

Chapter 20

Summary of OSPF Configuration Statements

The following sections explain each of the OSPF configuration statements. The statements are organized alphabetically.

area

Syntax	<code>area <i>area-id</i>;</code>
Hierarchy Level	[edit protocols ospf]
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 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 backbone area.
Usage Guidelines	See “Configure the Backbone Area and Other Areas” on page 223.
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 238, <code>virtual-link</code> on page 255

area-range

Syntax `area-range network/mask-length <restrict>;`

Hierarchy Level [edit protocols ospf area *area-id*],
[edit protocols ospf area *area-id* 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 *network*—IP address. You can specify one or more IP addresses.

mask-length—Number of significant bits in the network mask.

restrict—(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 228 and “Configure a Not-So-Stubby Area” on page 224.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—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]
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 the [edit 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-id identifier</i> —(For MD5 authentication only) Key identifier associated with the MD5 key. You can define multiple keys, each with a different key identifier. OSPF uses the key with the highest number. <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).
Usage Guidelines	See “Configure Authentication” on page 227.
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 240

authentication-type

Syntax authentication-type *authentication*;

Hierarchy Level [edit protocols ospf area *area-id*]

Description 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 the interface statement.

All routers connected to the same IP subnet must use the same authentication scheme and password.

Options *authentication*—Authentication scheme:

none—Disable authentication.

simple—Use a simple password. The password is included in the transmitted packet, making this method of authentication relatively insecure.

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.

Default: none (No authentication is performed.)

Usage Guidelines See “Configure Authentication” on page 227.

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

See Also authentication-key on page 239

dead-interval

Syntax dead-interval *seconds*;

Hierarchy Level [edit protocols ospf area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* 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 *seconds*—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 231.

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 242

default-metric

Syntax	default-metric <i>metric</i> ;
Hierarchy Level	[edit protocols ospf area <i>area-id</i> stub], [edit protocols ospf area <i>area-id</i> nssa]
Description	On area border routers only, for a stub area or an NSSA, 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 223 and “Configure a Not-So-Stubby Area” on page 224.
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 245, stub on page 250

disable

Syntax	disable;
Hierarchy Level	[edit protocols ospf], [edit protocols ospf area <i>area-id</i> interface], [edit protocols ospf virtual-link]
Description	Disable OSPF, an OSPF interface, or an OSPF virtual link.
Default	The configured object is enabled (operational) unless explicitly disabled.
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]
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 233 and “Configure Routing Policy” on page 35.
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 *preference*;

Hierarchy Level [edit protocols ospf]

Description Set the route preference for OSPF external routes.

Options *preference*—Preference value.

Range: 0 through 255

Default: 150

Usage Guidelines See “Configure Route Preferences” on page 229.

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 247

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit protocols ospf area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* 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 *seconds*—Length of time between hello packets.

Range: 1 through 255

Default: 10 seconds

Usage Guidelines See “Modify the Hello Interval” on page 230.

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 240

interface

Syntax	<pre>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>; }</pre>
Hierarchy Level	[edit protocols ospf area <i>area-id</i>]
Description	<p>Enable OSPF routing on a router interface.</p> <p>You must include at least one interface statement in the configuration to enable OSPF on the router.</p>
Options	<p><i>interface-name</i>—Name of the interface. To configure all interfaces, you can specify all. For details about specifying interfaces, see interface naming in the <i>JUNOS Internet Software Configuration Guide: Interfaces and Chassis</i>.</p> <p>For nonbroadcast interfaces, specify the address of the neighboring routers with the neighbor statement.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Minimum OSPF Configuration” on page 222 and “Configure OSPF on Router Interfaces” on page 225.
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	neighbor on page 245

interface-type

Syntax	interface-type (nbma p2mp);
Hierarchy Level	[edit protocols ospf area <i>area-id</i> interface <i>interface-name</i>]
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.
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 OSPF on Router Interfaces” on page 225.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

metric

Syntax	metric <i>metric</i> ;
Hierarchy Level	[edit protocols ospf 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 OSPF 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 228.
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 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 its neighbors. To specify multiple neighbors, include multiple neighbor statements.
Options	<p><i>address</i>—IP address of a neighboring router.</p> <p><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.</p>
Usage Guidelines	See “Configure an Interface on a Nonbroadcast, Multiaccess Network” on page 226.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

no-summaries

See Also summaries on page 250

nssa

Syntax	<pre>nssa { area-range network/mask-length <restrict>; default-metric metric; (no-summaries summaries) }</pre>
Hierarchy Level	[edit protocols ospf area <i>area-id</i>]
Description	<p>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.</p> <p>You cannot configure an area as being both a stub area and an NSSA.</p>
Options	The statements are explained separately.
Usage Guidelines	See “Configure a Not-So-Stubby Area” on page 224.
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	stub on page 250

ospf

Syntax	ospf { ... }
Hierarchy Level	[edit 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 222.
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 interface <i>interface-name</i>]
Description	Advertise the direct interface addresses into 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 one on which the protocol does not run. Enabling OSPF on an interface (by including the interface statement at the [edit 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 232.
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 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 226.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

preference

Syntax	<code>preference preference;</code>
Hierarchy Level	[edit protocols ospf]
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 229.
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 242

priority

Syntax	<code>priority number;</code>
Hierarchy Level	[edit protocols ospf 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 220 and “Configure the Priority for Becoming the Designated Router” on page 228.
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>ref-bandwidth</i> ;
Hierarchy Level	[edit protocols ospf]
Description	Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula: $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$
Options	<i>ref-bandwidth</i> —Reference bandwidth in bits per second. Default: 100 Mbps
Usage Guidelines	See “Modify the Interface Metric” on page 228.
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 area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf 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 231.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rib-group

Syntax	<code>rib-group <i>routing-table-group-name</i>;</code>
Hierarchy Level	[edit protocols ospf]
Description	Install routes learned from OSPF instances into routing tables in the OSPF routing table group.
Usage Guidelines	See “Configure BGP Routing Table Groups” on page 335, “Configure How Interface Routes Are Imported into Routing Tables” on page 124, and “Create Routing Table Groups” on page 123.
Options	<i>routing-table-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 141, rib-group on page 150

shortcuts

Syntax	<code>shortcuts;</code>
Hierarchy Level	[edit protocols ospf traffic-engineering]
Description	Configure OSPF 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 “Configure Route Preferences” on page 181.
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 *metric*> <(no-summaries | summaries)>;

Hierarchy Level [edit protocols ospf area *area-id*]

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** *metric*—(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

Default: 1

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 223.

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 245

summaries

Syntax (summaries | no-summaries);

Hierarchy Level [edit protocols ospf area *area-id* 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.

Usage Guidelines See “Configure a Not-So-Stubby Area” on page 224.

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 245, stub on page 250

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]
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 is to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—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> <p>OSPF Tracing Flags</p> <p>database-description—Database description packets, which are used in synchronizing the OSPF topological database.</p> <p>error—OSPF errored packets.</p> <p>event—OSPF state transitions.</p> <p>flooding—Link-state flooding packets.</p> <p>hello—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.</p> <p>lsa-ack—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.</p> <p>lsa-request—Link-state request packets, which are used in synchronizing the OSPF topological database.</p>

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 233.

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;
}

Hierarchy Level [edit protocols ospf]

Description Enable OSPF's traffic engineering features.

Default Traffic engineering support is disabled.

Options no-topology—(Optional) Disable the traffic engineering link state topology dissemination.
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 232.

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 *seconds*;

Hierarchy Level [edit protocols ospf area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* 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 *seconds*—Estimated time.
Range: 1 through 65,535
Default: 1 second

Usage Guidelines See “Specify the Transit Delay” on page 231.

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 *milliseconds*;

Hierarchy Level [edit protocols ospf area *area-id* interface *interface-name*]

Description Set the interval at which OSPF packets are transmitted on an interface.

Options *milliseconds*—Transmission interval
Range: 1 through 4,294,967
Default: 30 milliseconds

Usage Guidelines See “Modify the LSA Transmission Interval” on page 231.

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> <key-id <i>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 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 a Virtual Link” on page 225.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

