

Open Shortest Path First (OSPF)

Open Shortest Path First (OSPF) is a link-state interior gateway protocol. OSPF routers flood link-state information to neighboring routers so that all routers in an OSPF area have a complete view of the network topology.

There are separate OSPF versions based on the IP version: OSPFv2 for IPv4 networks, and OSPFv3 for IPv6 networks. These versions are independent; that is, OSPFv3 is not a replacement for OSPFv2.

You can configure OSPFv2 using Smart CLI objects to integrate your device into the OSPFv2 network topology. You cannot configure OSPFv3.

- Configure the OSPFv2 Process and Areas, on page 1
- Customizing OSPF Process and Area Characteristics, on page 3
- Configure OSPFv2 Interface Settings and OSPF Authentication, on page 15
- Monitoring OSPF, on page 19

Configure the OSPFv2 Process and Areas

You can configure up to 2 OSPFv2 processes. The process numbers are purely internal indicators; they do not need to match any process numbers used on other devices, although you can make the numbers consistent for your own tracking purposes.

If you use private network numbering, such as 192.168.1.0/24, for any internal networks, you might need to segregate private addresses from public addresses, using one OSPFv2 process for these internal networks, and a second process for the external, publicly-addressable networks. Even if you do not use private numbering, 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. If you use NAT, and 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.

Area numbers, in contrast, do exist in the network and you must use the same numbers used by other adjacent routers. If you are configuring a single-area network, use Area 0, also known as the backbone area. For multiple-area networks, where you have a hierarchical network design, you must understand the areas defined in the network, and know which areas this device is supposed to participate in.

If you are using virtual routers, you can configure 2 OSPFv2 processes per virtual router.

The following procedure explains how to create a single OSPFv2 process. Repeat the procedure to create a second process.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new process, click **+** > **OSPF** or click the **Create OSPF Object** > **OSPF** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need a process, click the trash can icon for the object to delete it.

Step 5 Enter a name for the object, and optionally, a description.

Step 6 Configure the basic process properties:

- **router ospf** *process-id*. Click *process-id* and enter a number from 1-65535. This number has meaning within this device only and does not need to match any process numbers configured on other routers. The number must be unique within a virtual router.
- **log-adj-changes** *log-state*. Click *log-state* and select one of the following options:
  - **enable** (recommended)—The system generates a syslog message when an OSPFv2 neighbor goes up or down. If you select this option, an additional **log-adj-changes** *log-type* line is added to the object. Click *log-type* and select **detail** if you want to generate a syslog message for each state change, not just when a neighbor goes up or down.

    If you do not want detailed messages, simply leave *log-type* as the option. Do not delete this line from the object.
  - **disable**—No syslog message is generated. The **no log-adj-changes** line is added to the object: do not delete this line.

Step 7 Click the **Show Disabled** link above the object body to add all other possible configuration lines.

Step 8 Configure the area number.

a) Click the **+** to the left of the **area** *area-id* line to enable the command. You cannot configure a command until you enable it.

b) Click *area-id* and enter the number of the area. This area number needs to be the same as the one used by the other routers that define the OSPFv2 area. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.

Step 9 Configure the networks and interfaces that should be routed within the area.

a) Click the **+** to the left of the **configure area** *area-id* *options* line.

b) Click *area-id* and enter the same area number from the **area** command.

c) Click *options* and select **properties**. This action adds several lines, including one that is enabled by default, the **network** command.

d) In the **network** command, click *network-object* and select the object that defines a network that should be included in this area. Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

e) (Optional.) In the **network** command, click *tag-interface* and select the interface that hosts or routes to the network. If you select the interface, the system can prevent you from changing the address on the interface because it is used in the routing process. This helps remind you that any change to interface addressing can impact your routing configuration.

If you select an interface here, before you can change the address on an interface, you must first remove it from the routing process. Then, after the changing the IP address, remember to return here and select the new networks and interface to ensure a correctly-configured routing process.

f) All other new area lines are optional and disabled by default. Configure them only if you need these services. For more information, see Customizing OSPF Process and Area Characteristics, on page 3.

Step 10 If you are configuring the process for a multiple-area network, mouse over the area to the left of the circled **-** on the **area** and **configure area** lines and click **...** > **duplicate**. Then, configure the new area and its networks as explained above. Repeat this process until you have defined all areas in which this routing process should participate.

Step 11 Click **OK**.

Customizing OSPF Process and Area Characteristics

OSPF includes many options that have default values. These values work well for many networks. However, you might need to adjust one or more settings to get the precise behavior you need. The following topics explain the various ways you can customize your OSPFv2 routing process.

Configure Advanced Settings for an OSPF Process

You can configure several settings that control the overall behavior of an OSPFv2 process, including distance metrics, timers, graceful restart, and the router ID used for sending link state advertisements and other routing updates. Many of these settings have defaults that are appropriate for most networks.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Look for the **setup ospf** line.

When adding an object, you must click the **Show Disabled** link to see the line. Then, click the **+** for the command to enable it, and click *configuration* and select **advanced**. The commands that will be enabled by default are already enabled with their default values.

When editing an object, the line will already be enabled.

The remainder of this procedure assumes you have clicked **Show Disabled**. If you cannot see a command, make sure you expose the disabled commands.

Step 6 (Optional.) Configure the router ID.

Click **+** to enable the **router-id** command, then click the variable and enter the IPv4 address that should be used when sending any router updates from this device. No two routers in an OSPF system can have the same router ID, so ensure that it is unique in the area.

If you do not explicitly specify a router ID for the process, the system uses the highest IP address assigned to an active interface. Thus, the router ID can change if you disable the selected interface, or you change its addresses. By assigning a router ID explicitly, you ensure consistency for your process.

Step 7 (Optional.) Configure RFC 1583 compatibility when calculating summary route costs.

Click **+** to enable the **configure summary-route-cost** command, then click the variable and select either **any**, which turns off RFC 1583 compatibility, or **rfc1583**, which turns it on.

Even though this command is not enabled by default in the OSPF object, in fact RFC 1583 compatibility is the default method used when calculating summary route costs. If you examine the configuration defined in the CLI, only the disabled setting is shown.

Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. Ensure that you set RFC 1583 compatibility the same on all OSPF routers in an OSPF routing domain.

Step 8 (Optional.) Suppress syslog messages for multicast OSPF (MOSPF) link state advertisements (LSA).

Click **+** to enable the **ignore lsa mospf** command.

The system does not support LSA Type 6 MOSPF packets. You can enable this command so the system does not send syslog messages when it receives these packets, to reduce the noise in your syslog server.

Step 9 Configure the distance metrics.

The following **distance** commands are enabled by default. You can change OSPF route administrative distances based on route type. The distances are 1 to 255, with the higher numbers being less trusted than lower numbers. These metrics are used to judge the relative value of a learned route when comparing similar routes from different processes.

- **distance ospf inter-area 110**. Click the number and set the distance for all routes from one area to another area.
- **distance ospf intra-area 110**. Click the number and set the distance for all routes within an area.
- **distance ospf external 110**. Click the number and set the distance for routes from other routing domains that are learned by redistribution.

Step 10 Configure the route calculation timers for the OSPF process.

The following timer commands are enabled with these default values.

- **timers lsa arrival 1000**. Click the number and set the minimum interval at which the system accepts the same link-state advertisement (LSA) from OSPF neighbors, from 0 to 600000 milliseconds. Use this command to indicate the minimum interval that must pass between acceptance of the same LSA that is arriving from neighbors. LSAs arriving before this minimum time are ignored.

- **timers pacing flood 33**. Click the number and set the time at which LSAs in the flooding queue are paced in-between updates, from 5 to 100 milliseconds.
- **timers pacing lsp-group 240**. Click the number and set the interval at which OSPF link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, from 10 to 1800 seconds.
- **timers pacing retransmission 66**. Click the number and set the time interval at which LSAs in the retransmission queue are paced, 5 milliseconds to 200 milliseconds. We recommend that you do not change the packet retransmission pacing timer unless all other options to meet OSPF packet flooding requirements have been exhausted. Specifically, configure summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers.
- **timers throttle lsa 0 5000 5000**. Click the numbers and set rate-limiting values for Open Shortest Path First (OSPF) link-state advertisement (LSA) generation. LSA and SPF throttling provide a dynamic mechanism to slow down LSA updates in OSPF during times of network instability and allow faster OSPF convergence. The values are:
  - **Start Interval** (first number)—The minimum delay to generate the first occurrence of an LSA, from 1 to 600000 milliseconds. The first instance of LSA is generated immediately after a local OSPF topology change. The next LSA is generated only after this start interval. Specify 0 to have LSAs generated without delay.
  - **Hold Time** (second number)—The minimum delay to generate the LSA again, from 1 to 600000 milliseconds. This value is used to calculate the subsequent rate limiting times for LSA generation.
  - **Maximum Interval** (third number)—The maximum delay to generate the LSA again, from 1 to 600000 milliseconds.
- **timers throttle spf 5000 10000 10000**. Click the numbers and set rate-limiting values for the shortest path first (SPF) generation. The values are:
  - **Start Interval** (first number)—The delay to receive a change to the SPF calculation, from 1 to 600000 milliseconds.
  - **Hold Time** (second number)—The delay between the first and second SPF calculations, from 1 to 600000 milliseconds.
  - **Maximum Interval** (third number)—The maximum wait time for SPF calculations, from 1 to 600000 milliseconds.

Step 11 (Optional.) Generate a default external route into an OSPF routing domain.

Click + to enable the **default-information originate** command. You can optionally enable and configure the following commands to fine-tune the feature:

- **default-information originate always**. Always advertise a default route even if there is no default route.
- **default-information originate metric 1 metric-type** *metric-type-value*. The metric type and value for generating the default route.
  - Click the **metric** number and enter the OSPF default metric value, from 0 to 16777214. Unless you know you need a different value, enter 10.
  - Click the **metric-type** number and select 1 or 2 as the external link type associated with the default route advertised into the OSPF routing domain. The default is 2.

- **default-information originate route-map** *route-map*. Select a route map that specifies the routing process that generates the default route if the route map is satisfied.

Step 12 (Optional.) Configure Non-Stop Forwarding (NSF) graceful restart if the device is configured for High Availability (HA).

The system can experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a component failure (for example, in HA, the active unit fails over to the standby unit, or in a cluster, the primary unit fails with a secondary unit being elected as new primary), or when there is a scheduled hitless software upgrade.

You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623).

You can configure 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.

- You can configure a device as NSF-aware irrespective of the mode in which it operates.
- A device has to be in either High Availability (failover) or Spanned Etherchannel (L2) cluster mode for you to configure it as NSF-capable.

Note
You must not configure the OSPF process to use fast hello packets if you also configure graceful restart. Graceful restart cannot occur with fast hello packets, because the time taken for the role change between the active and standby units is more than the configured dead interval.

To configure graceful restart:

a) Click **+** to enable the **configure nsf graceful-restart** command.
b) Click the *mechanism* variable and select one of the following:
   - **cisco** to configure an NSF-capable device according to Cisco RFC 4811 and RFC 4812.
   - **ietf** to configure an NSF-capable device according to IETF RFC 3623.
   - **both** to configure the device as an NSF-aware helper rather than an NSF-capable device.
   - **none** to disable graceful restart if you have configured it previously.
c) Your selection in the previous step adds the commands required to implement graceful restart according to your specifications. Do not disable these commands. There is only one command that optionally needs further configuration. Following is an explanation of the added commands; a **no** form of the command turns off the related feature.
   - **nsf cisco helper**. Enable Cisco nonstop forwarding (NSF) helper mode. When the NSF-capable Firewall Threat Defense device is performing graceful restart, the helper Firewall Threat Defense devices assist in the nonstop forwarding recovery process.
   - **nsf ietf helper** *mode-option*. Enable IETF nonstop forwarding (NSF) helper mode. When the NSF-capable Firewall Threat Defense device is performing graceful restart, the helper Firewall Threat Defense devices assist in the nonstop forwarding recovery process. Optionally, you can click *mode-option* and enable strict link-state advertisement (LSA) checking. With strict LSA checking enabled, the helper system will terminate the helping process of the restarting system if it detects

that there is a change to an LSA that would be flooded to the restarting system or if there is a changed LSA on the retransmission list of the restarting system when the graceful restart process is initiated.

- **capability lls**. Enables Link Local Signaling (LLS), which is needed for Cisco graceful restart.
- **capability opaque**. Enables opaque Link State Advertisements (LSAs), which is needed for IETF graceful restart.

Step 13 Click **OK**.

Configure OSPF Area Properties

You can configure several OSPF area parameters. You can define the networks to advertise in the area, plus filtering and virtual links. In addition, these area parameters 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.

When you configure area parameters, you need to know how the system functions within the area.

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 link-state advertisements (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 system 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.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Configure the area number.

a) Click the + to the left of the **area** *area-id* line to enable the command. You cannot configure a command until you enable it.

b) Click *area-id* and enter the number of the area. This area number needs to be the same as the one used by the other routers that define the OSPFv2 area. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.

Step 6 Configure the networks and interfaces that should be routed within the area.

a) Click the + to the left of the **configure area** *area-id options* line.

b) Click *area-id* and enter the same area number from the **area** command.

c) Click *options* and select **properties**. This action adds several lines, including one that is enabled by default, the **network** command.

d) In the **network** command, click *network-object* and select the object that defines a network that should be included in this area. Typically, this would be a directly-connected network. For example, if the IP address of the inside interface is 192.168.1.1/24, the associated network object for this command would contain 192.168.1.0/24. If the object does not already exist, click **Create New Network** and create it now.

e) (Optional.) In the **network** command, click *tag-interface* and select the interface that hosts or routes to the network. If you select the interface, the system can prevent you from changing the address on the interface because it is used in the routing process. This helps remind you that any change to interface addressing can impact your routing configuration.

   If you select an interface here, before you can change the address on an interface, you must first remove it from the routing process. Then, after the changing the IP address, remember to return here and select the new networks and interface to ensure a correctly-configured routing process.

Step 7 (Optional.) Configure the cost for the default summary route sent to a stub area or a not-so-stubby area (NSSA).

This option is meaningful only if you configure the area to be a stub or NSSA, as explained below. Click + to enable to following command in the area properties:

area *area-id* **default-cost 1**

If necessary, enter the correct area ID. Then, click the number and enter the relative cost of the route, from 0 to 16777214. The default is 1. The higher the number, the less likely the route will be used over another route that applies for the destination.

Step 8 (Optional.) Configure prefix filtering for the area.

You can filter prefixes advertised in Type 3 link-state advertisements (LSAs) between OSPFv2 areas of an area border router (ABR). Prefix filtering improves your control of route distribution between OSPF areas. With prefix filtering, you can allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. You can apply this type of area filtering out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF area at the same time.

Before configuring this command, you must create the prefix lists, which are Smart CLI objects, on the **Device** > **Advanced Configuration** page. You can configure separate prefix lists for inbound or outbound advertisements: select the direction for the filter-direction parameter.

area *area-id* **filter-list prefix** *prefix-list filter-direction*

Step 9 (Optional.) Configure the area as a stub 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. You must use default routing in the stub area for it to work properly. To further reduce the number of LSAs sent into a stub area, you can use the **no-summary** keyword of the **area stub** command on the ABR to prevent it from sending a summary link advertisement (LSA Type 3) into the stub area.

To configure the area as a stub:

a) Click the + to the left of the setup area-id as type line.
b) Click *type* and select **stub**. This will add the area stub command after the setup line.
c) Optionally, in the **area stub** command, click *stub-parameters* and select **no-summary**.

Step 10 (Optional.) Configure the area as a not-so-stubby area (NSSA).

A not-so-stubby area (NSSA) is similar to a 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 area border routers (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 by running the connecting area as an NSSA. The connection between the corporate site border router and the remote router cannot be run as an OSPFv2 stub area because routes for the remote site cannot be redistributed into the stub area, which means that two routing protocols would have to be maintained. A simple protocol such as RIP would usually be run to handle 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.

To configure the area as an NSSA:

a) Click the **+** to the left of the **setup** *area-id* **as** *type* line.
b) Click *type* and select **nssa**. This will add several commands after the setup line, including the **area nssa** command, which you must leave enabled.
c) (Optional.) To generate a Type 7 default route into the NSSA, click the + to enable the following command:

   area *area-id* **nssa default-information-originate metric 1 metric-type 2**

   You can optionally adjust the following values:

   - Click the **metric** number and enter the OSPF default metric value, from 0 to 16777214. Unless you know you need a different value, enter 10.
   - Click the **metric-type** number and select 1 or 2 as the external link type associated with the default route advertised into the OSPF routing domain. The default is 2.

d) (Optional.) If the system is an ABR, and you want redistribution from other routing protocols to import routes into normal areas only and not into the NSSA, click the **+** to enable the following command:

   area *area-id* **nssa no-redistribution**

e) (Optional.) If you do not want to inject summary routes into the NSSA, click **+** to enable the following command:

   area *area-id* **nssa no-summary**

Step 11 (Optional.) Configure virtual links for the area.

In OSPF, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link. You can configure virtual links to routers that are connected to the backbone area.

a) Click the **+** to the left of the **configure area** *area-id* **virtual-link** *ip_address option* line.
b) Click *ip_address* and enter the router ID of the router to which you are establishing the virtual link.

c) (Optional.) Click *option* and select **properties** to adjust the following attributes, which all have default values appropriate for most networks. The first part of these commands is omitted, because these are just parameters on the same command:

- **authentication** *auth-type*. Click **+** to enable the command, then click *auth-type* and select **none**, **password**, or **message-digest**. Configure the key options if you select something other than none. The options are the same as those you would configure on an OSPF interface, as explained in Configure OSPFv2 Interface Settings and OSPF Authentication, on page 15. Configure authentication only if the other router uses authentication.
- **hello-interval 10**. Click the number and enter interval between hello packets sent on the interface, from 1 to 65535 seconds.
- **retransmit-interval 5**. Click the number and enter the time between LSA retransmissions for the virtual link, from 1 to 65535 seconds.
- **transmit-delay 1**. Click the number and enter the delay time between when OSPF receives a topology change and when it starts a shortest path first (SPF) calculation, from 0 to 65535 seconds.

d) You can click **...** > **Duplicate** next to the **configure area virtual-link** command to define another virtual link. Define as many as you require.

Step 12 (Optional.) If the system is an area border router (ABR), configure ranges to consolidate or summarize routes for the area.

When you configure the **area range** command, the result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This behavior is called route summarization. You can configure multiple **area range** commands for an area. In this way, OSPF can summarize addresses for many different sets of address ranges.

To configure route summarization:

a) Click the **+** to the left of the **area** *area-id* **range** *network-object range-parameters* line.
b) Click *network-object* and select the network object that defines the address range whose routes you want summarized.
c) (Optional.) Click *range-parameters* and select one of the following attributes:

- **advertise**. Sets the address range status to advertise and generates Type 3 summary link-state advertisements (LSAs). This is the default if you select no option.
- **not-advertise**. Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.

d) You can click **...** > **Duplicate** next to the **area range** command to define another route summarization. Define as many as you require.

Step 13 If you are configuring the process for a multi-area network, mouse over the area to the left of the circled - on the **area** and **configure area** lines and click **...** > **Duplicate**. Then, configure the new area and its networks and other settings as explained above. Repeat this process until you have defined all areas in which this routing process should participate.

Step 14 Click **OK**.

Configure Static OSPF Neighbors

You need to define static OSPF neighbors to advertise OSPF routes over a point-to-point, non-broadcast network, that is, a VPN tunnel.

You do not need to define static neighbors that are on regular broadcast networks, as these routers can form adjacencies themselves.

Before you begin

Determine the interface through which the system should reach the neighbor. You must configure the OSPF settings of this interface before you can define the neighbor router.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit the OSPF interface object, and enable the **ospf network point-to-point non-broadcast** command for the selected interface. Save your changes.

Step 5 Add or edit an OSPF process object.

Step 6 Click **Show Disabled** to expose all commands, then click **+** to enable the **neighbor** command.

Step 7 Configure the neighbor address.

neighbor *ip-address* **interface** *interface*

- Click *ip-address* and enter the IP address of the neighbor router.
- Click *interface* and select the interface through which the system can reach the router.

Step 8 If necessary, configure a static route for the neighbor router.

If the IP address of the router is on the same network as the selected interface, a static route is not necessary. For example, if you select an interface whose IP address is 10.100.10.1/24, and the neighbor address is 10.100.10.2/24, you do not need a static route.

Step 9 You can click **...** > **Duplicate** next to the **neighbor** command to define another static neighbor. Define as many as you require.

Step 10 Click **OK**.

Configure OSPF Summary Addresses

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the system 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. You can suppress routes that match the specified IP address mask pair. You can use the tag value as a match value for controlling redistribution through route maps.

Route summarization is the consolidation of advertised addresses. You can summarize routes learned from other routing protocols. 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.

Before you begin

Create network objects for all addresses that you want to summarize.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure** *network-object* **as** *option* **summary-address** command.

Step 6 Click *network-object* and select the object that defines the address space that you want to summarize.

Step 7 Click *option* and select one of the following:

- **advertising**. Advertise routes that match the address.
- **non-advertising**. Suppress routes that match the address.

Step 8 (Optional.) To add a tag value to the summarized route, click **+** to enable the **summary-address tag** command, click the *tag-number* variable, and enter the tag number, from 0 to 4294967295.

This value is not used by OSPF itself. It may be used to communicate information between autonomous system boundary routers (ASBR). If none is specified, then the remote autonomous system number is used for routes from BGP and EGP; for other protocols, zero (0) is used.

The primary reason to use tag values is for controlling redistribution based on the tag number. If you do not use it in your redistribution route maps, there is no need to configure it here.

Step 9 You can click **...** > **Duplicate** next to the **configure summary-address** command to define another route summarization. Define as many as you require.

Step 10 Click **OK**.

Configure OSPF Filter Rules

Create the Smart CLI standard access list objects you need for each filter rule. Use deny access control entries (ACEs) to filter out routes that match the entry, and permit ACEs for the routes that should be updated.

Before you begin

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

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click **+** to enable the **configure filter-rules** *direction* command.

Step 6 Click *direction* and select **in**, for filtering incoming updates, or **out**, for filtering outbound updates.

Step 7 For inbound filters, you can optionally specify the interface on which to filter updates. If you do not specify an interface, the filter applies to all updates received on any interface.

a) Click **+** to enable the **distribute-list** *acl-name* **in interface** *interface* command.

b) Click the *interface* variable and select the interface.

Step 8 For outbound filters, you can optionally specify the protocol, to limit the filter to routes advertised to that routing process.

There are two forms of the **distribute-list out** command, one with an *identifier* variable after the *protocol* variable, and one without the identifier. You can select the following protocols, but they are divided between these command versions based on whether you must provide the additional identifier information.

- **connected**. For routes established for networks that are directly connected to the system's interfaces.
- **static**. For static routes you manually created.
- **rip**. For routes advertised to RIP.
- **bgp** *automonous-system*. For routes advertised to BGP. Click *identifier* and enter the autonomous system number for the BGP process defined on the system.
- **eigrp** *automonous-system*. For routes advertised to EIGRP. Click *identifier* and enter the autonomous system number for the EIGRP process defined on the system.
- **ospf** *process-id*. For routes advertised to OSPF. Click *identifier* and enter the process ID for the other OSPF process defined on the system.

Step 9 You can click **...** > **Duplicate** next to the **configure filter-rules** command to define another filter rule. Define as many as you require.

Step 10 Click **OK**.

Configure OSPF Redistribution

You can control the redistribution of routes into an OSPF process from other routing protocols, connected routes, and static routes.

Before you begin

It is best practice to configure the routing process from which you will redistribute routes, and deploy your changes, before you configure redistribution into OSPF.

If you want to apply a route map to fine-tune which routes are redistributed, create the Smart CLI route map object. Routes that match the route map are redistributed, and all non-matching routes are not redistributed.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Add or edit an OSPF process object.

Step 5 Click **Show Disabled** to expose all commands, then click + to enable the **configure redistribution** command.

Step 6 Click the *protocol* variable and select the source process from which you are redistributing routes. You can redistribute **connected** and **static** routes, or routes generated by **bgp**, **eigrp**, **isis**, **ospf**, or **rip**.

Step 7 If you select a routing process, click the *identifier* variable and enter the required value:

- **bgp**, **eigrp**. Enter the autonomous system number.
- **ospf**. Enter the process ID number.
- **connected**, **static**, **isis**, **rip**. Enter **none**. Even if you enter a different value, it will be ignored.

Step 8 (Optional; IS-IS only.) On the **redistribute isis level-2** command, click **level-2** and select whether you are redistributing routes learned only within an IS-IS area (**level-1**), between IS-IS areas (**level-2**) or both (**level-1-2**).

Step 9 (Optional; all protocols.) If you apply tags to routes in order to control redistribution, click + to enable the **redistribute tag** *tag-number* command, then click the variable and enter the tag associated with routes you want to redistribute. The tag number is from 0 to 4294967295.

Step 10 (Optional; all protocols.) If you want to redistribute routes for all subnets, not just those that abide by the standard class, click + to enable the **redistribute subnets** command.

For example, if you do not enable this command, a specific route for 10.100.10.0/24 would not be redistributed; instead, only a route for 10.0.0.0/8 would be redistributed.

Step 11 (Optional; all protocols.) To fine-tune which routes are redistributed based on a route map, click + to enable the **redistribute** *route-map* command, click the variable, and select the route map that defines your restrictions.

If you do not apply a route map, all routes for the process (that fit the other commands configured for redistribution), are redistributed.

Step 12 (Optional; all protocols.) To fine-tune the metrics for redistributed routes, click + to enable the following command and configure the options:

redistribute *protocol* **metric** *metric-value* **metric-type** *metric-type-value*

Click the variables and configure the following:

- **metric**. The metric value for the routes being distributed, from 0 to 16777214. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if you do not specify a metric value. When redistributing other processes to an OSPF process, the default metric is 20.
- **metric-type**. The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route. The default is 2.

Step 13 (Optional; OSPF only.) The following commands are enabled by default when you redistribute routes from another OSPF process. You can click - to disable unwanted commands.

These commands specify the criteria by which OSPF routes are redistributed into other routing domains.

- **redistribute ospf match external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- **redistribute ospf match external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- **redistribute ospf match internal**. Routes that are internal to a specific autonomous system.
- **redistribute ospf match nssa-external 1**. Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes and marked as Not-So-Stubby-Area (NSSA) only.
- **redistribute ospf match nssa-external 2**. Routes that are external to the autonomous system, but are imported into OSPF as Type 2external routes and marked as Not-So-Stubby-Area (NSSA) only.

Step 14 You can click **...** > **Duplicate** next to the **configure redistribution** command to configure redistribution for another protocol. Configure redistribution for each protocol that makes sense for your network.

Step 15 Click **OK**.

Configure OSPFv2 Interface Settings and OSPF Authentication

Any interface that faces a neighbor OSPF router communicates with the router using hello packets and other methods to verify the health of the neighbor and to share routing updates. While some of these characteristics have default settings, it is best practice to set the options explicitly using an OSPF Interface Settings object. Create an object for each interface that is adjacent to an OSPF neighbor router.



Note The routers on a given network must have the same values for authentication and for lost neighbor detection hello and dead intervals.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new object, click **+** > **OSPF Interface Settings** or click the **Create OSPF Object** > **OSPF Interface Settings** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

If you no longer need an interface settings object, click the trash can icon for the object to delete it.

Step 5 Enter a name for the object, and optionally, a description.

Step 6 Configure authentication for the interface.

configure authentication *auth-type*

To configure OSPF authentication, you must configure the password or authentication key on each of the OSPF interfaces, then enable authentication on the area itself. You must choose the same authentication method on the interfaces and the area.

You can select the following options by clicking *auth-type*.

- **none**—Do not use OSPF authentication. Any OSPF router operating on the link can establish an adjacency with this router. The following command is added to the object: **ospf authentication null**.
- **password**—Authenticate the OSPF connection using a shared password. You can configure a separate password to each network on a per-interface basis. However, all neighboring routers on the same network must have the same password to be able to exchange OSPF information.

  When you select this option, two commands are added: **ospf authentication** and **ospf authentication-key** *key*. Click the variable to configure the following:

  - *key*—Select the secret key object that contains the password. The password can be up to 8 characters. You can include blank space between two characters. Spaces at the beginning or end of the password are ignored. If the object does not yet exist, click **Create New Secret Key** at the bottom of the list and create it now.
- **message-digest**—Authenticate the OSPF connection using message digest (MD5). MD5 authentication verifies the integrity of the communication, authenticates the origin, and checks for timeliness. Both routers must be configured to use the same MD5 key.

  When you select this option, two commands are added: **ospf authentication message-digest** and **ospf message-digest-key** *key-id* **md5** *key*. Click the variables to configure the following:

  - *key-id*—The authentication key ID number, from 1 to 255. You must configure the neighbor router with the same key ID and associated MD5 key.

- *key*—Select the secret key object that contains the MD5 key. The key is an alphanumeric password up to 16 characters. You can include spaces between characters. Spaces at the beginning or end of the key are ignored. If the object does not yet exist, click **Create New Secret Key** at the bottom of the list and create it now.

Step 7 (Optional.) Configure link-state advertisement (LSA) timers.

These timers have default values, so you need to change them only if your network requires different settings. Configure the following commands:

- **ospf retransmit interval 5**—The number of seconds between LSA retransmissions for adjacencies belonging to an OSPF interface. The number of seconds must be greater than the expected round-trip delay between any two routers on the attached network. The range is from 1 to 8192 seconds. The default value is 5 seconds. Click the 5 and type in a new number to change the value.
- **ospf transmit-delay 1**—The estimated number of seconds required to send a link-state update packet on an OSPF interface, from 1 to 8192 seconds. The default value is 1 second. Click the 1 and type in a new number to change the value.

Step 8 (Optional.) All other settings have default values or they are optional. Change them, or enable them, only if you need different behavior. Click the **Show Disabled** link to expose the options.

Following are the additional interface settings. To enable a setting, click the + to the left of the command, then configure the command (if necessary).

- **ospf cost** *value*—The cost (a link-state metric) of sending a packet on an OSPF interface, from 1 to 65535. The value 1 represents a network that is directly connected to the interface. Click the variable and enter the cost that represents the capability of the interface based on the numbers you are using in your network.

  When deciding on a value, the higher the interface bandwidth, the lower the associated cost to send packets across that interface should be. In other words, a large cost value represents a low bandwidth interface and a small cost value represents a high bandwidth interface. The specific number you select has no inherent meaning: the value is relative to the other values you configure for interface across the OSPF area. These values then affect the calculation of the best route for a destination.

  The OSPF interface default cost on the Firewall Threat Defense device is 10. This default differs from Cisco IOS software, where the default cost is 1 for Fast Ethernet and Gigabit Ethernet and 10 for 10BaseT. This is important to take into account if you are using ECMP in your network.
- **ospf database-filter all out**—Filters out all outgoing LSAs to an OSPF interface during synchronization and flooding.
- **ospf mtu-ignore**—Disables OSPF maximum transmission unit (MTU) mismatch detection on receiving database packets. OSPF checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange Database Descriptor (DBD) packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency will not be established. If you cannot fix the MTU values on the interfaces to be the same, you can disable MTU checking.
- **ospf network point-to-point non-broadcast**—Configures the OSPF interface as a point-to-point, non-broadcast network. This lets you transmit OSPF routes over VPN tunnels. If you configure this option, dynamic discover of neighbors is not possible. You must also:

- Update the OSPF process object to define a single static neighbor for this interface. Also, update the OSPF process of the neighbor router to define this device as its static neighbor.
- Create static routes (on each router) that point to the neighbor router.

- **ospf priority** *value*—The priority of the router relative to the other routers in the network, from 0 to 255. The default priority is 1. When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Click the variable and select the priority based on the relative numbering system you use in your network.
- **ospf lost-neighbor-detection** *detection-mechanism*—Defines how the system determines if a neighbor router is down. OSPF must recalculate routes whenever an OSPF router is declared down. For detailed information on configuring lost neighbor detection, see Configure OSPFv2 Lost Neighbor Detection and Fast Hello Packets (OSPF Interface Settings), on page 18.

Step 9 Click **OK**.

Configure OSPFv2 Lost Neighbor Detection and Fast Hello Packets (OSPF Interface Settings)

The OSPF process regularly sends hello packets to each neighbor router to verify that the neighbor can still respond. Continued failure to respond indicates that the neighbor router (either entirely or just the adjacent interface) is not available for routing, and OSPF must recalculate routes and the OSPF system must converge on the updated routing table.

You can adjust the following values to fine-tune your network. Ideally, you want to minimize how often neighbors are declared down and routes recalculated. On the other hand, you also want to minimize how long it takes for the network to reconverge on a good routing table when an OSPF router (or interface) truly is down.

- Hello interval—This is the time between sending hello packets. The default is every 10 seconds. If desired, you can configure fast hello packets, where hellos are sent at sub-second intervals. Fast hello packets provide the quickest detection of a down neighbor and reconvergence of the routing table.
- Dead interval—The length of time during which, if no hello packets are seen from a neighbor, the neighbor is declared dead. The default is 40 seconds (4 times the default hello interval), unless you are using fast hello packets, in which case the dead interval is always 1 second. Specifying a smaller dead interval will give faster detection of a neighbor being down and improve convergence, but might cause more routing instability. In any case, you must configure the dead interval to be larger than the hello interval. You must set the same dead interval on all OSPF routers in the network.

You configure lost neighbor detection in the OSPF Interface Settings object.

Procedure

Step 1 Click **Device**, then click the **Routing** summary.

Step 2 If you enabled virtual routers, click the view icon ( ) for the router in which you are configuring OSPF.

Step 3 Click the **OSPF** tab.

Step 4 Do one of the following:

- To create a new object, click **+** > **OSPF Interface Settings** or click the **Create OSPF Object** > **OSPF Interface Settings** button.
- Click the edit icon ( ) for the object you want to edit. Note that when you edit an object, you might see lines that you did not directly configure. These lines are exposed to show you the default values that are being configured.

Step 5 If the **ospf lost-neighbor-detection** *detection-mechanism* command is not displayed, click the **Show Disabled** link.

Step 6 Click the **+** to the left of the command to enable it.

Step 7 Click *detection-mechanism* and select the mechanism that you want to implement:

- **dead-interval**—To configure a standard hello interval in seconds. The following commands are added; adjust their values as needed:
  - **ospf hello-interval 10**—The hello interval, from 1 to 8192 seconds. The default is 10. This value must be less than the dead interval. Click the value to enter the desired number.
  - **ospf dead-interval 40**—The dead interval, from 1 to 8192 seconds. The recommended value is 4 times the hello interval, but you can configure a shorter time for faster convergence.
- **hello-multiplier**—To configure sub-second fast hello packets. The following command is added, you must configure the value.
  - **ospf dead-interval minimal hello-multiplier** *value*—Click the variable and enter the number of hello packets that should be sent each second, from 3 to 20. The dead interval is set to 1 second by the **minimal** keyword.

Step 8 Click **OK**.

Monitoring OSPF

To monitor and troubleshoot OSPF, open the CLI console or log into the device CLI and use the following commands. You can also select some of these commands from the **Commands** menu on the Routing page.

Use **show ospf ?** to get lists of additional options. For example, you can specify process ID, area ID, and virtual router to limit the information you see, as well as other options to target just the information you are looking for. The following list is a summary only.

- **show ospf**

  Displays general information about OSPFv2 routing processes.

- **show ospf border-routers**

  Displays the internal OSPFv2 routing table entries to the ABR and ASBR.

- **show ospf database**

  Displays lists of information related to the OSPFv2 database for a specific router.

- **show ospf events**

  Displays OSPF internal event information.

- **show ospf flood-list**

  Displays a list of LSAs waiting to be flooded over an interface, to observe OSPF v2packet pacing. OSPFv2 update packets are automatically paced so they are not sent less than 33 milliseconds apart. Without pacing, some update packets could get lost in situations where the link is slow, a neighbor could not receive the updates quickly enough, or the router could run out of buffer space.

  Pacing is also used between resends to increase efficiency and minimize lost retransmissions. You also can display the LSAs waiting to be sent out of an interface. Pacing enables OSPFv2 update and retransmission packets to be sent more efficiently.

- **show ospf interface**

  Displays OSPFv2-related interface information.

- **show ospf neighbor**

  Displays OSPFv2 neighbor information on a per-interface basis.

- **show ospf nsf**

  Displays the OSPFv2 related Non-Stop Forwarding (NSF) information.

- **show ospf request-list**

  Displays a list of all LSAs requested by a router.

- **show ospf retransmission-list**

  Displays a list of all LSAs waiting to be resent.

- **show ospf rib**

  Displays the OSPF Router Information Base (RIB).

- **show ospf statistics**

  Displays various OSPF statistics, such as the number of times SPF was executed, the reasons, and the duration.

- **show ospf summary-addresses**

  Displays a list of all summary address redistribution information configured under an OSPFv2 process.

- **show ospf traffic**

  Displays a list of different types of packets being sent or received by a specific OSPFv2 instance.

- **show ospf virtual-links**

  Displays OSPFv2-related virtual links information.