



OSPF

This chapter describes how to configure the threat defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

- [OSPF, on page 1](#)
- [Requirements and Prerequisites for OSPF, on page 4](#)
- [Guidelines for OSPF, on page 4](#)
- [Configure OSPFv2, on page 7](#)
- [Configure OSPFv3, on page 19](#)
- [History for OSPF, on page 28](#)

OSPF

This chapter describes how to configure the threat defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

About OSPF

OSPF is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements rather than routing table updates. Because only LSAs are exchanged instead of the entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

The advantages of OSPF over RIP include the following:

- OSPF link-state database updates are sent less frequently than RIP updates, and the link-state database is updated instantly, rather than gradually, as stale information is timed out.
- Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The threat defense device calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.

The disadvantage of shortest path first algorithms is that they require a lot of CPU cycles and memory.

The threat defense device can run two processes of OSPF protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to coexist, but OSPF does not allow overlapping addresses). Or you might want to run one process on the inside and another on the outside, and redistribute a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPF routing process from another OSPF routing process, a RIP routing process, or from static and connected routes configured on OSPF-enabled interfaces.

The threat defense device supports the following OSPF features:

- Intra-area, inter-area, and external (Type I and Type II) routes.
- Virtual links.
- LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Configuring the threat defense device as a designated router or a designated backup router. The threat defense device also can be set up as an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

OSPF supports MD5 and clear text neighbor authentication. Authentication should be used with all routing protocols when possible because route redistribution between OSPF and other protocols (such as RIP) can potentially be used by attackers to subvert routing information.

If NAT is used, if OSPF is operating on public and private areas, and if address filtering is required, then you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR Type 3 LSA filtering, you can have separate private and public areas with the ASA acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other, which allows you to use NAT and OSPF together without advertising private networks.



Note Only Type 3 LSAs can be filtered. If you configure the threat defense device as an ASBR in a private network, it will send Type 5 LSAs describing private networks, which will get flooded to the entire AS, including public areas.

If NAT is employed but OSPF is only running in public areas, then routes to public networks can be redistributed inside the private network, either as default or Type 5 AS external LSAs. However, you need to configure static routes for the private networks protected by the threat defense device. Also, you should not mix public and private networks on the same threat defense device interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the threat defense device at the same time.

OSPF Support for Fast Hello Packets

The OSPF Support for Fast Hello Packets feature provides a way to configure the sending of hello packets in intervals less than one second. Such a configuration would result in faster convergence in an Open Shortest Path First (OSPF) network.

Prerequisites for OSPF Support for Fast Hello Packets

OSPF must be configured in the network already or configured at the same time as the OSPF Support for Fast Hello Packets feature.

OSPF Hello Interval and Dead Interval

OSPF hello packets are packets that an OSPF process sends to its OSPF neighbors to maintain connectivity with those neighbors. The hello packets are sent at a configurable interval (in seconds). The defaults are 10 seconds for an Ethernet link and 30 seconds for a non broadcast link. Hello packets include a list of all neighbors for which a hello packet has been received within the dead interval. The dead interval is also a configurable interval (in seconds), and defaults to four times the value of the hello interval. The value of all hello intervals must be the same within a network. Likewise, the value of all dead intervals must be the same within a network.

These two intervals work together to maintain connectivity by indicating that the link is operational. If a router does not receive a hello packet from a neighbor within the dead interval, it will declare that neighbor to be down.

OSPF Fast Hello Packets

OSPF fast hello packets refer to hello packets being sent at intervals of less than 1 second. To understand fast hello packets, you should already understand the relationship between OSPF hello packets and the dead interval. See [OSPF Hello Interval and Dead Interval, on page 3](#).

OSPF fast hello packets are achieved by using the `ospf dead-interval` command. The dead interval is set to 1 second, and the hello-multiplier value is set to the number of hello packets you want sent during that 1 second, thus providing subsecond or "fast" hello packets.

When fast hello packets are configured on the interface, the hello interval advertised in the hello packets that are sent out this interface is set to 0. The hello interval in the hello packets received over this interface is ignored.

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval.

Benefits of OSPF Fast Hello Packets

The benefit of the OSPF Fast Hello Packets feature is that your OSPF network will experience faster convergence time than it would without fast hello packets. This feature allows you to detect lost neighbors within 1 second. It is especially useful in LAN segments, where neighbor loss might not be detected by the Open System Interconnection (OSI) physical layer and data-link layer.

Implementation Differences Between OSPFv2 and OSPFv3

OSPFv3 is not backward compatible with OSPFv2. To use OSPF to route both IPv4 and IPv6 traffic, you must run both OSPFv2 and OSPFv3 at the same time. They coexist with each other, but do not interact with each other.

The additional features that OSPFv3 provides include the following:

- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Requirements and Prerequisites for OSPF

Model Support

Threat Defense

Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for OSPF

Firewall Mode Guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

High Availability Guidelines

OSPFv2 and OSPFv3 support Stateful High Availability.

IPv6 Guidelines

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The threat defense device installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.

OSPFv3 Hello Packets and GRE

Typically, OSPF traffic does not pass through GRE tunnel. When OSPFv3 on IPv6 is encapsulated inside GRE, the IPv6 header validation for security check such as Multicast Destination fails. The packet is dropped due to the implicit security check validation, as this packet has destination IPv6 multicast.

You may define a pre-filter rule to bypass GRE traffic. However, with pre-filter rule, inner packets would not be interrogated by the inspection engine.

Clustering Guidelines

- OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment.
- In Spanned interface mode, dynamic routing is not supported on management-only interfaces.
- In Individual interface mode, make sure that you establish the control and data units as either OSPFv2 or OSPFv3 neighbors.
- In Individual interface mode, OSPFv2 adjacencies can only be established between two contexts on a shared interface on the control unit. Configuring static neighbors is supported only on point-to-point-links; therefore, only one neighbor statement is allowed on an interface.
- When a control role change occurs in the cluster, the following behavior occurs:
 - In spanned interface mode, the router process is active only on the control unit and is in a suspended state on the data units. Each cluster unit has the same router ID because the configuration has been synchronized from the control unit. As a result, a neighboring router does not notice any change in the router ID of the cluster during a role change.
 - In individual interface mode, the router process is active on all the individual cluster units. Each cluster unit chooses its own distinct router ID from the configured cluster pool. A control role change in the cluster does not change the routing topology in any way.

Multiprotocol Label Switching (MPLS) and OSPF Guidelines

When a MPLS-configured router sends Link State (LS) update packets containing opaque Type-10 link-state advertisements (LSAs) that include an MPLS header, authentication fails and the appliance silently drops the update packets, rather than acknowledging them. Eventually the peer router will terminate the neighbor relationship because it has not received any acknowledgments.

Make sure that non-stop forwarding (NSF) is disabled on the appliance to ensure that the neighbor relationship remains stable:

- Navigate to the **Non Stop Forwarding** page in management center(**Devices > Device Management (select the desired device) > Routing > OSPF > Advanced > Non Stop Forwarding**).

Ensure the **Non Stop Forwarding Capability** boxes are not checked.



Note The Firepower 4100/9300 models may have high latency when using MPLS because they lack load balancing across multiple receiving queues.

Bidirectional and Forwarding Detection (BFD) and OSPF Guidelines

- You can enable BFD on OSPFv2 and OSPFv3 interfaces (Physical Interfaces, Sub-Interfaces, and Port-Channels).
- BFD is not supported on VTI Tunnels, DVTI Tunnels, Loopback, Switchport, VNI, VTEP, and IRB interfaces.

Route Redistribution Guidelines

- Redistribution of route maps with IPv4 or IPv6 prefix list on OSPFv2 or OSPFv3 is not supported. Use an access list in the route map on OSPF for redistribution.
- When OSPF is configured on a device that is a part of EIGRP network or vice versa, ensure that OSPF-router is configured to tag the route (EIGRP does not support route tag yet).

When redistributing OSPF into EIGRP and EIGRP into OSPF, a routing loop occurs when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Additional Guidelines

- OSPFv2 and OSPFv3 support multiple instances on an interface.
- OSPFv3 supports encryption through ESP headers in a non-clustered environment.
- OSPFv3 supports Non-Payload Encryption.
- OSPFv2 supports Cisco NSF Graceful Restart and IETF NSF Graceful Restart mechanisms as defined in RFCs 4811, 4812 & 3623 respectively.
- OSPFv3 supports Graceful Restart mechanism as defined in RFC 5187.
- There is a limit to the number of intra area (type 1) routes that can be distributed. For these routes, a single type-1 LSA contains all prefixes. Because the system has a limit of 35 KB for packet size, 3000 routes result in a packet that exceeds the limit. Consider 2900 type 1 routes to be the maximum number supported.

- For a device using virtual routing, you can configure OSPFv2 and OSPFv3 for a global virtual router. However, you can configure only OSPFv2 for a user-defined virtual router.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.

Configure OSPFv2

This section describes the tasks involved in configuring an OSPFv2 routing process. For a device using virtual routing, you can configure OSPFv2 for global as well as for user-defined virtual routers.

Configure OSPF Areas, Ranges, and Virtual Links

You can configure several OSPF area parameters, which include setting authentication, defining stub areas, and assigning specific costs to the default summary route. You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area.

Procedure

-
- Step 1** Choose **Devices > Device Management**, and edit the threat defense device.
- Step 2** Click **Routing**.
- Step 3** (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.
- Step 4** Click **OSPF**.
- Step 5** Check the check box of **Process 1**. You can enable up to two OSPF process instances for each context/virtual router. You must choose an OSPF process to be able to configure the Area parameters.
- If the device is using virtual routing, the ID fields display the unique process IDs generated for the chosen virtual router.
- Step 6** Choose the **OSPF Role** from the drop-down list, and enter a description for it in the next field. The options are Internal, ABR, ASBR, and ABR & ASBR. See [About OSPF, on page 1](#) for a description of the OSPF roles.
- Step 7** Select **Area > Add**.
- You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.
- Step 8** Configure the following area options for each OSPF process:
- **OSPF Process**—Choose the process ID. For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
 - **Area ID**—Designation of the area for which routes are to be summarized.
 - **Area Type**—Choose one of the following:

- **Normal**—(Default) Standard OSPF area.
- **Stub**—A stub area does not have any routers or areas beyond it. Stub areas prevent Autonomous System (AS) External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Types 3 and 4) from being flooded into the area by NOT checking the **Summary Stub** check box.
- **NSSA**—Makes the area a not-so-stubby area (NSSA). NSSAs accept Type 7 LSAs. You can disable route redistribution by NOT checking the **Redistribute** check box and checking the **Default Information Originate** check box. You can prevent summary LSAs from being flooded into the area by NOT checking the **Summary NSSA** check box.
- **Metric Value**—The metric used for generating the default route. The default value is 10. Valid metric values range from 0 through 16777214.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Available Network**—Choose one of the available networks and click **Add**, or click **Add (+)** to add a new network object. See [Network](#) for the procedure for adding networks.
- **Authentication**—Choose the OSPF authentication:
 - **None**—(Default) Disables OSPF area authentication.
 - **Password**—Provides a clear text password for area authentication, which is not recommended where security is a concern.
 - **MD5**—Allows MD5 authentication.
- **Default Cost**—The default cost for the OSPF area, which is used to determine the shortest paths to the destination. Valid values range from 0 through 65535. The default value is 1.

Step 9 Click **OK** to save the area configuration.

Step 10 Select **Range > Add**.

- Choose one of the available networks and whether to advertise, or,
- Click **Add (+)** to add a new network object. See [Network](#) for the procedure for adding networks.

Step 11 Click **OK** to save the range configuration.

Step 12 Select **Virtual Link**, click **Add (+)**, and configure the following options for each OSPF process:

- **Peer Router**—Choose the IP address of the peer router. To add a new peer router, click **Add (+)**. See [Network](#) for the procedure for adding networks.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. Valid values range from 1 through 65535. The default is 10.

The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.

- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 through 8192. The default is 1.

LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 through 65535. The default is 5.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The dead interval is an unsigned integer. The default is four times the hello interval, or 40 seconds. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 through 65535.

- **Authentication**—Choose the OSPF virtual link authentication from the following:

- **None**—(Default) Disables virtual link area authentication.
- **Area Authentication**—Enables area authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
- **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
- **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

Note Ensure to enter only numbers as the MD5 key ID.

- **Key Chain**—Allows key chain authentication. Click **Add**, and create the key chain, and then click **Save**. For detailed procedure, see [Creating Key Chain Objects](#). Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.

Step 13 Click **OK** to save the virtual link configuration.

Step 14 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPF Redistribution](#).

Configure OSPF Redistribution

The threat defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can

redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

-
- Step 1** Choose **Devices > Device Management**, and edit the threat defense device.
 - Step 2** Click **Routing**.
 - Step 3** (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.
 - Step 4** Click **OSPF**.
 - Step 5** From **OSPF Role** drop-down, choose role .
 - Step 6** Click **Redistribution > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.

- Step 7** Configure the following redistribution options for each OSPF process:
 - **OSPF Process**—Choose the process ID. For a device using virtual routing, this drop-down list displays the unique process IDs generated for the selected virtual router.
 - **Route Type**—Choose one of the following types:
 - **Static**—Redistributes static routes to the OSPF routing process.
 - **Connected**—Redistributes connected routes (routes established automatically by virtue of having the IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the device. You can select whether to use subnets under the Optional list.
 - **OSPF**—Redistributes routes from another OSPF routing process, for example, internal, external 1 and 2, NSSA external 1 and 2, or whether to use subnets. You can select these options under the Optional list.
 - **BGP**—Redistribute routes from the BGP routing process. Add the AS number and whether to use subnets.
 - **RIP**—Redistributes routes from the RIP routing process. You can select whether to use subnets under the Optional list.

Note As a user-defined virtual router does not support RIP, you cannot redistribute routes from RIP.

- **EIGRP**—Redistribute routes from the EIGRP routing process. Add the AS number and whether to use subnets.
- **Metric Value**—Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.

- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag Value**—Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **RouteMap**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** (+). See [Configure Route Map Entry](#) to add a new route map.

Step 8 Click **OK** to save the redistribution configuration.

Step 9 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPF Inter-Area Filtering, on page 11](#).

Configure OSPF Inter-Area Filtering

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **InterArea > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete inter-areas.

Step 6 Configure the following inter-area filtering options for each OSPF process:

- **OSPF Process**—For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Area ID**—The area for which routes are to be summarized.
- **PrefixList**—The name of the prefix. To add a new prefix list object, see Step 5.
- **Traffic Direction**—Inbound or outbound. Choose Inbound to filter LSAs coming into an OSPF area, or Outbound to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.

Step 7 Click **Add** (+), and enter a name for the new prefix list, and whether to allow overrides.

You must configure a prefix list before you can configure a prefix rule.

Step 8 Click **Add** to configure prefix rules, and configure the following parameters:

- **Action**—Select **Block** or **Allow** for the redistribution access.
- **Sequence No**—The routing sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.
- **IP Address**—Specify the prefix number in the format of IP address/mask length.
- **Min Prefix Length**—(Optional) The minimum prefix length.
- **Max Prefix Length**—(Optional) The maximum prefix length.

Step 9 Click **OK** to save the inter-area filtering configuration.

Step 10 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPF Filter Rules, on page 12](#).

Configure OSPF Filter Rules

You can configure ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. You can apply this type of area filtering out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF area at the same time. OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Filter Rule > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete filter rules.

Step 6 Configure the following filter rule options for each OSPF process:

- **OSPF Process**—For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Access List**—The access list for this OSPF process. To add a new standard access list object, click **Add** (+) and see [Configure Standard ACL Objects](#).
- **Traffic Direction**—Choose In or Out for the traffic direction being filtered. Choose In to filter LSAs coming into an OSPF area, or Out to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- **Interface**—The interface for this filter rule.

Step 7 Click **OK** to save the filter rule configuration.

Step 8 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPF Summary Addresses, on page 13](#).

Configure OSPF Summary Addresses

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the threat defense device to advertise a single route for all the redistributed routes that are included for a specified network address and mask. This configuration decreases the size of the OSPF link-state database. Routes that match the specified IP address mask pair can be suppressed. The tag value can be used as a match value for controlling redistribution through route maps.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Summary Address > Add**.

You can click **Edit** (✎) to edit, or use the right-click menu to cut, copy, past, insert, and delete summary addresses.

Step 6 Configure the following summary address options for each OSPF process:

- **OSPF Process**—For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Available Network**—The IP address of the summary address. Select one from the Available networks list and click **Add**, or to add a new network, click **Add** (+). See [Network](#) for the procedure for adding networks.
- **Tag**—A 32-bit decimal value that is attached to each external route. This value is not used by OSPF itself, but may be used to communicate information between ASBRs.
- **Advertise**—**Advertises** the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 7 Click **OK** to save the summary address configuration.

Step 8 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPF Interfaces and Neighbors, on page 14](#).

Configure OSPF Interfaces and Neighbors

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval, the dead interval, and the authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

You need to define static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv2 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Interface > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Configure the following Interface options for each OSPF process:

- **Interface**—The interface you are configuring.

Note If the device is using virtual routing, this drop-down list displays only those interfaces that belong to the router.

- **Default Cost**—The cost of sending a packet through the interface. The default value is 10.

- **Priority**—The designated router for a network. Valid values range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router.

When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router. This setting does not apply to interfaces that are configured as point-to-point interfaces.

- **MTU Ignore**—OSPF checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency is not established.

- **Database Filter**—Use this setting to filter the outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this flooding can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents OSPF flooding of the LSA on the selected interface.

- **Hello Interval**—Specifies the interval, in seconds, between hello packets sent on an interface. Valid values range 1–8192 seconds. The default value is 10 seconds.

The smaller the hello interval, the faster topological changes are detected, but more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface.

- **Transmit Delay**—Estimated time in seconds to send an LSA packet on the interface. Valid values range 1–65535 seconds. The default is 1 second.

LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range 1–65535.

- **Hello Multiplier**—Specifies the number of Hello packets to be sent per second. Valid values are 3–20.

- **Point-to-Point**—Lets you transmit OSPF routes over VPN tunnels.

- **Authentication**—Choose the OSPF interface authentication from the following:

- **None**—(Default) Disables interface authentication.
- **Area Authentication**—Enables interface authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
- **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
- **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

Note Ensure to enter only numbers as the MD5 key ID.

- **Key Chain**—Allows key chain authentication. Click **Add**, and create the key chain, and then click **Save**. For detailed procedure, see [Creating Key Chain Objects](#). Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.

- **Enable BFD**—Allows you to enable BFD on this interface.
- **Enter Password**—The password you configure if you choose Password as the type of authentication.
- **Confirm Password**—Confirm the password that you chose.

Step 7 Select **Neighbor > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 8 Configure the following parameters for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Neighbor**—Choose one of the neighbors in the drop-down list, or click **Add** (+) to add a new neighbor; enter the name, description, network, whether to allow overrides, and then click **Save**.
- **Interface**—Choose the interface associated with the neighbor.

Step 9 Click **OK** to save the neighbor configuration.

Step 10 Click **Save** on the Routing page to save your changes.

Configure OSPF Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, an LSA timer, and graceful restarts.

Graceful Restarts

The threat defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623).



Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

- Step 1** Choose **Devices > Device Management**, and edit the threat defense device.
- Step 2** Click **Routing**.
- Step 3** (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.
- Step 4** Click **OSPF > Advanced**.
- Step 5** Select **General**, and configure the following:
 - **Router ID**—Choose Automatic or IP Address (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode) for the router ID. If you choose IP address, enter the IP address in the adjacent field. If you choose Cluster Pool, choose the IPv4 cluster pool value in the adjacent drop-down field. For information on creating the cluster pool address, see [Address Pools](#).
 - **Ignore LSA MOSPF**—Suppresses syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.
 - **RFC 1583 Compatible**—Configures RFC 1583 compatibility as the method used to calculate summary route costs. Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. All OSPF routers in an OSPF routing domain should have RFC compatibility set identically.
 - **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.
By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.
 - **Log Adjacency Changes**—Causes the threat defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
 - **Log Adjacency Change Details**—Causes the threat defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distance**—Allows you to modify the settings that were used to configure administrative route distances for **inter-area**, **intra-area**, and **external** IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **LSA Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Enable Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPF routing domain and configure the following options:
 - **Always advertise the default route**—Ensures that the default route is always advertised.
 - **Metric Value**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
 - **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
 - **RouteMap**—Choose the routing process that generates the default route if the route map is satisfied or click **Add (+)** to add a new one. See [Configure Route Map Entry](#) to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Non Stop Forwarding**, and configure Cisco NSF graceful restart for OSPFv2, for an NSF-capable or NSF-aware device:

Note There are two graceful restart mechanisms for OSPFv2, Cisco NSF and IETF NSF. Only one of these graceful restart mechanisms can be configured at a time for an OSPF instance. An NSF-aware device can be configured as both Cisco NSF helper and IETF NSF helper but a NSF-capable device can be configured in either Cisco NSF or IETF NSF mode at a time for an OSPF instance.

- a) Check the **Enable Cisco Non Stop Forwarding Capability** check box.
- b) (Optional) Check the **Cancel NSF restart when non-NSF-aware neighboring networking devices are detected** check box if required.
- c) (Optional) Make sure the **Enable Cisco Non Stop Forwarding Helper** mode check box is unchecked to disable the helper mode on an NSF-aware device.

Step 8 Configure IETF NSF Graceful Restart for OSPFv2, for an NSF-capable or NSF-aware device:

- a) Check the **Enable IETF Non Stop Forwarding Capability** check box.
- b) In the **Length of graceful restart interval (seconds)** field, enter the restart interval in seconds. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
- c) (Optional) Make sure the **Enable IETF nonstop forwarding (NSF) for helper mode** check box is unchecked to disable the IETF NSF helper mode on an NSF-aware device.
- d) **Enable Strict Link State advertisement checking**—When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.
- e) **Enable IETF Non Stop Forwarding**—Enables non stop forwarding, which allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. OSPF uses extensions to the OSPF protocol to recover its state from neighboring OSPF devices. For the recovery to work, the neighbors must support the NSF protocol extensions and be

willing to act as "helpers" to the device that is restarting. The neighbors must also continue forwarding data traffic to the device that is restarting while protocol state recovery takes place.

Configure OSPFv3

This section describes the tasks involved in configuring an OSPFv3 routing process. For a device using virtual routing, you can configure OSPFv3 only for its global virtual router and not for its user-defined virtual router.

Configure OSPFv3 Areas, Route Summaries, and Virtual Links

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, and then redistribute the route into the targeted OSPFv3 routing process.

Procedure

- Step 1** Choose **Devices > Device Management**, and edit the threat defense device.
- Step 2** Select **Routing > OSPFv3**.
- Step 3** By default **Enable Process 1** is selected. You can enable up to two OSPF process instances.
- Step 4** Choose the OSPFv3 role from the drop-down list, and enter a description for it. The options are Internal, ABR, ASBR, and ABR and ASBR. See [About OSPF, on page 1](#) for descriptions of the OSPFv3 roles.
- Step 5** Select **Area > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.
- Step 6** Select **General**, and configure the following options for each OSPF process:
 - **Area ID**—The area for which routes are to be summarized.
 - **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
 - **Type**—Specifies Normal, NSSA, or Stub. If you select Normal, there are no other parameters to configure. If you select Stub, you can choose to send summary LSAs in the area. If you select NSSA, you can configure the next three options:
 - **Allow Sending summary LSA into this area**—Allows the sending of summary LSAs into the area.
 - **Imports routes to normal and NSSA area**—Allows redistribution to import routes to normal and not to stubby areas.
 - **Defaults information originate**—Generates a default external route into an OSPFv3 routing domain.
 - **Metric**—Metric used for generating the default route. The default value is 10. Valid metric values range from 0 to 16777214.

- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.

Step 7 Click **OK** to save the general configuration.

Step 8 (Not applicable for Internal OSPFv3 Role) Select **Route Summary > Add Route Summary**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete route summaries.

Step 9 Configure the following route summary options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix. To add a new network object, click **Add** (+). See [Network](#) for the procedure for adding networks.
- **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Advertise**—Advertises the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 10 Click **OK** to save the route summary configuration.

Step 11 (Not applicable for Internal OSPFv3 Role) Select **Virtual Link**, click **Add Virtual Link**, and configure the following options for each OSPF process:

- **Peer RouterID**—Choose the IP address of the peer router. To add a new network object, click **Add** (+). See [Network](#) for the procedure for adding networks.
- **TTL Security**—Enables TTL security check. The value for the hop-count is a number from 1 to 254. The default is 1.

OSPF sends outgoing packets with an IP header Time to Live (TTL) value of 255 and discards incoming packets that have TTL values less than a configurable threshold. Because each device that forwards an IP packet decrements the TTL, packets received via a direct (one-hop) connection have a value of 255. Packets that cross two hops have a value of 254, and so on. The receive threshold is configured in terms of the maximum number of hops that a packet may have traveled.
- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The default is four times the hello interval, or 40 seconds. Valid values range from 1 to 65535.

The dead interval is an unsigned integer. The value must be the same for all routers and access servers that are attached to a common network.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. Valid values range from 1 to 65535. The default is 10.

The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 to 65535. The default is 5.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 to 8192. The default is 1.

LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

Step 12 Click **OK** to save the virtual link configuration.

Step 13 Click **Save** on the Router page to save your changes.

What to do next

Continue with [Configure OSPFv3 Redistribution](#).

Configure OSPFv3 Redistribution

The Secure Firewall Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Select **Routing > OSPF**.

Step 3 Select **Redistribution**, and click **Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following redistribution options for each OSPF process:

- **Source Protocol**—The source protocol from which routes are being redistributed. The supported protocols are connected, OSPF, Static, EIGRP, and BGP. If you choose OSPF, you must enter the Process ID in the **Process ID** field. If you choose BGP, you must add the AS number in the **AS Number** field.
- **Metric** —Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.

- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag** —Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **Route Map**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add (+)**. See [Route Map](#) for the procedure to add a new route map.
- **Process ID**—The OSPF process ID, either 1 or 2.

Note The Process ID is enabled the OSPFv3 process is redistributing a route learned by another OSPFv3 process.
- **Match**—Enables OSPF routes to be redistributed into other routing domains:
 - **Internal** for routes that are internal to a specific autonomous system.
 - **External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 1 external routes.
 - **External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 2 external routes.
 - **NSSA External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 1 external routes.
 - **NSSA External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 2 external routes.

Step 5 Click **OK** to save the redistribution configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPFv3 Summary Prefixes, on page 22](#).

Configure OSPFv3 Summary Prefixes

You can configure the threat defense device to advertise routes that match a specified IPv6 prefix and mask pair.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Select **Routing > OSPFv3**.

Step 3 Select **Summary Prefix > Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete summary prefixes.

Step 4 Configure the following summary prefix options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix and prefix length label. Select one from the list or click **Add** (+) to add a new network object. See [Network](#) for the procedure for adding networks.
- **Advertise**—Advertises routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.
- (Optional) **Tag**—A value that you can use as a match value for controlling redistribution through route maps.

Step 5 Click **OK** to save the summary prefix configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with [Configure OSPFv3 Interfaces, Authentication, and Neighbors, on page 23](#).

Configure OSPFv3 Interfaces, Authentication, and Neighbors

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

Procedure

Step 1 Choose **Devices > Device Management**, and edit the threat defense device.

Step 2 Select **Routing > OSPFv3**.

Step 3 Select **Interface > Add**.

You can click **Edit** to edit, or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following interface options for each OSPFv3 process:

- **Interface**—The interface you are configuring.
- **Enable OSPFv3**—Enables OSPFv3.
- **OSPF Process**—Choose 1 or 2.
- **Area**—The area ID for this process.
- **Instance**—Specifies the area instance ID to be assigned to the interface. An interface can have only one OSPFv3 area. You can use the same area on multiple interfaces, and each interface can use a different area instance ID.

Step 5

Select **Properties**, and configuring the following options for each OSPFv3 process:

- **Filter Outgoing Link Status Advertisements**—Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.
- **Disable MTU mismatch detection**—Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.
- **Flood Reduction**—Changes normal LSAs into Do Not Age LSAs, so that they don't get flooded every 3600 seconds across areas.

OSPF LSAs are refreshed every 3600 seconds. In large OSPF networks, this can lead to large amounts of unnecessary LSA flooding from area to area.

- **Point-to-Point Network**—Lets you transmit OSPF routes over VPN tunnels. When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:
 - You can define only one neighbor for the interface.
 - You need to manually configure the neighbor.
 - You need to define a static route pointing to the crypto endpoint.
 - If OSPF over a tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
 - You should bind the crypto map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so that the OSPF adjacencies can be established over the VPN tunnel.
- **Broadcast**— Specifies that the interface is a broadcast interface. By default, this check box is checked for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, nonbroadcast interface. Specifying an interface as point-to-point, nonbroadcast lets you transmit OSPF routes over VPN tunnels.
- **Cost**—Specifies the cost of sending a packet on the interface. Valid values for this setting range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, nonbroadcast interfaces.

When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
- **Priority**—Determines the designated router for a network. Valid values range from 0 to 255.
- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range from 1 to 65535.
- **Hello Interval**— Time period in seconds between OSPF packets that the router will send before adjacency is established with a neighbor. Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval. Valid values range from 1 to 65535 seconds.

- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.
- **Transmit Delay**—Estimated time in seconds to send a link-state update packet on the interface. Valid values range from 1 to 65535 seconds. The default is 1 second.
- **Enable BFD**—Allows you to enable BFD on this interface.

Step 6 Click **OK** to save the properties configuration.

Step 7 Select **Authentication**, and configure the following options for each OSPFv3 process:

- **Type**—Type of authentication. The available options are Area, Interface, and None. The None option indicates that no authentication is used.
- **Security Parameters Index**— A number from 256 to 4294967295. Configure this if you chose Interface as the type.
- **Authentication**—Type of authentication algorithm. Supported values are SHA-1 and MD5. Configure this if you chose Interface as the type.
- **Authentication Key**— When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.
- **Encrypt Authentication Key**—Enables encryption of the authentication key.
- **Include Encryption**— Enables encryption.
- **Encryption Algorithm**—Type of encryption algorithm. Supported value is DES. The NULL entry indicates no encryption. Configure this if you chose **Include Encryption**.
- **Encryption Key**—Enter the encryption key. Configure this if you chose **Include Encryption**.
- **Encrypt Key**—Enables the key to be encrypted.

Step 8 Click **OK** to save the authentication configuration.

Step 9 Select **Neighbor**, click **Add**, and configure the following options for each OSPFv3 process:

- **Link Local Address**—The IPv6 address of the static neighbor.
- **Cost**—Enables cost. Enter the cost in the **Cost** field, and check the **Filter Outgoing Link State Advertisements** if you want to advertise.
- (Optional) **Poll Interval**—Enables the poll interval. Enter the **Priority** level and the **Poll Interval** in seconds.

Step 10 Click **Add** to add the neighbor.

Step 11 Click **OK** to save the Interface configuration.

Configure OSPFv3 Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, passive OSPFv3 routing, LSA timers, and graceful restarts.

Graceful Restarts

The threat defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).



Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

-
- Step 1** Choose **Devices > Device Management**, and edit the threat defense device.
- Step 2** Choose **Routing > OSPFv3 > Advanced**.
- Step 3** For **Router ID**, choose Automatic or IP Address (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode). If you choose IP Address, enter the IPv6 address in the **IP Address** field. If you choose Cluster Pool, choose the IPv6 cluster pool value from the **Cluster Pool** down-down field. For information on creating the cluster pool address, see [Address Pools](#).
- Step 4** Check the **Ignore LSA MOSPF** check box if you want to suppress syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.
- Step 5** Select **General**, and configure the following:
- **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.
By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.
 - **Adjacency Changes**—Causes the threat defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
 - **Include Details**—Causes the threat defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.
 - **Administrative Route Distances**—Allows you to modify the settings that were used to configure administrative route distances for inter-area, intra-area, and external IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.

- **Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPFv3 routing domain and configure the following options:
 - **Always Advertise**—Will always advertise the default route whether or not one exists.
 - **Metric**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
 - **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
 - **Route Map**—Choose the routing process that generates the default route if the route map is satisfied or click **Add** (+) to add a new one. See [Route Map](#) to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Passive Interface**, select the interfaces on which you want to enable passive OSPFv3 routing from the Available Interfaces list, and click **Add** to move them to the Selected Interfaces list.

Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface.

Step 8 Click **OK** to save the passive interface configuration.

Step 9 Select **Timer**, and configure the following LSA pacing and SPF calculation timers:

- **Arrival**—Specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6000,000 milliseconds. The default is 1000 milliseconds.
- **Flood Pacing**—Specifies the time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds.
- **Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Retransmission Pacing**—Specifies the time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds.
- **LSA Throttle**—Specifies the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 millisecond. The minimum specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds. The maximum specifies the maximum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.

Note For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- **SPF Throttle**—Specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds. The minimum specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds. The maximum specifies the maximum wait time in milliseconds for SPF calculations. The default value is 10000 milliseconds.

Note For SPF throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- Step 10** Click **OK** to save the LSA timer configuration.
- Step 11** Select **Non Stop Forwarding**, and check the **Enable graceful-restart helper** check box. This is checked by default. Uncheck this to disable the graceful-restart helper mode on an NSF-aware device.
- Step 12** Check the **Enable link state advertisement** check box to enable strict link state advertisement checking.
- When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.
- Step 13** Check the **Enable graceful-restart (Use when Spanned Cluster or Failover Configured)** and enter the graceful-restart interval in seconds. The range is 1-1800. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
- Step 14** Click **OK** to save the graceful restart configuration.
- Step 15** Click **Save** on the Routing page to save your changes.

History for OSPF

Table 1: Feature History for OSPF

Feature	Minimum Management Center	Minimum Threat Defense	Details
BFD Support for OSPF v2 and v3	7.4	7.4	<p>You can enable BFD on OSPFv2 and OSPFv3 interfaces.</p> <p>New/modified screens:</p> <ul style="list-style-type: none"> • Configuration > Device Setup > Routing > OSPFv2 • Configuration > Device Setup > Routing > OSPFv3