

Configuring OSPF

SYSTEM ADMINISTRATOR GUIDE

Copyright

© Ericsson AB 2009–2011. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

SmartEdge is a registered trademark of Telefonaktiebolaget LM Ericsson.

NetOp is a trademark of Telefonaktiebolaget LM Ericsson.



Contents

1	Overview	1
1.1	Areas	2
1.2	Non-Stop Routing	3
1.3	Router Functions	4
1.4	Route Selection Process	4
1.5	Packet Types	5
1.6	Link-State Advertisements	6
1.7	Sham Links	7
1.8	Virtual Links	7
1.9	Passive Mode	8
1.10	OSPFv3	9
2	Configuration and Operations Tasks	11
2.1	Configuring OSPF	11
2.2	Configuring OSPFv3	24
2.3	OSPF Operations Tasks	30
3	Configuration Examples	35
3.1	Basic OSPF	35
3.2	Configuring LDP-IGP Synchronization Between LDP and OSPF	39
3.3	OSPF Route Redistribution	40
3.4	OSPFv3 Route Redistribution and Aggregation	43
3.5	MD5 Authentication	43
3.6	Simple Key Chain	44





1 Overview

This document provides an overview of the Open Shortest Path First (OSPF) and describes the tasks and commands used to configure, monitor, troubleshoot, and administer OSPF features through the SmartEdge router.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the **Product Overview** folder of this Customer Product Information library.

OSPF is an Interior Gateway Protocol (IGP) that uses link-state advertisements (LSAs) to inform other routers of the state of the sender's links. In a link-state routing protocol, each router distributes information about its interfaces and neighbor relationships. The collection of the link states of individual routers forms a database that describes the autonomous system (AS) topology. As OSPF routers accumulate link-state information, they use the Shortest Path First (SPF) algorithm to calculate the shortest path to each node, which forms the basis for developing routing information for that autonomous system.

Ericsson supports multiple OSPF features, including those specified in the following IETF drafts and RFCs:

- RFC 2328, *OSPF Version 2*
- RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
- RFC 5250, *The OSPF Opaque LSA Option*
- RFC 1793, *Extending OSPF to support Demand Circuits*
- RFC3623, *Graceful OSPF Restart*
- RFC 3630, *Traffic Engineering Extensions to OSPF Version 2*
- RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*



- RFC 5309, *Point-to-point Operation over LAN in Link State Routing Protocols*

In OSPF, the autonomous system can be hierarchically organized by partitioning it into areas; see Figure 1.

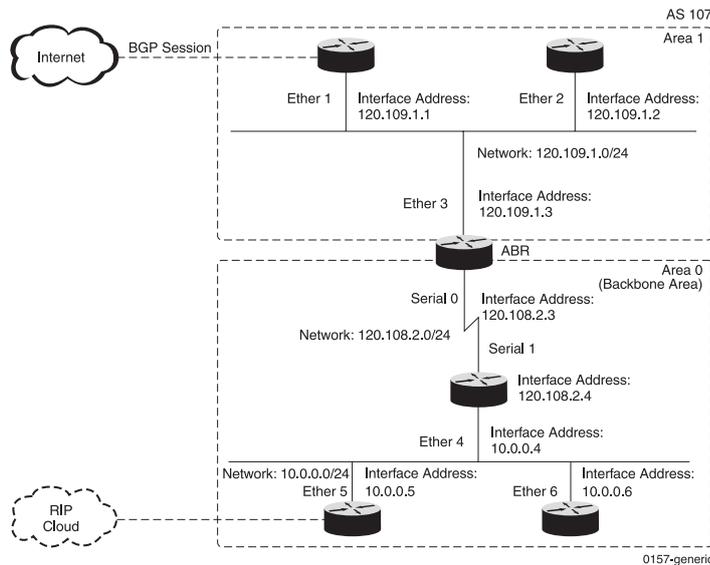


Figure 1 OSPF Hierarchy

Externally derived routes, also called AS-external routes, are routes learned from other routing protocols that are redistributed into the OSPF routing process. These AS-external routes are advertised to all areas, except for stub areas and not-so-stubby-areas (NSSAs). AS-external routes can also be forwarded out to another AS through routers on its boundary.

In-depth information on how OSPF is structured, and how it operates, is described in the following sections.

1.1 Areas

Each area can contain a group of contiguous networks and hosts. An area border router (ABR) communicates routing information between the areas.

Because routers within the same area share the same information, they have identical topological databases. An area's topology is invisible to entities outside the area. By keeping area topologies separate, OSPF passes less routing traffic than it would if an autonomous system were not partitioned.

Area partitioning creates two different types of OSPF routing, depending on whether the source and destination are in the same or different areas. Intra-area routing occurs when the source and destination are in the same area; inter-area routing occurs when they are in different areas.

The different area types are described in the following sections.



1.1.1 Normal and Backbone

A normal OSPF area, including the backbone area, is distinguished by the fact that it can carry transit traffic, allowing LSAs from outside the autonomous system (type 5 AS-external-LSAs) to be flooded throughout the area. Type 5 AS-external-LSAs can be originated both by routers internal to the area or by ABRs.

Hierarchical organization of an OSPF autonomous system requires one of the areas to be configured as the backbone area. The backbone area is configured with an identity of 0 and must be contiguous, contain all area border routers, and be responsible for distributing routing information to all other nonbackbone areas.

1.1.2 Stub

OSPF also allows some areas to be configured as stub areas. Type 5 AS-external LSAs are not flooded into a stub area, thereby reducing the link state database size and the processor and memory usage of routers inside stub areas. While a stub area cannot propagate routes external to the autonomous system in which it resides, it can propagate a default route, intra-area routes, and inter-area routes. A stub area relies on default routing to forward traffic addressed to external destinations. The backbone area cannot be configured as a stub area.

The SmartEdge router also supports totally-stubby areas which suppress the advertisement of type 3 and 4 LSAs. To create a totally-stubby area, you must enter the `area-type stub no-summary` command on the border routers of an area.

1.1.3 Not-So-Stubby-Area

Not-so-stubby-areas (NSSAs) are an extension of OSPF stub areas. Their intent is to preserve the properties of a stub area, while allowing limited import of external routes from other routing domains. These routes are imported as Type 7 NSSA-external LSAs, which are flooded only within the NSSA. For propagation of these routes to other areas, type 7 LSAs must be translated into type 5 external LSAs by the NSSA ABR. NSSA ABRs will also advertise a type 7 default route into the NSSA, and can be configured to summarize and to filter the translation of type 7 NSSA-external LSAs into Type 5 external LSAs.

1.2 Non-Stop Routing

OSPF non-stop routing (NSR) maintains OSPF neighbor relationships and operation in steady state when the active XCRP crashes and switches over to the standby XCRP. Consequently, other OSPF routers in the network are not aware of the active XCRP failure, and the OSPF routing domain continues to operate in steady state. All other routing domains will be disrupted.



NSR is available only when the status of the standby XCRP is in a ready state as displayed in the *show ospf command*.

IPsec multi-bind interfaces and other pseudo circuits are not supported.

If synchronization was not completed prior to the switchover, the OSPF instance will restart from scratch and without graceful-restart. Neighbors that are in the process of database exchange or the loading states will not have their progress saved across the switchover and will need to re-start the process..

1.3 Router Functions

Depending on its location in the OSPF hierarchy, an OSPF router can provide one or more of the following functions:

- Internal router—A router with all directly connected networks belonging to the same area. An internal router maintains a single topological database.
- Backbone router—A router that has one or more interfaces to the backbone area. The OSPF backbone is responsible for distributing routing information between areas.
- ABR—A router that attaches to multiple areas. ABRs maintain a separate topological database for each attached area and summarize the information for distribution to the backbone. The backbone in turn distributes the information to the other areas.
- ASBR—An autonomous system border router (ASBR) exchanges routing information with routers belonging to other autonomous systems, and advertises external routing information throughout its local autonomous system. The paths to each AS boundary router are known by every router in the autonomous system. ASBRs can be internal or area border routers, and may or may not participate in the backbone. ASBRs cannot be part of a stub area unless they are also ABRs; that is, connected to other non-stub areas.
- Designated router and backup designated router—On multi-access networks with more than one router, a designated router is responsible for generating the LSAs for the network. The designated router is elected by the Hello protocol. Designated routers allow a reduction in network traffic and in the size of the topological database. Backup designated routers provide a failsafe in case the designated router is not operational.

1.4 Route Selection Process

A routing table contains all the information necessary to forward an IP packet to a destination. When forwarding an IP data packet, the routing table entry providing the best match for the packet's IP destination is located. In the case of OSPF, the best path to a destination is determined via the SPF computation performed on the link-state database.



From the link-state database, the router uses the Dijkstra algorithm to construct a tree of shortest paths with itself as root. This shortest-path tree gives the route to each destination in the autonomous system. A separate SPF computation is performed and a different tree is constructed for each area in which the router resides. Externally derived routing information appears on the tree as leaves. OSPF intra-area and inter-area paths are preferred over external paths.

1.5 Packet Types

OSPF runs directly on top of IP (protocol 89). There are five types of packets specified in OSPF:

- Hello—The SmartEdge router sends Hello packets to its neighbors and receives their Hello packets. In this manner, adjacencies between neighbors are established. (Not all neighboring routers are adjacent.)
- Database description—Sent by adjacent routers when an adjacency is initialized, database description packets describe the contents of the respective database to synchronize the two neighboring databases.
- Link-state request—Requests pieces of the topological database from neighbor routers. These messages are sent after a router discovers (by examining database-description packets) that parts of its topological database are out of date.
- Link-state update—Responds to a link-state request packet. These messages are also used for the regular flooding of LSAs. Several LSAs can be included within a single link-state update packet.
- Link-state acknowledgment—Acknowledges link-state update packets.

Each packet includes a common header; see Figure 2.

Version Number	Type	Packet Length	Router ID	Area ID	Checksum	Authentication Type	Authentication	Data

0544

Figure 2 OSPF Packet Header

The OSPF packet header contains the following fields:

- Version Number—Identifies the OSPF version.
- Type—Identifies the OSPF packet type; for example, Hello, database description, link-state request, link-state update, and link-state acknowledgement.
- Packet Length—Specifies the packet length, including the OSPF header, in bytes.
- Router ID—Identifies the source of the packet.



- Area ID—Identifies the area to which the packet belongs. A packet is associated with a single area.
- Checksum—Checks the entire packet contents for any damage that may have occurred in transit.
- Authentication Type—Contains the authentication type. All OSPF protocol exchanges are authenticated. The authentication type is configurable on a per-area basis.
- Authentication—Contains authentication information.
- Data—Contains packet data.

1.6 Link-State Advertisements

Table 1 provides each LSA type and its description.

Table 1 LSA Type and Description

ID	Type	Description
1	Router-LSA	Originated by all routers. Describes the collected states of the router's interfaces to an area. Flooded throughout a single area only.
2	Network-LSA	Originated by the designated router. Contains the list of routers connected to the network. Flooded throughout a single area only.
3	Summary-LSA (networks)	Flooded throughout a single area only. Describes routes to networks. Each summary LSA describes a route to a destination outside the area, but still inside the autonomous system.
4	Summary-LSA (routers)	Flooded throughout a single area only. Describes routes to ASBRs. Each summary LSA describes a route to a destination outside the area, but still inside the autonomous system.
5	AS-external-LSA	Originated by ASBRs and flooded throughout the autonomous system. Each AS-external LSA describes a route to a destination in another autonomous system. Default routes for the AS can also be described by AS-external LSAs.
7	NSSA-external-LSA	Flooded throughout a single area only. Type 7 LSAs are advertised only within an NSSA. When forwarded outside the NSSA to non-stub areas, Type 7 LSAs are converted into Type 5 LSAs by an ABR configured to perform translation, or by the ABR with the highest router ID. ABRs can be configured to summarize and filter Type 7 LSAs.



Table 1 LSA Type and Description

ID	Type	Description
9	Link local scope opaque LSA	Type 9 Opaque LSAs are not flooded beyond the local (sub)network.
10	Area local scope opaque LSA	Type 10 Opaque LSAs are not flooded beyond the borders of their associated area.
11	AS scope opaque LSA	<p>The flooding scope of Type 11 LSAs are equivalent to the flooding scope of AS-external (Type 5) LSAs. Specifically, Type 11 Opaque LSAs are:</p> <ul style="list-style-type: none"> • Flooded throughout all transit areas • Not flooded into stub areas from the backbone • Not originated by routers into their connected stub areas

Note: Type 8 LSAs are not supported.

1.7 Sham Links

A sham link is an OSPF adjacency tunneled over a VPN backbone. Sham links allow the VPN backbone path to be preferred when there are intra-area backdoor links between customer edge (CE) routers in the VPN.

The local connected route corresponding to the source IP address for the sham link must be redistributed into BGP and advertised over the VPN infrastructure to a provider edge (PE) router containing the other end of the sham link.

The route corresponding the remote end of the sham link must be redistributed into the corresponding OSPF instance in the VPN context. VPN routing must be enabled for the OSPF instance.

The cost of the sham link can be configured or will inherit the BGP Multi-Exit Discriminator (MED) from the VPN route.

For more information on sham links, see RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs).

1.8 Virtual Links

The single backbone area (0.0.0.0) cannot be disconnected, or some areas of the autonomous system will become unreachable. To establish and maintain connectivity of the backbone, virtual links can be configured through non-backbone areas. Virtual links serve to connect physically separate components of the backbone. The two endpoints of a virtual link are area border routers. The virtual link must be configured in both routers. The



configuration information in each router consists of the other virtual endpoint (the other area border router), and the non-backbone area the two routers have in common, which is called the transit area. Virtual links cannot be configured through stub areas.

The virtual link is treated as if it were an unnumbered point-to-point network belonging to the backbone and joining the two area border routers. An attempt is made to establish an adjacency over the virtual link. When this adjacency is established, the virtual link is included in backbone router LSAs, and OSPF packets pertaining to the backbone area flow over the virtual adjacency.

In each endpoint router, the cost and viability of the virtual link is discovered by examining the routing table entry for the other endpoint router. An InterfaceUp event occurs for a virtual link when its corresponding routing table entry becomes reachable, and an InterfaceDown event occurs when its routing table entry becomes unreachable.

The other details concerning virtual links are as follows:

- AS-external-LSAs are NEVER flooded over virtual adjacencies.
- The cost of a virtual link is not configured.
- The IP interface address for the virtual interface and the virtual neighbor's IP address are determined by the routing table build process.
- In each endpoint's router-LSA for the backbone, the virtual link is represented as a Type 4 link whose link ID is set to the virtual neighbor's OSPF router ID and whose link data is set to the virtual interface's IP address.
- A non-backbone area can carry transit data traffic only if it serves as the transit area for one or more fully adjacent virtual links.
- The time between link state retransmissions, is configured for a virtual link.

For more information on virtual links, see RFC 2328, *OSPF Version 2*.

1.9 Passive Mode

Under normal operation, OSPF sends and receives OSPF packets on an interface, and advertises the interface's IP subnet as an intra-area stub network in the OSPF routing domain. When OSPF passive mode is enabled, OSPF continues to advertise the interface's IP subnet, but it does not send OSPF packets and drops all received OSPF packets. OSPF passive mode can be enabled for either of the following:

- An individual OSPF interface
- All interfaces within a defined OSPF area



For more information about configuring OSPF passive mode, see the description for the *passive* command.

1.10 OSPFv3

OSPF Version 3 (OSPFv3) is the version of OSPF that supports IP Version 6 (IPv6). The fundamental mechanisms of OSPF (flooding, area support, and Shortest Path First [SPF] calculations) remain unchanged in OSPFv3; however, because of changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6, the following changes have been made in OSPFv3:

- Addressing semantics has been removed from OSPF packets and basic LSAs.
- New LSAs exist to carry IPv6 addresses and prefixes.
- OSPFv3 runs on a per-link basis, instead of on a per-IP-subnet basis.
- Flooding scope for LSAs has been generalized.
- Authentication has been removed from OSPFv3; it is now handled by the IPv6 authentication header and encapsulating security payload.

OSPFv3 also supports all optional OSPF capabilities, including on-demand circuits, NSSA areas, and multicast extensions.

For a description of IPv6 addressing and the types of IPv6 addresses, see RFC 3513, *Internet Protocol Version 6 (IPv6) Addressing Architecture*.

For more information on OSPF support for IPv6, see RFC 5340, *OSPF for IPv6*.

Note: When IP Version 6 (IPv6) addresses are not referenced or explicitly specified, the term, IP address, can refer generally to IP Version 4 (IPv4) addresses, IPv6 addresses, or IP addressing. In instances where IPv6 addresses are referenced or explicitly specified, the term, IP address, refers only to IPv4 addresses.





2 Configuration and Operations Tasks

To configure OSPF or OSPFv3, perform the tasks described in the following sections.

Note: In this section, the command syntax in the task tables displays only the root command.

2.1 Configuring OSPF

To configure OSPF, perform the tasks described in the sections that follow.

2.1.1 Configure an OSPF Routing Instance

To configure an OSPF routing instance, perform the tasks described in Table 2. Enter all commands in OSPF router configuration mode, unless otherwise noted.

Table 2 Configure an OSPF Routing Instance

Task	Root Command	Notes
Create an OSPF routing instance and enter OSPF router configuration mode.	<i>router ospf</i>	Enter this command in context configuration mode.
Specify that the OSPF interface cost is computed automatically and to configure the reference bandwidth that is used in the interface cost computation.	<i>auto-cost</i>	The interface cost is computed by dividing the reference bandwidth by the interface speed. A cost of one is assigned if the interface speed is greater than the reference bandwidth. You can override the automatic cost setting on individual interfaces by issuing the <code>cost</code> command in OSPF interface configuration mode. For more information, see Configure an OSPF Interface.
Enable the advertisement of router capabilities using OSPF opaque LSAs.	<i>capabilities</i>	—

*Table 2 Configure an OSPF Routing Instance*

Task	Root Command	Notes
Configure a default metric that is used for redistributed OSPF routes when no metric is specified.	<i>default-metric (OSPF)</i>	—
Modify the OSPF distance value of one or more of these route types.	<i>distance (OSPF)</i>	The distance value of a route is used to select the preferred route when there are equivalent routes from multiple protocols. When a distance comparison is made the route with the lowest distance is selected. By default, OSPF external, inter-area, and intra-area routes are set to a distance value of 110.



Table 2 Configure an OSPF Routing Instance

Task	Root Command	Notes
<p>Enable fast convergence for the OSPF instance.</p>	<p><i>fast-convergence (OSPF)</i></p>	<p>Use the <i>spf-delay-interval</i> argument to set an SPF delay that is less than one second. When fast convergence is enabled, the <i>spf-delay-interval</i> argument provides an SPF delay with sub-second (millisecond) granularity, and the value for the <i>delay</i> argument of the <i>spf-timers</i> command (in OSPF router configuration mode) is ignored, regardless of whether it has been configured. Otherwise, under normal convergence, the <i>delay</i> argument value (in seconds) is used.</p> <p>Use the <i>max-spf-count</i> argument to allow additional SPF calculations within the SPF hold time specified by the <i>spf-timers</i> command. Specifying a value greater than zero effectively squeezes additional SPF calculations into the SPF time interval; specifying a value of zero does not allow for squeezing additional SPF calculations into the SPF hold time and returns OSPF to the standard SPF hold time behavior.</p>



Table 2 Configure an OSPF Routing Instance

Task	Root Command	Notes
Enable OSPF fast LSA origination for an OSPF instance.	<i>fast-lsa-origination</i>	<p>Normally, OSPF originates an LSA every five seconds. Because there can be multiple changes to router or network LSAs during that five-second interval, the five-second LSA origination limit can slow network convergence. When fast LSA origination is enabled, up to four instances of the same LSA can be originated in the same five-second interval.</p> <p>Likewise, LSA reception is normally rate limited to one new LSA instance per second. LSA instances received in less than one second after the previous LSA instance are dropped. When fast LSA origination is enabled, LSA reception is not restricted to one new instance per second.</p>
Enable graceful restart for an OSPF instance.	<i>graceful-restart (OSPF)</i>	—
Enables LDP-IGP synchronization between OSPF and LDP.	<i>ldp-igp-synchronization</i> [<i>timeout seconds</i>]	Use the optional <i>timeout seconds</i> construct to set the maximum number of seconds the interface waits before transporting traffic without receiving notification from LDP that label exchange is completed. ⁽¹⁾
Log neighbor transitions to and from the full neighbor adjacency state.	<i>log-neighbor-up-down</i>	—



Table 2 Configure an OSPF Routing Instance

Task	Root Command	Notes
Configure OSPF IGP shortcuts	<i>mpls igp-shortcut</i>	OSPF IGP shortcuts allow OSPF to use an MPLS LSP as the next hop. You must also configure the igp-shortcut command in RSVP router configuration mode or RSVP LSP router configuration mode for the shortcuts to work.
Enable the use of MPLS LSPs for tunneling LDP over Resource Reservation Protocol (RSVP).	<i>mpls tunnel-shortcut</i>	To configure LDP over RSVP, use the mpls tunnel-shortcut command in conjunction with the tunnel-shortcut command in both RSVP configuration mode (or RSVP LSP configuration mode) and LDP router configuration mode. For more information, see <i>Enabling LDP over RSVP</i> .
Enable the advertisement of OSPF Traffic Engineering (TE) metrics.	<i>mpls traffic-engineering</i>	—
Maintain OSPF neighbor relationships and operation in steady state when the active XCRP Controller card switches over to the standby XCRP Controller card.	<i>nonstop-routing</i>	To enable NSR, use the nonstop-routing command. The debug ospf nsr command can be used to track the synchronization between the active and standby XCRP Controller cards prior to switchover.
Originate the default route advertisement in the OSPF routing domain.	<i>originate-default</i>	—



Table 2 Configure an OSPF Routing Instance

Task	Root Command	Notes
Configure a fixed OSPF router ID for the SmartEdge router.	<i>router-id (OSPF)</i>	The router ID is used by OSPF to identify the originating router for packets and link-state advertisements (LSAs). If the OSPF router ID is not configured, OSPF chooses the lowest loopback interface address. If there are no loopback interfaces, OSPF chooses the lowest interface address. The default OSPF router ID is selected when OSPF is started initially or restarted using the <code>process restart ospf</code> command in exec mode.
Configure the delay time between the receipt of a topology change and the start of the SPF calculation, as well as to determine the hold time between two consecutive SPF calculations.	<i>spf-timers</i>	For OSPF, the default delay time is 2 seconds and the default hold time is 5 seconds.
Configure the SmartEdge router as an OSPF stub router.	<i>stub-router</i>	—
Configure the redistribution of routes into the OSPF routing instance.	For the complete list of tasks used to configure the redistribution of routes into the OSPF routing instance, see Configure the Redistribution of Routes into OSPF.	—
Configure an OSPF area.	For the complete list of tasks used to configure an OSPF area, see Configure an OSPF Area.	—

(1) Use the ***ldp-igp-synchronization timeout*** command in LDP configuration mode to set the maximum number of seconds LDP waits before notifying the IGP that label exchange is completed, so that IGP can start advertising the normal metric for the link. See *Configuring LDP* for more information.



2.1.2 Configure the Redistribution of Routes into OSPF

You can redistribute routes learned from other protocols into the OSPF routing instance, set a limit on the number of routes that can be redistributed into the OSPF routing instance, and set a limit on the number of routes per second that can be redistributed into the OSPF routing instance.

Note: IPv6 routes cannot be redistributed into an OSPF routing instance.

To configure the redistribution of routes into the OSPF routing instance, perform the tasks described in Table 3. Enter all commands in OSPF router configuration mode.

Table 3 Configure the Redistribution of Routes into OSPF

Task	Root Command	Notes
Redistribute routes learned from other protocols into the OSPF routing instance.	<i>redistribute (OSPF)</i>	—
Set a maximum limit on the number of routes that can be redistributed into the specified OSPF instance.	<i>maximum redistribute (OSPF)</i>	—
Set a maximum limit on the number of routes that can be redistributed per second into the OSPF routing instance.	<i>maximum redistribute-quantity</i>	—
Summarize external routes that are redistributed into the OSPF routing instance.	<i>summary-address (OSPF)</i>	—

2.1.3 Configure an OSPF Area

To configure an OSPF area, perform the tasks described in Table 4. Enter all commands in OSPF area configuration mode, unless otherwise noted.

Table 4 Configure an OSPF Area

Task	Root Command	Notes
Create an OSPF area and enter OSPF area configuration mode.	<i>area</i>	Enter this command in OSPF router configuration mode.
Define an OSPF area as a stub area or as an NSSA.	<i>area-type</i>	—
Change the attributes of a default route originated into a stub area or an NSSA.	<i>default-route</i>	—



Table 4 Configure an OSPF Area

Task	Root Command	Notes
Summarize NSSA routes advertised by an ABR.	<i>nssa-range</i>	—
Set all interfaces configured in the specified OSPF area to passive mode.	<i>passive</i>	<p>OSPF passive mode disables OSPF interfaces from sending OSPF packets.</p> <p>Setting all interfaces in an OSPF area to passive mode is useful for large, pure edge aggregation applications, where there may be hundreds, or perhaps thousands, of customer-facing circuits. To distribute routes for the customer-facing interfaces to the upstream routers, you can enable OSPF on the customer-facing interfaces, and then set them all to passive mode using the passive command in OSPF area configuration mode.</p>
Summarize inter-area routes advertised by an ABR.	<i>range</i>	—
Configure an OSPF interface.	For the complete list of tasks used to configure an OSPF interface, see Configure an OSPF Interface.	—
Configure an OSPF sham link.	For the complete list of tasks used to configure an OSPF sham link, see Configure an OSPF Sham Link.	—
Configure an OSPF virtual link.	For the complete list of tasks used to configure an OSPF virtual link, see Configure an OSPF Virtual Link.	—



2.1.4 Configure an OSPF Interface

To configure an OSPF interface, perform the tasks described in Table 5. Enter all commands in OSPF interface configuration mode, unless otherwise noted.

Table 5 Configure an OSPF Interface

Task	Root Command	Notes
Enable OSPF routing on an interface and enter OSPF interface configuration mode.	<i>interface (OSPF)</i>	Enter this command in OSPF area configuration mode.
Enable authentication and specify the authentication scheme for an OSPF interface.	<i>authentication (OSPF)</i>	Routes within the same area are not required to use the same authentication scheme and key ID; however, if two routers directly exchange updates, they must have the same authentication scheme and key ID.
Block the flooding of LSAs that are not self-originated.	<i>block-flooding</i>	Blocking flooding on an interface can result in inconsistencies between OSPF routers and their respective route tables. Exercise caution before blocking the flooding of LSAs that are not self-originated.
Configure the cost used in SPF computation for the specified OSPF-enabled interface.	<i>cost (OSPF)</i>	The lower the cost, the more likely the interface is to be used to forward data traffic.
Configure OSPF to treat a point-to-point (P2P) or a point-to-multipoint (P2MP) interface as a demand circuit.	<i>demand-circuit</i>	<p>Demand circuits are network segments whose costs vary with usage; charges can be based both on connect time and on bytes or packets transmitted. OSPF routing usually requires a demand circuit's underlying data-link connection to be constantly open, resulting in unwanted usage charges. Using the demand-circuit command enables OSPF Hello packets and the refresh of OSPF routing information to be suppressed on-demand circuits, allowing the underlying data-link connections to be closed when not carrying traffic.</p> <p>Hello suppression is not negotiated unless demand circuit support is enabled.</p>



Table 5 Configure an OSPF Interface

Task	Root Command	Notes
Disable Bidirectional Forwarding Detection (BFD) for an OSPF interface.	<i>disable-bfd (OSPF)</i>	Use the no or default form of this command to reenables BFD on an OSPF interface.
Enable the sending of more than one OSPF Hello packet per second on the interface.	<i>fast-hello</i>	Using this command results in faster OSPF convergence. The following restrictions apply to this command: <ul style="list-style-type: none">• After the fast-hello command is configured, you cannot use the hello-interval or router-dead interval commands until the fast-hello command has been disabled.• After the hello-interval or router-dead interval command has been configured, you cannot use the fast-hello command until the hello-interval or router-dead interval command has been disabled.
Suppress the periodic LSA refresh in stable topologies.	<i>flood-reduction</i>	If demand circuit operation is implicitly or explicitly enabled, LSAs are flooded as DoNotAge LSAs on the OSPF interface, and will not be reflooded until the network topology changes.
Configure the interval at which OSPF hello packets are sent on the interface.	<i>hello-interval (OSPF)</i>	—



Table 5 Configure an OSPF Interface

Task	Root Command	Notes
Enables or disables LDP-IGP synchronization between LDP and OSPF on this particular interface.	<i>ldp-igp-synchronization</i>	<p>When you use the ldp-igp-synchronization command in OSPF router configuration mode, LDP-IGP synchronization between LDP and OSPF is enabled for all interfaces on the router. You can then use the no ldp-igp-synchronization command in OSPF interface configuration mode to disable LDP-IGP synchronization between LDP and OSPF on an individual interface.</p> <p>Use the ldp-igp-synchronization command in OSPF interface configuration mode to reenable LDP-IGP synchronization between LDP and OSPF on an individual interface that has LDP-IGP synchronization between LDP and OSPF disabled.</p>
Configure an OSPF neighbor.	<i>neighbor (OSPF)</i>	—
Configure the OSPF network type.	<i>network-type</i>	<p>You can specify any of the following network types:</p> <ul style="list-style-type: none"> • Broadcast network—Broadcast networks support multiple routers and have the ability to address a single physical message to all attached routers. • Nonbroadcast multiaccess (NBMA)—A nonbroadcast network, such as frame relay, that simulates an OSPF broadcast network. • Point-to-point (P2P) network—A P2P network joins a single pair of routers. • Point-to-multipoint (P2MP) network—Acts as though the nonbroadcast network is a collection of P2P links.



Table 5 Configure an OSPF Interface

Task	Root Command	Notes
Set the OSPF interface to passive mode.	<i>passive</i>	OSPF passive mode disables OSPF interfaces from sending OSPF packets.
Modify the interval at which LSAs are retransmitted in link state update packets on an interface.	<i>retransmit-interval</i>	—
Modify the amount of time the OSPF routing process waits to receive an OSPF Hello packet from a neighbor before determining that the neighbor is not operational.	<i>router-dead-interval</i>	—
Modify the OSPF preference value for the SmartEdge router to act as the designated router on the network.	<i>router-priority</i>	—
Set a delay value, increasing the age of LSAs sent out through the OSPF interface.	<i>transmit-delay</i>	—

2.1.5 Configure an OSPF Sham Link

To configure an OSPF sham link, perform the tasks described in Table 6. Enter all commands in OSPF sham link configuration mode, unless otherwise noted.

Table 6 Configure an OSPF Sham Link

Task	Root Command	Notes
Create an OSPF adjacency tunneled over a VPN backbone (sham link).	<i>sham-link</i>	Enter this command in OSPF area configuration mode.
Enable authentication and specify the authentication scheme for an OSPF sham link.	<i>authentication (OSPF)</i>	Routes within the same area are not required to use the same authentication scheme and key ID; however, if two routers directly exchange updates, they must have the same authentication scheme and key ID.
Configure the cost used in SPF computation for the an OSPF sham link.	<i>cost (OSPF)</i>	The lower the cost, the more likely the sham link is to be used to forward data traffic.



Table 6 Configure an OSPF Sham Link

Task	Root Command	Notes
Configure the interval at which OSPF hello packets are sent out through an OSPF sham link.	<i>hello-interval (OSPF)</i>	—
Modify the interval at which LSAs are retransmitted in link state update packets on an OSPF sham link.	<i>retransmit-interval</i>	—
Modify the amount of time the OSPF routing process waits to receive an OSPF Hello packet from a neighbor before determining that the neighbor is not operational.	<i>router-dead-interval</i>	—
Set a delay value, increasing the age of LSAs sent out through an OSPF sham link.	<i>transmit-delay</i>	—

2.1.6 Configure an OSPF Virtual Link

To configure an OSPF virtual link, perform the tasks described in Table 7. Enter all commands in OSPF virtual link configuration mode, unless otherwise noted.

Table 7 Configure an OSPF Virtual Link

Task	Root Command	Notes
Create a virtual link through the specified transit area.	<i>virtual-link</i>	Enter this command in OSPF area configuration mode.
Enable authentication and specify the authentication scheme for an OSPF virtual link.	<i>authentication (OSPF)</i>	Routes within the same area are not required to use the same authentication scheme and key ID; however, if two routers directly exchange updates, they must have the same authentication scheme and key ID.
Configure the interval at which OSPF hello packets are sent out through an OSPF virtual link.	<i>hello-interval (OSPF)</i>	—
Modify the interval at which LSAs are retransmitted in link state update packets on an OSPF virtual link.	<i>retransmit-interval</i>	—



Table 7 Configure an OSPF Virtual Link

Task	Root Command	Notes
Modify the amount of time the OSPF routing process waits to receive an OSPF Hello packet from a neighbor before determining that the neighbor is not operational.	<i>router-dead-interval</i>	—
Set a delay value, increasing the age of LSAs sent out through an OSPF virtual link.	<i>transmit-delay</i>	—

2.2 Configuring OSPFv3

To configure OSPFv3, perform the tasks described in the following sections.

Note: If OSPFv3 is configured in a context that has no IPv4 addresses, the context or router interface must be configured with a router ID for OSPFv3 to operate properly.

2.2.1 Configure an OSPFv3 Routing Instance

To configure an OSPFv3 routing instance, perform the tasks described in Table 8. Enter all commands in OSPFv3 router configuration mode, unless otherwise noted.

Table 8 Configure an OSPFv3 Routing Instance

Task	Root Command	Notes
Create an OSPFv3 routing instance and enter OSPFv3 router configuration mode.	<i>router ospf3</i>	Enter this command in context configuration mode.
Specify that the OSPFv3 interface cost is computed automatically and to configure the reference bandwidth that is used in the interface cost computation.	<i>auto-cost</i>	The interface cost is computed by dividing the reference bandwidth by the interface speed. A cost of one is assigned if the interface speed is greater than the reference bandwidth. You can override the automatic cost setting on individual interfaces by issuing the <i>cost</i> command in OSPFv3 interface configuration mode. For more information, see Configure an OSPFv3 Interface.
Configure a default metric that is used for redistributed OSPFv3 routes when no metric is specified.	<i>default-metric (OSPF)</i>	—



Table 8 *Configure an OSPFv3 Routing Instance*

Task	Root Command	Notes
Modify the OSPFv3 distance value of one or more of these route types.	<i>distance (OSPF)</i>	The distance value of a route is used to select the preferred route when there are equivalent routes from multiple protocols. When a distance comparison is made the route with the lowest distance is selected. By default, OSPFv3 external, inter-area, and intra-area routes are set to a distance value of 110.
Enable graceful restart for an OSPFv3 instance.	<i>graceful-restart (OSPF)</i>	—
Log neighbor transitions to and from the full neighbor adjacency state.	<i>log-neighbor-up-down</i>	—
Originate the default route advertisement in the OSPFv3 routing domain.	<i>originate-default</i>	—
Configure a fixed OSPFv3 router ID for the SmartEdge router.	<i>router-id (OSPF)</i>	The router ID is used by OSPFv3 to identify the originating router for packets and link-state advertisements (LSAs). If the OSPFv3 router ID is not configured, OSPFv3 chooses the lowest loopback interface address. If there are no loopback interfaces, OSPFv3 chooses the lowest interface address. The default OSPFv3 router ID is selected when OSPFv3 is started initially or restarted using the <code>process restart ospf</code> command in exec mode.
Configure the delay time between the receipt of a topology change and the start of the OSPF calculation, as well as to determine the hold time between two consecutive SPF calculations.	<i>spf-timers</i>	The default delay time and hold time are both 0 seconds to enable fast convergence. Changing the time values to non-zero values disables fast convergence.
Configure the SmartEdge router as an OSPFv3 stub router.	<i>stub-router</i>	—



Table 8 Configure an OSPFv3 Routing Instance

Task	Root Command	Notes
Configure the redistribution of routes into the OSPFv3 routing instance.	For the complete list of tasks used to configure the redistribution of routes into the OSPFv3 routing instance, see OSPFv3 Route Redistribution and Aggregation.	—
Configure an OSPFv3 area.	For the complete list of tasks used to configure an OSPFv3 area, see Configure an OSPFv3 Area.	—

2.2.2 Configure the Redistribution of Routes into OSPFv3

You can redistribute routes learned from other protocols into the OSPFv3 routing instance, set a limit on the number of routes that can be redistributed into the OSPFv3 routing instance, and set a limit on the number of routes per second that can be redistributed into the OSPFv3 routing instance.

Note: IPv4 routes cannot be redistributed into an OSPFv3 routing instance.

To configure the redistribution of routes into the OSPFv3 routing instance, perform the tasks described in Table 9. Enter all commands in OSPF3 router configuration mode.

Table 9 Configure the Redistribution of Routes into OSPFv3

Task	Root Command
Redistribute routes learned from other protocols into the OSPFv3 routing instance.	<i>redistribute (OSPFv3)</i>
Set a maximum limit on the number of routes that can be redistributed into the specified OSPFv3 instance.	<i>maximum redistribute (OSPF)</i>
Set a maximum limit on the number of routes that can be redistributed per second into the OSPFv3 routing instance.	<i>maximum redistribute-quantum</i>
Summarize external routes that are redistributed into the OSPFv3 routing instance.	<i>summary-address (OSPF)</i>



2.2.3 Configure an OSPFv3 Area

To configure an OSPFv3 area, perform the tasks described in Table 10. Enter all commands in OSPF3 area configuration mode, unless otherwise noted.

Table 10 Configure an OSPFv3 Area

Task	Root Command	Notes
Create an OSPFv3 area and enter OSPF3 area configuration mode.	<i>area</i>	Enter this command in OSPF3 router configuration mode.
Define an OSPFv3 area as a stub area or as an NSSA.	<i>area-type</i>	—
Change the attributes of a default route originated into a stub area or an NSSA.	<i>default-route</i>	—
Summarize NSSA routes advertised by an ABR.	<i>nssa-range</i>	—
Summarize inter-area routes advertised by an ABR.	<i>range</i>	—
Configure an OSPFv3 interface.	For the complete list of tasks used to configure an OSPFv3 interface, see Configure an OSPFv3 Interface.	—

2.2.4 Configure an OSPFv3 Interface

To configure an OSPFv3 interface, perform the tasks described in Table 11. Enter all commands in OSPF3 interface configuration mode, unless otherwise noted.

Table 11 Configure an OSPFv3 Interface

Task	Root Command	Notes
Enable OSPFv3 routing on an interface and enter OSPF3 interface configuration mode.	<i>interface (OSPF)</i>	Enter this command in OSPF3 area configuration mode.
Block the flooding of LSAs that are not self-originated.	<i>block-flooding</i>	Blocking flooding on an interface can result in inconsistencies between OSPFv3 routers and their respective route tables. Exercise caution before blocking the flooding of LSAs that are not self-originated.



Table 11 Configure an OSPFv3 Interface

Task	Root Command	Notes
Configure the cost used in SPF computation for the specified OSPFv3-enabled interface.	<i>cost (OSPF)</i>	The lower the cost, the more likely the interface is to be used to forward data traffic.
Configure OSPFv3 to treat a P2P or a P2MP interface as a demand circuit.	<i>demand-circuit</i>	<p>Demand circuits are network segments whose costs vary with usage; charges can be based both on connect time and on bytes or packets transmitted. OSPFv3 routing usually requires a demand circuit's underlying data-link connection to be constantly open, resulting in unwanted usage charges. Using the demand-circuit command enables OSPFv3 Hello packets and the refresh of OSPFv3 routing information to be suppressed on-demand circuits, allowing the underlying data-link connections to be closed when not carrying traffic.</p> <p>Hello suppression is not negotiated unless demand circuit support is enabled.</p>
Suppress the periodic LSA refresh in stable topologies.	<i>flood-reduction</i>	If demand circuit operation is implicitly or explicitly enabled, LSAs are flooded as DoNotAge LSAs on the OSPFv3 interface, and will not be reflooded until the network topology changes.
Configure the interval at which OSPFv3 hello packets are sent on the interface.	<i>hello-interval (OSPF)</i>	—
Configure an OSPFv3 neighbor.	<i>neighbor (OSPF)</i>	—



Table 11 Configure an OSPFv3 Interface

Task	Root Command	Notes
Configure the OSPFv3 network type.	<i>network-type</i>	You can specify any of the following network types: <ul style="list-style-type: none"> • Broadcast network—Broadcast networks support multiple routers and have the ability to address a single physical message to all attached routers. • Nonbroadcast multiaccess (NBMA)—A nonbroadcast network, such as frame relay, that simulates an OSPFv3 broadcast network. • Point-to-point (P2P) network—A P2P network joins a single pair of routers.
Set the OSPFv3 interface to passive mode.	<i>passive</i>	OSPF passive mode disables OSPFv3 interfaces from sending OSPF packets.
Modify the interval at which LSAs are retransmitted in link-state update packets on an interface.	<i>retransmit-interval</i>	—
Modify the amount of time the OSPFv3 routing process waits to receive an OSPFv3 Hello packet from a neighbor before determining that the neighbor is not operational.	<i>router-dead-interval</i>	—
Modify the OSPFv3 preference value for the SmartEdge router to act as the designated router on the network.	<i>router-priority</i>	—
Set a delay value, increasing the age of LSAs sent out through the OSPFv3 interface.	<i>transmit-delay</i>	—

2.2.5 Configure an OSPFv3 Virtual Link

To configure an OSPFv3 virtual link, perform the tasks described in Table 12. Enter all commands in OSPF3 virtual link configuration mode, unless otherwise noted.



Table 12 Configure an OSPFv3 Virtual Link

Task	Root Command	Notes
Create an OSPFv3 virtual link through the specified transit area.	<i>virtual-link</i>	Enter this command in OSPF3 area configuration mode.
Configure the interval at which OSPFv3 hello packets are sent out through an OSPFv3 virtual link.	<i>hello-interval (OSPF)</i>	—
Modify the interval at which LSAs are retransmitted in link state update packets on an OSPFv3 virtual link.	<i>retransmit-interval</i>	—
Modify the amount of time the OSPFv3 routing process waits to receive an OSPFv3 Hello packet from a neighbor before determining that the neighbor is not operational.	<i>router-dead-interval</i>	—
Set a delay value, increasing the age of LSAs sent out through an OSPFv3 virtual link.	<i>transmit-delay</i>	—

2.3 OSPF Operations Tasks

To manage OSPF and OSPFv3 functions, perform the appropriate tasks described in Table 13. Enter the **show** commands in any mode; enter the **clear**, **debug**, and **monitor** commands in exec mode.

Table 13 OSPF and OSPFv3 Operations Tasks

Task	Root Command
Clear OSPF neighbor adjacencies, routes redistributed into OSPF, all routes, or statistics.	<i>clear ospf</i>
Enable the generation of debug messages for all OSPF events.	<i>debug ospf</i>
Enable the generation of debug messages for OSPF flooding events.	<i>debug ospf flooding</i>
Enable the generation of debug messages for a specific OSPF interface.	<i>debug ospf interface</i>
Enable the generation of debug messages for OSPF LSDB events.	<i>debug ospf lsdb</i>
Enable the generation of debug messages for OSPF neighbor events.	<i>debug ospf neighbor</i>
Enable the generation of debug messages for OSPF packet events.	<i>debug ospf packet</i>



Table 13 OSPF and OSPFv3 Operations Tasks

Task	Root Command
Enable the generation of debug messages for interactions between OSPF and the Router Configuration Manager (RCM).	<i>debug ospf rcm</i>
Enable the generation of debug messages for OSPF redistribution events.	<i>debug ospf redistribution</i>
Enable the generation of debug messages for interactions between OSPF and the Routing Information Base (RIB).	<i>debug ospf rib</i>
Enable the generation of debug messages for OSPF SPF calculations.	<i>debug ospf spf</i>
Display continuously updated information about OSPF interfaces.	<i>monitor ospf interface</i>
Display continuously updated information about OSPF neighbors.	<i>monitor ospf neighbor</i>
Display continuously updated information about the most recent OSPF SPF calculation.	<i>monitor ospf spf last</i>
Display continuously updated information about OSPF statistics.	<i>monitor ospf statistics</i>
Display the current OSPF configuration information for the current context.	<i>show configuration ospf</i>
Display high-level information for all OSPF instances, or optionally, for a specific OSPF instance.	<i>show ospf</i>
Display OSPF area information.	<i>show ospf area</i>
Display OSPF ABR and ASBR information.	<i>show ospf border-routers</i>
Display information stored in the OSPF link-state database (LSDB).	<i>show ospf database</i>
Display information about OSPF advertising router LSAs.	<i>show ospf database advertising-router</i>
Display information about OSPF opaque Type 10 link-state advertisements (LSAs).	<i>show ospf database area-scope-opaque</i>
Display information about OSPF opaque Type 11 LSAs.	<i>show ospf database as-scope-opaque</i>
Display a count, grouped by type, of OSPF LSAs.	<i>show ospf database database-summary</i>
Display information about OSPF Type 5 AS external LSAs.	<i>show ospf database external</i>
Display information about the OSPF interface LSD.	<i>show ospf database interface</i>
Display information about OSPF opaque Type 9 Shortest Path First (SPF).	<i>show ospf database link-scope-opaque</i>
Display information about OSPF network LSAs.	<i>show ospf database network</i>
Display information about OSPF not-so-stubby-area (NSSA) LSAs.	<i>show ospf database nssa</i>



Table 13 OSPF and OSPFv3 Operations Tasks

Task	Root Command
Display information about OSPF router LSAs.	<i>show ospf database router</i>
Display information about OSPF Type 4 summary autonomous system boundary routers (ASBRs) and other OSPFv3 routers.	<i>show ospf database summary-asbr</i>
Display information about OSPF Type 3 summary network LSAs.	<i>show ospf database summary-network</i>
Display the OSPF debug settings that have been enabled.	<i>show ospf debug</i>
Display OSPF interface information.	<i>show ospf interface</i>
Display OSPF neighbor information.	<i>show ospf neighbor</i>
Display OSPF route information.	<i>show ospf route</i>
Display OSPF SPF computation information.	<i>show ospf route vpn</i>
Display OSPF statistics.	<i>show ospf statistics</i>
Display OSPF summary address information.	<i>show ospf summary-address</i>
Display high-level information for all OSPFv3 instances, or optionally, for a specific instance.	<i>show ospf3</i>
Display information about OSPFv3 areas.	<i>show ospf3 area</i>
Display routes to ASBRs and other OSPFv3 routers.	<i>show ospf3 asbr</i>
Display information stored in the OSPFv3 LSD.	<i>show ospf3 database</i>
Display information about OSPFv3 advertising router LSAs.	<i>show ospf3 database advertising-router</i>
Display information about OSPFv3 grace LSA database entries.	<i>show ospf3 database grace</i>
Display information about OSPFv3 inter-area prefix LSA database entries.	<i>show ospf3 database inter-prefix</i>
Display information about OSPFv3 inter-area router LSA database entries.	<i>show ospf3 database inter-router</i>
Display information about OSPFv3 intra-area prefix LSA database entries.	<i>show ospf3 database intra-prefix</i>
Display information about OSPFv3 link LSAs.	<i>show ospf3 database link</i>
Display information about OSPFv3 network LSAs.	<i>show ospf3 database network</i>
Display information about OSPFv3 NSSA LSAs.	<i>show ospf3 database nssa</i>
Display information about OSPFv3 router LSAs.	<i>show ospf3 database router</i>
Display OSPFv3 debug information.	<i>show ospf3 debug</i>
Display OSPFv3 interface information.	<i>show ospf3 interface</i>
Display OSPFv3 intra-RIB information.	<i>show ospf3 intra-rib</i>
Display OSPFv3 malform log information.	<i>show ospf3 malform</i>

*Table 13 OSPF and OSPFv3 Operations Tasks*

Task	Root Command
Display OSPFv3 neighbor information.	<i>show ospf3 neighbor</i>
Display OSPFv3 route information.	<i>show ospf3 route</i>
Display OSPFv3 SPF calculation statistics.	<i>show ospf3 spf</i>
Display OSPFv3 statistics.	<i>show ospf3 statistics</i>
Display OSPFv3 summary address information.	<i>show ospf3 summary-address</i>





3 Configuration Examples

The sections that follow provide OSPF configuration examples.

Figure 3 illustrates the base OSPF topology for the examples provided in this section.

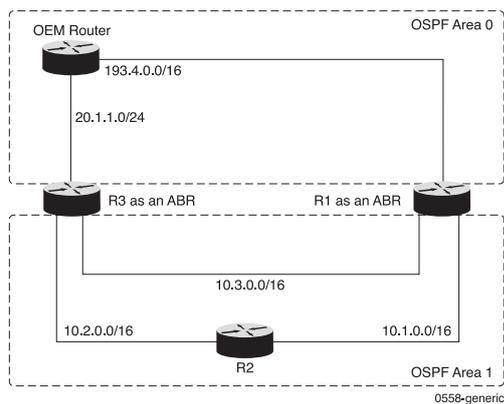


Figure 3 OSPF Topology

3.1 Basic OSPF

This section contains the basic OSPF configuration for the three routers, **R1**, **R2**, and **R3**, illustrated in Figure 3. Examples in proceeding sections contain only the configuration sections different from the examples here.

The basic configuration for **R1** is as follows. Because no router ID is explicitly configured, the loopback address is used as the OSPF router ID for **R1**:

```
[local]R1(config)#context local
[local]R1(config-ctx)#ip domain-lookup
[local]R1(config-ctx)#interface one
[local]R1(config-if)#ip address 193.4.5.2/16
[local]R1(config-if)#exit
[local]R1(config-ctx)#interface two
[local]R1(config-if)#ip address 10.1.1.1/16
[local]R1(config-if)#exit
```



```
[local]R1 (config-ctx) #interface three
[local]R1 (config-if) #ip address 10.3.1.1/16
[local]R1 (config-if) #exit
[local]R1 (config-ctx) #interface lo1 loopback
[local]R1 (config-if) #ip address 193.10.25.7/32
[local]R1 (config-if) #exit
[local]R1 (config-ctx) #router ospf 1
[local]R1 (config-ospf) #area 0.0.0.0
[local]R1 (config-ospf-area) #interface 193.4.5.2
[local]R1 (config-ospf-if) #exit
[local]R1 (config-ospf-area) #interface 193.10.25.7
[local]R1 (config-ospf-area) #exit
[local]R1 (config-ospf) #area 0.0.0.1
[local]R1 (config-ospf-area) #interface two
[local]R1 (config-ospf-if) #exit
[local]R1 (config-ospf-area) #interface three
[local]R1 (config-ospf-if) #exit
[local]R1 (config-ospf-area) #exit
[local]R1 (config-ospf) #exit
[local]R1 (config-ctx) #exit
[local]R1 (config) #port pos 5/1
[local]R1 (config-port) #bind interface one local
[local]R1 (config-port) #no shutdown
[local]R1 (config-port) #exit
[local]R1 (config) #port pos 5/2
[local]R1 (config-port) #bind interface two local
```



```
[local]R1(config-port)#no shutdown
[local]R1(config-port)#exit
[local]R1(config)#port pos 5/3
[local]R1(config-port)#bind interface three local
[local]R1(config-port)#no shutdown
```

The basic configuration for **R2** is as follows:

```
[local]R2(config)#context local
[local]R2(config-ctx)#ip domain-lookup
[local]R2(config-ctx)#interface one
[local]R2(config-if)#ip address 10.1.2.2/16
[local]R2(config-if)#exit
[local]R2(config-ctx)#interface two
[local]R2(config-if)#ip address 10.2.1.1/16
[local]R2(config-if)#exit
[local]R2(config-ctx)#router ospf 1
[local]R2(config-ospf)#router-id 22.22.22.22
[local]R2(config-ospf)#area 0.0.0.1
[local]R2(config-ospf-area)#interface 10.1.2.2
[local]R2(config-ospf-if)#exit
[local]R2(config-ospf-area)#interface 10.2.1.1
[local]R2(config-ospf-if)#exit
[local]R2(config-ospf-area)#exit
[local]R2(config-ospf)#exit
[local]R2(config-ctx)#exit
```



```
[local]R2 (config)#port pos 3/1
[local]R2 (config-port)#bind interface one local
[local]R2 (config-port)#no shutdown
[local]R2 (config-port)#exit
[local]R2 (config)#port ethernet 4/1
[local]R2 (config-port)#bind interface two local
[local]R2 (config-port)#no shutdown
```

The basic configuration for **R3** is as follows:

```
[local]R3 (config)#context local
[local]R3 (config-ctx)#ip domain-lookup
[local]R3 (config-ctx)#interface one
[local]R3 (config-if)#ip address 10.3.2.2/16
[local]R3 (config-if)#exit
[local]R3 (config-ctx)#interface two
[local]R3 (config-if)#ip address 10.2.2.2/16
[local]R3 (config-if)#exit
[local]R3 (config-ctx)#interface three
[local]R3 (config-if)#ip address 20.1.1.1/24
[local]R3 (config-if)#exit
[local]R3 (config-ctx)#router ospf 1
[local]R3 (config-ospf)#router-id 33.33.33.33
[local]R3 (config-ospf)#area 0.0.0.0
[local]R3 (config-ospf-area)#interface 20.1.1.1
[local]R3 (config-ospf-if)#exit
[local]R3 (config-ospf-area)#exit
```



```
[local]R3(config-ospf)#area 0.0.0.1
[local]R3(config-ospf-area)#interface 10.2.2.2
[local]R3(config-ospf-if)#exit
[local]R3(config-ospf-area)#interface 10.3.2.2
[local]R3(config-ospf-if)#exit
[local]R3(config-ospf-area)#exit
[local]R3(config-ospf)#exit
[local]R3(config-ctx)#exit
[local]R3(config)#port pos 3/1
[local]R3(config-port)#bind interface one local
[local]R3(config-port)#no shutdown
[local]R3(config-port)#exit
[local]R3(config)#port ethernet 1/1
[local]R3(config-port)#bind interface two local
[local]R3(config-port)#no shutdown
[local]R3(config-port)#exit
[local]R3(config)#port pos 3/2
[local]R3(config-port)#bind interface three local
[local]R3(config-port)#no shutdown
```

3.2 Configuring LDP-IGP Synchronization Between LDP and OSPF

The following example illustrates how to configure LDP-IGP synchronization between LDP and OSPF. In this example, the user enables LDP-IGP synchronization between LDP and OSPF on all interfaces configured in the context called OSPF1. The user then disables LDP-IGP synchronization with OSPF on one particular interface, called `ospf-int-1`. Finally, the user sets the maximum number of seconds LDP waits before notifying OSPF that label exchange is completed to 10 seconds:



Enable LDP-IGP synchronization with OSPF on all interfaces configured on the router:

```
[local]Redback(config)#context OSPF1
[local]Redback(config-ctx)#router ospf 1
[local]Redback(config-ospf)#ldp-igp-synchronization timeout 1000
```

Disable LDP-IGP synchronization with OSPF on the interface ospf-int-1:

```
[local]Redback(config-ospf)#area 0.0.0.0
[local]Redback(config-ospf-area)#interface ospf-int-1
[local]Redback(config-ospf-if)#no ldp-igp-synchronization
```

Set the maximum number of seconds LDP waits before notifying OSPF that label exchange is completed to 10 seconds:

```
[local]Redback(config-ospf-if)#exit
[local]Redback(config-ospf-area)#exit
[local]Redback(config-ospf)#exit
[local]Redback(config-ctx)#ldp
[local]Redback(config-ldp)#ldp-igp synchronization timeout 10
```

3.3 OSPF Route Redistribution

The following example illustrates how to redistribute static routes into the OSPF routing instance and how to modify the attributes of the redistributed routes. Only the routes matching the **122-nets-only** IP prefix list are selected for redistribution. These routes are **122.1.1.0/24**, **122.1.2.0/24**, and **122.1.3.0/24**. Once redistributed to OSPF, the routes are advertised with metric type 1 and metric value of 500. All modifications are accomplished by using the route map, **static-to-ospf**:

```
[local]Redback(config)#context local
```



```
[local]Redback (config-ctx)#ip domain-lookup
[local]Redback (config-ctx)#interface one
[local]Redback (config-if)#ip address 10.1.2.2/16
[local]Redback (config-if)#exit
[local]Redback (config-ctx)#interface two
[local]Redback (config-if)#ip address 10.2.1.1/16
[local]Redback (config-if)#exit
[local]Redback (config-ctx)#interface three
[local]Redback (config-if)#ip address 10.5.1.1/30
[local]Redback (config-if)#exit
[local]Redback (config-ctx)#router ospf 1
[local]Redback (config-ospf)#router-id 22.22.22.22
[local]Redback (config-ospf)#area 0.0.0.1
[local]Redback (config-ospf-area)#interface 10.1.2.2
[local]Redback (config-ospf-if)#exit
[local]Redback (config-ospf-area)#interface 10.2.1.1
[local]Redback (config-ospf-if)#exit
[local]Redback (config-ospf-area)#exit
[local]Redback (config-ospf)#redistribute static route-map static-to-ospf
[local]Redback (config-ospf)#exit
[local]Redback (config-ctx)#ip prefix-list 122-nets-only
[local]Redback (config-prefix-list)#seq 10 permit 122.0.0.0/8 le 24
[local]Redback (config-prefix-list)#seq 20 deny 0.0.0.0/0
[local]Redback (config-prefix-list)#exit
[local]Redback (config-ctx)#route-map static-to-ospf permit 10
[local]Redback (config-route-map)#match ip address prefix-list 122-nets-only
```



```
[local]Redback(config-route-map)#set metric 500
[local]Redback(config-route-map)#set metric-type type-1
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#ip route 50.0.0.0/8 three
[local]Redback(config-ctx)#ip route 121.1.1.0/24 three
[local]Redback(config-ctx)#ip route 121.1.2.0/24 three
[local]Redback(config-ctx)#ip route 121.1.3.0/24 three
[local]Redback(config-ctx)#ip route 121.1.5.0/24 three
[local]Redback(config-ctx)#ip route 122.1.1.0/24 three
[local]Redback(config-ctx)#ip route 122.1.2.0/24 three
[local]Redback(config-ctx)#ip route 122.1.3.0/24 three
[local]Redback(config-ctx)#exit
[local]Redback(config)#port pos 3/1
[local]Redback(config-port)#bind interface one local
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#exit
[local]Redback(config)#port ethernet 4/1
[local]Redback(config-port)#bind interface two local
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#exit
[local]Redback(config)#port pos 3/2
[local]Redback(config-port)#bind interface three local
[local]Redback(config-port)#no shutdown
```



3.4 OSPFv3 Route Redistribution and Aggregation

The following example configures route redistribution and aggregation for an OSPFv3 routing instance. First, configure a list of aggregate IP prefixes:

```
[local]Router(config-ctx)#ipv6 prefix-list test1-aggregate
[local]Router(config-ipv6-prefix-list)#seq 10 permit 4001:101:101:106::/64 ge 64
[local]Router(config-ipv6-prefix-list)#seq 20 permit 5001:101:101:106::/64 ge 64
[local]Router(config-ipv6-prefix-list)#seq 30 permit 6001:101:101:106::/64 ge 64
[local]Router(config-ipv6-prefix-list)#seq 40 permit 7001:101:101:106::/64 ge 64
[local]Router(config-ipv6-prefix-list)#seq 50 permit 2001:101:101::/48 ge 48
```

Next, configure a route map called `test1` that aggregates the IPv6 prefixes in the aggregate prefix list called `test1-aggregate`:

```
[local]Router(config-ctx)#route-map test1 permit 10
[local]Router(config-route-map)#match ipv6 address prefix-list test1-aggregate
[local]Router(config-route-map)#set ipv6 aggregate test1-aggregate
```

Specify that routes selected for redistribution are summarized only if they contain any of the prefixes specified in the IPv6 prefix list called `test1`:

```
[local]Redback(config-ctx)#router ospf3 1
[local]Redback(config-ospf3)#redistribute subscriber static route-map test1
```

Configure the static routes. In this example, the routes match the aggregate prefix `2001:101:101::/48`:

```
[local]Redback(config-ctx)#ipv6 route 2001:101:101:303::/64 80::2
[local]Redback(config-ctx)#ipv6 route 2001:101:101:304::/64 80::2
[local]Redback(config-ctx)#ipv6 route 2001:101:101:305::/64 80::2
[local]Redback(config-ctx)#ipv6 route 2001:101:101:306::/64 80::2
[local]Redback(config-ctx)#ipv6 route 2001:101:101:307::/64 80::2
```

Note: When an IP prefix list is used for aggregation, the `ge` and `le` parameters (configured with the `seq` command) are ignored and the prefix list entries match any route subsumed by the prefix. In such cases, the `ge` parameter is implicit.

3.5 MD5 Authentication

The following example shows how to use MD5 to provide authentication between two routers. Authentication is only configured at the interface level. A different type of authentication can be used on each interface and no area configuration is required.



The configuration for **SE1** is as follows:

```
[local]SE1(config-ctx)#router ospf 1
[local]SE1(config-ospf)#area 0.0.0.0
[local]SE1(config-ospf-area)#interface 193.4.5.2
[local]SE1(config-ospf-if)#exit
[local]SE1(config)#interface 193.10.25.7
[local]SE1(config-ospf-if)#exit
[local]SE1(config-ospf-area)#exit
[local]SE1(config-ospf)#area 0.0.0.1
[local]SE1(config-ospf-area)#interface two
[local]SE1(config-ospf-if)#authentication md5 ospf-key-chain
[local]SE1(config-ospf-if)#exit
[local]SE1(config-ospf-area)#interface three
```

The configuration for **SE2** is as follows:

```
[local]SE2(config-ctx)#router ospf 1
[local]SE2(config-ospf)#router-id 22.22.22.22
[local]SE2(config-ospf)#area 0.0.0.1
[local]SE2(config-ospf-area)#interface 10.1.2.2
[local]SE2(config-ospf-if)#authentication md5 ospf-key-chain
[local]SE2(config-ospf-if)#exit
[local]SE2(config-ospf-area)#interface 10.2.1.1
```

3.6 Simple Key Chain

This example shows how key chain lifetimes can be used to non-disruptively switch from one key string to another. OSPF always sends the key with the



most recent send-lifetime start time which is not greater than the current time. It accepts any key whose accept lifetime value includes the current time.

The configuration for both **SE1** and **SE2** is as follows:

```
[local]Redback (config-ctx) #key-chain ospf-key-chain key-id 1
[local]Redback (config-key-chain) #key-string secret
[local]Redback (config-key-chain) #accept-lifetime 2001:09:07:00:00:00
2002:09:07:12:00:00
[local]Redback (config-key-chain) #send-lifetime 2001:09:07:00:00:00
2002:09:07:08:00:00
[local]Redback (config-key-chain) #exit
[local]Redback (config-ctx) #key-chain ospf-key-chain key-id 2
[local]Redback (config-key-chain) #key-string psst
[local]Redback (config-key-chain) #accept-lifetime 2002:09:07:00:00:00
2003:09:07:12:00:00
[local]Redback (config-key-chain) #send-lifetime 2002:09:07:08:00:00
2003:09:07:07:00:00
```