



Configuring OSPF

This chapter provides configuration and viewing information for Open Shortest Path First (OSPF). 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 sections:

- [OSPF Overview](#)
- [CSS OSPF Quick Configuration](#)
- [Configuring OSPF on the CSS](#)
- [Configuring OSPF on a CSS IP Interface](#)
- [Displaying 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 its neighbors.

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 its IP address and subnet mask, and its relationship to its 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 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

The following sections provide brief descriptions about:

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

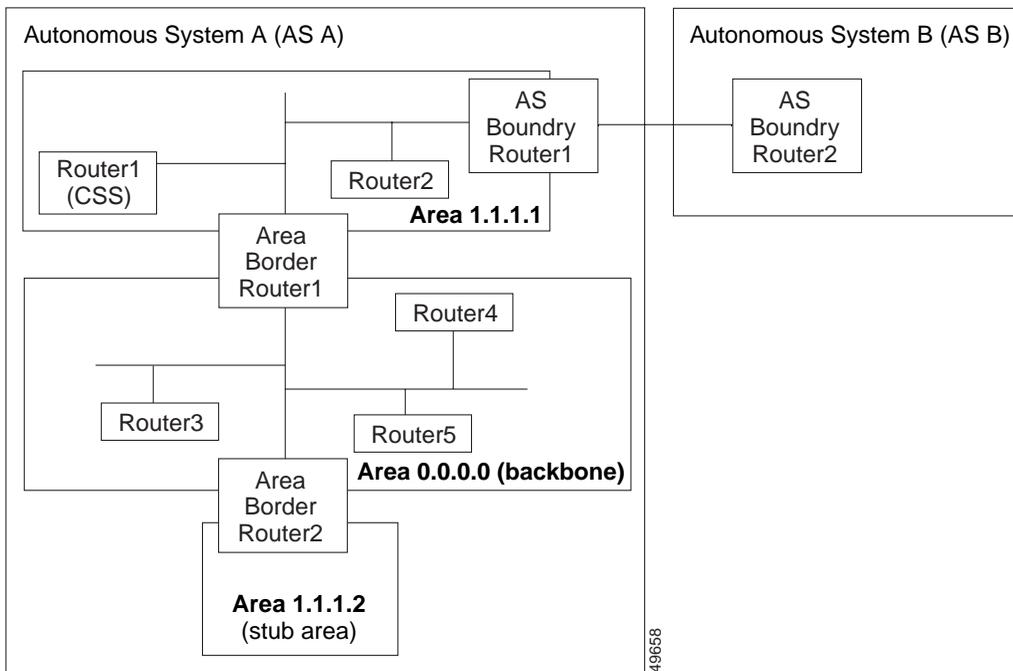
OSPF Routing Hierarchy

The OSPF routing hierarchy includes the following entities:

- Autonomous systems
- Areas, including the backbone and stub areas
- Area Border Router (ABR)
- AS Boundary Router (ASBR)

Figure 4-1 provides an example of an OSPF network topology.

Figure 4-1 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 4-1](#) 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 4-1](#), AS A consists of three areas: area 0.0.0.0, 1.1.1.1, and 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; it is not visible to routers outside that area. When OSPF routing is within an area, it is called intra-area routing. This routing limits the amount of link-state information flooding onto the network, thereby reducing routing traffic. It 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, it is called inter-area routing.

The backbone itself has all the properties of an area. It consists of area border routers, and routers and networks only on the backbone. As shown in [Figure 4-1](#), area 0.0.0.0 is an OSPF backbone area. Note that any OSPF backbone area has a reserved ID of 0.0.0.0.

Area Border Routers

Area border routers (ABRs) are routers with 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. Area border routers 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 4-1](#), 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 them 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 area border router 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 link-state advertisement (LSA) database space that would otherwise be used to store external LSAs flooded into the area. In [Figure 4-1](#), area 1.1.1.2 is a stub area that is reached only through area border router 2.

Autonomous System Boundary Routers

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 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.

Link-State Databases

OSPF routers advertise routes using LSAs. The link-state database stores the link-state advertisements 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 [“Setting the Priority of the CSS”](#) later in this chapter), 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 provides quick configuration information including:

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

Global OSPF Quick Configuration

To perform the global OSPF configuration for the CSS, see the steps in [Table 4-1](#). In the most basic global configuration where the CSS is just 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 area border router by configuring route summarization.
- As an AS boundary router, 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.



Note

After performing the global OSPF configuration, you must configure an OSPF IP interface (as described in [“OSPF IP Interface Quick Configuration”](#)) before the CSS can participate in OSPF routing.

For more detailed information on configuring these global OSPF parameters, see [“Configuring OSPF on the CSS”](#) later in this chapter.

Table 4-1 CSS Global OSPF Quick Configuration

Task and Command Example

1. In global configuration mode, configure the area router ID for the CSS. 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
```

Note that 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 it 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 that 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 next section.

OSPF IP Interface Quick Configuration

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

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

This configuration example assumes that you will want to 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, link-state update packet, and determining that a neighbor router is dead
- Authentication password
- CSS router priority
- Interface cost

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

Table 4-2 *OSPF on a CSS Interface Quick Configuration*

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 For more information on how to configure the CSS interfaces and circuits, and bridge interfaces to VLANs, refer to Chapter 2, [Configuring Interfaces and Circuits](#).

Table 4-2 OSPF on a CSS Interface Quick Configuration (continued)

Task and Command Example
<p>3. Create the IP interface to the circuit. To create an IP address of 3.1.2.2, enter:</p>
<pre>(config-circuit[VLAN6])# ip address 3.1.2.2/24 Create ip interface <3.1.2.2>, [y/n]: y</pre>
<p>4. Configure the IP interface as an OSPF interface. Enter:</p>
<pre>(config-circuit-ip[VLAN6-3.1.2.2])# ospf</pre>
<p>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.</p>
<pre>(config-circuit-ip[VLAN6-3.1.2.2]) ospf area 1.1.1.1</pre>
<p>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:</p>
<pre>(config-circuit-ip[VLAN6-3.1.2.2])# ospf priority 0</pre>
<p>For more information on designated routers, see “Setting the Priority of the CSS” later in this chapter.</p>
<p>7. Enable OSPF on the interface. Enter:</p>
<pre>(config-circuit-ip[VLAN6-3.1.2.2])# ospf enable</pre>

Quick Configuration Verification

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

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

```
# show ospf global
```

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

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

```
# show ospf redistribute
```

- To display 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

The global configuration of OSPF on the CSS includes:

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

Configuring the OSPF Router ID

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

To configure the OSPF router ID for the CSS, use the **ospf router-id** command. 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, you must disable OSPF by using the **(config) no ospf enable** command to 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. To enable OSPF, use the **ospf enable** command. 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 an area.

To assign an area, use the **ospf area** command. 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 either:

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

To set a metric for the default route advertised into the stub area, include the **default-metric** option. By default, the metric equals the least metric among the interfaces to other areas. You can assign an integer from 1 to 16,777,215.

To assign a metric of 200, enter:

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

To propagate summary LSAs into 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 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. To change the number of routes, use the **ospf equal-cost** command. 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 area border router (ABR), you can configure it to advertise a single summary route or network ranges that cover all the individual networks within its area that fall into 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.

To specify the IP address range to summarize routes at the ABR, use the **ospf range** command. 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.

**Note**

OSPF must be disabled before you can enter this 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 bitcount 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 192.172.0.0 and 192.172.255.255, enter:

```
(config)# ospf range 0.1.0.1 192.172.0.0 255.255.0.0
```

To remove the range, enter:

```
(config)# no ospf range 0.1.0.1 192.172.0.0 255.255.0.0
```

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 0.1.0.1 192.172.0.0 255.255.0.0 block
```

Configuring the CSS as an Autonomous System Boundary Router

If you want the CSS to be an AS boundary router that exchanges routing information with routers belonging to other Autonomous Systems, use the **ospf as-boundary** command.

**Note**

OSPF must be disabled before you can enter this command.

For example:

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

To unassign 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 AS boundary router can perform external route summarization to consolidate multiple routes into a single advertisement. For a CSS, this is useful when you want to advertise VIP addresses for content as OSPF AS external (ASE) through all OSPF interfaces. To advertise a route as OSPF ASE through all OSPF interfaces, use the **ospf advertise** command. To stop the advertisement of the route, enter the **no 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 will advertise the VIP address when the virtual router state is Master, but will stop advertising this VIP address when the virtual router state is either Down or Backup.
- If the advertised IP address is a service IP in a service record configuration, OSPF will advertise the IP address when the service state is Alive, Dying, or Suspend, but will stop 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 will treat it the same as the service IP described above. The VIP redundancy state will 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, which does not match any service IP and redundant VIP, OSPF will advertise the IP address all the time. If this address is used as a service IP at a later point, OSPF will determine whether to advertise the IP address based on the service state. When the service state for the service IP is Down, OSPF will not advertise the service IP. In addition, OSPF will not revert the service IP back to the original IP address and resume advertising it when the IP address is not overlapped with a service IP.

The following running-config 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

```

**Note**

Before you enter the **ospf advertise** command, you must configure the CSS as an AS boundary router. For more information, see [“Configuring the CSS as an Autonomous System Boundary Router”](#) earlier in this chapter.

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 bitcount notation format (for example, /24).

In this example, to advertise VIP addresses from 193.23.44.0 to 193.21.44.255, define the range by entering the IP address and subnet mask of 193.23.44.0 255.255.255.0:

```
(config)# ospf advertise 193.23.44.0 255.255.255.0
```

You can optionally define any of the following:

- Define the network cost for the route by including the **metric** option. Enter a number from 1 to 16,777,215. The default is 1.
- Define a 32-bit tag value to advertise each external route by including the **tag** option. This is not used by the OSPF protocol itself. You can use it to communicate information between AS boundary routers.
- Advertise the 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 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 193.23.44.255.255.255.0
```

Advertising a Default ASE Route

Routers use default routes when no more specific routes exist to a particular AS external destination. By default, an AS boundary router does not generate a default route into the OSPF routing domain. To force the CSS to generate a default ASE route and advertise it through OSPF, use the **ospf default** command.



Note

Before you enter this command, you must configure the CSS as an AS boundary router. For more information, see [“Configuring the CSS as an Autonomous System Boundary Router”](#) earlier in this chapter.

For example:

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

You can optionally:

- Define 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.
- Define a 32-bit tag value to advertise each external route by including the **tag** option. This is not used by the OSPF protocol itself. You can use it to communicate information between AS boundary routers.
- Advertise the 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 default metric 10 type1
```

To shut off the advertising of default ASE routes originated through OSPF, enter:

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

Advertising Other Routes through OSPF

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

To advertise a route protocol, include one of the following options:

- **firewall** - Firewall route
- **local** - Local route
- **rip** - RIP route
- **static** - Static route

To advertise a firewall route, enter:

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

You can optionally:

- Define the network cost for the route by including the **metric** option. Enter a number from 1 to 16,777,215. The default is 1.
- Define a 32-bit tag value to advertise each external route by including the **tag** option. This is not used by the OSPF protocol itself. You can use it to communicate information between AS boundary routers.
- Advertise the 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 shut off the advertising of 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 provides information for:

- [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 this command before the **(config-circuit-ip) ospf enable** command can take effect.

For example:

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

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

**Note**

For more information on how to configure the CSS interfaces and circuits, and bridge interfaces to VLANs, refer to Chapter 2, [Configuring Interfaces and Circuits](#).

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. To 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 [“Configuring the Interface Attributes”](#), do not enable OSPF until you are finished.

To enable OSPF on the IP interface, use the **ospf enable** command. 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 default values. You can elect to use these values for the CSS IP interface or configure your own settings. The following sections provides information for:

- [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

To set the cost for sending a data packet on this interface, use the **ospf cost** command. The cost for the interface is a number from 0 to 65535. The default value of 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.

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 its default value, enter:

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

Setting the Dead Router Interval

The interface declares that a neighbor router is dead if the interface does not receive hello packets from the router before the dead interval expires. To set the dead router interval for an interface, use the **ospf dead** command. The dead router interval is in seconds. This value must be a multiple of the hello interval, and it 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.

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. To set the hello interval for the IP interface, use the **ospf hello** command. 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 its 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 so that only known, trusted routers can participate in routing updates. The OSPF password is used for authentication of all OSPF protocol exchanges.

To set the password for an interface, use the **ospf password** command. 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.

To set the password of quotation, 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 hello packets that the CSS sends to an assumed inactive neighbor router in a non-broadcast, multi-access network. To set the poll interval for the interface, use the **ospf poll** command. This interval should be much larger 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).

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 its 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. A 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.

To set the router priority for the interface, use the **ospf priority** command. 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.



Note

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 its 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.

To set the retransmit interval for the interface, use the **ospf retransmit** command. 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 its 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 to transmit a link-state update packet over the OSPF interface. To set the transit delay for an interface, use the **ospf transit-delay** command. 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 its default value of 1 second, enter:

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

Displaying OSPF Information

You can view OSPF information on the CSS by using the **show ospf** command. This command is available in all modes. The following sections provides information on:

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

Displaying OSPF Area Information

To display information about OSPF areas, enter:

```
# show ospf areas
```

[Table 4-3](#) describes the fields in the **show ospf areas** output.

Table 4-3 *Field Descriptions for the show ospf areas Command*

Field	Description
Area ID	The ID for the area
Type	The area type. The type can be either transit or stub
SPF Runs	The number of times that the area has calculated the SPF
Area Border Routers	The number of area border routers including the CSS
AS Boundary Routers	The number of AS boundary routers including the CSS, if applicable

Table 4-3 *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 into the stub area, if applicable

Displaying Global Statistics

To display OSPF global statistics, enter:

```
# show ospf global
```

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

Table 4-4 *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	Whether the CSS is an ABR. True indicates that the CSS is an ABR, otherwise, the field displays False
AS Boundary Router	Whether the CSS is an ASBR. True indicates that 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

Displaying IP Interface Information

To display OSPF interfaces, enter:

```
# show ospf interfaces
```

[Table 4-5](#) describes the fields in the **show ospf interfaces** output.

Table 4-5 *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 is always broadcast

Table 4-5 *Field Descriptions for show ospf interfaces Command (continued)*

Field	Description
State	<p>The functional level of an interface. 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 have indicated that the interface is unusable. No protocol traffic will be sent or received on the interface. • Waiting - The router is trying to determine the identity of the (backup) designated router for the network. To do this, 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. It will be the designated router if and 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 will contain links to all routers, including the designated router itself, attached to the network.

Table 4-5 *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 designated router, the one with the highest priority takes precedence. A router whose priority is set to 0 is ineligible to become 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: its Router ID and its 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 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 will declare it down, when they stop hearing the router's hello packets. This interval is advertised in hello packets sent out on this interface.

Table 4-5 *Field Descriptions for show ospf interfaces Command (continued)*

Field	Description
Transit Delay	The number of seconds it takes to transmit a Link State Update packet over an interface. LSAs contained in the Link State Update packet will have their age incremented by this amount before transmission. This value should take into account transmission and propagation delays; it 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. This is advertised as the link cost for the interface in the router LSA. The cost of an interface must be greater than zero.

Displaying Link-State Databases

You can display the entire OSPF link-state database (LSDB) or its specific entry types with the **show ospf lsdb** command. To display entries for:

- Router LSAs that describe the states of the router interfaces, use the **show ospf lsdb router** command.
- Network LSAs that describe the set of routers attached to the network, use the **show ospf lsdb network** command.
- AS-external LSAs that describe routes to destinations external to the AS, use the **show ospf lsdb external** command.
- Summary LSAs that describe summarized routes to the network, use the **show ospf lsdb summary** command.
- Summary LSAs that describe routes to AS boundary routers, use the **show ospf lsdb asbr_summ** command.

To display the entire database, enter:

```
# show ospf lsdb
```

[Table 4-6](#) describes the fields in the **show ospf lsdb** output.

Table 4-6 *Field Descriptions for the show ospf lsdb Command*

Field	Description
Area	The ID for the area.
Type	The link-state type. The types are: <ul style="list-style-type: none"> • ASB-Summary for summary LSAs originated by area border routers. The LSAs describe routes to AS boundary routers. • 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 area border routers. The LSAs describe routes to networks.

Table 4-6 Field Descriptions for the `show ospf lsdb` Command (continued)

Field	Description
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 AS boundary router. • 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.
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 area border routers • AS-external LSAs, the originators are AS boundary routers • 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 area border routers
Age	<p>The age of the LSA in seconds. The age is set to 0 when the LSA is originated.</p>

Table 4-6 Field Descriptions for the `show ospf lsdb` Command (continued)

Field	Description
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>This field is the checksum of the complete contents of the LSA, excluding the age field. The age field is excluded so that an LSA age can be incremented 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 it is being held in a router's memory. The LS checksum field cannot take on the value of zero; the occurrence of this value is a checksum failure.</p>

Displaying ASE Entries

To display AS-external entries in the LSDB, enter:

```
# show ospf ase
```

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

Table 4-7 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, either 1 for aseType1 or 2 for aseType2
Tag	The tag for the route

Table 4-7 Field Descriptions for the `show ospf ase` Command (continued)

Field	Description
Metric	The network cost for the route
FwdAddr	The external destination (forwarding address) for the packets

**Note**

To find specific entries, pipe the output to `grep`. For example: **show ospf ase|grep 10.10.10.0**

Displaying the Configured Advertised ASE Routes

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

```
# show ospf advertise
```

To display the configuration of ASE routes into OSPF for a specific host, include the IP address or host and its 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 bitcount 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 4-8 describes the fields in the **show ospf advertise** output.

Table 4-8 *Field Descriptions for the show ospf advertise Command*

Field	Description
Prefix	The IP address for the route. For the CSS, the prefix will be 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 type 2 that is the external cost to reach the route. ASE type 1 combines the external and internal costs.
Tag	The 32-bit tag value to advertise the route. The value is not used by OSPF.

Displaying Redistribution Policy

To display the configured redistribution policy into OSPF, enter:

```
# show ospf redistribute
```

Table 4-9 describes the fields in the **show ospf redistribute** output.

Table 4-9 *Field Descriptions for the show ospf redistribute Command*

Static, RIP, Local, or Firewall Field	Description
Routes Redistribution	Displays 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.

Table 4-9 *Field Descriptions for the show ospf redistribute Command (continued)*

Static, RIP, Local, or Firewall Field	Description
Route Type (displayed when redistribution is enabled)	The ASE type, either aseType1 or aseType2. By default, the type is aseType2. 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.
Route Tag (displayed when redistribution is enabled)	The 32-bit tag value to advertise the external route. This is not used by the OSPF protocol itself. It is used to communicate information between AS boundary routers.

Displaying Summary Route Configuration Information

To display the summary-route configuration information, enter:

```
# show ospf range
```

[Table 4-10](#) describes the fields in the **show ospf range** output.

Table 4-10 *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.

Displaying OSPF Neighbors

To display the OSPF neighbors, enter:

```
# show ospf neighbors
```

Table 4-11 describes the fields in the **show ospf neighbors** output.

Table 4-11 *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. It is also used in router LSAs as the Link ID for the attached network if the neighboring router is selected to be 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 when selecting the designated router for the attached network.

Table 4-11 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 - This is the initial state of a neighbor conversation. It 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 in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description (DD) sequence number. Neighbor conversations in this state or greater are called adjacencies.

Table 4-11 Field Descriptions for show ospf neighbors Command (continued)

Field	Description
State (cont.)	<ul style="list-style-type: none"> • Exchange - In this state, the CSS is sending Database Description (DD) packets to the neighbor, describing its entire link-state database. Each Database Description packet has a DD sequence number and is explicitly acknowledged. Only one Database Description packet is allowed 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 3.1.2.2 255.255.255.0
ospf
ospf area 1.1.1.1
ospf priority 0
ospf enable
```

Where to Go Next

Chapter 5, [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.