
Configuration Commands

Generic Commands

shutdown

Syntax [no] shutdown

Context config>router>ospf
config>router>ospf3
config>router>ospf>area>interface
config>router>ospf3>area>interface
config>router>ospf>area>segment-routing
config>router>ospf>area>virtual-link
config>router>ospf3>area>virtual-link

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default **OSPF Protocol** — The Open Shortest Path First (OSPF) protocol is created in the **no shutdown** state.
OSPF Interface — When an IP interface is configured as an OSPF interface, OSPF on the interface is in the **no shutdown** state by default.

OSPF Global Commands

ospf

Syntax [no] **ospf** *ospf-instance* [*instance-id*] [*router-id*]

Context config>router

Description This command creates an OSPF routing instance and then enters the associated context to configure the associated protocol parameters.

Additionally, the router ID can be specified as another parameter of the OSPDF command. This parameter is required for all non-base OSPF instances.

The default value for the base instance is inherited from the configuration in the config>router context. When that is not configured the following applies:

1. The system uses the system interface address (which is also the loopback address).
2. If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

This is a required command when configuring multiple instances and the instance being configured is not the base instance. When configuring multiple instances of OSPF there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To avoid this from happening all routers in a domain should be configured with the same domain-id. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.

The default value for non-base instances is 0.0.0.0 and is invalid, in this case the instance of OSPF will not start. When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

Issue the shutdown and no shutdown commands for the instance for the new router ID to be used, or reboot the entire router.

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

Default **no ospf**

Parameters *instance-id* — Specifies a unique integer that identifies a specific instance of a version of the OSPF protocol running in the router instance specified by the router ID.

Values 1 — 31

router-id — Specifies the OSPF router ID to be used with the associated OSPF instance. The router-id must be given a dot decimal notation format.

Values 1 — 31

ospf3

Syntax **ospf3** [*instance-id*] [*router-id*]
[no] ospf3 *instance-id*

Context config>router

Description This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.

When an OSPFv3 instance is created, the protocol is enabled. To start or suspend execution of the OSPF.

The **no** form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.

Default **no default**

Parameters *instance-id* — Specify the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family. For ipv6-unicast, the instance id must be between 0 and 31. For ipv4-unicast the instance id must be between 64-95.

Values 0 — 31: IPV6 unicast

Values 64—95: IPV4 unicast

router-id — Specifies the OSPF router ID to be used with the associated OSPF instance. The router-id must be given a dot decimal notation format.

asbr

Syntax **[no] asbr** [*trace-path domain-id*]

Context config>router>ospf
 config>router>ospf3


Description This command configures the router as a Autonomous System Boundary Router (ASBR) if the router is to be used to export routes from the Routing Table Manager (RTM) into this instance of OSPF. Once a router is configured as an ASBR, the export policies into this OSPF domain take effect. If no policies are configured no external routes are redistributed into the OSPF domain.

The **no** form of the command removes the ASBR status and withdraws the routes redistributed from the Routing Table Manager into this instance of OSPF from the link state database.

When configuring multiple instances of OSPF there is a risk of loops because networks are advertised by multiple domains configured with multiple interconnections to one another. To avoid this from happening all routers in a domain should be configured with the same domain-id. Each domain (OSPF-instance) should be assigned a specific bit value in the 32-bit tag mask.

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding bit is set in the AS-external LSA. As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there

OSPF Global Commands



Domain-IDs are incompatible with any other use of normal tags. The domain ID should be configured with a value between 1 and 31 by each router in a given OSPF domain (OSPF Instance).

When an external route is originated by an ASBR using an internal OSPF route in a given domain, the corresponding (1-31) bit is set in the AS-external LSA.

As the route gets redistributed from one domain to another, more bits are set in the tag mask, each corresponding to the OSPF domain the route visited. Route redistribution looping is prevented by checking the corresponding bit as part of the export policy; if the bit corresponding to the announcing OSPF process is already set, the route is not exported there.

Default **no asbr** — The router is not an ASBR.

Parameters *domain-id* — Specifies the domain ID.

Values 1 — 31

Default 0

compatible-rfc1583

Syntax **[no] compatible-rfc1583**

Context config>router>ospf

Description This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.

RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliancy level, this command allows the router to use obsolete methods of calculation.

The **no** form of the command enables the post-RFC1583 method of summary and external route calculation.

Default **compatible-rfc1583** — RFC1583 compliance is enabled.

disable-ldp-sync

Syntax [no] disable-ldp-sync

Context config>router>ospf

Description This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF 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 advertised cost is different. It will then disable IGP-LDP synchronization 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 export *policy-name* [*policy-name*...]
no export

Context config>router>ospf
config>router>ospf3

Description This command associates export route policies to determine which routes are exported from the route table to OSPF. Export policies are only in effect if OSPF is configured as an ASBR.

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

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 will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default no export — No export route policies specified.

Parameters *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

export-limit

Syntax	export-limit <i>number</i> [<i>log percentage</i>] no export-limit
Context	config>router>ospf config>router>ospf3 config>service>vprn>ospf config>service>vprn>ospf3 config>router>ospf3
Description	<p>This command configures the maximum number of routes (prefixes) that can be exported into OSPF from the route table. After the maximum is reached, a warning log message is sent and additional routes are ignored.</p> <p>The no form of the command removes the parameters from the configuration.</p>
Default	no export-limit, the export limit for routes or prefixes is disabled.
Parameters	<p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into OSPF from the route table.</p> <p>Values 1 — 4294967295</p> <p><i>log percentage</i> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p>Values 1 — 100</p>

external-db-overflow

Syntax	external-db-overflow <i>limit interval</i> no external-db-overflow
Context	config>router>ospf config>router>ospf3
Description	<p>This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.</p> <p>The <i>limit</i> value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the <i>limit</i>, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.</p> <p>The <i>interval</i> specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.</p>

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of the command disables limiting the number of non-default AS-external-LSA entries.

Default **no external-db-overflow** — No limit on non-default AS-external-LSA entries.

Parameters *limit* — The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values 0 — 2147483674

interval — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 — 2147483674

external-preference

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

Context config>router>ospf
config>router>ospf3

Description This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference is used to decide 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 Table 6, “Route Preference Defaults by Route Type,” on page 342. 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 what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

Default **external-preference 150** — OSPF external routes have a default preference of 150.

OSPF Global Commands

Parameters *preference* — The preference for external routes expressed as a decimal integer. Defaults for different route types are listed in Table 6.

Table 6: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes*
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
TMS	167	No
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

*. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

graceful-restart

Syntax [no] graceful-restart

Context config>router>ospf
config>router>ospf3

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

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

Default no graceful-restart

helper-disable

Syntax [no] helper-disable

Context config>router>ospf>graceful-restart
config>router>ospf3>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

import

Syntax import *policy-name* [*policy-name...*(up to 5 max)]
no import

Context config>router>ospf
config>router>ospf3

Description This command applies one or more (up to 5) route policies as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions.

Default If an OSPF route has the lowest preference value among all routes to a destination it is installed in the routing table.

ldp-over-rsvp

Syntax [no] ldp-over-rsvp

Context config>router>ospf

Description This command allows LDP-over-RSVP processing in this OSPF instance.

loopfree-alternate

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

Context `config>router>ospf`
`config>router>ospf3`

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

lfa-policy-map

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

Context `config>router>ospf>area>interface`
`config>router>ospf3>area>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 **loopfree-alternate-exclude prefix-policy** *prefix-policy* [*prefix-policy...* **up to 5**]
no loopfree-alternate-exclude

Context config>router>ospf
config>router>ospf3

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.

mcast-import-ipv6

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

Context configure>router>ospf3

Description This command administratively enables the submission of routes into the IPv6 multicast RTM by OSPF3. The no form of the command disables the submission of the routes.

multicast-import

Syntax	[no] multicast-import
Context	config>router>ospf
Description	This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF. The no form of the command disables the submission of routes into the multicast RTM.
Default	no multicast-import

overload

Syntax	overload [timeout seconds] no overload
Context	config>router>ospf config>router>ospf3 config>router>ospf3
Description	<p>This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continues to reach the router.</p> <p>To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a no overload command is executed.</p> <p>If the overload command is encountered during the execution of an overload-on-boot command then this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the overload-on-boot command is saved after the overload command. However, when overload-on-boot is configured under OSPF with no timeout value configured, the router will remain in overload state indefinitely after a reboot.</p> <p>Use the no form of this command to return to the default. When the no overload command is executed, the overload state is terminated regardless of the reason the protocol entered overload state.</p>
Default	no overload
Parameters	timeout seconds — Specifies the number of seconds to reset overloading.
	Values 1 — 1800
	Default indefinite

overload-include-ext-2

Syntax	[no] overload-include-ext-2
Context	config>router>ospf config>router>ospf3
Description	This command is used to control if external type-2 routes should be re-advertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes will be advertised with the maximum metric.
Default	no overload-include-ext-2

overload-include-stub

Syntax	[no] overload-include-stub
Context	config>router>ospf config>router>ospf3
Description	This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.
Default	no overload-include-stub

overload-on-boot

Syntax	overload-on-boot [timeout <i>seconds</i>] no overload
Context	config>router>ospf config>router>ospf3
Description	<p>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:</p> <ul style="list-style-type: none"> • The timeout timer expires. • A manual override of the current overload state is entered with the no overload command. <p>The no overload command does not affect the overload-on-boot function.</p> <p>The no form of the command removes the overload-on-boot functionality from the configuration.</p> <p>The default timeout value is 60 seconds, which means after 60 seconds overload status the SR will recover (change back to non-overload status). However, when overload-on-boot is configured under OSPF with no timeout value the router will remain in overload state indefinitely after a reboot.</p>

OSPF Global Commands

Parameters `timeout seconds` — Specifies the number of seconds to reset overloading.

Values 1 — 1800

Default indefinitely in overload.

preference

Syntax `preference preference`
`no preference`

Context `config>router>ospf`
`config>router>ospf3`

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs the preference is used to decide 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 [Table 7](#). 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 what route to use is determined by the configuration of the **ecmp** in the `config>router` context.

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

Default `preference 10` — OSPF internal routes have a preference of 10.

Parameters `preference` — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in [Table 7](#).

Table 7: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes*
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
TMS	167	No
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

*. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

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>ospf
config>router>ospf3

Description This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

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

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100

OSPF Global Commands

- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** *ip-int-name* context.

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

Default **reference-bandwidth 100000000** — Reference bandwidth of 100 Gbps.

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**
no rib-priority

Context config>router>ospf
config>router>ospf3

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

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

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.

rtr-adv-lsa-limit

Syntax **rtr-adv-lsa-limit** *limit* [**log-only**] [**threshold** *percent*] [**overload-timeout** {*seconds* | **forever**}]
no rtr-adv-lsa-limit

Context config>router>ospf [instance-id]
 config>router>ospf3 [instance-id]

Description This command configures the maximum number of LSAs OSPF can learn from another router, in order to protect the system from a router that accidentally advertises a large number of LSAs. When the number of advertised LSAs reaches the configured percentage of this limit, an SNMP trap is sent. If the limit is exceeded, OSPF goes into overload.

The **overload-timeout** option allows the administrator to control how long OSPF is in overload as a result of the advertised LSA limit being reached. At the end of this duration of time the system automatically attempts to restart OSPF. One possible value for the **overload-timeout** is **forever**, which means OSPF is never restarted automatically and this corresponds to the default behavior when the **overload-timeout** option is not configured.

The **no** form of the command removes the **rtr-adv-lsa-limit**.

Default **forever**

Parameters **log-only** — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, overload is not set.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

limit — The number of LSAs that can be learned expressed as a decimal integer.

Values 1 - 4294967296

second — Specifies duration in minutes before restarting OSPF.

Values Values 1 - 1800

forever — Specifies that OSPF is restarted only after the clear router ospf|ospf3 overload rtr-adv-lsa-limit command is executed.

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>router>ospf
 config>router>ospf3

Description This command configures the router ID for the OSPF instance. This command configures the router ID for the OSPF instance.

When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the **config>router** context.

OSPF Global Commands

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **tconfig>router** context is not configured, the following applies:

- The system uses the system interface address (which is also the loopback address).
- If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

This is a **required** command when configuring multiple instances and the instance being configured is not the base instance.

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

It is possible to configure an SR OS node to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

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

advertise-router-capability

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

Context config>router>ospf
config>router>ospf3

Description This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:

- OSPF graceful restart capable: no
- OSPF graceful restart helper: yes, when enabled
- OSPF Stub Router support: yes
- OSPF Traffic Engineering support: yes, when enabled
- OSPF point-to-point over LAN: yes
- OSPF Experimental TE: no

The parameters (link, 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 **link** — are only advertised over local link and not flooded beyond

area — are only advertised within the area of origin.

as — are only advertised throughout the entire autonomous system

rsvp-shortcut

Syntax [no] **rsvp-shortcut**

Context config>router>ospf

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 endpoints 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:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index
 - within ip next-hops:
 - i. select lowest downstream router-id
 - ii. 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>ospf

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. Advertise 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 advertise 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>ospf>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:

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

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.

OSPF Global Commands

Parameters	start-label <i>label-value</i> — the label offset for the SR label range of this IGP instance. Values 0 — 524287 Default none
	max-index <i>index-value</i> — the maximum value of the prefix SID index range for this IGP instance. Values 1 — 524287 Default none

tunnel-mtu

Syntax	tunnel-mtu <i>bytes</i> no tunnel-mtu
Context	config>router>ospf>segment-routing
Description	<p>This command configures configure the MTU of all SR tunnels within each IGP instance.</p> <p>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.</p> <p>The MTU of the SR tunnel is then determined as follows:</p> $SR_Tunnel_MTU = MIN \{Cfg_SR_MTU, IGP_Tunnel_MTU - 3 \text{ labels}\}$ <p>Where,</p> <p><i>Cfg_SR_MTU</i> 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.</p> <p><i>IGP_Tunnel_MTU</i> 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.</p> <p>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.</p>
Parameters	<i>bytes</i> — 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>ospf>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 **advertise-tunnel-link** {link | area | as}
no advertise-tunnel-link

Context config>router>ospf
config>router>ospf3

Description This command enables the advertisement of router as defined in the IETF RFC 4970. This adds a new TLV based mechanism, allowing OSPF (OSPFv2 and OSPFv3) router to advertise specific including Traffic Engineering capability, graceful restart helper support and stub router support.
The parameters (**link**, **area**, and **as**) control the scope of the capabilities advertisements.
The no form of this command, disables this capability.

Default **no advertise-tunnel-link**

Parameters **link** — are only advertised over local link and not flooded beyond.
area — are only advertised within the area of origin.
as — are only advertised throughout the entire autonomous system.

super-backbone

Syntax [**no**] **super-backbone**

Context config>service>vprn>ospf

Description This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.
Refer to the OS Services Guide for syntax and command usage information.
The **no** form of the command disables the the super-backbone functionality.

Default **no super-backbone**

timers

Syntax **timers**

Context config>router>ospf
config>router>ospf3

Description This command enables the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

Default none

incremental-spf-wait

Syntax **incremental-spf-wait** *inc-spf-wait*
no incremental-spf-wait

Context config>router>ospf>timers
 config>router>ospf3>timers

Description This command sets the delay before an incremental SPF calculation is performed when LSA types 3, 4, 5, or 7 are received. This allows multiple updates to be processed in the same SPF calculation. Type 1 or type 2 LSAs are considered a topology change and will always trigger a full SPF calculation.

The **no incremental-spf-wait** form of the command resets the timer value back to the default value.

Default 1000ms (1 second)

Parameters *inc-spf-wait* — Specifies the OSPF incremental SPF calculation delay.

Values 0 — 10000

lsa-accumulate

Syntax **lsa-accumulate** *lsa-accum-time*
no lsa-accumulate

Context config>router>ospf>timers
 config>router>ospf3>timers

Description This command sets the internal OSPF delay to allow for the accumulation of multiple LSA so OSPF messages can be sent as efficiently as possible.

Shorting this delay can speed up the advertisement of LSAs to OSPF neighbors but may increase the number of OSPF messages sent.

Default 1000ms (1 second)

Parameters *lsa-accum-time* — Specifies the LSA accumulation delay in milliseconds.

Values 0 — 10000

lsa-arrival

Syntax	lsa-arrival <i>lsa-arrival-time</i> no lsa-arrival
Context	config>router>ospf>timers config>router>ospf3
Description	<p>This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors.</p> <p>It is recommended that the neighbors configured (lsa-generate) <i>lsa-second-wait</i> interval is equal or greater than the lsa-arrival timer configured here.</p> <p>Use the no form of this command to return to the default.</p>
Default	no lsa-arrival
Parameters	<p><i>lsa-arrival-time</i> — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected.</p> <p>Values 0 — 600000</p>

lsa-generate

Syntax	lsa-generate <i>max-lsa-wait</i> [<i>lsa-initial-wait</i> [<i>lsa-second-wait</i>]] no lsa-generate-interval
Context	config>router>ospf>timers config>router>ospf3
Description	<p>This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the <i>lsa-second-wait</i> timer until a maximum value is reached.</p> <p>Configuring the lsa-arrival interval to equal or less than the <i>lsa-second-wait</i> interval configured in the lsa-generate command is recommended.</p> <p>Use the no form of this command to return to the default.</p>
Default	no lsa-generate
Parameters	<p><i>max-lsa-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated.</p> <p>Values 10 — 600,000</p> <p>Default 5,000 milliseconds</p> <p><i>lsa-initial-wait</i> — Specifies the first waiting period between link-state advertisements (LSA) originate(s), in milliseconds. When the LSA exceeds the <i>lsa-initial-wait</i> timer value and the topology changes, there is no wait period and the LSA is immediately generated.</p>

When an LSA is generated, the initial wait period commences. If, within the specified *lsa-initial-wait* period and another topology change occurs, then the *lsa-initial-wait* timer applies.

Values 10 — 600000

Default 5,000 milliseconds

lsa-second-wait — Specifies the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time.). This assumes that each failure occurs within the relevant wait period.

Values 10 — 600000

Default 5,000 milliseconds

redistribute-delay

Syntax **redistribute-delay** *redist-wait*
no redistribute-delay

Context config>router>ospf>timers
config>router>ospf3>timers

Description This command sets the internal OSPF hold down timer for external routes being redistributed into OSPF. Shorting this delay can speed up the advertisement of external routes into OSPF but can result in additional OSPF messages if that source route is not yet stable. The **no redistribute-delay** form of the command resets the timer value back to the default value.

Default 1000ms (1 second)

Parameters *redist-wait* — Specifies the OSPF redistribution hold down timer for external routes being advertised into OSPF.

Values 0 — 1000

spf-wait

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

Context config>router>ospf>timers
config>router>ospf3>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. 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 the **spf-wait** value until

OSPF Global Commands

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*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.

Default **no spf-wait**

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

Values 10 — 120000

Default 10000

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 10 — 100000

Default 1000

traffic-engineering

Syntax **[no] traffic-engineering**

Context config>router>ospf

Description This command enables traffic engineering route calculations constrained by nodes or links.

Traffic engineering enables the router to perform route calculations constrained by nodes or links. The traffic engineering of this router are limited to calculations based on link and nodal constraints.

The **no** form of the command disables traffic engineered route calculations.

Default **no traffic-engineering** — Traffic engineered route calculations is disabled.

unicast-import-disable

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

Context config>router>ospf

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**

OSPF Area Commands

area

Syntax	<code>[no] area area-id</code>
Context	<code>config>router>ospf</code> <code>config>router>ospf3</code>
Description	<p>This command creates the context to configure an OSPF or OSPF3 area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.</p> <p>The no form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, and address-ranges etc., that are currently assigned to this area.</p>
Default	no area — No OSPF areas are defined.
Parameters	<p><i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.</p> <p>Values 0.0.0.0 — 255.255.255.255 (dotted decimal), 0 — 4294967295 (decimal integer)</p>

area-range

Syntax	<code>area-range ip-prefix/mask [advertise not-advertise]</code> <code>no area-range ip-prefix/mask</code>
Context	<code>config>router>ospf>area</code> <code>config>router>ospf>area>nssa</code>
Description	<p>This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.</p> <p>ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.</p> <p>The no form of the command deletes the range (non) advertisement.</p>
Default	no area-range — No range of addresses are defined.
Special Cases	<p>NSSA Context — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.</p> <p>Area Context — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.</p>

OSPF Area Commands

- Parameters** *ip-prefix* — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.
- Values** ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0)
- mask* — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.
- Values** 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal)
- advertise** | **not-advertise** — Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised. The default is **advertise**.

area-range

- Syntax** **area-range** *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]
no area-range *ip-prefix/prefix-length*
- Context** config>router>ospf3>area
config>router>ospf3>area>nssa
- Description** This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.
- ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.
- The **no** form of the command deletes the range (non) advertisement.
- Default** **no area-range** — No range of addresses are defined.
- Special Cases** **NSSA Context** — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.
- Area Context** — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.
- Parameters** *ip-prefix/prefix-length* — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.
- Values** ip-prefix/mask: ip-prefix a.b.c.d (host bits must be 0)
ipv6-prefix: x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D
prefix-length: 0 — 128

advertise | **not-advertise** — Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.
The default is **advertise**.

blackhole-aggregate

Syntax [no] **blackhole-aggregate**

Context config>router>ospf>area
config>router>ospf3>area

Description This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.

The **no** form of this command removes this option.

Default **blackhole-aggregate**

default-metric

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

Context config>router>ospf>area>stub
config>router>ospf3>area

Description This command configures the metric used by the area border router (ABR) for the default route into a stub area.

The default metric should only be configured on an ABR of a stub area.

An ABR generates a default route if the area is a **stub** area.

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

Default **default-metric 1**

Parameters *metric* — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 — 16777215

loopfree-alternate-exclude

Syntax	[no] loopfree-alternate-exclude
Context	config>router>ospf>area config>router>ospf>area>interface config>router>ospf3>area config>router>ospf3>area>interface
Description	<p>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.</p> <p>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.</p> <p>The no form of this command re-instates the default value for this command.</p>
Default	no loopfree-alternate-exclude.

lsa-filter-out

Syntax	lsa-filter-out [all except-own-rtrlsa except-own-rtrlsa-and-defaults] no lsa-filter-out
Context	config>router>ospf>area>interface config>router>ospf3>area>interface config>service>vprn>ospf>area>interface config>service>vprn>ospf3>area>interface
Description	<p>This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:</p> <ul style="list-style-type: none">• Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next-hop.• Flood the router's own router-LSA out the interface and suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and has a statically configured default route with a loopback or system interface address (contained in the router-LSA) as next-hop.• Flood the router's own router-LSA and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route. <p>The no form of this command disables OSPF LSA filtering (normal operation).</p>
Default	no lsa-filter-out

nssa

Syntax `[no] nssa`

Context `config>router>ospf>area`
`config>router>ospf3>area`

Description This command creates the context to configure an OSPF or OSPF3 Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF or OSPF3 domain.

Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of the command removes the NSSA designation and configuration context from the area.

Default `no nssa` — The OSPF or OSPF3 area is not an NSSA.

originate-default-route

Syntax `originate-default-route [type-7] [no-adjacency-check]`
`no originate-default-route`

Context `config>router>ospf>area>nssa`
`config>router>ospf3>area>nssa`

Description This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR) or Autonomous System Border Router (ASBR).

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of the command disables origination of a default route.

Default `no originate-default-route` — A default route is not originated.

Parameters `type-7` — Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default Type 3 LSA for the default route.

no-adjacency-check — Specifies whether or not adjacency checks shall be performed for the NSSA.

redistribute-external

Syntax [no] redistribute-external

Context config>router>ospf>area>nssa
config>router>ospf3>area>nssa

Description This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF or OSPF3 areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF or OSPF3 domain.

The **no** form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default redistribute-external — External routes are redistributed into the NSSA.

stub

Syntax [no] stub

Context config>router>ospf>area
config>router>ospf3>area

Description This command enables access to the context to configure an OSPF or OSPF3 stub area and adds/removes the stub designation from the area.

External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF or OSPF3 area cannot be both an NSSA and a stub area.

Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of the command removes the stub designation and configuration context from the area.

Default no stub — The area is not configured as a stub area.

summaries

Syntax [no] summaries

Context config>router>ospf>area>stub
config>router>ospf3>area>stub
config>router>ospf>area>nssa
config>router>ospf3>area>nssa

Description This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR).

This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or NSSA area. (Default: summary)

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default **summaries** — Summary routes are advertised by the ABR into the stub area or NSSA.

Interface/Virtual Link Commands

advertise-subnet

Syntax [no] advertise-subnet

Context config>router>ospf>area>interface *ip-int-name*

Description This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

The **no** form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default **advertise-subnet** — Advertises point-to-point interfaces as subnet routes.

authentication

Syntax **authentication** [inbound *sa-name* outbound *sa-name*]
authentication bidirectional *sa-name*
no authentication

Context config>router>ospf3>area>interface *ip-int-name*
config>router>ospf3>area>virtual-link >*if*

Description This command configures the password used by the OSPF3 interface or virtual-link to send and receive OSPF3 protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication.

By default, no authentication key is configured.

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

Default **no authentication** — No authentication is defined.

Parameters **inbound** *sa-name* — Specifies the inbound sa-name for OSPF3 authentication.

outbound *sa-name* — Specifies the outbound sa-name for OSPF3 authentication.

bidirectional *sa-name* — Specifies bidirectional OSPF3 authentication.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>ospf>area>interface <i>ip-int-name</i> config>router>ospf>area>virtual-link > <i>if</i> >
Description	<p>This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.</p> <p>All neighboring routers must use the same type of authentication and password for proper protocol communication. If the authentication-type is configured as password, then this key must be configured.</p> <p>By default, no authentication key is configured.</p> <p>The no form of the command removes the authentication key.</p>
Default	no authentication-key — No authentication key is defined.
Parameters	<p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p><i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p>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.</p> <p>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.</p>

authentication-type

Syntax	authentication-type { password message-digest } no authentication-type
Context	config>router>ospf>area>interface <i>ip-int-name</i> config>router>ospf>area>virtual-link <i>router-id</i>
Description	<p>This command enables authentication and specifies the type of authentication to be used on the OSPF interface.</p> <p>Both simple password and message-digest authentication are supported.</p> <p>By default, authentication is not enabled on an interface.</p> <p>The no form of the command disables authentication on the interface.</p>

Interface/Virtual Link Commands

Default **no authentication** — No authentication is enabled on an interface.

Parameters **password** — This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest — This keyword enables message digest MD5 authentication in accordance with RFC1321. If this option is configured, then at least one message-digest-key must be configured.

auth-keychain

Syntax **auth-keychain**

Context
config>router>ospf>areas>interface
config>router>ospf>areas>virtual-link
config>service>vprn>ospf>areas>interface
config>service>vprn>ospf>areas>sham-link
config>service>vprn>ospf>areas>virtual-link

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.

bfd-enable

Syntax [**no**] **bfd-enable** [**remain-down-on-failure**]

Context
config>router>ospf>area>interface
config>router>ospf3>area>interface

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given 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 via the BFD command under the IP interface.

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

Default **no bfd-enable**

Parameters **remain-down-on-failure** — Forces adjacency down on BFD failure.

dead-interval

Syntax	dead-interval <i>seconds</i> no dead-interval
Context	config>router>ospf>area>interface config>router>ospf3>area>interface config>router>ospf>area>virtual-link config>router>ospf3>area>virtual-link
Description	This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval. The no form of the command reverts to the default value.
Default	40 seconds
Special Cases	OSPF Interface — If the dead-interval configured applies to an interface, then all nodes on the subnet must have the same dead interval. Virtual Link — If the dead-interval configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.
Parameters	<i>seconds</i> — The dead interval expressed in seconds. Values 1 — 65535

hello-interval

Syntax	hello-interval <i>seconds</i> no hello-interval
Context	config>router>ospf>area>interface config>router>ospf3>area>interface config>router>ospf>area>virtual-link config>router>ospf3>area>virtual-link
Description	This command configures the interval between OSPF hellos issued on the interface or virtual link. The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent. Reducing the interval, in combination with an appropriate reduction in the associated dead-interval , allows for faster detection of link and/or router failures at the cost of higher processing costs. The no form of this command reverts to the default value.
Default	hello-interval 10 — A 10-second hello interval.
Special Cases	OSPF Interface — If the hello-interval configured applies to an interface, then all nodes on the subnet must have the same hello interval.

Interface/Virtual Link Commands

Virtual Link — If the **hello-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.

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

Values 1 — 65535

interface

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

Context config>router>ospf>area
config>router>ospf3>area

Description This command creates a context to configure an OSPF interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.

The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the config>router>ospf>interface context can be used to disable an interface without removing the configuration for the interface.

Default **no interface** — No OSPF interfaces are defined.

Parameters *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

secondary — Allows multiple secondary adjacencies to be established over a single IP interface.

interface-type

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

Context config>router>ospf>area>interface
config>router>ospf3>area>interface

Description This command configures the interface type to be 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 broadcast adjacency maintenance overhead of the Ethernet link provided the link is used as a point-to-point.

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

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

Default **point-to-point** if the physical interface is SONET.
broadcast if the physical interface is Ethernet or unknown.

Special Cases **Virtual-Link** — A virtual link is always regarded as a point-to-point interface and not configurable.

Parameters **broadcast** — Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.
point-to-point — Configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax **message-digest-key** *keyid* **md5** [*key* | *hash-key*] [**hash**]
no message-digest-key *keyid*

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.

The **no** form of the command removes the message digest key identified by the *key-id*.

Default No message digest keys are defined.

Parameters **keyid** — The *keyid* is expressed as a decimal integer.

Values 1 — 255

md5 *key* — The MD5 key. The *key* can be any alphanumeric string up to 16 characters in length.

md5 *hash-key* — The MD5 hash key. The key can be any combination of ASCII characters up to 32 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.

Interface/Virtual Link Commands

metric

Syntax	metric <i>metric</i> no metric
Context	config>router>ospf>area>interface config>router>ospf3>area>interface
Description	<p>This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.</p> <p>The no form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the reference-bandwidth command setting and the speed of the underlying link.</p>
Default	no metric — The metric is based on reference-bandwidth setting and the link speed.
Parameters	<i>metric</i> — The metric to be applied to the interface expressed as a decimal integer.
Values	1 — 65535

mtu

Syntax	mtu <i>bytes</i> no mtu
Context	config>router>ospf>area>interface config>router>ospf3>area>interface
Description	<p>This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:</p> <pre>config>port>ethernet config>port>sonet-sdh>path config>port>tdm>t3-e3 config>port>tdm>t1-e1>channel-group</pre> <p>If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.</p> <p>To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.</p> <p>Use the no form of this command to revert to default.</p>
Default	no mtu — Uses the value derived from the MTU configured in the config>port context.
Parameters	<i>bytes</i> — The MTU to be used by OSPF for this logical interface in bytes.
Values	512 — 9198 (9212 — 14) (Depends on the physical media)

node-sid

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

Context config>router>ospf>area>interface

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.

Default none

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

Values 0 — 4294967295

Default none

label *value* — integer.

Values 0 — 4294967295

Default none

passive

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

Context config>router>ospf>area>interface
 config>router>ospf3>area>interface

Description This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

The **no** form of the command removes the passive property from the OSPF interface.

Interface/Virtual Link Commands

Default Service interfaces defined in **config>router>service-prefix** are passive.
All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>router>ospf>area>interface
config>router>ospf3>area>interface

Description This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

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

Default **priority 1**

Parameters *number* — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router or Backup Designated Router on the interface subnet.

Values 0 — 255

retransmit-interval

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

Context config>router>ospf>area>interface
config>router>ospf3>area>interface
config>router>ospf>area>virtual-link
config>router>ospf3>area>virtual-link

Description This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgement has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

Default **retransmit-interval 5**

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

Values 1 — 1800

rib-priority

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

Context config>router>ospf>area>interface
 config>router>ospf3>area>interface

Description This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface will be processed through the OSPF route calculation process at a higher priority.

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

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.

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>router>ospf>area>interface
 config>router>ospf3>area>interface
 config>router>ospf>area>virtual-link
 config>router>ospf3>area>virtual-link

Description This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link.

The **no** form of this command reverts to the default delay time

Default **transit-delay 1**

Parameters *seconds* — The transit delay in seconds expressed as a decimal integer.

Values 1 — 1800

virtual-link

Syntax [**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context config>router>ospf>area
 config>router>ospf3>area

Description This command configures a virtual link to connect area border routers to the backbone via a virtual link. The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in the picture below) then the

Interface/Virtual Link Commands

area border routers (routers 1 and 2 in the picture below) must be connected via a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area. (area 0.0.0.1 in the picture below). A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of the command deletes the virtual link. (*Default: none specified*)

Default No virtual link is defined.

Parameters *router-id* — The router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area area-id — The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in Figure 13) then the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point like adjacency across the transit area (see Area 0.0.0.4).




Figure 13: OSPF Areas