b
   ```

12 MPLS Traffic Engineering and IS-IS Integration

Topics:

- MPLS Traffic Engineering
- MPLS TE adjacencies
- MPLS label distribution protocol IGP synchronization
- MPLS TE interarea tunnels
- Label distribution protocol IGP auto-configurations
- MPLS LDP-IGP synchronization scenarios with LDP graceful restart
- MPLS LDP-IGP synchronization compatibilities with IGP NSF

Outlines MPLS traffic engineering integration with IS-IS, including adjacency support, label distribution protocol synchronization, interarea tunnels, auto-configuration, and compatibility with graceful restart and nonstop forwarding for resilient network operations.

MPLS Traffic Engineering

Introduces MPLS traffic engineering support, guiding users through the configuration of MPLS TE features for IS-IS environments to optimize network path selection and resource utilization.

Multiprotocol Label Switching Traffic Engineering (MPLS TE) is a set of network capabilities that

- replicates and expands the traffic engineering abilities of traditional Layer 2 ATM and Frame Relay networks
- integrates Layer 2 and Layer 3 technologies to optimize resource management, and
- enables IS-IS to automatically establish and maintain label-switched paths using protocols like Resource Reservation Protocol (RSVP).

Table 13: Feature History Table

Feature Name	Release Name	Description
Multiprotocol Label Switching Traffic Engineering	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
Multiprotocol Label Switching Traffic Engineering	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
Multiprotocol Label Switching Traffic Engineering	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) This feature enhances network efficiency by using RSVP to automatically establish and maintain label-switched paths, selecting routes based on resource requirements and available bandwidth. *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

MPLS TE allows the MPLS backbone to emulate and extend the traffic engineering methods used in ATM and Frame Relay. With IS-IS, MPLS TE uses RSVP to automatically build and maintain label-switched paths (LSPs) across the backbone. The path selection depends on each LSP's resource requirements and the current network conditions, such as bandwidth availability. IS-IS advertises available resources through special TLV extensions. The label-switched paths are explicit routes, commonly called traffic engineering (TE) tunnels.

Configure MPLS TE for IS-IS

Guides you through configuring MPLS TE support for IS-IS, including prerequisites, step-by-step configuration, and verification using show commands to enable scalable and efficient traffic engineering in the network.

Enable MPLS Traffic Engineering (TE) within the IS-IS routing protocol to allow the network to support scalable and efficient path management for traffic engineering.

MPLS TE support in IS-IS lets routers advertise traffic engineering information, which is critical for creating and maintaining MPLS TE Label Switched Paths (LSPs). You should configure these settings on every IS-IS router participating in MPLS TE.

Before you begin

- Make sure your devices support the MPLS feature.
- Do not configure MPLS TE over unnumbered IP links, as it is not supported.

Procedure

1. Enter IS-IS router configuration mode, and specify the address family (IPv4 or IPv6).

   Example

   ```
   Router(config)# router isis isp
   Router(config-isis)# address-family ipv4 unicast
   ```

2. Enable MPLS TE link-state flooding in the desired IS-IS level.

   Example

   ```
   Router(config-isis-af)# mpls traffic-eng level 1
   ```

3. Set the MPLS TE router identifier.

   Example

   ```
   Router(config-isis-af)# mpls traffic-eng router-id loopback0
   ```

4. Configure wide metrics (required for MPLS TE).

   Example

   ```
   Router(config-isis-af)# metric-style wide level 1
   ```

   Configures a router to generate and accept only wide link metrics in the Level 1 area.

5. Use the **show isis instance isp mpls traffic-eng tunnel** command to verify MPLS TE tunnel information.

   Example

   ```
   Router# show isis instance isp mpls traffic-eng tunnel
   ```

6. Use the **show isis instance isp mpls traffic-eng adjacency-log** command to view MPLS TE IS-IS adjacency logs.

   Example

   ```
   Router# show isis instance isp mpls traffic-eng adjacency-log
   ```

7. Use the **show isis instance isp mpls traffic-eng advertisements** command to view the latest MPLS TE flooded advertisements.

   Example

   ```
   Router# show isis instance isp mpls traffic-eng advertisements
   ```

MPLS TE adjacencies

Explains MPLS TE adjacency support, including its integration with IS-IS, logical link operation for LSP tunnels, creation and advertisement of forwarding adjacencies, and technical prerequisites needed before configuring or verifying IS-IS behavior.

MPLS TE adjacency is a traffic engineering feature that

- enables a Label Switch Path (LSP) tunnel to operate as a logical link within an Interior Gateway Protocol (IGP) network using the Shortest Path First (SPF) algorithm,
- allows creation of forwarding adjacencies between routers in the same IS-IS level, even if the routers are several hops apart, and
- advertises TE tunnels as links with specific cost metrics, helping routers outside the TE domain calculate shortest paths for routing traffic.

MPLS TE adjacency support is considered in IS-IS SPF only if a two-way connectivity check is achieved. This check can be satisfied if the forwarding adjacency is bidirectional or if the head-end and tail-end routers of the MPLS TE tunnel are adjacent.

The MPLS TE adjacency support feature is available in IS-IS. For configuration instructions, see the *MPLS Configuration Guide*.

MPLS label distribution protocol IGP synchronization

Describes MPLS LDP IGP synchronization, providing instructions for configuring MPLS LDP IS-IS synchronization to ensure label distribution protocol and IGP alignment, preventing traffic loss during state changes.

MPLS label distribution protocol IGP synchronization is a network feature that

- ensures Label Distribution Protocol (LDP) completes label exchange before an Interior Gateway Protocol (IGP) path is used for MPLS switching
- coordinates LDP session status with IS-IS link metrics to minimize MPLS packet loss, and
- improves network stability and reliability by preventing traffic loss during adjacency changes.

Table 14: Feature History Table

Feature Name	Release Information	Feature Description
MPLS Label Distribution Protocol IGP Synchronization	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
MPLS Label Distribution Protocol IGP Synchronization	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
MPLS Label Distribution Protocol IGP Synchronization	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) You can now prevent traffic loss during IGP adjacency establishment or LDP session closure by synchronizing LDP with IS-IS. This feature ensures that LDP completes its label exchange before the IGP path is used for switching, thus preventing MPLS packet loss. This synchronization enhances network stability and reliability by coordinating LDP and IGP operations, particularly in dynamic network environments. *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Causes of MPLS traffic loss

MPLS traffic loss can happen in two main situations:

- When an IGP adjacency is established, a router might start forwarding packets over the new link before LDP has exchanged labels with the peer on that link.
- When an LDP session closes, the router may keep forwarding traffic across the affected link, instead of switching to an alternate path with an established LDP session.

Synchronization mechanism

This feature provides a synchronization mechanism between LDP and IS-IS by dynamically changing the link metric in IS-IS according to the LDP session state. When an IS-IS adjacency is present but the LDP session is down or label exchange is incomplete, IS-IS advertises the maximum metric for that link. Synchronization is considered achieved when IS-IS returns to advertising the regular configured or default metric.

Metric advertisement behavior

When an IS-IS adjacency is established on a link but the LDP session is lost or label exchange is incomplete, IS-IS advertises the maximum metric for that link. Synchronization is considered achieved once IS-IS resumes advertising the normal configured metric.

> **Note**
>
> In IS-IS, a link with the maximum wide metric (0xFFFFFF) is not considered for shortest path first (SPF). For MPLS LDP IGP synchronization, a maximum wide metric value of -1 (0xFFFFFE) is used with this feature.

Configure MPLS LDP IS-IS synchronization

Introduces the MPLS LDP sync level 1 feature, explaining its configuration and role in synchronizing LDP and IGP adjacencies to maintain optimal MPLS forwarding.

Enable MPLS LDP synchronization alongside IS-IS routing to ensure robust MPLS operations and minimize traffic loss.

MPLS LDP IS-IS synchronization coordinates label distribution with routing updates, reducing the risk of routing black holes and optimizing network stability. Only IPv4 unicast address family is supported.

Before you begin

- Ensure IS-IS routing is enabled.
- Confirm access to interface configuration mode.

Procedure

1. Enter IS-IS router configuration mode.

   Example

   ```
   Router(config)# router isis isp
   ```

2. Enter interface configuration mode for the desired interface.

   Example

   ```
   Router(config-isis)# interface HundredGigE 0/1/0/3
   ```

3. Specify the IPv4 address family.

   Example

   ```
   Router(config-isis-if)# address-family ipv4 unicast
   ```

4. Enable MPLS LDP synchronization for the interface (optionally specify level 1 or 2).

Example

```
Router(config-isis-if-af)# mpls ldp sync level 1
```

MPLS TE interarea tunnels

Details MPLS TE interarea tunnels, covering deployment and configuration strategies that enable scalable traffic engineering across multiple IS-IS areas for extended network reach and flexibility.

MPLS TE interarea tunnel is a network feature that

- allows the creation of MPLS TE tunnels across multiple IGP areas (such as OSPF) and levels (such as IS-IS)
- removes the restriction requiring both tunnel headend and tailend routers to reside in the same area, and
- uses IS-IS or OSPF to establish tunnels, increasing network flexibility and scalability.

Table 15: Feature History Table

Feature Name	Release Name	Description
MPLS TE Interarea Tunnels	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
MPLS TE Interarea Tunnels	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) This feature allows tunnel creation across multiple IGP areas, eliminating the need for the tunnel headend and tailend routers to reside in the same area. It uses IS-IS or OSPF to establish MPLS TE tunnels across areas, enhancing network flexibility and scalability. *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

For details on configuring MPLS TE interarea tunnels, refer to the *MPLS Configuration Guide.*

Label distribution protocol IGP auto-configurations

Introduces label distribution protocol IGP auto-configuration, outlining how to automate LDP enablement across multiple IGP interfaces, support IPv4 in default VRF, and selectively control auto-configuration for IS-IS environments.

Label distribution protocol IGP auto-configuration is a network automation feature that

- simplifies enabling Label Distribution Protocol (LDP) on all interfaces used by an Interior Gateway Protocol (IGP) instance
- supports the IPv4 address family for the default VPN routing and forwarding (VRF) instance, and
- allows auto-configuration on multiple interfaces and IGP instances at once, while letting you explicitly disable auto-configuration on selected interfaces.

This feature can be applied to a large number of interfaces, which is especially useful when LDP is used for transport in the network core. You can selectively disable LDP IGP auto-configuration on individual interfaces using the **igp auto-config disable** command. In this mode, LDP is automatically enabled on all IGP interfaces except those where you have explicitly disabled it.

Refer to the *MPLS configuration guide* for more details on configuring label distribution protocol IGP auto-configuration.

MPLS LDP-IGP synchronization scenarios with LDP graceful restart

Explains MPLS LDP-IGP synchronization scenarios involving LDP graceful restart, covering how traffic is protected during LDP failures, synchronization maintenance during restart, and conditions under which IGP synchronization may be lost.

MPLS LDP-IGP synchronization is a protocol interaction that

- leverages LDP graceful restart to protect MPLS traffic when an LDP session fails
- maintains MPLS LDP IS-IS synchronization on the interface during the graceful restart protection window, and
- loses IGP synchronization if LDP fails to restart within the defined timer intervals.

LDP graceful restart protects traffic when an LDP session is lost. If an LDP session with graceful restart enabled fails, MPLS LDP IS-IS synchronization is maintained on the protected interface while graceful restart is active. MPLS LDP IGP synchronization is eventually lost in these circumstances:

- LDP fails to restart before the LDP graceful restart reconnect timer expires.
- The LDP session on the protected interface fails to recover before the LDP graceful restart recovery timer expires.

MPLS LDP-IGP synchronization compatibilities with IGP NSF

Explains MPLS LDP-IGP synchronization compatibility with IGP NSF, including requirements for LDP graceful restart, effects on LDP synchronization state during IS-IS NSF, and considerations for maintaining state across restarts and failovers.

MPLS LDP-IGP synchronization compatibility is a protocol feature that

- enables LDP synchronization to operate together with IS-IS nonstop forwarding (NSF)
- requires LDP graceful restart to be enabled over the interface for full compatibility, and
- does not retain LDP synchronization state across restarts and failovers unless IS-IS NSF is enabled.

IS-IS nonstop forwarding (NSF) protects traffic during IS-IS process restarts and route processor (RP) failovers. LDP IS-IS synchronization is supported with IS-IS NSF only if LDP graceful restart is also enabled over the interface. Without IS-IS NSF, LDP synchronization state is not preserved across restarts and failovers.

If IS-IS NSF is enabled on the device and LDP graceful restart is active on the interface, LDP synchronization maintains state during RP failover and IS-IS process restarts. If IS-IS NSF is not enabled, synchronization will not persist across processes.

13 Fast Reroute and Load Balancing

Topics:

- IP fast reroute
- Unequal-cost multipaths for IS-IS

Explains fast reroute concepts and load balancing strategies, covering IP and LDP fast reroute mechanisms, configuration of loop-free alternate paths, and unequal-cost multipath approaches for IS-IS environments to enhance network resilience and traffic distribution.

IP fast reroute

Introduces the principles and mechanisms of IP fast reroute, outlining how rapid failover techniques preserve network connectivity and minimize traffic disruption during link or node failures.

A fast reroute mechanism is a network resiliency technique that

- computes alternate repair paths locally to protect against link failures
- prevents packet loss caused by loops during network reconvergence, and
- applies to IP routing networks without requiring MPLS.

Table 16: Feature History Table

Feature Name	Release Name	Description
IP Fast Reroute	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
IP Fast Reroute	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
IP Fast Reroute	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) This feature protects against link failure by using locally computed repair paths, preventing packet loss from loops during network reconvergence. *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

IP fast reroute (IPFRR) uses loop-free alternate (LFA) computation to protect against link failure. It employs locally computed repair paths to prevent packet loss caused by loops during network reconvergence after a failure. Unlike Multiprotocol Label Switching (MPLS), IPFRR is applicable to networks using conventional IP routing and forwarding.

For further technical details, refer to:

- IETF draft-ietf-rtgwg-ipfrr-framework-06.txt
- IETF draft-ietf-rtgwg-lf-conv-frmwk-00.txt

For guidance on configuring MPLS IPFRR, see the *MPLS Configuration Guide*.

Configure IP or LDP fast reroute loop-free alternate for an IS-IS interface

Guides configuration of IP and LDP fast reroute using loop-free alternate paths, providing step-by-step instructions for implementing high-availability redundancy and seamless traffic recovery in the event of failures.

Improve network resiliency by enabling IP or LDP fast reroute loop-free alternate so your router automatically redirects traffic after a link or node failure.

Use this task on Cisco IOS XR devices to configure IP/LDP fast reroute loop-free alternate (LFA) for IS-IS interfaces. IPFRR and MPLS FRR must not be enabled on the same interface at the same time. For node protection on broadcast links, enable both IPFRR (fast reroute) and BFD under IS-IS for the interface.

Before you begin

- Confirm your device runs Cisco IOS XR and supports IPFRR for IPv4 single-level interfaces.
- Do not configure MPLS FRR and IPFRR at the same time on an interface.

Procedure

1. Enable IS-IS routing and select the target interface.

   Example

   ```
   Router(config)# router isis isp
   RouterRouter(config-isis)# interface GigabitEthernet 0/1/0/3
   ```

2. (Optional) Set the IS-IS circuit type (adjacency level).

   Example

   ```
   RouterRouter(config-isis-if)# circuit-type level-1
   ```

3. Enter address family configuration mode for IPv4 unicast and specify the fast reroute computation method and level.

   Example

   ```
   Router(config-isis-if)# address-family ipv4 unicast
   Router(config-isis-if-af)# fast-reroute per-link level 1
   Router(config-isis-if-af)# fast-reroute per-prefix level 2
   ```

   per-link: Prefix-independent; per-prefix: Prefix-dependent.

4. (Optional) Exclude interfaces from FRR computation or define LFA candidates.

   Example

   ```
   Router(config-isis-if-af)# fast-reroute per-link exclude interface Loopback0
    level 1
   ```

```
Router(config-isis-if-af)# fast-reroute per-prefix lfa-candidate interface
MgmtEth0/RP0/CPU0/0 level 1
```

5. Enable IPFRR LFA for the chosen level and (optionally) exclude further interfaces.

Example

```
Router(config-isis-if-af)# ipfrr lfa level 1
Router(config-isis-if-af)# ipfrr lfa exclude interface POS0/1/0/4
```

Unequal-cost multipaths for IS-IS

Introduces unequal-cost multipath (UCMP) for IS-IS, detailing its load-balancing mechanism, configuration options, computation scope, path controls, timer behaviors, weighting adjustments, support for static routes, and comparison with ECMP to guide proper deployment and verification.

A UCMP mechanism is a load-balancing feature for IS-IS that

- allows traffic to be distributed proportionally across multiple paths with unequal costs
- installs paths in the forwarding information base (FIB) with a load metric or weight, and
- enables use of both high and low bandwidth links for forwarding based on their configured metrics.
- IS-IS (Intermediate System to Intermediate System): a routing protocol used for exchanging network topology information.
- UCMP (Unequal-Cost Multipath): a mechanism that enhances traditional ECMP (Equal-Cost Multipath) by including paths with varying metrics or weights.

How UCMP load balancing works in IS-IS

UCMP load-balancing enables IS-IS to split traffic across multiple paths, even when those paths have different costs. Higher bandwidth links typically have lower IGP metrics and form the shortest paths. When UCMP is enabled, IS-IS can also include lower bandwidth or higher cost links, installing these paths in the forwarding information base (FIB) and assigning each a "load metric" or weight. The FIB then uses these weights to distribute traffic proportionally across all eligible paths.

UCMP computation scope and prefix-list options

UCMP computation can be enabled per address family. It can be applied to all reachable prefixes in IS-IS or limited to specific prefixes using a prefix-list. If no prefix-list is specified, all IS-IS prefixes are included in UCMP calculation. This flexibility allows fine-grained traffic engineering based on network policy.

Controlling the number of UCMP paths

The number of UCMP nexthops installed in the routing table is controlled by the variance setting. Variance defines a metric range (expressed as a percentage over the primary path's metric) so that only paths within this range are installed for forwarding. In addition, the total number of ECMP and UCMP paths is limited by either the max-path configuration command or the hardware platform's path capability.

Key UCMP timer and behavior

UCMP computation is triggered after primary SPF (Shortest Path First) and route calculation are complete, with a default delay of 100 ms (configurable with the **ucmp delay-interval** command). UCMP is computed before fast reroute. Backup paths for both ECMP and UCMP routes are then calculated. If an interface should not be used for UCMP, use the **ucmp exclude interface** command to omit it from calculations.

Adjusting UCMP path weighting

You can adjust UCMP ratios by:

- Using the **bandwidth** command in interface configuration mode to alter the perceived speed of links,
- Adjusting IS-IS metrics directly on the interfaces.

These changes update the weights assigned to each path and thus the amount of traffic each path carries.

UCMP support for static routes

UCMP can also be configured for IPv4 or IPv6 static routes. This allows traffic to be load-balanced in a proportional manner across static paths with unequal metrics, extending UCMP capabilities beyond dynamic IGP routes.

If UCMP is enabled, traffic to a destination can be split across three links—a 10 Gbps, 1 Gbps, and 100 Mbps link—each carrying a proportional amount of the traffic according to its metric and assigned weight.

14 IS-IS Overload Bits, Avoidance, and Shutdown Modes

Topics:

- Overload bits on routers
- IS-IS overload bit avoidance
- Configure IS-IS router overload handling
- IID-TLV behavior in IS-IS protocol instances
- IS-IS protocol shutdown mode

Describes IS-IS overload bit mechanisms, avoidance techniques, and shutdown modes, highlighting configuration procedures, operational considerations, and best practices for maintaining stable network routing and service continuity.

Overload bits on routers

Explains the function and impact of IS-IS overload bits, outlining configuration steps, implications for route advertisement, and guidelines for using overload bits to manage network state transitions.

A overload bit is a state indicator that

- is included in a router's Link State Protocol (LSP) data,
- signals to other routers in the area that the router is temporarily unavailable for transit traffic, and
- enables routing protocol stability during specific operational circumstances.

The overload bit is useful in these situations:

- During a serious but nonfatal error, such as limited memory.
- During the startup and restart of the routing process; the overload bit can be set until the routing protocol has converged. However, it is not used during a standard NSF restart or failover, as this would cause routing flaps.
- During the trial deployment of a new router; it can be enabled until deployment is verified, then cleared.
- During the shutdown of a router; the overload bit can be set to remove the router from the topology before decommissioning.

Overload bit behavior in multitopology mode

Describes how the overload bit functions within multitopology mode to manage traffic forwarding. This information assists in the configuration and verification of IS-IS behavior across different address families.

Independent configuration for IPv4 and IPv6

Configuration mode

Impact on traffic forwarding

Operational considerations

The overload bit applies separately to IPv4 and IPv6 address families in IS-IS multitopology mode. You can configure and clear the overload bit independently for each topology.

You set or clear the overload bit from the router’s address family configuration mode. This allows precise control over routing behavior for each protocol.

- If the IPv4 overload bit is set, all routers in the area avoid using the router for IPv4 transit traffic but can still use it for IPv6 transit traffic.
- If the IPv6 overload bit is set, the router is avoided for IPv6 transit traffic and remains available for IPv4 transit traffic.

Understanding this separation helps you configure and verify IS-IS operation across different address families and ensures correct routing decisions in multitopology environments.

IS-IS overload bit avoidance

Details techniques for avoiding IS-IS overload conditions, including configuration methods and operational strategies to ensure network stability and minimize routing disruptions.

IS-IS overload bit avoidance feature is a routing protocol enhancement that

- allows continued use of label switched paths (LSPs) even if routers within the path are flagged with the IS-IS overload bit,
- enables administrators to ignore overload bit status for head, mid, and tail nodes during path selection, and
- provides command-line control for activation and deactivation, ensuring flexibility in traffic engineering.

Table 17: Feature History Table

Feature Name	Release Information	Feature Description
IS-IS Overload Bit Avoidance	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on Cisco 8711-48Z-M routers.
IS-IS Overload Bit Avoidance	Release 7.0.1	You can keep label switched paths (LSPs) available, even if some routers in the path have their IS-IS overload bit set. This feature prevents unnecessary LSP disablement. When you activate this feature, routers with the overload bit set, whether they are head, mid, or tail nodes, are still considered for LSPs. To enable this feature, use the **mpls traffic-eng path-selection ignore overload** command. If you disable the feature, the system uses the default behavior and excludes these nodes from LSPs except as nodes of last resort.

- When you activate overload bit avoidance, label switched paths remain available even if routers in the path have their overload bit set. This prevents unnecessary disablement of paths.
- If you deactivate overload bit avoidance, routers with the overload bit set cannot be used as nodes of last resort, which may cause LSPs to be disabled.

Note

The IS-IS overload bit avoidance feature does not change the default behavior on nodes with their overload bit set if those nodes are not included in the path calculation.

Configure IS-IS overload bit avoidance

Guides you through enabling and disabling IS-IS overload bit avoidance to maintain MPLS TE tunnel traffic during overload or maintenance, including prerequisites and recommended usage scenarios.

Enable the IS-IS overload bit avoidance feature so MPLS traffic engineering tunnels do not reroute traffic from routers marked as overloaded, reducing disruption during planned maintenance.

Use this task in networks running MPLS and IS-IS when you need tunnels to ignore routers marked with the overload bit, typically during maintenance or administrative events.

Before you begin

- Confirm your devices support MPLS and IS-IS.

Procedure

1. Enable IS-IS overload bit avoidance:

   Example

   ```
   Router# configure
   Router(config)# mpls traffic-eng path-selection ignore overload
   ```

   This enables MPLS TE tunnels to ignore the IS-IS overload bit when selecting paths.

2. (Optional) Use the **no mpls traffic-eng path-selection ignore overload** command to disable IS-IS overload bit avoidance.

   Example

   ```
   Router(config)# no mpls traffic-eng path-selection ignore overload
   ```

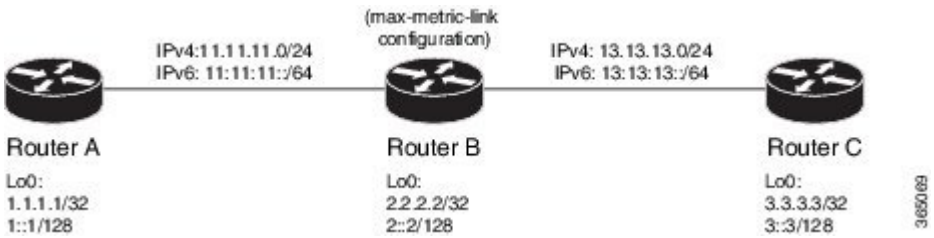
Configure IS-IS router overload handling

Guides you through configuring IS-IS overload handling using the max-metric feature, including interface setup, address family configuration, link metric adjustment, verification of routing and database changes, and ensuring routers remain reachable without becoming preferred transit nodes.

Use this procedure to configure IS-IS overload handling with the max-metric feature.

IS-IS overload handling provides a way to make routers participate in routing without forwarding transit traffic. Setting the overload bit removes the router from transit, but can lead to reachability issues. Using max-metric allows the router to remain reachable as the last-resort transit node while advertising maximum link metrics.

Figure 2: IS-IS overload handling topology



Before you begin

Ensure that you are familiar with configuring router interfaces for a given topology.

Procedure

1. Configure Routers A, B, and C according to your network topology.

   These are sample IP addresses:

   - Router A Loopback0: 1.1.1.1/32 and 1::1/128
   - Router A to Router B: 11.11.11.2/24 and 11:11:11::2/64
   - Router B Loopback0: 2.2.2.2/32 and 2::2/128

- Router B to Router A: 11.11.11.1/24 and 11:11:11::1/64
- Router B to Router C: 13.13.13.1/24 and 13:13:13::1/64
- Router C Loopback0: 3.3.3.3/32 and 3::3/128
- Router C to Router B: 13.13.13.2/24 and 13:13:13::2/64

2. Configure IS-IS and the corresponding NET addresses on Routers A, B, and C.

   Example

```
! Router A
RouterA(config)# router isis ring
RouterA(config-isis)# net 00.0000.0000.0001.00
RouterA(config-isis)# address-family ipv4 unicast
RouterA(config-isis)# metric-style wide
RouterA(config-isis-af)# exit

! Router B
RouterB(config)# router isis ring
RouterB(config-isis)# net 00.0000.0000.0002.00
RouterB(config-isis)# address-family ipv4 unicast
RouterB(config-isis-af)# exit

! Router C
RouterC(config)# router isis ring
RouterC(config-isis)# net 00.0000.0000.0003.00
RouterC(config-isis)# address-family ipv4 unicast
RouterC(config-isis)# metric-style wide
RouterC(config-isis-af)# exit
```

3. Configure IPv4 and IPv6 address families on the loopback interfaces of Routers A, B, and C.

   Example

```
Router(config-isis)# interface loopback0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# exit
Router(config-isis-if)# address-family ipv6 unicast
Router(config-isis-if-af)# exit
Router(config-isis-if)# exit
```

4. Configure the link metrics on the router interfaces.

   Example

```
! Router A interface GigabitEthernet 0/0/0/0 toward Router B

RouterA(config-isis)# interface GigabitEthernet 0/0/0/0
RouterA(config-isis-if)# address-family ipv4 unicast
RouterA(config-isis-if-af)# metric 10
RouterA(config-isis-if-af)# exit
RouterA(config-isis-if)# address-family ipv6 unicast
RouterA(config-isis-if-af)# exit
RouterA(config-isis-if)# exit
```

5. Confirm the route prefixes on Routers A, B, and C before enabling overload handling.

Example

```
! Router A sample output
RouterA# show route
L    1.1.1.1/32 is directly connected, Loopback0
i L1 2.2.2.2/32 [115/20] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 3.3.3.3/32 [115/30] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 13.13.13.0/24 [115/20] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 15.15.15.0/24 [115/30] via 11.11.11.2, GigabitEthernet0/0/0/0

RouterA# show route ipv6
L 1::1/128 is directly connected, Loopback0
i L1 2::2/128 [115/20] via fe80::e9:45ff:fe22:5326, GigabitEthernet0/0/0/0
i L1 3::3/128 [115/30] via fe80::e9:45ff:fe22:5326, GigabitEthernet0/0/0/0
i L1 13:13:13::/64 [115/20] via fe80::e9:45ff:fe22:5326, GigabitEthernet0/0/0/0
i L1 15:15:15::/64 [115/30] via fe80::e9:45ff:fe22:5326, GigabitEthernet0/0/0/0
```

6. Confirm the link metrics on Router B before configuring **max-metric**.

Example

```
RouterB# show isis database
RouterB.00-00
 Metric: 10 IS RouterB.01
 Metric: 10 IS RouterA.00
 Metric: 10 IP 2.2.2.2/32
 Metric: 10 IP 11.11.11.0/24
 Metric: 10 IP 13.13.13.0/24
 MT (IPv6 Unicast) IPv6 2::2/128 Metric: 10
 MT (IPv6 Unicast) IPv6 11:11:11::/64 Metric: 10
 MT (IPv6 Unicast) IPv6 13:13:13::/64 Metric: 10
```

7. Configure **max-metric** on Router B.

Example

```
RouterB(config)# router isis ring
RouterB(config-isis)# max-metric
RouterB(config-isis)# exit
```

8. Confirm the change in link metrics on Router B.

Example

```
RouterB# show isis database
RouterB.00-00
 Metric: 63 IS RouterB.01
 Metric: 63 IS RouterA.00
 Metric: 63 IP 2.2.2.2/32
 Metric: 63 IP 11.11.11.0/24
 Metric: 63 IP 13.13.13.0/24
 MT (IPv6 Unicast) IPv6 2::2/128 Metric: 16777214
 MT (IPv6 Unicast) IPv6 11:11:11::/64 Metric: 16777214
 MT (IPv6 Unicast) IPv6 13:13:13::/64 Metric: 16777214
```

9. Optional: verify the route changes on Routers A and C.

Example

```
! Router A sample output after max-metric
RouterA# show route
i L1 2.2.2.2/32 [115/73] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 3.3.3.3/32 [115/83] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 13.13.13.0/24 [115/73] via 11.11.11.2, GigabitEthernet0/0/0/0
i L1 15.15.15.0/24 [115/83] via 11.11.11.2, GigabitEthernet0/0/0/0

RouterA# show route ipv6
i L1 2::2/128 [115/16777224] via fe80::e9:45ff:fe22:5326,
GigabitEthernet0/0/0/0
i L1 3::3/128 [115/16777234] via fe80::e9:45ff:fe22:5326,
GigabitEthernet0/0/0/0
i L1 13:13:13::/64 [115/16777224] via fe80::e9:45ff:fe22:5326,
GigabitEthernet0/0/0/0
i L1 15:15:15::/64 [115/16777234] via fe80::e9:45ff:fe22:5326,
GigabitEthernet0/0/0/0
```

The output shows the higher route costs caused by the maximum metric configuration.

IID-TLV behavior in IS-IS protocol instances

Explains IID-TLV behavior in IS-IS protocol instances, including its role in uniquely identifying instances, controlling PDU applicability to topologies, and recent feature updates that enable disabling IID-TLV and relevant CLI and YANG data model enhancements.

IID-TLV is a protocol attribute that

- uniquely identifies IS-IS protocol instances when multiple instances are configured
- defines how the Instance Identifier Type-Length-Value (IID-TLV) is included in Hello and LSP packets, and
- allows granular control over protocol data unit (PDU) applicability to specific topologies.

Table 18: Feature History Table

Feature Name	Release Information	Feature Description
Disable IID-TLV of IS-IS Protocol Instance	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8712-MOD-M • 8011-4G24Y4H-I

Feature Name	Release Information	Feature Description
Disable IID-TLV of IS-IS Protocol Instance	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM
Disable IID-TLV of IS-IS Protocol Instance	Release 7.10.1	You can now disable Instance Identifier Type-Length-Value (IID-TLV) in the Hello and LSP packets when multiple IS-IS protocol instances are configured on the router. Each IS-IS instance has a unique instance-ID set, the TLV of which is sent in the Hello and LSP packets. The IID-TLV attribute helps in uniquely identifying the IS-IS protocol instance as well as the topologies to which the Protocol Data Units (PDUs) apply. The feature introduces these changes: CLI New Command: **iid disable** Modified Commands: The **hello-padding** command is extended to IS-IS process configuration mode. The **disable** (IS-IS) command is modified with a new **level** keyword, and also extended to interface configuration mode. **YANG Data Model** New XPaths for openconfig-isis.yang (refer to GitHub and YANG Data Models Navigator)

The IID-TLV attribute provides a way to distinguish between multiple IS-IS protocol instances on the same device and the specific topologies to which their protocol data units (PDUs) apply. By disabling IID-TLV in Hello and LSP packets, you can control how instances are identified and ensure proper isolation in complex configurations.

IS-IS protocol shutdown mode

Explains IS-IS protocol shutdown mode, its operation and technical context, supported platforms and releases, and outlines its benefits for simplifying maintenance, reducing errors, and ensuring stable and consistent network operations.

IS-IS protocol shutdown mode is a router feature that

- enables a graceful shutdown of the IS-IS protocol using a single command

- automatically sets the overload-bit and disables all active IS-IS interfaces (except for Non Stop Routing interfaces), and
- preserves the local link-state packet (LSP) database and other local data structures to ensure a smooth return to operation.

Table 19: Feature History Table

Feature Name	Release Information	Feature Description
IS-IS protocol shutdown mode	Release 26.1.1	Introduced in this release on: Centralized Systems (8400 [ASIC: K100])(select variants only*) * This feature is now supported on Cisco 8404-SYS-D routers.
IS-IS protocol shutdown mode	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8712-MOD-M • 8011-4G24Y4H-I
IS-IS protocol shutdown mode	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) You can now gracefully shut down IS-IS on an interface or router without abruptly interrupting network operations. This feature simplifies operations by consolidating multiple steps into a single command, ensuring network stability during maintenance or configuration changes. This feature introduces these changes: CLI: protocol shutdown **YANG Data Model:** Cisco-IOS-XR-um-router-isis-cfg (refer to GitHub and YANG Data Models Navigator) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Benefits of IS-IS protocol shutdown mode

- Simplifies maintenance by consolidating multiple steps into a single command.
- Reduces the risk of manual errors during configuration changes.

- Ensures network stability during maintenance or configuration events.
- Provides a consistent operational procedure across multiple protocols (similar to OSPF and OSPFv3).

Configure IS-IS protocol shutdown mode

Guides you through enabling or disabling IS-IS protocol shutdown mode to control routing operation, including prerequisites, configuration steps, and verification procedures for maintenance or troubleshooting.

Allow you to control IS-IS protocol operation by activating or deactivating protocol shutdown mode.

Use this task to temporarily disable IS-IS routing without removing its configuration, such as during troubleshooting or maintenance.

Before you begin

- Access the router CLI with administrative privileges.
- Verify that IS-IS routing is already configured on your device.

Procedure

1. Enable IS-IS protocol shutdown mode.

   Example

   ```
   Router# config
   Router(config)# router isis 1
   Router(config-isis)# protocol shutdown
   ```

   To disable protocol shutdown mode and restore IS-IS routing, use:

   ```
   Router(config-isis)# no protocol shutdown
   ```

2. Verify that protocol shutdown mode is enabled for IS-IS.

   Example

   ```
   Router# show isis protocol
   Protocol shutdown: Yes
   ```

3. Confirm that the IS-IS interface is disabled in protocol shutdown state.

   Example

   ```
   Router# show isis interface 1
   Disabled (Protocol shutdown in IS-IS cfg)
   ```

15 Multicast Features

Topics:

- Multicast-intact
- IS-IS multicast topologies

Outlines multicast capabilities, focusing on multicast-intact functionality and topology support using IS-IS to enhance network performance and reliability for multicast traffic.

Multicast-intact

Explains the multicast-intact feature, covering its role in enabling multicast and unicast routing coexistence, protocol support, interaction with MPLS tunnels, management of next-hops, and implementation considerations for IS-IS and OSPFv2.

Multicast-intact is a routing protocol capability that

- enables multicast routing protocols (such as PIM) to coexist with IGP shortcuts (such as MPLS tunnels) on a router
- provides a parallel set of equal-cost next-hops for use by multicast protocols, and
- ensures these next-hops are guaranteed not to contain any IGP shortcuts and are not used for unicast routing or published to the FIB.

Supported protocols

Multicast-intact feature is supported by both OSPFv2 and IS-IS.

Coexistence with MPLS TE and IP multicast

Multicast-intact feature allows MPLS Traffic Engineering (TE) and IP multicast to coexist on the same router.

Publishing multicast-intact next-hops

When multicast-intact is enabled, all IPv4 destinations learned through link-state advertisements are published with a set of equal-cost mcast-intact next-hops to the RIB.This applies even if the native next-hops have no IGP shortcuts.

Next-hop limits in IS-IS and OSPFv2

- In IS-IS, the max-paths limit is applied by counting both native and multicast-intact next-hops together.
- OSPFv2 handles multicast-intact next-hops with its own unique behavior.

Enabling multicast-intact allows PIM protocols to function correctly by providing eligible next-hops, even when the router uses MPLS tunnels as IGP shortcuts.

Enable multicast-intact for IS-IS routes

Guides you through enabling multicast-intact for IS-IS routes to optimize multicast traffic handling and MPLS traffic engineering with IPv4 and IPv6.

Ensure IS-IS routes are configured with multicast-intact to optimize MPLS traffic engineering and improve multicast traffic management for both IPv4 and IPv6 networks.

Before you begin

- Confirm you have access to the router's CLI and are in configuration mode.
- Identify the IS-IS routing instance you need to configure.

Procedure

1. Enter IS-IS routing configuration mode for the appropriate instance.

Example

```
Router(config)# router isis 1
```

2. Specify the address family (IPv4 or IPv6) and enter address family configuration mode.

   Example

   ```
   Router(config-isis)# address-family ipv4 unicast
   ```

   Replace ipv4 unicast with ipv6 unicast or multicast as needed.

3. Enable multicast-intact within the address family configuration.

   Example

   ```
   Router(config-isis-af)# mpls traffic-eng multicast-intact
   ```

   This command enables multicast-intact for the chosen address family.

IS-IS multicast topologies

Explains multicast topology support leveraging IS-IS, highlighting how IS-IS protocols facilitate efficient multicast routing and network stability.

IS-IS multicast topologies are routing configurations that

- enable IS-IS to maintain separate topologies specifically for multicast traffic,
- run a distinct Shortest Path First (SPF) calculation for multicast routes, and
- insert routes from the multicast topology into the multicast-unicast Routing Information Base (muRIB) for each address family.

When you configure IS-IS to support multicast routing for IPv4 or IPv6, IS-IS creates a dedicated multicast topology. This topology operates separately from the unicast topology and uses its own SPF algorithm. IS-IS multicast routes are placed into the muRIB, which is accessed by Protocol Independent Multicast (PIM). As a result, PIM forwards multicast traffic using routes derived from the multicast topology, not the unicast topology. This separation improves multicast routing efficiency and provides independent control over multicast traffic handling.

16 Management Information Base for IS-IS

Topics:

- Management information base (MIB)

Outlines the Management Information Base (MIB) for IS-IS, detailing its structure and role in facilitating effective network monitoring and protocol management.

Management information base (MIB)

Introduces the management information base (MIB) for IS-IS, outlines its structure and function, summarizes supported platforms and release history, and highlights integration with SNMP and RFC 4444 for IS-IS monitoring and management.

Management information base (MIB) is a database structure that

- stores network resources and protocol data for monitoring and management
- organizes managed elements as collections of standard-defined objects, and
- enables external applications and SNMP agents to access device information and control functions.

Table 20: Feature History Table

Feature Name	Release Information	Feature Description
Management Information Base (MIB) for IS-IS	Release 25.4.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is now supported on: • 8011-12G12X4Y-A • 8011-12G12X4Y-D
Management Information Base (MIB) for IS-IS	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on Cisco 8011-4G24Y4H-I routers.
Management Information Base (MIB) for IS-IS	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Cisco IOS XR fully supports MIBs and SNMP traps for IS-IS as defined in RFC 4444. RFC 4444 specifies the objects used to represent IS-IS Routing Protocol data through SNMP.

To find additional details about available MIBs, use the Cisco MIB Locator.