

Configuring OSPF

7

This chapter provides information for configuring the Open Shortest Path First (OSPF) routing protocol on your ERX system.

Topic	Page
Overview	7-2
References	7-4
Features	7-5
Configuration Tasks	7-8
Enabling OSPF	7-9
Aggregating OSPF Networks	7-11
Configuring OSPF Interfaces	7-13
Configuring OSPF Areas	7-19
Configuring Authentication	7-22
Configuring Additional Parameters	7-26
Configuring OSPF for NBMA Networks	7-33
Traffic Engineering	7-34
Using OSPF Routes for Multicast RPF Checks	7-35
OSPF and BGP/MPLS VPNs	7-36
Remote Neighbors	7-36
Monitoring OSPF	7-39

Overview

OSPF is an interior gateway protocol (IGP) that runs within a single autonomous system (AS). Exterior gateway protocols (EGP), such as Border Gateway Protocol (BGP), exchange routing information between ASs.

OSPF is a link state routing protocol, similar to the Intermediate System-to-Intermediate System (IS-IS) routing protocol. It advertises the states of its local network links. This distinguishes OSPF from some IGPs such as Routing Information Protocol (RIP). A distance vector protocol, such as RIP, advertises the distances (that is, the number of hops) to each known destination within the network.

Each participating OSPF router within the AS has an identical database describing the AS's topology. Each individual piece of this database is a particular router's local state. From this database, OSPF calculates a routing table by constructing a shortest-path tree.

OSPF learns the best routes to reachable destinations. It can quickly perceive changes in the topology of an AS and, after a short convergence period, calculate new loop-free routes. This protocol has been designed expressly for the TCP/IP Internet environment, including explicit support for classless interdomain routing (CIDR) and the tagging of externally derived routing information.

This chapter provides direction for customizing basic OSPF settings if you need to do so. For detailed information on the OSPF commands, see the *ERX Command Reference Guide*.

Terms

Table 7-1 defines commonly used OSPF terms.

Table 7-1 OSPF-related terms

Term	Meaning
adjacency	The relationship between selected neighboring routers for exchanging routing information. Not every pair of neighboring routers is adjacent.
area	An OSPF area is a collection of network segments interconnected by routers. It is a region in an OSPF routing domain.
area border router (ABR)	A router that sits on the edge of an OSPF area and routes link state advertisements (LSAs) between areas.
area ID	A unique number that identifies an area. Typically, an IP address.

Table 7-1 OSPF-related terms (continued)

Term	Meaning
authentication	Contains verification information and is 64-bits in length.
authentication type	A simple password is an example. All OSPF protocol exchanges are authenticated. The authentication type is configurable on a per-area basis.
autonomous system (AS)	A group of routers using a common routing protocol to exchange routing information.
autonomous system boundary router (ASBR)	An OSPF router that redistributes routes from other autonomous systems.
classless interdomain routing (CIDR)	CIDR replaces the traditional class structure of IP addresses. In CIDR, an IP network is represented by a prefix and a notation that indicate the IP address and mask—for example, 10.12.8.3/16
designated router	Generates an LSA for the network. Performs other special responsibilities in running the protocol.
domain	A collection of routers that use a common IGP constitutes an OSPF routing domain.
flooding	The distribution and synchronization of the link-state database between OSPF routers.
hello protocol	Establishes and maintains neighbor relationships. Also, dynamically discovers neighboring routers on broadcast networks.
interior gateway protocol (IGP)	A routing protocol that routers within an AS use to exchange information.
link state advertisement (LSA)	A unit of data that describes the local state of a router or network. LSAs are flooded throughout the routing domain.
neighboring routers	Routers that have interfaces to a common network.
nonbroadcast network	A network that has no broadcast capability but supports more than two routers.
Not-so-stubby area (NSSA)	Similar to a stub area but can also import selected external LSAs.
router ID	A 32-bit number that uniquely identifies a router within an AS—for example, 10.10.1.5
stub area	An area that does not get flooded with external LSAs but does carry intra-area and interarea routes and a default route.
virtual link	A logical link between two backbone routers where the link tunnels through a nonbackbone area.

Figure 7-1 illustrates the topology of an OSPF routing domain.

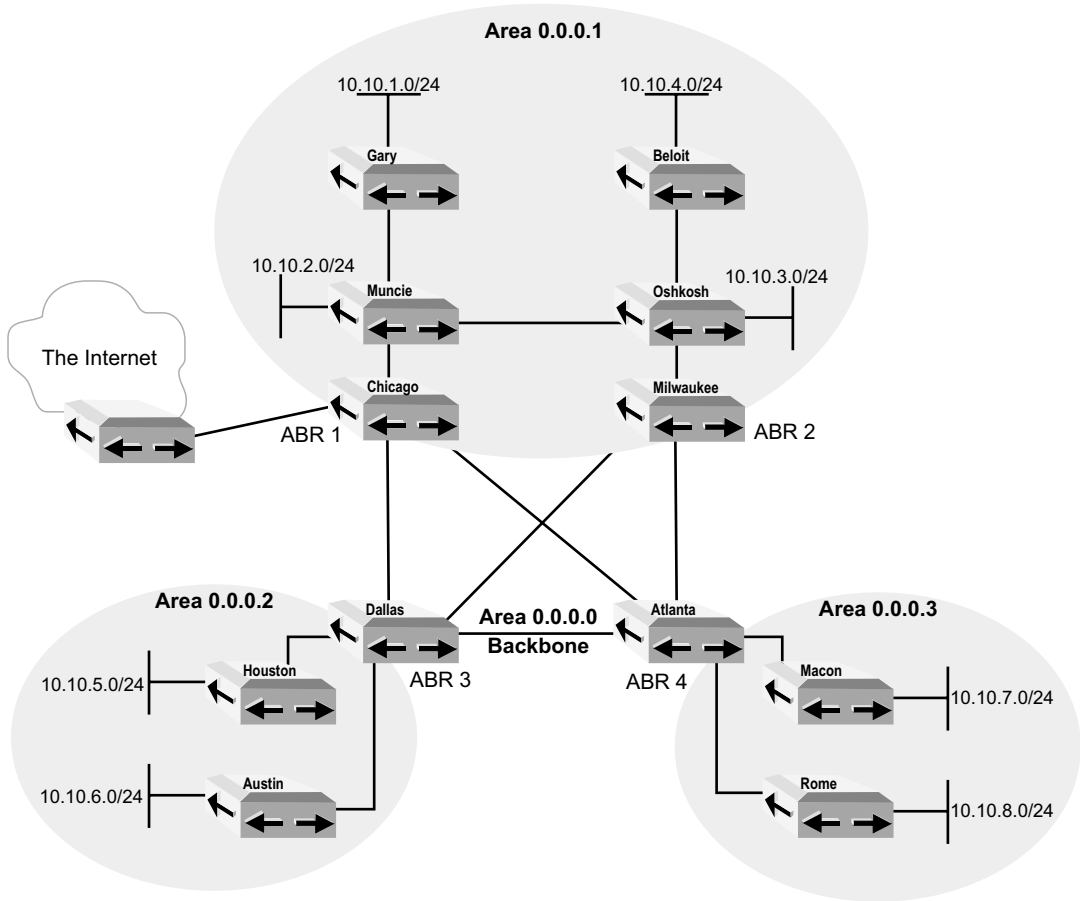


Figure 7-1 OSPF topology

References

If you need more information about the OSPF protocol, consult the following resources:

- *ERX Release Notes, Appendix A, System Maximums* – refer to the Release Notes corresponding to your software release for information on maximum values.
- RFC 2328 – OSPF Version 2 (April 1998)
- RFC 2370 – The OSPF Opaque LSA Option (July 1998)
- Traffic Engineering Extensions to OSPF – draft-katz-yeung-ospf-traffic-06.txt (April 2002 expiration)



Note: IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF Web site at <http://www.ietf.org> for the latest drafts.

Features

The following sections provide a brief description of key OSPF features supported in our implementation of OSPF.

Intra-area, Interarea, and External Routes

You can split up an OSPF AS into areas. Doing this reduces the size of the link state database (LSDB). Each OSPF area runs as a separate network and maintains its own LSDB. OSPF computes routes only to destinations within the area and does not flood routes beyond the area boundaries.

Routing Priority

OSPF areas receive routes based on priority. Table 7-2 lists the routing priority.

Table 7-2 Routing priority

Priority	Type	Description
1 (highest)	Intra-area	Intra-area routing. Refers to routing within a single OSPF area.
2	Interarea	Interarea routing. Refers to routing between OSPF areas within a single OSPF routing domain.
3	External	External type 1. Refers to routing from other protocols that can be imported into the OSPF domain and readvertised by OSPF as type 1 external. Type 1 metric is comparable to the link state metric; the cost is equal to the sum of the internal costs plus the external cost.
4 (lowest)	External	External type 2. Refers to routing from other protocols that can be imported into the OSPF domain and readvertised by OSPF as type 2 external. Type 2 metric is much larger than the cost of any intra-AS path; the cost is equal to the external cost. This is the OSPF default.

If you use the **redistribute** command to import routes from other protocols or sources, the routes default to external type 2. You can specify a route map with the **redistribute** command to modify the type. Alternatively, you can use the **metric-type** keyword with the **redistribute** command to specify the type.

Virtual Links

Each OSPF area must be directly connected to the backbone area. The backbone is responsible for distributing routing information between nonbackbone areas. All routers in the backbone must be contiguous, but they need not be physically adjacent. You can configure backbone routers to be logically adjacent by creating OSPF virtual links.

Authentication

OSPF supports three modes of authentication:

- Null authentication – implies no authentication is in use
- Simple password authentication – requires a 64-bit unencrypted password in each OSPF packet
- Cryptographic authentication – uses a shared secret key that is configured on each router on a network. RFC 2328 defines the use of OSPF cryptographic authentication with the MD5 algorithm.

Opaque LSAs

OSPF opaque LSAs provide a generalized way of extending OSPF. The system generates opaque LSAs to carry traffic engineering (TE) information, accepts them from other routers, and floods them accordingly. OSPF uses the TE information to build a database from which paths can be computed for MPLS label-switched paths.

Route Leakage

Routes can be leaked into OSPF or from OSPF as follows:

- Route leakage into OSPF – When another routing protocol adds a new route to the routing table, or when a static route is added to the routing table, OSPF can be informed through the **redistribute** commands. When OSPF learns the new route, it floods the information into the routing domain using external LSAs.
- Route leakage from OSPF – OSPF adds routing information to the routing table, which is used in forwarding IP packets.

Equal-Cost Multipath

OSPF inherently supports the notion of equal-cost multipath (ECMP). When building the shortest-path tree, OSPF calculates all paths of equal cost to a given destination. If equal-cost paths exist, OSPF inserts into the routing table the next hops for all equal-cost paths to a destination.

OSPF MIB

See the *ERX System Software CD*, shipped with your system, for complete information on the OSPF Management Information Base (MIB) supported by your system. In the MIBs folder you will find information on all supported standard and Juniper Networks ERX enterprise (proprietary) MIBs. OSPF does not act as a host within the system and therefore does not support the `ospfIfMetric` and `ospfHost` tables.

Interacting with Other Routing Protocols

OSPF interacts seamlessly with other routing protocols.

With IS-IS

OSPF was developed originally from an early version of the IS-IS intradomain routing protocol. OSPF can receive IS-IS routing information. See *Chapter 8, Configuring IS-IS*.

With RIP

ERX systems can simultaneously run OSPF and RIP. When doing so, OSPF routes are preferred over RIP. In general, use of the OSPF protocol is preferred because of its robustness, responsiveness, and decreased bandwidth requirements. See *Chapter 6, Configuring RIP*.

With BGP

The default expectation is that your routing environment is an AS running OSPF and exchanging BGP routes with other ASs. See *ERX Routing Protocols Configuration Guide, Vol. 2, Chapter 1, Configuring BGP Routing*.

Configuration Tasks

Configuring OSPF calls for careful coordination among a variety of routers:

- Routers internal to a single area
- Routers that link multiple areas within a single routing domain; these routers are called area border routers (ABRs).
- Routers that link multiple routing domains; these routers are called autonomous system boundary routers (ASBRs).

To configure OSPF:

- 1 Enable OSPF.
- 2 Configure and aggregate network ranges.
- 3 Create the router's OSPF network interfaces.
- 4 Define the OSPF areas attached to the router.

Configuring OSPF Interfaces

You can create OSPF interfaces in the following ways:

- You can issue the **network area** command, which creates OSPF interfaces for all IP interfaces with IP addresses within the specified range.
- You can issue the **address area** command, which creates an OSPF interface in the specified area that sits on top of the IP interface at the given IP address (or on the unnumbered interface, if that is specified).

You can delete OSPF interfaces in the following ways:

- You can issue the **no network area** command, which deletes all OSPF interfaces within the specified range.
- If the OSPF interface was created with the **address area** command, you can issue the **no address area** command to delete the specified interface.
- You can issue the **no ip address** command to delete the IP interface associated with the OSPF interface and also the OSPF interface itself.



Note: *If an OSPF interface is configured on top of an IP interface and you delete the IP interface, the corresponding OSPF interface is also deleted. The previously configured network range, however, is not deleted. You must issue the **no network area** command to delete the range.*

Enabling OSPF

When you enable OSPF on your system, you can create either a range of OSPF interfaces or a single OSPF interface.

Creating a Range of OSPF Interfaces

To create a range of OSPF interfaces:

- 1 Create an OSPF routing process.
- 2 Create the range of IP addresses associated with the routing process and the corresponding OSPF interfaces.
- 3 Assign an area ID associated with each range of IP addresses.

Each router running OSPF has a database describing a map of the routing domain. This map needs to be identical in all participating routers.

network area

- Use to configure a range of OSPF interfaces and their related area.
- If the specified range matches one or more of the IP addresses configured for IP interfaces, one or more corresponding OSPF interface(s) is created and placed in the specified area.
- Create address ranges that do not overlap; you can attach only the same range of interfaces to a single area.
- You cannot use this command for unnumbered interfaces.
- Example 1 shows the creation of one OSPF interface in the backbone area:

```

host1(config-if)#ip address 2.2.2.1 255.255.0.0
host1(config-if)#ip address 2.2.1.1 255.255.0.0 secondary
host1(config)#router ospf 2
host1(config-router)#network 2.2.2.0 0.0.0.255 area 0

```

- Example 2 shows the creation of two OSPF interfaces, one in the backbone area and one in a non-backbone area:

```

host1(config-if)#ip address 2.2.2.1 255.255.255.0
host1(config-if)#ip address 2.2.1.1 255.255.255.0 secondary
host1(config)#router ospf 2
host1(config-router)#network 2.2.2.0 0.0.0.255 area 0
host1(config-router)#network 2.2.1.0 0.0.0.255 area 1

```

This sequence of commands creates two OSPF ranges (2.2.2.0/24 and 2.2.1.0/24), with each range belonging to a different area. Area 0 is configured for 2.2.2.0/24, and area 1 is configured for 2.2.1.0/24. This sequence also creates two OSPF interfaces: one in the backbone area (area 0) using IP address 2.2.2.1, the second in a nonbackbone area (area 1) using IP address 2.2.1.1. This command also creates the two areas if they do not already exist.

- Use the **no** version to delete OSPF interfaces, ranges, and areas.



Note: The configured network range is not active for summarization until you activate this range for summaries by issuing the area range command. The only range that is active by default if you do not issue the **area range** command is the network that matches the IP interface's network exactly. (In other words, by default the exact network of the IP interface is going to be summarized into other areas.)



Note: Active for summarization means that the network range is summarized through area summaries—for ABRs only. See the *Aggregating OSPF Networks* section in this chapter.

ospf enable

- Use to enable OSPF on the router.
- OSPF is enabled by default.
- Example

```
host1(config-router)#ospf enable
```
- Use the **no** version to disable OSPF on the router.

router ospf

- Use to set an OSPF process ID.
- The process ID can be any positive integer 1–65535.
- You must assign a unique ID for each OSPF routing process.
- From a virtual router context you can specify a VRF name. Doing so changes the context to that of the specified VRF and remains so until you exit from the OSPF router context.
- Example

```
host1(config)#router ospf 5
```
- Use the **no** version to end the designated OSPF routing process.

Creating a Single OSPF Interface

To create a single OSPF interface:

- 1 Create an OSPF routing process.
- 2 Create the OSPF interface associated with the IP interface at the specified address.

Each router running OSPF has a database describing a map of the routing domain. This map needs to be identical in all participating routers.

address area

- Use to create an interface in an area on which OSPF runs, on top of the IP interface at the specified IP address.
- You can specify either an IP address or an unnumbered interface.
- Configures OSPF with the default values. You can configure the interface with a nondefault value by using the other **address** commands. You must first issue the **address area** command before issuing any other **address** commands. See *Configuring OSPF Interfaces* in this chapter for more information.
- Example

```
host1(config-router)#address 10.10.32.100 area 0.0.0.0
```

- Use the **no** version to delete the OSPF interface.

ospf enable

- Use to enable OSPF on the router.
- OSPF is enabled by default.
- Example

```
host1(config-router)#ospf enable
```

- Use the **no** version to disable OSPF on the router.

router ospf

- Use to set an OSPF process ID.
- The process ID can be any positive integer 1–65535.
- You must assign a unique ID for each OSPF routing process.
- Example

```
host1(config)#router ospf 5
```
- Use the **no** version to end the designated OSPF routing process.

Aggregating OSPF Networks

You can aggregate OSPF networks at the border of an OSPF area by using the **area range** command. You can also aggregate OSPF networks when entering the border of the OSPF domain by using the **summary-address** command.

area range

- Use to aggregate OSPF routes at an OSPF area border.
- Use only for ABRs.
- You can configure multiple instances of the **area range** command for a single OSPF area.
- The following example shows the steps you must take to create an area range.
 - 1 Configure the interface's IP address(es) using the **ip address** command.
 - 2 Enable OSPF using the **router ospf** command.
 - 3 Configure the network area with the **network area** command.
 - 4 Configure the area range with the **area range** command.
- Example

```
host1(config-if)#ip address 2.2.10.1 255.255.255.0
host1(config-if)#ip address 2.2.11.1 255.255.255.0 secondary
host1(config)#router ospf 2
host1(config-router)#network 2.2.0.0 0.0.255.255 area 0
```

At this point, the OSPF process is configured with two OSPF interfaces. If your router is an ABR, two networks must be summarized: 2.2.10.0/24 and 2.2.11.0/24.

```
host1(config-router)#area 0 range 2.2.0.0 255.255.0.0
```

After you enter this **area range** command, only the aggregated range 2.2.0.0/16 is going to be summarized.

- By default, the range of configured networks is advertised in type 3 (summary) LSAs.
- Use the **do-not-advertise** keyword to prevent advertisement of configured networks.
- Use the command **no area area-id** (with no other keywords) to remove the specified area from the configuration.
- Use the **summary-address** command to summarize external routes being redistributed into OSPF.
- Use the **no** version to disable the aggregation of routes at the OSPF area border.

summary-address

- Use to aggregate external routes at the border of the OSPF routing domain.
- Use only for ASBRs.
- The ASBR advertises one external route as an aggregate for all redistributed routes that are covered by the address.
- For OSPF, this command summarizes only routes from other routing protocols that are being redistributed into OSPF.
- With this command, you can reduce the load of advertising many OSPF external routes by specifying a range that includes some (or all) of these external routes.
- Example

```
host1(config-router)#summary-address 10.1.0.0 255.255.0.0
```

- Use the **area range** command for route summarization between OSPF areas.
- Use the **no** version to restore the default.

Configuring OSPF Interfaces

You can configure OSPF attributes for either a single OSPF network using the **address** commands, or for all OSPF networks on a particular media interface using the **ip ospf** commands.

OSPF MTU size is negotiated rather than configured. OSPF database description exchange uses the interface MTU to signal the largest OSPF MTU that can be sent over an OSPF interface without fragmentation.

Perform the following tasks to configure OSPF attributes for the OSPF network(s):

- Set the cost
- Set the dead interval
- Set the hello interval
- Set router priority
- Set the retransmit interval
- Set the transmit delay



Note: Before using the **address** or **ip ospf** commands, see *Precedence of Commands* on p 7-18 for information on the relationship between these commands.

address Commands

You can use the **address area** command to create a new OSPF interface. Use the other **address** commands to configure parameters for OSPF interfaces that already exist.



Note: You must first issue the **address area** command before issuing any other **address** command.



Note: The **address** commands configure OSPF attributes for a single OSPF network. The **ip ospf** commands configure OSPF attributes for all OSPF networks in the given interface context—for example, in a multinet environment where multiple IP networks sit on top of an Ethernet interface.

address area

- Use to create a new OSPF interface and configure the area ID.
- The interface can have an IP address, or it can be unnumbered.
- Example


```
host1(config-router)#address 10.12.10.2 area 3
```
- You must first issue the **address area** command before issuing any other **address** commands.
- Use the **no** version to delete the area ID from the specified interface.

address cost

- Use to specify the cost metric for the interface. The cost is used in calculating the SPF routing table and can range from 0–65535 seconds.
- The interface can have an IP address, or it can be unnumbered.
- Example


```
host1(config-router)#address unnumbered area 3
host1(config-router)#address unnumbered atm 4/0.1 cost 50
```
- Use the **no** version to reset the path cost to the default value, 10.

address dead-interval

- Use to specify the time period that the system's neighbors should wait without seeing hello packets from the system before they declare the system to be down.
- The dead interval can range from 1–65535 seconds, and is advertised by the system's hello packets.
- For the OSPF routers to become adjacent, the dead interval has to be identical on each router.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 192.168.10.32 area 6
host1(config-router)#address 192.168.10.32 dead-interval 60
```

- Use the **no** version to reset the dead interval to the default value, 40.

address hello-interval

- Use to specify the interval between hello packets that the system sends on the interface.
- The hello interval can range from 1–65535 seconds.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 192.168.1.1 area 5
host1(config-router)#address 192.168.1.1 hello-interval 25
```

- Use the **no** version to reset the hello interval to the default value, 10.

address passive-interface

- Use to disable the transmission of routing updates on the interface, meaning that OSPF routing information is neither sent by nor received through the interface.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 192.168.100.20 area 5
host1(config-router)#address 192.168.100.20
passive-interface
```

- Use the **no** version to reenables the transmission of routing updates.

address priority

- Use to specify the router priority, an 8-bit number from 1 to 255. Used in determining the designated router for the particular network.
- Applies only to nonbroadcast multiaccess (NBMA) networks. Every broadcast and NBMA network has a designated router.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address unnumbered area 6
```

```
host1(config-router)#address unnumbered loopback 0 priority
```

- Use the **no** version to restore the default value, 1.

address retransmit-interval

- Use to specify the time between LSA retransmissions for the interface when an acknowledgment for the LSA is not received.
- Specify an interval in the range 1–65535 seconds; the default value is 5.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 192.168.10.200 area 6
host1(config-router)#address 192.168.10.200
retransmit-interval
```

- Use the **no** version to restore the default value, 5.

address transmit-delay

- Use to specify the estimated time it takes to transmit a link state update packet on the interface.
- Specify an interval in the range 1–65535 seconds; the default value is 1.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 10.100.25.38 area 7
host1(config-router)#address 10.100.25.38 transmit-delay
```

- Use the **no** version to restore the default value, 1.

ip ospf Commands

The **ip ospf** commands have two effects on interface configuration. These effects apply to all **ip ospf** commands:

- Configuration per logical IP interface (for example, Fast Ethernet 0/1.3 or ATM 5/0.1):

The **ip ospf** command configures the specified OSPF parameter(s) for all networks configured on the given IP interface—for example, all multinetted addresses on an interface.

The **no** version of the command resets the specified parameter(s) to *unspecified*.

If the **no** version of the command takes effect, then for the specified IP interface there is no default value for the specified parameter(s). The parameter is set back to unspecified values. However, the value of the specified parameter for the OSPF interface is set back to the default value or the value previously specified by the **address** command.

- Configuration per OSPF interface:

The **ip ospf** command configures the specified OSPF parameter(s) for *each* OSPF interface which sits on top of the IP interface.

The **no** version of the command restores the specified parameter(s) to the default value(s).



Note: If you use **ip ospf** commands to configure OSPF interfaces, you must first have created the interfaces with the **network area** command.



Note: The **ip ospf** commands configure OSPF attributes for all OSPF networks in the given interface context—for example, in a multinet environment where multiple IP networks sit on top of an Ethernet interface. The **address** commands configure OSPF attributes for a single OSPF interface.

ip ospf cost

- Use to configure the cost of sending a packet on the network.
- Cost is a metric value in the range 0–65535; the default value is 1.
- The router LSA advertises the link state metric as the link cost.
- Example

```
host1(config)#interface fastethernet 0/0
host1(config-if)#ip ospf cost 50
```

- Use the **no** version to reset the path cost to the default value, 1.

ip ospf dead-interval

- Use to configure the interval since the last hello packet was seen.
- Specify an interval in the range 1–65535 seconds; the default value is 40.
- For the OSPF routers to become adjacent, the dead interval has to be identical on each router.
- The router's hello packets advertise this interval.
- Example

```
host1(config-if)#ip ospf dead-interval 60
```

- Use the **no** version to restore the default value, 40.

ip ospf hello-interval

- Use to configure the interval between hello packets.
- Specify an interval in the range 1–65535 seconds; the default value is 10.
- For the OSPF routers to become adjacent, the hello interval has to be identical on each router.
- Example

```
host1(config-if)#ip ospf hello-interval 8
```

- Use the **no** version to restore the default value, 10.

ip ospf priority

- Use to configure the router's priority.
- Select a priority level in the range 0–255; the default value is 1.
- This setting determines the designated router for the particular network.
- A router whose priority is set to 0 cannot be a designated router.
- Configure priority only for interfaces to multiaccess networks.
- Example

```
host1(config-if)#ip ospf priority 2
```
- Use the **no** version to restore the default value, 1.

ip ospf retransmit-interval

- Use to configure the time interval between retransmission of an LSA.
- Specify an interval in the range 1–65535 seconds; the default value is 5.
- Example

```
host1(config-if)#ip ospf retransmit-interval 10
```
- Use the **no** version to return to the default value, 5.

ip ospf transmit-delay

- Use to configure the time it takes to transmit a link state update on the interface.
- This is the time between transmissions of LSAs.
- Specify an interval in the range 1–65535 seconds; the default value is 1.
- In setting the time, consider the interface's transmission and propagation delays.
- Example

```
host1(config-if)#ip ospf transmit-delay 4
```
- Use the **no** version to return to the default value, 1.

Comparison Example

Suppose you configure a range of OSPF interfaces with the **network area** command as follows:

```
host1(config)#interface fastEthernet 0/0
host1(config-if)#ip address 1.1.1.1 255.255.255.0
host1(config-if)#ip address 2.2.2.2 255.255.255.0 secondary
host1(config-if)#exit
host1(config)#router ospf 1
host1(config-router)#network 1.1.1.0 0.0.0.255 area 0
host1(config-router)#network 2.2.2.0 0.0.0.255 area 0
```

If you want to specify the cost, you can do so for both interfaces simultaneously:

```
host1(config)#interface fastEthernet 0/0
```

```
host1(config-if)#ip ospf cost 30
```

You can use **address** commands to create a third OSPF interface over the Ethernet interface. When you specify a cost, it is set for only that interface:

```
host1(config)#interface fastEthernet 0/0
host1(config-if)#ip address 3.3.3.3 255.255.255.0 secondary
host1(config-if)#exit
host1(config)#router ospf 1
host1(config-router)#address 3.3.3.3 area 0
host1(config-router)#address 3.3.3.3 cost 25
```

Precedence of Commands

For a single OSPF interface, when you modify the same OSPF attribute by issuing both the **ip ospf** command and the **address** command, the value configured with the **address** command takes precedence. In other words, the most specific command for a single OSPF interface takes precedence.

Consider the following example. Suppose you have a numbered IP interface with an IP address of 10.10.1.1/24 sitting on top of Fast Ethernet interface 0/0. Configure a single OSPF interface on top of the IP interface.

```
host1(config)#router ospf 100
host1(router-config)#address 10.10.1.1 area 0
```

The default cost for this OSPF interface is 10. Change the cost for this OSPF interface using the **address cost** command:

```
host1(router-config)#address 10.10.1.1 cost 45
```

The cost for OSPF interface 10.10.1.1 is now 45.

Now use the **ip ospf cost** command to change the cost for this OSPF interface:

```
host1(config)#int FastEthernet 0/0
host1(config-if)#ip ospf cost 23
```

The cost of OSPF interface 10.10.1.1 does *not* change. The previously issued **address cost** command is more specific for the interface and takes precedence over the **ip ospf cost** command. You must use the **address cost** command if you want to change the cost again:

```
host1(router-config)#address 10.10.1.1 cost 23
```

Configuring OSPF Areas

You can divide your OSPF routing domain into OSPF areas. Doing this provides the following benefits:

- Reduces resource demands placed on routers and links
- Reduces the router CPU usage by the OSPF routing calculation
- Reduces the amount of memory used for link state databases
- Hides subnets within areas from the rest of the routing domain
- Increases routing security within the area

You must attach each area in your routing domain to an area called the backbone area (0.0.0.0).

Disadvantages to using OSPF areas include the following:

- Areas hide information, which may result in less-than-optimal data paths.
- Creating areas complicates the task of configuring OSPF routing domains.

You can optionally define an area to be a stub area or a not-so-stubby area. You can configure virtual links for areas that are not directly connected to a backbone area.

area ***no area***

- There is no affirmative version of this command; there is only a **no** version.
- Example

```
host1(config-router)#no area 47.0.0.0
```
- Use the **no** version to remove the specified area only if no OSPF interfaces are configured in the area.

area default-cost

- Use to configure the cost for the default summary route sent into a stub area.
- Cost is a metric value in the range 1–65535; the default value is 1.
- Use only on an ABR attached to a stub area.
- Provides the metric for the summary default route that the ABR generates into the stub area.
- Example

```
host1(config-router)#area 47.0.0.0 default-cost 1
```
- Use the **no** version to remove the configured default route cost.

area nssa

- Use to configure the area as an NSSA.
- An NSSA is like a stub area, but it can also import external AS routes in a limited way.
- To cause NSSA border routers to generate a type 7 default LSA in the OSPF database if there is a default route in the routing table, you must specify the **default-information-originate** option.
- You can specify a metric cost, metric type, or a route map to be applied to the generated type 7 default LSAs.
- Example

```
host1(config-router)#area 35.0.0.0 nssa
```

- Use the **no** version to remove the NSSA designation from the area, to stop the generation of type 7 default LSAs, or to stop the application of the specified metric cost, metric type, or a route map to the type 7 default LSAs.

area stub

- Use to configure a stub area. Stub areas do not get flooded with external LSAs but do carry a default route, intra-area routes, and interarea routes. This reduces the size of the area's OSPF database and decreases memory usage for external routers in the stub area.
- You must configure each router in a stub area as belonging to the stub area.
- You cannot configure virtual links across a stub area.
- Stub areas cannot contain ASBRs.
- Example

```
host1(config-router)#area 47.0.0.0 stub
```

- Use the **no** version to disable this function.

area virtual-link

- Use to configure an OSPF virtual link.
- A virtual link is used for areas that do not have a direct connection to the backbone area.
- In order to have configured virtual links, the router itself must be an ABR.
- Virtual links are identified by the router ID of the other endpoint, which is also an ABR.
- The two endpoint routers must be attached to a common area, called the virtual link's transit area.
- Virtual links are part of the backbone and behave as if they were unnumbered point-to-point networks between the two routers.
- A virtual link uses the intra-area routing of its transit area to forward packets.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.8.4.2
```

- Use the **no** version to remove an OSPF virtual link.

area virtual-link dead-interval

- Use to set the time in seconds to wait before declaring a neighbor down after not receiving packets from that neighbor.
- Specify an interval in the range 1–65535 seconds; the default value is 40.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.8.4.2  
dead-interval 10
```

- Use the **no** version to remove the virtual link's dead interval.

area virtual-link hello-interval

- Use to configure the hello interval on an OSPF virtual link.
- Specify an interval in the range 1–65535 seconds; the default value is 10.
- The hello interval is the time between the transmission of hello packets.
- The hello interval must be the same for all routers attached to a common network.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.8.4.2  
hello-interval 10
```

- Use the **no** version to remove the virtual link's hello interval.

area virtual-link retransmit-interval

- Use to configure the retransmission interval on an OSPF virtual link.
- The retransmit interval is the time between retransmissions of link state advertisements for adjacencies belonging to the interface.
- Specify an interval in the range 1–65535 seconds; the default value is 5.
- Set the value greater than the expected round-trip delay.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.8.4.2  
retransmit-interval 6
```

- Use the **no** version to remove the interface's retransmit interval.

area virtual-link transmit-delay

- Use to configure the estimated time it takes to transmit a link state update packet on the virtual link.
- Specify an interval in the range 1–65535 seconds; the default value is 1.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.8.4.2  
transmit-delay 1
```

- Use the **no** version to remove the interface's transmit delay.

automatic-virtual-link

- Use to enable an automatic virtual link configuration.

- If enabled, then backbone connectivity is ensured by the automatic creation of a virtual link between this backbone router that has an interface to a common nonbackbone area and other backbone routers that have interfaces to a common nonbackbone area.
- Example


```
host1(config-router)#automatic-virtual-link
```
- Use the **no** version to disable an automatic virtual link.

Configuring Authentication

The system supports the following authentication capabilities:

- Null authentication
- Simple password authentication
- MD5 authentication

The MD5 algorithm takes as input a message of arbitrary length and produces a 128-bit *fingerprint* or *message digest* of the input. MD5 is used to create digital signatures. It is a one-way *hash* function, meaning that it takes a message and converts it into a fixed string of digits, called a message digest.

When using a one-way hash function, you can compare a calculated message digest with the message digest that is decrypted using a public key (password). The key verifies that the message has not been tampered with. This comparison process is called a hashcheck.



Note: You must first issue the **address area** command before issuing any other **address** command.

address authentication-key

- Use to assign a password used by neighboring routers for OSPF simple password authentication.
- The interface can have an IP address, or it can be unnumbered.
- You can specify whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.
- The password, or key, is a character string up to 8 characters in length.
- Example


```
host1(config-router)#address 10.12.10.2 authentication-key  
9rdf7
```
- Use the **no** version to delete the password from the specified interface.

address authentication message-digest

- Use to specify that MD5 authentication is used for the OSPF interface.

- You must configure the MD5 key ID and password with the **address message-digest-key md5** command.
- Switching between authentication types does not delete a configured MD5 key ID or password; only using the **no** version of that configuration command can delete the MD5 key ID and password.
- Example

```
host1(config-router)#address 10.12.10.2 authentication
message-digest
```
- Use the **no** version to set authentication for the interface to none without removing any configured MD5 key. You could subsequently apply MD5 authentication to the interface without having to reconfigure the key.

address authentication-none

- Use to disable authentication on the interface.
- The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 192.168.10.32
authentication-none
```
- The **no** version has no effect.

address message-digest-key md5

- Use to enable OSPF MD5 authentication and configure the MD5 key.
- The MD5 key is a character string up to 16 characters long. You must also specify a key identifier and whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.
- Configures an interface already created, or creates a new OSPF interface and configures the MD5 key. The interface can have an IP address, or it can be unnumbered.
- Example

```
host1(config-router)#address 10.1.1.1 message-digest-key 1
md5 0 9mwk6gdr76
```
- Use the **no** version to delete the MD5 key.

area virtual-link authentication-key

- Use to configure a simple password for a virtual link.
- You can specify whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.
- The password can be up to eight characters long.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.3.4.5
authentication-key sadsa29c
```
- Use the **no** version to remove the password.

area virtual-link authentication message-digest

- Use to specify that MD5 authentication is used for the particular virtual link.
- You must configure the MD5 key ID and password with the **area virtual-link message-digest-key md5** command.
- Switching between authentication types does not delete a configured MD5 key ID or password; only using the **no** version of that configuration command can delete the MD5 key ID and password.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.2.3.4
authentication message-digest
```

- Use the **no** version to set authentication for the virtual link to none without removing any configured MD5 key. You could subsequently apply MD5 authentication to the virtual link without having to reconfigure the key.

area virtual-link authentication-none

- Use to specify that no authentication is used for the particular virtual link.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 27.2.3.4
authentication-none
```

- The **no** version has no effect.

area virtual-link message-digest-key md5

- Use to enable MD5 authentication and to configure MD5 keys for virtual links.
- The MD5 key is a character string up to 16 characters long. You must also specify a key identifier and whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.
- Example

```
host1(config-router)#area 27.0.0.0 virtual-link 327.3.4.5
message-digest-key 2 md5 rc45lsm2c
```

- Use the **no** version to remove the password.

ip ospf authentication-key

- Use to configure a type 1 authentication (a simple password) on the interface.
- Neighboring OSPF routers use the password to access the router's interface.
- Use the same password on all neighboring routers on the same network.
- Use this password only when you enable authentication for the interface.
- You can specify whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.
- Use a password that is a continuous string up to 8 characters long.
- Example

```
host1(config-if)#ip ospf authentication-key yourpwd
```

- Use the **no** version to remove the password on the interface.

ip ospf authentication message-digest

- Use to specify the authentication method for the interface as MD5.
- You must configure the MD5 key ID and password with the **ip ospf message-digest-key md5** command.
- Switching between authentication types does not delete a configured MD5 key ID or password; only using the **no** version of that configuration command can delete the MD5 key ID and password.

- Example

```
host1(config-if)#ip ospf authentication message-digest
```

- Use the **no** version to set authentication for the interface to none without removing any configured MD5 key. You could subsequently apply MD5 authentication to the interface without having to reconfigure the key.

ip ospf authentication-none

- Use to specify that no authentication is used for the OSPF interface.
- Example

```
host1(config-if)#ip ospf authentication-none
```

- The **no** version has no effect.

ip ospf message-digest-key md5

- Use to enable MD5 authentication on the OSPF interface and configure the MD5 key.
- The MD5 key is a character string up to 16 characters long. You must also specify a key identifier and whether the key is entered in unencrypted or encrypted format. If you do not specify which, the string is assumed to be unencrypted.

- Example

```
host1(config-if)#ip ospf message-digest-key 3 md5 0 tre987is
```

- Use the **no** version to delete an MD5 key from the OSPF interface.



Note: If all the MD5 keys have been deleted, the authentication type is still MD5, but you need to configure MD5 keys.



Note: To disable MD5 authentication for the interface, use the **ip ospf authentication-none** command.



Note: To display the password only in encrypted text, use the **service password-encryption** command.

Configuring Additional Parameters

The commands presented in this section include both OSPF-specific commands and routing protocol-independent commands that are not limited to OSPF. You can use these commands to perform the following tasks:

- Add an access list and a route map.
- Set a baseline for statistics.
- Set the redistribution routes.
- Set the distance for OSPF routes.
- Enable logs for OSPF neighbor changes.
- Administratively disable OSPF.
- Set the maximum paths (maximum number of parallel routes).
- Set auto-cost.
- Set SPF hold time.
- Set a default metric.

access-list ***route-map***

- Use the **access-list** command to create a standard or extended access list.
- Use the **route-map** command to create a route map.
- For detailed information on configuring access lists and route maps, see *Chapter 1, Configuring Routing Policy*.
- Example
 - 1 Configure three static routes:

```
host1(config)#ip route 20.20.20.0 255.255.255.0 192.168.1.0
host1(config)#ip route 20.20.21.0 255.255.255.0 192.168.1.0
host1(config)#ip route 20.21.0.0 255.255.255.0 192.168.1.0
```
 - 2 Configure an access list with filters on routes 20.20.20.0/24 and 20.20.21.0/24:

```
host1(config)#access-list boston permit 20.20.0.0
0.0.255.255
```
 - 3 Configure a route map which matches the previous access list and applies a metric type 1 (OSPF):

```
host1(config)#route-map boston
host1(config-route-map)#match ip address boston
host1(config-route-map)#set metric-type type-1
```
 - 4 Configure redistribution of the static routes into OSPF with route map boston:

```
host1(config)#router ospf 2
host1(config-router)#redistribute static route-map boston
```
 - 5 Use the **show ip ospf database** command to verify the effect of the redistribution (the two static routes matching the route map are redistributed as external type 1):

```
host1#show ip ospf database
```

```
OSPF Database
```

```
Router Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.1.250	192.168.1.250	3	0x80000006	0x39a1
192.168.254.7	192.168.254.7	220	0x80000169	0xd2b5

```
Network Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.1.214	192.168.254.7	220	0x80000001	0xe0f2

```
AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum
20.20.20.0	192.168.1.250	3	0x80000001	0x6045
20.20.21.0	192.168.1.250	3	0x80000001	0x554f

- Use the **no** version of the **access-list** command to remove the access list or the specified entry in the access list.
- Use the **no** version of the **route-map** command to remove an entry.

baseline ip ospf

- Use to set a baseline for OSPF statistics and counters.
- The following example first displays the output of the **show ip ospf** command shown before you run the **baseline ip ospf** command; then, the execution of the **baseline ip ospf** command; and finally, the display of the **show ip ospf** command run after you execute the **baseline ip ospf** command.
 - › The output of the **show ip ospf** command run before the **baseline ip ospf** command reflects the up-to-date packet counters.
 - › The output of the **show ip ospf delta** command run after you run the **baseline ip ospf** command reflects the baseline set for OSPF statistics and counters.
- There is no **no** version.
- Example

```
host1#show ip ospf
```

```
Routing Process OSPF 1 with Router ID 5.106.7.1
  OSPF Statistics:
    Rcvd: 217935 total, 0 checksum errors
          8987 hello, 8367 database desc, 188 link state req
          159898 link state updates, 40484 link state acks
    Sent: 265026 total, 0 pkts dropped
          8927 hello, 8341 database desc, 53 link state req
          158571 link state updates, 89134 link state acks
  Supports only single TOS(TOS0) routes
  SPF schedule delay 0 secs, Hold time between two SPFs 3 secs
  Maximum path splits 1
  Area BACKBONE(0.0.0.0)
    Area is a transit area
    SPF algorithm executed 425 times
    ABR count 0
    ASBR count 1
    LSA Count 12
    Number of interfaces in this area is 24
    Area ranges are:
  Number of active areas in this router is 1
  1 normal, 0 stub, 0 NSSA.
```

```
host1#baseline ip ospf
host1#show ip ospf delta
```

```
Routing Process OSPF 1 with Router ID 5.106.7.1
  OSPF Statistics:
```

```
Rcvd: 0 total, 0 checksum errors
      0 hello, 0 database desc, 0 link state req
      0 link state updates, 0 link state acks
Sent: 0 total, 0 pkts dropped
      0 hello, 0 database desc, 0 link state req
      0 link state updates, 0 link state acks
Supports only single TOS(TOS0) routes
SPF schedule delay 0 secs, Hold time between two SPFs 3 secs
Maximum path splits 1
Area BACKBONE(0.0.0.0)
  Area is a transit area
  SPF algorithm executed 425 times
  ABR count 0
  ASBR count 1
  LSA Count 12
  Number of interfaces in this area is 24
  Area ranges are:
Number of active areas in this router is 1
1 normal, 0 stub, 0 NSSA.
```

clear ip ospf redistribution

- Use to clear all the routes that have been previously redistributed into OSPF.
- Example

```
host1#clear ip ospf redistribution
```
- There is no **no** version.

default-information originate

- Use to generate a default route into an OSPF routing domain.
- When you use this command to redistribute routes into an OSPF routing domain, the system automatically becomes an ASBR.
- An ASBR, however, does not, by default, generate a default route into the OSPF routing domain. The software must have a default route before it generates one, except when you have specified the **always** keyword.
- You can specify a metric for the route or specify that the route be OSPF external type 1 or 2.
- Example

```
host1(config)#router ospf 1
host1(config-router)#default-information originate
route-map 5
```
- Use the **no** version to disable this feature.

disable-dynamic-redistribute

- Use to halt the dynamic redistribution of routes that are initiated by changes to a route map.
- Dynamic redistribution is enabled by default.
- Example

```
host1(config-router)#disable-dynamic-redistribute
```

- Use the **no** version to reenable dynamic redistribution.

distance

- Use to configure the administrative distance for OSPF routes.
- Example

```
host1(config-router)#distance ospf external 60
```

- Default settings:
 - › Intra-area routes – 110
 - › Interarea routes – 112
 - › External routes – 114
- Use the **no** version to restore the default values noted above.

ip ospf shutdown

- Use to disable OSPF on the interface.
- Example

```
host1(config-if)#ip ospf shutdown
```

- Use the **no** version to enable OSPF on the interface.

maximum-paths

- Use to control the maximum number of parallel routes that OSPF can support.
- The maximum number of routes can range from 1–16.
- The default for OSPF is 4 paths.
- To enable equal-cost multipath (ECMP) for OSPF, you need to specify a value for *maxPaths* greater than 1.
- Example

```
host1(config-router)#maximum-paths 2
```

- Use the **no** version to restore the default value, 4.

ospf auto-cost reference-bandwidth

- Use to calculate OSPF interface cost according to bandwidth.
- Sets the OSPF metric for an interface according to the bandwidth specified.
- Affects OSPF metrics only for OSPF interfaces created after its execution.
- This command's value overrides the cost resulting from the command.

- If you want this command to apply to OSPF interfaces already configured, you need to bounce the existing interfaces: Use the **no network** and then the **network** command for the selected OSPF interfaces.
- Example

```
host1(config-router)#ospf auto-cost reference-bandwidth 1000
```
- If you issue this command, the metric is calculated as follows:
OSPF metric = bandwidth*1,000,000/link speed
For the previous example, a 64K link will get a metric of 15625, while a T1 link will have a metric of 647. The minimum value for the metric is 1.
- If you never issue the **ospf auto-cost reference-bandwidth** command, OSPF calculates the cost as 10^8 /link speed.
- Use the **no** version to assign cost based only on the interface type.

ospf log-adjacency-changes

- Use to configure the router to send a log message when the state of an OSPF neighbor changes.
- Example

```
host1(config-router)#ospf log-adjacency-changes severity 3
verbosity low
```
- Use the **no** version to halt logging of neighbor changes.

ospf shutdown

- Use to administratively disable OSPF on the router.
- Example

```
host1(config-router)#ospf shutdown
```
- Use the **no** version to reenabling OSPF on the interface.

passive-interface

- Use to disable the transmission of routing updates on the interface, meaning that OSPF routing information is neither sent by nor received through the interface.
- The specified interface appears as a stub network in the OSPF domain.
- By default, OSPF is enabled on a configured OSPF interface.
- Example

```
host1(config-router)#passive-interface ethernet 1/0
```
- Use the **no** version to reenabling the transmission of OSPF routing updates on the specified interface.

redistribute

- Use to redistribute information from a routing domain other than OSPF into the OSPF domain.
- You can set the OSPF metric type—type 1 or type 2—and set a metric for all redistributed routes.

• Example 1

```
host1(config)#router ospf 5
host1(config-router)#redistribute bgp route-map 4
```

- If you do not specify **route-map**, all routes are redistributed. By default, all routes are imported as external type 2 routes.
- If you do specify **route-map**, but do not list any route map tags, no routes are imported.
- Use to redistribute routes from OSPF into other non-OSPF routing domains.

• Example 2

```
host1(config)#router bgp 100
host1(config-router)#redistribute ospf 5
```

- Use the **no** version to disable redistribution.

timers spf

- Use to configure the time between two consecutive SPF calculations.
- Set the time (in seconds) in the range 1–5; the default value is 3.
- If you set the hold time to 0, there is no delay between two consecutive SPF calculations. They can be done one immediately after the other.

• Example

```
host1(config-router)#timers spf 2
```

- Use the **no** version to return to the default value, 3.

Default Metrics

Although the system does not support a **default-metric** command, the **redistribute** command provides two ways to set a default metric for redistributed routes.

You can simply configure a metric with the **redistribute** command to apply to all routes redistributed from the specified source protocol:

```
host1(config)#router ospf 5
host1(config-router)#redistribute bgp metric 5
```

Alternatively, you can create one or more route maps that set the metric and apply them selectively to redistributed routes:

```
host1(config)#access-list 1 permit any any
host1(config)#route-map defmetric
host1(config-route-map)#match ip address 1
```

```
host1(config-route-map)#set metric 10
host1(config-route-map)#exit
host1(config)#router ospf 5
host1(config-router)#redistribute bgp route-map defmetric
host1(config-router)#redistribute isis route-map defmetric
```

See *Chapter 1, Configuring Routing Policy*, for information on configuring route maps.

Configuring OSPF for NBMA Networks

You can configure OSPF nonbroadcast multiaccess (NBMA) networks. You can configure your OSPF network type as NBMA, regardless of the default medium. This is useful when, for example, you have routers in your network that do not support multicast addressing.

You must use the **neighbor** command to specify the router's OSPF neighbors.

To configure OSPF for an NBMA network:

- 1 Configure an interface network type as NBMA for OSPF.

```
host1(config-subif)#ip ospf network non-broadcast
```
- 2 Exit Interface Configuration mode. Enter Global Configuration mode.

```
host1(config-subif)#exit
```
- 3 Configure an OSPF routing process, and enter Router Configuration mode.

```
host1(config)#router ospf 5
```
- 4 Specify an OSPF neighbor, and optionally assign a priority number or poll interval to the neighbor.

```
host1(config-router)#neighbor 10.12.14.1 priority 5
poll-interval 180
```
- 5 Repeat step 4 for each neighbor in the OSPF network.

If you want to configure the network type for a specific interface or OSPF area, rather than for all OSPF interfaces, you can use the **address network** command rather than the **ip ospf network** command.

address network

- Use to configure the network type on a specific OSPF interface or for a specific OSPF area to a type other than the default for the medium.
- You must first issue the **address area** command before issuing the **address network** command.
- Use the **no** version to restore the default value for the medium.

ip ospf network

- Use to configure the network type on all OSPF interfaces on the OSPF network to a type other than the default for the medium.
- Use the **no** version to restore the default value for the medium.

neighbor

- Use to configure OSPF neighbors on the NBMA network.
- Specify priority and poll interval only for routers that are eligible to become the designated router or backup designated router—that is, a router with a nonzero router priority value. The default priority value is 0, and the default polling interval is 120 seconds.
- Use the **no** version to remove the neighbor or restore the default values 0 and 120.

Traffic Engineering

Traffic engineering (TE) enables more effective use of network resources by providing for the setup of explicitly routed Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) that satisfy resource and administrative constraints. You can use OSPF to exchange link resource and traffic engineering administrative information between routers. OSPF uses this information to calculate paths in the network that satisfy the administrative constraints. MPLS can then set up LSPs along these paths. Refer to *ERX Routing Protocols Configuration Guide, Vol. 2, Chapter 2, Configuring MPLS*, for a detailed discussion of MPLS.

Configuring OSPF for TE

For OSPF to support TE, you must issue both of the following commands:

- **mpls traffic-eng area** – enables the router to flood TE resource and administrative information in the specified area using type 10 opaque LSAs. These LSAs have an area-wide scope and therefore are flooded only within the indicated area.
- **mpls traffic-eng router-id** – designates a router as TE capable and specifies the address of a stable router interface as the router ID of the node for TE purposes. The TE router ID serves as the tunnel endpoint

for tunnels terminating at the node. Each node advertises its TE router ID in type 10 LSAs.

You can use the **show ip ospf database opaque-area** command to display information about TE opaque LSAs.

In order for OSPF routes to use established MPLS tunnels as next hops—so that traffic can be mapped to use these tunnels—you must configure the tunnels with the **tunnel mpls autoroute announce ospf** command. See *ERX Routing Protocols Configuration Guide, Vol. 2, Chapter 2, Configuring MPLS*, for information on configuring MPLS on a system.

mpls traffic-eng area

- Use to enable flooding of MPLS TE link information into the specified OSPF area. Flooding is disabled by default.
- Example

```
host1(config-router)#mpls traffic-eng area 0
```
- Use the **no** version to disable flooding.

mpls traffic-eng router-id

- Use to specify a stable interface to be used as a router ID for MPLS TE. Typically you specify a loopback interface to provide the greatest stability, because this is flooded to all nodes. The interface acts as the destination node for tunnels originating at other nodes.
- Example

```
host1(config-router)#mpls traffic-eng router-id loopback 0
```
- Use the **no** version to remove the interface as a router ID.

Using OSPF Routes for Multicast RPF Checks

You can use the **ip route-type** command to specify whether OSPF routes are available for only unicast forwarding protocols or only multicast reverse path forwarding (RPF) checks. Routes available for unicast forwarding appear in the unicast view of the routing table, whereas routes available for multicast RPF checks appear in the multicast view of the routing table.

ip route-type

- Use to specify whether OSPF routes are available only for unicast forwarding, only for multicast reverse path forwarding checks, or for both.
- Use the **show ip route** command to view the routes available for unicast forwarding.

- Use the **show ip rpf-routes** command to view the routes available for multicast reverse path forwarding checks.
- By default, OSPF routes are available for both unicast forwarding and multicast reverse path forwarding checks.
- Example

```
host1(config)#router ospf
host1(config-router)#ip route-type unicast
```
- Use the **no** version to restore the default value, both.

OSPF and BGP/MPLS VPNs

Some network topologies will use OSPF as the routing protocol between CE and PE routers in BGP/MPLS VPNs. See *ERX Routing Protocols Configuration Guide, Vol. 2, Chapter 3, Configuring BGP/MPLS VPNs*, for information on configuring OSPF for this purpose.

Remote Neighbors

You can create OSPF remote neighbors to enable the router to establish neighbor adjacencies through unidirectional interfaces, such as MPLS tunnels, rather than the standard practice of using the same interface for receipt and transmission of OSPF packets. The remote neighbor can be more than one hop away through intermediate routers that are not running OSPF. OSPF uses the interface associated with the best route to reach the remote neighbor. A best route to the neighbor must exist in the IP routing table.

You must explicitly configure a remote neighbor on an OSPF router. You must specify the remote neighbor with which the router will form an adjacency and the source IP address the router will use for OSPF packets destined to its peer remote neighbor.

In order to form an adjacency with its remote neighbor, all OSPF packets are sent to the remote neighbor as unicast packets with the destination IP address equal to the source IP address of the remote neighbor. Use the **update-source loopback** command to assign the source IP address to a remote neighbor.

The connection between two remote neighbors is treated as an unnumbered point-to-point link that resides in the same area as that to which the pair of remote neighbors belongs.

The rules of OSPF adjacency must be followed for remote neighbors to form an adjacency with each other; for example, the neighbors must be in

the same OSPF area and have the same hello interval and dead interval, and so on.

Once you have used the **remote-neighbor** command to specify the remote neighbors and the update-source loopback to assign the source IP address, you must set a TTL value with the **tll** command, because a remote neighbor can be more than one hop away. Configuration of all other remote-neighbor attributes is optional.

authentication-key

- Use to enable simple password authentication and assign a password for communication with OSPF remote neighbors.
- Example

```
host1(config-router-rn)#authentication-key 0 br549hee
```
- Use the **no** version to delete the password.

authentication message-digest

- Use to specify that MD5 authentication is to be used on the OSPF remote neighbor interface.
- Example

```
host1(config-router-rn)#authentication message-digest
```
- There is no **no** version.

authentication-none

- Use to specify that no authentication is to be used on the OSPF remote neighbor interface.
- Example

```
host1(config-router-rn)#authentication-none
```
- There is no **no** version.

cost

- Use to specify a cost metric for the OSPF remote-neighbor interface; the metric is used in the calculation of the SPF routing table.
- The default value is 10 if there is no route to the remote neighbor; otherwise, the default is calculated based on the bandwidth of the physical interface used to reach the remote neighbor and the OSPF auto-cost reference bandwidth.
- Example

```
host1(config-router-rn)#cost 235
```
- Use the **no** version to restore the default value, as described above.

dead-interval

- Use to set the time period that the OSPF router waits without seeing hello packets from a remote neighbor before declaring the neighbor to be down.
- Example

```
host1(config-router-rn)#dead-interval 180
```
- Use the **no** version to restore the default value, 40 seconds.

hello-interval

- Use to set the interval between hello packets that the router sends on the OSPF remote-neighbor interface.
- Specify a value in the range 1–65535 seconds; the default value is 40.
- Example

```
host1(config-router-rn)#hello-interval 15
```
- Use the **no** version to restore the default value, 40.

message-digest-key md5

- Use to enable MD5 authentication for the OSPF remote-neighbor interface and configure the MD5 key.
- Example

```
host1(config-router-rn)#message-digest-key 42 md5 0 sal29ute
```
- If you delete all MD5 keys, MD5 authentication is still enabled; you must either configure an MD5 key or disable MD5 authentication with the **authentication-none** command.
- Use the **no** version to delete the MD5 key.

remote-neighbor

- Use to configure an OSPF remote neighbor.
- Example

```
host1(config-router)#remote-neighbor 10.25.100.14 area 35672
```
- Use the **no** version to remove the remote neighbor and any attributes configured for the remote neighbor.

retransmit-interval

- Use to set the time between LSA retransmissions for the OSPF remote-neighbor interface when an acknowledgment for the LSA is not received.
- Specify a value in the range 1–3600 seconds; the default value is 5.
- Example

```
host1(config-router-rn)#retransmit-interval 10
```
- Use the **no** version to restore the default value, 5.

transmit-delay

- Use to set the estimated time it takes to transmit a link state update packet on the OSPF remote-neighbor interface.
- Specify a value in the range 0–3600 seconds; the default value is 1 second.
- Example

```
host1(config-router-rn)#transmit-delay 3
```
- Use the **no** version to restore the default value, 1.

ttl

- Use to configure a hop count by setting the value of the time-to-live field used by packets sent to an OSPF remote neighbor.
- Specify a value in the range 1–255 seconds; the default value is 1 second.
- Example

```
host1(config-router-rn)#ttl 35
```
- Use the **no** version to restore the default value, 1.

update-source

- Use to specify the loopback interface whose local IP address is used as the source address for the OSPF connection to a remote neighbor.
- Example

```
host1(config-router-rn)#update-source loopback 1
```
- Use the **no** version to delete the source address from the connection to the remote neighbor.

Monitoring OSPF

Two sets of commands enable you to monitor OSPF operation on your system: the **debug** and the **show** commands. Both sets of commands provide information on your router's OSPF state and configuration.

The task you are performing with each of these monitoring commands is basically the same for each command; that is, you are requesting information. The results of this request may vary. For instance, the **debug** commands provide information on problems with the network or the system, whereas the **show** commands provide information on the actual state and configuration of your system.

debug Commands

The **debug** commands provide information on the following OSPF items:

- Adjacencies
- Designated router

- General events
- Link state advertisements
- Neighbors
- Packets received
- Packets sent
- Route events
- SPF events

debug ip ospf

- Use to display information on selected OSPF events. This command has many keywords that allow you to specify a variety of OSPF events.
- You can set the level of severity for the events you want displayed: 0–7.
- You can set the verbosity of the messages you want displayed: low, medium, high.
- Example

```
host1#debug ip ospf adj
```
- Use the **no** version to cancel the display of any information on the designated variable.

ospf log-adjacency-changes

- Use to enable the logging of changes in the state of an OSPF neighbor.
- Example

```
host1(config-router)#ospf log-adjacency-changes
```
- Use the **no** version to disable the logging of changes in the state of an OSPF neighbor.

undebug ip ospf

- Use to cancel the display of information on a selected event.
- The same OSPF variables can be designated as in the **debug ip ospf** command.
- Example

```
host1#undebug ip ospf adj
```
- There is no **no** version.

show Commands

The **show** commands provide information on the following OSPF items:

- Routing processes
- Border routers

- Database
- Interfaces
- Neighbors
- Virtual links
- Internal statistics
- MPLS tunnels and opaque LSAs

You can use the output filtering feature of the **show** command to include or exclude lines of output based on a text string you specify. Refer to *ERX System Basics Configuration Guide, Chapter 2, Command Line Interface*, for details.

show ip ospf

- Use to display general information about OSPF routing processes.
- Field descriptions
 - › Routing process – process ID
 - › Router – router types: internal, area border, or autonomous system boundary routers
 - › OSPF Statistics – packets sent and received
 - › TOS type – number of types of service supported
 - › SPF timers – timers configured on the router
 - › Maximum path splits – maximum equal-cost paths supported
 - › Areas – areas configured and their parameters
 - › Number of areas – number of areas in the router
- Example

```
host1#show ip ospf 2
Routing Process OSPF 2 with Router ID 5.5.0.250
Router is an Area Border Router (ABR)
Router is an Autonomous System Border Router (ASBR)
OSPF Statistics:
  Rcvd: 1406 total, 0 checksum errors
        465 hello, 11 database desc, 3 link state req
        77 link state updates, 21 link state acks
  Sent: 442 total
Supports only single TOS(TOS0) routes
SPF schedule delay 0 secs, Hold time between two SPFs 3 secs
Maximum path splits 16
Area BACKBONE(0.0.0.0)
Area is a transit area
  Area has no authentication
  SPF algorithm executed 7 times
  ABR count 2
```

```

ASBR count 1
LSA Count 37
Area ranges are:
  56.56.56.0/255.255.255.0
Area 0.0.0.1
  Area is a transit area
  Area has no authentication
  SPF algorithm executed 7 times
  ABR count 2
  ASBR count 1
  LSA Count 34
  Area ranges are:
    4.4.4.0/255.255.255.0
Number of active areas in this router is 2
2 normal, 0 stub, 0 NSSA.

```

show ip ospf border-routers

- Use to display a list of OSPF border routers.
- Field descriptions
 - › Destination – destination's router ID
 - › Next Hop – next hop toward the destination
 - › Router Type – router type of the destination: either an ABR or ASBR or both
 - › Route Type – type of this route: either an intra-area or inter-area route
 - › Area – area ID of the area that this route is learned from
- Example

```

host1#show ip ospf border-routers
Destination NEXT HOP   Interface      Router Type  Route Type Area
5.5.0.250   5.5.6.250   fastethernet0 ABR/ASBR    INTRA     0.0.0.0
5.5.0.250   4.4.4.250   fastethernet0 ABR/ASBR    INTRA     0.0.0.1
6.6.6.250   4.4.4.13    fastethernet0 ABR          INTRA     0.0.0.1

```

show ip ospf database

- Use to display either the full OSPF database or a summary of it.
- Field descriptions
 - › Link ID – link state ID of the LSA:
 - for router links, set to the router's OSPF router ID
 - for network links, set to the IP interface address of the network's designated router
 - for type 3 summary LSAs, set to an IP network number
 - for type 4 summary LSAs, set to an ASBR's router ID
 - for type 5 externals, set to an IP network number
 - › Adv Router – advertising router's ID
 - › Age – link state age
 - › Seq# – link state sequence number (detects old or duplicate LSAs)

› Checksum – Fletcher checksum of the complete contents of the LSA

- Example

```
host1#show ip ospf database
```

```
OSPF Database
```

```
Router Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
5.1.101.1	5.1.101.1	932	0x80000069	0x102f
192.168.1.13	192.168.1.13	1763	0x80000099	0xaa4e
192.168.1.10	192.168.1.10	285	0x80000087	0xada6
192.168.1.11	192.168.1.11	401	0x80000087	0xaba5
192.168.24.6	192.168.24.6	622	0x800005bf	0x6087

```
Network Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
56.56.56.220	5.6.6.1	499	0x80000069	0x26a0
192.168.1.12	192.168.254.6	622	0x8000009e	0xebc2

```
Summary Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
4.4.4.0	5.5.0.250	497	0x8000005a	0x2ca6
4.4.4.0	192.168.1.13	528	0x80000059	0x 45d

```
AS Summary Link States (Area 0.0.0.0)
```

Link ID	ADV Router	Age	Seq#	Checksum
5.5.0.250	192.168.1.13	491	0x80000002	0xe9d4

```
AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum
8.8.8.0	5.5.0.250	502	0x8000005f	0x2d67

```
Router Link States (Area 0.0.0.1)
```

Link ID	ADV Router	Age	Seq#	Checksum
5.5.0.250	5.5.0.250	498	0x80000067	0xdec1
192.168.1.13	192.168.1.13	505	0x800000a5	0x3b32

```
Network Link States (Area 0.0.0.1)
```

Link ID	ADV Router	Age	Seq#	Checksum
4.4.4.13	192.168.1.13	505	0x80000001	0x410b
5.1.0.0	192.168.1.13	940	0x80000059	0x82c4
5.2.0.0	5.5.0.250	495	0x80000001	0x51bf

```

5.2.0.0      192.168.1.13    932 0x80000059   0x76cf
5.3.0.0      5.5.0.250       495 0x80000001    0x45ca
5.3.0.0      192.168.1.13    932 0x80000059   0x6ada
56.56.56.0   5.5.0.250       495 0x80000062    0xc469

```

AS Summary Link States (Area 0.0.0.1)

```

Link ID     ADV Router Age Seq#           Checksum
5.5.0.250  5.5.0.250   496 0x80000001    0x51c0

```

show ip ospf database opaque-area

- Use to display lists of information about the TE opaque LSAs.
- The TE router address LSA describes a stable IP address on the originating router that can be used for TE purposes—such as setting up TE LSPs to this address.
- The TE link LSA describes TE information about an interface on the originating router.
- Field descriptions
 - › LS age – age of LSA
 - › Options – optional capabilities supported by the described portion of the routing domain
 - › LS Type – type of LSA; opaque area TE router address or opaque area TE link LSA
 - › Link State ID – link state ID of the opaque LSA
 - › Advertising Router – router ID of the router that originated the LSA
 - › LS Seq Number – link state sequence number to identify duplicate or old LSIDs
 - › Checksum – checksum of the complete contents of the LSA
 - › Length – length of the LSA in bytes
 - › TE router ID – TE router ID of the originating router
 - › Link type – point-to-point or multiaccess
 - › Link ID – for point-to-point interfaces, this is the router ID of the router at the remote end; for multiaccess interfaces, this is the address of the DR
 - › Local Address – IP address of the local interface for the link
 - › Remote Address – IP address of the remote (neighbor's) interface for the link
 - › TE Metric – link metric for traffic engineering purposes; can be different from the standard OSPF link
 - › Max BW – maximum bandwidth that can be used on this link in this direction
 - › Max Reservable BW – maximum bandwidth that can be reserved on this link; can exceed the maximum bandwidth in the event of oversubscription
 - › Max Unreserved BW – amount of bandwidth not yet reserved at each of the eight priority levels; each value is less than or equal to the maximum reservable bandwidth
 - › Color – bitmask that specifies the administrative group membership for this link; a link that is a member of more than one group will have multiple bits set
- Example

```
host1#show ip ospf database opaque-area
```

```
    Opaque-area Link States (Area 0.0.0.0)

    LS age: 914
    Options: (TOS-capable, No Type7-LSA,
    ExternalRoutingCapability, No
    Multicast Capability, No External Attributes   LSA)
    LS Type: Opaque-Area (TE Router Address)
    Link State ID: 1.0.0.0(Instance)
    Advertising Router: 100.1.1.1
    LS Seq Number: 0x80000003
    Checksum: 0xd293
    Length: 28
    TE Router-ID: 100.1.1.1

    LS age: 919
    Options: (TOS-capable, No Type7-LSA,
    ExternalRoutingCapability, No
    Multicast Capability, No External Attributes   LSA)
    LS Type: Opaque-Area (TE Links)
    Link State ID: 1.0.0.1(Instance)
    Advertising Router: 100.1.1.1
    LS Seq Number: 0x80000003
    Checksum: 0xf66e
    Length: 124
    Link Type: P2P
    Link ID: 1744896257
    Local Address 14.1.1.2
    Remote Address 14.1.1.1
    TE Metric 0
    Max BW 1000 kb/sec (125000 Bps)
    Max Reservable BW 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 0 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 1 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 2 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 3 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 4 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 5 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 6 1000 kb/sec (125000 Bps)
    Max Unreserved BW : pri 7 1000 kb/sec (125000 Bps)
    Color 0
```

show ip ospf interface

- Use to display a list of OSPF interfaces.
- Field descriptions

- › Interface value (FastEthernet) – status of the physical link and the operational status of the protocol
- › Internet Address – interface IP address
- › Area – area identifier: IP address
- › Network type – broadcast, NBMA, Point-to-Point, or Point-to-Multipoint
- › Authentication type – none, simple, or MD5
- › Interface cost – metric for OSPF transmission
- › Transmit Delay – time between transmissions from the specified interface
- › Interface State – current state of the specified interface
- › Priority – router's priority on the specified interface
- › Designated Router – designated router ID and respective interface IP address
- › Backup Designated Router – designated router ID and respective interface IP address of the backup router
- › Interface timer intervals – configuration of timer intervals: Hello, Dead, Wait, and Retransmit
- › Neighbors – number of neighbors and their state; adjacent neighbors
- Example

```

host1#show ip ospf interface
FastEthernet0 is up, OSPF line protocol is up
OSPF interface configuration:
  Internet Address 192.168.1.250, Area 0.0.0.0
  Network type BROADCAST, No authentication
  Cost: 1
  Transmit Delay is 1 sec, Interface State DROTHER, Priority
  1
  Designated Router (Interface address) 192.168.1.107
  Backup Designated Router (Interface address) 192.168.1.214
  Timer intervals configured, Hello 10, Dead 40, Wait 120,
  Retransmit 5
  Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 192.168.1.107 (Designated Router)
  Adjacent with neighbor 192.168.254.7 (Backup Designated
  Router)

```

show ip ospf internal-statistics

- Use to display internal OSPF statistics, such as allocation failures for different OSPF components.
- Field descriptions
 - › LSA bytes allocated – number of bytes allocated for LSAs
 - › Router LSA bytes allocated – number of bytes allocated for router LSAs
 - › Summary bytes allocated – number of bytes allocated for summary LSAs
 - › Neighbor RTX bytes allocated – number of bytes allocated for neighbor retransmissions
 - › Timers bytes allocated – number of bytes allocated for OSPF timers

- › Ospf total bytes free – total number of bytes free
 - › Ospf heap total bytes allocated – total number of bytes allocated from the OSPF heap
 - › Neighbor allocation failures – number of neighbor allocation failures
 - › LSA allocation failures – number of LSA allocation failures
 - › LSA HDR allocation failures – number of LSA header allocation failures
 - › DB Request allocation failures – number of database request allocation failures
 - › RTX allocation failures – number of neighbor retransmission allocation failures
 - › LS Ack allocation failures – number of LSA acknowledgment packet allocation failures
 - › DD pkt allocation failures – number of database description packet allocation failures
 - › OSPF interface allocation failures – number of interface allocation failures
 - › OSPF general packet allocation failures – number of general packet allocation failures
- Example

```
host1#show ip ospf internal-statistics
Routing Process OSPF 1 with Router ID 5.72.3.1
Internal OSPF Statistics, bytes allocated/free:
  LSA bytes allocated:216
  Router LSA bytes allocated:936
  Summary bytes allocated:0
  Neighbor RTX bytes allocated:0
  Timers bytes allocated:352
  Ospf total bytes free:824368
  Ospf heap total bytes allocated:1048576
Internal OSPF Statistics, allocation failures:
  Neighbor allocation failures:0
  LSA allocation failures:0
  LSA HDR allocation failures:0
  DB Request allocation failures:0
  RTX allocation failures:0
  LS Ack allocation failures:0
  DD pkt allocation failures:0
  OSPF interface allocation failures:0
  OSPF general packet allocation failures:0
```

show ip ospf neighbors

- Use to display information on OSPF neighbors on a per-interface basis.
- Field descriptions
 - › Neighbor ID – neighbor's router ID
 - › Pri – router priority of neighbor
 - › State – OSPF neighbor's state

- DR – designated router
- BDR – backup designated router
- DR Other – neighbor to a designated router or a backup designated router
- › Dead Time – interval since last hello packet from neighbor
- › Address – IP address of the neighbor's interface
- › Interface – name of the specified interface and its port number; for example, FastEthernet0
- Example

```
host1#show ip ospf neighbors
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.8.1	1	TWO-WAY/DR Other	00:00:39	10.0.76.1	fastEthernet11/0
10.0.71.1	1	FULL/DR	00:00:42	10.0.76.2	fastEthernet11/0
10.0.96.1	1	FULL/BDR	00:00:28	10.0.76.4	fastEthernet11/0

show ip ospf remote-neighbor interface

- Use to display all interfaces that are associated with OSPF remote neighbors.
- Example

```
host1#show ip ospf remote-neighbor interface
OSPF remote-neighbor 221.221.221.221 interface
configuration:
Update-source loopback0
Remote-neighbor reachable: yes
Area 0.0.0.0
Network type POINT-TO-POINT, No authentication
Cost: 1
Transmit Delay is 1 sec, Interface State POINT-TO-POINT,
Priority 1
No designated router on this network
No backup designated router on this network
Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 221.221.221.221
```

show ip ospf spf-log

- Use to display how often and why the router has run a full SPF calculation.
- Field descriptions
 - › Intra SPF log – log for SPF calculations run to compute intra-area LSAs
 - › Inter SPF log – log for SPF calculations run to compute inter-area LSAs
 - › External SPF log – log for SPF calculations run to compute routes outside the OSPF routing domain
 - › When – amount of time since a full SPF calculation took place given in hours:minutes:seconds; the previous 20 calculations are logged

- › Duration – number of milliseconds to complete this SPF run; the elapsed time is in actual clock time, not CPU time
- › LSA Router Id – whenever a full SPF calculation is triggered by a new LSA, the router ID is stored in the router
- › Triggers – list of reasons that triggered a full SPF calculation
- Example

```
host1#show ip ospf spf-log
```

Intra SPF log

When	Duration	LSA Router Id	Triggers
00:04:42	0.000	23.23.23.3	Protocol Off
00:04:38	0.000	23.23.23.3	LSA Add
00:04:34	0.000	12.12.12.2	LSA Add
00:04:30	0.010	23.23.23.3	LSA Update
00:03:51	0.000	23.23.23.3	Protocol Off
00:03:47	0.000	23.23.23.3	LSA Add
00:03:43	0.000	12.12.12.2	LSA Add
00:03:39	0.000	23.23.23.3	LSA Update

Inter SPF log

When	Duration	LSA Router Id	Triggers
00:04:46	0.010	23.23.23.3	Protocol Off
00:04:42	0.000	23.23.23.3	LSA Add
00:04:38	0.000	12.12.12.2	LSA Add
00:04:34	0.000	23.23.23.3	LSA Update
00:03:55	0.000	23.23.23.3	Protocol Off
00:03:51	0.000	23.23.23.3	LSA Add
00:03:47	0.000	12.12.12.2	LSA Add
00:03:43	0.000	23.23.23.3	LSA Update

External SPF log

When	Duration	LSA Router Id	Triggers
00:04:47	0.000	23.23.23.3	Protocol Off
00:04:43	0.000	23.23.23.3	LSA Add
00:04:39	0.000	12.12.12.2	LSA Add
00:04:35	0.010	23.23.23.3	LSA Update
00:03:56	0.000	23.23.23.3	Protocol Off
00:03:52	0.000	23.23.23.3	LSA Add
00:03:48	0.000	12.12.12.2	LSA Add
00:03:44	0.000	23.23.23.3	LSA Update

show ip ospf virtual-links

- Use to display the parameters and the current state of OSPF virtual links.
- Field descriptions

- › Virtual link to router – identifies the OSPF neighbor and the current state of the virtual link
 - › Transmit Delay – time between transmissions from the specified interface
 - › Timer intervals – timer intervals configured for the link: Hello, Dead, and Retransmit
- Example

```
host1#show ip ospf virtual-links
Virtual link to router 192.168.1.13 in state POINT-TO-POINT
Transmit Delay is 1 sec
Timer intervals configured, Hello 10 sec, Dead 40 sec,
Retransmit 5 sec
```