

# Chapter 6: OSPF

## Instructor Materials

CCNP Enterprise: Advanced Routing



# Chapter 6 Content

## This chapter covers the following content:

- **OSPF Fundamentals** - This section provides an overview of the OSPF routing protocol.
- **OSPF Configuration** - This section explains how to configure a router with basic OSPF functionality.
- **The Designated Router and Backup Designated Router** - This section describes the function of the designated router and how it provides scalability for broadcast network segments.
- **OSPF Network Types** - This section provides an overview of the OSPF network types and their impact to OSPF's behavior.
- **Failure Detection** - This section explains how OSPF detects and verifies the health of OSPF neighbor routers.
- **Authentication** - This section explains how OSPF authentication functions and is configured.

# OSPF Fundamentals

- OSPF is a nonproprietary Interior Gateway Protocol (IGP) that overcomes the deficiencies of distance vector routing protocols and distributes routing information within a single OSPF routing domain.
- There are two main versions of OSPF in production networks today:
  - OSPFv2 - Originally defined in RFC 2328 with IPv4 support
  - OSPFv3 - Modifies the original structure to support IPv6
- OSPF advertises link-state advertisements (LSAs) that contain the link state and link metric to neighboring routers.

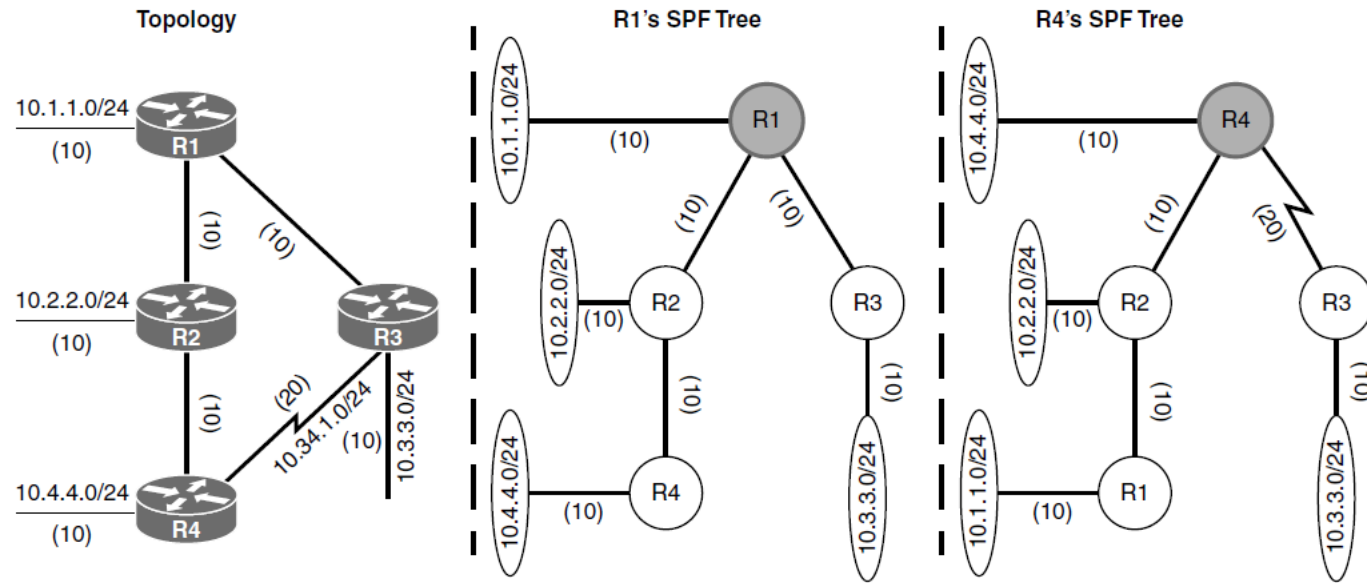
# Foundation Topics

- OSPF advertises link-state advertisements (LSAs) that contain the link state and link metric to neighboring routers. Received LSAs are stored in a local database called the link-state database (LSDB) and advertise the link-state information to neighboring routers exactly as the original advertising router advertised it. All OSPF routers maintain a synchronized identical copy of the LSDB within an area.
- The LSDB provides the topology of the network, in essence providing the router a complete map of the network. All OSPF routers run Dijkstra's shortest path first (SPF) algorithm to construct a loop-free topology of shortest paths.
- Each router sees itself as the root or top of the SPF tree (SPT), and the SPT contains all network destinations within the OSPF domain.
- A router can run multiple OSPF processes. Each process maintains its own unique database, and routes learned in one OSPF process are not available to a different OSPF process without redistribution of routes between processes. The OSPF process numbers are locally significant and do not have to match among routers.

# OSPF Fundamentals

## Foundation Topics (Cont.)

Figure 6-1 demonstrates a simple OSPF topology and the SPT from R1's and R4's perspective. Notice that the local router's perspective is always that of the root (or top of the tree).



**Figure 6-1** OSPF Shortest Path First (SPF) Tree

## OSPF Fundamentals

# Areas

OSPF provides scalability for the routing table by splitting segments of the topology into multiple OSPF areas within the routing domain. Area membership is set at the interface level, and the area ID is included in the OSPF hello packet. An interface can belong to only one area. All routers within the same OSPF area maintain an identical copy of the LSDB.

An OSPF area grows in size as the number of network links and number of routers increase in the area. While using a single area simplifies the topology, there are trade-offs:

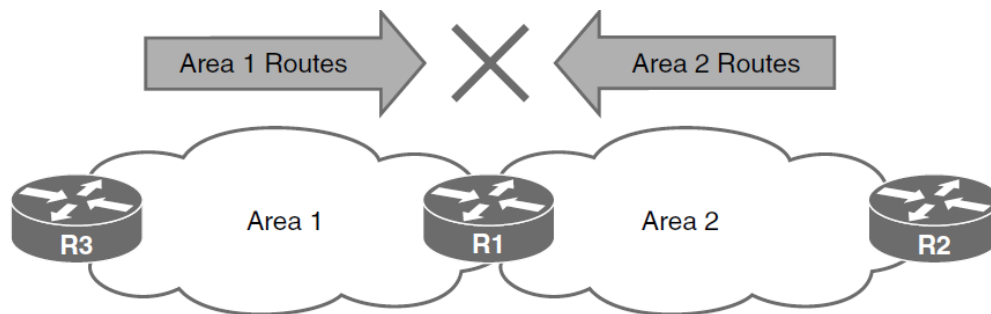
- A full SPT calculation runs when a link flaps within the area.
- With a single area, the LSDB increases in size and becomes unmanageable.
- The LSDB for the single area grows, consumes more memory, and takes longer during the SPF computation process.
- With a single area, no summarization of route information occurs.

Proper design addresses each of these issues by segmenting the routers into multiple OSPF areas, thereby keeping the LSDB to a manageable size.

## OSPF Fundamentals

# Areas (Cont.)

- If a router has interfaces in multiple areas, the router has multiple LSDBs (one for each area). The internal topology of one area is invisible from outside that area.
- Segmenting the OSPF domain into multiple areas reduces the size of the LSDB for each area, making SPT calculations faster and decreasing LSDB flooding between routers when a link flaps.
- Just because a router connects to multiple OSPF areas does not mean the routes from one area will be injected into another area. Figure 6-2 shows router R1 connected to Area 1 and Area 2. Routes from Area 1 do not advertise into Area 2 and vice versa.



**Figure 6-2** *Failed Route Advertisement Between Areas*

## OSPF Fundamentals

# Areas (Cont.)

- Area 0 is a special area called the backbone. By design, OSPF uses a two-tier hierarchy in which all areas must connect to the upper tier, Area 0, because OSPF expects all areas to inject routing information into Area 0. Area 0 advertises the routes into other non-backbone areas. The backbone design is crucial to preventing routing loops.
- The area identifier (also known as the area ID) is a 32-bit field and can be formatted in simple decimal (0 through 4294967295) or dotted decimal (0.0.0.0 through 255.255.255.255). When configuring routers in an area, if you use decimal format on one router and dotted-decimal format on a different router, the routers will be able to form an adjacency. OSPF advertises the area ID in the OSPF packets.
- Area border routers (ABRs) are OSPF routers connected to Area 0 and another OSPF area, per Cisco definition and according to RFC 3509. ABRs are responsible for advertising routes from one area and injecting them into a different OSPF area. Every ABR needs to participate in Area 0 to allow for the advertisement of routes into another area. ABRs compute an SPT for every area in which they participate.

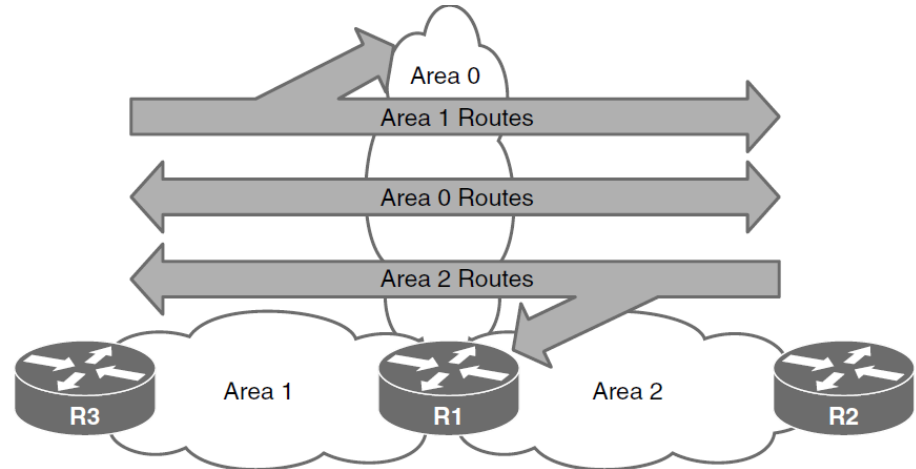
# OSPF Fundamentals

## Areas (Cont.)

Figure 6-3 shows that R1 is connected to Area 0, Area 1, and Area 2. R1 is a proper ABR router because it participates in Area 0. The following occurs on R1:

- Routes from Area 1 advertise into Area 0.
- Routes from Area 2 advertise into Area 0.
- Routes from Area 0 advertise into Areas 1 and 2. This includes the local Area 0 routes, in addition to the routes that were advertised into Area 0 from Area 1 and Area 2.

The topology in Figure 6-3 is a larger-scale OSPF multi-area topology that is used throughout this chapter to describe various OSPF concepts.



**Figure 6-3** *Successful Route Advertisement Between Areas*

# Inter-Router Communication

OSPF runs directly over IPv4, using its own protocol 89, which is reserved for OSPF by the Internet Assigned Numbers Authority (IANA). OSPF uses multicast where possible to reduce unnecessary traffic. There are two OSPF multicast addresses:

- **AllSPFRouters** - IPv4 address 224.0.0.5 or MAC address 01:00:5E:00:00:05. All routers running OSPF should be able to receive these packets.
- **AllIDRouters** - IPv4 address 224.0.0.6 or MAC address 01:00:5E:00:00:06. Communication with designated routers (DRs) uses this address.

Within the OSPF protocol, five types of packets are communicated. Table 6-2 briefly describes the OSPF packet types.

Table 6-2 OSPF Packet Types

Type	Packet Name	Functional Overview
1	Hello	Packets are sent out periodically on all OSPF interfaces to discover new neighbors while ensuring that other neighbors are still online.
2	Database description (DBD or DDP)	Packets are exchanged when an OSPF adjacency is first being formed. These packets are used to describe the contents of the LSDB.
3	Link-state request (LSR)	When a router thinks that part of its LSDB is stale, it may request a portion of a neighbor's database using this packet type.
4	Link-state update (LSU)	This is an explicit LSA for a specific network link, and normally it is sent in direct response to an LSR.
5	Link-state acknowledgment	These packets are sent in response to the flooding of LSAs, thus making the flooding a reliable transport feature.

## OSPF Fundamentals

# Router ID

The OSPF router ID (RID) is a 32-bit number that uniquely identifies an OSPF router. In some OSPF output commands, *neighbor ID* refers to the RID; the terms are synonymous. The RID characteristics are:

- The RID must be unique for each OSPF process in an OSPF domain and must be unique between OSPF processes on a router.
- The RID is dynamically allocated by default using the highest IP address of any *up* loopback interfaces. If there are no up loopback interfaces, the highest IP address of any active *up* physical interfaces becomes the RID when the OSPF process initializes.
- The OSPF process selects the RID when the OSPF process initializes, and it does not change until the process restarts.
- Setting a static RID helps with troubleshooting and reduces LSAs when an RID changes in an OSPF environment. The RID is four octets in length and is configured with the command **router-id** *router-id* under the OSPF process.

# OSPF Fundamentals

## OSPF Hello Packets

OSPF hello packets are responsible for discovering and maintaining neighbors. In most instances, a router sends hello packets to the AllSPFRouters address (224.0.0.5). Table 6-3 lists some of the data contained within an OSPF hello packet.

**Table 6-3** OSPF Hello Packet Fields

Data Field	Description
Router ID (RID)	A unique 32-bit ID within an OSPF domain.
Authentication Options	A field that allows secure communication between OSPF routers to prevent malicious activity. Options are none, plaintext, or Message Digest 5 (MD5) authentication.
Area ID	The OSPF area that the OSPF interface belongs to. It is a 32-bit number that can be written in dot-decimal format (0.0.1.0) or decimal (256).
Interface Address Mask	The network mask for the primary IP address for the interface out which the hello is sent.
Interface Priority	The router interface priority for DR elections.
Hello Interval	The time interval, in seconds, at which a router sends out hello packets on the interface.
Dead Interval	The time interval, in seconds, that a router waits to hear a hello from a neighbor router before it declares that router down.
Designated Router and Backup Designated Router	The IP address of the DR and backup DR (BDR) for that network link.
Active Neighbor	A list of OSPF neighbors seen on that network segment. A router must have received a hello from the neighbor within the dead interval.

# OSPF Fundamentals

## Neighbors

- An OSPF neighbor is a router that shares a common OSPF-enabled network link.
- OSPF routers discover other neighbors through the OSPF hello packets.
- An adjacent OSPF neighbor is an OSPF neighbor that shares a synchronized OSPF database.
- Each OSPF process maintains a table for adjacent OSPF neighbors and the state of each router.

**Table 6-4** OSPF Neighbor States

State	Description
Down	The initial state of a neighbor relationship. It indicates that the router has not received any OSPF hello packets.
Attempt	A state that is relevant to nonbroadcast multi-access (NBMA) networks that do not support broadcast and that require explicit neighbor configuration. This state indicates that no recent information has been received, but the router is still attempting communication.
Init	A state in which a hello packet has been received from another a router, but bidirectional communication has not been established.
2-Way	A state in which bidirectional communication has been 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	A state during which routers are exchanging link states by using DBD packets.
Loading	A state in which LSR packets are sent to the neighbor, asking for the more recent LSAs that have been discovered (but not received) in the Exchange state.
Full	A state in which neighboring routers are fully adjacent.

Table 6-4 briefly describes the OSPF neighbor states.

# Requirements for Neighbor Adjacency

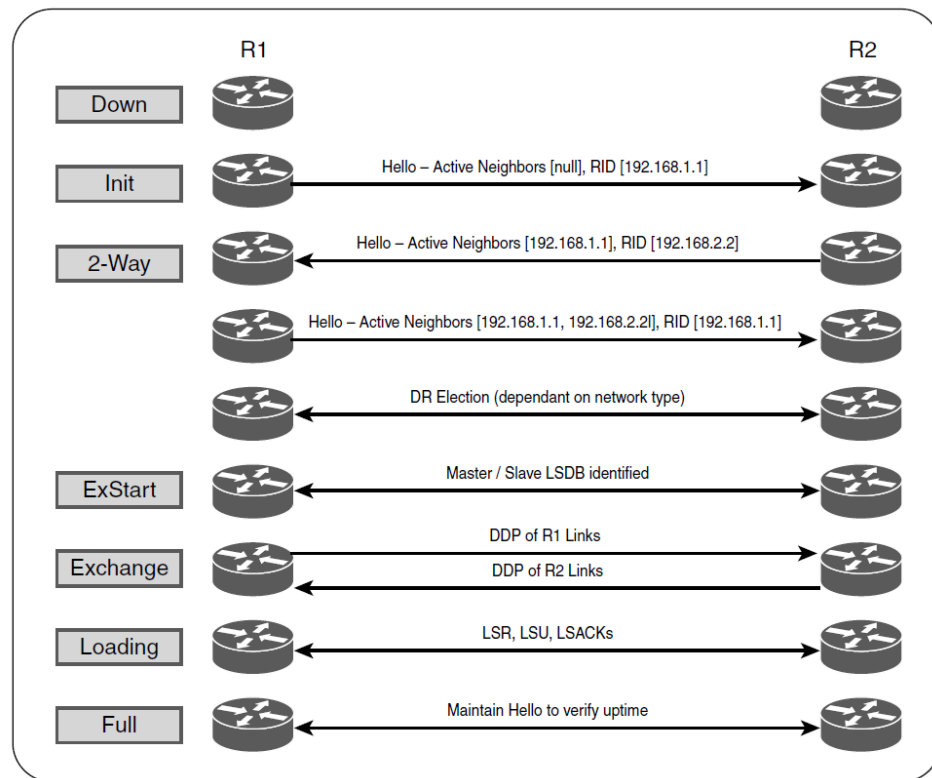
The following list of requirements must be met for an OSPF neighborship to be formed:

- The RIDs must be unique between the two devices. To prevent errors, they should be unique for the entire OSPF routing domain.
- The interfaces must share a common subnet. OSPF uses the interface's primary IP address when sending out OSPF hellos. The network mask (netmask) in the hello packet is used to extract the network ID of the hello packet.
- The interface maximum transmission unit (MTU) must match because the OSPF protocol does not support fragmentation.
- The area ID must match for that segment.
- The need for a DR must match for that segment.
- OSPF hello and dead timers must match for that segment.
- The authentication type and credentials (if any) must match for that segment.
- Area type flags must be identical for that segment (stub, NSSA, and so on).

# Requirements for Neighbor Adjacency (Cont.)

Figure 6-4 illustrates the states and packets exchanged when two routers, R1 and R2, form an OSPF adjacency.

To view the process of forming an adjacency, use the **debug ip ospf adj** command to get detailed information for all of the states.



**Figure 6-4** *Process for Forming OSPF Neighbor Adjacencies*

# OSPF Configuration

- The configuration process for OSPF occurs mostly under the OSPF process, but some OSPF options go directly on the interface configuration submode.
- The OSPF process ID is locally significant but is generally kept the same for operational consistency.
- OSPF is enabled on an interface using two methods:
  - OSPF network statement
  - Interface-specific configuration

# OSPF Configuration

## Enabling OSPF

### OSPF Network Statement

The command **router ospf** *process-id* defines and initializes the OSPF process. The OSPF network statement identifies the interfaces that the OSPF process will use and the area that those interfaces participate in. The network statements select and enable OSPF on the interfaces. The selection of interfaces within the OSPF process is accomplished by using the command **network** *ip-address wildcard-mask area area-id*.

### Interface-Specific Configuration

The second method for enabling OSPF on an interface is to configure it specifically on the interface with the command **ip ospf** *process-id area area-id [secondaries none]*. This method allows for secondary connected networks to be added to the LSDB, unless the secondaries none option is used. While this method provides explicit control for enabling OSPF, it scatters the configuration and increases complexity.

### Passive Interfaces

The command **passive** *interface-id* under the OSPF process makes the interface passive, and the command **passive interface default** makes all interfaces passive.

# OSPF Configuration

## Multi-area OSPF Configuration

Figure 6-5 displays a reference topology for a basic multi-area OSPF configuration. R1, R2, and R3 belong to Area 1234. R4 and R5 belong to Area 0. R5 and R6 belong to Area 56. R4 and R5 are ABRs, the other routers are member (internal) routers.

Example 6-2 provides the OSPF configurations for all six routers.

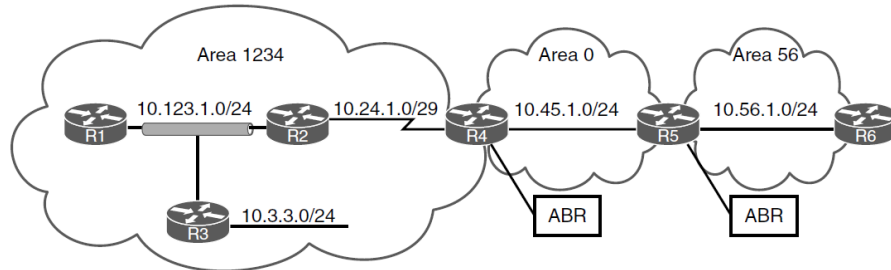


Figure 6-5 Basic Multi-Area OSPF Topology

### Example 6-2 OSPF Configurations for Topology Example

```
R1
router ospf 1
router-id 192.168.1.1
network 0.0.0.0 255.255.255.255 area 1234
```

```
R2
router ospf 1
router-id 192.168.2.2
network 10.123.1.2 0.0.0.0 area 1234
network 10.24.1.2 0.0.0.0 area 1234
```

```
R3
router ospf 1
router-id 192.168.1.1
network 0.0.0.0 255.255.255.255 area 1234
passive interface GigabitEthernet0/1
```

```
R4
router ospf 1
router-id 192.168.4.4
!
interface GigabitEthernet0/0
ip ospf 1 area 0
interface Serial1/0
ip ospf 1 area 1234
```

```
R5
router ospf 1
router-id 192.168.5.5
network 10.45.1.0 0.0.0.255 area 0
network 0.0.0.0 255.255.255.255 area 56
```

```
R6
router ospf 1
router-id 192.168.6.6
network 0.0.0.0 255.255.255.255 area 56
```

# OSPF Configuration

## Confirming OSPF Interfaces

You view OSPF-enabled interfaces by using the command **show ip ospf interface [brief | interface-id]**.

Example 6-3 shows output from using the **show ip ospf interface** command on R4.

The output lists all the OSPF-enabled interfaces, the IP address associated with each interface, the RID for the DR and BDR (and their associated interface IP addresses for that segment), and the OSPF timers for that interface.

**Example 6-3** Detailed OSPF Interface Output

```
R4# show ip ospf interface
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 10.45.1.4/24, Area 0, Attached via Interface Enable
  Process ID 1, Router ID 192.168.4.4, Network Type BROADCAST, Cost: 1
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                1          no            no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 192.168.5.5, Interface address 10.45.1.5
  Backup Designated router (ID) 192.168.4.4, Interface address 10.45.1.4
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  ..
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 192.168.5.5  (Designated Router)
  Suppress hello for 0 neighbor(s)
Serial1/0 is up, line protocol is up
  Internet Address 10.24.1.4/29, Area 1234, Attached via Interface Enable
  Process ID 1, Router ID 192.168.4.4, Network Type POINT_TO_POINT, Cost: 64
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                64          no            no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  ..
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 192.168.2.2
  Suppress hello for 0 neighbor(s)
```

# OSPF Configuration

## Confirming OSPF Interfaces (Cont.)

Example 6-4 shows the command with the **brief** keyword for R1, R2, R3, and R4. Table 6-5 provides an overview of the fields in the output shown in Example 6-4.

**Example 6-4** OSPF Interface Output in Brief Format

R1# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Gi0/0	1	1234	10.123.1.1/24	1	DROTH	2/2	
R2# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se1/0	1	1234	10.24.1.1/29	64	P2P	1/1	
Gi0/0	1	1234	10.123.1.2/24	1	BDR	2/2	
R3# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Gi0/1	1	1234	10.3.3.3/24	1	DR	0/0	
Gi0/0	1	1234	10.123.1.3/24	1	DR	2/2	
R4# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Gi0/0	1	0	10.45.1.4/24	1	BDR	1/1	
Se1/0	1	1234	10.24.1.4/29	64	P2P	1/1	

**Table 6-5** OSPF Interface Columns

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	A factor the SPF algorithm uses to calculate a metric for a path
State	The current interface state for segments with a designated router (DR, BDR, or DROTHER), P2P, LOOP, or Down
Nbrs F	The number of neighbor OSPF routers for a segment that are fully adjacent
Nbrs C	The number of neighbor OSPF routers for a segment that have been detected and are in a 2-WAY state

# OSPF Configuration

## Verification of OSPF Neighbor Adjacencies

The command **show ip ospf neighbor [detail]** provides the OSPF neighbor table. Example 6-5 displays the OSPF neighbors for R1 and R2.

Table 6-6 provides a brief overview of the fields used in Example 6-5. The neighbor state on R1 identifies R3 as the DR and R2 as the BDR for the 10.123.1.0 network segment. R2 identifies R1 as DROTHER for that network segment.

**Example 6-5** OSPF Neighbor Output

R1# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.2.2	1	FULL/BDR	00:00:34	10.123.1.2	GigabitEthernet0/0
192.168.3.3	1	FULL/DR	00:00:37	10.123.1.3	GigabitEthernet0/0
R2# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.4.4	0	FULL/ -	00:00:38	10.24.1.4	Serial1/0
192.168.1.1	1	FULL/DROTHER	00:00:37	10.123.1.1	GigabitEthernet0/0
192.168.3.3	1	FULL/DR	00:00:34	10.123.1.3	GigabitEthernet0/0

**Table 6-6** Fields from the OSPF Neighbor State Output

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 State field is the neighbor state, as described in Table 6-4. The second State field is the DR, BDR, or DROTHER role if the interface requires a DR. For non-DR network links, the second field shows just a hyphen (-).
Dead Time	The dead 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.

# Viewing OSPF Installed Routes

You display OSPF routes installed in the Routing Information Base (RIB) by using the command **show ip route ospf**. In the output, two sets of numbers are in the brackets and look like [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 along with the next-hop IP address.

Example 6-7 provides the routing table for R4 from Figure 6-5. Notice that R4's OSPF routing table shows the routes from within Area 1234 and Area 0 as intra-area and routes from Area 56 as interarea because R4 does not connect to Area 56.

## Example 6-7 OSPF Routing Tables for ABR R4

```
R4# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
O       10.3.3.0/24 [110/66] via 10.24.1.2, 00:03:45, Serial1/0
O IA    10.56.1.0/24 [110/2]  via 10.45.1.5, 00:04:56, GigabitEthernet0/0
O       10.123.1.0/24 [110/65] via 10.24.1.2, 00:13:19, Serial1/0
```

# External OSPF Routes

When a router redistributes routes into an OSPF domain, the router is called an *autonomous system boundary router (ASBR)*. An ASBR can be any OSPF router, and the ASBR function is independent of the ABR function. An OSPF domain can have an ASBR without having an ABR. An OSPF router can be an ASBR and an ABR at the same time. External routes are classified as Type 1 or Type 2. The main differences between Type 1 and Type 2 external OSPF routes are as follows:

- Type 1 routes are preferred over Type 2 routes.
- The Type 1 metric equals the redistribution metric plus the total path metric to the ASBR. In other words, as the LSA propagates away from the originating ASBR, the metric increases.
- The Type 2 metric equals only the redistribution metric. The metric is the same for the router next to the ASBR as the router 30 hops away from the originating ASBR. This is the default external metric type used by OSPF.

# OSPF Configuration

## External OSPF Routes (Cont.)

Figure 6-6 revisits the previous topology where R6 is redistributing two networks in to the OSPF domain.

Example 6-9 shows only the OSPF routes in the routing table from R1 and R2. The 172.16.6.0/24 network is redistributed as a Type 1 route, and the 172.31.6.0/24 network is redistributed as a Type 2 route.

External OSPF network routes are marked as O E1 and O E2 in the routing table and correlate with OSPF Type 1 and Type 2 external routes.

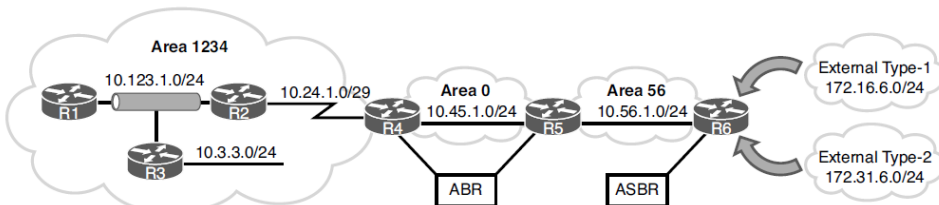


Figure 6-6 OSPF Multi-Area Topology with External Routes

### Example 6-9 Examining OSPF External Route Metrics on R1 and R2

```
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
       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, 6 subnets, 3 masks
O       10.3.3.0/24 [110/2] via 10.123.1.3, 23:20:25, GigabitEthernet0/0
O       10.24.1.0/29 [110/65] via 10.123.1.2, 23:20:15, GigabitEthernet0/0
O IA    10.45.1.0/24 [110/66] via 10.123.1.2, 23:13:25, GigabitEthernet0/0
O IA    10.56.1.0/24 [110/67] via 10.123.1.2, 23:13:25, GigabitEthernet0/0
       172.16.0.0/24 is subnetted, 1 subnets
O E1    172.16.6.0 [110/87] via 10.123.1.2, 00:01:00, GigabitEthernet0/0
       172.31.0.0/24 is subnetted, 1 subnets
O E2    172.31.6.0 [110/20] via 10.123.1.2, 00:01:00, GigabitEthernet0/0

R2# show ip route ospf | begin Gateway
Gateway of last resort is not set

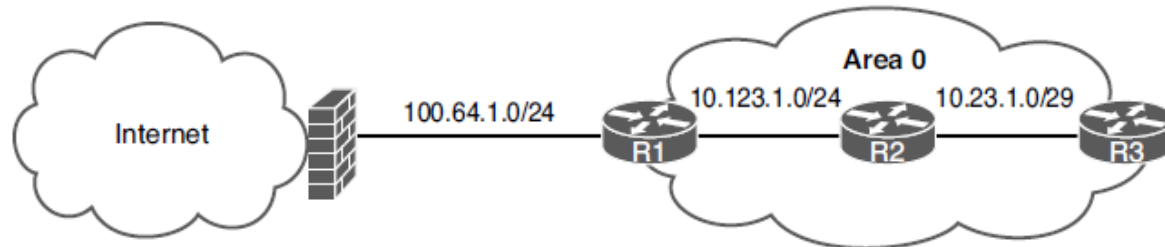
    10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
O       10.3.3.0/24 [110/2] via 10.123.1.3, 23:24:05, GigabitEthernet0/0
O IA    10.45.1.0/24 [110/65] via 10.24.1.4, 23:17:11, Serial1/0
O IA    10.56.1.0/24 [110/66] via 10.24.1.4, 23:17:11, Serial1/0
       172.16.0.0/24 is subnetted, 1 subnets
O E1    172.16.6.0 [110/86] via 10.24.1.4, 00:04:45, Serial1/0
       172.31.0.0/24 is subnetted, 1 subnets
O E2    172.31.6.0 [110/20] via 10.24.1.4, 00:04:45, Serial1/0
```

# Default Route Advertisement

OSPF supports advertising the default route into the OSPF domain. To advertise the default route, you use the command **default-information originate** [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] underneath the OSPF process. The **always** optional keyword causes the default route to be advertised even if a default route does not exist in the RIB.

The route metric can be changed with the **metric** *metric-value* option, and the metric type can be changed with the **metric-type** *type-value* option.

Figure 6-7 illustrates a common situation, where R1 has a static default route to the firewall, which is connected to the internet. To provide connectivity to other parts of the network (that is, R2 and R3), R1 advertises a default route into OSPF.



**Figure 6-7** *Default Route Topology*

# Default Route Advertisement (Cont.)

Example 6-10 provides the relevant configuration on R1. R1 has a static default route to the firewall (100.64.1.2) to satisfy the requirement of having the default route in the RIB.

Example 6-11 shows the routing tables of R2 and R3. Notice that OSPF advertises the default route as an external OSPF route.

**Example 6-10** *OSPF default-information originate 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
```

**Example 6-11** *R2's and R3's Routing Tables*

```
R2# show ip route | begin Gateway
Gateway of last resort is 10.12.1.1 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 10.12.1.1, 00:02:56, GigabitEthernet0/1
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C      10.12.1.0/24 is directly connected, GigabitEthernet0/1
C      10.23.1.0/24 is directly connected, GigabitEthernet0/2

R3# show ip route | begin Gateway
Gateway of last resort is 10.23.1.2 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 10.23.1.2, 00:01:47, GigabitEthernet0/1
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O      10.12.1.0/24 [110/2] via 10.23.1.2, 00:05:20, GigabitEthernet0/1
C      10.23.1.0/24 is directly connected, GigabitEthernet0/1
```

# The Designated Router and Backup Designated Router

- Multi-access networks such as Ethernet (LANs) and Frame Relay networks allow more than two routers to exist on a network segment.
- Additional routers flood more LSAs on the segment, and OSPF traffic becomes excessive as OSPF neighbor adjacencies increase.
- The Designated Router (DR) reduces the number of OSPF adjacencies on a multi-access network segment because routers form full OSPF adjacencies only with the DR and not each other.

# The Designated Router and Backup Designated Router Concepts

Having too many adjacencies per segment consumes more bandwidth, more CPU processing, and more memory to maintain each of the neighbor states.

OSPF overcomes this inefficiency by creating a pseudonode (that is, a virtual router) to manage the adjacency state with all the other routers on that broadcast network segment. A router on the broadcast segment, known as the designated router (DR), assumes the role of the pseudonode.

The DR is then responsible for flooding the update to all OSPF routers on that segment as updates occur. Figure 6-8 demonstrates how this simplifies a four-router topology using only three neighbor adjacencies.

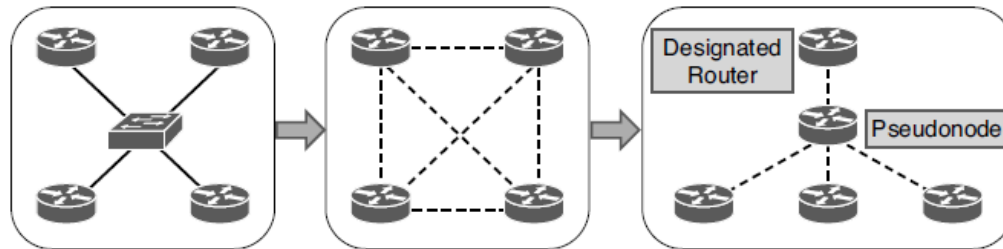


Figure 6-8 OSPF DR Concept

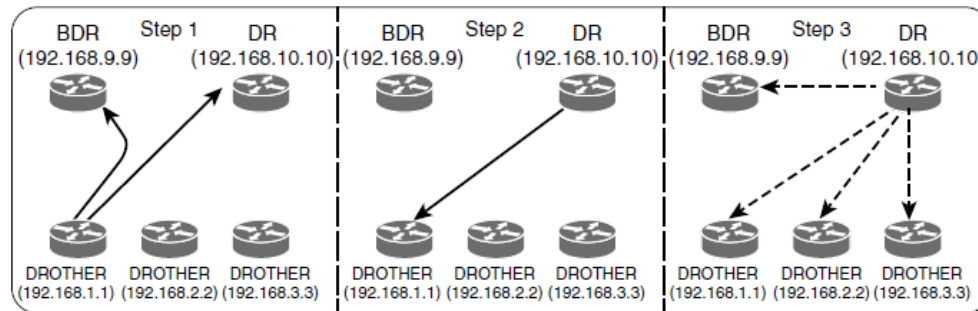
# The Designated Router and Backup Designated Router Concepts (Cont.)

The DR/BDR process distributes LSAs in the following manner:

**Step 1.** All OSPF routers (DR, BDR, and DROTHER) on a segment form a full OSPF adjacency with the DR and BDR. As an OSPF router learns of a new route, it sends the updated LSA to the AllDRouters (224.0.0.6) address, which only the DR and BDR receive and process, as illustrated in Step 1 in Figure 6-9.

**Step 2.** The DR sends a unicast acknowledgment to the router that sent the initial LSA update, as illustrated in Step 2 in Figure 6-9.

**Step 3.** The DR floods the LSA to all the routers on the segment via the AllSPFRouters 224.0.0.5) address, as shown in Step 3 in Figure 6-9.



**Figure 6-9** Network Prefix Advertisement with DR Segments

# The Designated Router and Backup Designated Router

## Designated Router Elections

The DR/BDR election occurs during OSPF neighborship. Specifically, during the last phase of the 2-Way neighbor state and just before the ExStart state.

- Any router with the 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.
- Routers receive and examine OSPF hellos from neighboring routers. If a router identifies itself as a more favorable router than the OSPF hellos it receives, it continues to send out hellos with its RID and priority listed.
- 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.
- When all the routers have agreed 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 as the DR election, except that the DR does not add its RID to the BDR field of the hello packet.

# The Designated Router and Backup Designated Router

## DR and BDR Placement

In Example 6-12, R3 wins the DR election, and R2 is elected the BDR because all the OSPF routers have the same OSPF priority, and the next decision is to use the higher RID. The RIDs match the Loopback 0 interface IP addresses, and R3's loopback address is the highest on that segment; R2's is the second highest.

Modifying a router's RID for DR placement is a bad design strategy. It is preferable to set the interface priority to a higher value than that of the existing DR. 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 over interfaces with the default value.

**Example 6-12** OSPF Interface State

R1# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	1	0	192.168.1.1/32	1	LOOP	0/0	
Gi0/0	1	0	10.123.1.1/24	1	DROTH	2/3	
R2# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	1	0	192.168.2.2/32	1	LOOP	0/0	
Se1/0	1	1234	10.24.1.1/29	64	P2P	1/1	
Gi0/0	1	1234	10.123.1.2/24	1	BDR	2/2	
R3# show ip ospf interface brief							
Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	1	0	192.168.3.3/32	1	LOOP	0/0	
Gi0/0	1	0	10.123.1.3/24	1	DR	2/2	
Gi0/1	1	0	10.3.3.3/24	1	DR	0/0	

# OSPF Network Types

- The default OSPF network type is set based on the media used for the connection and can be changed independently of the actual media type used.
- Cisco's implementation of OSPF considers the various media and provides five OSPF network types: Broadcast, Nonbroadcast, Point-to-point, Point-to-Multipoint and Loopback.

# OSPF Network Types

## Concepts

Table 6-7 provides detail on the five OSPF network types that exist in Cisco's implementation of OSPF.

**Table 6-7** OSPF Network Types

Type	Description	DR/BDR Field in OSPF Hellos	Timers
Broadcast	Default setting on OSPF-enabled Ethernet links.	Yes	Hello: 10 Wait: 40 Dead: 40
Nonbroadcast	Default setting on enabled OSPF Frame Relay main interface or Frame Relay multipoint subinterfaces.	Yes	Hello: 30 Wait: 120 Dead: 120
Point-to-point	Default setting on enabled OSPF 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 is advertised as a host route (/32), and sets the next-hop address to the outbound interface. Primarily used for hub-and-spoke topologies.	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

## OSPF Network Types

# Broadcast

- Broadcast networks are multi-access in that they are capable of connecting more than two devices, and broadcasts sent out one interface are capable of reaching all interfaces attached to that segment.
- The OSPF network type is set to broadcast by default for Ethernet interfaces. A DR is required for this OSPF network type because of the possibility that multiple nodes can exist on a segment and LSA flooding needs to be controlled.
- The hello timer defaults to 10 seconds, as defined in RFC 2328.
- The interface parameter command **ip ospf network broadcast** overrides the automatically configured setting and statically sets an interface as an OSPF broadcast network type.

# OSPF Network Types

## Nonbroadcast

- Frame Relay, ATM, and X.25 are considered nonbroadcast multi-access (NBMA) in that they can connect more than two devices, and broadcasts sent out one interface might not always be capable of reaching all the interfaces attached to the segment.
- Frame Relay interfaces set the OSPF network type to nonbroadcast by default.
- Multiple routers can exist on a segment, so the DR functionality is used.
- Neighbors are statically defined with the **neighbor ip-address** command because multicast and broadcast functionality do not exist on this type of circuit. Configuring a static neighbor causes OSPF hellos to be sent using unicast.
- The interface parameter command **ip ospf network non-broadcast** manually sets an interface as an OSPF nonbroadcast network type.

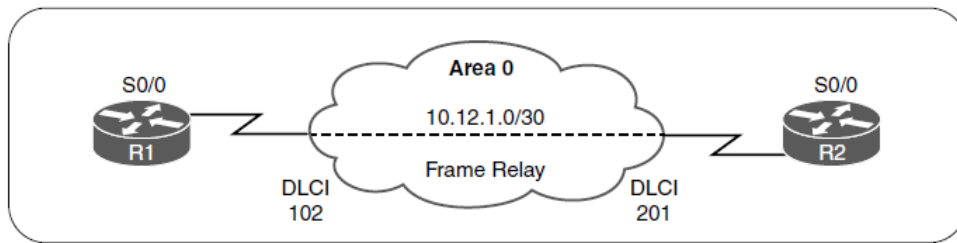


Figure 6-10 OSPF Topology Using Frame Relay

## OSPF Network Types

# Nonbroadcast (Cont.)

Example 6-13 provides the OSPF configuration over a Frame Relay interface. Notice that the static neighbor configuration is required when OSPF packets cannot be received through broadcast (multicast) discovery.

**Example 6-13** *OSPF Configuration for Frame Relay Interfaces*

```
R1
interface Serial 0/0
 ip address 10.12.1.1 255.255.255.252
 encapsulation frame-relay
 no frame-relay inverse-arp
 frame-relay map ip address 10.12.1.2 102
!
router ospf 1
 router-id 192.168.1.1
 neighbor 10.12.1.2
 network 0.0.0.0 255.255.255.255 area 0
```

The nonbroadcast network type is verified by filtering the output of the **show ip ospf interface** command with the *Type* keyword. The following snippet confirms that the interfaces operate as nonbroadcast:

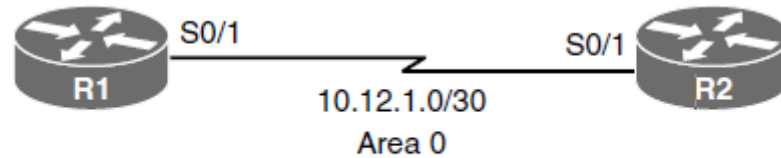
```
R1# show ip ospf interface Serial 0/0 | include Type
Process ID 1, Router ID 192.168.1.1, Network Type NON _ BROADCAST,
Cost: 64
```

# OSPF Network Types

## Point-to-Point

- A network circuit that allows only two devices to communicate is considered a point-to-point (P2P) network.
- The OSPF network type is set to point-to-point by default for serial interfaces (HDLC or PPP encapsulation), Generic Routing Encapsulation (GRE) tunnels, and point-to-point Frame Relay subinterfaces.
- There are no special configuration commands to set a serial interface to point-to-point.

Figure 6-11 shows a serial connection between R1 and R2.



**Figure 6-11** *OSPF Topology with Serial Interfaces*

# OSPF Network Types

## Point-to-Point (Cont.)

To verify that the OSPF network type is set to point-to-point, use the **show ip ospf interface** command with the *Type* keyword.

Example 6-15 verifies that the OSPF network type is set to POINT\_TO\_POINT, indicating the OSPF point-to-point network type.

**Example 6-15** *Verification of OSPF P2P Interfaces*

```
R1# show ip ospf interface s0/1 | include Type
  Process ID 1, Router ID 192.168.1.1, Network Type POINT_TO_POINT, Cost: 64
```

---

```
R2# show ip ospf interface s0/1 | include Type
  Process ID 1, Router ID 192.168.2.2, Network Type POINT_TO_POINT, Cost: 64
```

Example 6-16 shows that point-to-point OSPF network types do not use a DR. Notice the hyphen (-) in the State field.

**Example 6-16** *Verification of OSPF Neighbors on P2P Interfaces*

```
R1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.2.2	0	FULL/ -	00:00:36	10.12.1.2	Serial0/1

# OSPF Network Types

## Point-to-Multipoint Networks

- The OSPF network type point-to-multipoint is not enabled by default for any medium. It requires manual configuration. The IOS interface parameter command **ip ospf network point-to-multipoint** manually sets an interface as an OSPF point-to-multipoint network type.
- A point-to-multipoint OSPF network type supports hub-and-spoke connectivity while using the same IP subnet and is commonly found in Frame Relay and Layer 2 VPN (L2VPN) topologies.
- Interfaces set for the OSPF point-to-multipoint network type add the interface's IP address to the OSPF LSDB as a /32 network.
- When advertising routes to OSPF peers on that interface, the next-hop address is set to the IP address of the interface even if the next-hop IP address resides on the same IP subnet.

Figure 6-12 provides a topology example with R1, R2, and R3 all using Frame Relay point-to-multipoint subinterfaces using the same subnet.

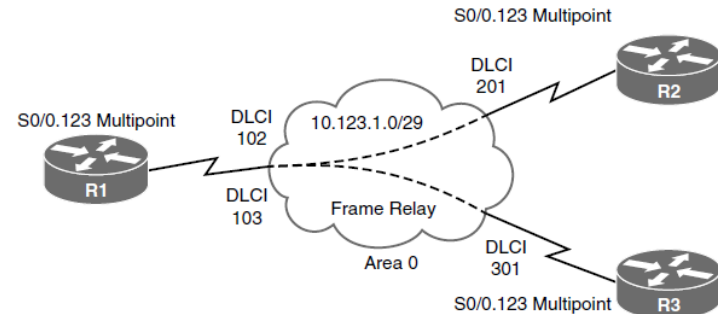


Figure 6-12 OSPF Topology with Frame Relay Multipoint Interfaces

# OSPF Network Types

## Point-to-Multipoint Networks (Cont.)

This image from Example 6-17 demonstrates the relevant configuration for router R1. The configurations for R2 and R3 are similar.

**Example 6-17** *OSPF Point-to-Multipoint Configuration*

```
R1
interface Serial 0/0
  encapsulation frame-relay
  no frame-relay inverse-arp
!
interface Serial 0/0.123 multipoint
  ip address 10.123.1.1 255.255.255.248
  frame-relay map ip 10.123.1.2 102 broadcast
  frame-relay map ip 10.123.1.3 103 broadcast
  ip ospf network point-to-multipoint
!
router ospf 1
  router-id 192.168.1.1
  network 0.0.0.0 255.255.255.255 area 0
```

Example 6-18 verifies that the interfaces are the OSPF point-to-multipoint network type.

**Example 6-18** *Verification of OSPF Network Type Point-to-Multipoint*

```
R1# show ip ospf interface Serial 0/0.123 | include Type
  Process ID 1, Router ID 192.168.1.1, Network Type POINT_TO_MULTIPOINT, Cost: 64

R2# show ip ospf interface Serial 0/0.123 | include Type
  Process ID 1, Router ID 192.168.2.2, Network Type POINT_TO_MULTIPOINT, Cost: 64

R3# show ip ospf interface Serial 0/0.123 | include Type
  Process ID 1, Router ID 192.168.3.3, Network Type POINT_TO_MULTIPOINT, Cost: 64
```

# OSPF Network Types

## Point-to-Multipoint Networks (Cont.)

Example 6-19 shows that OSPF does not use a DR for the OSPF point-to-multipoint network type. All three routers are on the same subnet, but R2 and R3 do not establish an adjacency with each other.

The snip from Example 6-20 shows that all the Serial 0/0.123 interfaces are advertised into OSPF as a /32 network and that the next-hop address is set (by R1) when advertised to the spokes nodes.

**Example 6-19** OSPF Neighbor Adjacency on a Hub-and-Spoke Topology

R1# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.3.3	0	FULL/ -	00:01:33	10.123.1.3	Serial0/0.123
192.168.2.2	0	FULL/ -	00:01:40	10.123.1.2	Serial0/0.123

R2# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.1	0	FULL/ -	00:01:49	10.123.1.1	Serial0/0.123

R3# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.1	0	FULL/ -	00:01:46	10.123.1.1	Serial0/0.123

**Example 6-20** OSPF Point-to-Multipoint Routing Tables

R1# show ip route ospf   begin Gateway					
Gateway of last resort is not set					
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks					
O	10.123.1.2/32	[110/64]	via 10.123.1.2,	00:07:32,	Serial0/0.123
O	10.123.1.3/32	[110/64]	via 10.123.1.3,	00:03:58,	Serial0/0.123
192.168.2.0/32 is subnetted, 1 subnets					
O	192.168.2.2	[110/65]	via 10.123.1.2,	00:07:32,	Serial0/0.123
192.168.3.0/32 is subnetted, 1 subnets					
O	192.168.3.3	[110/65]	via 10.123.1.3,	00:03:58,	Serial0/0.123

# OSPF Network Types

## Loopback Networks

- 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 indicates 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.
- Changing the loopback network type to point-to-point advertises the configured prefix length, rather than the /32 prefix length. Examples 6-22 and 6-23 show the effect of a change in R2's network type from loopback to point-to-point.

**Example 6-22** *Display of OSPF Network Type for Loopback Interfaces*

```
R1# show ip ospf interface Loopback 0 | include Type
Process ID 1, Router ID 192.168.1.1, Network Type LOOPBACK, Cost: 1

R2# show ip ospf interface Loopback 0 | include Type
Process ID 1, Router ID 192.168.2.2, Network Type POINT_TO_POINT, Cost: 1
```

**Example 6-23** *OSPF Routing Table for OSPF Loopback Network Types*

```
R1# show ip route ospf
! Output omitted for brevity
Gateway of last resort is not set

O          192.168.2.0/24 [110/65] via 10.12.1.2, 00:02:49, Serial0/0

R2# show ip route ospf
! Output omitted for brevity
Gateway of last resort is not set

          192.168.1.0/32 is subnetted, 1 subnets
O          192.168.1.1 [110/65] via 10.12.1.1, 00:37:15, Serial0/0
```

# Failure Detection

- A secondary function of OSPF hello packets is to ensure that adjacent OSPF neighbors are still healthy and available.
- The hello timer and the dead interval timer determine how long the OSPF process waits before declaring a neighbor state to be down.

# Failure Detection

## Timers

A secondary function of OSPF hello packets is to ensure that adjacent OSPF neighbors are still healthy and available. OSPF sends hello packets at set intervals, according to the hello timer. OSPF uses a second timer called the OSPF dead interval timer, which defaults to 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.

### Hello Timer

The default OSPF hello timer interval varies based on the OSPF network type. OSPF allows modification to the hello timer interval with values between 1 and 65,535 seconds. Changing the hello timer interval modifies the default dead interval, too. The OSPF hello timer is modified with the interface configuration submode **command ip ospf hello-interval 1-65,535**.

### Dead Interval Timer

You can change the dead interval timer to a value between 1 and 65,535 seconds. You change the OSPF dead interval timer by using the command `ip ospf dead-interval 1-65,535` under the interface configuration submode.

# Failure Detection Timers

## Verifying OSPF Timers

You view the timers for an OSPF interface by using the command **show ip ospf interface**, as demonstrated in Example 6-24. Notice the highlighted hello and dead timers.

### Example 6-24 *OSPF Interface timers*

```
R1# show ip ospf interface | i Timer|line
Loopback0 is up, line protocol is up
GigabitEthernet0/2 is up, line protocol is up
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
GigabitEthernet0/1 is up, line protocol is up
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

# Authentication

- An attacker can forge OSPF packets or gain physical access to the network.
- OSPF authentication is enabled on an interface-by-interface basis or for all interfaces in an area.
- OSPF supports either plaintext authentication or MD5 cryptographic hash.

# Authentication Types

OSPF supports two types of authentication:

**Plaintext** - Provides little security, as anyone with access to the link can see the password by using a network sniffer.

- Plaintext authentication for an OSPF area is enabled with the command **area area-id authentication**.
- The interface parameter command **ip ospf authentication** sets plaintext authentication only on that interface.
- The plaintext password is configured by using the interface parameter command **ip ospf authentication-key password**.

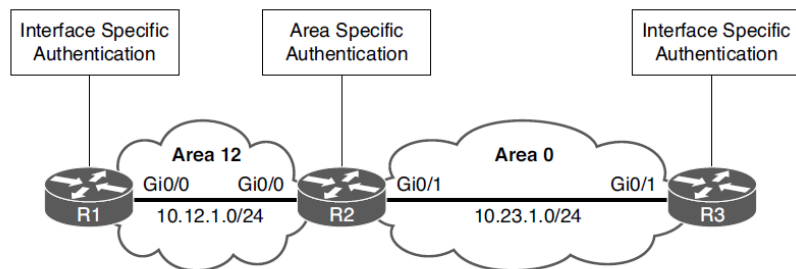
**MD5 cryptographic hash** - This type of authentication uses a hash, so the password is never sent out the wire. This technique is widely accepted as being more secure.

- MD5 authentication for an OSPF area is enabled using the command **area area-id authentication message-digest**.
- The interface parameter command **ip ospf authentication message-digest** sets MD5 authentication for that interface.
- The interface parameter command **ip ospf message-digest-key key-number md5 password** sets the MD5 password.

# Authentication Types

Figure 6-13 provides a simple topology to demonstrate the OSPF authentication configuration. Area 12 uses plaintext authentication, and Area 0 use MD5 authentication. R1 and R3 use interface-based authentication, and R2 uses area-specific authentication.

Example 6-25 provides the OSPF authentication configuration.



**Figure 6-13** Authentication Topology

## Example 6-25 OSPF Authentication Configuration

```

R1
interface GigabitEthernet0/0
 ip address 10.12.1.1 255.255.255.0
 ip ospf authentication
 ip ospf authentication-key CISCO
!
router ospf 1
 network 10.12.1.0 0.0.0.255 area 12

R2
interface GigabitEthernet0/0
 ip address 10.12.1.2 255.255.255.0
 ip ospf authentication-key CISCO
!
interface GigabitEthernet0/1
 ip address 10.23.1.2 255.255.255.0
 ip ospf message-digest-key 1 md5 CISCO
!
router ospf 1
 area 0 authentication message-digest
 area 12 authentication
 network 10.12.1.0 0.0.0.255 area 12
 network 10.23.1.0 0.0.0.255 area 0

R3
interface GigabitEthernet0/1
 ip address 10.23.1.3 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 CISCO
!
router ospf 1
 network 10.23.1.0 0.0.0.255 area 0
  
```

# Authentication Types (Cont.)

You verify the authentication settings by examining the OSPF interface *without* the **brief** option. Example 6-26 shows sample output from R1, R2, and R3, where the Gi0/0 interface uses MD5 authentication and the Gi0/1 interface uses plaintext authentication. MD5 authentication also identifies the key number that the interface uses.

## Example 6-26 IOS OSPF Authentication Verification

```
R1# show ip ospf interface | include line|authentication|key
```

```
GigabitEthernet0/0 is up, line protocol is up
```

```
Simple password authentication enabled
```

```
R2# show ip ospf interface | include line|authentication|key
```

```
GigabitEthernet0/1 is up, line protocol is up
```

```
Cryptographic authentication enabled
```

```
Youngest key id is 1
```

```
GigabitEthernet0/0 is up, line protocol is up
```

```
Simple password authentication enabled
```

```
R3# show ip ospf interface | include line|authentication|key
```

```
GigabitEthernet0/1 is up, line protocol is up
```

```
Cryptographic authentication enabled
```

```
Youngest key id is 1
```

# Prepare for the Exam

# Prepare for the Exam

## Key Topics for Chapter 6

Description	
OSPF Areas	Interface-specific configuration
OSPF backbone	External OSPF routes
Area border routers	The designated router
OSPF packet types	Designated router elections
OSPF neighbor states	DR and BDR placement
Requirements of neighbor adjacency	OSPF network types
OSPF network statement	Authentication

# Prepare for the Exam

## Key Terms for Chapter 6

Terms		
Router ID (RID)	Interface priority	Router LSA
Hello packets	Passive interface	Network LSA
Hello interval	Shortest path first tree (SPT)	Summary LSA
Dead interval	Area border router (ABR)	Inter-area route
Designated router (DR)	Backbone area	External OSPF route
Backup designated router (BDR)	Intra-area route	

# Prepare for the Exam

## Command Reference for Chapter 6

Task	Command Syntax
Initialize the OSPF process	<b>router ospf</b> <i>process-id</i>
Enable OSPF on network interfaces that match a specified network range for a specific OSPF area	<b>network</b> <i>ip-address wildcard-mask</i> <b>area</b> <i>area-id</i>
Enable OSPF on an explicit specific network interface for a specific OSPF area	<b>ip ospf</b> <i>process-id</i> <b>area</b> <i>area-id</i>
Configure a specific interface as passive	<b>passive</b> <i>interface-id</i>
Configure all interfaces as passive	<b>passive interface default</b>
Advertise a default route into OSPF	<b>default-information originate</b> [ <b>always</b> ] [ <b>metric</b> <i>metric-value</i> ] [ <b>metric-type</b> <i>type-value</i> ]
Modify the OSPF reference bandwidth for dynamic interface metric costing	<b>auto-cost reference-bandwidth</b> <i>bandwidth-in-mbps</i>

# Command Reference for Chapter 6 (Cont.)

Task	Command Syntax
Configure the OSPF priority for a DR/BDR election	<b>ip ospf priority</b> <i>0-255</i>
Statically configure an interface as a broadcast OSPF network type	<b>ip ospf network broadcast</b>
Statically configure an interface as a nonbroadcast OSPF network type	<b>ip ospf network non-broadcast</b>
Statically configure an interface as a point-to-point OSPF network type	<b>ip ospf network point-to-point</b>
Statically configure an interface as a point-to-multipoint OSPF network type	<b>ip ospf network point-to-multipoint</b>
Enable OSPF authentication for an area	<b>area</b> <i>area-id</i> <b>authentication [message-digest]</b>
Define the plaintext password for an interface	<b>ip ospf authentication-key</b> <i>password</i>

## Command Reference for Chapter 6 (Cont.)

Task	Command Syntax
Define the MD5 password for an interface	<b>ip ospf message-digest-key</b> <i>key-number</i> <b>md5</b> <i>password</i>
Restart the OSPF process	<b>clear ip ospf process</b>
Display the OSPF interfaces on a router	<b>show ip ospf interface</b> [ <b>brief</b>   <i>interface-id</i> ]
Display the OSPF neighbors and their current states	<b>show ip ospf neighbor</b> [ <b>detail</b> ]
Display the OSPF routes that are installed in the RIB	<b>show ip route ospf</b>

