IS-IS Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>router>isis
config>router>isis>interface *ip-int-name*
config>router>isis>if>level *level-number*
config>router>isis>if>segment-routing

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Special Cases **IS-IS Global —** In the **config>router>isis** context, the **shutdown** command disables the IS-IS protocol instance. By default, the protocol is enabled, **no shutdown**.

IS-IS Interface — In the **config>router>isis>interface** context, the command disables the IS-IS interface. By default, the IS-IS interface is enabled, **no shutdown**.

IS-IS Interface and Level — In the **config>router>isis>interface *ip-int-name*>level** context, the command disables the IS-IS interface for the level. By default, the IS-IS interface at the level is enabled, **no shutdown**.

Default **no shutdown** — IS-IS entity is administratively enabled.

IS-IS Commands

isis

Syntax **[no] isis** [*isis-instance*]

Context config>router

Description This command creates the context to configure the Intermediate-System-to-Intermediate-System (IS-IS) protocol instance.

The IS-IS protocol instance is enabled with the **no shutdown** command in the **config>router>isis** context. Alternatively, the IS-IS protocol instance is disabled with the **shutdown** command in the **config>router>isis** context.

The **no** form of the command deletes the IS-IS protocol instance. Deleting the protocol instance removes all configuration parameters for this IS-IS instance.

Parameters *isis-instance* — Specifies the instance ID for an IS-IS instance.

Values 1–31

Default 0

tag

Syntax **tag** *tag*
no tag

Context config>router>isis>interface

Description This command configures a route tag to the specified IP address of an interface.

Parameters *tag* — Assigns a route tag.

Values 1 — 4294967295

all-l1isis

Syntax **all-l1isis** *ieee-address*
no all-l1isis

Context config>router>isis

Description This command enables you to specify the MAC address to use for all L1 IS-IS routers. The MAC address should be a multicast address. You should shut/no shut the IS-IS instance to make the change operational.

Default all-l1isis 01-80-C2-00-01-00

Parameters *ieee-address* — Specifies the destination MAC address for all L1 I-IS neighbors on the link for this IS-IS instance.

all-l2isis

Syntax **all-l2isis** *ieee-address*
no all-l2isis

Context config>router>isis

Description This command enables you to specify the MAC address to use for all L2 IS-IS routers. The MAC address should be a multicast address. You should shut/no shut the IS-IS instance to make the change operational.

Default all-l2isis 01-80-C2-00-02-11

Parameters *ieee-address* — Specifies the destination MAC address for all L2 IS-IS neighbors on the link for this IS-IS instance.

authentication-check

Syntax [**no**] **authentication-check**

Context config>router>isis

Description This command sets an authentication check to reject PDUs that do not match the type or key requirements.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.

The **no** form of this command allows authentication mismatches to be accepted and generate a log event.

Default **authentication-check** — Rejects authentication mismatches.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>router>isis
config>router>isis>level *level-number*

Description This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The **authentication-type** statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of the command removes the authentication key.

Default **no authentication-key** — No authentication key is configured.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax **authentication-type** {**password | message-digest**}
no authentication

Context config>router>isis
config>router>isis>level *level-number*

Description This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.

Both the authentication key and the authentication type on a segment must match. The **authentication-key** statement must also be included.

Configure the authentication type on the global level in the **config>router>isis** context.

Configure or override the global setting by configuring the authentication type in the **config>router>isis>level** context.

The **no** form of the command disables authentication.

Default **no authentication-type** — No authentication type is configured and authentication is disabled.

Parameters **password —** Specifies that simple password (plain text) authentication is required.

message-digest — Specifies that MD5 authentication in accordance with RFC2104 is required.

auth-keychain

Syntax **auth-keychain** *name*

Context config>router>isis>
config>router>isis>level
config>service>vprn>isis>
config>service>vprn>isis>level

Description This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default no auth-keychain

Parameters *name* — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

hello-auth-keychain

Syntax **hello-auth-keychain** *name*

Context config>router>isis
config>router>isis>level
config>service>vprn>isis>interface
config>service>vprn>isis>interface>level

Description This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default no hello-auth-keychain

Parameters *name* — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

default-route-tag

Syntax **default-route-tag** *tag*
no default-route-tag

Context config>router>isis

Description This command configures the route tag for default route.

Parameters *tag* — Assigns a default tag.

Values 1 — 4294967295

csnp-authentication

Syntax [**no**] **csnp-authentication**

Context config>router>isis
config>router>isis>level *level-number*

Description This command enables authentication of individual IS-IS packets of complete sequence number PDUs (CSNP) type.

The **no** form of the command suppresses authentication of CSNP packets.

csnp-interval

Syntax **csnp-interval** *seconds*
no csnp-interval

Context config>router>isis>interface *ip-int-name*

Description This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically.

The **no** form of the command reverts to the default value.

Default **csnp-interval 10** — CSN PDUs are sent every 10 seconds for LAN interfaces.

csnp-interval 5 — CSN PDUs are sent every 5 seconds for point-to-point interfaces.

Parameters *seconds* — The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer.

Values 1 — 65535

link-group

Syntax **link-group** *link-group-name*
no link-group

Context config>router>isis

Description This command specifies the IS-IS link group associated with this particular level of the interface.

Parameters *link-group-name* — Specifies an IS-IS link group on the system up to 32 characters in length.

description

Syntax **description** *string*
no description

Context config>router>isis>link-group

Description This command adds a description string to the associated link-group. The string can be up to 256 characters long and can only contain printable characters. If the command is issued in the context of a link-group that already contains a description then the previous description string is replaced.

The **no** form of the command removes the description from the associated link-group.

Parameters *string* — Character string to be associated with the associated link-group.

member

Syntax [**no**] **member** *interface-name*

Context config>router>isis>link-group

Description This command adds or removes a links to the associated link-group. The interface name should already exist before it is added to a link-group.

The **no** form of the command removes the specified interface from the associated link-group.

Parameters *interface-name* — Name of the interface to be added or removed from the associated link-group.

oper-members

Syntax **oper-members** [**0-8**]
no oper-members

Context config>router>isis>link-group

Description This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3.

The **no** form of the command reverts the oper-members limit to 1.

Default oper-members 0

revert-members

Syntax **revert-members** [**0-8**]
no revert-members

Context config>router>isis>link-group

Description This command sets the threshold for the minimum number of operational links to return the associated link-group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured revert-member threshold then the configured offsets are removed.

The **no** form of the command reverts the revert-members threshold back to the default which is equal to the oper-member threshold value.

Default revert-members *oper-members*

ipv4-unicast-metric-offset

Syntax **ipv4-unicast-metric-offset** *offset-value*
no ipv4-unicast-metric-offset

Context config>router>isis>link-group

Description This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric.

The **no** form of the command reverts the offset value to 0.

Default no ipv4-unicast-metric-offset

Parameters *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 — 6777215

ipv6-unicast-metric-offset

Syntax **ipv6-unicast-metric-offset** *offset-value*
no ipv6-unicast-metric-offset

Context config>router>isis>link-group

Description This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the **oper-members** threshold, the configured offset is applied to the interface metric for the IPv6 topology.

The **no** form of the command reverts the offset value to 0.

Default no ipv6-unicast-metric-offset

Parameters *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

Values 0 — 6777215

ipv4-multicast-metric-offset

Syntax **ipv4-multicast-metric-offset** *offset-value*
no ipv4-multicast-metric-offset

Context config>router>isis>link-group

Description This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology

The **no** form of the command reverts the offset value to 0.

Default no ipv4-multicast-metric-offset

Parameters *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

Values 0 — 6777215

ipv6-multicast-metric-offset

Syntax **ipv6-multicast-metric-offset** *offset-value*
no ipv6-multicast-metric-offset

Context config>router>isis>link-group

Description This command sets the offset value for the IPv6 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv6 multicast topology.

The no form of the command reverts the offset value to 0.

Default no ipv6-multicast-metric-offset

Parameters *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold

Values 0 — 6777215

default-metric

Syntax **default-metric** *ipv4 metric*
no default-metric

Context config>router>isis>level

Description This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.

Default 10

ipv4 metric — Specifies the default metric for IPv4 unicast.

Values 1 — 16777215

default-ipv4-multicast-metric

Syntax **default-ipv4-multicast-metric** *metric*
no default-ipv4-multicast-metric

Context config>router>isis>level

Description This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3).

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

Default 10

Parameters *metric* — Specifies the default metric for interfaces in the IPv4 multicast topology (MT3)

Values 1 — 16777215

default-ipv6-multicast-metric

Syntax **default-ipv6-multicast-metric** *metric*
no default-ipv6-multicast-metric

Context config>router>isis>level

Description This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4).

The **no** form of this command deletes the specified default metric and reverts to using the system default of 10.

Default 10

Parameters *metric* — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4).

1 — 16777215

default-ipv6-unicast-metric

Syntax **default-ipv6-unicast-metric** *ipv6 metric*
no default-ipv6-unicast-metric

Context config>router>isis>level

Description This command specifies the default metric for IPv6 unicast.

Default no default-ipv6-unicast-metric

Parameters *ipv6-metric* — Specifies the default metric for IPv6 unicast.

Values 1 — 16777215

disable-ldp-sync

Syntax [**no**] **disable-ldp-sync**

Context config>router>isis

Description This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF or IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertized cost is different. It will then disable IGP-LDP synchornization for all interfaces. This command does not delete the interface configuration. The **no** form of this command has to be entered to re-enable IGP-LDP synchronization for this routing protocol.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the ldp-sync-timer is configured.

Default no disable-ldp-sync

export

Syntax [**no**] **export** *policy-name* [*policy-name*...up to 5 max]

Context config>router>isis

Description This command configures export routing policies that determine the routes exported from the routing table to IS-IS.

If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

If an **aggregate** command is also configured in the **config>router** context, then the aggregation is applied before the export policy is applied.

Routing policies are created in the **config>router>policy-options** context.

The **no** form of the command removes the specified *policy-name* or all policies from the configuration if no *policy-name* is specified.

Default **no export** — No export policy name is specified.

Parameters *policy-name* — The export policy name. Up to five *policy-name* arguments can be specified.

export-limit

Syntax **export-limit** *number* [**log** *percentage*]
no export-limit

Context config>router>isis
config>service>vprn>isis

Description This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.

The **no** form of the command removes the parameters from the configuration.

Default no export-limit, the export limit for routes or prefixes is disabled.

Parameters *number* — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.

Values 1 — 4294967295

log *percentage* **—** Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.

Values 1 — 100

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>router>isis>level *level-number*

Description This command configures the external route preference for the IS-IS level.

The **external-preference** command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the **ecmp** in the **config>router** context.

Default Default preferences are listed in the following table:

Route Type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
OSPF internal routes	10	No
IS-IS Level 1 internal	15	Yes*
IS-IS Level 2 internal	18	Yes*
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
TMS	167	No
BGP	170	Yes

*. Internal preferences are changed using the **preference** command in the config>router>isis>level *level-number* context

Parameters *preference* — The preference for external routes at this level as expressed.

Values 1 — 255

graceful-restart

Syntax [**no**] **graceful-restart**

Context config>router>isis
config>service>vprn>isis

Description This command enables graceful-restart helper support for IS-IS. The router will act as a helper to neighbors who are graceful-restart-capable and are restarting.

When the control plane of a graceful-restart-capable router fails, the neighboring routers (graceful-restart helpers) temporarily preserve adjacency information so packets continue to be forwarded through the failed graceful-restart router using the last known routes. If the control plane of the graceful-restart router comes back up within the timer limits, then the routing protocols re-converge to minimize service interruption.

The **no** form of the command disables graceful restart and removes all graceful restart configurations in the IS-IS instance.

Default **disabled**

helper-disable

Syntax [**no**] **helper-disable**

Context config>router>isis>graceful-restart
config>service>vprn>isis>graceful-restart

Description This command disables the helper support for graceful restart.

When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the router will not help the neighbors to restart).

The **no helper-disable** command enables helper support and is the default when graceful-restart is enabled.

Default disabled

hello-authentication

Syntax [**no**] **hello-authentication**

Context config>router>isis
config>router>isis>level *level-number*
config>service>vprn>isis
config>service>vprn>isis>interface
config>service>vprn>isis>level

Description This command enables authentication of individual IS-IS packets of HELLO type.

The **no** form of the command suppresses authentication of HELLO packets.

hello-padding

Syntax [no] **hello-padding {adaptive | loose | strict}**

Context config>router>isis
config>service>vprn>isis

Description This command enables IS-IS Hello (IIH) message padding to ensure that IS-IS LSPs can traverse the link. When this option is enabled, IS-IS Hello messages are padded to the maximum LSP MTU value, which can be set with the "lsp-mtu-size" command.

The **no** form of the command disables IS-IS hello padding.

Default **no hello-padding** — hello padding is not configured

Parameters **adaptive —** Specifies the adaptive padding option; this option is able to detect LSP MTU asymmetry from one side of the connection but uses more overhead than loose padding.

1. point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in state up. If the implementation supports RFC 3373/5303, “Three-Way Handshake for IS-IS Point-to- Point Adjacencies” then this is when the three-way state is Up. If the implementation use the “classic” algorithm described in ISO 10589, this is when adjacency state is Up. If the neighbor does not support the adjacency state TLV, then padding continues.
2. broadcast interface—Padding starts until at least one adjacency is up on the interface.

loose — Specifies the loose padding option; the loose padding may not be able to detect certain situations such as asymmetrical LSP MTUs between the routing devices.

1. point-to-point interface—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state.
2. broadcast interface—Padding starts until there is at least one adjacency (broadcast only has up/down) is up on the interface.

strict — Specifies the strict padding option; this option is the most overhead-intensive but detects LSP MTU issues on both sides of a link.

1. point-to-point interface—Padding is done for all adjacency states, and is continuous.
2. broadcast interface—Padding is done for all adjacency states, and is continuous.

ignore-lsp-errors

Syntax [no] **ignore-lsp-errors**

Context config>router>isis
config>service>vprn>isis

Description This command specifies that IS-IS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged.

The **no** form of the command specifies that IS-IS will not ignore LSP errors.

ignore-narrow-metric

Syntax [**no**] **ignore-narrow-metric**

Context config>router>isis

Description This command specifies that IS-IS will ignore links with narrow metrics when wide-metrics support has been enabled.

The **no** form of the command specifies that IS-IS will not ignore these links.

iid-tlv-enable

Syntax [**no**] **iid-tlv-enable**

Context config>router>isis

Description This command specifies whether Instance Identifier (IID) TLV has been enabled or disabled for this IS-IS instance.

When enabled, each IS-IS instance marks its packets with the IID TLV containing its unique 16-bit IID for the routing domain. You should shut/no shut the isis instance to make the change operational.

Default no iid-tlv-enable

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>isis

Description This command creates the context to configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas.

When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted.

The **no** form of the command removes IS-IS from the interface.

The **shutdown** command in the **config>router>isis>interface** context administratively disables IS-IS on the interface without affecting the IS-IS configuration.

Default **no interface** — No IS-IS interfaces are defined.

Parameters *ip-int-name* — Identify the IP interface name created in the **config>router>interface** context. The IP interface name must already exist.

bfd-enable

Syntax [**no**] **bfd-enable** {**ipv4 | ipv6**} *[include-bfd-tlv]*

Context config>router>isis>interface

Description This command enables the use of bi-directional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on an IPv4 or IPv6 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set by the BFD command under the IP interface. This command must be given separately to enable/disable BFD for both IPv4 and IPv6.

The **no** form of this command removes BFD from the associated adjacency.

Default no bfd-enable ipv4

Parameters *include-bfd-tlv* — Enables support for the IS-IS BFD TLV options, specified in RFC 6213, which specifies that a BFD session must be established before an IS-IS adjacency can transition to the established state. This option should be enabled on all IS-IS neighbors on a shared interface.

default-instance

Syntax [**no**] **default-instance**

Context config>router>isis>interface *ip-int-name*
config>service>vprn>isis>interface i*p-int-name*

Description This command enables a non-MI capable router to establish an adjacency and operate with an SR OS router in a non-zero instance. If the router does not receive IID-TLVs, it will establish an adjacency in a single instance. Instead of establishing an adjacency in the standard instance 0, the router will establish an adjacency in the configured non-zero instance. The router will then operate in the configured non-zero instance so that it appears to be in the standard instance 0 to its neighbor. This feature is supported on point-to-point interfaces, broadcast interfaces are not supported.

The **no** form of this command disables the functionality so that the router can only establish adjacencies in the standard instance 0.

Default no default-instance

hello-authentication-key

Syntax **hello-authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no hello-authentication-key

Context config>router>isis>interface *ip-int-name*
config>router>isis>if>level *level-number*
config>service>vprn>isis>interface
config>service>vprn>isis>level

Description This command configures the authentication key (password) for hello PDUs. Neighboring routers use the password to verify the authenticity of hello PDUs sent from this interface. Both the hello authentication key and the hello authentication type on a segment must match. The **hello-authentication-type** must be specified.

To configure the hello authentication key in the interface context use the **hello-authentication-key** in the **config>router>isis>interface** context.

To configure or override the hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>interface>level** context.

If both IS-IS and hello-authentication are configured, hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including hello) protocol PDUs.

When the hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

The **no** form of the command removes the authentication-key from the configuration.

Default **no hello-authentication-key** — No hello authentication key is configured.

Parameters *authentication-key* — The hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

hello-authentication-type

Syntax **hello-authentication-type** {**password** | **message-digest**}
no hello-authentication-type

Context config>router>isis>interface *ip-int-name*
config>router>isis>if>level *level-number*
config>service>vprn>isis>interface
config>service>vprn>isis>level

Description This command enables hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello **authentication-key** statement must also be included.

To configure the hello authentication type at the interface context, use **hello-authentication-type** in the **config>router>isis>interface** context.

To configure or override the hello authentication setting for a given level, configure the **hello-authentication-type** in the **config>router>isis>interface>level** context.

The **no** form of the command disables hello authentication.

Default **no hello-authentication-type** — Hello authentication is disabled.

Parameters **password —** Specifies simple password (plain text) authentication is required.

message-digest — Specifies MD5 authentication in accordance with RFC2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>router>isis>if>level *level-number*

Description This command configures the interval in seconds between hello messages issued on this interface at this level.

The **no** form of the command to reverts to the default value.

Default **3** — Hello interval default for the designated intersystem.

9 — Hello interval default for non-designated intersystems.

Parameters *seconds* — The hello interval in seconds expressed as a decimal integer.

Values 1 — 20000

hello-multiplier

Syntax **hello-multiplier** *multiplier*
no hello-multiplier

Context config>router>isis>if>level *level-number*

Description This command configures the number of missing hello PDUs from a neighbor after the router declares the adjacency down.

The **no** form of the command reverts to the default value.

Default **3** — The router can miss up to 3 hello messages before declaring the adjacency down.

Parameters *multiplier* — The multiplier for the hello interval expressed as a decimal integer.

Values 2 — 100

ipv6-unicast-metric

Syntax **ipv6-unicast-metric** *metric*
no ipv6-unicast-metric

Context config>router>isis>if>level

Description This command configures IS-IS interface metric for IPv6 unicast.

The **no** form of this command removes the metric from the configuration.

Parameters *metric* — Specifies the IS-IS interface metric for IPv6 unicast.

Values 1 — 16777215

if-topology

Syntax **if-topology** *mt-id*
no if-topology *mt-id*

Context config>router>is-is>interface>level

Description This command links the associated interface with the specified IS-IS topology.

By default all IS-IS interfaces should be associated the respective unicast topology. To exclude an interface from the respective unicast topology use the command no if-topology <0|2>.

The **no** form of this command deletes the specified IS-IS topology.

Default No default

Parameters *mt-id* — Specify the topology to be created.

Values 3 or 4

interface-type

Syntax **interface-type** {**broadcast | point-to-point**}
no interface-type

Context config>router>isis>interface *ip-int-name*

Description This command configures the IS-IS interface type as either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually.

The **no** form of the command reverts to the default value.

Special Cases **SONET —** Interfaces on SONET channels default to the point-to-point type.

Ethernet or Unknown — Physical interfaces that are Ethernet or unknown default to the broadcast type.

Default **point-to-point** — For IP interfaces on SONET channels.

broadcast — For IP interfaces on Ethernet or unknown type physical interfaces.

Parameters **broadcast —** Configures the interface to maintain this link as a broadcast network.

point-to-point — Configures the interface to maintain this link as a point-to-point link.

ipv4-node-sid

Syntax **ipv4-node-sid index** *value*
ipv4-node-sid label *value*
no ipv4-node-sid

Context config>router>isis>interface

Description This command assigns a node SID index or label value to the prefix representing the primary address of an IPv4 network interface of type loopback. Only a single node SID can be assigned to an interface. The secondary address of an IPv4 interface cannot be assigned a node SID index and does not inherit the SID of the primary IPv4 address.

The above command should fail if the network interface is not of type loopback or if the interface is defined in an IES or a VPRN context. Also, assigning the same SID index/label value to the same interface in two different IGP instances is not allowed within the same node.

The value of the label or index SID is taken from the range configured for this IGP instance. When using the global mode of operation, a new segment routing module checks that the same index or label value cannot be assigned to more than one loopback interface address. When using the per-instance mode of operation, this check is not required since the index and thus label ranges of the various IGP instance are not allowed to overlap.

Parameters **index** *value* **—** integer.

Values 0 — 4294967295

Default none

label *value* **—** integer.

Values 0 — 4294967295

Default none

ipv4-multicast-routing

Syntax **ipv4-multicast-routing {native | mt}**
[no] ipv4-multicast-routing

Context config>router>isis

Description The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM.

The **no** ipv4-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.

Default ipv4-multicast-routing native

Parameters **native —** Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.

mt — Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

ipv6-multicast-routing

Syntax **ipv6-multicast-routing {native | mt}**
[no] ipv6-multicast-routing

Context config>router>isis

Description The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv6 multicast RTM.

The **no** ipv6-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.

Default ipv6-multicast-routing native

Parameters **native —** Causes IPv6 routes from the MT0 topology to be added to the multicast RTM for RPF checks.

mt — Causes IPv6 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

ipv4-routing

Syntax [**no**] **ipv4-routing**

Context config>router>isis

Description This command specifies whether this IS-IS instance supports IPv4.

The **no** form of the command disables IPv4 on the IS-IS instance.

Default ipv4-routing

ipv6-routing

Syntax [**no**] **ipv6-routing** {**native** | **mt**}

Context config>router>isis

Description This command enables IPv6 routing.

The **no** form of the command disables support for IS-IS IPv6 TLVs for IPv6 routing.

Default disabled

Parameters **native —** Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.

mt — Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

ldp-over-rsvp

Syntax [**no**] **ldp-over-rsvp**

Context config>router>isis

Description This command allows LDP over RSVP processing in IS-IS.

The **no** form of the command disables LDP over RSVP processing.

Default no ldp-over-rsvp

level

Syntax **level {1 | 2}**

Context config>router>isis
config>router>isis>interface *ip-int-name*

Description This command creates the context to configure IS-IS Level 1 or Level 2 area attributes.

A router can be configured as a Level 1, Level 2, or Level 1-2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies will not established over this interface.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

level> no hello-authentication-key
level> no hello-authentication-type
level> no hello-interval
level> no hello-multiplier
level> no metric
level> no passive
level> no priority

Special Cases **Global IS-IS Level —** The **config>router>isis** context configures default global parameters for both Level 1 and Level 2 interfaces.

IS-IS Interface Level — The **config>router>isis>interface** context configures IS-IS operational characteristics of the interface at Level 1 and/or Level 2. A logical interface can be configured on one Level 1 and one Level 2. In this case, each level can be configured independently and parameters must be removed independently.

By default an interface operates in both Level 1 and Level 2 modes.

Default level **1** or level **2**

Parameters 1 — Specifies the IS-IS operational characteristics of the interface at level 1.

2 — Specifies the IS-IS operational characteristics of the interface at level 2.

level-capability

Syntax **level-capability** {**level-1** | **level-2** | **level-1/2**}
no level-capability

Context config>router>isis
config>router>isis>interface *ip-int-name*

Description This command configures the routing level for an instance of the IS-IS routing process.

An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 *and* 2.

Table 10 displays configuration combinations and the potential adjacencies that can be formed.

Table 10: Potential Adjacency

Global Level	Interface Level	Potential Adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

The **no** form of the command removes the level capability from the configuration.

Special Cases **IS-IS Router —** In the **config>router>isis** context, changing the **level-capability** performs a restart on the IS-IS protocol instance.

IS-IS Interface — In the **config>router>isis>interface** context, changing the **level-capability** performs a restart of IS-IS on the interface.

Default **level-1/2**

Parameters **level-1 —** Specifies the router/interface can operate at Level 1only.

level-2 — Specifies the router/interface can operate at Level 2 only.

level-1/2 — Specifies the router/interface can operate at both Level 1 and Level 2.

loopfree-alternate

Syntax **loopfree-alternate [remote-lfa [max-pq-cost *value*]]**
no loopfree-alternate

Context config>router>isis

Description This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol or under the OSPF routing protocol instance.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The user enables the remote LFA next-hop calculation by the IGP LFA SPF by appending the **remote-lfa** option. When this option is enabled in an IGP instance, SPF performs the remote LFA additional computation following the regular LFA next-hop calculation when the latter resulted in no protection for one or more prefixes which are resolved to a given interface.

Remote LFA extends the protection coverage of LFA-FRR to any topology by automatically computing and establishing/tearing-down shortcut tunnels, also referred to as repair tunnels, to a remote LFA node which puts the packets back into the shortest without looping them back to the node which forwarded them over the repair tunnel. The remote LFA node is referred to as PQ node. A repair tunnel can in theory be an RSVP LSP, a LDP-in-LDP tunnel, or a SR tunnel. In this feature, it is restricted to use SR repair tunnel to the remote LFA node.

The remote LFA algorithm is a per-link LFA SPF calculation and not a per-prefix like the regular LFA one. So, it provides protection for all destination prefixes which share the protected link by using the neighbor on the other side of the protected link as a proxy for all these destinations.

The no form of this command disables the LFA computation by IGP SPF.

Default no loopfree-alternate

Parameters **max-pq-cost** *value* — integer used to limit the search of candidate P and Q nodes in remote LFA by setting the maximum IGP cost from the router performing remote LFA calculation to the candidate P or Q node.

Values 0 – 4294967295

Default none

loopfree-alternate-exclude

Syntax **loopfree-alternate-exclude prefix-policy** *prefix-policy* [*prefix-policy*... **up to 5**]
no loopfree-alternate-exclude

Context config>router>isis
config>service>vprn>isis

Description This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF. Note that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **loopfree-alternate-exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form deletes the exclude prefix policy.

Parameters **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

lfa-policy-map

Syntax **lfa-policy-map route-nh-template** *template-name*
no lfa-policy-map

Context config>router>isis>interface

Description This command applies a route next-hop policy template to an OSPF or IS-IS interface.

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

Parameters *template-name* — Specifies the name of the template, up to 32 characters.

loopfree-alternate-exclude

Syntax [**no**] **loopfree-alternate**

Context configure>router>isis>level
configure>router>isis>interface
configure>service>vprn>isis>level
configure>service>vprn>isis>interface

Description This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command re-instates the default value for this command.

Default no loopfree-alternate-exclude

lsp-pacing-interval

Syntax **lsp-pacing-interval** *milliseconds*
no lsp-pacing-interval

Context config>router>isis>interface

Description This command configures the interval between LSP PDUs sent from this interface.

To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSPs). If a value of zero is configured, no LSPs are sent from the interface.

The **no** form of the command reverts to the default value.

Default **100** — LSPs are sent in 100 millisecond intervals.

Parameters *milliseconds* — The interval in milliseconds that IS-IS LSPs can be sent from the interface expressed as a decimal integer.

Values 0 — 65535

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
no lsp-lifetime

Context config>router>isis

Description This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula: lsp-lifetime/2

The **no** form of the command reverts to the default value.

Default **1200** — LSPs originated by the router should be valid for 1200 seconds (20 minutes).

Parameters *seconds* — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

Values 350 — 65535

lsp-mtu-size

Syntax **lsp-mtu-size** *size*
no lsp-mtu-size

Context config>router>isis

Description This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then IS-IS must be restarted in order for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context. Note: Using the **exec** command to execute a configuration file to change the LSP MTU-size from its default value will automatically bounce IS-IS for the change to take effect.

The **no** form of the command reverts to the default value.

Default 1492

Parameters *size* — Specifies the LSP MTU size.

Values 490 — 9190

lsp-refresh-interval

Syntax **lsp-refresh-interval** *seconds*
no lsp-refresh-interval

Context config>router>isis

Description This command configures the IS-IS LSP refresh timer interval. When configuring the LSP refresh interval, the value that is specified for **lsp-lifetime** must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.

The no form of the command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the no lsp-refresh-interval command will be rejected.

Default 600 seconds

Parameters *seconds* — Specifies the refresh interval.

Values 150— 65535

lsp-wait

Syntax **lsp-wait** *lsp-wait* [*lsp-initial-wait* [*lsp-second-wait*]]

Context config>router>isis

Description This command is used to customize the throttling of IS-IS LSP-generation. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second **lsp-wait** timer until a maximum value is reached.

Parameters *lsp-max-wait* — Specifies the maximum interval in seconds between two consecutive ocurrences of an LSP being generated.

Values 1 — 120

Default 5

lsp-initial-wait — Specifies the initial LSP generation delay in seconds.

Values 0 — 100

Default 0

lsp-second-wait — Specifies the hold time in seconds between the first and second LSP generation.

Values 1 — 100

Default 1

ipv4-multicast

Syntax **[no] ipv4-multicast**

Context config>router>is-is>multi-topology

Description This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance.

The **no** form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.

Default no ipv4-multicast

ipv6-multicast

Syntax **[no] ipv6-multicast**

Context config>router>is-is>multi-topology

Description This command enables support for the IPv6 topology (MT4) within the associate IS-IS instance.

The **no** form of this command disables support for the IPv6 topology (MT4) within the associated IS-IS instance.

Default no ipv6-multicast

mcast-import-ipv6

Syntax [**no**] **mcast-import-ipv6**

Context configure>router>isis

Description This command administratively enables/disables submission of routes into the IPv6 multicast RTM by IS-IS.

multi-topology

Syntax [**no**] **multi-topology**

Context config>router>isis

Description This command enables IS-IS multi-topology support.

Default disabled

topology

Syntax **topology** *mt-id* **rtm** *rtm-id* **|** *rtm-name*
no topology *mt-id*

Context config>router>is-is>multi-topology

Description This command creates a new topology within the associate IS-IS instance. In addition, it associates the IS-IS topology with the specified RTM instance. Routes generated from the topology SPF calculation are in turn added to this associate RTM instance.

The **no** form of this command deletes the specified IS-IS topology.

Default No default

Parameters *mt-id* — Specify the topology to be created (Note: in Release 11.0 this parameters is limited to 3 or 4.

rtm-id — RTM Instance ID that is to be associated with the new IS-IS topology.

Values integer: 3 — 32

rtm-name — string name given to the RTM instance.

ipv6-unicast

Syntax [**no**] **ipv6-unicast**

Context config>router>isis>multi-topology

Description This command enables multi-topology TLVs.

The no form of the command disables multi-topology TLVs.

multicast-import

Syntax [**no**] **multicast-import**

Context config>router>isis

Description This command enables the submission of routes into the multicast Route Table Manager (RTM) by IS-IS.

The **no** form of the command disables the submission of routes into the multicast RTM.

Default no multicast-import

mesh-group

Syntax **mesh-group** {**value** | **blocked**}
no mesh-group

Context config>router>isis>interface *ip-int-name*

Description This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to created isolated islands that do not receive updates as (other) links fail.

The **no** form of the command removes the interface from the mesh group.

Default **no mesh-group** — The interface does not belong to a mesh group.

Parameters **value —** The unique decimal integer value distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.

Values 1 — 2000000000

blocked — Prevents an interface from flooding LSPs.

ipv6-unicast-disable

Syntax [**no**] **ipv6-unicast-disable**

Context config>router>isis>if

Description This command disables IS-IS IPv6 unicast routing for the interface.

By default IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the **config>router>isis>ipv6-routing mt** command is configured.

The **no** form of the command enables IS-IS IPv6 unicast routing for the interface.

metric

Syntax **metric** *metric*
no metric

Context config>router>isis>if>level *level-number*

Description This command configures the metric used for the level on the interface.

In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default of 10 is used unless reference bandwidth is configured.

The **no** form of the command reverts to the default value.

Default **10** — A metric of 10 for the level on the interface is used.

Parameters *metric* — The metric assigned for this level on this interface.

Values 1 — 16777215

advertise-passive-only

Syntax [**no**] **advertise-passive-only**

Context config>router>isis

Description This command enables and disables IS-IS to advertise only prefixes that belong to passive interfaces.

advertise-router-capability

Syntax **advertise-router-capability** {**area** | **as**}
no advertise-router-capability

Context config>router>isis

Description This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability.

The parameters (area and as) control the scope of the capability advertisements.

The **no** form of this command, disables this capability.

Default **no advertise-router-capability**

Parameters **area —** are only advertised within the area of origin.

as — are only advertised throughout the entire autonomous system

area-id

Syntax [**no**] **area-id** *area-address*

Context config>router>isis

Description This command was previously named the **net** *network-entity-title* command. The **area-id** command allows you to configure the area ID portion of NSAP addresses which identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.

A maximum of 3 area addresses can be configured.

NSAP addresses are divided into three parts. Only the area ID portion is configurable.

- Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.
- System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.
- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first area address.

The **no** form of the command removes the area address.

Default **none** — No area address is assigned.

Parameters *area-address* — The 1 — 13-byte address. Of the total 20 bytes comprising the NET, only the first 13 bytes can be manually configured. As few as one byte can be entered or, at most, 13 bytes. If less than 13 bytes are entered, the rest is padded with zeros.

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>router>isis

Description This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The overload command is cleared from the configuration after a reboot if **overload-on-boot** is configured with or without a timeout value. To keep the IS-IS router in the overload state indefinitely after rebooting, configure overload-on-boot with no timeout value or configure the overload command with no overload-on-boot command.

The **overload** command can be useful in circumstances where the router is overloaded or used prior to executing a **shutdown** command to divert traffic around the router.

The **no** form of the command causes the router to exit the overload state.

Default **no overload**

Parameters *seconds* — The time, in seconds, that this router must operate in overload state.

Default infinity (overload state maintained indefinitely)

Values 60 — 1800

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload-on-boot

Context config>router>isis

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

1. The timeout timer expires.
2. A manual override of the current overload state is entered with the **config>router>isis>no overload** command.

The **no overload** command does not affect the **overload-on-boot** function.

If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

L1 LSDB Overload : Manual on boot (Indefinitely in overload)

L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **config>router>isis>no overload** command.

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

L1 LSDB Overload : Manual on boot (Overload Time Left : 17)

L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default no overload-on-boot

Use the **show router isis status** command to display the administrative and operational state as well as all timers.

Parameters **timeout** *seconds* **—** Configure the timeout timer for overload-on-boot in seconds.

Values 60 — 1800

poi-tlv-enable

Syntax **poi-tlv-enable**
no poi-tlv-enable

Context config>router>isis

Description Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge.

The **no** form of the command removes the POI functionality from the configuration.

Default no poi-tlv-enable

passive

Syntax [**no**] **passive**

Context config>router>isis>if
config>router>isis>if>level

Description This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.

The **no** form of the command removes the passive attribute.

Special Cases **Service Interfaces —** Service interfaces (defined using the service-prefix command in **config>router**) are passive by default.

All other Interfaces — All other interfaces are not passive by default.

Default **passive** — Service interfaces are passive.
no passive — All other interfaces are not passive.

preference

Syntax **preference** *preference*
no preference

Context config>router>isis>level

Description This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the **config>router** context.

Default Default preferences are listed in the following table:

Route Type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
OSPF internal routes	10	No
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes*
IS-IS level 2 external	165	Yes*
BGP	170	Yes

*. External preferences are changed using the **external-preference** command in the config>router>isis>level *level-number* context.

Parameters *preference* — The preference for external routes at this level expressed as a decimal integer.

Values 1 — 255

priority

Syntax **priority** *number*
no priority

Context config>router>isis>if>level *level-number*

Description This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.

This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.

The **no** form of the command reverts to the default value.

Default **64**

Parameters *number* — The priority for this interface at this level.

Values 0 — 127

sd-offset

Syntax **sd-offset** *offset-value*
no sd-offset

Context config>router>isis>if>level

Description If the pre-FEC error rate of the associated DWDM port crosses the configured **sd-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sd-threshold** value is configured under that port.

The **no** form of the command reverts the offset value to 0.

Default no sd-offset

Parameters *offset-value* — Specifies the amount the interface metric is increased by if the **sd-threshold** is crossed.

Values 0 — 16777215

sf-offset

Syntax **sf-offset** *offset-value*
no sf-offset

Context config>router>isis>if>level

Description If the pre-FEC error rate of the associated DWDM port crosses the configured **sf-threshold**, this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the **sf-threshold** value is configured under that port.

The **no** form of the command reverts the offset value to 0.

Default no sf-offset

Parameters *offset-value* — Specifies the amount the interface metric is increased by if the **sf-threshold** is crossed.

Values 0 — 16777215

psnp-authentication

Syntax [**no**] **psnp-authentication**

Context config>router>isis
config>router>isis>level

Description This command enables authentication of individual IS-IS packets of partial sequence number PDU (PSNP) type.

The **no** form of the command suppresses authentication of PSNP packets.

reference-bandwidth

Syntax **reference-bandwidth** *bandwidth-in-kbps*
reference-bandwidth [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]
no reference-bandwidth

Context config>router>isis

Description This command configures the reference bandwidth that provides the basis of bandwidth relative costing.

In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$$

If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 M/bps interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed. (See wide-metrics-only in the **config>router>isis** context.)

If the reference bandwidth is not configured, then all interfaces have a default metric of 10.

The **no** form of the command reverts to the default value.

Default **no reference-bandwidth** — No reference bandwidth is defined. All interfaces have a metric of 10.

Parameters *bandwidth-in-kbps* — The reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 1000000000

tbps *Tera-bps* — The reference bandwidth in terabits per second expressed as a decimal integer.

Values 1 — 4

gbps *Giga-bps* **—** The reference bandwidth in gigabits per second expressed as a decimal integer.

Values 1 — 999

mbps *Mega-bps* **—** The reference bandwidth in megabits per second expressed as a decimal integer.

Values 1 — 999

kbps *Kilo-bps* **—** reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 999

rib-priority

Syntax **rib-priority** {**high**} *prefix-list-name* **| tag** *tag-value*
no rib-priority

Context config>router>isis

Description This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.

The no rib-priority form of command disables RIB prioritization.

Default no rib-priority

Parameters *prefix-list-name* — specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

tag *tag-value* — specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.

Values 1 – 4294967295

rsvp-shortcut

Syntax [**no**] **rsvp-shortcut**

Context config>router>isis

Description This command enables the use of an RSVP-TE shortcut for resolving IGP routes by IS-IS or OSPF routing protocols.

This command instructs IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. Note that Dijkstra will always use the IGP metric to build the SPF tree and the LSP metric value does not update the SPF tree calculation. During the IP reach to determine the reachability of nodes and prefixes, LSPs are then overlaid and the LSP metric is used to determine the subset of paths which are equal lowest cost to reach a node or a prefix. If the user enabled the **relative-metric** option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix which is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When rsvp-shortcut is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **configure>router>mpls>lsp>to**, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the **config>router>mpls>lsp>no igp-shortcut** command.

The SPF in OSPF or IS-IS will only use RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as end-points for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
  - → select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
  - → exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
  - → prefer tunnel next-hop over ip next-hop
  - → within tunnel next-hops:
    1. select lowest endpoint to destination cost
    2. if same endpoint to destination cost, select lowest endpoint node router-id
    3. if same router-id, select lowest tunnel-index
  - → within ip next-hops:
    1. select lowest downstream router-id
    2. if same downstream router-id, select lowest interface-index

- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

The ingress IOM will spray the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the **rsvp-shortcut** and the **multicast-import** options are enabled in IGP. The unicast RTM can still make use of the tunnel next-hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.

The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

Default **no rsvp-shortcut**

segment-routing

Syntax **segment-routing**
no segment-routing

Context config>router>isis

Description This command enables the context to configure segment routing parameters within a given IGP instance.

Segment routing adds to IS-IS and OSPF routing protocols the ability to perform shortest path routing and source routing using the concept of abstract segment. A segment can represent a local prefix of a node, a specific adjacency of the node (interface/next-hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as Segment ID (SID).

When segment routing is used together with MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will thus push one or more MPLS labels.

Segment routing using MPLS labels can be used in both shortest path routing applications and in traffic engineering applications. This feature implements the shortest path forwarding application.

After segment routing is successfully enabled in the IS-IS or OSPF instance, the router will perform the following operations:

1. Advertize the Segment Routing Capability Sub-TLV to routers in all areas/levels of this IGP instance. However, only neighbors with which it established an adjacency will interpret the SID/label range information and use it for calculating the label to swap to or push for a given resolved prefix SID.
2. Advertise the assigned index for each configured node SID in the new prefix SID sub-TLV with the N-flag (node-SID flag) set. Then the segment routing module programs the incoming label map (ILM) with a pop operation for each local node SID in the data path.
3. Assign and advertize automatically an adjacency SID label for each formed adjacency over a network IP interface in the new Adjacency SID sub-TLV. The segment routing module programs the incoming label map (ILM) with a pop operation, in effect with a swap to an implicit null label operation, for each advertised adjacency SID.

4. Resolve received prefixes and if a prefix SID sub-TLV exists, the Segment Routing module programs the ILM with a swap operation and also an LTN with a push operation both pointing to the primary/LFA NHLFE. An SR tunnel is also added to the TTM.

When the user enables segment routing in a given IGP instance, the main SPF and LFA SPF are computed normally and the primary next-hop and LFA backup next-hop for a received prefix are added to RTM without the label information advertised in the prefix SID sub-TLV.

prefix-sid-range

Syntax **prefix-sid-range** {**global | start-label** *label-value* **max-index** *index-value*}
no prefix-sid-range

Context config>router>isis>segment-routing

Description This command configures the prefix SID index range and offset label value for a given IGP instance.

The key parameter is the configuration of the prefix SID index range and the offset label value which this IGP instance will use. Since each prefix SID represents a network global IP address, the SID index for a prefix must be network-wide unique. Thus, all routers in the network are expected to configure and advertise the same prefix SID index range for a given IGP instance. However, the label value used by each router to represent this prefix; that is, the label programmed in the ILM can be local to that router by the use of an offset label, referred to as a start label:

Local Label (Prefix SID) = start-label + {SID index}

The label operation in the network becomes thus very similar to LDP when operating in the independent label distribution mode (RFC 5036) with the difference that the label value used to forward a packet to each downstream router is computed by the upstream router based on advertised prefix SID index using the above formula.

There are two mutually exclusive modes of operation for the prefix SID range on the router. In the **global** mode of operation, the user configures the global value and this IGP instance will assume the start label value is the lowest label value in the SRGB and the prefix SID index range size equal to the range size of the SRGB. Once one IGP instance selected the global option for the prefix SID range, all IGP instances on the system will be restricted to do the same. The user must shutdown the segment routing context and delete the **prefix-sid-range** command in all IGP instances in order to change the SRGB. Once the SRGB is changed, the user must re-enter the **prefix-sid-range** command again. The SRGB range change will be failed if an already allocated SID index/label goes out of range.

In the per-instance mode of operation, the user partitions the SRGB into non-overlapping sub-ranges among the IGP instances. The user thus configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. Note that all resulting net label values (start-label + index} must be within the SRGB or the configuration will be failed. Furthermore, the code checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce that these ranges do not overlap. The user must shutdown the segment routing context of an IGP instance in order to change the SID index/label range of that IGP instance using the **prefix-sid-range** command. In addition, any range change will be failed if an already allocated SID index/label goes out of range. The user can however change the SRGB on the fly as long as it does not reduce the current per IGP instance SID index/label range defined with the **prefix-sid-range**. Otherwise, the user must shutdown the segment routing context of the IGP instance and delete and re-configure the **prefix-sid-range** command.

Parameters **start-label** *label-value* **—** the label offset for the SR label range of this IGP instance.

Values 0 — 524287

Default none

max-index *index-value* **—** the maximum value of the prefix SID index range for this IGP instance.

Values 1 — 524287

Default none

tunnel-mtu

Syntax **tunnel-mtu** *bytes*
no tunnel-mtu

Context config>router>isis>segment-routing

Description This command configures configure the MTU of all SR tunnels within each IGP instance.

The MTU of a SR tunnel populated into TTM is determined like in the case of an IGP tunnel; for example, LDP LSP, based on the outgoing interface MTU minus the label stack size. Remote and directed LFA can add at least two more labels to the tunnel for a total of three. There is no default value for this new command. If the user does not configure a SR tunnel MTU, the MTU will be fully determined by IGP as explained below.

The MTU of the SR tunnel is then determined as follows:

SR_Tunnel_MTU = MIN {Cfg_SR_MTU, IGP_Tunnel_MTU- 3 labels}

Where,

Cfg_SR_MTU is the MTU configured by the user for all SR tunnels within a given IGP instance using the above CLI. If no value was configured by the user, the SR tunnel MTU will be fully determined by the IGP interface calculation explained next.

IGP_Tunnel_MTU is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.

The SR tunnel MTU is dynamically updated anytime any of the above parameters used in its calculation changes. This includes when the set of the tunnel next-hops changes or the user changes the configured SR MTU or interface MTU value.

Parameters *bytes* **—** the size of the Maximum Transmission Unit (MTU) in bytes.

Values 512— 9198

Default none

tunnel-table-pref

Syntax **tunnel-table-pref** *preference*
no tunnel-table-pref

Context config>router>isis>segment-routing

Description This command configures the TTM preference of SR tunnels created by the IGP instance. This is used in the case of BGP shortcuts, VPRN auto-bind, or BGP transport tunnel when the new tunnel binding commands are configured to the **any** value which parses the TTM for tunnels in the protocol preference order. The user can choose to either go with the global TTM preference or list explicitly the tunnel types they want to use. When they list the tunnel types explicitly, the TTM preference will still be used to select one type over the other. In both cases, a fallback to the next preferred tunnel type is performed if the selected one fails. Also, a reversion to a more preferred tunnel type is performed as soon as one is available.

The segment routing module adds to TTM a SR tunnel entry for each resolved remote node SID prefix and programs the data path with the corresponding LTN with the push operation pointing to the primary and LFA backup NHLFEs .

The default preference for SR tunnels in the TTM is set lower than LDP tunnels but higher than BGP tunnels to allow controlled migration of customers without disrupting their current deployment when they enable segment routing. The following is the setting of the default preference of the various tunnel types. This includes the preference of SR tunnels based on shortest path (referred to as **SR-ISIS** and **SR-OSPF**).

The global default TTM preference for the tunnel types is as follows:

- ROUTE_PREF_RSVP 7
- ROUTE_PREF_SR_TE 8
- ROUTE_PREF_LDP 9
- ROUTE_PREF_OSPF_TTM 10
- ROUTE_PREF_ISIS_TTM 11
- ROUTE_PREF_BGP_TTM 12
- ROUTE_PREF_GRE 255

The default value for SR-ISIS or SR-OSPF is the same regardless if one or more IS-IS or OSPF instances programmed a tunnel for the same prefix. The selection of a SR tunnel in this case will be based on lowest IGP instance-id.

Parameters *preference* — integer value to represent the preference of IS-IS or OSPF SR tunnels in TTM.

Values	1— 255
Default	11

advertise-tunnel-link

Syntax [**no**] **advertise-tunnel-link**

Context config>router>isis

Description This command enables the forwarding adjacency feature. With this feature, IS-IS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs per RFC 3906.

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both **rsvp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win.

When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router-id of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA/LSP advertising the corresponding P2P unnumbered link is installed in the local routing database, then the node performs an SPF using it like any other link LSA/LSP. The link bi-directional check requires that a link, regular link or tunnel link, exists in the reverse direction for the tunnel to be used in SPF.

Note that the **igp-shortcut** option under the LSP name governs the use of the LSP with both the **rsvp-shortcut** and the **advertise-tunnel-link** options in IGP. In other words, the user can exclude a specific RSVP LSP from being used as a forwarding adjacency by entering the command **config>router>mpls>lsp>no igp-shortcut**.

The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.

The **no** form of this command disables forwarding adjacency and hence disables the advertisement of RSVP LSP into IGP.

Default no advertise-tunnel-link

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>router>isis>interface *ip-int-name*

Description This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

The **no** form of the command reverts to the default value.

Default **100**

Parameters *seconds* — The interval in seconds that IS-IS LSPs can be sent on the interface.

Values 1 — 65535

spf-wait

Syntax [**no**] **spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

Context config>router>isis

Description This command defines the maximum interval between two consecutive SPF calculations in seconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

Default no spf-wait

Parameters *spf-wait* — Specifies the maximum interval in seconds between two consecutive spf calculations.

Values 1 — 120

Default 10

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 — 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 1 — 100000

Default 1000

strict-adjacency-check

Syntax [**no**] **strict-adjacency-check**

Context config>router>isis

Description This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are different. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well.

When disabled (**no strict-adjacency-check**), both routers only need to have one common address family to establish the adjacency.

Default no strict-adjacency-check

summary-address

Syntax **summary-address** {*ip-prefix/mask* | *ip-prefix* [*netmask*]} *level* [**tag** *tag*]
no summary-address {*ip-prefix/mask* | *ip-prefix* [*netmask*]}

Context config>router>isis

Description This command creates summary-addresses.

Default none

Parameters *ip-prefix/mask* — Specifies information for the specified IP prefix and mask length.

Values	ipv4-prefix:	a.b.c.d (host bits must be 0)
	ipv4-prefix-length:	0 — 32
	ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D
	ipv6-prefix-length:	[0 — 128]

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

level — Specifies IS-IS level area attributes.

Values level-1, level-2, level-1/2

tag *tag* — Assigns a route tag to the summary address.

Values 1 – 4294967295

ignore-attached-bit

Syntax **ignore-attached-bit**
[**no**] **ignore-attached-bit**

Context config>router>isis

Description This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes.

suppress-attached-bit

Syntax **suppress-attached-bit**
no suppress-attached-bit

Context config>router>isis

Description This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it.

traffic-engineering

Syntax [**no**] **traffic-engineering**

Context config>router>isis

Description This command configures traffic-engineering and determines if IGP shortcuts are required by BGP.

Default **disabled**

unicast-import-disable

Syntax [**no**] **unicast-import-disable**

Context config>router>isis

Description This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured.

Default disabled

wide-metrics-only

Syntax [**no**] **wide-metrics-only**

Context config>router>isis>level *level-number*

Description This command enables the exclusive use of wide metrics in the LSPs for the level number.. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When wide-metrics-only is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

The **no** form of the command reverts to the default value.