

7 Configuring OSPF

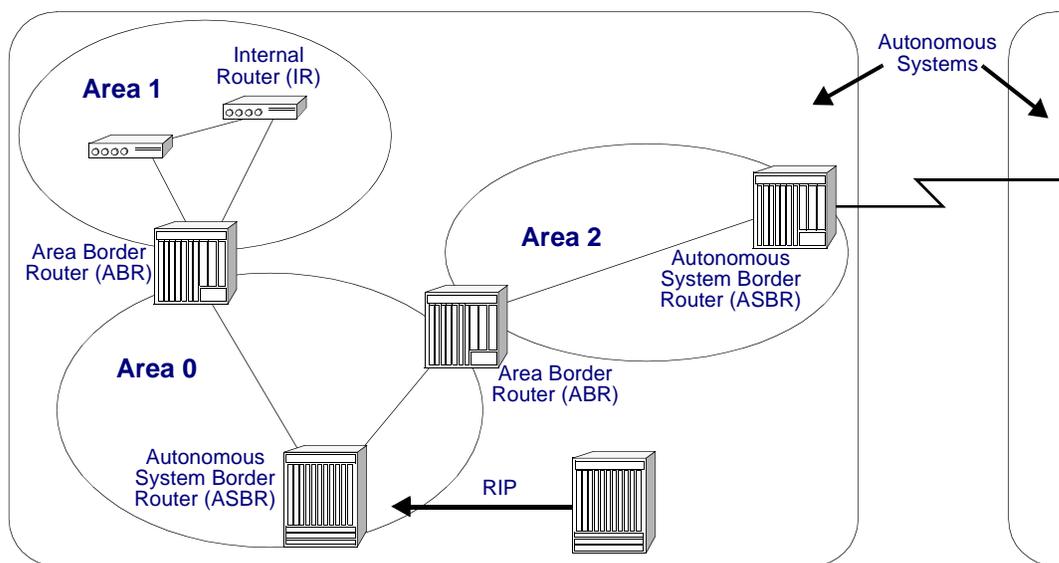
Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) used with larger autonomous system (AS) networks. An AS is considered to be a group of routers and hosts that use a common IGP and are administrated by one entity. OSPF determines the shortest path to a destination through the AS by using an algorithm that places the router at the root of a tree, and then calculating the lowest cost to reach the desired destination. Thus each router has a different view of the network's topology, even if they have been supplied with the same link-state database.

OSPF is a link-state protocol, and as such, it uses link state advertisements (LSAs) to inform other routers in the network of changes it has detected. Unlike RIP, which advertises its entire routing table every 30 seconds, OSPF immediately multicasts a LSA when it detects a change in the network. However, it sends only those portions of the routing table that have changed.

To further limit LSA broadcasts, the AS can be divided into separate areas. Areas should, in general, be defined along natural network divisions such as subnet assignments. When areas are configured, link state advertisements and shortest path configuration on a router are limited to changes within an area.

OSPF routers can be identified as one of several different types of routers.

- Routers whose interfaces all belong to only one area are known as internal routers (IRs). Since Internal Routers belong to the same area, they all have exactly the same link-state database as the other routers in their area.
- If a router belongs to more than one area, it is an area border router (ABR). ABRs have the responsibility of distributing route advertisements among its interfaces' respective areas.
- A router that is running multiple protocols or acting as a gateway to other exterior routers is called an autonomous system border router (ASBR). Using a process known as redistribution, ASBRs are able to translate and inject routes learned through other protocols into OSPF areas. In addition, ASBRs can be configured to export or redistribute OSPF routes into other routing protocols, such as RIP.



- In addition, on each multi-access segment (i.e., segments that allow multiple, simultaneous communication and connections), OSPF elects a designated router (DR) and a backup designated router (BDR). The DR provides a central point of contact, where all routers on a segment can exchange information, instead of exchanging updates with every router. The

DR then relays route information or LSAs to other routers, thus drastically reducing network traffic. The BDR also receives the same route update information, allowing it to provide redundancy should the DR fail. The election of the DR and BDR is determined by router priority. The router with the highest priority will become the DR. In the event that two routers have the same priority, the router with the highest router ID (the router ID allows a router to be identified in the AS) will be elected as the DR. A router with a priority of 0 is ineligible to become the designated router.

OSPF Commands

The major OSPF global and interface commands in the OmniCore CLI are listed in the following tables. Other commands are available for fine-tuning your OSPF configuration. To see a complete list of these commands, or for more information regarding these commands, see the *OmniCore CLI Reference Manual*.

OSPF Global Commands

Command	Default	Description
ip ospf area	no default	Creates an OSPF area.
ip ospf area default-metric	no default	Creates a TOS instance for an area.
ip ospf area default-metric cost	1	Configures cost advertised for TOS value.
ip ospf area default-metric type	ospf	Determines type of metric advertised for a TOS entry.
ip ospf area-type	normal	Enables the area type: normal, stub, or NSSA.
ip ospf area range	no default	Creates a summarized route.
ip ospf area range effect	adMatching	Specifies advertisement mode for summarized LSAs.
ip ospf area summary	enable	Enables Type-3 LSAs into a stub area.
ip ospf asbr	disable	Configures the OmniCore routing switch as an ASBR.
ip ospf ext-lsa limit	0	Defines the total number of allowed external LSAs for OSPF areas.
ip ospf host	1	Creates a host with specified TOS value.
ip ospf host metric	1	Defines a cost metric value for a specified host.
ip ospf log-status	disable	Specifies the OSPF message logging status.
ip ospf log-type	no default	Specifies the types of OSPF messages that will be logged.
ip ospf neighbor		Displays neighbor state.
ip ospf overflow-interval	0 seconds	Defines the overflow interval for attempting to leave the database overflow state.

OSPF Global Commands (Continued)

ip ospf router	IP address of first identified interface	Determines the router ID.
ip ospf route-tag	0	Specifies tag value for AS external routes.
ip ospf spf-timer	10 seconds (delay value) 5 seconds (hold value)	Configures timers for SPF calculation.
ip ospf status	disable	Enables OSPF.
ip ospf virtual-link	no default	Create a virtual link.
ip ospf virtual-link auth-key	null string	Configures text string used for authentication.
ip ospf virtual-link auth-type	no default	Determines the authentication type used.
ip ospf virtual-link dead-interval	40 seconds	Determines time until a router's neighbors will assume it is down.
ip ospf virtual-link hello-interval	10 seconds	Determines time between the transmission of hello packets.
ip ospf virtual-link retrans-interval	5 seconds	Determines time between link state advertisement retransmissions.
ip ospf virtual-link transit-delay	1 second	Configures time needed to transmit a link state update packet.

OSPF Interface Commands

Command	Default	Description
vlan ip ospf area	area 0 (0.0.0.0)	Sets the area to which an interface belongs.
vlan ip ospf auth-key	null string	Configures text string used for authentication.
vlan ip ospf auth-type	none	Determines the authentication type used.
vlan ip ospf dead-interval	40 seconds	Determines time until a router's neighbors will assume it is down.
vlan ip ospf hello-interval	10 seconds	Determines time between the transmission of hello packets.
vlan ip ospf metric	1	Creates a TOS metric entry.
vlan ip ospf metric cost	1	Associates a cost with a TOS entry.
vlan ip ospf poll-interval	120 seconds	Determines time allowed between hello packets.
vlan ip ospf priority	1	Determines router priority.
vlan ip ospf retrans-interval	5 seconds	Determines time between link state advertisement retransmissions.

OSPF Interface Commands (Continued)

vlan ip ospf status	disable	Enables OSPF on an interface basis.
vlan ip ospf transit-delay	1 second	Configures time needed to transmit a link state update packet.

Configuring OSPF

Configuring OSPF on the OmniCore routing switch requires several steps. However, depending on your requirements, you may not need to perform all of the steps listed below. By default, OSPF is disabled on the OmniCore routing switch. The total number of OSPF neighbors cannot exceed 64. Configuring OSPF consists of these tasks:

- Set the router ID. (Optional)
- Create an OSPF area(s).
- Set area parameters. (Optional)
- Configure OSPF interface parameters.
- Configure virtual links. (Optional)
- Create a redistribution policy for the desired protocol. (Optional)
- Create redistribution filters for the desired protocols. (Optional)
- Globally enable OSPF.

Follow these steps to configure OSPF:

1. Set the router ID. (Optional step)

The router ID, which identifies the OmniCore routing switch in an autonomous system (AS), must be in dotted decimal format. The OmniCore routing switch is automatically assigned a router ID if you do not manually set one. A router ID of 10.100.100.101 is used for this example.

```
OmniCore> ip ospf
OmniCore/ip/ospf> router 10.100.100.101
OmniCore/ip/ospf> router show
Router ID                               :10.100.100.101
```

2. Creates an OSPF area. Areas should be created along logical network divisions such as subnet assignments. Area IDs can be in decimal format or in 32-bit dotted decimal format. The following commands create areas 0.0.0.20 and 0.0.0.100.

```
OmniCore/ip/ospf> area 0.0.0.20 create
OmniCore/ip/ospf> area 0.0.0.100 create
OmniCore/ip/ospf> area show
Area Id      Summary Status      Area-type      Current State
-----
0.0.0.20    disable              Normal         enable
0.0.0.100   disable              Normal         enable
```

3. (Optional step) Insert a default route into the area. This can be done on any area type. For example:

```
OmniCore/ip/ospf> area default metric 0 create
```

4. Configure a route summary. This enables the area border router to advertise a consolidated route summary that represents an entire address range. An IP address of 172.17.0.0 and a mask of 255.255.252.0 are used in the following example. This will summarize the following networks; 172.17.0.0/24, 172.17.1.0/24, 172.17.2.0/24, 172.17.3.0/24

```
OmniCore/ip/ospf> area 0.0.0.100 range summary 172.17.0.0 255.255.252.0
create

OmniCore/ip/ospf> area 0.0.0.100 range summary 172.17.0.0 255.255.252.0 show
LSDB Type                :summary
Net Address for Range    :172.17.0.0
Subnet Mask              :255.255.252.0
Advertise Effect        :admatching
Current State            :enable
```

5. Configure an area's area-type as a normal OSPF, stub, or not-so-stubby (NSSA) area.
- For normal OSPF, use the following command.
 - For a stub area-type see step 6.
 - For a NSSA area-type see step 7.

```
OmniCore/ip/ospf> area 0.0.0.100 area-type normal

OmniCore/ip/ospf> area 0.0.0.100 area-type show
Area Type                :normal
```

6. For configuring the area-type as a stub use the following command. This is necessary if you do not wish to import AS external link-state advertisements into a specific area. A stub area is an AS which has only one connection to the OSPF domain. One default route summarizes all external routes. Configuring a stub area consists of specifying the area as a stub, and enabling the stub area as shown in the following commands.

```
OmniCore/ip/ospf> area 0.0.0.100 area-type stub enable

OmniCore/ip/ospf> area 0.0.0.100 area-type show
Area Type                :stub
```

- a. Enable the summary link-state advertisements:

```
OmniCore/ip/ospf> area 0.0.0.1000 summary enable
```

This enables the ABR to propagate Type-3 LSA (network summaries) into a stub area and is generally useful only when there is more than one ABR connection to a stub area. This will allow a router within the stub area to choose which ABR would provide the shortest path to any given network. (After enabling, go to step 8.)

7. For configuring the area-type as a as NSSA, use the following command. See [Not-So-Stubby-Areas](#) on page 7-7 for and explanation of NSSA:

```
OmniCore> ip ospf area 0.0.0.100 area-type nssa

OmniCore> ip ospf area 0.0.0.100 area-type show
Area Type                :nssa
```

◆ Note ◆

All routers in an NSSA must have their OSPF area defined as an NSSA. To configure otherwise will ensure that the router will be unsuccessful in establishing an adjacent in the OSPF domain. Also, routers to be able to originate Type-7 LSA or ABRs that need to translate the Type-7 LSA must have their ASBR status enabled

- a. Configure a NSSA route summary. This enables the area border router to advertise a consolidated route summary that represents an entire NSSA range. An IP address of 172.17.0.0 and a mask of 255.255.252.0 are used in the following example. This will summarize the following networks; 172.17.0.0/24, 172.17.0.1/24,

172.17.0.2/24, 172.17.0.3/24

```
OmniCore/ip/ospf> area 0.0.0.100 range NSSA 172.17.0.0 255.255.252.0 create
OmniCore/ip/ospf> area 0.0.0.100 range NSSA 172.17.0.0 255.255.252.0 show
LSDB Type                :NSSA
Net Address for Range    :172.17.0.0
Subnet Mask              :255.255.252.0
Advertise Effect        :admatching
Current State            :enable
```

8. Configure OSPF interface parameters.

- a. Associate a specific IP interface with an existing OSPF area.

By default, all OSPF interfaces on the OmniCore routing switch belong to Area 0.0.0.0. For more information regarding VLAN configuration, see [Chapter 3, "Configuring VLANs and Priority"](#). The following example associates an IP address of 10.0.23.111 with area 0.0.0.100. Note that the area to which the IP interface is assigned must have been previously created.

```
OmniCore/ip/ospf> home
OmniCore> vlan 2
OmniCore/vlan=2> ip 10.0.23.111 ospf area 0.0.0.100
```

- b. Enable OSPF on the interface (by default this is enabled). Please note that global OSPF must also be enabled.

```
OmniCore/vlan=2> ip 10.0.23.111 ospf status enable
```

- c. (Optional step) Configure individual interface commands. All interfaces, when assigned to an area, are created with default values for each parameter. While the default values should be sufficient for most networks, you may need to modify some parameters, such as authentication type and key, router priority, and metric path cost. See the *ip ospf* and *vlan ip ospf* commands in the *OmniCore CLI Reference Manual* for more information.

9. (Optional step) Configure a virtual link. A virtual link must be established when an ABR is not directly connected to the backbone area (area 0). When configuring a virtual link, the designated transition area must be defined on both virtual link neighbors in the specified transition area.

- a. Create the virtual link. This example establishes a virtual link through area 0.0.0.10 for a router with ID 10.0.52.11.

```
OmniCore/vlan=2> ..
OmniCore> ip ospf
OmniCore/ip/ospf> virtual-link 0.0.0.10 10.0.52.11 create
```

- b. Configure virtual link parameters. All virtual links are created with default values for each parameter. While the default values should be sufficient for most networks, you may need to modify some parameters, such as authentication type and key. See the *ip ospf virtual-link* commands in the *OmniCore CLI Reference Manual* for more information.

10. Globally enable OSPF:

```
OmniCore/ip/ospf> status enable
```

Further optional OSPF configuration options follow.

Configuring a Stub Area

A stub area is an AS which has only one connection to the OSPF domain. One default route summarizes all external routes. Configuring a stub area consists of specifying the area a stub, and enabling the stub area:

```
OmniCore> ip ospf area 0.0.0.10 area-type stub
OmniCore> ip ospf area 0.0.0.10 area-type show
Area Type                               :stub
```

Not-So-Stubby-Areas

NSSA, or not-so-stubby area, is an extension to the base OSPF specification and is defined in RFC 1587. An NSSA is similar to a stub areas in many ways: AS-external LSAs are not flooded into NSSAs and virtual links are not allowed in an NSSA. The primary difference is that selected external routing information can be imported into an NSSA and then redistributed into the rest of the OSPF routing domain. These routes are imported into the NSSA using a new LSA type; Type-7 LSA. These LSAs are flooded within the NSSA and are translated at the NSSA boundary into AS-external LSAs so as to convey the external routing information to other areas.

NSSAs enable routers with limited resources to participate in OSPF routing while also allowing the import of a selected number of external routes into the area. For example, an area which connects to a small external routing domain running RIP may be configured as an NSSA. This will allow the import of RIP routes into this area and the rest of the OSPF routing domain and at the same time, prevent the flooding of other external routing information (learned, for example, through BGP) into this area.

To set the area-type, enter these commands:

```
OmniCore> ip ospf area 0.0.0.100 area-type nssa
OmniCore> ip ospf area 0.0.0.100 area-type show
Area Type                               :nssa
```

All routers in an NSSA must have their OSPF area defined as an NSSA. To configure otherwise will ensure that the router will be unsuccessful in establishing an adjacent in the OSPF domain. Also, routers to be able to originate Type-7 LSA or ABRs that need to translate the Type-7 LSA must have their ASBR status enabled

OSPF Redistribution

Normally, routing information from one routing protocol, such as RIP, will be transparent to another routing protocol, such as OSPF. Through route redistribution, we are able to pass routing information from local or static routes or from other routing protocols such as RIP or BGP, into a OSPF domain. This is done through route redistribution. The route redistribution commands are shown below.

OSPF Route Distribution Commands

Command	Default	Description
ip ospf redistribution	no default	Creates a redistribution entry for non-OSPF routes.
ip ospf redistribution metric	1	Sets the metric value that will be assumed upon the reception of external routes.
ip ospf redistribution metric-type	type2	Specifies the redistribution metric type used.
ip ospf redistrib-filter route-tag	0	Specifies tag value for filtered route.
ip ospf redistrib-filter metric	Value set at redistribution level	Overrides default cost metric specified at ip/ospf/redistribution level.
ip ospf redistrib-filter redistrib-control	all-subnets	Specifies how subnet routes are redistributed.
ip ospf redistrib-filter effect	permit	Specifies affect of a given redistribution filter.
ip ospf redistrib-filter	no default	Creates or deletes a redistribution filter.
ip ospf redistrib-filter effect	permit	Specifies the redistribution filter action for route importation.
ip ospf redistrib-filter metric	0	Sets the cost metric for the redistribution filter.

Redistribution Procedure

This procedure is optional. To create a redistribution policy for a desired protocol, perform the following step.

◆ Note ◆

In order to make changes to redistribution policies or redistribution filters, it is necessary to either:

- Disable the redistribution for the given protocol or filter:

```
OmniCore/ip/ospf> redistribution rip metric type1 disable
```
- Globally disable OSPF:

```
OmniCore/ip/ospf> status disable
```

Then make your changes and re-enable.

1. Enable the OmniCore routing switch as an ASBR. Only an ASBR can redistribute non-OSPF routes into an OSPF network.

```
OmniCore/ip/ospf> asbr enable
```

2. To import and translate routes from non-OSPF protocols, you must enable route redistribution. Routes supported are static, local, RIP and BGP. In addition, both OSPF and the protocol being redistributed must be running on the ASBR, and a redistribution filter must be created. The following example enables the redistribution of routes learned through RIP into OSPF. Please note that when a redistribution policy is created, its status is enabled and subnet routing is also enabled.

```
OmniCore/ip/ospf> redistribution rip create
OmniCore/ip/ospf> redistribution show
Routing      Metric Type  Metric Value Current State
-----
RIP          type2       0           enable
Number of Entries Displayed: 1
```

- a. Specify the metric-type to be used for redistributed routes (see example below). Note that if you want both internal and external cost to be calculated, metric-type must be set to "type1".

```
OmniCore/ip/ospf> redistribution rip metric-type type1 enable
```

Type 2 metrics, which calculate only external cost (calculated weight value from non-OSPF protocols) are enabled by default.

- b. Specifying (see example below) the metric cost at this level will carry down to the redist-filter level as the default.

```
OmniCore/ip/ospf> redistribution rip metric-type type1 metric 10
```

3. Create a redistribution filter. This filter defines which routes will be forwarded to the OSPF domain. To forward all routes:

```
OmniCore/ip/ospf> redist-filter rip 0.0.0.0 0.0.0.0 create
```

- a. For finer granularity, individual routes can be defined. For example, rather than issuing the above statement, repeat for each range or route that will be forwarded:

```
Range: OmniCore/ip/ospf> redist-filter rip 172.17.0.0 255.255.240.0 create
```

```
Route: OmniCore/ip/ospf> redist-filter rip 172.17.4.0 255.255.255.0 create
```

- b. Specify the cost metric to be used for the filter. The metric command is only supported under the type1 metric-type.

```
OmniCore/ip/ospf> redistribution rip metric 10 type1 enable
```