



Configuring Open Shortest Path First (OSPF)

This chapter provides configuration and viewing information for the Open Shortest Path First (OSPF) protocol. Information in this chapter applies to all CSS models, except where noted.



Note

The CSS supports OSPF Version 2, as defined in RFC 2178. For detailed information about OSPF MIB objects, refer to RFC 1850.

This chapter contains the following major sections:

- [OSPF Overview](#)
- [CSS OSPF Quick Configuration](#)
- [Configuring OSPF on the CSS](#)
- [Configuring OSPF on a CSS IP Interface](#)
- [Showing OSPF Information](#)
- [OSPF Configuration in a Startup-Config File](#)

OSPF Overview

OSPF is a link-state routing protocol that:

- Provides network topology discovery within a group of routers and networks called an autonomous system (AS)
- Calculates the shortest path to destinations within the AS

As a link-state protocol, OSPF routers flood any change in routing information throughout the network. This action differs from a distance vector protocol, such as RIP, which periodically exchanges routing information only with neighboring devices.

Within an AS, each OSPF router builds and synchronizes a database of the AS network topology. The routers synchronize their databases by requesting information from other AS routers. Each router sends its information as link-state advertisements (LSAs) that include information about the state of each router and link in the AS. A link is an interface on the router. The state of the link is the description of the interface, including the router's IP address and subnet mask, and its relationship to the neighboring router.

Then, the router uses its database and the Shortest Path First (SPF) algorithm to calculate the shortest path to every destination in the AS and stores this information in a dynamic table. When changes occur, the router calculates new paths.

The CSS, operating as an OSPF router, provides:

- Intra-area route support for routing in a single area between other OSPF routers
- Inter-area route support for routing between multiple OSPF areas
- Route summarization between areas as an Area Border Router (ABR)
- Stub area and AS boundary router support
- Redistribution of local, RIP, static, and firewall routes into an OSPF domain
- Advertisement of VIP addresses for content as AS external routes
- Simple authentication

This section includes the following topics:

- [OSPF Routing Hierarchy](#)
- [Link-State Databases](#)

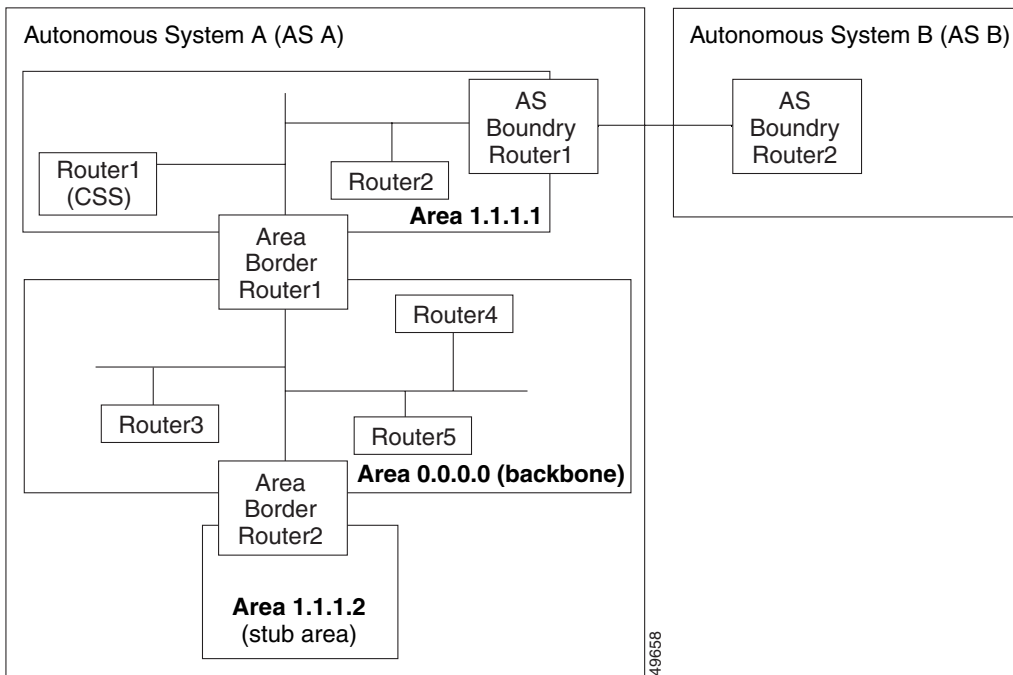
OSPF Routing Hierarchy

The OSPF routing hierarchy includes the following functions:

- Autonomous systems
- Areas, including the backbone and stub areas
- Area Border Routers (ABRs)
- Autonomous System Boundary Routers (ASBRs)

Figure 7-1 illustrates an OSPF network topology.

Figure 7-1 Basic OSPF Network Topology



Autonomous System

The autonomous system (AS) is a collection of networks, under the same administrative control, that share the same routing information with each other. An AS is also referred to as a routing domain. [Figure 7-1](#) shows two ASs: AS A and AS B. An AS can consist of one or more OSPF areas.

Areas

Areas allow the subdivision of an AS into smaller, more manageable networks or sets of adjacent networks. As shown in [Figure 7-1](#), AS A consists of three areas: area 0.0.0.0, area 1.1.1.1, and area 1.1.1.2.

OSPF hides the topology of an area from the rest of the AS. An area's network topology is visible only to routers inside that area; the network topology is not visible to routers outside the area. When OSPF routing is within an area, this is called intra-area routing. This routing limits the amount of link-state information flooding onto the network, thereby reducing routing traffic. OSPF routing also reduces the size of the topology information in each router, which conserves processing and memory requirements in each router.

Conversely, the routers within an area cannot see detailed network structures outside the area. Because of this restriction of topological information, you can control traffic flow between areas and reduce routing traffic when the entire autonomous system is a single routing domain.

Backbone Area

A backbone area is responsible for distributing routing information between the areas of an autonomous system. When OSPF routing occurs outside of an area, this is called inter-area routing.

The backbone itself has all the properties of an area. It consists of ABRs, and routers and networks only on the backbone. As shown in [Figure 7-1](#), area 0.0.0.0 is an OSPF backbone area. Note that a designated OSPF backbone area has a reserved ID of 0.0.0.0.

Area Border Routers

ABRs have multiple interfaces that connect directly to networks in two or more areas. An ABR runs a separate copy of the OSPF algorithm and maintains separate routing data for each area that is connected to it, including the backbone area. ABRs also send configuration summaries for their attached areas to the backbone area, which distributes this information to other OSPF areas in the autonomous system. In [Figure 7-1](#), there are two ABRs. ABR 1 interfaces area 1.1.1.1 to the backbone area. ABR 2 interfaces the backbone area to area 1.1.1.2, a stub area.

**Note**

ABRs are always backbone routers. You must configure ABRs to the backbone area.

Stub Area

A stub area is an area that does not accept or distribute detailed network information external to the area. A stub area has only one router that interfaces the area to the rest of the AS. The ABR attached to the stub area advertises a single default external route into the area. Routers within a stub area use this route for destinations outside the autonomous system, as well as for inter-area routes. This relationship conserves LSA database space that would otherwise be used to store external LSAs flooded into the area. As shown in [Figure 7-1](#), area 1.1.1.2 is a stub area that is reached only through ABR 2.

Autonomous System Boundary Routers

ASBRs provide connectivity from one autonomous system to another system. ASBRs exchange their autonomous system routing information with boundary routers in other autonomous systems. Every router inside an autonomous system knows how to reach the boundary routers for its autonomous system.

ASBRs can import external routing information from other protocols like RIP and redistribute them as AS-external LSAs to the OSPF network. If the CSS is an ASBR, you can configure it to advertise VIP addresses for content as AS external routes. In this way, ASBRs flood information about external networks to routers within the OSPF network.

ASBR routes can be advertised as type1 or type2 ASE. The difference between type1 and type2 is how the cost is calculated. For a type2 ASE, only the external cost (metric) is used when comparing multiple paths to the same destination. For type1 ASE, the combination of the external cost and the cost to reach the ASBR is used.

Link-State Databases

OSPF routers advertise routes using LSAs. The link-state database stores the LSAs from routers throughout the area. The advertisements depict the topology of the autonomous system. They could include:

- Router links that describe the state and cost of each router's interface to an area
- Network links from the designated router (see the [“Setting the Priority of the CSS”](#) section) that describe all routes on a segment for multi-access segments with more than one attached router
- Summarized links from ABRs that describe networks in the AS but outside an area
- External links from ASBRs that describe destinations external to the AS

All routers that are connected to an area maintain identical routing databases about the area. Routers that are connected to multiple areas maintain a separate routing database for each attached area.

Instead of each router sending routing information to every other router on the network, OSPF routers establish adjacencies among neighboring routers. When the link-state databases of two neighboring routers are synchronized, they are considered adjacent.

OSPF routers collect raw topological data from the LSAs that they receive. Each router then prunes this data down to a tree of the shortest network paths centered on itself. The router examines the total cost to reach each router or network node in its domain. By discarding all but the lowest-cost path to each destination, the router builds a shortest-path tree to each destination, which it uses until the network topology changes. It is possible to have multiple lowest-cost paths to a destination.

CSS OSPF Quick Configuration

This section includes the following topics:

- [Global OSPF Quick Configuration](#)
- [OSPF IP Interface Quick Configuration](#)
- [Verifying Your Quick Configuration](#)

Global OSPF Quick Configuration

To perform the global OSPF configuration for the CSS, see the steps in [Table 7-1](#). In the most basic global configuration, where the CSS functions as a router in the OSPF backbone area, you need to perform only steps 1 and 2 to:

- Define the CSS router ID
- Enable OSPF

Optionally, you can define the CSS:

- In an area other than the backbone, including a stub area.
- As an ABR, by configuring route summarization.
- As an ASBR, to advertise non-OSPF routes through OSPF, as AS-external routes such as static and RIP routes. You could also advertise VIP addresses for content as AS external routes.

After performing the global OSPF configuration, you must configure an OSPF IP interface (see the [“OSPF IP Interface Quick Configuration”](#) section) before the CSS can participate in OSPF routing. For more information on configuring global OSPF parameters, see the [“Configuring OSPF on the CSS”](#) section.

Table 7-1 CSS Global OSPF Quick Configuration**Task and Command Example**

1. Configure the area router ID for the CSS in global configuration mode. In this example, the CSS router ID is 121.23.21.1.

```
(config) ospf router-id 121.23.21.1
```

2. Enable OSPF on the CSS.

```
(config) ospf enable
```

3. (Optional) If the CSS area is other than the backbone area, enter the area ID for the CSS. In this example, the area ID is 1.1.1.1.

```
(config) ospf area 1.1.1.1
```

The default ID is 0.0.0.0 for the backbone area. To define a stub area, enter the stub option after the area ID.

4. (Optional) If you want the CSS to advertise external routes, define the CSS as an AS boundary router. For example:

```
(config) ospf as-boundary
```

5. (Optional) If the CSS is an ABR, you can advertise VIP addresses for content as OSPF ASE routes. To advertise the VIP address 192.168.4.15 with a default cost of 1 and the default type of ASE type2, enter:

```
(config) ospf advertise 192.168.4.15 255.255.255.255
```

6. (Optional) To advertise routes other than OSPF, such as a firewall, local, RIP or static route, configure OSPF to redistribute routes from the specific protocol. To advertise static routes through OSPF with a default cost of 1 and default type of ASE type2, enter:

```
(config) ospf redistribute static
```

After you complete the global OSPF configuration, configure OSPF on CSS IP interfaces as described in the [“OSPF IP Interface Quick Configuration”](#) section.

OSPF IP Interface Quick Configuration

To configure OSPF on a CSS IP interface, see the steps in [Table 7-2](#). In the most basic IP interface configuration, you need to perform only steps 1 through 4, and step 7 to:

- Assign OSPF to the IP interface
- Associate OSPF with the globally defined area, if this is an area other than the backbone area (0.0.0.0)
- Enable OSPF on the interface

This configuration example assumes you will accept the default OSPF configuration settings for the interface, except the router priority. The interface OSPF configuration settings include:

- Intervals for the hello packet, LSA retransmission, and link-state update packet
- Authentication password
- CSS router priority
- Interface cost

For more information on configuring these OSPF IP interface settings, see the [“Configuring OSPF on a CSS IP Interface”](#) section.

Table 7-2 Quick Configuration for OSPF on a CSS Interface

Task and Command Example

1. Access global configuration mode. Enter:

```
# config
```

2. Access the circuit configuration mode for a preconfigured circuit on which you want to create the IP interface. For example, if circuit VLAN6 already exists, enter:

```
(config)# circuit VLAN6  
(config-circuit[VLAN6])#
```

Note Refer to [Chapter 5, Configuring Interfaces and Circuits](#) for information on how to configure the CSS interfaces and circuits and the bridge interfaces to VLANs.

Table 7-2 Quick Configuration for OSPF on a CSS Interface (continued)**Task and Command Example**

3. Create the IP interface to the circuit. To create an IP address of 3.1.2.2 with a subnet mask of /24, enter:

```
(config-circuit[VLAN6])# ip address 3.1.2.2/24
Create ip interface <3.1.2.2>, [y/n]: y
```

4. Configure the IP interface as an OSPF interface. Enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf
```

5. (Optional) If the globally configured area is other than the backbone area, enter the configured area ID. In this example, the globally configured area ID is 1.1.1.1.

```
(config-circuit-ip[VLAN6-3.1.2.2]) ospf area 1.1.1.1
```

6. (Optional) With a default setting of 1, the CSS is set to a priority that allows it to become the designated router. If you do not want the CSS to become the designated router, you can change its priority or disable it from eligibility. For example, if you want the CSS to be ineligible to become a designated router, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf priority 0
```

For more information on designated routers, see the [“Setting the Priority of the CSS”](#) section.

7. Enable OSPF on the interface. Enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf enable
```

Verifying Your Quick Configuration

To verify the OSPF global and interface quick configurations, use the **show ospf** command and its options. For example:

- To show the OSPF global configuration, use the **show ospf global** command. For example:

```
# show ospf global
```

If the Admin Status field is disabled, use the **ospf enable** command to enable OSPF.

- To show the route redistribution policy into OSPF, use the **show ospf redistribute** command. To show the configured static route redistribution policy, enter:

```
# show ospf redistribute
```

- To show the VIP addresses advertised as ASE routes, use the **show ospf advertise** command. For example:

```
# show ospf advertise
```

- To view the CSS IP interface configuration, use the **show ospf interfaces** command. For example:

```
# show ospf interfaces
```

Configuring OSPF on the CSS

This section includes the following topics:

- [Configuring the OSPF Router ID](#)
- [Enabling OSPF](#)
- [Configuring an Area](#)
- [Configuring Equal-Cost Routes](#)
- [Configuring Summarized Routes at an ABR](#)
- [Configuring the CSS as an Autonomous System Boundary Router](#)

Configuring the OSPF Router ID

Before you enable OSPF on the CSS, configure the router ID. Assigning a router ID to the CSS uniquely identifies it to other routers within the autonomous system. In addition, in the case of a priority tie when determining which router is the designated router, the ID serves as a tie-breaker in the designated router election. For more information on designated routers, see the [“Setting the Priority of the CSS”](#) section.

Use the **ospf router-id** command to configure the OSPF router ID for the CSS. A router ID is a 32-bit number in dotted-decimal notation.

To assign the router ID of 121.23.21.1 to the CSS, enter:

```
(config)# ospf router-id 121.23.21.1
```



Note

If OSPF is globally enabled, use the **no** form of the **ospf enable** command to disable OSPF and change the router ID.

To delete the router ID on the CSS, disable OSPF and enter:

```
(config)# no ospf router-id
```

Enabling OSPF

After you assign the router ID to the CSS, globally enable OSPF on the CSS. Use the **ospf enable** command to enable OSPF. For example:

```
(config)# ospf enable
```

To disable OSPF, enter:

```
(config)# no ospf enable
```

Configuring an Area

By default, the CSS is configured to the backbone area automatically. The backbone area has a reserved ID of 0.0.0.0. If the CSS is part of an area other than the backbone area, assign the CSS to that area.

Use the **ospf area** command to assign an area. Enter the ID in dotted-decimal notation (for example, 0.0.0.1). Although an area ID has the same form as an IP address, the area ID address space is its own distinct address space.

For example, if the CSS is in area 0.0.0.1, enter:

```
(config)# ospf area 0.0.0.1
```

If the CSS is in a stub area, include the **stub** option.

For example, if area 0.0.0.1 is a stub area, enter:

```
(config)# ospf area 0.0.0.1 stub
```

Optionally, for a stub area you can also:

- Set a metric for the default route advertised in the stub area.
- Propagate summary LSAs into the stub area.

To set a metric for the default route advertised in the stub area, include the **default-metric** option. By default, the metric equals the smallest metric among the interfaces to other areas. You can assign an integer from 1 to 16777215.

For example, to assign a metric of 200, enter:

```
(config)# ospf area 0.0.0.1 stub default-metric 200
```

To propagate summary LSAs in the stub area, include the **send-summaries** option. For example:

```
(config)# ospf area 0.0.0.1 stub send-summaries
```

Removing an Area

To remove an OSPF area, disable OSPF, then use the **no** form of the **ospf area** command. For example:

```
(config)# no ospf enable  
(config)# no ospf area 0.0.0.1
```

Configuring Equal-Cost Routes

By default, the OSPF CSS is configured to use 15 equal-cost routes. Use the **ospf equal-cost** command to change the number of routes. Enter a number from 1 to 15.

To configure 10 equal-cost routes for use by the CSS, enter:

```
(config)# ospf equal-cost 10
```

To reset the equal-cost routes to its default value of 15, enter:

```
(config)# no ospf equal-cost
```

Configuring Summarized Routes at an ABR

If the CSS is an ABR, you can configure it to advertise a single summary route or network ranges that cover all the individual networks within the specified range. This summarization helps control routing table sizes and prevents the constant changing of routes whenever an interface within an area comes online or goes offline. These route changes do not cause route changes in backbone ABRs and other area routers.

Use the **ospf range** command to specify the IP address range to summarize routes at the ABR. This summarization applies to inter-area paths that are paths to destinations in other OSPF areas. You can also determine whether you want to advertise this range. Disable OSPF before you enter the **ospf range** command.

Define an address range by specifying an IP address and subnet mask that represents networks in the area being summarized. Enter the IP address and subnet mask in dotted-decimal notation (for example, 192.168.128.0 255.255.224.0). You can also enter the mask in CIDR bit-count notation format (for example, /24).

To configure the CSS as an ABR with an area ID of 0.1.0.1 with a collection of destinations between 192.168.0.0 and 192.168.255.255, enter:

```
(config)# no ospf enable
(config)# ospf range 0.1.0.1 192.168.0.0 255.255.0.0
```

To remove the range, enter:

```
(config)# no ospf range 0.1.0.1 192.168.0.0 255.255.0.0
```

By default, the ABR advertises this range. If you want to hide the range from the rest of the AS, include the **block** option. For example:

```
(config)# ospf range 0.1.0.1 192.168.0.0 255.255.0.0 block
```

Configuring the CSS as an Autonomous System Boundary Router

Use the **ospf as-boundary** command if you want the CSS to be an ASBR that exchanges routing information with routers belonging to other autonomous systems. Disable OSPF before you enter the **ospf as-boundary** command.

For example:

```
(config)# no ospf enable
(config)# ospf as-boundary
```

To remove the CSS as an AS boundary router, enter:

```
(config)# no ospf as-boundary
```

To advertise a route as OSPF ASE through all OSPF interfaces or generate a default route, see the following sections.

- [Advertising a Route as an OSPF ASE Route](#)
- [Advertising a Default ASE Route](#)
- [Advertising Other Routes Through OSPF](#)

Advertising a Route as an OSPF ASE Route

The ASBR can perform external route summarization to consolidate multiple routes into a single advertisement. For a CSS, this consolidation is useful when you want to advertise VIP addresses for content as OSPF AS external (ASE) through all OSPF interfaces. Use the **ospf advertise** command to advertise a route as OSPF ASE through all OSPF interfaces. To stop the advertisement of the route, use the **no** form of the **ospf advertise** command (as described later in this section).

When you configure OSPF advertising of an IP address as an ASE route, note that:

- If the advertised IP address is the redundant VIP in a VIP redundancy virtual router configuration environment, OSPF advertises the VIP address when the virtual router state is Master, but OSPF stops advertising this VIP address when the virtual router state is Down or Backup.
- If the advertised IP address is a service IP in a service record configuration, OSPF advertises the IP address when the service state is Alive, Dying, or Suspend, but OSPF stops advertising this IP address when the service state is Down.
- If the advertised IP address is the service IP for a critical service, in which the virtual router has a dependency on the critical service, OSPF treats the IP address the same as the service IP described above. The VIP redundancy state does not have an impact on whether OSPF advertises the IP address; only the service state determines the impact on OSPF advertising.
- If OSPF is configured to advertise an IP address that does not match any service IP and redundant VIP, OSPF advertises the IP address all the time. If this address is later used as a service IP, OSPF determines whether to advertise the IP address based on the service state. When the service state for the service IP is Down, OSPF does not advertise the service IP. In addition, OSPF does not revert the service IP back to the original IP address and resume advertising it when the IP address does not overlap with a service IP.

First, before you enter the **ospf advertise** command, configure the CSS as an ASBR. For more information, see the “[Configuring the CSS as an Autonomous System Boundary Router](#)” section.

Define an address range for the **ospf advertise** command by specifying an IP address and subnet mask that represents networks in the area being summarized. Enter the IP address and subnet mask in dotted-decimal notation (for example, 192.168.128.0 255.255.224.0). You can also enter the mask in CIDR bit-count notation format (for example, /24).

For example, to advertise VIP addresses from 192.168.44.0 to 192.168.44.255, define the range by entering the IP address and subnet mask of 192.168.44.0 255.255.255.0:

```
(config)# ospf advertise 192.168.44.0 255.255.255.0
```

Optionally, you can define any of the following:

- The network cost for the route by including the **metric** option. Enter a number from 1 to 16777215. The default is 1.
- A 32-bit tag value to advertise each external route by including the **tag** option. The 32-bit tag value is not used by the OSPF protocol itself. You can use the tag value to communicate information between ASBRs.
- The advertised routes as ASE type1 by including the **type1** option. By default, the type is ASE type2. The difference between type1 and type2 is how the cost is calculated. For a type2 ASE, only the external cost (metric) is used when comparing multiple paths to the same destination. For type1 ASE, the combination of the external cost and the cost to reach the ASBR is used.

For example:

```
(config)# ospf advertise 193.23.44.0 255.255.255.0 metric 3 type1
```

To stop advertising of the route as OSPF ASE through all OSPF interfaces, enter:

```
(config)# no ospf advertise 193.23.44.0 255.255.255.0
```

The following running configuration example illustrates the **ospf advertise** command for OSPF advertising of VIP addresses and an IP address. Comments are preceded by an exclamation point (!).

```

!***** GLOBAL *****
ospf enable

ospf advertise 1.1.1.10
!advertise redundant VIP
ospf advertise 2.1.1.1
!advertise IP address of service s1
ospf advertise 1.1.1.100
!advertise IP address of critical service c100
ospf advertise 99.99.99.99
!advertise simple IP address, not tied to anything
record

!***** CIRCUIT *****
circuit VLAN1

ip address 1.1.1.200 255.0.0.0
ip virtual-router 1
ip redundant-vip 1 1.1.1.10
!redundant VIP
ip critical-service 1 c100

!***** SERVICE *****
service c100
ip address 1.1.1.100
!IP address for critical service
active

service s1
ip address 2.1.1.1
!IP address for service s1
keepalive method get
keepalive type http
active

service s2
ip address 2.1.1.2
keepalive method get
keepalive type http
active

```

```
!***** OWNER *****
owner admin1

content r1
add service s1
add service s2
vip address 1.1.1.10
!redundant VIP equals content VIP

active
```

Advertising a Default ASE Route

Routers use default routes when no additional routes exist to a particular AS external destination. By default, an ASBR does not generate a default route into the OSPF routing domain. Use the **ospf default** command to force the CSS to generate a default ASE route and advertise the route through OSPF.

Before you enter the **ospf default** command, configure the CSS as an ASBR. For more information, see the [“Configuring the CSS as an Autonomous System Boundary Router”](#) section.

For example:

```
(config)# ospf default
```

Optionally, you can define any of the following:

- The network cost for an OSPF default route by including the **metric** option. If a default route metric is defined, the router advertises itself as the default router to the area. Enter a number from 1 to 16,777,215. The default is 1.
- A 32-bit tag value to advertise each external route by including the **tag** option. The 32-bit tag value is not used by the OSPF protocol itself. You can use the tag value to communicate information between ASBRs.
- The advertised routes as ASE type1 by including the **type1** option. By default, the type is ASE type2. The difference between type1 and type2 is how the cost is calculated. For a type2 ASE, only the external cost (metric) is used when comparing multiple paths to the same destination. For type 1 ASE, the combination of the external cost and the cost to reach the ASBR is used.

For example:

```
(config)# ospf default metric 10 type1
```

To stop advertising the default ASE routes originated through OSPF, enter:

```
(config)# no ospf default
```

Advertising Other Routes Through OSPF

Use the **ospf redistribute** command to advertise routes from other protocols, such as firewall, local, RIP, and static routes through OSPF. Redistribution of these routes makes them OSPF external routes.

To redistribute routes from other protocols, include one of the following options:

- **firewall** - Advertises firewall routes through OSPF
- **local** - Advertises local routes (interfaces *not* running OSPF)
- **rip** - Advertises RIP routes through OSPF
- **static** - Advertises static routes configured for the Ethernet interface ports. The **ospf redistribute static** command does not advertise static routes configured for the Ethernet management port.

To advertise a firewall route, enter:

```
(config)# ospf redistribute firewall
```

Optionally, you can define any of the following:

- The network cost for the route by including the **metric** option. Enter a number from 1 to 16,777,215. The default is 1.
- A 32-bit tag value to advertise each external route by including the **tag** option. The 32-bit tag value is not used by the OSPF protocol itself. You can use the tag value to communicate information between AS boundary routers.
- The advertised routes as ASE type1 by including the **type1** option. By default, the type is ASE type2. The difference between type1 and type2 is how the cost is calculated. For a type2 ASE, only the external cost (metric) is considered when comparing multiple paths to the same destination. For type1 ASE, the combination of the external cost and the cost to reach the ASBR is used.

For example:

```
(config)# ospf redistribute rip metric 3 type1
```

To stop advertising the RIP routes via OSPF, enter:

```
(config)# no ospf redistribute rip
```

Configuring OSPF on a CSS IP Interface

When you configure a CSS IP interface as an OSPF interface, you define its behavior and role within the OSPF routing domain. This section includes the following topics:

- [Configuring the CSS IP Interface as an OSPF Interface](#)
- [Assigning an OSPF Area to the Interface](#)
- [Enabling OSPF on the Interface](#)
- [Configuring the Interface Attributes](#)

Configuring the CSS IP Interface as an OSPF Interface

An OSPF interface is an IP interface that you configure to send and receive OSPF traffic. To configure the CSS IP interface as an OSPF interface, use the **ospf** command.



Note

You must enter the **ospf** command before the **ospf enable** command can take effect.

To configure the CSS IP interface as an OSPF interface:

1. Access the circuit configuration mode for the preconfigured circuit on which you want to create the IP interface. For example, if circuit VLAN6 already exists, enter:

```
(config)# circuit VLAN6  
(config-circuit[VLAN6])#
```



Note Refer to [Chapter 5, Configuring Interfaces and Circuits](#) for information on how to configure the CSS interfaces and circuits, and bridge interfaces to VLANs.

2. Create the IP interface to the circuit. To create an IP address of 3.1.2.2, enter:

```
(config-circuit[VLAN6])# ip address 3.1.2.2/24
Create ip interface <3.1.2.2>, [y/n]:y
```

3. Configure this circuit as an OSPF circuit. Enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf
```

Assigning an OSPF Area to the Interface

After you configure the IP interface as an OSPF interface, assign it to the area that you globally configured to the CSS. The default area is the backbone area with the ID of 0.0.0.0. If the area is other than the backbone, use the **ospf area** command to assign the interface to an OSPF area. For example, if the area is 0.0.0.1, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf area 0.0.0.1
```

To reset the interface to the default backbone area, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf area
```

Enabling OSPF on the Interface



Note If you need to configure the interface attributes as described in the [“Configuring the Interface Attributes”](#) section, do not enable OSPF on the IP interface until you finish configuring the attributes.

Use the **ospf enable** command to enable OSPF on the IP interface. By default, OSPF is disabled on an IP interface. For example:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf enable
```

To disable OSPF on the interface, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf enable
```

Configuring the Interface Attributes

The OSPF interface attributes are set to a series of default values. You can elect to use these values for the CSS IP interface or configure your own settings. This section includes the following topics:

- [Setting the Cost](#)
- [Setting the Dead Router Interval](#)
- [Setting the Hello Packet Interval](#)
- [Setting the Password](#)
- [Setting the Poll Interval](#)
- [Setting the Priority of the CSS](#)
- [Setting the Retransmission Interval](#)
- [Setting the Transit-Link Delay](#)

Setting the Cost

Use the **ospf cost** command to set the cost for sending a data packet on this interface. The cost for the interface is a number from 0 to 65535. The default value of the cost for a given type of circuit is 108/interface speed. For a Gigabit Ethernet interface, the value is 1. For a 10/100-Mbps Fast Ethernet interface, the value is 10.

For example, to set a cost of 25, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf cost 25
```

To reset the packet cost for the interface to the default value, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf cost
```

Setting the Dead Router Interval

The interface declares a neighbor router is dead if the interface does not receive hello packets from the router before the dead interval expires. Use the **ospf dead** command to set the dead router interval for an interface. The dead router interval is in seconds. This value must be a multiple of the hello interval, and the value must be the same for all routers attached to a common network. Enter a number from 1 to 2,147,483,647. The default is 40.

For example, to set the dead router interval to 100 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf dead 100
```

To reset the dead router interval to its default of 40 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf dead
```

Setting the Hello Packet Interval

Router interfaces periodically transmit hello packets to identify and maintain communications with their neighbors. When a router detects its own address in another router's hello packet, the two routers establish two-way communications as neighbors.

The hello interval is the length of time, in seconds, between hello packets that the interface sends to its neighbor routers. The hello interval must be the same value for all routers attached to a common network. Use the **ospf hello** command to set the hello interval for the IP interface. Enter an integer from 1 to 65535. The default is 10 seconds.

To set a hello interval of 25 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf hello 25
```

To reset the hello interval to the default value of 10 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf hello
```

Setting the Password

All OSPF protocol exchanges can be authenticated to ensure only known, trusted routers participate in routing updates. The OSPF password is used for authentication of all OSPF protocol exchanges.

Use the **ospf password** command to set the password for an interface. This password must be the same for all routers attached to a common network. Enter a quoted text string with a maximum of eight characters.

For example, to set the password of *quota*, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf password "quota"
```

To remove the OSPF password from the interface, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf password
```

Setting the Poll Interval

The poll interval is the length of time, in seconds, between the transmittal of hello packets by the CSS to an assumed inactive neighbor router in a non-broadcast, multi-access network. Use the **ospf poll** command to set the poll interval for the interface. The poll interval should be a value that is greater than the hello time interval. Enter a number from 1 to 2,147,483,647. The default is 120 seconds.



Note

The **ospf poll** command has no effect when you operate the CSS over a broadcast LAN (that is, an Ethernet network).

For example, to set the poll interval to 200 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf poll 200
```

To reset the poll interval to the default value of 120 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf poll
```

Setting the Priority of the CSS

To avoid the need for each router on a LAN to talk to every router on a network that has more than two attached routers, one router is elected as the designated router. Designated routers advertise network link states for attached network segments. An LSA lists all routers that are connected to a segment.

The priority determines which router is the designated router. The router with the highest priority becomes the designated router. In case of a tie, routers use their router ID as a tie breaker.

Use the **ospf priority** command to set the router priority for the interface. The priority of the interface is an integer from 0 to 255. The default is 1, which is the highest router priority. A value of 0 signifies that the CSS is not eligible to become the designated router on a particular network.

If a designated router exists on the network, it remains the designated router regardless of its router priority.

To make the interface ineligible to become a designated router, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf priority 0
```

To reset the router priority to the default value of 1, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf priority
```

Setting the Retransmission Interval

The retransmission interval is the number of seconds between link-state advertisement retransmissions for adjacencies belonging to an interface. OSPF creates adjacencies between neighboring routers for the purpose of exchanging routing information. The CSS also uses the interval when retransmitting database descriptions and link-state request packets.

Use the **ospf retransmit** command to set the retransmit interval for the interface. Enter a number from 1 to 3600 seconds (1 hour). The default is 5 seconds.

To set the retransmission interval to 10 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf retransmit 10
```

To reset the retransmit interval to the default value of 5 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf retransmit
```

Setting the Transit-Link Delay

Transit delay is the estimated number of seconds the CSS waits to transmit a link-state update packet over the OSPF interface. Use the **ospf transit-delay** command to set the transit delay for an interface. Enter a number from 0 to 3600 seconds (1 hour). The default is 1 second.

To set the transit delay to 3 seconds, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# ospf transit-delay 3
```

To reset the transit delay to the default value of 1 second, enter:

```
(config-circuit-ip[VLAN6-3.1.2.2])# no ospf transit-delay
```

Showing OSPF Information

Use the **show ospf** command to view OSPF information on the CSS. This command is available in all modes. This section includes the following topics:

- [Showing OSPF Area Information](#)
- [Showing Global Statistics](#)
- [Showing IP Interface Information](#)
- [Showing Link-State Databases](#)
- [Showing ASE Entries](#)
- [Showing the Configured Advertised ASE Routes](#)
- [Showing the Redistribution Policy](#)
- [Showing Summary Route Configuration Information](#)
- [Showing OSPF Neighbors](#)

Showing OSPF Area Information

To show information about OSPF areas, enter:

```
# show ospf areas
```

[Table 7-3](#) describes the fields in the **show ospf areas** command output.

Table 7-3 *Field Descriptions for the show ospf areas Command*

Field	Description
Area ID	The ID for the area
Type	The area type: Transit or Stub
SPF Runs	The number of times the area calculated the SPF
Area Border Routers	The number of ABRs, including the CSS
AS Boundary Routers	The number of ASBRs, including the CSS, if applicable

Table 7-3 Field Descriptions for the `show ospf areas` Command (continued)

Field	Description
LSAs	The number of link-state advertisements in the database
Summaries	The capability of summarized LSAs in the stub area, if applicable

Showing Global Statistics

To show OSPF global statistics, enter:

```
# show ospf global
```

[Table 7-4](#) describes the fields in the `show ospf global` command output.

Table 7-4 Field Descriptions for the `show ospf global` Command

Field	Description
Router ID	The router ID of the CSS.
Admin Status	The state of OSPF on the CSS: Enabled or Disabled.
Area Border Router	Indicates whether the CSS is an ABR. True indicates the CSS is an ABR; otherwise, the field displays False.
AS Boundary Router	Indicates whether the CSS is an ASBR. True indicates the CSS is an ASBR; otherwise, the field displays False.
External LSAs	The number of external LSAs currently contained in the database.
LSA Sent	The number of LSAs sent by the CSS.
LSA Received	The number of LSAs received by the CSS.

Showing IP Interface Information

To show OSPF interfaces, enter:

```
# show ospf interfaces
```

[Table 7-5](#) describes the fields in the **show ospf interfaces** command output.

Table 7-5 *Field Descriptions for show ospf interfaces Command*

Field	Description
IP Address	The IP address for the OSPF IP interface
Admin State	Administrative state of OSPF on the interface, as affected by the IP interface ospf enable command
Area	The area assigned to the interface
Type	The OSPF interface type; always broadcast

Table 7-5 Field Descriptions for show ospf interfaces Command (continued)

Field	Description
State	<p>The functional level of an interface. The state determines whether full adjacencies are allowed to form over the interface. The states include:</p> <ul style="list-style-type: none"> • Down - The initial interface state. In this state, the lower-level protocols indicate the interface is unusable. No protocol traffic is sent or received on the interface. • Waiting - The router is trying to determine the identity of the (backup) designated router for the network. To determine the router identify, the router monitors the hello packets it receives. The router is not allowed to elect a backup designated router nor a designated router until it transitions out of the Waiting state. • DR Other - The interface is on a network on which another router has been selected to be the designated router. In this state, the router itself has not been selected as the backup designated router. The router forms adjacencies to both the designated router and the backup designated router. • Backup - The router itself is the backup designated router on the attached network. The router is the designated router when the present designated router fails. The router establishes adjacencies to all other routers attached to the network. The backup designated router performs slightly different functions during the flooding procedure, as compared to the designated router. • DR - The router itself is the designated router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network LSA for the network node. The network LSA contains links to all routers, including the designated router itself, attached to the network.

Table 7-5 *Field Descriptions for show ospf interfaces Command (continued)*

Field	Description
Priority	The priority assigned to the interface advertised in the hello packets. When two routers attached to a network both attempt to become the designated router, the router with the highest priority takes precedence. A router whose priority is set to 0 is ineligible to become the designated router on the attached network.
DR	The IP interface address of the designated router selected for the attached network. The designated router is selected on broadcast networks by the hello protocol. Two pieces of identification are kept for the designated router: the Router ID and the IP interface address on the network. The designated router advertises the link state for the network. This network LSA is labeled with the designated router's IP address. The designated router is initialized to 0.0.0.0, which indicates the lack of a designated router.
BR	The backup designated router selected for the attached network. The backup designated router is also selected on all broadcast networks by the hello protocol. All routers on the attached network become adjacent to both the designated router and the backup designated router. The backup designated router becomes the designated router when the current designated router fails. The backup designated router is initialized to 0.0.0.0, indicating the lack of a backup designated router.
Hello	The length of time, in seconds, between the hello packets that the router sends on the interface. This interval is advertised in hello packets sent out on this interface.
Dead	The number of seconds before the router's neighbors declare that the router is down, and when they stop receiving the router's hello packets. This interval is advertised in hello packets sent out on this interface.

Table 7-5 Field Descriptions for `show ospf interfaces` Command (continued)

Field	Description
Transit Delay	The number of seconds to transmit a Link State Update packet over an interface. LSAs contained in the Link State Update packet have their age incremented by this amount before transmission. This value should take into account transmission and propagation delays; the value must be greater than zero.
Retransmit	The number of seconds between LSA retransmissions for adjacencies belonging to an interface. Also, the interval is used when retransmitting Database Description and Link State Request packets.
Cost	The cost of sending a data packet on the interface, expressed in the link-state metric. The cost of sending a packet is advertised as the link cost for the interface in the router LSA. The cost of an interface must be greater than zero.

Showing Link-State Databases

You can show the entire OSPF link-state database (LSDB) or its specific entry types with the `show ospf lsdb` command. The options for the `show ospf lsdb` command are as follows:

- **show ospf lsdb router** - Displays router LSAs that describe the states of the router interfaces
- **show ospf lsdb network** - Displays network LSAs that describe the set of routers attached to the network
- **show ospf lsdb external** - Displays AS-external LSAs that describe routes to destinations external to the AS
- **show ospf lsdb summary** - Displays summary LSAs that describe summarized routes to the network
- **show ospf lsdb asbr_summ** - Displays summary LSAs that describe routes to AS boundary routers

To show the entire database, enter:

```
# show ospf lsdb
```

Table 7-6 describes the fields in the **show ospf lsdb** command output.

Table 7-6 Field Descriptions for the show ospf lsdb Command

Field	Description
Area	The ID for the area.
Type	<p>The link-state type. The types are as follows:</p> <ul style="list-style-type: none"> • ASB-Summary for summary LSAs originated by ABRs. The LSAs describe routes to ASBRs. • ASE for AS-external LSAs that describe routes to destinations external to the autonomous system. • Network for the network LSAs that describe the set of routers attached to the network. • Router for router LSAs that describe the collected states of the router interfaces. • Summary-Net for summary LSAs originated by ABRs. The LSAs describe routes to networks.
Link State ID	<p>This field identifies the piece of the routing domain that is being described by the LSA. Depending on the link-state type, the Link State ID has following values:</p> <ul style="list-style-type: none"> • For the ASB-Summary type, the ID is the router ID of the ASBR. • For the ASE type, the ID is the destination network IP address. • For Network type, the ID is the IP interface address of the network designated router. • For Router type, the ID is the originating router's Router ID. • For Summary-Net type, the ID is the destination network IP address.

Table 7-6 Field Descriptions for the `show ospf lsdb` Command (continued)

Field	Description
ADV Router	<p>This field specifies the OSPF Router ID of the LSA originator, as follows:</p> <ul style="list-style-type: none"> • ASB-Summary LSAs, the originators are the ABRs • AS-external LSAs, the originators are ASBRs • Network LSAs, the originators are network-designated routers • Router LSAs, this field is identical to the Link State ID field • Summary LSAs, the originators are the ABRs
Age	<p>The age of the LSA, in seconds. The age is set to 0 when the LSA is originated.</p>
Sequence	<p>A signed 32-bit integer to detect old and duplicate LSAs. The space of sequence numbers is linearly ordered. The larger the sequence number (when compared as signed 32-bit integers), the more recent the LSA.</p> <p>The sequence number 0x80000000 is reserved and unused.</p>
Checksum	<p>The checksum of the complete contents of the LSA, excluding the age field. The age field is excluded to allow the LSA age to increment without updating the checksum.</p> <p>The checksum is used to detect data corruption of an LSA. This corruption can occur while an LSA is being flooded, or while an LSA is being held in a router's memory. The LSA checksum field cannot take on the value of zero; the occurrence of this value is a checksum failure.</p>

Showing ASE Entries

To show AS-external (ASE) entries in the LSDB, enter:

```
# show ospf ase
```

To find specific entries, pipe the output through the **grep** command. For example:
show ospf aselgrep 10.10.10.0

[Table 7-7](#) describes the fields in the **show ospf ase** command output.

Table 7-7 *Field Descriptions for the show ospf ase Command*

Field	Description
Link State ID	The network destination for the advertisement
Router ID	The advertising router
Age	The age, in seconds, of the ASE LSA
T	The ASE type of the route; 1 for ASE Type1 or 2 for ASE Type2
Tag	The tag for the route
Metric	The network cost for the route
FwdAddr	The external destination (forwarding address) for the packets

Showing the Configured Advertised ASE Routes

To show the configuration of ASE routes into OSPF, enter:

```
# show ospf advertise
```

To show the configuration of ASE routes into OSPF for a specific host, include the IP address or host and the subnet mask. Enter the address in dotted-decimal format (for example, 192.168.11.1) or mnemonic host-name format (for example, myname.mydomain.com). Enter the mask either:

- As a prefix length in CIDR bit-count notation (for example, /24). Do not enter a space to separate the IP address from the prefix length.
- In dotted-decimal notation (for example, 255.255.255.0).

For example:

```
# show ospf advertise 192.168.11.1/24
```

Table 7-8 describes the fields in the **show ospf advertise** command output.

Table 7-8 *Field Descriptions for the show ospf advertise Command*

Field	Description
Prefix	The IP address for the route. For the CSS, the prefix is predominately VIP addresses.
Prefix Length	The prefix length for the IP address.
Metric	The network cost for the route. The range is from 1 to 16777215. The default is 1.
Type	The ASE type for the route. By default, the ASE type is ASE type2, which is the external cost to reach the route. ASE type1 combines the external and internal costs.
Tag	The 32-bit tag value to advertise the route. The value is not used by OSPF.

Showing the Redistribution Policy

To show the configured redistribution policy into OSPF, enter:

```
# show ospf redistribute
```

Table 7-9 describes the fields in the **show ospf redistribute** command output.

Table 7-9 *Field Descriptions for the show ospf redistribute Command*

Static, RIP, Local, or Firewall Field	Description
Routes Redistribution	Indicates whether the redistribution of static, RIP, local or firewall routes is enabled or disabled. If route redistribution is enabled, the configured metric, type, and tag fields are displayed.
Route Metric (displayed when redistribution is enabled)	The external cost for the route. The cost can range from 1 to 16777215. The default is 1.
Route Type (displayed when redistribution is enabled)	The ASE type, either ASE Type1 or ASE Type2. By default, the type is aseType2. The difference between type1 and type2 is how the cost is calculated. For a type 2 ASE, only the external cost (metric) is used when comparing multiple paths to the same destination. For type1 ASE, the combination of the external cost and the cost to reach the ASBR is used.
Route Tag (displayed when redistribution is enabled)	The 32-bit tag value to advertise the external route. The route tag value is not used by the OSPF protocol itself. It is used to communicate information between AS boundary routers.

Showing Summary Route Configuration Information

To show the summary-route configuration information, enter:

```
# show ospf range
```

[Table 7-10](#) describes the fields in the **show ospf range** command output.

Table 7-10 Field Descriptions for the `show ospf range` Command

Field	Description
Area ID	The ID for the area.
Lsdb Type	The type of link-state database. For an ABR, the type is summaryLink.
Addr Range Mask Range	The address range for the summary route as specified by the IP address (Addr Range) and mask (Mask Range) pair.
Effect	Displays whether the range is advertised or block.

Showing OSPF Neighbors

To show the OSPF neighbors, enter:

```
# show ospf neighbors
```

Table 7-11 describes the fields in the `show ospf neighbors` command output.

Table 7-11 Field Descriptions for `show ospf neighbors` Command

Field	Description
Address	The IP address of the neighboring router's interface to the attached network. This address is used as the destination IP address when protocol packets are sent as unicasts along this adjacency. The IP address is also used in router LSAs as the Link ID for the attached network if the neighboring router is selected to be the designated router. The CSS learns the neighbor IP address when it receives hello packets from the neighbor.
Neighbor ID	The OSPF Router ID of the neighboring router. The CSS learns the Neighbor ID when it receives hello packets from the neighbor.
Prio	The router priority of the neighboring router. Contained in the neighbor's hello packets, this value is used by OSPF to select the designated router for the attached network.

Table 7-11 Field Descriptions for *show ospf neighbors* Command (continued)

Field	Description
State/Dr	<p>The state of a conversation being held with a neighboring router. The following states are listed in order of their progression.</p> <ul style="list-style-type: none"> • Down - The initial state of a neighbor conversation. The Down state indicates that the CSS has received no recent information from the neighbor. • Init - In this state, the CSS has seen a hello packet from the neighbor. However, the CSS has not established bidirectional communication with the neighbor (the router itself did not appear in the neighbor's hello packet). All neighbors in this state (or higher) are listed in the hello packets sent from the associated interface. • 2-Way - In this state, communication between the two routers is bidirectional. The designated router is selected from the set of neighbors in state 2-Way (or greater). • ExStart - This is the first step to create an adjacency between the two neighboring routers. The goal is to decide which router is the master, and to determine the initial Database Description (DD) sequence number. Neighbor conversations in this state (or greater) are called adjacencies.

Table 7-11 Field Descriptions for show ospf neighbors Command (continued)

Field	Description
State/Dr (cont.)	<ul style="list-style-type: none"> • Exchange - In this state, the CSS sends DD packets to the neighbor to describe its entire link-state database. Each DD packet has a DD sequence number and is explicitly acknowledged. Only one DD packet is allowed to be outstanding at any one time. In this state, the CSS may also send Link State Request packets, requesting the neighbor's more recent LSAs. All adjacencies in Exchange state (or greater) are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets. • Loading - In this state, the CSS sends Link State Request packets to the neighbor, requesting the more recent LSAs that have been discovered (but not yet received) in the Exchange state. • Full - In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router LSAs and network LSAs.
Type	Always dynamic.
Rxmt_Q	The number of LSAs to retransmit to the neighbors.

OSPF Configuration in a Startup-Config File

The following example shows an OSPF configuration in a startup-config file.

```
!***** GLOBAL *****
  ospf router-id 121.23.21.1
  ospf enable
  ospf area 1.1.1.1
  ospf as-boundary
  ospf advertise 192.168.4.15 255.255.255.0
  ospf redistribute static
!***** INTERFACE *****
interface ethernet-10
  bridge vlan 6
!***** CIRCUIT *****
circuit VLAN6
ip address 192.168.2.2 255.255.255.0
  ospf
  ospf area 1.1.1.1
  ospf priority 0
  ospf enable
```

Where to Go Next

[Chapter 8, Using the CSS Logging Features](#) describes how to enable logging, set up the log buffer, and determine where to send the activity information. This chapter also provides information on interpreting sys.log messages and a description of frequently queried messages.

■ Where to Go Next