

Chapter 15

Summary of OSPF Configuration Statements

The following sections explain each of the OSPF configuration statements, which are organized alphabetically. The term OSPF refers to both OSPF (OSPF version 2) and OSPFv3 (OSPF version 3).

area

Syntax	<code>area <i>area-id</i>;</code>
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	<p>Specify the area identifier for this router to use when participating in OSPF routing. All routers in an area must use the same area identifier to establish adjacencies.</p> <p>Specify multiple area statements to configure the router as an area border router. An area border router automatically summarizes routes between areas; use the <code>area-range</code> statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, use the <code>virtual-link</code> statement.</p> <p>To specify that the router is directly connected to the OSPF and OSPFv3 backbone, include the <code>area 0.0.0.0</code> statement.</p> <p>All routers on the backbone must be contiguous. If they are not, use the <code>virtual-link</code> statement to create the appearance of connectivity to the backbone.</p>
Options	<code><i>area-id</i></code> —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number 0.0.0.0 is reserved for the OSPF and OSPFv3 backbone area.
Usage Guidelines	See “Configure the Backbone Area and Other Areas” on page 245.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
See Also	<code>area-range</code> on page 264, <code>virtual-link</code> on page 288

area-range

Syntax	<code>area-range network/mask-length <restrict>;</code>
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i>], [edit protocols (ospf ospf3) area <i>area-id</i> nssa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa]
Description	(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple area-range statements. For an NSSA, summarize a range of IP addresses when sending NSSA LSAs. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other area. To specify multiple prefixes, include multiple area-range statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.
Default	By default, area border routers do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.
Options	<i>network</i> —IP address. You can specify one or more IP addresses. <i>mask-length</i> —Number of significant bits in the network mask. <i>restrict</i> —(Optional) Do not advertise the configured summary. This has the effect of hiding all routes that are contained within the summary, effectively creating a route filter.
Usage Guidelines	See “Configure Route Summarization” on page 252 and “Configure a Not-So-Stubby Area” on page 246.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.

authentication-key

Syntax	authentication-key <i>key</i> key-id <i>identifier</i> ;
Hierarchy Level	[edit protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> virtual-link], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link]
Description	Configure an authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must specify an authentication scheme using the authentication-type statement at either the [edit protocols ospf area <i>area-id</i>] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>] hierarchy level. All routers that are connected to the same IP subnet must use the same authentication scheme and password.
Options	<i>key</i> —Authentication password. It can be 1 to 8 characters long. Characters can include any ASCII strings. If you include spaces, enclose the string in quotation marks (double quotes). <i>key-id identifier</i> —(For MD5 authentication only) Key identifier associated with the MD5 key.
Usage Guidelines	See “Configure Authentication” on page 250.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	authentication-type on page 266

authentication-type

Syntax	authentication-type <i>authentication</i> ;
Hierarchy Level	[edit protocols ospf area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>]
Description	<p>Enable authentication and specify the authentication scheme for the backbone or area. If you enable authentication, you must specify a password by including the authentication-key statement within [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i>] hierarchy level.</p> <p>All routers connected to the same IP subnet must use the same authentication scheme and password.</p>
Options	<p><i>authentication</i>—Authentication scheme:</p> <p> none—Disable authentication.</p> <p> simple—Use a simple password. The password is included in the transmitted packet, making this method of authentication relatively insecure.</p> <p> md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving router uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</p> <p>Default: none (No authentication is performed.)</p>
Usage Guidelines	See “Configure Authentication” on page 250.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
See Also	authentication-key on page 265

dead-interval

Syntax	dead-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link]
Description	Specify how long OSPF waits before declaring that a neighboring router is unavailable. This is an interval during which the router receives no hello packets from the neighbor.
Options	<i>seconds</i> —Interval to wait. Range: 1 through 65,535 Default: 40 seconds (four times the hello interval)
Usage Guidelines	See “Modify the Router Dead Interval” on page 254.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	hello-interval on page 271

default-lsa

Syntax	default-lsa { default-metric <i>metric</i> ; metric-type <i>type</i> ; type-7; }
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> nssa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa]
Description	On area border routers only, for an NSSA, inject a default LSA with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.
Options	The statements are explained separately.
Usage Guidelines	See “Configure a Not-So-Stubby Area” on page 246.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	nssa on page 275, stub on page 283

default-metric

Syntax	default-metric <i>metric</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> stub], [edit protocols (ospf ospf3) area <i>area-id</i> nssa default-isa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> stub], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-isa]
Description	On area border routers only, for a stub area, inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.
Options	<i>metric</i> —Metric value. Range: 1 through 16,777,215
Usage Guidelines	See “Configure a Stub Area” on page 246.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	nssa on page 275, stub on page 283

disable

Syntax	disable;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) virtual-link], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link]
Description	Disable OSPF, an OSPF interface, or an OSPF virtual link.
Default	The configured object is enabled (operational) unless explicitly disabled.
Usage Guidelines	See “Minimum OSPF Configuration” on page 244.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

domain-id

Syntax	domain-id <i>domain-id</i> ;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Specify a domain ID for a route. The domain ID identifies the OSPFv2 domain from which the route originated.
Options	<i>domain-id</i> —IP address. Default: If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.
Usage Guidelines	See “Configure an OSPF Domain ID” on page 165.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

domain-vpn-tag

Syntax	domain-vpn-tag <i>number</i> ;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider-edge (PE) router.
Options	<i>number</i> —VPN tag
Usage Guidelines	See “Configure an OSPF Domain ID” on page 165.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

export

Statement	export [<i>policy-names</i>];
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Apply one or more policies to routes being exported from the routing table into OSPF.
Options	<i>policy-names</i> —Name of one or more policies.
Usage Guidelines	See “Configure OSPF Routing Policy” on page 259 and the <i>JUNOS Internet Software Configuration Guide: Policy Framework</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

external-preference

Syntax	external-preference <i>preference</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Set the route preference for OSPF external routes.
Options	<i>preference</i> —Preference value. Range: 0 through 255 Default: 150
Usage Guidelines	See “Configure Route Preferences” on page 253.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	preference on page 279

graceful-restart

Syntax	graceful-restart { disable; helper-disable; notify-duration <i>seconds</i> ; restart-duration <i>seconds</i> ; }
Hierarchy Level	[edit protocols (ospf ospf3)]
Description	Configure graceful restart for OSPF.
Options	disable—Disable graceful restart for OSPF. notify-duration <i>seconds</i> —Estimated time to send out purged grace LSAs over all the interfaces, in seconds. Range: 1 through 3,600 Default: 30 restart-duration <i>seconds</i> —Estimated time to reacquire a full OSPF neighbor from each area, in seconds. Range: 1 through 3,600 Default: 60 helper-disable—Disables graceful restart helper capability.
Usage Guidelines	See “Configure Graceful Restart” on page 255 and “Configure Graceful Restart” on page 87.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hello-interval

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link]
Description	Specify how often the router sends hello packets out the interface. The hello interval must be the same for all routers on a shared logical IP network.
Options	<i>seconds</i> —Time between hello packets. Range: 1 through 255 Default: 10 seconds; 120 seconds (nonbroadcast networks)
Usage Guidelines	See “Modify the Hello Interval” on page 253.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	dead-interval on page 267

interface

Syntax	interface <i>interface-name</i> { disable; authentication-key <i>key</i> < key-id <i>identifier</i> > ; dead-interval <i>seconds</i> ; hello-interval <i>seconds</i> ; interface-type <i>type</i> ; metric <i>metric</i> ; neighbor <i>address</i> < eligible> ; passive; poll-interval <i>seconds</i> ; priority <i>number</i> ; retransmit-interval <i>seconds</i> ; transit-delay <i>seconds</i> ; transmit-interval <i>seconds</i> ; }
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>]
Description	Enable OSPF routing on a router interface. You must include at least one interface statement in the configuration to enable OSPF on the router.

Options *interface-name*—Name of the interface. To configure all interfaces, you can specify all. For details about specifying interfaces, see interface naming in the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* .



Note

For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

The remaining statements are explained separately.

Usage Guidelines See “Minimum OSPF Configuration” on page 244 and “Configure an Interface on a Broadcast or Point-to-Point Network” on page 249.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

See Also neighbor on page 275

interface-type

Syntax interface-type (nbma | p2mp);

Hierarchy Level [edit protocols ospf area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id*
interface *interface-name*]

Description Specify the type of interface.

By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in NBMA mode, configure the nbma interface type, using the IP address of the local ATM interface.

Default The software chooses the correct interface type based on the type of physical interface.

Options nbma—Nonbroadcast multiaccess (NBMA) interface.

p2mp—Point-to-multipoint interface.

Usage Guidelines See “Configure an Interface on a Broadcast or Point-to-Point Network” on page 249.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

label-switched-path

Syntax	label-switched-path <i>name</i> metric <i>metric</i> ;
Hierarchy Level	[edit protocols ospf area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>]
Description	Advertise label-switched paths into OSPF as point-to-point links. The label-switched path is advertised in the appropriate OSPF levels as a point-to-point link and contains a local address and a remote address.
Options	<i>name</i> —Name of the label-switched path. <i>metric</i> —Metric value. Range: 1 through 65535 Default: 1
Usage Guidelines	See “Advertise Label-Switched Paths into OSPF” on page 256.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

lsp-metric-into-summary

Syntax	lsp-metric-into-summary;
Hierarchy Level	[edit protocols ospf traffic-engineering shortcuts], [edit routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts]
Description	Advertise the LSP metric in summary LSAs.
Usage Guidelines	See “Enable OSPF Traffic Engineering Support” on page 258.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

metric

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]
Description	Cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation. To set the cost of routes exported into OSPF, configure the appropriate routing policy.
Options	<i>metric</i> —Cost of the route. Range: 1 through 65,535 Default: 1
Usage Guidelines	See “Modify the Interface Metric” on page 252.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

metric-type

Syntax	<code>metric-type <i>type</i>;</code>
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa]
Description	External metric type for the default LSA.
Options	<i>type</i> —Metric type. Range: 1 or 2
Usage Guidelines	See “Configure a Not-So-Stubby Area” on page 246.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

neighbor

Syntax	<code>neighbor address <eligible>;</code>
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]
Description	For nonbroadcast interfaces only, specify neighboring routers. On a nonbroadcast interface, you must specify neighbors explicitly because OSPF does not send broadcast packets to dynamically discover their neighbors. To specify multiple neighbors, include multiple neighbor statements.
Options	<i>address</i> —IP address of a neighboring router. <i>eligible</i> —(Optional) Allow the neighbor to become a designated router. Default: If you omit this option, the neighbor is not considered eligible to become a designated router.
Usage Guidelines	See “Configure an Interface on a Nonbroadcast, Multiaccess Network” on page 249.
Required Privilege Level	<i>routing</i> —To view this statement in the configuration. <i>routing-control</i> —To add this statement to the configuration.

no-summaries

See summaries on page 283

nssa

Syntax	<code>nssa { area-range network/mask-length <restrict>; default-lsa { default-metric metric; metric-type type; type-7; } (no-summaries summaries) }</code>
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>]
Description	Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas. You cannot configure an area as being both a stub area and an NSSA.
Options	The statements are explained separately in this chapter.
Usage Guidelines	See “Configure a Not-So-Stubby Area” on page 246.
Required Privilege Level	<i>routing</i> —To view this statement in the configuration. <i>routing-control</i> —To add this statement to the configuration.

See Also stub on page 283

ospf

Syntax	ospf { ... }
Hierarchy Level	[edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Description	Enable OSPF routing on the router. You must include the ospf statement to enable OSPF on the router.
Default	OSPF is disabled on the router.
Usage Guidelines	See “Minimum OSPF Configuration” on page 244.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ospf3

Syntax	ospf3 { ... }
Hierarchy Level	[edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Description	Enable OSPFv3 routing on the router. You must include the ospf statement to enable OSPFv3.
Default	OSPFv3 is disabled.
Usage Guidelines	See “Minimum OSPF Configuration” on page 244.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

overload

Syntax	overload { <timeout <i>seconds</i> >; }
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Configure the local router so that it appears to be overloaded. You might do this when you want the router to participate in OSPF routing, but do not want it to be used for transit traffic. Note that traffic destined to immediately attached interfaces continues to reach the router.
Option	timeout <i>seconds</i> —(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the router will remain in overload state until the overload statement is deleted or a timeout is set. Range: 60 through 1800 seconds Default: 0 seconds
Usage Guidelines	See “Configure the Router to Appear Overloaded” on page 257.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

peer-interface

Syntax	peer-interface <i>interface-name</i> { disable; dead-interval <i>seconds</i> ; hello-interval <i>seconds</i> ; retransmit-interval <i>seconds</i> ; transit-delay <i>seconds</i> ; }
Hierarchy Level	[edit protocols ospf area <i>area-id</i>]
Description	Configure a peer interface.
Options	<i>interface-name</i> —Name of the peer interface. To configure all interfaces, you can specify all. For details about specifying interfaces, see the <i>JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service</i> . The remaining statements are explained separately.
Usage Guidelines	See “Configure a Peer Interface” on page 259.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

passive

Syntax	passive;
Hierarchy Level	[edit protocols (ospf ospf3) interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) interface <i>interface-name</i>]
Description	Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run. Enabling OSPF on an interface (by including the interface statement at the [edit protocols (ospf ospf3)] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf] hierarchy level), disabling it (by including the disable statement), and not actually having OSPF run on an interface (by including the passive statement) are mutually exclusive states.
Usage Guidelines	See “Advertise Interface Addresses without Running OSPF” on page 256.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

poll-interval

Syntax	poll-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]
Description	For nonbroadcast interfaces only, specify how often the router sends hello packets out of the interface before it establishes adjacency with a neighbor.
Options	<i>seconds</i> —Frequency at which to send hello packets. Range: 1 through 255 Default: 120 seconds
Usage Guidelines	See “Configure an Interface on a Nonbroadcast, Multiaccess Network” on page 249.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

preference

Syntax	preference <i>preference</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Set the route preference for OSPF internal routes.
Options	<i>preference</i> —Preference value. Range: 0 through 255 Default: 10
Usage Guidelines	See “Configure Route Preferences” on page 253.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	external-preference on page 270

prefix-export-limit

Syntax	prefix-export-limit <i>number</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Configure a limit to the number of prefixes exported into OSPF.
Options	<i>number</i> —Prefix limit. Range: 0 through 4294967295 Default: None
Usage Guidelines	See “Configure a Prefix Export Limit” on page 251 .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

priority

Syntax	priority <i>number</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]
Description	Specify the router's priority for becoming the designated router. The router that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one router on each logical IP network or subnet to be the designated router. You also should specify a router's priority for becoming the designated router on point-to-point interfaces.
Options	<i>number</i> —Router's priority for becoming the designated router. A priority value of 0 means that the router never will become the designated router. A value of 1 means that the router has the least chance of becoming a designated router. Range: 0 through 255 Default: 128
Usage Guidelines	See "Designated Router" on page 239 and "Configure the Priority for Becoming the Designated Router" on page 251.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

reference-bandwidth

Syntax	reference-bandwidth <i>reference-bandwidth</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula: $cost = ref\text{-}bandwidth / bandwidth$
Options	<i>ref-bandwidth</i> —Reference bandwidth in bits per second. Default: 100 Mbps
Usage Guidelines	See "Modify the Interface Metric" on page 252.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link]
Description	Specify how long the router waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors.
Options	<i>seconds</i> —Interval to wait. Range: 1 through 65,535 Default: 5 seconds
Usage Guidelines	See “Control the LSA Retransmission Interval” on page 254.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rib-group

Syntax	rib-group <i>group-name</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.
Usage Guidelines	See “Create Routing Table Groups” on page 79, “Configure How Interface Routes Are Imported into Routing Tables” on page 81, and “Configure BGP Routing Table Groups” on page 412.
Options	<i>group-name</i> —Name of the routing table group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	interface-routes on page 108, rib-group on page 120

route-type-community

Syntax	route-type-community (iana vendor);
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either IANA or vendor-specific route type.
Options	iana—Encode route type with the value 0x0306. This is the default value. vendor—Encode route type with the value 0x8000.
Usage Guidelines	See “Configure an OSPF Domain ID” on page 165.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

shortcuts

Syntax	shortcuts;
Hierarchy Level	[edit protocols ospf traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering]
Description	Configure OSPFv2 to use MPLS label-switched paths (LSPs) as next hops if possible when installing routing information into the inet.3 routing table.
Usage Guidelines	See “Enable OSPF Traffic Engineering Support” on page 258.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

spf-delay

Syntax	spf-delay <i>milliseconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	Configure the shortest path first (SPF) delay.
Options	<i>milliseconds</i> —Number of milliseconds between the detection of a topology change and when the SPF algorithm runs. Range: 50 through 1000 Default: 1000 milliseconds
Usage Guidelines	See “Configure the SPF Delay” on page 255.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

stub

Syntax	stub <default-metric <i>metric</i> > <(no-summaries summaries)>;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i>] [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>]
Description	Indicate that this area should not be flooded with AS external link-state advertisements. You must include the stub statement when configuring all routers that are in the stub area. The backbone cannot be configured as a stub area. You cannot configure an area as being both a stub area and an NSSA. Options default-metric <i>metric</i> —(Optional) Inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area. Range: 1 through 16,777,215 no-summaries—(Optional) Do not advertise routes into the stub area. If you include the default-metric option, only the default route is advertised. summaries—(Optional) Flood summary LSAs into the stub area. Usage Guidelines See “Configure a Stub Area” on page 246. Required Privilege Level routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. See Also nssa on page 275

summaries

Syntax	(summaries no-summaries);
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> nssa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa]
Description	Configure whether area border routers advertise summary routes into an NSSA: summaries—Flood summary LSAs into the NSSA. no-summaries—Prevent area border routers from advertising summaries into an NSSA. If default-metric is configured for an NSSA, a Type 3 LSA is injected into the area by default. Usage Guidelines See “Configure a Not-So-Stubby Area” on page 246. Required Privilege Level routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. See Also nssa on page 275, stub on page 283

traceoptions

Syntax	<pre>traceoptions { file <i>name</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Description	<p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	The default OSPF protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>. We recommend that you place OSPF tracing output in the file <code>ospf-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files Default: 2 files</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p>

OSPF and OSPFv3 Tracing Flags

database-description—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.

error—OSPF and OSPFv3 error packets.

event—OSPF and OSPFv3 state transitions.

flooding—Link-state flooding packets.

hello—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.

lsa-ack—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.

lsa-request—Link-state request packets, which are used in synchronizing the OSPF topological database.

lsa-update—Link-state updates packets, which are used in synchronizing the OSPF topological database.

packets—All OSPF packets.

packet-dump—Dump the contents of selected packet types.

spf—Shortest-path-first (SPF) calculations.

Global Tracing Flags

all—All tracing operations.

general—A combination of the normal and route trace operations.

normal—All normal operations.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

policy—Policy operations and actions.

route—Routing table changes.

state—State transitions.

task—Interface transactions and processing.

timer—Timer usage.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

detail—Provide detailed trace information.

receive—Packets being received.

send—Packets being transmitted.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Disallow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the *files* option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB
Range: 10 KB through the maximum file size supported on your system
Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See “Trace OSPF Protocol Traffic” on page 260.

Required Privilege Level routing and trace—To view this statement in the configuration.
 routing-control and trace-control—To add this statement to the configuration.

traffic-engineering

Syntax traffic-engineering {
 <no-topology>;
 <shortcuts> {
 <lsp-metric-into-summary>;
 }
 }

Hierarchy Level [edit protocols ospf],
 [edit routing-instances *routing-instance-name* protocols ospf]

Description Enable OSPF's traffic engineering features.

Default Traffic engineering support is disabled.

Options lsp-metric-into-summary—(Optional) Advertise LSP metric in Summary LSAs.
 no-topology—(Optional) Disable the dissemination of the link state topology information.
 shortcuts—(Optional) Use label-switched paths as next hops, if label-switched paths are configured.

Usage Guidelines See “Enable OSPF Traffic Engineering Support” on page 258.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

transit-delay

Syntax	transit-delay <i>seconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link]
Description	Set the estimated time required to transmit a link-state update on the interface. When calculating this time, you should account for transmission and propagation delays. You should never have to modify the transit delay time.
Options	<i>seconds</i> —Estimated time. Range: 1 through 65,535 Default: 1 second
Usage Guidelines	See “Specify the Transit Delay” on page 255.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

transmit-interval

Syntax	transmit-interval <i>milliseconds</i> ;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>]
Description	Set the interval at which OSPF packets are transmitted on an interface.
Options	<i>milliseconds</i> —Transmission interval Range: 1 through 4,294,967 Default: 30 milliseconds
Usage Guidelines	See “Control the LSA Retransmission Interval” on page 254.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

type-7

Syntax	type-7;
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa]
Description	Flood Type 7 default LSA if the no-summaries statement is configured. By default, when the no-summaries statement is configured, a Type 3 LSA is injected into NSSA areas for JUNOS Release 5.0 and later. To support backward compatibility with previous JUNOS releases, include the type-7 statement. This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the no-summaries statement in the configuration.
Usage Guidelines	See “Configure a Not-So-Stubby Area” on page 246.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

virtual-link

Syntax	virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> { disable; authentication-key <i>key</i> < <i>key-id identifier</i> >; dead-interval <i>seconds</i> ; hello-interval <i>seconds</i> ; retransmit-interval <i>seconds</i> ; transit-delay <i>seconds</i> ; }
Hierarchy Level	[edit protocols (ospf ospf3) area <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>]
Description	For backbones only, create a virtual link to use in place of an actual physical link. All area border routers and other routers on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routers that form the end points of the link, and both these two routers must be area border routers. You cannot configure links through stub areas.
Options	neighbor-id <i>router-id</i> —IP address of the router at the remote end of the virtual link. transit-area <i>area-id</i> —Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area. The remaining statements are explained separately.
Usage Guidelines	See “Configure an RSVP LSP” on page 247.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.