

IS-IS Commands

- address-family (IS-IS), on page 4
- address-family multicast topology (IS-IS), on page 5
- advertise route-policy, on page 6
- authentication-check disable, on page 7
- auto-cost (IS-IS), on page 8
- circuit-type, on page 10
- clear isis process, on page 12
- clear isis route, on page 13
- clear isis statistics, on page 15
- csnp-interval, on page 16
- default-information originate (IS-IS), on page 17
- fast-reroute per-link (IS-IS), on page 19
- fast-reroute per-prefix (IS-IS), on page 21
- fast-reroute per-link priority-limit (IS-IS), on page 23
- fast-reroute per-prefix load-sharing disable (IS-IS), on page 24
- fast-reroute per-prefix tiebreaker (IS-IS), on page 25
- hello-interval (IS-IS), on page 26
- hello-multiplier, on page 28
- hello-padding, on page 30
- hello-password, on page 32
- hello-password keychain, on page 34
- interface (IS-IS), on page 35
- is-type, on page 36
- iid disable, on page 38
- log adjacency changes (IS-IS), on page 39
- log pdu drops, on page 40
- lsp-gen-interval, on page 41
- lsp-mtu, on page 43
- lsp-fast-flooding, on page 44
- lsp-password, on page 46
- lsp-refresh-interval, on page 48
- maximum-paths (IS-IS), on page 49
- maximum-redistributed-prefixes (IS-IS), on page 50

- maximum-propagated-prefixes (IS-IS), on page 51
- max-lsp-lifetime, on page 53
- max-lsp-tx, on page 54
- metric (IS-IS), on page 55
- metric fallback anomaly, on page 57
- metric-style wide, on page 59
- mpls ldp auto-config , on page 61
- mpls ldp sync (IS-IS), on page 62
- mpls traffic-eng (IS-IS), on page 63
- mpls traffic-eng router-id (IS-IS), on page 64
- oor-set-overload-bit disable, on page 65
- net, on page 66
- nsf (IS-IS), on page 68
- passive (IS-IS), on page 69
- point-to-point, on page 70
- protocol shutdown (IS-IS), on page 71
- psnp-interval, on page 72
- redistribute (IS-IS), on page 73
- remote-psnp-delay, on page 77
- router isis, on page 79
- router isis static neighbor, on page 80
- set default-isis-instance, on page 81
- set-overload-bit, on page 83
- show isis, on page 85
- show isis adjacency, on page 87
- show isis adjacency-log, on page 89
- show isis checkpoint adjacency, on page 91
- show isis checkpoint interface, on page 93
- show isis checkpoint lsp, on page 94
- show isis database, on page 96
- show isis database-log, on page 105
- show isis fast-reroute, on page 107
- show isis hostname, on page 110
- show isis interface, on page 112
- show isis lsp-log, on page 116
- show isis mesh-group, on page 119
- show isis mpls traffic-eng adjacency-log, on page 121
- show isis mpls traffic-eng advertisements, on page 123
- show isis mpls traffic-eng tunnel, on page 125
- show isis neighbors, on page 127
- show isis protocol, on page 130
- show isis purge-log, on page 135
- show isis route, on page 137
- show isis spf-log, on page 141
- show isis statistics, on page 148
- show isis topology, on page 152

- show protocols (IS-IS), on page 155
- single-topology, on page 158
- snmp-server traps isis, on page 159
- spf-interval, on page 161
- spf-interval ietf, on page 163
- spf prefix-priority (IS-IS), on page 165
- summary-prefix (IS-IS), on page 167

address-family (IS-IS)

To enter address family configuration mode for configuring Intermediate System-to-Intermediate System (IS-IS) routing that use standard IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes, use the **address-family** command in XR Config mode or interface configuration mode. To disable support for an address family, use the **no** form of this command.

address-family { **ipv4** | **ipv6** } { **unicast** | **multicast** }

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.

Command Default

An address family is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

XR Config mode

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **address family** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can configure routing that uses standard IPv4 or IPv6 address prefixes. An address family must be specified in interface configuration mode. In interface address family configuration mode, you can alter interface parameters for IPv4or IPv6.

You must specify an address family in order to configure parameters that pertain to a single address family.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the IS-IS router process with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)#
```

address-family multicast topology (IS-IS)

To enable a multicast topology when configuring Intermediate System-to-Intermediate System (IS-IS) routing (or to place a given topology within the IS-IS interface), use the **address-family multicast topology** command with either IPv4 or IPv6 address prefix in the appropriate configuration mode. To disable a multicast topology in IS-IS, use the **no** form of this command.

address-family { **ipv4** | **ipv6** } [**maximum prefix** *prefix-limit*]
no address-family

Syntax Description

ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
topology	Specifies the name of the topology.
maximum prefix	Specifies maximum number of prefixes that a routing table can have.
prefix-limit	Maximum number of prefixes. Range is from 32 to 2,000,000.

Command Default

An address family for multicast topology is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration

Interface configuration

Task ID

Task ID	Operations
isis	read, write

advertise route-policy

To specify the route policy that sets the conditions under which the static neighbor is advertised, use the **advertise** command on the IS-IS static neighbor configuration mode.

advertise route-policy *route-policy*

Syntax Description

route-policy	Name of the configured policy.

Command Default

No route policy is applied for advertising the static neighbor.

Command Modes

IS-IS static neighbor configuration

Command History

Release	Modification
Release 26.1.1	The command was introduced.

Usage Guidelines

When **passive** command is enabled on the interface with the static neighbor, the static adjacency and the local IP subnet on the interface are advertised.

When **suppressed** command is configured on the interface with the static neighbor, advertising of the local subnet stops, but the static adjacency continues to be advertised.

When **shutdown** command is configured on the interface with the static adjacency, all advertising, including the static adjacency, stops.

Task ID

Task ID	Operation
isis	read and write

Example

This example shows how to use a route policy to control the advertisement of the static neighbor.

```
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)# 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)# 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
```

authentication-check disable

To suppress Intermediate System-to-Intermediate System (IS-IS) authentication check, use the **authentication-check disable** command in configuration mode. To remove this function, use the **no** form of this command.

authentication-check disable
no authentication-check disable

Command Default

Authentication check is enabled

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.7.1	This command was introduced.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to disable authentication check for IS-IS.

```
Router# configure
Router(config)# router isis 1
Router(config)# authentication-check disable
Router(config)# commit
```

auto-cost (IS-IS)

To control how the Intermediate System-to-Intermediate System (IS-IS) protocol calculates default metrics for the interface, use the **auto-cost** command in the appropriate mode. To revert to the default reference bandwidth, use the **no** form of this command.

auto-cost **reference-bandwidth** *bandwidth* [**granularity** *granularity*]

Syntax Description

reference-bandwidth *bandwidth*	Specifies a bandwidth in kilobits per second (kbps). The range is from 1 to 4294967295.
granularity *granularity*	Specifies a granularity in kilobits per second (kbps). The range is from 1 to 4294967295.

Command Default

mbps : 100 Mbps

Command Modes

Router configuration

Command History

Release	Modification
Release 25.4.1	This command was introduced.

Usage Guidelines

The interface-specific **address-family auto-cost reference-bandwidth** configuration overrides the router-level IS-IS **address-family auto-cost reference-bandwidth** configuration.

The router-level IS-IS **address-family auto-cost reference-bandwidth** configuration applies to all interfaces except those with an interface-specific **address-family auto-cost reference-bandwidth** configuration.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the reference value for the auto cost calculation to 30000000 kbps:

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

The following example shows how to set the reference value for auto-cost calculation on an interface to 30,000,000 kbps.:

```
Router# configure
  Router# (config)#router isis 1
```

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

circuit-type

To configure the type of adjacency used for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **circuit-type** command in interface configuration mode. To reset the circuit type to Level l and Level 2, use the **no** form of this command.

circuit-type { **level-1** | **level-1-2** | **level-2-only** }
no circuit-type

Syntax Description

level-1	Establishes only Level 1 adjacencies over an interface.
level-1-2	Establishes both Level 1 and Level 2 adjacencies, if possible.
level-2-only	Establishes only Level 2 adjacencies over an interface.

Command Default

Default adjacency types are Level 1 and Level 2 adjacencies.

Command Modes

Interface configuration

Usage Guidelines

Adjacencies may not be established even if allowed by the **circuit-type** command. The proper way to establish adjacencies is to configure a router as a Level 1, Level 1 and Level 2, or Level 2-only system using the IS type command. Only on networking devices that are between areas (Level 1 and Level 2 networking devices) should you configure some interfaces to be Level 2-only to prevent wasting bandwidth by sending out unused Level 1 hello packets. Remember that on point-to-point interfaces, the Level 1 and Level 2 hello packets are in the same packet.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a Level 1 adjacency with its neighbor on GigabitEthernetinterface 0/2/0/0 and Level 2 adjacencies with all Level 2-capable routers on GigabitEthernet interface 0/5/0/2:

```
RP/0//CPU0:router(config)# router isis isp
RP/0//CPU0:router(config-isis)# is-type level-1-2
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0//CPU0:router(config-isis-if)# circuit-type level-1
RP/0//CPU0:router(config-isis-if)# exit
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/5/0/2
RP/0//CPU0:router(config-isis-if)# circuit-type level-2-only
```

In this example, only Level 2 adjacencies are established because the **is-type** command is configured:

```
RP/0//CPU0:router(config)# router isis isp
```

```
RP/0//CPU0:router(config-isis)# is-type level-2-only
RP/0//CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0//CPU0:router(config-isis-if)# circuit-type level-1-2
```

Related Commands

Command	Description
IS type	Configures the routing level for an instance of the IS-IS routing process.
Net	Configures an IS-IS NET for the routing process.

clear isis process

To clear the link-state packet (LSP) database and adjacency database sessions for an Intermediate System-to-Intermediate System (IS-IS) instance or all IS-IS instances, use the **clear isis process** command in XR EXEC mode.

clear isis [**instance** *instance-id*] **process**

Syntax Description

instance *instance-id*	(Optional) Specifies IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **clear isis process** command without any keyword to clear all the IS-IS instances. Add the **instance** *instance-id* keyword and argument to clear the specified IS-IS instance.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows the IS-IS LSP database and adjacency sessions being cleared for instance 1:

```
RP/0/RP0/CPU0:router# clear isis instance 1 process
```

clear isis route

To clear the Intermediate System-to-Intermediate System (IS-IS) routes in a topology, use the **clear isis route** command in XR EXEC mode.

clear isis [**instance** *instance-id*] {**afi-all** | **ipv4** | **ipv6**} {**unicast** | **multicast** | **safi-all**} [**topology** *topo-name*] **route**

Syntax Description

instance *instance-id*	(Optional) Specifies IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
afi-all	Specifies IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes.
ipv4	Specifies IPv4 address prefixes.
ipv6	Specifies IPv6 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
safi-all	Specifies all secondary address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.

Command Default

No default behavior or value

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **clear isis route** command to clear the routes from the specified topology or all routes in all topologies if no topology is specified.

Task ID

Task ID	Operations
isis	execute
rib	read, write
basic-services	read, write

Examples

The following example shows how to clear the routes with IPv4 unicast address prefixes:

```
RP/0/RP0/CPU0:router# clear isis ipv4 unicast route
```

clear isis statistics

To clear the Intermediate System-to-Intermediate System (IS-IS) statistics, use the **clear isis statistics** command in XR EXEC mode.

clear isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

instance *instance-id*	(Optional) Clears IS-IS sessions for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **clear isis statistics** command to clear the information displayed by the **show isis statistics** command.

Task ID

Task ID	Operations
isis	execute
rib	read, write
basic-services	read, write

Examples

The following example shows the IS-IS statistics for a specified interface being cleared:

```
RP/0/RP0/CPU0:router# clear isis instance 23 statistics
```

csnp-interval

To configure the interval at which periodic complete sequence number PDU (CSNP) packets are sent on broadcast interfaces, use the **csnp-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

csnp-interval *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Interval (in seconds) of time between transmission of CSNPs on multiaccess networks. This interval applies only for the designated router. Range is 0 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the interval of time between transmission of CSNPs for Level 1 or Level 2 independently.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **csnp-interval** command applies only to the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization. The CSNP interval can be configured independently for Level 1 and Level 2.

Use of the **csnp-interval** command on point-to-point subinterfaces makes sense only in combination with the IS-IS mesh-group feature.

Task ID

Task ID	Operations
isis	execute
rib	read, write
basic-services	read, write

Examples

The following example shows how to set the CSNP interval for Level 1 to 30 seconds:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/0/2/0
RP/0/RP0/CPU0:router(config-isis-if)# csnp-interval 30 level 1
```

default-information originate (IS-IS)

To generate a default route into an Intermediate System-to-Intermediate System (IS-IS) routing domain, use the **default-information originate** command in address family configuration mode. To remove the **default-information originate** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

default-information originate [**external** | **route-policy** *route-policy-name*]

Syntax Description

external	(Optional) Enables default route to be originated as an external route.
route-policy	(Optional) Defines the conditions for the default route.
route-policy-name	(Optional) Name for the route policy.

Command Default

A default route is not generated into an IS-IS routing domain.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

If a router configured with the **default-information originate** command has a route to 0.0.0.0 in the routing table, IS-IS originates an advertisement for 0.0.0.0 in its link-state packets (LSPs).

Without a route policy, the default is advertised only in Level 2 LSPs. For Level 1 routing, there is another process to find the default route, which is to look for the closest Level 1 and Level 2 router. The closest Level 1 and Level 2 router can be found by looking at the attached-bit (ATT) in Level 1 LSPs.

A route policy can be used for two purposes:

- To make the router generate the default route in its Level 1 LSPs.
- To advertise 0.0.0.0/0 conditionally.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to generate a default external route into an IS-IS domain:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# default-information originate
```

fast-reroute per-link (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link [**exclude interface** *type interface-path-id* | **level** { **1** | **2** } | **lfa-candidate interface** *type interface-path-id*]

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

IP fast-reroute LFA per-link computation is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
isis	read, write

This example shows how to configure per-link fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-link level 1
```

fast-reroute per-prefix (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix dependent computation, use the **fast-reroute per-prefix** command in interface address family configuration mode. LFA is supported only on Enhanced Ethernet line card. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix [**exclude interface** *type interface-path-id* | **level** { **1** | **2** } | **lfa-candidate interface** *type interface-path-id* | **remote-lfa** { **maximum-metric** *metric-value* | **tunnel mpls-ldp** } **prefix-list** *prefix-list-name* [**level** { **1** | **2** }]]

Syntax Description

exclude	Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information
level {1 | 2}	Configures FRR LFA computation for one level only.
lfa-candidate	Specifies FRR LFA computation candidate information
interface	Specifies an interface that needs to be either excluded from FRR LFA computation (when used with **exclude** keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the **lfa-candidate** keyword).
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
remote-lfa	Enable remote LFA related configuration.
prefix-list *prefix-list-name*	Filter PQ node router ID based on prefix list.

Command Default

IP fast-reroute LFA per-prefix computation is disabled.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operation
isis	read, write

This example shows how to configure per-prefix fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix level 1
```

This example shows how to configure per-prefix remote-lfa prefix list. The prefix-list option filters PQ node router ID based on prefix list.

```
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix remote-lfa prefix-list
```

fast-reroute per-link priority-limit (IS-IS)

To enable the IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link priority-limit** command in address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link priority-limit { **critical** | **high** | **medium** } **level** { **1** | **2** }

Syntax Description

critical	Enables LFA omputation for critical priority prefixes only.
high	Enables LFA computation for for criticaland high priority prefixes.
medium	Enables LFA computation for for critical, high, and medium priority prefixes.
level {1|2}	Sets priority-limit for routing Level 1 or Level 2 independently.

Command Default

Fast-reroute per link priority limit LFA computation is disabled.

Command Modes

IPv4 unicast address family configuration

IPv6 unicast address family configuration

IPv4 multicast address family configuration

IPv6 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

This example shows how to configure fast-reroute prefix independent per-link computation for critical priority prefixes for level 1 only:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-link priority-limit critical level 1
```

fast-reroute per-prefix load-sharing disable (IS-IS)

To disable load sharing prefixes across multiple backups, use the **fast-reroute per-prefix load-sharing disable** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix load-sharing disable

Syntax Description

level {1|2}	Disables load-sharing for Level 1 or Level 2 independently.

Command Default

Load sharing is enabled.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

This example shows how to disable load-sharing prefixes across multiple backups for level 1 routes:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix load-sharing disable level 1
```

fast-reroute per-prefix tiebreaker (IS-IS)

To configure tie-breaker for multiple backups, use the **fast-reroute per-prefix tiebreaker** command in IPv4 address family configuration mode. To disable tie-breaker configuration, use the **no** form of this command.

fast-reroute per-prefix tiebreaker [**downstream** | **lc-disjoint** | **lowest-backup-metric** | **node-protecting** | **primary-path** | **secondary-path**] **index** *index_number* **level {1 | 2}**

Syntax Description

downstream	Configures to prefer backup path via downstream node, in case of tie-breaker.
lc-disjoint	Configures to prefer Prefer line card disjoint backup path.
lowest-backup-metric	Configures to prefer backup path with lowest total metric.
node-protecting	Configures to prefer node protecting backup path.
primary-path	Configures to prefer backup path from ECMP set.
secondary-path	Configures to prefer non-ECMP backup path.
index	Sets preference order among tie-breakers.
index_number	Value for the index. Range is 1-255.
level {1 | 2}	Configures tiebreaker for Level 1 or Level 2 independently.

Command Default

Tie-breaker for multiple backups is not configured.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

This example shows how to configure preference of backup path via downstream node in case of a tie-breaker for selection of backup path from multiple backup paths:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router isis isp_lfa
RP/0/RP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker downstream index
255
```

hello-interval (IS-IS)

To specify the length of time between consecutive hello packets sent by the Intermediate System-to-Intermediate System (IS-IS) protocol software, use the **hello-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-interval *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Integer value (in seconds) for the length of time between consecutive hello packets. By default, a value three times the hello interval *seconds* is advertised as the *hold time* in the hello packets sent. (That multiplier of three can be changed by using the **hello-multiplier** command.) With smaller hello intervals, topological changes are detected more quickly, but there is more routing traffic. Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies the hello interval for Level 1 and Level 2 independently. For broadcast interfaces only.

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because only a single type of hello packet is sent on serial links, it is independent of Level 1 or Level 2.) Configuring Level 1 and Level 2 independently is used on LAN interfaces.

Note A shorter hello interval gives quicker convergence, but increases bandwidth and CPU usage. It might also add to instability in the network.

A slower hello interval saves bandwidth and CPU. Especially when used in combination with a higher hello multiplier, this strategy may increase overall network stability.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Note Currently, a user can configure an aggressive hello-interval (lower than the default of 10 seconds for peer-to-peer session). But, if NSR or NSF is configured, the default hello interval has to be used so that the sessions do not run into the risk of flapping during switchover.

Using LAN adjacencies in high availability (HA) scenarios is not recommended, since there is no designated intermediate system (DIS) redundancy in the protocol and traffic will either drop or be rerouted temporarily during DIS re-election.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure HundredGigE 0/6/0/0 to advertise hello packets every 5 seconds for Level 1 topology routes. This situation causes more traffic than configuring a longer interval, but topological changes are detected more quickly.

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/6/0/0
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 5 level 1
```

hello-multiplier

To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the **hello-multiplier** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-multiplier *multiplier* [**level** { **1** | **2** }]

Syntax Description

multiplier	Advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Range is 3 to 1000. Neighbors declare an adjacency to this down router after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on an individual interface basis, and can be different between different networking devices in one area. Using a smaller hello multiplier gives faster convergence, but can result in more routing instability. Increase the hello multiplier to a larger value to help network stability when needed. Never configure a hello multiplier to a value lower than the default value of 3.
level { **1** | **2** }	(Optional) Specifies the hello multiplier independently for Level 1 or Level 2 adjacencies.

Command Default

multiplier : 3

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The “holding time” carried in an IS-IS hello packet determines how long a neighbor waits for another hello packet before declaring the neighbor to be down. This time determines how quickly a failed link or neighbor is detected so that routes can be recalculated.

Use the **hello-multiplier** command in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval (hello-interval (IS-IS), on page 26 command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.

On point-to-point links, there is only one hello for both Level 1 and Level 2. Separate Level 1 and Level 2 hello packets are also sent over nonbroadcast multiaccess (NBMA) networks in multipoint mode, such as X.25, Frame Relay, and ATM.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how the network administrator wants to increase network stability by making sure an adjacency goes down only when many (ten) hello packets are missed. The total time to detect link failure is 60 seconds. This strategy ensures that the network remains stable, even when the link is fully congested.

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-interval 6
RP/0/RP0/CPU0:router(config-isis-if)# hello-multiplier 10
```

hello-padding

To configure padding on Intermediate System-to-Intermediate System (IS-IS) hello protocol data units (IIH PDUs) for all IS-IS interfaces on the router, use the **hello-padding** command in interface configuration mode and process configuration mode. To suppress padding, use the **no** form of this command.

hello-padding { **adaptive** | **always** | **disable** | **sometimes** }

[**level** { **1** | **2** }]

Syntax Description

adaptive	Enable hello-padding till neighbor confirms adjacency up.
always	Always enable hello-padding.
disable	Suppresses hello-padding.
sometimes	Enables hello-padding during adjacency formation only.
level { **1** | **2** }	(Optional) Specifies hello padding for Level 1 or Level 2 independently.

Command Default

Hello padding is enabled.

Command Modes

Interface configuration and IS-IS process configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 7.10.1	This command was extended to IS-IS process configuration mode.

Usage Guidelines

You might want to suppress hello padding to conserve network resources. The lower the circuit speed, the higher the percentage of padding overhead. Before suppressing the hello padding, you should know your physical and data link layer configurations and have control over them, and also know your router configuration at the network layer.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to suppress IS-IS hello padding over local area network (LAN) circuits for interface HundredGigE 0/2/0/1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/1
RP/0/RP0/CPU0:router(config-isis-if)# hello-padding disable
```

The following example shows how to suppress IS-IS hello padding at the process level:

```
Router (config)#router isis 100
Router (config-isis)#hello-padding disable
Router (config)#commit
```

hello-password

To configure the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password** command in interface configuration mode. To disable authentication, use the **no** form of this command.

hello-password [**hmac-md5** | **text**] [**clear** | **encrypted**] *password* [**level** { **1** | **2** }] [**send-only**]

Syntax Description

hmac-md5	(Optional) Specifies that the password use HMAC-MD5 authentication.
text	(Optional) Specifies that the password use clear text password authentication.
clear	(Optional) Specifies that the password be unencrypted.
encrypted	(Optional) Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign for an interface.
level { **1** | **2** }	(Optional) Specifies whether the password is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the password applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When a **text** password is configured, it is exchanged as clear text. Therefore, the **hello-password** command provides limited security.

When an **hmac-md5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a password with HMAC-MD5 authentication for hello packets running on HundredGigE 0/2/0/3 interface:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/2/0/3
RP/0/RP0/CPU0:router(config-isis-if)# hello-password hmac-md5 clear mypassword
```

hello-password keychain

To configure the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password keychain** command in interface configuration mode. To disable the authentication password keychain, use the **no** form of this command.

hello-password keychain *keychain-name* [**level** { **1** | **2** }] [**send-only**]

Syntax Description

keychain	Keyword that specifies the keychain to be configured. An authentication password keychain is a sequence of keys that are collectively managed and used for authenticating a peer-to-peer group.
keychain-name	Specifies the name of the keychain.
level { **1** | **2** }	(Optional) Specifies whether the keychain is for a Level 1 or a Level 2 protocol data unit (PDU).
send-only	(Optional) Specifies that the keychain applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Specify a keychain to enable keychain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a password keychain for level 1, send only authentication on a HundredGigE:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-isis-if)# hello-password keychain mykeychain level 1 send-only
```

interface (IS-IS)

To configure the Intermediate System-to-Intermediate System (IS-IS) protocol on an interface, use the **interface** command in XR Config mode. To disable IS-IS routing for interfaces, use the **no** form of this command.

interface *type* *interface-path-id*

Syntax Description

type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default

No interfaces are specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

An address family must be established on the IS-IS interface before the interface is enabled for IS-IS protocol operation.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable an IS-IS multitopology configuration for IPv4 on HundredGigE interface 0/3/0/0:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# metric-style wide level 1
!
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-if)# ipv4 address 2001::1/64
```

is-type

To configure the routing level for an Intermediate System-to-Intermediate System (IS-IS) area, use the **is-type** command in XR Config mode. To set the routing level to the default level, use the **no** form of this command.

is-type { **level-1** | **level-1-2** | **level-2-only** }

Syntax Description

level-1	Specifies that the router perform only Level 1 (intra-area) routing. This router learns only about destinations inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.
level-1-2	Specifies that the router perform both Level 1 and Level 2 routing.
level-2-only	Specifies that the routing process acts as a Level 2 (interarea) router only. This router is part of the backbone, and does not communicate with Level 1-only routers in its own area.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When the router is configured with Level 1 routing only, this router learns about destinations only inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.

When the router is configured with Level 2 routing only, this router is part of the backbone, and does not communicate with Level 1 routers in its own area.

The router has one link-state packet database (LSDB) for destinations inside the area (Level 1 routing) and runs a shortest path first (SPF) calculation to discover the area topology. It also has another LSDB with link-state packets (LSPs) of all other backbone (Level 2) routers, and runs another SPF calculation to discover the topology of the backbone and the existence of all other areas.

We highly recommend that you configure the type of an IS-IS routing process to establish the proper level of adjacencies. If there is only one area in the network, there is no need to run both Level 1 and Level 2 routing algorithms.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify that the router is part of the backbone and that it does not communicate with Level 1-only routers:

```
RP/0/RP0/CPU0:router(config)# router isis isp
```

```
RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only
```

iid disable

To suppress instance-ID TLV for `hello` and `lsp` packets when multiple IS-IS instances are configured in a router, use the **iid disable** command in process configuration mode.

To remove this function, use the **no** form of this command.

iid disable

Syntax Description

This command has no keywords or arguments.

Command Default

Sending or receiving is enabled.

Command Modes

IS-IS process configuration

Command History

Release	Modification
Release 7.10.1	The command was introduced.

Usage Guidelines

Each IS-IS instance has a unique instance-ID set, the TLV of which is sent in the **hello** and **lsp** packets.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to suppress instance ID TLV at the IS-IS process configuration level:

```
Router(config-isis)#iid disable
Router(config-isis)#commit
```

log adjacency changes (IS-IS)

To cause an IS-IS instance to generate a log message when an Intermediate System-to-Intermediate System (IS-IS) adjacency changes state (up or down), use the **log adjacency changes** command in router configuration mode. To restore the default value, use the **no** form of this command.

log adjacency changes

Command Default

No IS-IS instance log messages are generated.

Command Modes

Router configuration

Usage Guidelines

Use the **log adjacency changes** command to monitor IS-IS adjacency state changes; it may be very useful when you are monitoring large networks. Messages are logged using the system error message facility. Messages can be in either of two forms:

```
%ISIS-4-ADJCHANGE: Adjacency to 0001.0000.0008 (Gi 0/2/1/0) (L2) Up, new adjacency
%ISIS-4-ADJCHANGE: Adjacency to router-gsr8 (Gi /2/1/0) (L1) Down, Holdtime expired
```

Using the **no** form of the command removes the specified command from the configuration file and restores the system to its default condition with respect to the command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to log adjacency changes:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# log adjacency changes
```

Related Commands

Command	Description
logging	Logs messages to a syslog server host.

log pdu drops

To log Intermediate System-to-Intermediate System (IS-IS) protocol data units (PDUs) that are dropped, use the **log pdu drops** command in XR Config mode. To disable this function, use the **no** form of this command.

log pdu drops

Command Default

PDU logging is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **log pdu drops** command to monitor a network when IS-IS PDUs are suspected of being dropped. The reason for the PDU being dropped and current PDU drop statistics are recorded.

The following are examples of PDU logging output:

```
%ISIS-4-ERR_LSP_INPUT_Q_OVERFLOW - An incoming LSP or SNP pdu was dropped because the input
 queue was full
%ISIS-3-ERR_SEND_PAK - The process encountered a software-error while sending the IS-IS
packet
```

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable PDU logging:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# log pdu drops
```

lsp-gen-interval

To customize IS-IS throttling of link-state packet (LSP) generation, use the **lsp-gen-interval** command in XR Config mode. To restore the default value, use the **no** form of this command.

lsp-gen-interval [**initial-wait** *initial*] [**secondary-wait** *secondary*] [**maximum-wait** *maximum*] [**level** { **1** | **2** }]

Syntax Description

initial-wait *initial*	Specifies the initial LSP generation delay (in milliseconds). Range is 0 to 120000 milliseconds.
secondary-wait *secondary*	Specifies the hold time between the first and second LSP generation (in milliseconds). Range is 1 to 120000 milliseconds.
maximum-wait *maximum*	Specifies the maximum interval (in milliseconds) between two consecutive occurrences of an LSP being generated. Range is 1 to 120000 milliseconds.
level { **1** | **2** }	(Optional) Specifies the LSP time interval for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

During prolonged periods of network instability, repeated recalculation of LSPs can cause increased CPU load on the local router. Further, the flooding of these recalculated LSPs to the other Intermediate Systems in the network causes increased traffic and can result in other routers having to spend more time running route calculations.

Use the **lsp-gen-interval** command to reduce the rate of LSP generation during periods of instability in the network. This command can help to reduce CPU load on the router and to reduce the number of LSP transmissions to its IS-IS neighbors.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the maximum interval between two consecutive occurrences of an LSP to 15 milliseconds and the initial LSP generation delta to 5 milliseconds:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5
```

lsp-mtu

To set the maximum transmission unit (MTU) size of Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs), use the **lsp-mtu** command in XR Config mode. To restore the default, use the **no** form of this command.

lsp-mtu *bytes* [**level** { **1** | **2** }]

Syntax Description

bytes	Maximum packet size in bytes. The number of bytes must be less than or equal to the smallest MTU of any link in the network. Range is from 128 to 4352 bytes. **Note** Range is 128 to 8979 bytes from Release 6.6.3 onwards.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Under normal conditions, the default MTU size should be sufficient. However, if the MTU size of a link is less than 1500 bytes, the LSP MTU size must be lowered accordingly on each router in the network. If this action is not taken, routing becomes unpredictable.

This guideline applies to all Cisco networking devices in a network. If any link in the network has a reduced MTU size, all devices must be changed, not just the devices directly connected to the link.

Note Do not set the **lsp-mtu** command (network layer) to a value greater than the link MTU size that is set with the **mtu** command (physical layer).

To be certain about a link MTU size, use the show isis interface, on page 112 command to display the value.

Examples

The following example shows how to set the MTU size to 1300 bytes:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-mtu 1300
```

lsp-fast-flooding

To enable the fast flooding of Link State Packets (LSPs) to improve network response to topology changes, use the **lsp-fast-flooding** command in IS-IS router configuration mode.

lsp-fast-flooding [**max-lsp-tx** *lsps-per-sec* | **remote-psnp-delay** *milliseconds*\

Syntax Description

max-lsp-tx *lsps-per-sec*	(Optional) Sets the maximum number of LSPs that can be transmitted per second. The default is 1000. The range is from 33 through 5000
remote-psnp-delay *milliseconds*	(Optional) Specifies the maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending an LSP. The default is 500 ms. The range is from 1 through 5000 ms.

Command Default

This command is disabled by default. IS-IS fast flooding of LSPs is not enabled.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to enable IS-IS fast flooding of LSPs.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
```

This example shows how to set the maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-router)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
```

lsp-password

To configure the link-state packet (LSP) authentication password, use the **lsp-password** command in XR Config mode. To remove the **lsp-password** command from the configuration file and disable link-state packet authentication, use the **no** form of this command.

lsp-password [[**hmac-md5** | **text**] [**clear** | **encrypted**] *password* | **keychain** *keychain-name*] [**level** { **1** | **2** }] [**send-only**] [**snp send-only**] [**enable-poi**]]

Syntax Description

hmac-md5	Specifies that the password uses HMAC-MD5 authentication.
text	Specifies that the password uses clear text password authentication.
clear	Specifies that the password be unencrypted.
encrypted	Specifies that the password be encrypted using a two-way algorithm.
password	Authentication password you assign.
keychain	(Optional) Specifies a keychain.
keychain-name	Name of the keychain.
level { **1** | **2** }	(Optional) Specifies the password for Level 1 or Level 2 independently.
send-only	(Optional) Adds passwords to LSP and sequence number protocol (SNP) data units when they are sent. Does not check for authentication in received LSPs or sequence number PDUs (SNPs).
snp send-only	(Optional) Adds passwords to SNP data units when they are sent. Does not check for authentication in received SNPs. This option is available when the **text** keyword is specified.
enable-poi	The **enable-poi** keyword inserts the purge originator identification (POI), if you are using cryptographic authentication. If you are not using cryptographic authentication, then the POI is inserted by default.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

When a **text** password is configured, it is exchanged as clear text. Therefore, the **lsp-password** command provides limited security.

When an **HMAC-MD5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

The recommended password configuration is that both incoming and outgoing SNPs be authenticated.

Note To disable SNP password checking, the **snp send-only** keywords must be specified in the **lsp-password** command.

To configure an additional password, use the **lsp-password accept** command.

Specify a key chain to enable key chain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure separate Level 1 and Level 2 LSP and SNP passwords, one with HMAC-MD5 authentication and encryption and one with clear text password authentication and no encryption:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1
RP/0/RP0/CPU0:router(config-isis)# lsp-password text clear password2 level 2
```

lsp-refresh-interval

To set the time between regeneration of link-state packets (LSPs) that contain different sequence numbers, use the **lsp-refresh-interval** command in XR Config mode. To restore the default refresh interval, use the **no** form of this command.

lsp-refresh-interval *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Refresh interval (in seconds). Range is 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 900 seconds (15 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The refresh interval determines the rate at which the software periodically sends the route topology information that it originates. This behavior is done to keep the information from becoming too old. By default, the refresh interval is 900 seconds (15 minutes).

LSPs must be refreshed periodically before their lifetimes expire. The refresh interval must be less than the LSP lifetime specified with this router command. Reducing the refresh interval reduces the amount of time that undetected link-state database corruption can persist at the cost of increased link utilization. (This event is extremely unlikely, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to change the LSP refresh interval to 10,800 seconds (3 hours):

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# lsp-refresh-interval 10800
```

maximum-paths (IS-IS)

To configure the maximum number of parallel routes that an IP routing protocol will install in the routing table, use the **maximum-paths** command in address family configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum-paths *maximum*

Syntax Description

maximum	Maximum number of parallel routes that IS-IS can install in a routing table. Range is 1 to 64

Command Default

No default Command

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to allow a maximum of 16 paths to a destination:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# maximum-paths 16
```

maximum-redistributed-prefixes (IS-IS)

To specify an upper limit on the number of redistributed prefixes (subject to summarization) that the Intermediate System-to-Intermediate System (IS-IS) protocol advertises, use the **maximum-redistributed-prefixes** command in address family mode. To disable this feature, use the **no** form of this command.

maximum-redistributed-prefixes *maximum* [**level** { **1** | **2** }]

Syntax Description

maximum	Maximum number of redistributed prefixes advertised. Range is 1 to 28000.
level { **1** | **2** }	(Optional) Specifies maximum prefixes for Level 1 or Level 2.

Command Default

maximum: 10000

level : 1-2

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **maximum-redistributed-prefixes** command to prevent a misconfiguration from resulting in redistribution of excess prefixes. If IS-IS encounters more than the maximum number of prefixes, it sets a bi-state alarm. If the number of to-be-redistributed prefixes drops back to the maximum or lower—either through reconfiguration or a change in the redistribution source—IS-IS clears the alarm.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify the number of redistributed prefixes at 5000 for Level 2:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# maximum-redistributed-prefixes 5000 level 2
```

maximum-propagated-prefixes (IS-IS)

To configure the maximum number of propagated prefixes, use the **maximum-propagated-prefixes** command in mode.

maximum-propagated-prefixes *limit* [**warning-only** | **warning-percentage** *percentage*] [**level** *value*]

Syntax Description

maximum-propagated-prefixes *limit*	Configures the maximum number of propagated prefixes. The number of propagated prefixes ranges from 1 to 100000.
warning-only	Configures the warning-only behavior. The default warning-only threshold is 20,000 prefixes.
warning-percentage *percent*	Configures the warning threshold percentage. The default percentage is 70%. The warning threshold percentage ranges from 20 percentage and 99 percentage.
level *value*	Applies the limit to level 1 or level 2 propagation.

Command Default

Limit: 20,000

Behavior: warning-only

Warning percentage: 70% when a hard limit applies

Command Modes

IS-IS address-family configuration

Command History

Release	Modification
Release 26.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to configure the maximum propagated prefix limit for the level 2:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 8000 level 2
```

This example shows how to set an early warning threshold as a percentage.

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 8000 warning-percentage 70 level 2
```

This example shows how to configure a warning-only threshold.

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# maximum-propagated-prefixes 20000 warning-only level 2
```

max-lsp-lifetime

To set the maximum time that link-state packets (LSPs) persist without being refreshed, use the **max-lsp-lifetime** command in XR Config mode. To restore the default time, use the **no** form of this command.

max-lsp-lifetime *seconds* [**level** { **1** | **2** }]

Syntax Description

seconds	Lifetime (in seconds) of the LSP. Range from 1 to 65535 seconds.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

seconds : 1200 seconds (20 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

You might need to adjust the maximum LSP lifetime if you change the LSP refresh interval with the **lsp-refresh-interval** command. The maximum LSP lifetime must be greater than the LSP refresh interval.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the maximum time that the LSP persists to 11,000 seconds (more than 3 hours):

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# max-lsp-lifetime 11000
```

max-lsp-tx

To set the maximum number of Link State Packets (LSPs) that can be transmitted per second in an IS-IS routing process, use the **max-lsp-tx** command in IS-IS router lsp-fast-flooding configuration mode. To remove the configured maximum LSP transmission rate, use the **no** form of this command.

max-lsp-tx *lsps-per-sec*

Syntax Description

lsps-per-sec	The maximum number of LSPs that can be transmitted per second. The default is 1000 ms. The range is from 33 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set maximum LSP transmission rate to 600 LSPs per second.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# max-lsp-tx 600
```

metric (IS-IS)

To configure the metric for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric** command in address family or interface address family configuration mode. To restore the default metric value, use the **no** form of this command.

metric { *default-metric* | **maximum** } [**level** { **1** | **2** }]

Syntax Description

default-metric	Metric assigned to the link and used to calculate the cost from each other router using the links in the network to other destinations. Range is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. **Note** Setting the default metric under address family results in setting the same metric for all interfaces that is associated with the address family. Setting a metric value under an interface overrides the default metric
maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

default-metric : Default is 10.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Set the default metric under address family to set the same metric for all interfaces that is associated with the address family. Set a metric value under an interface to override the default metric.

We highly recommend that you configure metrics on all interfaces.

Metrics of more than 63 cannot be used with narrow metric style.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure HundredGigE interface with a default link-state metric cost of 15 for Level 1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE0/1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# metric 15 level 1
```

The following example shows how to configure a metric cost of 15 for all interfaces under address family IPv4 unicast for level 2:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# metric 15 level 2
```

metric fallback anomaly

To configure the metric when a delay or loss anomaly occurs for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric fallback anomaly** command in address family or interface address family configuration mode.

metric fallback anomaly { [**delay** | **loss**] [**te-metric**] { **increment** *penalty_value* | **multiplier** *penalty_multiplier* | **maximum** } } [**level** { **1** | **2** }]

Syntax Description

maximum	Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF).
delay	Increases the link metric only when a delay anomaly occurs.
loss	Increases the link metric only when a loss anomaly occurs.
te-metric	Modifies the TE metric.
increment *penalty_value*	Increments the link metric by the specified amount. The range is from 1 to 16777213.
multiplier *penalty_multiplier*	Multiplies the link metric by the specified amount. The range is from 2 to 100000.
level { **1** | **2** }	(Optional) Specifies the SPF calculation for Level 1 or Level 2 independently.

Command Default

Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

Release	Modification
Release 25.1.1	The **loss** keyword was introduced.
Release 7.4.1	This command was introduced.

Usage Guidelines

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Use the **fallback anomaly** options to increase the metric when you configure link delay anomaly detection in Segment Routing Performance Measurement (SR-PM).

When you configure link delay anomaly detection, SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

From Release 25.1.1, you can configure link loss anomaly detection so that SR-PM sends an anomaly bit (A-bit) when the upper bound of the anomaly threshold is exceeded. When an anomaly is received for an interface, IS-IS can increase the IGP metric or TE metric to make the link less attractive for traffic.

You can specify how to increase the metric in the following ways:

- As a factor by which the base metric is multiplied
- As an increment that is added to the base metric
- As the maximum value. All routers will exclude this link from their SPF.

Task ID

Task ID	Operations
isis	read, write

Examples

In the following example, the base IGP metric value is 10. This example shows how to increase the IGP metrics for an interface. The IGP metric will be increased by 50. When an anomalous bit notification comes for interface HundredGigE 0/1/0/0, the value advertised in the TLVs will be IGP metric 60.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly delay increment 50
```

From Release 25.1.1, you can configure link loss anomaly detection. This example displays how to increase the IGP metrics for an interface.

Consider that the base IGP metric value is 10. You can configure the IGP metric to increase by 500, so that if an anomalous bit notification occurs for interface HundredGigE 0/1/0/0 due to packet loss, the value advertised in the TLVs will be IGP metric 510.

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/1/0/0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# metric fallback anomaly loss increment 500
```

metric-style wide

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept only new-style type, length, and value (TLV) objects, use the **metric-style wide** command in address family configuration mode. To remove the **metric-style wide** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style wide [**transition**] [**level** { **1** | **2** }]

Syntax Description

transition	(Optional) Instructs the router to generate and accept both old-style and new-style TLV objects. It generates only new-style TLV objects.
level { **1** | **2** }	(Optional) Specifies routing Level 1 or Level 2 independently.

Command Default

Old-style TLV lengths are generated, if this command is not configured.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

IS-IS traffic engineering extensions include new-style TLV objects with wider metric fields than old-style TLV objects. If you enter the **metric-style wide** command, a router generates and accepts only new-style TLV objects. Therefore, the router uses less memory and fewer other resources rather than generating both old-style and new-style TLV objects.

To perform MPLS traffic engineering, a router needs to generate new-style TLV objects.



Note This discussion of metric styles and transition strategies is oriented toward traffic engineering deployment. Other commands and models might be appropriate if the new-style TLV objects are desired for other reasons. For example, a network may require wider metrics, but might not use traffic engineering.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a router to generate and accept only new-style TLV objects on Level 1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
```

```
RP/0/RSP0RP0/CPU0:router(config-isis-af)# metric-style wide level 1
```

mpls ldp auto-config

To enable Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) interface auto-configuration, use the **mpls ldp auto-config** command in IPv4 address family configuration mode. To disable LDP IGP auto-configuration, use the **no** form of this command.

mpls ldp auto-config

Syntax Description

This command has no keywords or arguments.

Command Default

LDP IGP auto-configuration is disabled.

Command Modes

IPv4 address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **mpls ldp auto-config** command to automatically configure LDP on a set of interfaces associated with a specified IGP instance. Further, LDP IGP auto-configuration provides a means to block LDP from being enabled on a specified interface. If you do not want an IS-IS interface to have LDP enabled, use the **igp auto-config disable** command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable LDP IGP auto-configuration:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# mpls ldp auto-config
```

mpls ldp sync (IS-IS)

To configure Label Distribution Protocol (LDP) IS-IS synchronization, use the **mpls ldp sync** command in interface address family configuration mode. To disable LDP synchronization, use the **no** form of this command.

mpls ldp sync [**level** { **1** | **2** }]

Syntax Description

level { **1** | **2** }	(Optional) Sets LDP synchronization for the specified level.

Command Default

If a level is not specified, LDP synchronization is set for both levels.

Command Modes

Interface address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

MPLS VPN traffic forwarded using LDP labels can be dropped in the following instances:

- A new link is introduced in the network and IS-IS has converged before LDP establishes labels.
- An existing LDP session goes down while IS-IS adjacency is intact over the link.

In both instances, outbound LDP labels are not available for forwarding MPLS traffic. LDP IS-IS synchronization addresses the traffic drop. When the **mpls ldp sync** command is configured, IS-IS advertises the maximum possible link metric until LDP has converged over the link. The link is less preferred and least used in forwarding MPLS traffic. When LDP establishes the session and exchanges labels, IS-IS advertises the regular metric over the link.



Note IS-IS advertises the maximum metric –1 (16777214) if wide metrics are configured since the maximum wide metric is specifically used for link exclusion from the shortest path first algorithm (SPF) (RFC 3784). However, the maximum narrow metric is unaffected by this definition.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable LDP IS-IS synchronization:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/3/0/0
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# mpls ldp sync
```

mpls traffic-eng (IS-IS)

To configure a router running the Intermediate System-to-Intermediate System (IS-IS) protocol to flood Multiprotocol Label Switching traffic engineering (MPLS TE) link information into the indicated IS-IS level, use the **mpls traffic-eng** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng { **level-1** | **level-1-2** | **level-2-only** }

Syntax Description

level-1	Specifies routing level 1.
level-1-2	Specifies routing levels 1 and 2.
level-2-only	Specifies routing level 2.

Command Default

Flooding is disabled.

Command Modes

IPv4 address family configuration

Usage Guidelines

Use the **mpls traffic-eng** command, which is part of the routing protocol tree, to flood link resource information (such as available bandwidth) for appropriately configured links in the link-state packet (LSP) of the router.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to turn on MPLS traffic engineering for IS-IS level 1:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# mpls  traffic-eng level-1
```

mpls traffic-eng router-id (IS-IS)

To specify the Multiprotocol Label Switching traffic engineering (MPLS TE) router identifier for the node, use the **mpls traffic-eng router-id** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id { *ip-address* | *type* *interface-path-id* }

Syntax Description

ip-address	IP address in four-part, dotted-decimal notation.
type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

Global router identifier is used.

Command Modes

IPv4 address family configuration

Usage Guidelines

The identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router ID of the destination node, because that is the address used by the traffic engineering topology database at the tunnel head for its path calculation.

Note We recommend that loopback interfaces be used for MPLS TE, because they are more stable than physical interfaces.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# mpls traffic-eng router-id Loopback0
```

oor-set-overload-bit disable

To prevent the setting of the overload bit during out-of-resource (OOR) conditions, use **oor-set-overload-bit disable** command in router configuration mode. To allow the setting of the overload bit during OOR conditions, use the **no** form of the command.

oor-set-overload-bit disable

no oor-set-overload-bit disable

Syntax Description

This command has no keywords or arguments.

Command Default

The overload bit is set.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.2.1	This command was introduced.

Usage Guidelines

The setting of the overload bit is typically allowed only when a router experiences issues, such as memory shortages, which might result in an incomplete link-state database and inaccurate routing tables. When the overload bit is set in the link-state packets (LSPs) of an unreliable router, other routers can ignore this router in their shortest path first (SPF) calculations until it recovers. Consequently, no paths through the unreliable router are visible to other routers in the IS-IS area, though IP prefixes directly connected to this router remain reachable.

We recommend not disabling the setting of the overload-bit during OOR conditions. IS-IS will continue to forward traffic but will base decisions on an incomplete link-state database, potentially causing routing loops and other forwarding failures.

Task ID

Task ID	Operation
isis	read, write

Example

The following example shows how to disable the overload bit:

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

net

To configure an Intermediate System-to-Intermediate System (IS-IS) network entity title (NET) for the routing instance, use the **net** command in router configuration mode. To remove the **net** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

net *network-entity-title*

Syntax Description

network-entity-title	NET that specifies the area address and the system ID for an ISIS routing process.

Command Default

No NET is configured. The IS-IS instance is not operational, because a NET is mandatory.

Command Modes

Router configuration

Usage Guidelines

Under most circumstances, one and only one NET should be configured.

A NET is a network service access point (NSAP) where the last byte is always 0. On a Cisco router running IS-IS, a NET can be 8 to 20 bytes in length. The last byte is always the n-selector and must be 0. The n-selector indicates to which transport entity the packet is sent. An n-selector of 0 indicates no transport entity and means that the packet is for the routing software of the system.

The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2).

All bytes preceding the system ID are the area ID.

A maximum of three NETs for each router is allowed. In rare circumstances, it is possible to configure two or three NETs. In such a case, the area this router is in has three area addresses. Only one area still exists, but it has more area addresses.

Configuring multiple NETs can be temporarily useful in network reconfiguration in which multiple areas are merged, or in which one area is split into more areas. Multiple area addresses enable you to renumber an area individually as needed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a router with NET area ID 47.0004.004d.0001 and system ID 0001.0c11.1110:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
```

Related Commands

Command	Description
log adjacency changes (IS-IS), on page 39	Configures the routing level for an instance of the IS-IS routing process.
router isis, on page 79	Enables the IS-IS routing protocol and specifies an IS-IS instance.

nsf (IS-IS)

To enable nonstop forwarding (NSF) on the next restart, use the **nsf** command in XR Config mode. To restore the default setting, use the **no** form of this command.

nsf { **cisco** | **ietf** }

Syntax Description

cisco	Specifies Cisco-proprietary NSF restart.
ietf	Specifies Internet Engineering Task Force (IETF) NSF restart.

Command Default

NSF is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

NSF allows an Intermediate System-to-Intermediate System (IS-IS) instance to restart using checkpointed adjacency and link-state packet (LSP) information, and to perform restart with no impact on its neighbor routers. In other words, there is no impact on other routers in the network due to the destruction and recreation of adjacencies and the system LSP.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable Cisco proprietary NSF:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# nsf cisco
```

passive (IS-IS)

To suppress Intermediate System-to-Intermediate System (IS-IS) packets from being transmitted to the interface and received packets from being processed on the interface, use the **passive** command in interface configuration mode. To restore IS-IS packets coming to an interface, use the **no** form of this command.

passive

Command Default

Interface is active.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the router to suppress IS-IS packets on GigabitEthernet interface 0/1/0/1:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RP0/CPU0:router(config-isis-if)# passive
```

point-to-point

To configure a network of only two networking devices that use broadcast media and the integrated Intermediate System-to-Intermediate System (IS-IS) routing protocol to function as a point-to-point link instead of a broadcast link, use the **point-to-point** command in interface configuration mode. To disable the point-to-point usage, use the **no** form of this command.

point-to-point

Syntax Description

This command has no keywords or arguments.

Command Default

Interface is treated as broadcast if connected to broadcast media.

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **point-to-point** command only on broadcast media in a network with two networking devices. The command causes the system to issue packets point-to-point rather than as broadcasts. Configure the command on both networking devices in the network.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure a 10-Gb Ethernet interface to act as a point-to-point interface:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# interface HundredGigE 0/6/0/0
RP/0/RP0/CPU0:router(config-isis-if)# point-to-point
```

protocol shutdown (IS-IS)

To disable an instance of the Intermediate System-to-Intermediate System (IS-IS) protocol, so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the isis configuration mode. To re-enable the IS-IS protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default

None

Command Modes

isis configuration

Command History

Release	Modification
Release 24.4.1	This command was introduced.

Usage Guidelines

Use the **protocol shutdown** command to disable the IS-IS protocol for a specific routing instance without removing any existing IS-IS configuration parameters.

This command is similar to executing the following sample configuration:

```
router isis 1
 set-overload-bit
 exit
router isis 1
 interface HundredGigE 0/1/0/3
  shutdown
```

Task ID

Task ID	Operations
isis	read, write

Examples

This example shows how to disable the IS-IS.

```
Router# config
Router(config)# router isis 1
Router(config-isis)# protocol shutdown
```

psnp-interval

To configure the interval at which the Intermediate System to Intermediate System (IS-IS) protocol sends Partial Sequence Number PDUs (PSNPs) to acknowledge received Link State PDUs (LSPs), use the **psnp-interval** in router configuration mode.

psnp-interval *milliseconds*

Syntax Description

milliseconds	The interval, specified in milliseconds, at which PSNPs are sent to acknowledge received LSPs. The range is from 1 to 5000 milliseconds. The default value is 50 milliseconds.

Command Default

This command is disabled by default.

Command Modes

Router configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

None

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to configure the PSNP interval to 100 milliseconds for an IS-IS process tagged 1:

```
Router(config)# router isis 1
Router(config-isis)# psnp-interval 100
```

redistribute (IS-IS)

To redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS), use the **redistribute** command in address family configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)
redistribute bgp *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Connected Routes
redistribute connected [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Intermediate System-to-Intermediate System (IS-IS)
redistribute isis *process-id* [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*] [**down-flag-clear**]

Open Shortest Path First (OSPF)
redistribute ospf *process-id* [**level-1** | **level-2** | **level-1-2**] [**match** { **external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**] }] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Open Shortest Path First Version 3 (OSPFv3)
redistribute ospfv3 *process-id* [**level-1** | **level-2** | **level-1-2**] [**match** { **external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**] }] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Static Routes
redistribute static [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **1** { **internal** | **external** | **rib-metric-as-external** } | **2** **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

RIP
redistribute rip [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Subscriber Routes
redistribute subscriber [**level-1** | **level-2** | **level-1-2**] [**metric** *metric-value*] [**metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }] [**route-policy** *route-policy-name*]

Syntax Description

process-id	For the **bgp** keyword, an autonomous system number has the following ranges: • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. For the **isis** keyword, an IS-IS instance identifier from which routes are to be redistributed. For the **ospf** keyword, an OSPF process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the **ospfv3** keyword, an OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
level-1	(Optional) Specifies that redistributed routes are advertised in the Level-1 LSP of the router.
level-1-2	(Optional) Specifies that redistributed routes are advertised in the Level-1-2 LSP of the router.
level-2	(Optional) Specifies that redistributed routes are advertised in the Level-2 LSP of the router.
metric *metric-value*	(Optional) Specifies the metric used for the redistributed route. Range is 0 to 16777215. The *metric-value* must be consistent with the IS-IS metric style of the area and topology into which the routes are being redistributed.
metric-type { **internal** | **external** } **metric-type** { **internal** | **external** | **rib-metric-as-external** | **rib-metric-as-internal** }	(Optional) Specifies the external link type associated with the route advertised into the ISIS routing domain. It can be one of two four values: • external • **internal** –Use the **internal** keyword to set IS-IS internal metric-type. • **external** –Use the **external** keyword to set IS-IS external metric-type • rib-metric-as-external–Use the **rib-metric-as-external** keyword to use RIB metric and set IS-IS external metric-type • rib-metric-as-internal–Use the **rib-metric-as-internal** keyword to use RIB metric and set IS-IS internal metric-type Any route with an internal metric (however large the metric is) is preferred over a route with external metric (however small the metric is). Use the **rib-metric-as-external** and **rib-metric-as-internal** keywords to preserve RIB metrics when redistributing routes from another IS-IS router instance or another protocol.

route-policy *route-policy-name*	(Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to IS-IS.
match { **internal** | **external** [**1** | **2**] | **nsaa-external** [**1** | **2**]}	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal** —Routes that are internal to a specific autonomous system (intra- and interarea OSPF routes). • **external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1 | 2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.
down-flag-clear	(Optional) Specifies that routes redistributed from another IS-IS instance should be advertised with the up/down bit set to zero. This is contrary to the behavior specified by RFC 5305 and RFC 7775 and can lead to route loops.

Command Default

Level 2 is configured if no level is specified.

metric-type: internal

match : If no match keyword is specified, all OSPF routes are redistributed.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Note When redistributing routes (into IS-IS) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Use the **redistribute** command to control the redistribution of routes between separate IS-IS instances. To control the propagation of routes between the levels of a single IS-IS instance, use the propagate level command.

Only IPv4 OSPF addresses can be redistributed into IS-IS IPv4 address families and only IPv6 OSPFv3 prefixes can be distributed into IS-IS IPv6 address families.

Examples

In this example, IS-IS instance isp_A readvertises all of the routes of IS-IS instance isp_B in Level 2 LSP. Note that the **level-2** keyword affects which levels instance isp_A advertises the routes in and has no impact on which routes from instance isp_B are advertised. (Any Level 1 routes from IS-IS instance isp_B are included in the redistribution.

```
RP/0/RSP0RP0/CPU0:router(config)# router isis isp_A
RP/0/RSP0RP0/CPU0:router(config-isis)# net 49.1234.2222.2222.2222.00
RP/0/RSP0RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0RP0/CPU0:router(config-isis-af)# redistribute isis isp_B level-2
!
RP/0/RSP0RP0/CPU0:router(config)# router isis isp_B
RP/0/RSP0RP0/CPU0:router(config-isis)# is-type level 1
RP/0/RSP0RP0/CPU0:router(config-isis)# net 49.4567.2222.2222.2222.00
RP/0/RSP0RP0/CPU0:router(config-isis)# address-family ipv4 unicast
```

remote-psnp-delay

To specify the maximum delay, in milliseconds, that the router should wait for a Partial Sequence Number Protocol (PSNP) acknowledgment from neighbors after sending a Link State Packet (LSP), use the **remote-psnp-delay** command in appropriate mode. To remove the configured maximum delay, use the **no** form of this command.

remote-psnp-delay *milliseconds*

Syntax Description

milliseconds	Maximum delay, in milliseconds, that the router should wait for a PSNP acknowledgment from neighbors after sending a LSP. The range is from 1 through 5000 ms.

Command Default

The command is disabled by default.

Command Modes

IS-IS router lsp-fast-flooding configuration

Interface configuration

Command History

Release	Modification
Release 24.3.1	This command was introduced.

Usage Guidelines

The order of precedence of the configuration using the **remote-psnp-delay** command is as follows:

1. The **remote-psnp-delay** command in interface configuration mode.
2. The value in the PSNP sub-TLV received from the neighbour, if any.
3. The **remote-psnp-delay** command in router configuration mode.

Task ID

Task ID	Operation
isis	read, write

Example

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on a device.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# lsp-fast-flooding
Router(config-isis-lsp-fast-flooding)# remote-psnp-delay 600
Router(config-isis-lsp-fast-flooding)# commit
```

This example shows how to set the PSNP acknowledgment delay to 600 milliseconds on an interface.

```
Router# configure terminal
Router(config)# router isis 1
Router(config-isis)# interface GigabitEthernet0/0/0/0
Router(config-isis-if)# remote-psnp-delay 600
```

router isis

To enable the Intermediate System-to-Intermediate System (IS-IS) routing protocol and to specify an IS-IS instance, use the **router isis** command in mode. To disable IS-IS routing, use the **no** form of this command.

router isis *instance-id*

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

An IS-IS routing protocol is not enabled.

Command Modes

Usage Guidelines

Use the **router isis** command to create an IS-IS routing process. An appropriate network entity title (NET) must be configured to specify the address of the area (Level 1) and system ID of the router. Routing must be enabled on one or more interfaces before adjacencies may be established and dynamic routing is possible.

Multiple IS-IS processes can be configured. Up to eight processes are configurable. A maximum of five IS-IS instances on a system are supported.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IS-IS for IP routing:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0001.0000.0001.00
```

Related Commands

Command	Description
net, on page 66	Configures an IS-IS NET for the routing process.

router isis static neighbor

To enable the Intermediate System-to-Intermediate System (IS-IS) static neighbor and to configure its attributes, use the **router isis static-neighbor** command in XR Config mode. To disable IS-IS static neighbor, use the **no** form of this command.

router isis *instance-id* **static-neighbor**

Syntax Description

instance-id	Name of the routing process. Maximum number of characters is 40.

Command Default

None

Command Modes

XR Config mode

Command History

Release	Modification
25.2.1	This command was introduced.

Usage Guidelines

You can configure the following attributes by using the **static-neighbor** configuration mode.

system-id	ISIS System-id of the static neighbor.
remote ipv4/ipv6 address	Remote IPv4/IPv6 addresses.
affinity	Affinity link to advertise.
log advertisement changes	Enable logging static neighbor advertisement changes.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure IS-IS static neighbor and its attributes for a router:

```
Router#config
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 2.2.2.2
Router(config-isis-static-nbr)#affinity RED
Router(config-isis-static-nbr)#log-advertisement-changes
```

set default-isis-instance

When multiple IS-IS instances are configured on a router, by default, **show isis** commands display information from all IS-IS instances. To display information from only one specific IS-IS instance, use the **set default-isis-instance** *instance-name* command in the XR EXEC mode.

When the default IS-IS instance is set, all the show IS-IS commands during the terminal session display only information from that default IS-IS instance. To revert back to the default behaviour of displaying information from all IS-IS instances, use the **set default-isis-instance all** command.

set default-isis-instance { *instance-name* | **all** }

Syntax Description

instance-name	Displays information for the specified *instance-name* when running **show isis** commands.
all	Displays information from all configured IS-IS instances when running **show isis** commands.

Command Default

By default, all configured IS-IS instances are displayed when you run any **show isis** command.

Command Modes

XR EXEC mode

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

Task ID	Operations
isis	read, write

Examples

Example 1: This example shows how to set a specific IS-IS instance as the default instance.

```
RP/0/1/CPU0:ios#set default-isis-instance india
Tue Apr 23 11:47:04.353 CEST
%% Default isis instance is set to  india
```

Example 2: This example uses the **show isis summary** command to display information from the default instance set in Example 1.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:10.802 CEST

                                 (L1/L2)     (L1/L2)       (L1/L2)       (IPv4/IPv6)
Tag                    Level     Adj         Routers       LSPs          Routes
---                    -----     ---         -------       ----          ------
india                  L2        0/1         0/2           0/2           0/0
RP/0/1/CPU0:ios#
```

Example 3: This example shows how to set all configured IS-IS instances as default instances.

```
RP/0/1/CPU0:ios#set default-isis-instance all
Tue Apr 23 11:47:19.308 CEST
%% Default isis instance is set to  all
```

Example 4: This example uses the **show isis summary** command to display information from all the configured instances set in Example 3.

```
RP/0/1/CPU0:ios#show isis summary
Tue Apr 23 11:47:24.697 CEST

                                (L1/L2)     (L1/L2)      (L1/L2)      (IPv4/IPv6)
Tag                    Level    Adj         Routers      LSPs         Routes
---                    -----    ---         -------      ----         ------
eu                     L2       0/2         0/6          0/6          5/5
india                  L2       0/1         0/2          0/2          0/0
RP/0/1/CPU0:ios#
```

set-overload-bit

To configure the router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations, use the **set-overload-bit** command in XR Config mode. To remove the designation, use the **no** form of this command.

set-overload-bit [**on-startup** { *delay* | **wait-for-bgp** }] [**level** { **1** | **2** }] [**advertise** { **external** | **interlevel** }]

Syntax Description

on-startup	(Optional) Sets the overload bit only temporarily after reboot.
delay	(Optional) Time (in seconds) to advertise when the router is overloaded after reboot. Range is 5 to 86400 seconds (86400 seconds = 1 day).
wait-for-bgp	(Optional) Sets the overload bit on startup until the Border Gateway Protocol (BGP) signals converge or time out.
level { **1** | **2** }	(Optional) Specifies the overload bit for Level 1 or Level 2 independently.
advertise { **external** | **interlevel**	(Optional) Sets the overload bit set if the router advertises the following types of IP prefixes: • **external**—If overload-bit set advertises IP prefixes learned from other protocols • **interlevel**— If overload-bit set advertise IP prefixes learned from another ISI S level

Command Default

The overload bit is not set.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **set-overload-bit** command to force the router to set the overload bit in its nonpseudonode link-state packets (LSPs). Normally the setting of the overload bit is allowed only when a router experiences problems. For example, when a router is experiencing a memory shortage, the reason might be that the link-state database is not complete, resulting in an incomplete or inaccurate routing table. If the overload bit is set in the LSPs of the unreliable router, other routers can ignore the router in their SPF calculations until it has recovered from its problems. The result is that no paths through the unreliable router are seen by other routers in the Intermediate System-to-Intermediate System (IS-IS) area. However, IP prefixes directly connected to this router are still reachable.

The **set-overload-bit** command can be useful when you want to connect a router to an IS-IS network, but do not want real traffic flowing through it under any circumstances.

Routers with overload bit set are:

- A test router in the lab, connected to a production network.
- A router configured as an LSP flooding server, for example, on a nonbroadcast multiaccess (NBMA) network, in combination with the mesh group feature.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to configure the overload bit:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# set-overload-bit
```

show isis

The **show isis** command displays general information about an IS-IS instance and protocol operation. If the instance ID is not specified, the command shows information about all IS-IS instances.

show isis [**instance** *instance-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. **Note** The instance-id argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

For each instance, the first line of output lists the IS-IS instance ID with the following lines identifying the IS-IS system ID, supported levels (level 1, level 2, or level-1-2), configured area addresses, active area addresses, status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF), and the mode in which the last IS-IS process startup occurred.

Next, the status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not). Then redistributed protocols are listed, followed by the administrative distance applied to the redistributed routes.

Finally, the running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface is listed.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis** command:

```
RP/0/RP0/CPU0:router# show isis
Wed Aug 20 23:54:55.043 PST DST

IS-IS Router: lab
  System Id: 0000.0000.0002
  IS Levels: level-2-only
  Manual area address(es):
    49.1122
```

```
Routing for area address(es):
  49.1122
Non-stop forwarding: Disabled
Most recent startup mode: Cold Restart
Topologies supported by IS-IS:
  IPv4 Unicast
    Level-2
      Metric style (generate/accept): Narrow/Narrow
      Metric: 10
      ISPF status: Disabled
    No protocols redistributed
    Distance: 115
Interfaces supported by IS-IS:
  Loopback0 is running passively (passive in configuration)
  POS0/1/0/2 is running actively (active in configuration)
  POS0/1/0/3 is running actively (active in configuration
```

This table describes the significant fields shown in the display.

Table 1: show isis Field Descriptions

Field	Description
IS-IS Router	IS-IS instance ID.
System Id	IS-IS system ID.
IS Levels	Supported levels for the instance.
Manual area address(es)	Domain and area.
Routing for area address(es):	Configured area addresses and active area addresses.
Non-stop forwarding	Status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF).
Most recent startup mode	The mode in which the last IS-IS process startup occurred.
Topologies supported by IS-IS	The summary of the status of each configured address family (or just IPv4 unicast if none are configured).
Redistributed protocols	List of redistributed protocols, followed by the administrative distance applied to the redistributed routes.
Metric style (generate/accept)	The status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not).
Interfaces supported by IS-IS	The running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface.

show isis adjacency

To display Intermediate System-to-Intermediate System (IS-IS) adjacencies, use the **show isis adjacency** command in XR EXEC mode.

show isis [**instance** *instance-id*] **adjacency** [**level** {**1** | **2**}] [*type interface-path-id*] [**detail**] [**systemid** *system-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS adjacencies for Level 1 or Level 2 independently.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
detail	(Optional) Displays neighbor IP addresses and active topologies.
systemid *system-id*	(Optional) Displays the information for the specified router only.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis adjacency** command:

```
RP/0/RP0/CPU0:router# show isis adjacency

IS-IS p Level-1 adjacencies:
System Id      Interface        SNPA           State Hold     Changed  NSF      BFD
12a4           PO0/1/0/1        *PtoP*         Up    23       00:00:06 Capable  Init
12a4           Gi0/6/0/2        0004.2893.f2f6 Up    56       00:04:01 Capable  Up

Total adjacency count: 2

IS-IS p Level-2 adjacencies:
System Id      Interface        SNPA           State Hold     Changed  NSF      BFD
12a4           PO0/1/0/1        *PtoP*         Up    23       00:00:06 Capable  None
12a4           Gi0/6/0/2        0004.2893.f2f6 Up    26       00:00:13 Capable  Init

Total adjacency count: 2
```

This table describes the significant fields shown in the display.

Table 2: show isis adjacency Field Descriptions

Field	Description
Level-1	Level 1 adjacencies.
Level-2	Level 2 adjacencies.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or the **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Interface	Interface used to reach the neighbor.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
State	Adjacency state of the neighboring interface. Valid states are Down, Init, and Up.
Holdtime	Hold time of the neighbor.
Changed	Time the neighbor has been up (in hours:minutes:seconds).
NSF	Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism.
BFD	Specifies the Bidirectional Forwarding Detection (BFD) status for the interface. Valid status are: • None—BFD is not configured. • Init—BFD session is not up. One reason is that other side is not yet enabled. • Up—BFD session has been established. • Down—BFD session holdtime expired.

show isis adjacency-log

To display the Intermediate System-to-Intermediate System (IS-IS) adjacency log, use the **show isis adjacency-log** command in XR EXEC mode.

show isis adjacency-log [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

level { **1** | **2** }	(Optional) Displays the IS-IS adjacency log for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output is restricted to the last *number* of entries. Range is 1 to 100.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 100.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis adjacency-log** command:

```
RP/0/RP0/CPU0:router# show isis adjacency-log

   IS-IS 10 Level 1 Adjacency log
When           System          Interface          State   Details
4d00h          12a1            PO0/5/0/0          d -> i
4d00h          12a1            PO0/5/0/0          i -> u  New adjacency
                                                          IPv4 Unicast Up
4d00h          12a1            Gi0/6/0/0          d -> u  New adjacency
4d00h          12a1            Gi0/6/0/0          u -> d  Interface state
down
3d17h          12a1            Gi0/6/0/0          d -> u  New adjacency
3d17h          12a1            Gi0/6/0/0          u -> d  Interface state
down
01:44:07       12a1            Gi0/6/0/0          d -> u  New adjacency

   IS-IS 10 Level 2 Adjacency log
When           System          Interface          State   Details
4d00h          12a1            PO0/5/0/0          d -> i
4d00h          12a1            PO0/5/0/0          i -> u  New adjacency
```

```
                                                                IPv4 Unicast Up
4d00h           12a1            Gi0/6/0/0           d -> u  New adjacency
4d00h           12a1            Gi0/6/0/0           u -> d  Interface state
down
3d17h           12a1            Gi0/6/0/0           d -> u  New adjacency
3d17h           12a1            Gi0/6/0/0           u -> d  Interface state
down
01:44:07        12a1            Gi0/6/0/0           d -> u  New adjacency
```

This table describes the significant fields shown in the display.

Table 3: show isis adjacency-log Field Descriptions

Field	Description
When	Elapsed time (in hh:mm:ss) since the event was logged.
System	System ID of the adjacent router.
Interface	Specific interface involved in the adjacency change.
State	State transition for the logged event.
Details	Description of the adjacency change.

show isis checkpoint adjacency

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint adjacency database, use the **show isis checkpoint adjacency** command in XR EXEC mode.

show isis [**instance** *instance-id*] **checkpoint adjacency**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS checkpoint adjacencies for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint adjacencies for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show isis checkpoint adjacency** command to display the checkpointed adjacencies. With this information you can restore the adjacency database during a Cisco proprietary nonstop forwarding (NSF) restart. This command, with the **show isis adjacency** command, can be used to verify the consistency of the two databases.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint adjacency** command:

```
RP/0/RP0/CPU0:router# show
          isis
          checkpoint
          adjacency

 Interface        Level  System ID        State  Circuit ID          Chkpt ID
 Gi3/0/0/1           1    router-gsr8       Up     0001.0000.0008.04   80011fec
 Gi0/4/0/1           1    router-gsr9       Up     0001.0000.0006.01   80011fd8
 Gi3/0/0/1           2    router-gsr8       Up     0001.0000.0008.04   80011fc4
```

This table describes the significant fields shown in the display.

Table 4: show isis checkpoint adjacency Field Descriptions

Field	Description
Interface	Interface used to reach the neighbor.
Level	Lists either routers with Level 1 or Level 2 adjacency configured.
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
State	State of the neighboring interface.
Circuit ID	Unique ID issued to a circuit at its creation.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

show isis checkpoint interface

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint interfaces, use the **show isis checkpoint interface** command in XR EXEC mode.

show isis checkpoint interface

This command has no keywords or arguments.

Command Default

No default behavior or values

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint interface** command:

```
RP/0/RP0/CPU0:router# show isis checkpoint interface

IS-IS 10 checkpoint interface
Interface         Index  CircNum  DIS Areas    Chkpt ID
PO0/5/0/0         0      0        NONE         80002fe8
Gi0/6/0/0         1      3        L1L2         80002fd0
```

This table describes the significant fields shown in the display.

Table 5: show isis checkpoint interface Field Descriptions

Field	Description
Interface	Interface used to reach the neighbor.
Index	Interface index assigned to an interface upon its creation.
CircNum	Unique ID issued to a circuit internally.
DIS Areas	Designated Intermediate System area.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

show isis checkpoint lsp

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint link-state packet (LSP) protocol data unit (PDU) identifier database, use the **show isis checkpoint lsp** command in XR EXEC mode.

show isis [**instance** *instance-id*] **checkpoint lsp**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS checkpoint LSPs for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint LSPs for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The checkpointed LSPs displayed by this command are used to restore the LSP database during a Cisco-proprietary nonstop forwarding (NSF) restart. The **show isis checkpoint lsp** command, with the **show isis database** command, may be used to verify the consistency of the two databases.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis checkpoint lsp** command:

```
RP/0/RP0/CPU0:router#
            show isis checkpoint lsp


Level  LSPID              Chkpt ID
1      router-gsr6.00-00  80011f9c
1      router-gsr6.01-00  80011f88
1      router-gsr8.00-00  80011f74
1      router-gsr9.00-00  80011f60
2      router-gsr6.00-00  80011f4c
2      router-gsr6.01-00  80011f38
2      router-gsr8.00-00  80011f24
2      router-gsr9.00-00  80011f10
Total LSP count: 8 (L1: 4, L2 4, local L1: 2, local L2 2)
```

This table describes the significant fields shown in the display.

Table 6: show isis checkpoint lsp Field Descriptions

Field	Description
Level	Routers with Level 1 or Level 2 adjacency configured.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
Chkpt ID	Unique ID issued to the checkpoint at its creation.

show isis database

To display the Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) database, use the **show isis database** command in XR EXEC mode.

show isis [**instance** *instance-id*] **database** [**level** {**1** | **2**}] [**update**] [**summary**] [**detail**] [**verbose**] [*****|*lsp-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS LSP database for the specified instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the IS-IS LSP database for Level 1 or Level 2 independently.
update	(Optional) Displays contents of LSP database managed by update thread.
summary	(Optional) Displays the LSP ID number, sequence number, checksum, hold time, and bit information.
detail	(Optional) Displays the contents of each LSP.
verbose	(Optional) Displays the contents of each LSP.
* | *lsp-id*	(Optional) LSP protocol data units (PDUs) identifier. Displays the contents of a single LSP by its ID number or may contain an * as a wildcard character.

Command Default

No instance ID specified displays the IS-IS LSP database for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Each of the options for the **show isis database** command can be entered in an arbitrary string within the same command entry. For example, the following are both valid command specifications and provide the same output: **show isis database detail level 2** and **show isis database level 2 detail** .

The **summary** keyword used with this command allows you to filter through a large IS-IS database and quickly identify problematic areas.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis database** command with no keywords specified:

```
RP/0/RP0/CPU0:router# show isis database

IS-IS Area a1 (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
router-gsr6.00-00   * 0x00000016   0x62c8        896             0/0/0
router-gsr6.01-00   * 0x0000000f   0x56d9        902             0/0/0
router-gsr8.00-00     0x00000019   0x4b6d        1015            0/0/0
router-gsr9.00-00     0x00000016   0x33b7        957             0/0/0


 Total LSP count: 4 (L1: 4, L2 0, local L1: 2, local L2 0)
```

The following sample output shows the remaining lifetime value that is received in LSP database. The received value is shown in the output under the **Rcvd** field.

```
RP/0/0/CPU0:ios#show isis database
Thu Dec 14 16:03:45.131 EST

IS-IS 100 (Level-1) Link State Database
LSPID                 LSP Seq Num   LSP Checksum  LSP Holdtime/Rcvd  ATT/P/OL
1111.1111.1111.00-00* 0x0073000e     0x0001                1002/*           1/0/1
1111.1111.1112.00-00  0x0073004a     0x0001                1195/59          1/0/1
1111.1111.1112.01-00  0x007300b8     0x0001                1178/59          0/0/1
1111.1111.1112.03-00  0x007300b6     0x0001                1179/59          0/0/1
1111.1111.1113.00-00  0x0073000d     0x0001                533/1200         1/0/1

 Total Level-1 LSP count: 5     Local Level-1 LSP count: 1
```

This table describes the significant fields shown in the display.

Table 7: show isis database Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.

Field	Description
LSP Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
ATT/P/OL	ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.

The following is sample output from the **show isis database** command with the **summary** keyword:

```
RP/0/RP0/CPU0:router# show isis database summary

  IS-IS 10 Database Summary for all LSPs
                                Active              Purged               All
                           L1    L2   Total   L1    L2   Total   L1    L2   Total
                          ----- ----- -----  ----- ----- -----  ----- ----- -----
  Fragment 0 Counts
              Router LSPs:     1     1     2      0     0     0      1     1     2
        Pseudo-node LSPs:     0     0     0      0     0     0      0     0     0
                 All LSPs:     1     1     2      0     0     0      1     1     2
  Per Topology
      IPv4 Unicast
        ATT bit set LSPs:     0     0     0      0     0     0      0     0     0
        OVL bit set LSPs:     0     0     0      0     0     0      0     0     0

  All Fragment Counts
              Router LSPs:     1     1     2      0     0     0      1     1     2
        Pseudo-node LSPs:     0     0     0      0     0     0      0     0     0
                 All LSPs:     1     1     2      0     0     0      1     1     2
```

This table describes the significant fields shown in the display.

Table 8: show isis database summary Field Descriptions

Field	Description
Router LSPs:	Active, purged, and total LSPs associated with routers.
Pseudo-node LSPs:	Active, purged, and total LSPs associated with pseudonodes.
All LSPs:	Total active and purged LSPs.
ATT bit set LSPs	Attach bit (ATT). Indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router.

Field	Description
OVL bit set LSPs	Overload bit. Indicates if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.

The following is sample output from the **show isis database** command with the **detail verbose** keyword specified:

```
RP/0/RP0/CPU0:router# show isis instance isp database detail verbose

  IS-IS isp test (Level-1) Link State Database
      LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
      router-5router1.00-00       0x00000003 * 0x00000d44 0x8074460 0x7e2c 457 535
     0/0/0
        Area Address: 4900
    Area Address: 01
        NLPID:        0xcc
    NLPID:  Hostname:    router-50x8e
        IP Address:  172.1.1.5
        MetricMT: 0  IP 172.3.55.0/24Standard (IPv4 Unicast)
        Metric: 10         IP 172.6.1.0/24
    MT:           IPv6 Unicast                                       0/0/0
    MT:           IPv4 Multicast                  Metric: 10 IP 172.7.0./0/240
        Metric: 10         IS router-11.00
        Metric: 10         IS router-11.01
    MT:           IPv6 Multicast                router-11.00-00  * 0x0000000b   0x8074460
       1161            0/0/0
    Hostname:  Area Address: 49router1
        NLPID:       0xcc
        Hostname:    router-11
    IP Address:   192.168.0.145
        IP IPv6 Address:  172.1.11.11192:168::145     MetricRouter ID: 0        IP
172192.1168.1110.0/24145
        Metric: 10         IP 172IS-Extended router1.016.1.0/24
        Metric: 10         IP 172IS-Extended router2.007.0.0/24
        Metric: 10         IS routerIS-11Extended router2.0100
        Metric: 10         IS router-5.00
      router-11.01-00  * 0x00000001   0x80770ec        457             0/0/0
        Metric: 0          IS router-11.00
        Metric: 0          IS router-5.00
      Affinity: 0x00000000
      Interface IP Address: 10.3.11.145
      Neighbor IP Address: 10.3.11.143
      Physical BW: 155520 kbits/sec
       Total LSP count: 3 (L1: 3, L2 0, local L1: 2, local L2 0)
      Reservable Global pool BW: 0 kbits/sec
      Global Pool BW Unreserved:
      IS-IS isp (Level-2) Link State Database
      LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
      router-5.00-00     0x00000005   0x807997c        457             0/0/0
        [0]: 0        kbits/sec          [1]: 0        kbits/sec
        [2]: 0        kbits/sec          [3]: 0        kbits/sec
        [4]: 0        kbits/sec          [5]: 0        kbits/sec
        [6]: 0        kbits/sec          [7]: 0        kbits/sec
    MPLS SRLG: Area Address: 49router2.00
      Interface IP Address: 10.3.11.145
      Neighbor IP Address: 10.3.11.143
        NLPIDFlags:       0xcc0x1   HostnameSRLGs:    router-5IP Address[0]:  172.6.10,
[1.5]: 20
        Metric: 0 10  IP 172IP-Extended 10.3.5511.0/24
```

```
   Metric: 10         IP 172IP-Extended 192.1686.10.0145/2432
   Metric: 10         IS routerMT (IPv6 Unicast) IS-11Extended router1.0001
   Metric: 10         IP 172.1.0.0MT (IPv6 Unicast) IPv6 192:168::145/24128
   Metric: 10         IS routerMT (IPv4 Multicast) IS-11Extended router1.01
   Metric: 10         IP 172.8.111.0/24


 router-11.00-00  * 0x0000000d   0x807997c        1184            0/0/0
   Area Address: 49
   NLPID:       0xcc
   Hostname:    router-11
   IP Address:  172.28.111.111
   Metric: 0          IP 172.8.111.0/24
   Metric: 10         IP 172.6.1.0/24
   Metric: 10         IP 172MT (IPv4 Multicast) IP-Extended 192.7168.0./
   Metric: 10         IS router-11.01
   Metric: 10         IS router-5.00
   Metric: 10         IP 172.3.55.0MT (IPv6 Multicast) IPv6 192:168::145/24.01-00
0x0000013e 0x80770ec 0x3309 457 1159             0/0/0
   Metric: 0          IS routerIS-11Extended router1.00
   Metric: 0          IS routerIS-5Extended router2.00

  Total LSP count: 3 (L1: 0, L2 3, local L1: 0, local L2 2)
```

As the output shows, besides the information displayed with the **show isis database** command, the command with the **detail verbose** keyword displays the contents of each LSP.

Table 9: show isis instance isp database detail Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Amount of time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.

Field	Description
ATT/P/OL	ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area this route belongs to.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address:	Address of the node.
Metric	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a Connectionless Network Service (CLNS) prefix).

The following is additional sample output from the **show isis database detail** command. This is a Level 2 LSP. The area address 39.0001 is the address of the area in which the router resides.

```
RP/0/RSP0/CPU0:router# show isis database level 2 detail

IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.1111.00-00* 0x00000006   0x4DB3        1194          0/0/0
  Area Address: 39.0001
  NLPID:       0x81 0xCC
  IP Address:  172.18.1.17
  Metric: 10   IS 0000.0C00.1111.09
  Metric: 10   IS 0000.0C00.1111.08
  Metric: 10   IP 172.17.4.0 255.255.255.0
  Metric: 10   IP 172.18.8.0 255.255.255.0
  Metric: 0    IP-External 10.0.0.0 255.0.0.0
```

The IP entries are the directly connected IP subnets the router is advertising (with associated metrics). The IP-External entry is a redistributed route.

Table 10: show isis database level 2 detail Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
ATT/P/OL	ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area to which this route belongs.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address:	IP address of the node.
Metric:	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a Connectionless Network Service (CLNS) prefix).

Table 11: show isis database verbose Field Descriptions

Field	Description
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
LSP Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of zero indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
ATT/P/OL	ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair. OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.
Area Address	Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area to which this route belongs.
NLPID	Network Layer Protocol Identifier.
Hostname	Hostname of the node.
IP Address	IP address of the node.
Metric	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a Connectionless Network Service (CLNS) prefix).

Field	Description
MPLS SRLG	MPLS SRLG TLV information per neighbor, identified by hostname or system ID.
Interface IP Address	Local interface IP address.
Neighbor IP Address	Remote interface IP address.
Flags	Flags carried in SRLG TLV. The Least Significant Bit (LSB) is set if the interface is numbered.
SRLGs	SRLG values.

This is the sample output from the **show isis database verbose** command. The output shows IPv4 adjacency segment ID (SID), prefix (node) SID, and Segment Routing Global Block (SRGB) values.

```
RP/0/RP0/CPU0:router show isis database verbose
Fri May  2 17:53:44.575 PDT

IS-IS DEFAULT (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
p1.00-00              0x00000080   0x4780        1044            1/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     1/0/0
  Hostname:     p1
  IP Address:   172.16.255.101
  IPv6 Address: 2001:db8::ff:101
  Router Cap:   172.16.255.101, D:0, S:0
    Segment Routing: I:1 V:0, SRGB Base: 16000 Range: 7999
  Metric: 10         IS-Extended p2.00
    Interface IP Address: 172.16.2.4
    Neighbor IP Address: 172.16.2.5
    ADJ-SID: F:0 B:0 weight:0 Adjacency-sid:24002
  Metric: 10         IS-Extended pe1.00
    Interface IP Address: 172.16.1.1
    Neighbor IP Address: 172.16.1.0
    ADJ-SID: F:0 B:0 weight:0 Adjacency-sid:24003
  Metric: 10         IP-Extended 172.16.1.0/31
  Metric: 10         IP-Extended 172.16.2.2/31
  Metric: 10         IP-Extended 172.16.2.4/31
  Metric: 10         IP-Extended-Interarea 172.16.255.2/32
    Admin. Tag: 255
    Prefix-SID Index: 42, R:1 N:0 P:1
  Metric: 0          IP-Extended 172.16.255.101/32
    Prefix-SID Index: 141, R:0 N:0 P:0
  Metric: 10         MT (IPv6 Unicast) IS-Extended p2.00
  Metric: 10         MT (IPv6 Unicast) IS-Extended pe1.00
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::1:0/127
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::2:2/127
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:db8::2:4/127
  Metric: 10         MT (IPv6 Unicast) IPv6-Interarea 2001:db8::ff:2/128
    Admin. Tag: 255
  Metric: 0          MT (IPv6 Unicast) IPv6 2001:db8::ff:101/128
```

show isis database-log

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) database log, use the **show isis database-log** command in XR EXEC mode.

show isis database-log [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

level { **1** | **2** }	(Optional) Displays the database log for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output be restricted to the last *number* of entries. Range is 1 to 1000.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 1000.

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis database-log** command:

```
RP/0/RP0/CPU0:router# show isis database-log

   IS-IS 10 Level 1 Link State Database Log
                                  New LSP                 Old LSP
WHEN     LSPID                Op  Seq Num    Holdtime OL  Seq Num    Holdtime OL
01:17:19 12b1.03-00           REP 0x00000003 1200      0  0x00000002 340       0
001:06:20 12b1.00-00          REP 0x000001d8 1200      0  0x000001d7 375       0
01:06:00 12b1.03-00           REP 0x00000004 1200      0  0x00000003 520       0
01:05:46 12a1.00-00           REP 0x000001fc 1200      0  0x000001fb 425       0
00:55:01 12b1.00-00           REP 0x000001d9 1200      0  0x000001d8 520       0
00:53:39 12b1.03-00           REP 0x00000005 1200      0  0x00000004 459       0
00:53:19 12a1.00-00           REP 0x000001fd 1200      0  0x000001fc 453       0
00:42:12 12b1.00-00           REP 0x000001da 1200      0  0x000001d9 431       0
00:39:56 12b1.03-00           REP 0x00000006 1200      0  0x00000005 376       0
00:38:54 12a1.00-00           REP 0x000001fe 1200      0  0x000001fd 334       0
00:29:10 12b1.00-00           REP 0x000001db 1200      0  0x000001da 418       0
00:27:22 12b1.03-00           REP 0x00000007 1200      0  0x00000006 446       0
00:25:10 12a1.00-00           REP 0x000001ff 1200      0  0x000001fe 375       0
```

```
00:17:04 12b1.00-00           REP 0x000001dc 1200      0  0x000001db 473
```

This table describes the significant fields shown in the display.

Table 12: show isis database-log Field Descriptions

Field	Description
WHEN	Elapsed time (in hh:mm:ss) since the event was logged.
LSPID	LSP identifier. The first six octets form the system ID of the router that originated the LSP. The next octet is the pseudonode ID. When this byte is 0 zero, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router. For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN. The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.
New LSP	New router or pseudonode appearing in the topology.
Old LSP	Old router or pseudonode leaving the topology.
Op	Operation on the database: inserted (INS) or replaced (REP).
Seq Num	Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source.
Holdtime	Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed.
OL	Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.

show isis fast-reroute

To display per-prefix LFA information, use the **show isis fast-reroute** command in XR EXEC mode.

show isis fast-reroute
A.B.C.D/length | **detail** | **summary** | **sr-only**

Syntax Description

A.B.C.D/length	Network to show per-prefix LFA information.
detail	Use to display tiebreaker information about the backup.
summary	Use to display the number of prefixes having protection per priority.
sr-only	Use to display SR-labeled prefixes only.

Command Default

None

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

The following is sample output from **show isis fast-reroute** command that displays per-prefix LFA information:

```
RP/0/RP0/CPU0:router# show isis fast-reroute 10.1.6.0/24

L1 10.1.6.0/24 [20/115]
     via 10.3.7.47, POS0/3/0/1, router2
       FRR backup via 10.1.7.145, GigabitEthernet0/1/0/3, router3
```

The following is sample output from **show isis fast-reroute detail** command that displays tie-breaker information about the backup:

```
RP/0/RP0/CPU0:router# show isis fast-reroute 10.1.6.0/24 detail

L1 10.1.6.0/24 [20/115] low priority
     via 10.3.7.47, POS0/3/0/1, router2
       FRR backup via 10.1.7.145, HundredGigE 0/1/0/3, router3
       P: No, TM: 30, LC: Yes, NP: No, D: No
     src router2.00-00, 192.168.0.47
   L2 adv [20] native, propagated
```

The following is sample output from **show isis fast-reroute summary** command that displays the number of prefixes having protection per priority:

```
RP/0/RP0/CPU0:router#show isis fast-reroute summary
IS-IS frr IPv4 Unicast FRR summary

                                    Critical   High       Medium     Low        Total
                                    Priority   Priority   Priority      Priority

Prefixes reachable in L1
  All paths protected     0              0             2                8               10

  Some paths protected 0               0             1                3               4

  Unprotected                0              0             1                3
 4
  Protection coverage    0.00%     0.00%     75.00%     78.57%     77.78%
Prefixes reachable in L2
  All paths protected     0              0             0                0               0

  Some paths protected 0               0             1                0               1

  Unprotected                0              0             0                0
 0
  Protection coverage    0.00%     0.00%     100.00%   0.00%      100.00%
```

The following is sample output from **show isis fast-reroute sr-only** command that displays fast-reroute repair paths for prefixes associated with a segment routing prefix SID:

```
RP/0/RP0/CPU0:router#show isis fast-reroute sr-only
IS-IS 1 IPv4 Unicast FRR backups

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       E - EIGRP, A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)
       D - Downstream, LC - Line card disjoint, NP - Node protecting
       P - Primary path, SRLG - SRLG disjoint, TM - Total metric via backup

Maximum parallel path count: 8

L2 20.1.0.101/32 [10/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 10.4.1.103, GigabitEthernet0/0/0/1 r103, SRGB Base:
 16000, Weight: 0
           P node: r103.00 [20.1.0.103], Label: ImpNull
           Q node: r102.00 [20.1.0.102], Label: 24001
           Prefix label: 16101
           Backup-src: r101.00
L2 20.1.0.102/32 [30/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 10.4.1.103, GigabitEthernet0/0/0/1 r103, SRGB Base:
 16000, Weight: 0
           P node: r103.00 [20.1.0.103], Label: ImpNull
           Q node: r102.00 [20.1.0.102], Label: 24001
           Prefix label: ImpNull
           Backup-src: r102.00
L2 20.1.0.103/32 [20/115]
     via 10.4.1.103, GigabitEthernet0/0/0/1, r103, SRGB Base: 16000, Weight: 0
        Backup path: TI-LFA (link), via 10.1.1.101, GigabitEthernet0/0/0/2 r101, SRGB Base:
 16000, Weight: 0
           P node: r102.00 [20.1.0.102], Label: 16102
           Q node: r103.00 [20.1.0.103], Label: 24001
```

```
Prefix label: ImpNull
Backup-src: r103.00
```

show isis hostname

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) router name-to-system ID mapping table, use the **show isis hostname** command in XR EXEC mode.

show isis [**instance** *instance-id*] **hostname**

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS router name-to-system ID mapping table for the specified IS-IS instance only. The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS router name-to-system ID mapping table for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **show isis hostname** command does not display entries if the dynamic hostnames are disabled.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis hostname** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp hostname

ISIS isp hostnames
    Level  System ID      Dynamic Hostname
     1     0001.0000.0005 router
     2   * 0001.0000.0011 router-11
```

This table describes the significant fields shown in the display.

Table 13: show isis instance isp hostname Field Descriptions

Field	Description
Level	IS-IS level of the router.

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Dynamic Hostname	Hostname of the router.
*	Local router.

show isis interface

To display information about the Intermediate System-to-Intermediate System (IS-IS) interfaces, use the **show isis interface** command in XR EXEC mode.

show isis interface [*type interface-path-id* | **level** {**1** | **2**}] [**brief**]

Syntax Description

type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
level { **1** | **2** }	(Optional) Displays IS-IS interface information for Level 1 or Level 2 independently.
brief	(Optional) Displays brief interface output.

Command Default

Displays all IS-IS interfaces.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.
Release 25.4.1	The output was modified to display the field *Bandwidth* and its value.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis interface** command:

```
 Router# show isis interface
0/3/0/2
  tenGigE 0/3/0/2                 Enabled
   Adjacency Formation:       Enabled
   Prefix Advertisement:      Enabled
   Bandwidth:                 400000000

   BFD:                       Disabled
   BFD Min Interval:          150
```

```
BFD Multiplier:           3

Circuit Type:             level-2-only
Media Type:               P2P
Circuit Number:           0
Extended Circuit Number:  67111168
Next P2P IIH in:          4 s
LSP Rexmit Queue Size:    0

Level-2
  Adjacency Count:        1
  LSP Pacing Interval:    33 ms
  PSNP Entry Queue Size:  0

CLNS I/O
  Protocol State:         Up
  MTU:                    4469

IPv4 Unicast Topology:    Enabled
  Adjacency Formation:    Running
  Prefix Advertisement:   Running
  Metric (L1/L2):         80/80
  MPLS LDP Sync (L1/L2):  Disabled/Disabled
IPv6 Unicast Topology:    Disabled (Not cfg on the intf)

IPv4 Address Family:      Enabled
  Protocol State:         Up
  Forwarding Address(es): 10.3.10.143
  Global Prefix(es):      10.3.10.0/24
IPv6 Address Family:      Disabled (No topology enabled which uses IPv6)

LSP transmit timer expires in 0 ms
LSP transmission is idle
Can send up to 9 back-to-back LSPs in the next 0 ms
```

This table describes the significant fields shown in the display.

Table 14: show isis interface Field Descriptions

Field	Description
Bandwidth: 400000000	Bandwidth of the interface, in kilobits per second (kbps). The system automatically detects this value, or you can manually configure it, in kilobits per second (kbps).
HundredGigE 0/6/0/0	Status of the interface, either enabled or disabled.
Adjacency formation:	Status of adjacency formation, either enabled or disabled.
Prefix Advertisement:	Status of advertising connected prefixes, either enabled or disabled.
BFD:	Status of Bidirectional Forwarding Detection (BFD), either enabled or disabled.
BFD Min Interval:	BFD minimum interval.
BFD Multiplier:	BFD multiplier.
Circuit Type:	Levels the interface is running on (circuit-type configuration) which may be a subset of levels on the router.

Field	Description
Media Type:	Media type on which IS-IS is running.
Circuit Number:	Unique ID assigned to a circuit internally (8-bit integer).
Extended Circuit Number:	Valid only for point-to-point interfaces (32-bit integer).
LSP Rexmit Queue Size:	Number of LSPs pending retransmission on the interface.
Adjacency Count:	Number of adjacencies formed with a neighboring router that supports the same set of protocols.
PSNP Entry Queue Size:	Number of SNP entries pending inclusion in the next PSNP.
LAN ID:	ID of the LAN.
Priority (Local/DIS):	Priority of this interface or priority of the Designated Intermediate System.
Next LAN IIH in:	Time (in seconds) in which the next LAN hello message is sent.
LSP Pacing Interval:	Interval at which the link-state packet (LSP) transmission rate (and by implication the reception rate of other systems) is to be reduced.
Protocol State:	Running state of the protocol (up or down).
MTU:	Link maximum transmission unit (MTU).
SNPA:	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
All Level-n ISs:	Status of interface membership in Layer 2 multicast group. The status options are Yes or reason for not being a member of the multicast group.
IPv4 Unicast Topology:	Status of the topology, either enabled or disabled.
Adjacency Formation:	Status of adjacency formation. The status options are Running or a reason for not being ready to form adjacencies.
Prefix Advertisement:	Status of advertising prefixes, either enabled or disabled.
Metric (L1/L2):	IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a connectionless network service (CLNS) prefix).
MPLS LDP Sync (L1/L2)	Status of LDP IS-IS synchronization, either enabled or disabled. When enabled, the state of synchronization (Sync Status) is additionally displayed as either achieved or not achieved.
IPv4 Address Family:	Status of the address family, either enabled or disabled.
Protocol State:	State of the protocol.
Forwarding Address(es):	Addresses on this interface used by the neighbor for next-hop forwarding.

Field	Description
Global Prefix(es):	Prefixes for this interface included in the LSP.
LSP transmit timer expires in	LSP transmission expiration timer interval (in milliseconds).
LSP transmission is	State of LSP transmission. Valid states are: • idle • in progress • requested • requested and in progress

The following is sample output from the **show isis interface** command with the **brief** keyword:

```
RP/0/0/CPU0:router# show isis interface brief

    Interface      All     Adjs    Adj Topos  Adv Topos  CLNS   MTU    Prio
                   OK    L1   L2    Run/Cfg    Run/Cfg                L1   L2
-----------------  ---  ---------  ---------  ---------  ----  ----  --------
PO0/5/0/0          Yes    1    1      1/1        1/1      Up    4469    -    -
Gi0/6/0/0          Yes    1*   1*     1/1        1/1      Up    1497   64   64
```

This table describes the significant fields shown in the display.

Table 15: show isis interface brief Field Descriptions

Field	Description
Interface	Name of the interface.
All OK	Everything is working as expected for this interface.
Adjs L1 L2	Number of L1 and L2 adjacencies over this interface.
Adj Topos Run/Cfg	Number of topologies that participate in forming adjacencies. Number of topologies that were configured to participate in forming adjacencies.
Adv Topos Run/Cfg	Number of topologies that participate in advertising prefixes. Number of topologies that were configured to participate in advertising prefixes.
CLNS	Status of the Connectionless Network Service. Status options are Up or Down.
MTU	Maximum transfer unit size for the interface.
Prio L1 L2	Interface L1 priority. Interface L2 priority.

show isis lsp-log

To display link-state packet (LSP) log information, use the **show isis lsp-log** command in XR EXEC mode.

show isis [**instance** *instance-id*] **lsp-log** [**level** {**1** | **2**}] [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the LSP log information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
level { **1** | **2** }	(Optional) Displays the Intermediate System-to-Intermediate System (IS-IS) link-state database for Level 1 or Level 2 independently.
last *number*	(Optional) Specifies that the output be restricted to the last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output be restricted to the first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays the LSP log information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis lsp-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp lsp-log

ISIS isp Level 1 LSP log
  When       Count      Interface       Triggers
00:02:36         1
00:02:31         1                      LSPREGEN
00:02:26         1      PO4/1           DELADJ
00:02:24         1      PO4/1           NEWADJ
00:02:23         1      Gi5/0           DIS
00:01:27         1      Lo0             IPDOWN
```

```
00:01:12         1      Lo0             IPUP


ISIS isp Level 2 LSP log
  When       Count      Interface       Triggers
00:02:36         1
00:02:30         1                      LSPREGEN
00:02:26         1      PO4/1           DELADJ
00:02:24         1      PO4/1           NEWADJ
00:02:23         1      Gi5/0           DIS
00:02:21         1                      AREASET
00:01:27         1      Lo0             IPDOWN
00:01:12         1      Lo0             IPUP
```

This table describes the significant fields shown in the display.

Table 16: show isis instance isp lsp-log Field Descriptions

Field	Description
Level	IS-IS level of the router.
When	How long ago (in hh:mm:ss) an LSP rebuild occurred. The last 20 occurrences are logged.
Count	Number of events that triggered this LSP run. When there is a topology change, often multiple LSPs are received in a short period. A router waits 5 seconds before running a full LSP, so it can include all new information. This count denotes the number of events (such as receiving new LSPs) that occurred while the router was waiting its 5 seconds before running full LSP.
Interface	Interface that corresponds to the triggered reasons for the LSP rebuild.

Field	Description
Triggers	A list of all reasons that triggered an LSP rebuild. The triggers are: • AREASET—area set changed • ATTACHFLAG—bit attached • CLEAR— **clear** command • CONFIG—configuration change • DELADJ—adjacency deleted • DIS—DIS changed • IFDOWN—interface down • IPADDRCHG—IP address change • IPDEFORIG—IP def-orig • IPDOWN—connected IP down • IFDOWN—interface down • IPEXT—external IP • IPIA—nterarea IP • IPUP—connected IP up • LSPDBOL—LSPDBOL bit • LSPREGEN—LSP regeneration • NEWADJ— new adjacency

show isis mesh-group

To display Intermediate System-to-Intermediate System (IS-IS) mesh group information, use the **show isis mesh-group** command in XR EXEC mode.

show isis [**instance** *instance-id*] **mesh-group**

Syntax Description

instance *instance-id*	(Optional) Displays the mesh group information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS mesh group information for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mesh-group** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp mesh-group

 ISIS isp Mesh Groups

 Mesh group 6:
 HundredGigE 0/4/0/1
```

This table describes the significant fields shown in the display.

Table 17: show isis instance isp mesh-group Field Descriptions

Field	Description
Mesh group	Mesh group number to which this interface is a member. A mesh group optimizes link-state packet (LSP) flooding in nonbroadcast multiaccess (NBMA) networks with highly meshed, point-to-point topologies. LSPs that are first received on interfaces that are part of a mesh group are flooded to all interfaces except those in the same mesh group.
GigabitEthernet0/4/0/1	Interface belonging to mesh group 6.

show isis mpls traffic-eng adjacency-log

To display a log of Multiprotocol Label Switching traffic engineering (MPLS TE) adjacency changes for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng adjacency-log** command in XR EXEC mode.

show isis [**instance** *instance-id*] **mpls traffic-eng adjacency-log** [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the MPLS TE adjacency changes for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
last *number*	(Optional) Specifies that the output is restricted to last *number* of entries. Range is 1 to 20.
first *number*	(Optional) Specifies that the output is restricted to first *number* of entries. Range is 1 to 20.

Command Default

No instance ID specified displays MPLS TE adjacency changes for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show isis mpls traffic-eng adjacency-log** command to display the status of MPLS TE adjacencies.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mpls traffic-eng adjacency-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp mpls traffic-eng adjacency-log

IS-IS isp Level-2 MPLS Traffic Engineering adjacency log
When      Neighbor ID        IP Address       Interface Status
00:03:36  router-6           172.17.1.6         PO0/3/0/1 Up
00:03:36  router-6           172.17.1.6         PO0/3/0/1 Down
00:02:38  router-6           172.17.1.6         PO0/3/0/1 Up
```

This table describes the significant fields shown in the display.

Table 18: show isis instance isp mpls traffic-eng adjacency-log Field Descriptions

Field	Description
When	Time (in hh:mm:ss) since the entry was recorded in the log.
Neighbor ID	Identification value of the neighbor.
IP Address	Neighbor IP Version 4 (IPv4) address.
Interface	Interface from which a neighbor is learned.
Status	Up (active) or Down (disconnected).

Related Commands

Command	Description
show isis mpls traffic-eng advertisements, on page 123	Displays the last flooded record from MPLS traffic engineering.

show isis mpls traffic-eng advertisements

To display the latest flooded record from Multiprotocol Label Switching traffic engineering (MPLS TE) for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng advertisements** command in XR EXEC mode.

show isis [**instance** *instance-id*] **mpls traffic-eng advertisements**

Syntax Description

instance *instance-id*	(Optional) Displays the latest flooded record from MPLS TE for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the latest flooded record from MPLS TE for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show isis mpls traffic-eng advertisements** command to verify that MPLS TE is flooding its record and that the bandwidths are correct.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mpls traffic-eng advertisements** command with the **instance** and *instance-id* values specified:

```
RP/0/RP0/CPU0:router# show isis instance isp mpls traffic-eng advertisements

 ISIS isp Level-2 MPLS Traffic Engineering advertisements
       System ID: router-9
       Router ID: 172.18.0.9
       Link Count: 1
         Link[0]
           Neighbor System ID: router-gsr6 (P2P link)
           Interface IP address: 172.18.0.9
           Neighbor IP Address: 172.18.0.6
           Admin. Weight: 0
           Physical BW: 155520000 bits/sec
           Reservable BW global: 10000000 bits/sec
           Reservable BW sub: 0 bits/sec
           Global pool BW unreserved:
             [0]: 10000000 bits/sec, [1]: 10000000 bits/sec
             [2]: 10000000 bits/sec, [3]: 10000000 bits/sec
```

```
    [4]: 10000000 bits/sec, [5]: 10000000 bits/sec
    [6]: 10000000 bits/sec, [7]: 10000000 bits/sec
  Sub pool BW unreserved:
    [0]: 0 bits/sec, [1]: 0 bits/sec
    [2]: 0 bits/sec, [3]: 0 bits/sec
    [4]: 0 bits/sec, [5]: 0 bits/sec
    [6]: 0 bits/sec, [7]: 0 bits/sec
  Affinity Bits: 0x00000000
```

This table describes the significant fields shown in the display.

Table 19: show isis instance isp mpls traffic-eng advertisements Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or if the **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Router ID	MPLS TE router ID.
Link Count	Number of links that MPLS TE advertised.
Neighbor System ID	System ID of a neighbor number in an area. The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2). In an IS-IS routing domain, each router is represented by a 6-byte hexadecimal system ID. When network administrators maintain and troubleshoot networking devices, they must know the router name and corresponding system ID.
Interface IP address	IP address of the interface.
Neighbor IP Address	IP address of the neighbor.
Admin. Weight	Administrative weight associated with this link.
Physical BW	Link bandwidth capacity (in bits per second).
Reservable BW	Reservable bandwidth on this link.
Global pool BW unreserved	Unreserved bandwidth that is available in the global pool.
Sub pool BW unreserved	Amount of unreserved bandwidth that is available in the subpool.
Affinity Bits	Link attribute flags being flooded. Bits are MPLS-TE specific.

Related Commands

Command	Description
show isis mpls traffic-eng adjacency-log, on page 121	Displays a log of MPLS TE adjacency changes for IS-IS.

show isis mpls traffic-eng tunnel

To display Multiprotocol Label Switching traffic engineering (MPLS TE) tunnel information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng tunnel** command in XR EXEC mode.

show isis [**instance** *instance-id*] **mpls traffic-eng tunnel**

Syntax Description

instance *instance-id*	(Optional) Displays the MPLS TE tunnel information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the MPLS TE tunnel information for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show isis** command to find the current status of MPLS TE tunnels.

Tunnels are used in IS-IS next-hop calculations.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis mpls traffic-eng tunnel** command:

```
RP/0/RP0/CPU0:router# show isis mpls traffic-eng tunnel

ISIS isp Level-2 MPLS Traffic Engineering tunnels
    System Id           Tunnel Name   Bandwidth    Nexthop       Metric   Mode
    router-6            tu0           100000       172.18.1.6      0        Relative
```

This table describes the significant fields shown in the display.

Table 20: show isis mpls traffic-eng tunnel Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.

Field	Description
Tunnel Name	Name of the MPLS TE tunnel interface.
Bandwidth	MPLS TE-specified tunnel bandwidth of the tunnel.
Nexthop	MPLS TE destination IP address of the tunnel.
Metric	MPLS TE metric of the tunnel.
Mode	MPLS TE metric mode of the tunnel. It can be relative or absolute.

show isis neighbors

To display information about Intermediate System-to-Intermediate System (IS-IS) neighbors, use the **show isis neighbors** command in XR EXEC mode.

show isis [**instance** *instance-id*] **neighbors** [*type interface-path-id* | **summary**] [**detail**] [**systemid** *system-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS neighbor information for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.
summary	(Optional) Displays neighbor status count for each level.
detail	(Optional) Displays additional details.
systemid *system-id*	(Optional) Displays the information for the specified neighbor only.

Command Default

No instance ID specified displays neighbor information for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis neighbors** command with the **instance** and *instance-id* values specified:

```
Total neighbor count: 3
RP/0/RSP0RP0/CPU0:router# show isis instance isp neighbors detail

IS-IS isp neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
e222e          Gi0/1/0/0        *PtoP*         Up    23       L1   Capable
  Area Address(es): 00
  IPv4 Address(es): 10.1.0.45*
  IPv6 Address(es): fe80::212:daff:fe6b:68a8*
    Topologies: 'IPv4 Unicast' 'IPv6 Unicast'
  Uptime: 01:09:44
  IPFRR: LFA Neighbor: elise
         LFA IPv4 address: 10.100.1.2
         LFA Router address: 192.168.0.45
e333e LFA Interface: Gi0/1/0/0.1      0012.da6b.68a8 Up    8        L1   Capable 1
e333e          Gi0/1/0/0.1      0012.da6b.68a8 Up    8        L1   Capable
  Area Address(es): 00
  IPv4 Address(es): 10.100.1.2*
  Topologies: 'IPv4 Unicast'
  Uptime: 01:09:46
  IPFRR: LFA Neighbor: elise
         LFA IPv4 address: 10.1.0.45
         LFA Router address: 192.168.0.45
         LFA Interface: Gi0/1/0/0
m44i           Gi0/1/0/1        0012.da62.e0a8 Up    7        L1   Capable
  Area Address(es): 00 11
  IPv4 Address(es): 10.1.2.47*
  IPv6 Address(es): fe80::212:daff:fe62:e0a8*
    Topologies: 'IPv4 Unicast' 'IPv6 Unicast'
  Uptime: 01:09:33

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 21: show isis instance isp neighbors Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Interface	Interface through which the neighbor is reachable.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.
State	Adjacency state of the neighboring interface. Valid states are: Down, Init, and Up.
Holdtime	Hold time of the neighbor.
Type	Type of adjacency.

Field	Description
IETF-NSF	Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism. Valid states are Capable and Unable.
Area Address(es)	Number of area addresses on this router.
IPv4 Address(es)	IPv4 addresses configured on this router.
Topologies	Address and subaddress families for which IS-IS is configured.
Uptime	Time (in hh:mm:ss) that the neighbor has been up.
IPFRR: LFA Neighbor	IP fast reroute (IPFRR) loop-free alternate (LFA) neighbor.
LFA IPv4 address:	Address of the LFA.
LFA Interface:	LFA interface.

The following is sample output from the **show isis neighbors** command with the **summary** keyword specified:

```
RP/0/RP0/CPU0:router# show isis instance isp neighbors summary

ISIS isp neighbor summary:
    State          L1       L2     L1L2
    Up              0        0        2
    Init            0        0        0
    Failed          0        0        0
```

This table describes the significant fields shown in the display.

Table 22: show isis neighbors summary Field Descriptions

Field	Description
State	State of the neighbor is up, initialized, or failed.
L1	Number of Level 1 neighbors.
L2	Number of Level 2 neighbors.
L1L2	Number of Level 1 and 2 neighbors.

show isis protocol

To display summary information about an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis protocol** command in XR EXEC mode.

show isis [**instance** *instance-id*] **protocol**

Syntax Description

instance	*instance-id*	(Optional) Specifies the IS-IS instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 26.3.1	The command output was modified to contain the details from the **maximum-propagated-prefixes** configuration.
Release 24.2.1	This command was modified. The fields to indicate the memory state of the IS-IS protocol were included.
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

This example shows the **show isis protocol** command output.

```
RP/0/RP0/CPU0: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:      1
        Connected:               1
        Propagated:              0
          Warning when over:     20000
        Redistributed:           0
          Limit:                 10000
          Warning when over:     7000 (70%)
    Level-2
      Metric style:              Wide/Wide (generate/accept)
      Default Interface Metric:  10
      Advertising prefixes:      91
        Connected:               1
        Propagated:              90
          Warning when over:     20000
        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
  Level-1:                       configured
    Current status:              clear
    Clear since:                 Tue 2026 Sep 08 14:47:26.791 CEST
    Configuration:               on startup for 300 seconds
      Interarea prefixes:        suppressed
      External prefixes:         suppressed
  Level-2:                       configured
    Current status:              clear
    Clear since:                 Tue 2026 Sep 08 14:47:26.789 CEST
    Configuration:               on startup for 300 seconds
```

```
      Interarea prefixes:        suppressed
      External prefixes:         suppressed
  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:  22097440 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       4459
    Links:                       28       34048
    Prefixes:                    385      373440
    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)
```

This example shows the output containing propagated prefix limits configured using the **maximum-propagated-prefixes** command.

```
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)
```

show isis purge-log

To display the IS-IS LSP purge log to help troubleshoot purge events, use the **show isis purge-log** command in XR EXEC mode.

```
show isis purge-log [ last number | first number ][ level { 1 | 2 } ]
```

Table 23: Syntax Description

show isis purge-log	Executes the command without any additional parameters, displaying the default IS-IS purge log output.
first *number*	(Optional) Displays the first specified *number* of entries from the IS-IS purge log. The *number* can be from 1 to 50,000.
last *number*	(Optional) Displays the last specified *number* of entries from the IS-IS purge log. The *number* can be from 1 to 50,000.
level { **1** | **2** }	(Optional) Filters the display to show purge log entries for a specific IS-IS level, for example, level 1 or level 2.

Command Default

Displays the complete IS-IS purge log output without any filters.

Command Modes

XR EXEC mode

Table 24: Command History

Release	Modification
Release 25.4.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID	Operations
isis	read

This is a sample output of the **show isis purge-log** command.

```
Router#show isis purge-log

  IS-IS 1  Level 2  LSP Purge Log                                  Capacity: 500, Size:
4

When          LSP ID                 Received from    On             POI originator  POI
neighbor
--- Mon Jun 16 2025 ---
12:56:04.295  r104.01-00             r100             Gi0/0/0/1      r100            r103
12:56:05.297  r103.01-00             r100             Gi0/0/0/1      r100            r103
12:56:05.689  r104.01-00             r102             Gi0/0/0/2      r102            r103
12:56:07.297  r103.01-00             r102             Gi0/0/0/2      r102            r103
```

This table describes the significant fields shown in the display.

Table 25: Field Descriptions

Field	Description
When	Time when the LSP purge was logged.
LSP ID	Identifier of the LSP including originating router and LSP number.
Received from	Router ID of the neighbor that sent the purged LSP.
On	Interface on which the LSP was received.
POI originator	Router ID of the original LSP originator.
POI neighbor	Router ID of the neighbor related to the purge event.

show isis route

To display IP reachability information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis route** command in XR EXEC mode.

show isis [**instance** *instance-id*] [**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** {**all***topo-name*}] | **safi-all**] **route** [*ip-address mask* | *ip-address*/*length* [**longer-prefixes**]] [**summary**] **[multicast-intact]** [**backup**] [**detail**] [**sr-only**]

Syntax Description

instance *instance-id*	(Optional) Displays the IP reachability information for the specified IS-IS instance only. The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology	(Optional) Specifies IS-IS paths to intermediate systems.
all	(Optional) Specifies all topologies.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
ip-address	(Optional) Network IP address about which routing information should be displayed.
mask	(Optional) Network mask specified in either of two ways: • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address.
/ *length*	(Optional) Length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. Range is 0 to 32 for IPv4 prefix and 0 to 128 for IPv6 prefix.
longer-prefixes	(Optional) Displays route and more-specific routes.
summary	(Optional) Displays topology summary information.

multicast-intact	(Optional) Displays multicast intact information for this entry.
systemid	(Optional) Displays multicast information by system ID.
backup	(Optional) Displays backup information for this entry.
detail	(Optional) Displays link-state packet (LSP) details.
sr-only	(Optional) Displays SR-labeled prefixes only.

Command Default

No instance ID specified displays the IP reachability information for all the IS-IS instances.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 26.3.1	The command output was modified for a route that was not propagated or advertised for over-limit.
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis route** command:

```
RP/0/RSP0RP0/CPU0:router# show isis route

IS-IS isp IPv4 Unicast routes
Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
df - level 1 default (closest attached router), su - summary null
C - connected, S - static, R - RIP, B - BGP, O - OSPF
i - IS-IS (redistributed from another instance)

Maximum parallel path count: 8

L2 10.76.240.6/32 [4/115]
via 10.76.245.252, SRP0/1/0/2, isp2
via 10.76.246.252, SRP0/1/0/0, isp2
C 10.76.240.7/32
is directly connected, Loopback0
L2 10.76.240.9/32 [256/115]
via 10.76.249.2, HundredGigE 0/3/0/0, isp3
L2 10.76.240.10/32 [296/115]
via 10.76.249.2, HundredGigE 0/3/0/0, isp3
C 10.76.245.0/24
is directly connected, SRP0/1/0/2
C 10.76.246.0/24
is directly connected, SRP0/1/0/0
```

```
  C 10.76.249.0/26
  is directly connected, HundredGigE 0/3/0/0
  L2 10.101.10.0/24 [296/115]
  via 10.76.249.2, HundredGigE 0/3/0/0, isp3
```

This table describes the significant fields shown in the display.

Table 26: show isis route ipv4 unicast Field Descriptions

Field	Description
C172.18.0.0/24	Connected route for HundredGigE interface 0/5/0/0.
C 172.19.1.0/24	Connected route for HundredGigE interface 0/4/0/1.
L1 172.35.0.0/24 [10]	Level 1 route to network 172.35.0.0/24.
C 172.18.0/24	Connected route for loopback interface 0.

This is sample output from the **show isis route** command with **detail** keyword that shows prefix segment ID (SID) and Segment Routing Global Block (SRGB) values:

```
Sun May  4 13:05:11.073 PDT

L2 172.16.255.2/32 [10/115] medium priority
     via 172.16.2.2, HundredGigE 0/0/0/1, pe2 tag 255, SRGB Base: 16000, Weight: 0
     src pe2.00-00, 172.16.255.2, tag 255, prefix-SID index 42, R:0 N:0 P:0
   L1 adv [10] native, propagated, interarea, tag 255, prefix-SID index 42, R:0
          N:0 P:0
```

This is sample output from the **show isis route** command with **sr-only** keyword that shows only routes associated with a segment routing prefix SID:

```
RP/0/RP0/CPU0:router# show isis route sr-only

IS-IS 1 IPv4 Unicast routes

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)

Maximum parallel path count: 8

C  20.1.0.100/32
     is directly connected, Loopback0
L2 20.1.0.101/32 [10/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
L2 20.1.0.102/32 [30/115]
     via 10.1.1.101, GigabitEthernet0/0/0/2, r101, SRGB Base: 16000, Weight: 0
L2 20.1.0.103/32 [20/115]
     via 10.4.1.103, GigabitEthernet0/0/0/1, r103, SRGB Base: 16000, Weight: 0
```

This example shows the output for a route that is not propagated or advertised due to an exceeded limit.

```
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#
```

show isis spf-log

To display how often and why the router has run a full shortest path first (SPF) calculation, use the **show isis spf-log** command in XR EXEC mode.

show isis [**instance** *instance-id*] [[**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** {**all***topo-name*}] | **safi-all**]] **spf-log** [**level** {**1** | **2**}] [**ispf** | **fspf** | **prc** | **nhc**] [**detail** | **verbose** | **plfrr** | **ppfrr**] [**last** *number* | **first** *number*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS SPF log for the specified IS-IS instance only.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology all | *topo-name*	(Optional) Specifies topology table information for all topologies or for the specified topology table (*top-name*).
safi-all	(Optional) Specifies all secondary address prefixes.
level { **1** | **2** }	(Optional) Displays the IS-IS SPF log for Level 1 or Level 2 independently.
ispf	(Optional) Specifies incremental SPF entries only.
fspf	(Optional) Specifies full SPF entries only.
prc	(Optional) Specifies partial route calculations only.
nhc	(Optional) Specifies next-hop route calculations only.
plfrr	(Optional) Specifies per link fast-reroute calculations only.
ppfrr	(Optional) Specifies per prefix fast-reroute calculations only.
detail	(Optional) Specifies detailed output. Includes a breakdown of the time taken to perform the calculation and changes resulting from the calculation.
verbose	(Optional) Specifies verbose output.
last *number*	(Optional) Specifies that the output is restricted to the last *number* of entries. Range is 1 to 210.
first *number*	(Optional) Specifies that the output is restricted to the first *number* of entries. Range is 1 to 210.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.

Displays all types of route calculation (not just fspf, ispf and prc).

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis spf-log** command:

```
RP/0/RP0/CPU0:router# show isis spf-log

      IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
                  Time  Total Trig
 Timestamp   Type (ms)  Nodes Count First Trigger LSP Triggers
 ----------- ---- ----  ----- ----- ----- ------- --- --------
 --- Thurs Aug 19 2004 ---
 12:00:50.787  FSPF  1    1     3   ensoft-grs7.00-00 LSPHEADER TLVCODE
 12:00:52.846  FSPF  1    1     1   ensoft-grs7.00-00 LSPHEADER
 12:00:56.049  FSPF  1    1     1   ensoft-grs7.00-00 TLVCODE
 12:01:02.620  FSPF  1    1     2   ensoft-grs7.00-00 NEWADJ LINKTLV

      IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
                  Time  Total Trig
 Timestamp   Type (ms)  Nodes Count First Trigger LSP Triggers
 ----------- ---- ----  ----- ----- ----- ------- --- --------
 --- Mon Aug 19 2004 ---
 12:00:50.790  FSPF  0    1     4   ensoft-grs7.00-00 LSPHEADER TLVCODE
 12:00:54.043  FSPF  1    1     2   ensoft-grs7.00-00 NEWADJ LSPHEADER
 12:00:55.922  FSPF  1    2     1   ensoft-grs7.00-00 NEWLSPO
 12:00:56.724  FSPF  1    13    1   ensoft-grs7.00-00 NEWLSPO
```

This table describes the significant fields shown in the display.

Table 27: show isis spf-log ipv4 unicast Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time when the SPF calculation started.
Duration	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.

Field	Description
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

This table lists triggers of a full SPF calculation.

Table 28: List of Triggers

Trigger	Description
PERIODIC	Runs a full SPF calculation very 15 minutes.
NEWLEVEL	Configured new level (using is-type) on this router.
RTCLEARED	Cleared IS-IS topology on the router.
MAXPATHCHANGE	Changed IP maximum parallel path.
NEWMETRIC	Changed link metric.
ATTACHFLAG	Changed Level 2 Attach bit.
ADMINDIST	Configured another administrative distance for the IS-IS instance on this router.
NEWADJ	Created a new adjacency to another router.
DELADJ	Deleted adjacency.
BACKUP	Installed backup route.
SEEDISPF	Seed incremental SPF.
NEXTHOP	Changed IP next-hop address.
NEWLSP0	New LSP 0 appeared in the topology.
LSPEXPIRED	Some LSP in the link-state database (LSDB) has expired.
LSPHEADER	Changed important LSP header fields.

Trigger	Description
TLVCODE	Type, length, and value (TLV) objects code mismatch, indicating that different TLV objects are included in the newest version of an LSP.
LINKTV	Changed Link TLV content.
PREFIXTLV	Changed Prefix TLV content.
AREAADDRTLV	Changed Area address TLV content.
IP ADDRTLV	Changed IP address TLV content.
TUNNEL	Changed RRR tunnel.

The following is sample output from the **show isis spf-log** command with the **first** keyword specified:

```
RP/0/RP0/CPU0:router# show isis spf-log first 2

    IISIS isp Level 1 IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.140  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.646  FSPF  1    1     1                         NEWADJ


   IISIS isp Level 2  IPv4 Unicast Route Calculation Log
                 Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.139  FSPF  1    1     1             12a5.00-00 NEWLSP0
19:25:35.347  FSPF  1    1     2             12a5.00-00 NEWSADJ TLVCODE
```

This table describes the significant fields shown in the display.

Table 29: show isis spf-log first Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.

Field	Description
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.

The following is sample output from the **show isis spf-log** command with the **detail** keyword specified:

```
RP/0/RP0/CPU0:router# show isis spf-log detail

    IISIS isp Level 1 IPv4 Unicast Route Calculation Log
                Time  Total Trig
Timestamp   Type (ms)  Nodes Count First Trigger LSP   Triggers
  Mon Aug 16 2004
19:25:35.140  FSPF  1    1     1             12a5.00-00 NEWLSP0
  Delay:              51ms (since first trigger)
  SPT Calculation
    CPU Time:         0ms
    Real Time:        0ms
  Prefix Updates
    CPU Time:         1ms
    Real Time:        1ms
  New LSP Arrivals:    0
  Next Wait Interval: 200ms
                            Results
                       Reach Unreach Total
   Nodes:                  1       0     1
   Prefixes (Items)
     Critical Priority:    0       0     0
     High Priority:        0       0     0
     Medium Priority       0       0     0
     Low Priority          0       0     0

     All Priorities        0       0     0
   Prefixes (Routes)
     Critical Priority:    0       -     0
     High Priority:        0       -     0
     Medium Priority       0        -    0
     Low Priority:         0        -    0

     All Priorities        0        -    0
```

This table describes the significant fields shown in the display.

Table 30: show isis spf-log detail Field Descriptions

Field	Description
Level	IS-IS level of the router.
Timestamp	Time at which the SPF calculation started.
Type	Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC).
Time (ms)	Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time.
Nodes	Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run.
Trig Count	Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the **spf-interval** command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see *List of Triggers*.
First Trigger LSP	LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered.
Triggers	List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see *List of Triggers*.
Delay	Two different delays exist: **1.** The delay between the time when the route calculation was first triggered and the time when it was run. **2.** The delay between the end of the last route calculation and the start of this one. This is used to verify that the SPF-interval timers are working correctly, and is only reported for calculations after the first delay.
CPU Time	Two different CPU times exist: **1.** CPU time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** CPU time (in milliseconds) taken to perform the prefix updates.
Real Time	Two different real times exist: **1.** Real time (in milliseconds) taken to calculate the shortest path tree (SPT). **2.** Real time (in milliseconds) taken to perform the prefix updates.

Field	Description
New LSP Arrivals	Number of LSP arrivals since the start of this route calculation.
Next Wait Interval	Enforced delay until the next route calculation can be run, based on the **spf-interval** command configuration.
Reach	Number of reachable nodes or prefixes.
Unreach	Number of unreachable nodes or prefixes.
Total	Total number of nodes or prefixes at various priorities.

show isis statistics

To display Intermediate System-to-Intermediate System (IS-IS) traffic counters, use the **show isis statistics** command in XR EXEC mode.

show isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS traffic statistics for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default

No instance ID specified displays IS-IS traffic statistics for all the IS-IS instances.

IS-IS traffic statistics are displayed for all interfaces.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

The **show isis statistics** command displays IS-IS traffic counters for the specified interface or all traffic counters if no interface is specified.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis statistics** command that shows all traffic counters:

```
RP/0/RP0/CPU0:router#show isis statistics
IS-IS isp statistics:
    Fast PSNP cache (hits/tries): 164115/301454
    Fast CSNP cache (hits/tries): 41828/43302
    Fast CSNP cache updates: 2750
    LSP checksum errors received: 0
```

```
      LSP Dropped: 1441
      SNP Dropped: 1958
      UPD Max Queue size: 2431
      Average transmit times and rate:
        Hello:          0 s,     987947 ns,          4/s
        CSNP:           0 s,    1452987 ns,          0/s
        PSNP:           0 s,    1331690 ns,          0/s
        LSP:            0 s,    1530018 ns,          1/s
      Average process times and rate:
        Hello:          0 s,     874584 ns,         41/s
        CSNP:           0 s,     917925 ns,         29/s
        PSNP:           0 s,    1405458 ns,          0/s
        LSP:            0 s,    4352850 ns,          0/s
      Level-1:
        LSPs sourced (new/refresh): 3376/2754
        Level-1::LSPs sourced (new/refresh)SPF calculations          : 3376/2754520   ISPF
 calculations        IPv4 Unicast: 0SPF calculations Next Hop Calculations            :
5200ISPF calculations          Partial Route Calculations : 0
          NextIPFRR R-hop Calculations   SPF calculations    : 0
          Partial Route Calculations IPFRR Parallel calculations: 0
        IPv6 Unicast
          SPF calculations           : 527
          ISPF calculations          : 0
          Next Hop Calculations      : 13
          Partial Route Calculations : 1
          Level-2:
        LSPs sourced (new/refresh): 4255/3332
        IPv4 Unicast
          SPF calculations           : 432
          ISPF calculations          : 0
          Next Hop Calculations      : 8
          LSPs sourced (new/refresh)Partial Route Calculations: 4255/33320LSPs sourced
(new/refresh)IPFRR R-SPF calculations   : 4255/33320
          IPFRR Parallel calculations: 0
        IPv4 IPv6 Unicast
          SPF calculations           : 432444
          ISPF calculations          : 0
          Next-hop Next Hop Calculations      : 882
          Partial Route Calculations : 01      Interface GigabitEthernet0/1/0/1.1:
      Level-1 Hellos (sent/rcvd): 22398/25633
      Level-1 DR Elections      : 66
      Level-1 LSPs (sent/rcvd)  : 246/7077
      Level-1 CSNPs (sent/rcvd) : 0/33269
      Level-1 PSNPs (sent/rcvd) : 22/0
      Level-1 LSP Flooding Duplicates  : 25129
      Level-2 Hellos (sent/rcvd): 22393/67043
      Level-2 DR Elections      : 55
      Level-2 LSPs (sent/rcvd)  : 265/437
      Level-2 CSNPs (sent/rcvd) : 0/86750
      Level-2 PSNPs (sent/rcvd) : 0/0
      Level-2 LSP Flooding Duplicates  : 78690
```

This table describes the significant fields shown in the display.

Table 31: show isis statistics Field Descriptions

Field	Description
Fast PSNP cache (hits/tries)	Number of successful lookups (hits) along with the number of lookup attempts (tries). To save time or processing power when receiving multiple copies of the same LSP, IS-IS attempts to look up incoming LSPs to see if they have been received recently.
Fast CSNP cache (hits/tries):	Number of successful lookups (hits) along with the number of lookup attempts (tries). To reduce CSNP construction time, IS-IS maintains a cache of CSNPs and attempts to look up CSNP in this cache before transmission on the interface.
Fast CSNP cache updates:	Number of times the CSNP cache has been updated since the last clearing of statistics. The cache is updated on LSP addition or removal from the database.
LSP checksum errors received:	Number of internal checksum errors received in LSPs.
IIH (LSP/SNP) dropped:	Number of hello, LSP, and SNP messages dropped.
IIH (UPD) Max Queue size:	Maximum number of queued packets.
Average transmit times and rate:	Average time taken to transmit the pdu type across all interfaces and the corresponding rate at which the pdu type is being transmitted.
Average process times and rate:	Average time taken to process an incoming pdu type across all interfaces and the corresponding rate at which the pdu type is being received.
LSPs sourced (new/refresh):	Number of LSPs this IS-IS instance has created or refreshed. To find more details on these LSPs, use the **show isis lsp-log** command.
SPF calculations:	Number of shortest path first (SPF) calculations. SPF calculations are performed only when the topology changes. They are not performed when external routes change. The interval at which SPF calculations are performed is configured using the **spf-interval** command.
iSPF calculations:	Number of incremental shortest path first (iSPF) calculations. iSPF calculations are performed only when ISPF has been configured in the isis address family configuration submode.
Partial Route Calculations:	Number of partial route calculations (PRCs). PRCs are processor intensive. Therefore, it may be useful to limit their number, especially how often a PRC is done, especially on slower networking devices. Increasing the PRC interval reduces the processor load on the router, but might slow the rate of convergence. The interval at which PRC calculations are performed is configured using the **spf-interval** command.
Level-(1/2) (LSPs/CSNPs/PSNPs/Hellos) (sent/rcvd):	Number of LSPs, Complete Sequence Number Packets (CSNPs), Partial Sequence Number Packets (PSNPs), and hello packets sent or received on this interface.

Field	Description
PTP Hellos (sent/rcvd):	Point-to-point (PTP) hellos sent and received.
LSP Retransmissions:	Total number of retransmissions on each IS-IS LSP on a point-to-point interface. The LSP retransmission interval can be configured using the **retransmit-throttle-interval** command.
Level-(1.2) DRElections:	Total number of Designated Intermediate System elections that have taken place. These counts are maintained on an individual level basis.
LSP Flooding Duplicates:	Number of duplicate LSPs filtered from flooding to the neighbor. In case of parallel interfaces to the same neighbor, IS-IS optimizes the flooding by avoiding sending the same LSP copy on other interfaces.

show isis topology

To display a list of connected Intermediate System-to-Intermediate System (IS-IS) routers in all areas, use the **show isis topology** command in XR EXEC mode.

show isis [**instance instance-id**] [[**ipv4** | **ipv6** | **afi-all**] [**unicast** | **multicast** [**topology** {**all** | **topo-name**}] | **safi-all**]] | **summary** | **level** {**1** | **2**} **[multicast-intact]** [**systemid system-id**] [**detail**]

Syntax Description

instance *instance-id*	(Optional) Displays the IS-IS topology for the specified IS-IS instance only. • The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.
ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
afi-all	(Optional) Specifies all address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
topology *topo-name*	(Optional) Specifies topology table information and name of the topology table.
safi-all	(Optional) Specifies all secondary address prefixes.
summary	(Optional) Displays a brief list of the IS-IS topology.
level { **1** | **2** }	(Optional) Displays the IS-IS link-state topology for Level 1 or Level 2 independently.
multicast-intact	(Optional) Displays multicast intact information on the IS-IS topology.
systemid *system-id*	(Optional) Displays the information for the specified router only.
detail	(Optional) Displays detailed information on the IS-IS topology.

Command Default

No instance ID specified displays a list of connected routers in all areas for all the IS-IS instances.

Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **show isis topology** command to verify the presence and connectivity among all routers in all areas.

Task ID

Task ID	Operations
isis	read

Examples

The following is sample output from the **show isis topology** command:

```
RP/0/RP0/CPU0:router# show isis topology

IS-IS isp paths to (Level-1) routers
    System Id       Metric  Next-hop Interface       SNPA
    ensoft-5        10      ensoft-5        PO0/4/0/1       *PtoP*
    ensoft-5        10      ensoft-5        Gi0/5/0/0       0003.6cff.0680
    ensoft-11       --

    IS-IS isp paths to (Level-2) routers
    System Id       Metric  Next-hop Interface       SNPA
    ensoft-5        10      ensoft-5        PO0/4/0/1       *PtoP*
    ensoft-5        10      ensoft-5        Gi0/5/0/0       0003.6cff.0680
    ensoft-11       --
```

This table describes the significant fields shown in the display.

Table 32: show isis topology ipv4 unicast Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
Metric	Metric assigned to the link and used to calculate the cost from each router using the links in the network to other destinations. Range is 1 to 16777214. Default is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. 0 is set internally if no metric has been specified by the user.
Next-hop	Address of the next-hop.
Interface	Interface used to reach the neighbor.
SNPA	Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor.

The following is sample output from the **show isis topology** command with the **summary** keyword specified:

```
RP/0/RP0/CPU0:router# show isis topology summary

IS-IS 10 IS Topology Summary IPv4 Unicast
                        L1                              L2
                Reach  UnReach  Total       Reach  UnReach  Total
               ------- ------- -------     ------- ------- -------
Router nodes:        1       1       2           1       1       2
Pseudo nodes:        0       0       0           0       0       0

 Total nodes:        1       1       2           1       1       2
```

This table describes the significant fields shown in the display.

Table 33: show isis topology summary Field Descriptions

Field	Description
L1/L2	IS-IS level of the router.
Reach	Number of router nodes or pseudonodes that are reachable.
UnReach	Number of router nodes or pseudonodes that are unreachable.
Total	Total number of reachable and unreachable nodes.

show protocols (IS-IS)

To group a number of protocol show commands according to the specified address family, use the **show protocols** command in XR EXEC mode.

show protocols [**afi-all** | **ipv4** | **ipv6**] [**all***protocol*]

Syntax Description

afi-all	(Optional) Specifies all address families.
ipv4	(Optional) Specifies an IPv4 address family.
ipv6	(Optional) Specifies an IPv6 address family.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are: • **bgp** • **isis** • **ospf** • **rip** For the IPv6 address family, the options are: • **bgp** • **isis** • **ospfv3**

Command Default

If no address family is specified, the default is IPv4.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

If IPv6 is enabled on an IS-IS instance, the instance is displayed in the **show protocols ipv6** command output. IPv4 IS-IS instances are displayed in the **show protocols ipv4** command output.

When using the **show protocols** command with the **ipv6** or **ipv4** keyword, you get all routing instances in that particular address family—not only IS-IS instances.

Task ID

Task ID	Operations
isis	read

Task ID	Operations
rib	read

Examples

The following example shows the output for the **show protocols** command :

```
RP/0/RP0/CPU0:router# show protocols ipv4

IS-IS Router: uut
  System Id: 0000.0000.12a8
  IS Levels: level-1-2
  Manual area address(es):
    49.1515.1515
  Routing for area address(es):
    49.1515.1515
  Non-stop forwarding: Disabled
  Most recent startup mode: Cold Restart
  Topologies supported by IS-IS:
    IPv4 Unicast
      Level-1
        Metric style (generate/accept): Narrow/Narrow
        ISPF status: Disabled
      Level-2
        Metric style (generate/accept): Narrow/Narrow
        ISPF status: Disabled
      Redistributing:
        static
      Distance: 115
    IPv6 Unicast
      Level-1
        ISPF status: Disabled
      Level-2
        ISPF status: Disabled
      No protocols redistributed
      Distance: 45
  Interfaces supported by IS-IS:
    GigabitEthernet 0/6/0/0 is running actively (active in configuration)
```

This table describes the significant fields shown in the display.

Table 34: show protocols ipv4 Field Descriptions

Field	Description
System ID	Dynamic hostname of the system. The hostname is specified using the **hostname** command. If the dynamic hostname is not known or **hostname dynamic disable** command has been executed, the 6-octet system ID is used.
IS Levels	IS-IS level of the router.
Manual area address(es)	Area addresses configured manually on the originating router.
Routing for area address(es)	Area addresses for which this router provides the routing.

Field	Description
Non-stop forwarding	Status and name of NSF.
Most recent startup mode	Mode in which the most recent startup was performed.
Topologies supported by IS-IS	Address and subaddress family IS-IS are configured.
Metric style	Type, length, and value (TLV) objects accepted by IS-IS. To configure this value, see the metric-style wide, on page 59 command.
ISPF status	State of iSPF configuration for this IS-IS instance. Four states exist: • Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Dormant if iSPF has been configured but is awaiting initial convergence before initializing. • Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run.
Redistributing	IS-IS is configured to redistribute IP static routes into Level 1 or Level 2. The **redistribute** command is used to configure redistribution.
Distance	Administrative distance.
Interfaces supported by IS-IS	Interfaces and their states currently supported by IS-IS. Both operational and configuration status are displayed.

The following example shows how to disable the IPv4 address family, with no output shown for IS-IS IPv4 instances from the **show protocols ipv4** command:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router isis uut
RP/0/RP0/CPU0:router(config-isis)# no address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis)# commit

RP/0/RP0/CPU0:router# show protocols ipv4
```

single-topology

To configure the link topology for IP Version 4 (IPv4) when IP Version 6 (IPv6) is configured, use the **single-topology** command in address family configuration mode. To remove the **single-topology** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

single-topology

Command Default

Performs in multitopology mode in which independent topologies for IPv4 and IPv6 are running in a single area or domain.

Command Modes

IPv6 address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **single-topology** command to allow Intermediate System-to-Intermediate System (IS-IS) for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols, and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is being used, only old-style type, length, and value (TLV) objects may be used and a single shortest path (SPF) individual level is used to compute IPv4 (if configured) and IPv6 routes. The use of a single SPF means that both IPv4 IS-IS and IPv6 IS-IS routing protocols must share a network topology.

To allow link information to be shared between IPv4 and IPv6, you must configure the **single-topology** command for an address family. In single-topology IPv6 mode, the configured metric is always the same for both IPv4 and IPv6.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to enable single-topology mode for IPv6:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RP0/CPU0:router(config-isis)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)# single-topology
```

snmp-server traps isis

To enable the Simple Network Management Protocol (SNMP) server notifications (traps) available for IS-IS, use the **snmp-server traps isis** command in XR Config mode. To disable all available SNMP notifications, use the **no** form of this command.

snmp-server traps isis { **all** | *traps set* }

Syntax Description

all	Specifies all IS-IS SNMP server traps.
traps set	Specify any set of trap names.

Command Default

SNMP server traps notification is disabled.

Command Modes

XR Config mode

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
isis	read, write

Examples

The following examples show how to enable all SNMP server traps available for isis:

```
RP/0/RP0/CPU0:router(config)# snmp-server traps isis?

adjacency-change                isisAdjacencyChange
all                             Enable all IS-IS traps
area-mismatch                   isisAreaMismatch
attempt-to-exceed-max-sequence  isisAttemptToExceedMaxSequence
authentication-failure          isisAuthenticationFailure
authentication-type-failure     isisAuthenticationTypeFailure
corrupted-lsp-detected          isisCorruptedLSPDetected
database-overload               isisDatabaseOverload
id-len-mismatch                 isisIDLenMismatch
lsp-error-detected              isisLSPErrorDetected
lsp-too-large-to-propagate      isisLSPTooLargeToPropagate
manual-address-drops            isisManualAddressDrops
max-area-addresses-mismatch     isisMaxAreaAddressesMismatch
orig-lsp-buff-size-mismatch     isisOrigLSPBuffSizeMismatch
own-lsp-purge                   isisOwnLSPPurge
protocols-supported-mismatch    isisProtocolsSupportedMismatch
rejected-adjacency              isisRejectedAdjacency
```

```
  sequence-number-skip            isisSequenceNumberSkip
  version-skew                    isisVersionSkew

RP/0/RP0/CPU0:router(config)#snmp-server traps isis all
```

The following example shows how to enable area-mismatch lsp-error-detected trap:

```
RP/0/RP0/CPU0:router(config)# snmp-server traps isis area-mismatch
lsp-error-detected
```

spf-interval

To customize IS-IS throttling of shortest path first (SPF) calculations, use the **spf-interval** command in address family configuration mode. To restore default values, use the **no** form of this command.

spf-interval [**initial-wait** *initial* | **secondary-wait** *secondary* | **maximum-wait** *maximum*] . . . [**level** { **1** | **2** }]

Syntax Description

initial-wait *initial*	Initial SPF calculation delay (in milliseconds) after a topology change. Range is 0 to 120000.
secondary-wait *secondary*	Hold time between the first and second SPF calculations (in milliseconds). Range is 0 to 120000.
maximum-wait *maximum*	Maximum interval (in milliseconds) between two consecutive SPF calculations. Range is 0 to 120000.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

initial-wait *initial* : 50 milliseconds

secondary-wait *secondary* : 200 milliseconds

maximum-wait *maximum* : 5000 milliseconds

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Use the **spf-interval** command to control how often the software can perform the SPF calculation. The SPF calculation is processor intensive. Therefore, it may be useful to limit how often this calculation is done, especially when the area is large and the topology changes often. Increasing the SPF interval reduces the processor load of the router, but potentially slows the rate of convergence.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the initial SPF calculation delay to 10 milliseconds and the maximum interval between two consecutive SPF calculations to 5000 milliseconds:

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 5000
```

spf-interval ietf

To set an shortest path first (SPF) interval in IS-IS for SPF calculations, use the **spf-interval ietf** command in the System Admin Config mode. Use the **no** form of this command to enable the fabric bundle port.

spf-interval ietf [**initial-wait** *msec* | **short-wait** *msec* | **long-wait** *msec* | **learn-interval** *msec* | **holddown-interval** *msec*] [**level** { **1** | **2** }]

Syntax Description

spf-interval	Specifies the number of seconds between two consecutive SPF calculations.
ietf	Specifies Internet Engineering Task Force (IETF) RFC standard 8405.
initial-wait *msec*	Initial SPF calculation delay before running a route calculation. The initial-wait must be less than or equal to short-wait. Range is 0 to 120000. The default value is 50 milliseconds.
short-wait *msec*	Short SPF calculation delay before running a route calculation. The short-wait must be less than or equal to long-wait. Range is 0 to 120000. The default value is 200 milliseconds.
long-wait *msec*	Long SPF calculation delay before running a route calculation. Range is 0 to 120000. The default value is 5000 milliseconds.
learn-interval *msec*	Time To Learn interval for running a route calculation. The learn-interval must be less than or equal to holddown-interval. Range is 0 to 120000. The default value is 500 milliseconds.
holddown-interval *msec*	Hold-down interval for running a route calculation. Range is 0 to 120000. The default value is 10000 milliseconds.
level { **1** | **2** }	(Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently.

Command Default

None

Command Modes

System Admin Config mode

Command History

Release	Modification
Release 7.7.1	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Task ID

Task ID	Operations
is-is	read, write

Example

The following example shows how to configure IETF to set an SPF interval in IS-IS for SPF calculations.

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# spf-interval ietf?
initial-wait   Initial delay before running a route calculation [50]
short-wait     Short delay before running a route calculation [200]
long-wait      Long delay before running a route calculation [5000]
learn-interval Time To Learn interval for running a route calculation [500]
holddown-interval  Holddown interval for running a route calculation [10000]
level             Set SPF interval for one level only
Router(config-isis-af)# spf-interval ietf
Router(config-isis-af)# commit
```

The following **show** command displays the output with the new spf-interval algorithm. The output displays the actual delay taken to compute the SPF.

```
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
                         ----- -----
```

spf prefix-priority (IS-IS)

To assign a priority to an ISIS prefix for customizing the RIB update sequence, use the**spf prefix-priority** command in address family configuration mode. To restore default values, use the **no** form of this command.

spf prefix-priority [**level** { **1** | **2** }] { **critical** | **high** | **medium** } { *access-list-name* | **tag** *tag* }

Syntax Description

level { **1** | **2** }	(Optional) Enables the assignment of a priority to Level 1 or Level 2 independently.
critical	Assigns a critical priority.
high	Assigns a high priority.
medium	Assigns a medium priority.
access-list-name	Name of an access list.
tag *tag*	Specifies a tag to indicate priority. The *tag* argument range is 1 to 4294967295.

Command Default

By default, IPv4 prefixes with a length of 32 and IPv6 prefixes with a length of 128 are given medium priority. The remaining prefixes are given low priority.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **spf prefix-priority** command to change the sequence of prefix updates to the RIB after an SPF is run. ISIS installs prefixes in the RIB according to the following priority order:

Critical > High > Medium > Low

The **spf prefix-priority** command supports prefix lists for the first three priorities. The unmatched prefixes are updated with low priority.

If a **spf prefix-priority** is specified, the default behavior of prioritizing either length 32 or 128 prefixes for IPv4 or IPv6, respectively, as **medium** is disabled.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to set the prefix priorities:

```
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-critical-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 32
```

```
!
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-med-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 29
!
RP/0/RP0/CPU0:router(config)# ipv4 prefix-list isis-high-acl
RP/0/RP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 30
!
RP/0/RP0/CPU0:router(config)# router isis ring
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority critical isis-critical-acl
RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority high isis-high-acl
RP/0/RP0/CPU0:router(config-isis-af)# spf prefix-priority medium isis-med-acl
```

summary-prefix (IS-IS)

To create aggregate addresses for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **summary-prefix** command in address family configuration mode. To restore the default behavior, use the **no** form of this command.

Syntax Description

address	Summary address designated for a range of IPv4 addresses. The *address* argument must be in four-part, dotted-decimal notation.
/ *prefix-length*	Length of the IPv4 or IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value.
ipv6-prefix	Summary prefix designated for a range of IPv6 prefixes. The *ipv6-prefix* argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons.
level { **1** | **2** }	(Optional) Redistributes routes into Level 1 or Level 2 and summarizes them with the configured address and mask value.
tag *tag*	Sets a tag value. The value range is 1- 4294967295.

Command Default

All redistributed routes are advertised individually.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more-specific routes. Use the **summary-prefix** command to help reduce the size of the routing table.

This command also reduces the size of the link-state packets (LSPs) and thus the link-state database. It also helps ensure stability, because a summary advertisement depends on many more specific routes. If one more-specific route flaps, in most cases, this flap does not cause a flap of the summary advertisement.

The drawback of summary addresses is that other routes might have less information to calculate the most optimal routing table for all individual destinations.



Note When IS-IS advertises a summary prefix, it automatically inserts the summary prefix into the IP routing table but labels it as a “discard” route entry. Any packet that matches the entry is discarded to prevent routing loops. When IS-IS stops advertising the summary prefix, the routing table entry is removed.

Task ID

Task ID	Operations
isis	read, write

Examples

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS:

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS. In the OSPF routing table, IPv6 routes exist for 3ffe:f000:0001:0000::/64, 3ffe:f000:0002:0000::/64, 3ffe:f000:0003:0000::/64, and so on. This example shows only 3ffe:f000::/24 advertised into IPv6 IS-IS Level 2.

```
RP/0/RP0/CPU0:router(config)# router isis isp
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 ipv6 unicast
RP/0/RP0/CPU0:router(config-isis-af)# redistribute ospf ospfv3 2 level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 3ffe:f000::/24 level-2
RP/0/RP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 3ffe:f000::/24 tag
```