

OSPF

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

- OSPF routing, on page 1
- Guidelines for OSPF, on page 5
- Configure OSPFv2, on page 8
- Configure OSPFv2 Router ID, on page 10
- Customize OSPFv2, on page 11
- Configure OSPFv3, on page 29
- Configure Graceful Restart, on page 39
- Example for OSPFv2, on page 43
- Examples for OSPFv3, on page 45
- Monitoring OSPF, on page 46
- History for OSPF, on page 47

OSPF routing

OSPF is an interior gateway routing protocol that

- uses link states rather than distance vectors for path selection,
- propagates link-state advertisements rather than routing table updates, and
- converges more quickly than RIP networks because only LSAs are exchanged instead of the entire routing tables.

OSPF characteristics and features

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 these characteristics:

- 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 ASA 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 ASA 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 ASA supports these 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 ASA AS a designated router or a designated backup router. The ASA also can be set up AS an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

OSPF supports both 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 ASA 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 ASA. Also, you should not mix public and private networks on the same ASA interface.

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

OSPF support for fast hello packets

OSPF support for fast hello packets is a feature that

- provides a way to configure the sending of hello packets in intervals less than one second, and
- results in faster convergence in an Open Shortest Path First (OSPF) network.

OSPF fast hello packets configuration prerequisites

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

About OSPF Support for Fast Hello Packets

The key concepts related to OSPF support for fast hello packets and the benefits of OSPF Fast Hello Packets are described below:

OSPF hello interval and dead interval

OSPF hello interval and dead interval are configurable time parameters that

- maintain connectivity between OSPF neighbors through regular hello packet exchanges,
- detect neighbor failures when hello packets are not received within the specified timeframes, and
- must be consistent across all routers within the same network to ensure proper operation.

Hello interval and dead interval configuration details

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 are hello packets that

- are sent at intervals of less than 1 second,
- are achieved by using the OSPF dead-interval command with a dead interval set to 1 second, and
- use a hello-multiplier value set to the number of hello packets sent during that 1 second.

Fast hello packet behavior

When fast hello packets are configured on the interface, these behaviors occur:

- 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.

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.

Benefits of OSPF fast hello packets

A fast hello packet is an OSPF protocol feature that

- reduces network convergence time by enabling routers to detect changes in the network faster,
- allows rapid detection of lost neighbors within one second, and
- enhances reliability in LAN environments where neighbor loss may not be detected at the Open Systems Interconnection (OSI) physical or data-link layer.

The benefit of OSPF Fast Hello Packets is that your OSPF network can react swiftly to changes, minimizing downtime and improving overall network stability. This feature is particularly useful in scenarios where routers may not detect lost neighbors via traditional methods, allowing OSPF to maintain accurate routing tables.

If a router on a LAN segment is powered off or disconnected and the physical and data-link layers do not signal neighbor loss, OSPF fast hello packets will detect the neighbor's disappearance within one second.

OSPFv2 and OSPFv3 implementation differences

A protocol distinction is a protocol category that

- prevents backward compatibility between OSPFv2 and OSPFv3,
- requires both OSPF versions to run simultaneously for IPv4 and IPv6 routing, and
- allows OSPFv2 and OSPFv3 to coexist without interacting with each other.

OSPFv3 additional features

OSPFv3 provides additional features that include:

- 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.

Guidelines for OSPF

Context mode guidelines

OSPFv2 supports single and multiple context mode.

- OSPFv2 instances cannot form adjacencies with each other across shared interfaces because, by default, inter-context exchange of multicast traffic is not supported across shared interfaces. However, you can use the static neighbor configuration under OSPFv2 process configuration under OSPFv2 process to bring up OSPFv2 neighbourship on a shared interface.
- Inter-context OSPFv2 on separate interfaces is supported.
- **(For point-to-point topology)** In a multiple context mode, you can configure static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. For a successful OSPF adjacency formation, when you remove and reconfigure OSPF on a shared interface, execute the commands in this order: Remove the OSPF configuration using **no router <OSPF name>**, and then remove the OSPF point-to-point non-broadcast configuration on the interface using **no OSPF network point-to-point non-broadcast**. When you reconfigure OSPF on the shared interface, configure the OSPF router using **router <OSPF name>**, and then configure the interface with **OSPF network point-to-point non-broadcast**.
- OSPFv3 supports single mode only.

Key chain authentication guidelines

OSPFv2 supports key chain authentication on both single and multiple mode, both in physical and virtual modes. However, in multiple mode, you can configure the key chain only in context mode.

- The rotating keys are applicable only for OSPFv2 protocol. OSPF area authentication with key chain is not supported.
- The existing MD5 authentication without time range in OSPFv2 is still supported along with new rotating keys.
- Though the platform supports SHA1 and MD5 cryptographic algorithms, only MD5 cryptographic algorithm is used for authentication.

Firewall mode guidelines

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

Failover Guidelines

OSPFv2 and OSPFv3 support Stateful Failover.

IPv6 guidelines

OSPFv2 does not support IPv6, while OSPFv3 provides IPv6 support with specific authentication and routing capabilities.

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The ASA installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.
- OSPFv3 packets can be filtered out using IPv6 ACLs in the **capture** command.

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

OSPFv2 and OSPFv3 have specific limitations and behaviors in clustering environments.

- 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, specific behavior occurs based on interface mode. 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.

- Disable the opaque capability on the ASA to ensure that the neighbor relationship remains stable:

```
router ospf process_ID_number
no nsf ietf helper
no capability opaque
```

These guidelines apply to MPLS environments where opaque LSAs might cause authentication failures.



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 specific interface types with OSPFv2 and OSPFv3.

- 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

Route redistribution with prefix lists and between routing domains requires specific configuration considerations.

- Redistribution of route maps with IPv4 prefix list on OSPFv2 is supported. However, redistribution of route maps with IPv6 prefix list on 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 various features and have specific limitations in different deployment scenarios.

- 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.

- 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.
- The ASA virtual cannot use dynamic interior routing protocols like EIGRP and OSPF due to the nature of Azure cloud routing. The Effective Routing Table determines the next hop, regardless of whether a virtual client has any static/dynamic route configured. Currently, you cannot view either the Effective Routing Table or the System Routing Table.
- OSPFv3 drops the LS update if the packet size exceeds 8190. As a result, the neighbourship is terminated. Hence, configure the switch with the "ospfv3 MTU-ignore" command to avoid neighbourship termination.

Configure OSPFv2

This section describes how to enable an OSPFv2 process on the ASA.

After you enable OSPFv2, you need to define a route map. For more information, see Define a Route Map. Then you generate a default route. For more information, see Configure a Static Route.

After you have defined a route map for the OSPFv2 process, you can customize it for your particular needs, To learn how to customize the OSPFv2 process on the ASA, see Customize OSPFv2, on page 11.

To enable OSPFv2, you need to create an OSPFv2 routing process, specify the range of IP addresses associated with the routing process, then assign area IDs associated with that range of IP addresses.

You can enable up to two OSPFv2 process instances. Each OSPFv2 process has its own associated areas and networks.

To enable OSPFv2, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

In the OSPF Setup pane, you can enable OSPF processes, configure OSPF areas and networks, and define OSPF route summarization.

Step 2 The three tabs in ASDM used to enable OSPF are as follows:

- The Process Instances tab allows you to enable up to two OSPF process instances for each context. Single context mode and multiple context mode are both supported. After you check the **Enable Each OSPF Process** check box, you can enter a unique identifier numeric identifier for that OSPF process. This process ID is used internally and does not need to match the OSPF process ID on any other OSPF devices; valid values range from 1 to 65535. Each OSPF process has its own associated areas and networks.

  If you click **Advanced**, the Edit OSPF Process Advanced Properties dialog box appears. From here, you can configure the Router ID, cluster IP address pools in Spanned EtherChannel or Individual Interface clustering, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for each OSPF process. You can enable BFD on all interfaces that support OSPFv2 here, or enable BFD for specific OSPFv2 interfaces (see Configure OSPFv2 Interface Parameters, on page 16).
- The Area/Networks tab allows you to display the areas and the networks that they include for each OSPF process on the ASA. From this tab you can display the area ID, the area type, and the type of authentication

set for the area. To add or edit the OSPF area or network, see Configure OSPFv2 Area Parameters, on page 19 for more information.

- The Route Summarization tab allows you to configure an ABR. In OSPF, an ABR will advertise networks in one area into another area. If the network numbers in an area are assigned in a way so that they are contiguous, you can configure the ABR to advertise a summary route that includes all the individual networks within the area that fall into the specified range. See Configure Route Summarization Between OSPFv2 Areas, on page 15 for more information.

Configure a Key Chain for Authentication

To enhance data security and protection of devices, you can enable rotating keys for authenticating IGP peers. The rotating keys prevent any malicious user from guessing the keys used for routing protocol authentication and thereby protecting the network from advertising incorrect routes and redirecting traffic. Changing the keys frequently reduces the risk of them eventually being guessed. When configuring authentication for routing protocols that provide key chains, configure the keys in a key chain to have overlapping lifetimes. This helps to prevent loss of key-secured communication due to absence of an active key. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with the peers.

This section describes how to create a key chain for OSPF peer authentication. This section also covers steps to add or edit the key chain attributes. After configuring a key chain object, you can use it in defining the OSPFv2 authentication for an interface and for a virtual link. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency. To learn how to define authentication for an interface, see Configure OSPFv2 Interface Parameters, on page 16; for a virtual link, see Configure a Virtual Link in OSPF, on page 27.

To configure a key chain, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Key Chain** .

Step 2 In the **Configure Key Chain** section, click **Add**.

Step 3 Enter the key chain name in the **Add Key Chain** dialog box, and click **Ok**.

The created key chain name is listed in the **Configure Key Chain** grid.

Step 4 Select the key chain name from the **Configure Key Chain** section, and in the **Configure Key** section, click **Add**. To edit an existing key, seclect the key name and click **Edit**.

The **Add Key** or **Edit Key**dialog box appears, depending on the action that you have selected.

Step 5 Specify the key identifier in the **Key ID** field.

The key id value can be between 0 and 255. Use the value 0 only when you want to signal an invalid key.

Note
You cannot edit a saved key id.

Step 6 From the **Cryptographic Algorithm** drop-down, choose **MD5**. MD5 is the only algorithm supported for authenticating the key chain.

Step 7 Select the encryption type by clicking the **Plain Text** or **Encrypted** radio button, and then enter the password in the **Authentication Key** field.

- Even though the key-string can accommodate 80 characters, only a maximum of 64 characters is permitted for OSPF authentication.
- The passwords cannot be a single digit nor those starting with a digit followed by a white space. For example, "0 pass" or "1" are invalid.

Step 8 Provide the lifetime values in the **Accept Lifetime** and **Send Lifetime** fields:

You can specify the time interval for the device to accept/send the key during key exchange with another device. The end time can be the duration, the absolute time when the accept/send lifetime ends, or never expires.

Following are the validation rules for the start and end values:

- Start lifetime cannot be null when the end lifetime is specified.
- The start lifetime for accept or send lifetime must be earlier than the end lifetime.

Step 9 To save the key chain attibutes, click **Ok**. In the **Key Chain** page, click **Appy**.

What to do next

You can now apply the configured key chain to define the OSPFv2 authentication for an interface and for virtual link.

- Configure OSPFv2 Interface Parameters, on page 16
- Configure a Virtual Link in OSPF, on page 27

Configure OSPFv2 Router ID

The OSPF Router-ID is used to identify a specific device within an OSPF database. No two routers in an OSPF system can have the same router-id.

If a router-id is not configured manually in the OSPF routing process the router will automatically configure a router-id determined from the highest IP address of an active interface. When configuring a router-id, the neighbors will not be updated automatically until that router has failed or the OSPF process has been cleared and the neighbor relationship has been re-established.

Manually Configure OSPF Router-ID

This section describes how to manually configure router-id in OSPFv2 process on the ASA.

Procedure

Step 1 To use a fixed router ID, use the **router-id** command.

router-id *ip-address*

Example:

```
ciscoasa(config-router)# router-id 193.168.3.3
```

Step 2 To revert to the previous OSPF router ID behavior, use the **no router-id** command.

no router-id *ip-address*

Example:

```
ciscoasa(config-router)# no router-id 193.168.3.3
```

Router ID Behaviour while Migrating

While migrating OSPF configuration from one ASA , say ASA 1 to another ASA, say ASA 2, the following router id selection behaviour is observed:

1. ASA 2 does not use any IP address for OSPF router-id when all interfaces are in shutdown mode. The possibilities for configuring router-id when all interfaces are in "admin down" state or shutdown mode are:
   - If ASA 2 does not have any router-id configured before, you would see this message:

     ```
     %OSPF: Router process 1 is not running, please configure a router-id
     ```

     After the first interface is brought up, ASA 2 will take IP address of this interface as router id.
   - If ASA 2 had router-id configured before and all interfaces were in "admin down" state when "no router-id" command was issued, ASA 2 will use old router id. ASA 2 uses the old router id, even if IP addresses on the interface that is brought up is changed, until "clear ospf process" command is issued.
2. ASA 2 uses new router id, when ASA 2 had router-id configured before and at least one of interfaces were not in "admin down" state or shutdown mode when "no router-id" command was issued. ASA 2 will use new router id from the IP address of the interfaces even when interfaces are in "down/down" state.

Customize OSPFv2

This section explains how to customize the OSPFv2 processes.

Redistribute Routes Into OSPFv2

The ASA can control the redistribution of routes between OSPFv2 routing processes.

Note If you want to redistribute a route by defining which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process, you must first generate a default route. See Configure a Static Route, and then define a route map according to Define a Route Map.

To redistribute static, connected, RIP, or OSPFv2 routes into an OSPFv2 process, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Redistribution**.

The Redistribution pane displays the rules for redistributing routes from one routing process into an OSPF routing process. You can redistribute routes discovered by RIP and OSPF into the EIGRP routing process. You can also redistribute static and connected routes into the EIGRP routing process. You do not need to redistribute static or connected routes if they fall within the range of a network that has been configured through the Setup > Networks tab.

Step 2 Click **Add** or **Edit**.

Alternatively, double-clicking a table entry in the Redistribution pane (if any) opens the Add/Edit OSPF Redistribution Entry dialog box for the selected entry.

Note
All steps that follow are optional.

The Add/Edit OSPF Redistribution Entry dialog box lets you add a new redistribution rule or edit an existing redistribution rule in the Redistribution table. Some of the redistribution rule information cannot be changed when you are editing an existing redistribution rule.

Step 3 Choose the OSPF process associated with the route redistribution entry. If you are editing an existing redistribution rule, you cannot change this setting.

Step 4 Choose the source protocol from which the routes are being redistributed. You can choose one of the following options:

- Static—Redistributes static routes to the OSPF routing process.
- Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS.
- OSPF—Redistributes routes from another OSPF routing process. Choose the OSPF process ID from the list. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or EIGRP routes. Skip to Step 5.
- RIP—Redistributes routes from the RIP routing process.
- BGP— Redistribute routes from the BGP routing process.

- EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list.

Step 5 If you have chosen OSPF for the source protocol, choose the conditions used for redistributing routes from another OSPF routing process into the selected OSPF routing process. These options are not available when redistributing static, connected, RIP, or EIGRP routes. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions:

- Internal—The route is internal to a specific AS.
- External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
- NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Step 6 In the **Metric Value** field, enter the metric value for the routes being redistributed. Valid values range from 1 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.

Note
Redistribution of static routes with a route-map matching a metric is not supported.

Step 7 Choose one of the following options for the Metric Type.

- If the metric is a Type 1 external route, choose **1**.
- If the metric is a Type 2 external route, choose **2**.

Step 8 Enter the tag value in the **Tag Value** field.

The tag value is a 32-bit decimal value attached to each external route that is not used by OSPF itself, but may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.

Step 9 Check the **Use Subnets** check box to enable the redistribution of subnetted routes. Uncheck this check box to cause only routes that are not subnetted to be redistributed.

Step 10 Choose the name of the route map to apply to the redistribution entry from the Route Map drop-down list.

Step 11 If you need to add or configure a route map, click **Manage**.

The Configure Route Map dialog box appears.

Step 12 Click **Add** or **Edit** to define which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. For more information, see Define a Route Map.

Step 13 Click **OK**.

Configure Route Summarization When Redistributing Routes Into OSPFv2

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the ASA 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.

Add a Route Summary Address

The Summary Address pane displays information about the summary addresses configured for each OSPF routing process.

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.



Note OSPF does not support summary-address 0.0.0.0 0.0.0.0.

To configure the software advertisement on one summary route for all redistributed routes included for a network address and mask, perform the following steps:

Procedure

Step 1 In the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Summary Address**.

Step 2 Click **Add**.

The Add OSPF Summary Address Entry dialog box appears. You can add new entries to existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the OSPF Process drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Step 5 Choose the network mask for the summary address from the**Netmask** drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

The Tag value displays 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.

Step 7 Click **OK**.

Add or Edit an OSPF Summary Address

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the **OSPF Process** drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Step 5 Enter the network mask for the summary address from the **Netmask** drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Configure Route Summarization Between OSPFv2 Areas

Route summarization is the consolidation of advertised addresses. This feature causes a single summary route to be advertised to other areas by an area boundary router. In OSPF, an area boundary router advertises networks in one area into another area. If the network numbers in an area are assigned in a way so that they are contiguous, you can configure the area boundary router to advertise a summary route that includes all the individual networks within the area that fall into the specified range.

To define an address range for route summarization, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Enter the OSPF Area ID in the **Area ID** field. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Configure OSPFv2 Interface Parameters

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.

In ASDM, the Interface pane lets you configure interface-specific OSPF routing properties, such as OSPF message authentication and properties. There are two tabs that help you configure interfaces in OSPF:

- The Authentication tab displays the OSPF authentication information for the ASA interfaces.
- The Properties tab displays the OSPF properties defined for each interface in a table format.

To configure OSPFv2 interface parameters, perform the following steps:

Procedure

Step 1 Click the **Authentication** tab to display the authentication information for the ASA interfaces. Double-clicking a row in the table opens the Edit OSPF Authentication Interface dialog box for the selected interface.

Step 2 Click **Edit**.

The Edit OSPF Authentication Interface dialog box appears. The Edit OSPF Interface Authentication dialog box lets you configure the OSPF authentication type and parameters for the selected interface.

Step 3 Choose the Authentication type by clicking the relevant radio button:

- **No authentication** to disable OSPF authentication.
- **Area authentication, if defined** (Default) to use the authentication type specified for the area. See Configure OSPFv2 Area Parameters, on page 19 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure this setting.
- **Password authentication** to use clear text password authentication (not recommended where security is a concern).
- **MD5 authentication** to use MD5 authentication.
- **Key chain authentication** to use key chain authentication (recommended). See Configure a Key Chain for Authentication, on page 9 for information about configuring key chain for authentication.

Step 4 If you have chosen password authentication, in the Authentication Password area, enter the password:

a) In the Enter Password field, type a text string of up to eight characters.
b) In the Re-enter Password field, retype the password.

Step 5 If you have chosen Key chain authentication, enter the key chain name in the Enter Key chain name field.

Step 6 Choose the settings for MD5 IDs and keys in the ID area, which includes the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.

a) In the Key ID field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.

b) In the Key field, enter an alphanumeric character string of up to 16 bytes. The key displays for the selected interface.

c) Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 7 Click **OK.**

Step 8 Click the **Properties** tab.

Step 9 Choose the interface that you want to edit. Double-clicking a row in the table opens the Properties tab dialog box for the selected interface.

Step 10 Click **Edit**.

The Edit OSPF Interface Properties dialog box appears. The Interface field displays the name of the interface for which you are configuring OSPF properties. You cannot edit this field.

Step 11 Check or uncheck the **Broadcast** check box to specify 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.

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. See Define Static OSPFv2 Neighbors, on page 23 for more information.
- You need to define a static route pointing to the crypto endpoint. See Configure a Static Route for more information.
- 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.

Step 12 Configure the following options:

- Enter a value in the Cost field, which determines the cost of sending a packet through the interface. The default value is 10.
- In the Priority field, enter the OSPF router priority value.

  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.

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.

In multiple context mode, for shared interfaces, specify 0 to ensure the device does not become the designated router. OSPFv2 instances cannot form adjacencies with each other across shared interfaces.

- Check or uncheck the **MTU Ignore** check box.

  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 will not be established.

- Check or uncheck the **Database filter** check box.

  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.

Step 13 To enable BFD on this interface, from the **BFD** drop-down list, choose **Enable**. To enable BFD on all interfaces that are supporting OSPFv2, see Configure OSPFv2, on page 8.

Step 14 (Optional) Click **Advanced** to display the Edit OSPF Advanced Interface Properties dialog box, which lets you change the values for the OSPF hello interval, retransmit interval, transmit delay, and dead interval.

Typically, you only need to change these values from the defaults if you are experiencing OSPF problems on your network.

Step 15 In the Intervals section, enter values for the following:

- The Hello Interval, which specifies the interval, in seconds, between hello packets sent on an interface. 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. Valid values range from 1 to 8192 seconds. The default value is 10 seconds.
- The Retransmit Interval, which specifies the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 8192 seconds. The default value is 5 seconds.
- The Transmit Delay, which specifies the estimated time, in seconds, required to send an LSA packet on the interface. 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. Valid values range from 1 to 8192 seconds. The default value is 1 second.

Step 16 In the Detecting Lost Neighbors section, do one of the following:

- Click Configure interval within which hello packets are not received before the router declares the neighbor to be down. In the Dead Interval field, specify the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 8192 seconds. The default value of this setting is four times the interval that was set in the Hello Interval field.

- Click Send fast hello packets within 1 seconds dead interval. In the Hello multiplier field, specify the number of hello packets to be sent per second. Valid values are between 3 and 20.

Configure OSPFv2 Area Parameters

You can configure several OSPF area parameters. These area parameters (shown in the following task list) include setting authentication, defining stub areas, and assigning specific costs to the default summary route. 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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Area/Networks** tab.

The Add OSPF Area dialog box appears.

Step 3 Choose one of the following Area Type options:

- **Normal** to make the area a standard OSPF area. This option is selected by default when you first create an area.
- **Stub** to make the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you have the option of preventing summary LSAs (Types 3 and 4) from being flooded into the area by unchecking the Summary check box.
- **Summary** to prevent LSAs from being sent into the stub area when the area being defined is a stub area, uncheck this check box. By default, this check box is checked for stub areas.
- **NSSA** to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the NSSA, you have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and checking the Default Information Originate check box.

Step 4 Enter the IP address in the IP Address field of the network or host to be added to the area. Use **0.0.0.0** with a netmask of **0.0.0.0** to create the default area. You can only enter **0.0.0.0** in one area.

Step 5 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 6 Choose the OSPF Authentication type from the following options:

- **None** to disable OSPF area authentication. This is the default setting.
- **Password** to provide a clears text password for area authentication, which is not recommended where security is a concern.

- **MD5** to allow MD5 authentication.

Step 7 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 8 Click **OK**.

Configure OSPFv2 Filter Rules

Use the following procedure to filter routes or networks received or transmitted in OSPF updates.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Filter Rules**.

Step 2 Click **Add**.

Step 3 Select the OSPF process ID in **OSPF AS**.

Step 4 Choose a standard access list from the Access List drop-down list. Click **Manage** to add a new ACL.

Step 5 Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.

Step 6 For inbound filters, you can optionally specify an interface to limit the filter to updates received on that interface.

Step 7 For outbound filters, you can optionally specify what types of route are distributed.

a) Choose an option from the Protocol drop-down list.

You can choose a routing protocol, such as **BGP**, **EIGRP**, **OSPF**, or **RIP**.

Choose **Connected** to filter on peers and networks learned through connected routes.

Choose **Static** to filter on peers and networks learned through static routes.

b) If you chose BGP, EIGRP, or OSPF, also choose the **Process ID** for that protocol.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Configure an OSPFv2 NSSA

The OSPFv2 implementation of an NSSA is similar to an OSPFv2 stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but it can import autonomous system external routes in a limited way within the area.

NSSA imports Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA ABRs, which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

You can simplify administration if you are an ISP or a network administrator that must connect a central site using OSPFv2 to a remote site that is using a different routing protocol with NSSA.

Before the implementation of NSSA, the connection between the corporate site border router and the remote router could not be run as an OSPFv2 stub area because routes for the remote site could not be redistributed into the stub area, and two routing protocols needed to be maintained. A simple protocol such as RIP was usually run and handled the redistribution. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

Before you use this feature, consider these guidelines:

- You can set a Type 7 default route that can be used to reach external destinations. When configured, the router generates a Type 7 default into the NSSA or the NSSA area boundary router.
- Every router within the same area must agree that the area is NSSA; otherwise, the routers cannot communicate with each other.

Procedure

Step 1 From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Area/Networks** tab.

Step 3 Click **Add**.

The Add OSPF Area dialog box appears.

Step 4 Click the **NSSA** radio button in the Area Type area.

Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the NSSA, you have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and checking the Default Information Originate check box.

Step 5 Enter the IP address in the IP Address field of the network or host to be added to the area. Use **0.0.0.0** with a netmask of **0.0.0.0** to create the default area. You can only enter **0.0.0.0** in one area.

Step 6 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 7 In the Authentication area, click the **None** radio button to disable OSPF area authentication.

Step 8 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 9 Click **OK**.

Configure an IP Address Pool for Clustering (OSPFv2 and OSPFv3)

You can assign a range of IPv4 addresses for the router ID cluster pool if you are using Individual Interface clustering.

To assign a range of IPv4 addresses for the router ID cluster pool in Individual Interface for OSPFv2, perform the following steps:

Procedure

Step 1 From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 Click the **Cluster Pool** radio button. If you are using clustering, then you do not need to specify an IP address pool for the router ID (that is, leave the field blank). If you do not enter an IP address pool, then the ASA uses the automatically generated router ID.

Step 5 Enter the name of the IP address pool, or click the ellipses to display the Select IP Address Pool dialog box.

Step 6 Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

The Add IPv4 Pool dialog box appears.

Step 7 Enter the new IP address pool name in the **Name** field.

Step 8 Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.

Step 9 Double-click an entry to add it to the Starting IP Address field, then click **OK**.

Step 10 Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box

Step 11 Double-click an entry to add it to the Ending IP Address field, then click **OK**

Step 12 Choose the subnet mask from the drop-down list, then click **OK**.

The new IP address pool appears in the Select IP Address Pool list.

Step 13 Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

Step 14 Click **OK**.

Step 15 If you want to change the newly added IP address pool settings, click **Edit**.

The Edit IPv4 Pool dialog box appears.

Step 16 Repeat Steps 4 through 14.

Note
You cannot edit or delete an existing IP address pool that has been assigned and is already being used by one or more connection profiles.

Step 17 Click **OK**.

Step 18 To assign a range of IPv4 addresses for the router ID cluster pool in Individual Interface clustering for OSPFv3, perform the following steps:

a) From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
b) Click the **Process Instances** tab.
c) Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

d) Choose the Cluster Pool option from the Router ID drop-down list. If you do not need to specify an IP address pool for the router ID, choose the Automatic option. If you do not configure an IP address pool, then the ASA uses the automatically generated router ID.
e) Enter the IP address pool name. Alternatively, click the ellipses to display the Select IP Address Pool dialog box.
f) Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

   The Add IPv4 Pool dialog box appears.

g) Enter the new IP address pool name in the **Name** field.
h) Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.
i) Double-click an entry to add it to the Starting IP Address field, then click **OK**.
j) Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box.
k) Double-click an entry to add it to the Ending IP Address field, then click **OK**.
l) Choose the subnet mask from the drop-down list, then click **OK**.

   The new IP address pool appears in the Select IP Address Pool list.

m) Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

   The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

n) Click **OK**.
o) If you want to change the newly added cluster pool settings, click **Edit**.

   The Edit IPv4 Pool dialog box appears.

p) Repeat Steps 4 through 14.

   Note
   You cannot edit or delete an existing IP address pool that has been assigned and is already being used by another OSPFv3 process.

q) Click **OK**.

Define Static OSPFv2 Neighbors

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.

Before you begin, you must create a static route to the OSPFv2 neighbor. See Configure a Static Route for more information about creating static routes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add/Edit OSPF Neighbor Entry dialog box appears. This dialog box lets you define a new static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPF processes.
- You need to define a static route for each static neighbor.

Step 3 From the OSPF Process drop-down list, choose the OSPF process associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the **Neighbor** field, enter the IP address of the static neighbor.

Step 5 In the**Interface** field, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 6 Click **OK**.

Configure Route Calculation Timers

You can configure the delay time between when OSPFv2 receives a topology change and when it starts an SPF calculation. You also can configure the hold time between two consecutive SPF calculations.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA pacing and SPF calculation timers. In the Timers area, enter the following values:

- The Initial SPF Delay, specifies the time (in milliseconds) between when OSPF receives a topology change and when the SPF calculation starts. Valid values range from 0 to 600000 milliseconds.
- The Minimum SPF Hold Time, specifies the hold time (in milliseconds)between consecutive SPF calculations.Valid values range from 0 to 600000 milliseconds.
- The Maximum SPF Wait Time, specifies the maximum wait time between two consecutive SPF calculations. Valid values range from 0 to 600000 milliseconds.

Step 5 Click **OK**.

Log Neighbors Going Up or Down

By default, a syslog message is generated when an OSPFv2 neighbor goes up or down.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Adjacency Changes area includes settings that define the adjacency changes that cause syslog messages to be sent. In the Adjacency Changes area, enter the following values:

- Check the **Log Adjacency Changes** check box to cause the ASA to send a syslog message whenever an OSPFv2 neighbor goes up or down. This setting is checked by default.
- Check the **Log Adjacency Changes Detail** check box to cause the ASA 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.

Step 5 Click **OK**.

Note
Logging must be enabled for the neighbor up or down messages to be sent.

Configure a Key Chain for Authentication

To enhance data security and protection of devices, you can enable rotating keys for authenticating IGP peers. The rotating keys prevent any malicious user from guessing the keys used for routing protocol authentication and thereby protecting the network from advertising incorrect routes and redirecting traffic. Changing the keys frequently reduces the risk of them eventually being guessed. When configuring authentication for routing protocols that provide key chains, configure the keys in a key chain to have overlapping lifetimes. This helps to prevent loss of key-secured communication due to absence of an active key. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with the peers.

This section describes how to create a key chain for OSPF peer authentication. This section also covers steps to add or edit the key chain attributes. After configuring a key chain object, you can use it in defining the OSPFv2 authentication for an interface and for a virtual link. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency. To learn how to define authentication for an interface, see Configure OSPFv2 Interface Parameters, on page 16; for a virtual link, see Configure a Virtual Link in OSPF, on page 27.

To configure a key chain, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Key Chain** .

Step 2 In the **Configure Key Chain** section, click **Add**.

Step 3 Enter the key chain name in the **Add Key Chain** dialog box, and click **Ok**.

The created key chain name is listed in the **Configure Key Chain** grid.

Step 4 Select the key chain name from the **Configure Key Chain** section, and in the **Configure Key** section, click **Add**. To edit an existing key, seclect the key name and click **Edit**.

The **Add Key** or **Edit Key**dialog box appears, depending on the action that you have selected.

Step 5 Specify the key identifier in the **Key ID** field.

The key id value can be between 0 and 255. Use the value 0 only when you want to signal an invalid key.

Note
You cannot edit a saved key id.

Step 6 From the **Cryptographic Algorithm** drop-down, choose **MD5**. MD5 is the only algorithm supported for authenticating the key chain.

Step 7 Select the encryption type by clicking the **Plain Text** or **Encrypted** radio button, and then enter the password in the **Authentication Key** field.

- Even though the key-string can accommodate 80 characters, only a maximum of 64 characters is permitted for OSPF authentication.
- The passwords cannot be a single digit nor those starting with a digit followed by a white space. For example, "0 pass" or "1" are invalid.

Step 8 Provide the lifetime values in the **Accept Lifetime** and **Send Lifetime** fields:

You can specify the time interval for the device to accept/send the key during key exchange with another device. The end time can be the duration, the absolute time when the accept/send lifetime ends, or never expires.

Following are the validation rules for the start and end values:

- Start lifetime cannot be null when the end lifetime is specified.
- The start lifetime for accept or send lifetime must be earlier than the end lifetime.

Step 9 To save the key chain attibutes, click **Ok**. In the **Key Chain** page, click **Appy**.

What to do next

You can now apply the configured key chain to define the OSPFv2 authentication for an interface and for virtual link.

- Configure OSPFv2 Interface Parameters, on page 16

• Configure a Virtual Link in OSPF, on page 27

Configure Filtering in OSPF

The Filtering pane displays the ABR Type 3 LSA filters that have been configured 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. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Note Only Type 3 LSAs that originate from an ABR are filtered.

To configure filtering in OSPF, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Filtering**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Filtering Entry dialog box lets you add new filters to the Filter table or modify an existing filter. Some of the filtering information cannot be changed when you edit an existing filter.

Step 3 Choose the OSPF process that is associated with the filter entry from the OSPF Process drop-down list.

Step 4 Choose the Area ID that is associated with the filter entry from the Area ID drop-down list. If you are editing an existing filter entry, you cannot modify this setting.

Step 5 Choose a prefix list from the Prefix List drop-down list.

Step 6 Choose the traffic direction being filtered from the Traffic Direction drop-down list.

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 **Manage** to display the Configure Prefix Lists dialog box, from which you can add, edit, or delete prefix lists and prefix rules. For more information, see Configure Prefix Lists and the Configure the Metric Values for a Route Action.

Step 8 Click **OK**.

Configure a Virtual Link in OSPF

If you add an area to an OSPF network, and it is not possible to connect the area directly to the backbone area, you need to create a virtual link. A virtual link connects two OSPF devices that have a common area, called the transit area. One of the OSPF devices must be connected to the backbone area.

To define new virtual links or change the properties of existing virtual links, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Virtual Link**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Virtual Link dialog box appears, which allows you to define new virtual links or change the properties of existing virtual links.

Step 3 Choose the OSPF process ID that is associated with the virtual link from the OSPF Process drop-down list. If you are editing an existing virtual link entry, you cannot modify this setting.

Step 4 Choose the Area ID that is associated with the virtual link from the Area ID drop-down list.

Choose the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a Stub area. If you are editing an existing virtual link entry, you cannot modify this setting.

Step 5 In the **Peer Router ID** field, enter the router ID of the virtual link neighbor.

If you are editing an existing virtual link entry, you cannot modify this setting.

Step 6 Click **Advanced** to edit advanced virtual link properties,

The Advanced OSPF Virtual Link Properties dialog box appears. You can configure the OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings.

Step 7 In the Authentication area, choose the Authentication type by clicking the radio button next to one of the following options:

- **No authentication** to disable OSPF authentication.
- **Password authentication** to use clear text password authentication (not recommended where security is a concern).
- **MD5 authentication** to use MD5 authentication.
- **Key chain authentication** to use key chain authentication (recommended). See Configure a Key Chain for Authentication, on page 9 for information about configuring key chain for authentication.

Step 8 In the Authentication Password area, enter and re-enter a password when password authentication is enabled. Passwords must be a text string of up to 8 characters.

Step 9 In the MD5 IDs and Key area, enter the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID. Specify the following settings:

a) In the **Key ID**field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.
b) In the **Key** field, enter an alphanumeric character string of up to 16 bytes. The Key ID displays for the selected interface.
c) Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 10 In the Interval area, specify the interval timing for the packet by choosing from the following options:

- **Hello Interval** to specify the interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.

This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.

- **Retransmit Interval** to specify the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
- **Transmit Delay** to specify the estimated time, in seconds, required to send an LSA packet on the interface. 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. Valid values range from 1 to 65535 seconds. The default value is 1 second.
- **Dead Interval** to specify the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this field is four times the interval set by the Hello Interval field.

Step 11 Click **OK**.

Configure OSPFv3

This section describes the tasks involved in configuring an OSPFv3 routing process.

Enable OSPFv3

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

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 On the Process Instances tab, check the **Enable OSPFv3 Process** check box. You can enable up to two OSPF process instances. Only single context mode is supported.

Step 3 Enter a process ID in the Process ID field. The ID can be any positive integer.

Step 4 To enable BFD on all interfaces that support OSPFv3, click **Advanced**. In the **Edit OPPFv3 Process Advanced Properties** window, under **Enable BFD on all interfaces**, click the **Enable BFD** check box. To enable BFD on a specific OSPFv3 intreface, see Configure OSPFv3 Interface Parameters, on page 30.

Step 5 Click **Apply** to save your changes.

Step 6 To continue, see Configure OSPFv3 Area Parameters, on page 31.

Configure OSPFv3 Interface Parameters

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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Interfaces**.

Step 2 Click the **Authentication** tab.

Step 3 To specify the authentication parameters for an interface, select the interface and click **Edit**.

The **Edit OSPFv3 Interface Authentication** dialog box appears.

Step 4 Choose the authentication type from the **Authentication Type** drop-down list. The available options are Area, Interface, and None. The None option indicates that no authentication is used.

Step 5 Choose the authentication algorithm from the **Authentication Algorithm** drop-down list. Supported values are SHA-1 and MD5.

Step 6 Enter the authentication key in the **Authentication Key** field. 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.

Step 7 Choose the encryption algorithm from the **Encryption Algorithm** drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 8 Enter the encryption key in the **Encryption Key** field.

Step 9 Click **OK**.

Step 10 Click the **Properties** tab.

Step 11 Select the interface whose properties you want to modify, and click **Edit**.

The Edit OSPFv3 Interface Properties dialog box appears.

Step 12 Check the **Enable OSPFv3 on this interface** check box.

Step 13 Choose the process ID from the drop-down list.

Step 14 Choose the area ID from the drop-down list.

Step 15 (Optional) Specify 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 16 Choose the network type from the drop-down list. Supported options are Default, Broadcast, and Point-to-Point.

Step 17 Enter the cost of sending a packet on an interface in the Cost field.

Step 18 Enter the router priority, which helps determine the designated router for a network. in the Priority field. Valid values range from 0 to 255.

Step 19 To enable BFD on this interface, from the **BFD Configuration** drop-down list, choose **Enable**. To enable BFD on all interfaces that support OSPFv3, see Enable OSPFv3, on page 29.

Step 20 Check the **Disable MTU mismatch detection** check box to disable the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.

Step 21 Check the **Filter outgoing link state advertisements** check box to filter outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.

Step 22 Check the **OSPF Flood Reduction** check box to reduce unnecessary flooding and refreshing of LSAs to the interface.

Step 23 In the **Timers** area, in the **Dead Interval** field, enter the 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.

Step 24 In the **Hello Interval** field, enter the interval in seconds between hello packets sent on the interface. The value must be the same for all nodes on a specific network and can range from 1 to 65535. The default interval is 10 seconds for Ethernet interfaces and 30 seconds for non-broadcast interfaces.

Step 25 In the **Retransmit Interval** field, enter the 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.

Step 26 In the **Transmit Delay** field, enter the 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.

Step 27 Click **OK**.

Step 28 Click **Apply** to save your changes.

Configure OSPFv3 Area Parameters

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Areas** tab.

Step 3 To add a new area, click **Add**. To modify an existing area, click **Edit**. To remove a selected area, click **Delete**.

The Add OSPFv3 Area dialog box or Edit OSPFv3 Area dialog box appears.

Step 4 From the OSPFv3 Process ID drop-down list, choose the process ID.

Step 5 Enter the area ID, which specifies the area for which routes are to be summarized, in the Area ID field.

Step 6 Choose the area type from the Area Type drop-down list. Available options are Normal, NSSA, and Stub.

Step 7 To allow the sending of summary LSAs into the area, check the **Allow sending of summary LSAs into the area** check box.

Step 8 To allow redistribution to import routes to normal and not so stubby areas, check the **Redistribution imports routes to normal and NSSA areas** check box.

Step 9 To generate a default external route into an OSPFv3 routing domain, check the **Default information originate** check box.

Step 10 Enter the metric used for generating the default route in the Metric field. The default value is 10. Valid metric values range from 0 to 16777214.

Step 11 Choose the metric type from the Metric Type drop-down list. 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 12 Enter the cost in the Default Cost field.

Step 13 Click **OK**.

Step 14 Click the **Route Summarization** tab.

Step 15 To specify a new range for consolidating and summarizing routes, click **Add**. To modify an existing range for consolidating and summarizing routes, click **Edit**.

The Add Route Summarization dialog box or Edit Route Summarization dialog box appears.

Step 16 Choose the process ID from the Process ID drop-down list.

Step 17 Choose the area ID from the Area ID drop-down list.

Step 18 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 19 (Optional) Enter 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.

Step 20 Check the **Advertised** check box to set the address range status to advertised and generate a Type 3 summary LSA.

Step 21 Click **OK**.

Step 22 To continue, see Configure a Virtual Link Neighbor, on page 32.

Configure a Virtual Link Neighbor

To configure a virtual link neighbor, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Virtual Link**.

Step 2 To add a new virtual link neighbor, click **Add**. To modify an existing virtual link neighbor, click **Edit**. To remove a selected virtual link neighbor, click **Delete**.

The Add Virtual Link dialog box or Edit Virtual Link dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the area ID from the Area ID drop-down list.

Step 5 Enter the peer router ID (that is, the IP address) in the Peer Router ID field.

Step 6 (Optional) Enter the time-to-live (TTL) security hop count on a virtual link in the TTL Security field. The hop count value can range from 1 to 254.

Step 7 In the Timers area, enter the time in seconds that hello packets are not seen before a neighbor indicates that the router is down in the Dead Interval field. 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 to 8192.

Step 8 Enter the time in seconds between the hello packets that are sent on an interface in the Hello Interval field. 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 that are attached to a common network. Valid values range from 1 to 8192. The default is 10.

Step 9 Enter the time in seconds between LSA retransmissions for adjacencies that belong to the interface in the Retransmit Interval field. 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 8192. The default is 5.

Step 10 Enter the estimated time in seconds that is required to send a link-state update packet on the interface in the Transmit Delay field. The integer value must be greater than zero. LSAs in the update packet have their own ages incremented by this amount before transmission. The range of values can be from 1 to 8192. The default is 1.

Step 11 In the Authentication area, check the **Enable Authentication** check box to enable authentication.

Step 12 Enter the security policy index, which must be a number from 256 to 4294967295, in the Security Policy Index field.

Step 13 Choose the authentication algorithm from the Authentication Algorithm drop-down list. Supported values are SHA-1 and MD5. 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.

Step 14 Enter the authentication key in the Authentication Key field. The key must include 32 hexadecimal characters.

Step 15 Choose the encryption algorithm from the Encryption Algorithm drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 16 Enter the encryption key in the Encryption Key field.

Step 17 Click **OK**.

Step 18 Click **Apply** to save your changes.

Configure OSPFv3 Passive Interfaces

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Passive Interfaces area allows you to enable passive OSPFv3 routing on an interface. Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface. In the Passive Interfaces area, choose the following settings:

- Check the **Global passive** check box to make all of the interfaces listed in the table passive. Uncheck individual interfaces to make them non-passive.
- Uncheck the **Global passive** check box to make all of the interfaces non-passive. Check individual interfaces to make them passive.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure OSPFv3 Administrative Distance

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Administrative Route Distances area allows you to modify the settings that were used to configure administrative route distances. The administrative route distance is an integer from 10 to 254. In the Administrative Route Distances area, enter the following values:

- The Inter Area, which specifies the inter-area routes for OSPV for IPv6 routes.
- The Intra Area, which specifies the intra-area routes for OSPF for IPv6 routes.
- The External, which specifies the external type 5 and type 7 routes for OSPF for IP6 routes.

Step 4 Click **OK**.

Step 5 Click **Apply** to save your changes.

Configure OSPFv3 Timers

You can set LSA arrival, LSA pacing, and throttling timers for OSPFv3.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA arrival, LSA pacing, LSA retransmission, LSA throttle, and SPF throttle times. In the Timers area, enter the following values:

- The LSA Arrival, which 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.
- The LSA Flood Pacing, which 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.
- The LSA Group Pacing, which 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.

- The LSA Retransmission Pacing, which 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.
- The LSA Throttle Initial, which specifies the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 milliseconds.
- The LSA Throttle Min Hold, which specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.
- The LSA Throttle Max Wait, which 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.

- The SPF Throttle Initial, specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds.
- The SPF Throttle Min Hold, which specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds.
- The SPF Throttle Max Wait, which 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 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Define Static OSPFv3 Neighbors

You need to define static OSPFv3 neighbors to advertise OSPF routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv3 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Before you begin, you must create a static route to the OSPFv3 neighbor. See Configure a Static Route for more information about creating static routes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add or Edit Static Neighbor dialog box appears. This dialog box lets you define a new static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPFv3 processes.
- You need to define a static route for each static neighbor.

Step 3 From the Interface drop-down list, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the Link-local Address field, enter the IPv6 address of the static neighbor.

Step 5 (Optional) In the Priority field, enter the priority level.

Step 6 (Optional) In the Poll Interval field, enter the poll interval in seconds.

Step 7 Click **OK**.

Send Syslog Messages

Configure the router to send a syslog message when an OSPFv3 neighbor goes up or down.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Adjacency Changes area allows you to modify the settings for sending syslog messages when an OSPFv3 neighbor goes up or down. In the Adjacency Changes area, do the following:

- To send a syslog message when an OSPFv3 neighbor goes up or down, check the **Log Adjacency Changes** check box.
- To sends a syslog message for each state, not only when an OSPFv3 neighbor goes up or down, check the **Include Details** check box.

Step 4 Click **OK**.

Step 5 Click **Apply** to save your changes.

Suppress Syslog Messages

To suppress the sending of syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **Ignore LSA MOSPF** check box, then click **OK**.

Calculate Summary Route Costs

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **RFC1583 Compatible** check box, then click **OK**.

Generate a Default External Route into an OSPFv3 Routing Domain

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 In the Default Information Originate Area, do the following:

a) Check the **Enable** check box to enable the OSPFv3 routing process.

b) Check the **Always advertise** check box to always advertise the default route, whether or not one exists.

c) Enter the metric used for generating the default route in the Metric field. Valid metric values range from 0 to 16777214. The default value is 10.

d) From the Metric Type drop-down list, choose the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are the following:

- 1—Type 1 external route

- 2—Type 2 external route

The default is the Type 2 external route.

e) From the Route Map drop-down list, choose the routing process that generates the default route if the route map is satisfied.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure an IPv6 Summary Prefix

Procedure

Step 1 In the ASDM main window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Summary Prefix**.

Step 2 To add a new summary prefix, click **Add**. To modify an existing summary prefix, click **Edit**. To remove a summary prefix, click **Delete**.

The Add Summary Prefix dialog box or Edit Summary Prefix dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 5 Check the **Advertise** check box to advertise routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.

Step 6 Enter the tag value that you can use as a match value for controlling redistribution through route maps in the Tag field.

Step 7 Click **OK**.

Step 8 Click **Apply** to save your changes.

Redistribute IPv6 Routes

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Redistribution**.

Step 2 To add new parameters for redistributing connected routes into an OSPFv3 process, click **Add**. To modify existing parameters for redistributing connected routes into an OSPFv3 process, click **Edit**. To remove a selected set of parameters, click **Delete**.

The Add Redistribution dialog box or Edit Redistribution dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the source protocol from which routes are being redistributed from the Source Protocol drop-down list. The supported protocols are connected, static, and OSPF.

Step 5 Enter the metric value in the Metric field. When redistributing routes from one OSPF process into another OSPF process on the same router, the metric is carried through from one process to the other if no metric value is specified. When redistributing other processes into an OSPF process, the default metric is 20 when no metric value is specified.

Step 6 Choose the metric type from the Metric Type drop-down list. The available options are None, 1, and 2.

Step 7 (Optional) Enter the tag value in the Tag field. This parameter specifies the 32-bit decimal value attached to each external route, 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.

Step 8 Choose the route map from the Route Map drop-down list to check 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.

Step 9 To include connected routes in the redistribution, check the **Include connected** check box.

Step 10 Check the **Match** check box to redistribute routes into other routing domains, then check one of the following check boxes:

- **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 11 Click **OK**.

Step 12 Click **Apply** to save your changes.

Configure Graceful Restart

The ASA 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.

In a high availability mode, the OSPF process restarts when the active unit becomes inactive and the standby unit becomes the new active. Similarly, in a cluster mode, the OSPF process restarts when the control unit becomes inactive and the data unit is elected as the new control unit. Such OSPF transitioning processes involve a considerable amount of delay. You can configure NSF to avoid traffic loss during the OSPF process state change. The NSF capability is also useful when there is a scheduled hitless software upgrade.

Graceful restart is supported on both OSPFv2 and OSPFv3. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623). You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).

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.



Note When fast hellos are configured for OSPFv2, graceful restart does not occur when the active unit reloads and the standby unit becomes active. This is because the time taken for the role change is more than the configured dead interval.

Configuring Graceful Restart for OSPFv2

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.

Configure Cisco NSF Graceful Restart for OSPFv2

Configure Cisco NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

Step 1 In the main ASDM window, choose Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties.

Step 2 Under Configuring Cisco NSF, check the Enable Cisco nonstop forwarding (NSF) check box.

Step 3 (Optional) Check the Cancels NSF restart when non-NSF-aware neighboring networking devices are detected check box if required.

Step 4 (Optional) Under Configuring Cisco NSF helper, uncheck the Enable Cisco nonstop forwarding (NSF) for helper mode check box.

Note
This is checked by default. Uncheck this to disable the Cisco NSF helper mode on NSF-aware device.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure IETF NSF Graceful Restart for OSPFv2

Configure IETF NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties**.

Step 2 Under Configuring IETF NSF, check the Enable IETF nonstop forwarding (NSF) check box.

Step 3 (Optional) Enter the restart interval in seconds in the Length of graceful restart interval field.

Note
The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 4 (Optional) Under Configuring IETF NSF helper, uncheck the Enable IETF nonstop forwarding (NSF) for helper mode check box.

This is checked by default. Uncheck this to disable the IETF NSF helper mode on NSF-aware device.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configuring Graceful Restart for OSPFv3

Configuring the NSF graceful-restart feature for OSPFv3 involves two steps; configuring a device to be NSF-capable and then configuring a device to be NSF-aware.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup > Advanced > Add NSF Properties**.

Step 2 Under Configuring Graceful Restart, check the Enable Graceful Restart check box.

Step 3 (Optional) Enter a value for the restart interval in the Restart Interval field.

Note
The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 4 Under Configuring Graceful Restart Helper, check the Enable Graceful Restart Helper check box.

This is checked by default. Uncheck this to disable the Graceful-restart helper mode on a NSF-aware device.

Step 5 (Optional) Check the Enable LSA checking 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 6 Click **OK**.

Step 7 Click **Apply** to save your changes.

Configuring Graceful Restart Wait Timer for OSPF

OSPF routers are expected to set the RS-bit in the EO-TLV attached to a Hello packet when it is not known that all neighbors are listed in the packet, but the restarting routers require to preserve their adjacencies. However, the RS-bit value must not be longer than the RouterDeadInterval seconds. Hence the **timers nsf wait** command is introduced to set the RS-bit in Hello packets lesser than RouterDeadInterval seconds. The default value of NSF wait timer is 20 seconds.

Before you begin

- To configure Cisco NSF wait time for OSPF, the device must be NSF-aware or NSF-capable.

Procedure

Step 1 Enter into OSPF router configuration mode.

Example:

```
ciscoasa(config)# router ospf
```

Step 2 Enter timers and specify nsf.

Example:

```
ciscoasa(config-router)# timers?
router mode commands/options:
  lsa       OSPF LSA timers
  nsf       OSPF NSF timer
  pacing    OSPF pacing timers
  throttle  OSPF throttle timers
ciscoasa(config-router)# timers nsf ?
```

Step 3 Enter the graceful restart wait interval. This value can range between 1 and 65535.

Example:

```
ciscoasa(config-router)# timers nsf wait 200
```

By using the graceful restart wait interval, you can ensure that the wait interval is not longer than the router dead interval.

Remove the OSPFv2 Configuration

Remove the OSPFv2 configuration.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.
Step 2 Uncheck the **Enable this OSPF Process** check box.
Step 3 Click **Apply**.

Remove the OSPFv3 Configuration

Remove the OSPFv3 configuration.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
Step 2 Uncheck the **Enable OSPFv3 Process** check box.
Step 3 Click **Apply**.

Example for OSPFv2

The following example shows how to enable and configure OSPFv2 with various optional processes:

1. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.
2. Click the **Process Instances** tab and in the OSPF Process 1 field, type **2**.
3. Click the **Area/Networks** tab, and click **Add**.
4. Enter **0** in the Area ID field.
5. In the Area Networks area, enter **10.0.0.0** in the IP Address field.
6. Choose 255.0.0.0 from the Netmask drop-down list.
7. Click **OK**.
8. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Redistribution**.
9. Click **Add**.

   The Add/Edit OSPF Redistribution Entry dialog box appears.
10. In the Protocol area, click the **OSPF** radio button to choose the source protocol from which the routes are being redistributed. Choosing OSPF redistributes routes from another OSPF routing process.
11. Choose the OSPF process ID from the OSPF Process drop-down list.

12. In the Match area, check the **Internal** check box.
13. In the Metric Value field, enter **5** for the metric value for the routes being redistributed.
14. From the Metric Type drop-down list, choose 1 for the Metric Type value.
15. From the Route Map drop-down list, choose 1.
16. Click **OK**.
17. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Interface**.
18. From the Properties tab, choose the **inside** interface and click **Edit.**

    The Edit OSPF Properties dialog box appears.
19. In the Cost field, enter **20**.
20. Click **Advanced**.
21. In the Retransmit Interval field, enter **15**.
22. In the Transmit Delay field, enter **20**.
23. In the Hello Interval field, enter **10**.
24. In the Dead Interval field, enter **40**.
25. Click **OK**.
26. In the Edit OSPF Properties dialog box, enter **20** in the Priorities field, and click **OK**.
27. Click the **Authentication** tab.

    The Edit OSPF Authentication dialog box appears.
28. In the Authentication area, click the **MD5** radio button.
29. In the MD5 and Key ID area, enter **cisco** in the MD5 Key field, and **1** in the MD5 Key ID field.
30. Click **OK**.
31. Choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup,** and click the **Area/Networks** tab.
32. Choose the **OSPF 2** process and click **Edit**.

    The Edit OSPF Area dialog box appears.
33. In the Area Type area, choose **Stub**.
34. In the Authentication area, choose **None**, and enter **20** in the Default Cost field.
35. Click **OK**.
36. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.
37. Click the **Process Instances** tab and check the **OSPF process 2** check box.
38. Click **Advanced**.

    The Edit OSPF Area dialog box appears.

39. In the Timers area, enter **10** in the SPF Delay Time field and **20** in the SPF Hold Time field.
40. In the Adjacency Changes area, check the **Log Adjacency Change Details** check box.
41. Click **OK**.
42. Click **Reset**.

Examples for OSPFv3

The following example shows how to configure OSPFv3 routing in ASDM:

1. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
2. On the Process Instances tab, do the following:
   a. Check the **Enable OSPFv3 Process** check box.
   b. Enter **1** in the Process ID field.
3. Click the **Areas** tab. then click **Add** to display the Add OSPFv3 Area dialog box.
4. From the OSPFv3 Process ID drop-down list, choose **1**.
5. Enter **22** in the Area ID field.
6. Choose **Normal** from the Area Type drop-down list.
7. Enter **10** in the Default Cost field.
8. Check the **Redistribution imports routes to normal and NSSA areas** check box.
9. Enter **20** in the Metric field.
10. Choose **1** from the Metric Type drop-down list.
11. Check the **inside** check box as the specified interface being used.
12. Check the **Enable Authentication** check box.
13. Enter **300** in the Security Policy Index field.
14. Choose **SHA-1** from the Authentication Algorithm drop-down list.
15. Enter **12345ABCDE** in the Authentication Key field.
16. Choose **DES** from the Encryption Algorithm drop-down list.
17. Enter **1122334455aabbccddee** in the Encryption Key field.
18. Click **OK**.
19. Click the **Route Summarization** tab, then click **Add** to display the Add Route Summarization dialog box.
20. Choose **1** from the Process ID drop-down list.
21. Choose **22** from the Area ID drop-down list.

22. Enter **2000:122::/64** in the IPv6 Prefix/Prefix Length field.
23. (Optional) Enter **100** in the Cost field.
24. Check the **Advertised** check box.
25. Click **OK**.
26. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Interface**.
27. Click the **Properties** tab.
28. Check the **inside** check box and click **Edit** to display the Edit OSPF Properties dialog box.
29. In the Cost field, enter **20**.
30. Enter **1** in the Priority field.
31. Check the **point-to-point** check box.
32. In the Dead Interval field, enter **40**.
33. In the Hello Interval field, enter **10**.
34. In the Retransmit Interval field, enter **15**.
35. In the Transmit Delay field, enter **20**.
36. Click **OK**.
37. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Redistribution**.
38. Choose **1** from the Process ID drop-down list.
39. Choose **OSPF** from the Source Protocol drop-down list.
40. Enter **50** in the Metric field.
41. Choose **1** from the Metric Type drop-down list.
42. Click **OK**.
43. Click **Apply** to save your changes.

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can also use the information provided to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your device packets are taking through the network.

To monitor or display various OSPFv2 routing statistics in ASDM, perform the following steps:

1. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPF LSAs**.
2. You can select and monitor OSPF LSAs, Types 1 through 5 and 7. Each pane shows one LSA type, as follows:

- Type 1 LSAs represent the routes in an area under a process.
- Type 2 LSAs show the IP address of the designated router that advertises the routers.
- Type 3 LSAs show the IP address of the destination network.
- Type 4 LSAs show the IP address of the AS boundary router.
- Type 5 LSAs and Type 7 LSAs show the IP address of the AS external network.

3. Click **Refresh** to update each LSA type pane.
4. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPF Neighbors**.

   In the OSPF Neighbors pane, each row represents one OSPF neighbor. In addition, the OSPF Neighbors pane shows the network on which the neighbor is running, the priority, the state, the amount of dead time in seconds, the IP address of the neighbor, and the interface on which it is running. For a list of possible states for an OSPF neighbor, see RFC 2328.
5. Click **Refresh** to update the OSPF Neighbors pane.

To monitor or display various OSPFv3 routing statistics in ASDM, perform the following steps:

1. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPFv3 LSAs**.
2. You can select and monitor OSPFv3 LSAs. Choose a link-state type to display its status according to specified parameters from the Link State type drop-down list. The supported link-state types are router, network, inter-area prefix, inter-area router, AS external, NSSA, link, and intra-area prefix.
3. Click **Refresh** to update each link-state type.
4. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPFv3 Neighbors**.

   In the OSPFv3 Neighbors pane, each row represents one OSPFv3 neighbor. In addition, the OSPFv3 Neighbors pane shows the IP address of the neighbor, the priority, the state, the amount of dead time in seconds, and the interface on which it is running. For a list of possible states for an OSPFv3 neighbor, see RFC 5340.
5. Click **Refresh** to update the OSPFv3 Neighbors pane.

History for OSPF

Table 1: Feature History for OSPF

Feature Name	Platform Releases	Feature Information
OSPF Support	7.0(1)	Support was added for route data, authentication, and redistribution and monitoring of routing information using the Open Shortest Path First (OSPF) routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > OSPF.
Dynamic Routing in Multiple Context Mode	9.0(1)	OSPFv2 routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup

Feature Name	Platform Releases	Feature Information
Clustering	9.0(1)	For OSPFv2 and OSPFv3, bulk synchronization, route synchronization, and Spanned EtherChannel load balancing are supported in the clustering environment.
OSPFv3 Support for IPv6	9.0(1)	OSPFv3 routing is supported for IPv6. We introduced the following screens: Configuration > Device Setup > Routing > OSPFv3 > Setup, Configuration > Device Setup > Routing > OSPFv3 > Interface, Configuration > Device Setup > Routing > OSPFv3 > Redistribution, Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix, Configuration > Device Setup > Routing > OSPFv3 > Virtual Link, Monitoring > Routing > OSPFv3 LSAs, Monitoring > Routing > OSPFv3 Neighbors.
OSPF support for Fast Hellos	9.2(1)	OSPF supports the Fast Hello Packets feature, resulting in a configuration that results in faster convergence in an OSPF network. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Interface > Edit OSPF Interface Advanced Properties
Timers	9.2(1)	New OSPF timers were added; old ones were deprecated. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup > Edit OSPF Process Advanced Properties
Route filtering using access-list	9.2(1)	Route filtering using ACL is now supported. We introduced the following screen: Configuration > Device Setup > Routing > OSPF > Filtering Rules > Add Filter Rules
OSPF Monitoring enhancements	9.2(1)	Additional OSPF monitoring information was added.
OSPF redistribute BGP	9.2(1)	OSPF redistribution feature was added. We added the following screen: Configuration > Device Setup > Routing > OSPF > Redistribution
OSPF Support for Non-Stop Forwarding (NSF)	9.3(1)	OSPFv2 and OSPFv3 support for NSF was added. We added the following screens: Configuration > Device Setup > Routing > OSPF > Setup > NSF Properties, Configuration > Device Setup > Routing > OSPFv3 > Setup > NSF Properties
OSPF Support for Non-Stop Forwarding (NSF)	9.13(1)	NSF wait timer was added. We added a new command for setting the timer for the NSF restart interval. This command was introduced to ensure the wait interval is not longer that the router dead interval. We introduced the following command: **timers nsf wait <seconds>**