



Chapter 8: OSPF

Instructor Materials

CCNP Enterprise: Core Networking



Chapter 8 Content

This chapter covers the following content:

- **OSPF Fundamentals** - This section provides an overview of communication between OSPF routers.
- **OSPF Configuration** - This section describes the OSPF configuration techniques and commands that can be executed to verify the exchange of routes.
- **Default Route Advertisement** - This section explains how default routes are advertised in OSPF.
- **Common OSPF Optimizations** - This section reviews common OSPF settings for optimizing the operation of the protocol.

OSPF Fundamentals

- The Open Shortest Path First (OSPF) protocol is a nonproprietary Interior Gateway Protocol (IGP) that overcomes the deficiencies of other distance vector routing protocols and distributes routing information within a single OSPF routing domain.
- OSPF introduced variable-length subnet masking (VLSM), which supports classless routing, summarization, authentication, and external route tagging.
- There are two main versions of OSPF in production networks today: OSPF Version 2 (OSPFv2) which supports IPv4, OSPF Version 3 (OSPFv3) which supports IPv6.

OSPF Fundamentals

LSAs, LSDB, and SPT

- OSPF sends link-state advertisements (LSAs) to neighboring routers. LSAs contain the link state and link metric. The received LSAs are stored in a local database called the link-state database (LSDB). The LSDB provides the topology of the network. The SPT contains all network destinations within the OSPF domain.
- Figure 8-1 shows a simple OSPF topology and the SPT from R1's and R4's perspective. Notice that the local router's perspective will always be the root (top of the tree).

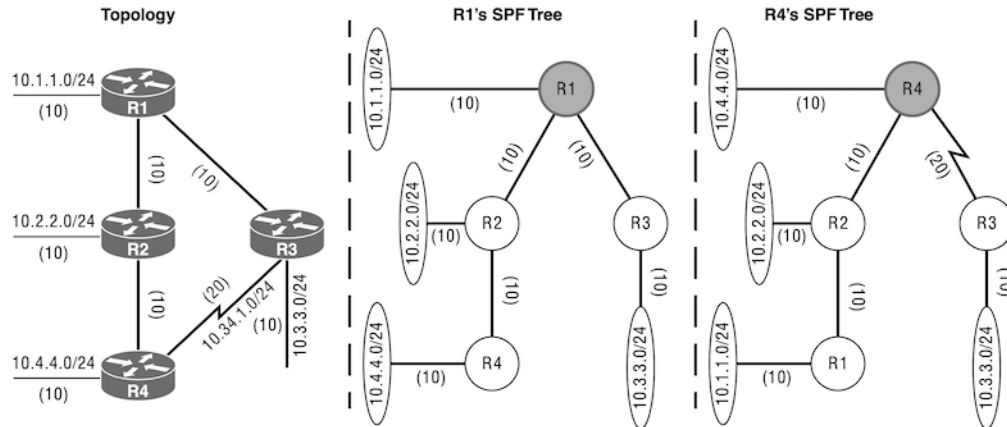


Figure 8-1 OSPF Shortest Path First (SPF) Tree

OSPF Fundamentals

OSPF Architecture

- OSPF uses multiple OSPF areas within the routing domain. OSPF uses a two-tier hierarchical architecture, where Area 0 is a special area known as the backbone, to which all other areas must connect. Nonbackbone areas advertise routes into the backbone. The backbone advertises routes into other nonbackbone areas.
- Figure 8-2 shows route advertisement into other areas. Area 12 routes are advertised to Area 0 and then into Area 34. Area 34 routes are advertised to Area 0 and then into Area 12. Area 0 routes are advertised into all other OSPF areas.

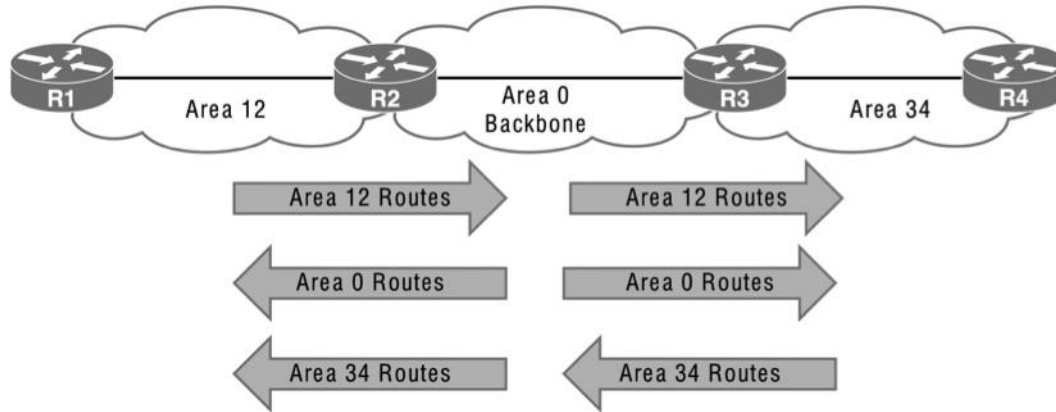


Figure 8-2 *Two-Tier Hierarchical Area Structure*

OSPF Fundamentals

Inter-Router Communication

OSPF uses the assigned IPv4 protocol 89 and multicast addresses 224.0.0.5 (All routers) and 224.0.0.6 (DR routers) where possible to reduce unnecessary traffic.

Table 8-2 briefly describes the five OSPF packet types.

Type	Packet Name	Functional Overview
1	Hello	Discover and maintain neighbors. Packets are sent periodically on all OSPF interfaces to discover neighbors and ensure that other adjacent neighbors are still online.
2	Database description (DBD) or (DDP)	Summarize database contents. Packets are exchanged when an OSPF adjacency is formed. They describe the LSDB contents.
3	Link-state request (LSR)	Download databases. If a router thinks that part of its LSDB is stale, it requests part of a neighbor's DB using this packet type.
4	Link-state update (LSU)	Update databases. This is an explicit LSA for a specific network link and normally is sent in direct response to an LSR.
5	Link-state ack	Flood acknowledgments. These are sent in response to LSA flooding, making flooding a reliable transport feature.

OSPF Fundamentals

OSPF Hello Packets

Table 8-3 lists some data found within an OSPF hello packet.

Data Field	Description
Router ID (RID)	A unique 32-bit ID within an OSPF domain.
Authentication options	Between OSPF routers: none, clear text, or MD5 authentication.
Area ID	An interface's OSPF area. A 32-bit number written in dotted-decimal format (0.0.1.0) or decimal (256).
Interface address mask	The interface's primary IP address network mask.
Interface priority	The router interface priority for DR elections.
Hello interval	The time span, in seconds, that a router sends out hello packets.
Dead interval	The time span, in seconds, that a router waits to hear a hello from a neighbor router before it declares that router down.
DR and BDR	IP address of the DR and backup DR (BDR) for the network link.
Active neighbor	A list of OSPF neighbors on the network segment. A router must have received a hello from the neighbor within the dead interval.

Table 8-4 OSPF neighbor states

State	Description
Down	The initial state of a neighbor relationship. Indicates that the router has not received any OSPF hello packets.
Attempt	Indicates that no information has been received recently, but the router is still attempting communication.
Init	Indicates that a hello packet has been received from another router, but bidirectional communication has not been established.
2-Way	Bidirectional communication established. If a DR or BDR is needed, the election occurs during this state.
ExStart	The first state in forming an adjacency. Routers identify which router will be the master or slave for the LSDB synchronization.
Exchange	Routers are exchanging link states by using DBD packets.
Loading	LSRs sent to neighbor asking for more recent LSAs that have been discovered (but not received) in the Exchange state.
Full	Neighboring routers are fully adjacent.

OSPF Fundamentals

DR and BDR

- If four routers share the same multi-access network, six OSPF adjacencies form, along with six occurrences of database flooding on a network. Figure 8-3 shows a simple four-router physical topology and the adjacencies established.
- Figure 8-5 shows how a DR simplifies a four-router topology with only three neighbor adjacencies.

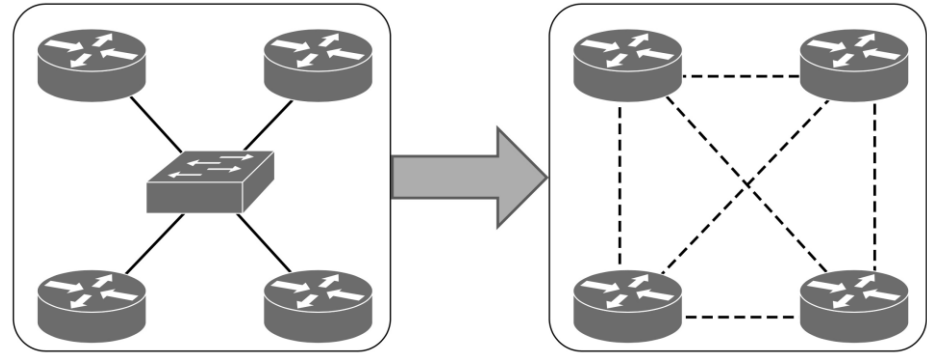


Figure 8-3 Multi-Access Physical Topology Versus Logical Topology

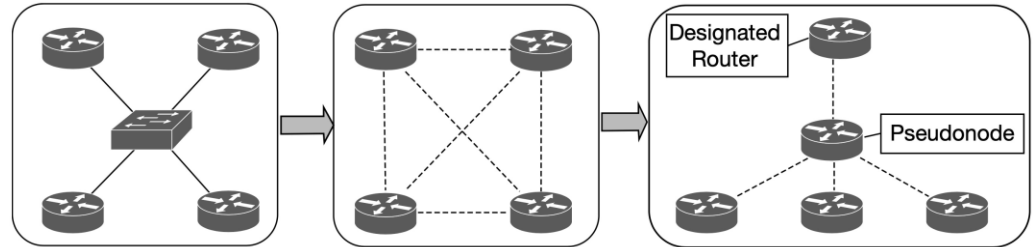


Figure 8-5 OSPF DR Concept

OSPF Configuration

The command **router ospf** *process-id* defines and initializes the OSPF process. OSPF is enabled on an interface using two methods:

- An OSPF network statement
- Interface-specific configuration

OSPF Configuration

OSPF Network Statement

- The OSPF network statement identifies the interfaces that the OSPF process will use and the area that those interfaces participate in. The network statements match against the primary IPv4 address and netmask associated with an interface.
- The selection of interfaces within the OSPF process is accomplished by using the command **network** *ip-address wildcard-mask* **area** *area-id*. This is similar to configuring EIGRP, except that the OSPF area is specified. Example 8-2 provides one method.
- The connected network for the OSPF-enabled interface is added to the OSPF LSDB under the corresponding OSPF area in which the interface participates.

Example 8-2 *Configuring OSPF with Explicit Subnet*

```
router ospf 1
  network 10.0.0.0 0.0.0.255 area 0
  network 10.0.10.0 0.0.0.255 area 0
  network 192.0.0.0 0.0.0.255 area 0
  network 192.10.0.0 0.0.0.255 area 0
```

OSPF Configuration

Interface-Specific Configuration

- The second method for enabling OSPF on an interface for IOS is to configure it specifically on an interface with the command **ip ospf process-id area area-id**. This configuration is not centralized and increases in complexity as the number of interfaces on the routers increases. If a hybrid configuration exists on a router, interface-specific settings take precedence over the network statement with the assignment of the areas.
- Example 8-5 provides a sample interface-specific configuration.

Example 8-5 *Configuring OSPF on IOS for a Specific Interface*

```
interface GigabitEthernet 0/0
  ip address 10.0.0.1 255.255.255.0
  ip ospf 1 area 0
```

Statically Set the RID and Passive Interfaces

- The OSPF topology is built on the RID. Setting a static RID helps with troubleshooting and reduces LSAs when a RID changes in an OSPF environment. The command **router-id** *router-id* statically assigns the OSPF RID under the OSPF process. The command **clear ip ospf process** restarts the OSPF process on a router so that OSPF can use the new RID.
- Making a network interface passive still adds the network segment into the LSDB but prohibits the interface from forming OSPF adjacencies. A passive interface does not send out OSPF hellos and does not process any received OSPF packets. The command **passive** *interface-id* under the OSPF process makes the interface passive, and the command **passive interface default** makes all interfaces passive. To allow for an interface to process OSPF packets, the command **no passive** *interface-id* is used.

OSPF Configuration

Requirements for Neighbor Adjacency

The following list of requirements must be met for an OSPF neighborship to be formed:

- RIDs must be unique between the two devices.
- The interfaces must share a common subnet.
- The MTUs on the interfaces must match.
- The area ID must match for the segment.
- The DR enablement must match for the segment.
- OSPF hello and dead timers must match for the segment.
- Authentication type and credentials (if any) must match for the segment.
- Area type flags must match for the segment (for example, Stub, NSSA).

OSPF Configuration

Sample Topology and Interface Confirmation

- Figure 8-7 shows a topology of a basic OSPF configuration. All routers have loopback IP addresses matching their RIDs. On R1 and R2, OSPF is enabled on all interfaces, R3 uses specific network-based statements, R4 uses interface-specific commands. R1 and R2 set Gi0/2 interface as passive, and R3 and R4 make all interfaces passive by default but make Gi0/1 active.
- Verify that the correct interfaces are running OSPF after making changes to the OSPF configuration. The command **show ip ospf interface [brief | interface-id]** displays the OSPF-enabled interfaces.

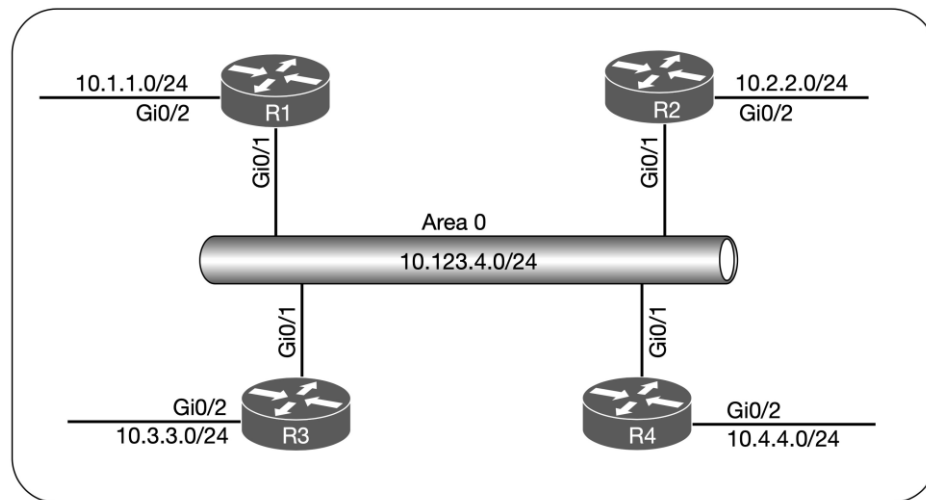


Figure 8-7 Sample OSPF Topology

OSPF Configuration

OSPF Interface Columns

Table 8-6 OSPF Interface Columns displayed with the **show ip ospf interface brief** command

Field	Description
Interface	Interfaces with OSPF enabled
PID	The OSPF process ID associated with this interface
Area	The area that this interface is associated with
IP Address/Mask	The IP address and subnet mask for the interface
Cost	The cost metric assigned to an interface that is used to calculate a path metric
State	The current interface state, which could be DR, BDR, DROTHER, LOOP, or Down
Nbrs F	The number of neighbor OSPF routers for a segment that are fully adjacent
Nbrs C	The number of detected neighbor OSPF routers for a segment in a 2-Way state

OSPF Configuration

Verification of OSPF Neighbor Adjacencies

Table 8-7 OSPF Neighbor State Fields displayed with the **show ip ospf neighbor** command

Field	Description
Neighbor ID	The router ID (RID) of the neighboring router.
PRI	The priority for the neighbor's interface, which is used for DR/BDR elections.
State	The first field is the neighbor state. The second field is the DR, BDR, or DROTHER role if the interface requires a DR.
Dead Time	The time left until the router is declared unreachable.
Address	The primary IP address for the OSPF neighbor.
Interface	The local interface to which the OSPF neighbor is attached.

OSPF Configuration

Verification of OSPF Routes

Verify OSPF routes that install into the RIB with the command **show ip route ospf**.

Example 8-10 provides sample output of the OSPF routing table for R1. In the output, two sets of numbers are in the brackets (for example, [110/2]), the first number is the administrative distance (AD), which is 110 by default for OSPF, and the second number is the metric of the path used for that network.

Example 8-10 OSPF Routes Installed in the RIB

```
R1# show ip route ospf
! Output omitted for brevity
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O        10.2.2.0/24 [110/2] via 10.123.4.2, 00:35:03, GigabitEthernet0/1
O        10.3.3.0/24 [110/2] via 10.123.4.3, 00:35:03, GigabitEthernet0/1
O        10.4.4.0/24 [110/2] via 10.123.4.4, 00:35:03, GigabitEthernet0/1
    192.168.2.0/32 is subnetted, 1 subnets
O        192.168.2.2 [110/2] via 10.123.4.2, 00:35:03, GigabitEthernet0/1
    192.168.3.0/32 is subnetted, 1 subnets
O        192.168.3.3 [110/2] via 10.123.4.3, 00:35:03, GigabitEthernet0/1
    192.168.4.0/32 is subnetted, 1 subnets
O        192.168.4.4 [110/2] via 10.123.4.4, 00:35:03, GigabitEthernet0/1
```

Default Route Advertisement

- OSPF supports advertising the default route into the OSPF domain. The default route is advertised by using the command **default-information originate** [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] underneath the OSPF process.

Default Route Topology & Configuration

The **always** optional keyword advertises a default route even if a default route does not exist in the RIB. The route metric can be changed with the **metric** *metric-value* option. The metric type can be changed with the **metric-type** *type-value* option.

Figure 8-8 illustrates a scenario for providing connectivity to other parts of the network by having R1 advertise a default route into OSPF as shown in Example 8-11.

Example 8-11 OSPF Default Information Origination Configuration

```
R1
ip route 0.0.0.0 0.0.0.0 100.64.1.2
!
router ospf 1
 network 10.0.0.0 0.255.255.255 area 0
 default-information originate
```

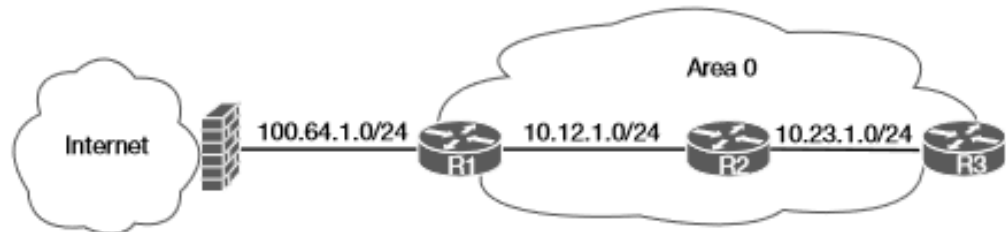


Figure 8-8 Default Route Topology

Common OSPF Optimizations

- Almost every network requires tuning based on the equipment, technical requirements, or a variety of other factors. This section explains common concepts involved with the tuning of an OSPF network.

Common OSPF Optimizations

Link Costs

OSPF assigns the OSPF link cost for an interface using the formula in Figure 8-9. The default reference bandwidth is 100 Mbps. Table 8-8 provides the OSPF cost for common network interface types using the default reference bandwidth. The command **auto-cost reference-bandwidth** *bandwidth-in-mbps* changes the reference bandwidth for all OSPF interfaces associated with that process. The OSPF cost can be set manually with the command **ip ospf cost** 1–65535 underneath the interface.

If the value is not changed a FastEthernet interface would have the same cost as a 10 Gb interface resulting in poor routing decisions.

$$\text{Cost} = \frac{\text{Reference Bandwidth}}{\text{Interface Bandwidth}}$$

Figure 8-9 OSPF Interface Cost Formula

Table 8-8 OSPF Interface Costs Using Default Settings

Interface Type	OSPF Cost
T1	64
Ethernet	10
FastEthernet	1
GigabitEthernet	1
10 GigabitEthernet	1

Common OSPF Optimizations

Failure Detection

- OSPF sends hello packets at set intervals based on the hello timer. OSPF uses a dead interval timer, which is four times the hello timer. If a router does not receive a hello before the OSPF dead interval timer reaches 0, the neighbor state is changed to down. The OSPF router sends out an LSA, with the topology change, and the SPF algorithm processes all routers in the area.
- OSPF allows modification to the hello timer interval with values between 1 and 65,535 seconds. The OSPF hello timer is modified with the interface configuration sub mode command **ip ospf hello-interval 1–65535**.
- The OSPF dead interval timer can be changed with the command **ip ospf dead-interval 1–65535** under the interface configuration sub mode. The dead interval timer setting must be greater than the hello timer setting to ensure that the dead interval timer does not reach 0 in between hello packets.
- The timers for an OSPF interfaces are shown with the command **show ip ospf interface**. The timers need to match on the neighboring interface.

Common OSPF Optimizations

DR and BDR Elections

- Any router with OSPF priority of 1 to 255 on its OSPF interface attempts to become the DR. By default, all OSPF interfaces use a priority of 1. The routers place their RID and OSPF priorities in their OSPF hellos for that segment.
- Routers then receive and examine OSPF hellos from neighboring routers. If the hello received is more favorable, the router updates its OSPF hello packet to use the more preferable RID in the DR field. OSPF deems a router more preferable if the priority for the interface is the highest for that segment. If the OSPF priority is the same, the higher RID is more favorable.
- After all routers agree on the same DR, all routers for that segment become adjacent with the DR. Then the election for the BDR takes place. The election follows the same logic for the DR election, except that the DR does not add its RID to the BDR field of the hello packet.
- The OSPF DR and BDR roles cannot be preempted after the DR/BDR election except for failure (or process restart of the DR or BDR).
- Determine the interface role by viewing the OSPF interface with the command **show ip ospf interface brief**.

Common OSPF Optimizations

DR and BDR Placement

- To change DR placement, modify the interface priority to a higher value than the existing DR has. The priority can be set manually under the interface configuration with the command **ip ospf priority 0–255** for IOS nodes. Setting an interface priority to 0 removes that interface from the DR/BDR election immediately. Raising the priority above the default value (1) makes that interface more favorable compared to interfaces with the default value.
- Figure 8-10 provides a topology example to illustrate modification of DR/BDR placement in a network segment.

To accomplish the placement in the topology, the OSPF priority for R1 is set to 100, R2 and R3 are left at the default priority of 1, and R4 is set to 0 so that it will never become DR or BDR.

```
R1(config)# interface G0/1
R1(config-if)# ip ospf priority 100
R4(config)# interface G0/1
R4(config-if)# ip ospf priority 0
```

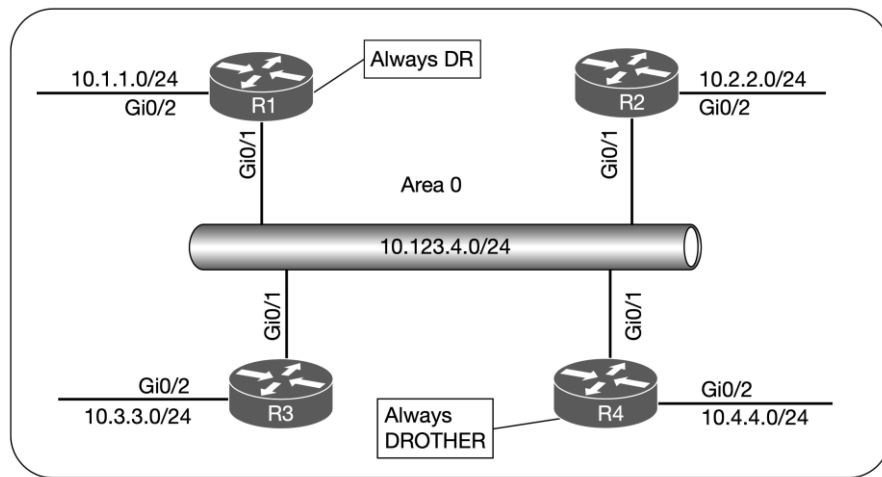


Figure 8-10 OSPF Topology for DR/BDR Placement

Common OSPF Optimizations

OSPF Network Types

The default OSPF network type is based on media used for the connection. Can be changed independently of media used. Cisco provides five OSPF network types, as listed in Table 8-9.

Type	Description	DR/BDR Field in OSPF Hellos	Timers
Broadcast	Default setting on OSPF-enabled Ethernet links	Yes	Hello: 10, Wait: 40, Dead: 40
Non-broadcast	Default setting on OSPF-enabled Frame Relay main interface or Frame Relay multipoint subinterfaces	Yes	Hello: 30, Wait: 120, Dead: 120
Point-to-point	Default setting on OSPF-enabled Frame Relay point-to-point subinterfaces.	No	Hello: 10, Wait: 40, Dead: 40
Point-to-multipoint	Not enabled by default on any interface type. Interface advertised as a host route. Sets the next-hop address to the outbound interface.	No	Hello: 30, Wait: 120, Dead: 120
Loopback	Default setting on OSPF-enabled loopback interfaces. Interface is advertised as a host route (/32).	N/A	N/A

Broadcast, Point-to-Point and Loopback Networks

- **Broadcast:** Broadcast networks are multiaccess in that they are capable of connecting more than two devices. A DR is required for OSPF broadcast networks because of the possibility that multiple nodes can exist on a segment, and LSA flooding needs to be controlled. The interface parameter command **ip ospf network broadcast** overrides the automatically configured setting and statically sets an interface as an OSPF broadcast network type.
- **P2P:** A network circuit that allows only two devices to communicate is considered a point-to-point (P2P) network. Only two nodes can exist on this type of network so no DR is required. The hello timer is set to 10 seconds and there is no wait timer. The interface parameter command **ip ospf network point-to-point** sets an interface as an OSPF point-to-point network type.
- **Loopback:** The OSPF network type loopback is enabled by default for loopback interfaces and can be used only on loopback interfaces. The OSPF loopback network type states that the IP address is always advertised with a /32 prefix length, even if the IP address configured on the loopback interface does not have a /32 prefix length.

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 8

Description
OSPF backbone
Inter-router communication
OSPF Packet Types
OSPF Neighbor States
Designated router
OSPF network statement
Interface specific enablement
Passive interfaces

Key Topics for Chapter 8 (Cont.)

Description
Requirements for neighbor adjacency
OSPF Interface Columns
OSPF Neighbor State Fields
Default route advertisement
Link costs
Failure detection
Designated router elections
OSPF Network Types

Prepare for the Exam

Key Terms for Chapter 8

Key Terms

backup designated router (BDR)

dead interval

designated router (DR)

hello interval

hello packets

interface priority

passive interface

router ID (RID)

shortest path tree (SPT)

Prepare for the Exam

Command Reference for Chapter 8

Task	Command Syntax
Initialize the OSPF process	router ospf <i>process-id</i>
Enable OSPF on network interfaces matching a specified network range for a specific OSPF area	network <i>ip-address wildcard-mask</i> area <i>area-id</i>
Enable OSPF on an explicit specific network interface for a specific OSPF area	ip ospf <i>process-id</i> area <i>area-id</i>
Configure a specific interface as passive	passive <i>interface-id</i>
Configure all interfaces as passive	passive interface default
Advertise a default route into OSPF	default-information originate [always] [metric <i>metric-value</i>] [metric-type <i>type-value</i>]

Command Reference for Chapter 8 (Cont.)

Task	Command Syntax
Modify the OSPF reference bandwidth for dynamic interface metric costing	auto-cost reference-bandwidth <i>bandwidth-in-mbps</i>
Statically set the OSPF metric for an interface	ip ospf cost 1–65535
Configure the OSPF priority for a DR/BDR election	ip ospf priority 0–255
Statically configure an interface as a broadcast OSPF network type	ip ospf network broadcast
Statically configure an interface as a point-to-point OSPF network type	ip ospf network point-to-point
Restart the OSPF process	clear ip ospf process

Command Reference for Chapter 8 (Cont.)

Task	Command Syntax
Display the OSPF interfaces on a router	show ip ospf interface [brief <i>interface-id</i>]
Display the OSPF neighbors and their current states	show ip ospf neighbor [detail]
Display the OSPF routes that are installed in the RIB	show ip route ospf

