

Module 2: Single-Area OSPFv2 Configuration

Enterprise Networking, Security, and Automation v7.0
(ENSA)



Module Objectives

Module Title: Single-Area OSPFv2 Configuration

Module Objective: Implement single-area OSPFv2 in both point-to-point and broadcast multiaccess networks.

Topic Title	Topic Objective
OSPF Router ID	Configure an OSPFv2 router ID.
Point-to-Point OSPF Networks	Configure single-area OSPFv2 in a point-to-point network.
Multiaccess OSPF Networks	Configure the OSPF interface priority to influence the DR/BDR election in a multiaccess network.
Modify Single-Area OSPFv2	Implement modifications to change the operation of single-area OSPFv2.
Default Route Propagation	Configure OSPF to propagate a default route.
Verify Single-Area OSPFv2	Verify a single-area OSPFv2 implementation.

Module Overview/Summary

Don't stress about there being 70 slides in this deck. Once you understand the theory, it all boils down to about a dozen commands:

1. Router ID (1 command)
2. Activating Interfaces / Including Subnets in OSPF (1+ command)
3. Multi-Access links - OSPF router priority (1 command)
4. Tuning timers and fixing the Cost reference (4 commands)
5. Propagating a Default Route (1 command)
6. Verification (show) commands (4 commands)

New Terms and Commands

- **router ospf process-id**
- router ID
- router-id rid
- **show ip protocols**
- **show ip ospf**
- **network** *network-address wildcard-mask* **area** *area-id*
- **ip ospf** *process-id* **area** *area-id*
- **passive interface**
- passive-interface intf-id
- **show ip ospf interface** *intf-id*
- **ip ospf network point-to-point**
- host route
- designated router (DR)
- backup designated router (BDR)
- DROTHER
- **show ip ospf neighbor**
- **ip ospf priority**
- **clear ip ospf process**
- metric
- cost
- **auto-cost reference-bandwidth** Mbps
- **ip ospf cost** *value*
- **ip ospf hello-interval** *value*
- **ip ospf dead-interval** *value*
- **default-information originate**

2.1 OSPF Router ID


Summary of next 8 slides:
router-id x.x.x.x

OSPF Reference Topology

The figure shows the example topology we'll use for OSPFv2

The routers in the topology have a basic config, including interface addresses, but there is no other routing configured

All interfaces are within the OSPF backbone area (except the loopback 1 on R2 which simulates the internet). The ISP router is used as the default gateway for this topology.



Router Configuration Mode for OSPF

OSPFv2 is enabled using the **router ospf process-id** global configuration mode command. The *process-id* value represents a number between 1 and 65,535 and is selected by the network administrator. The *process-id* value is locally significant. It is considered best practice to use the same *process-id* on all OSPF routers (i.e. **so we don't drive YOU or other human administrators crazy!!**)

```
R1(config)# router ospf 10
R1(config-router)# ?
  area                OSPF area parameters
  auto-cost           Calculate OSPF interface cost according to bandwidth
  default-information Control distribution of default information
  distance            Define an administrative distance
  exit               Exit from routing protocol configuration mode
  log-adjacency-changes Log changes in adjacency state
  neighbor           Specify a neighbor router
  network            Enable routing on an IP network
  no                 Negate a command or set its defaults
  passive-interface  Suppress routing updates on an interface
  redistribute       Redistribute information from another routing protocol
  router-id          router-id for this OSPF process
R1(config-router)#
```

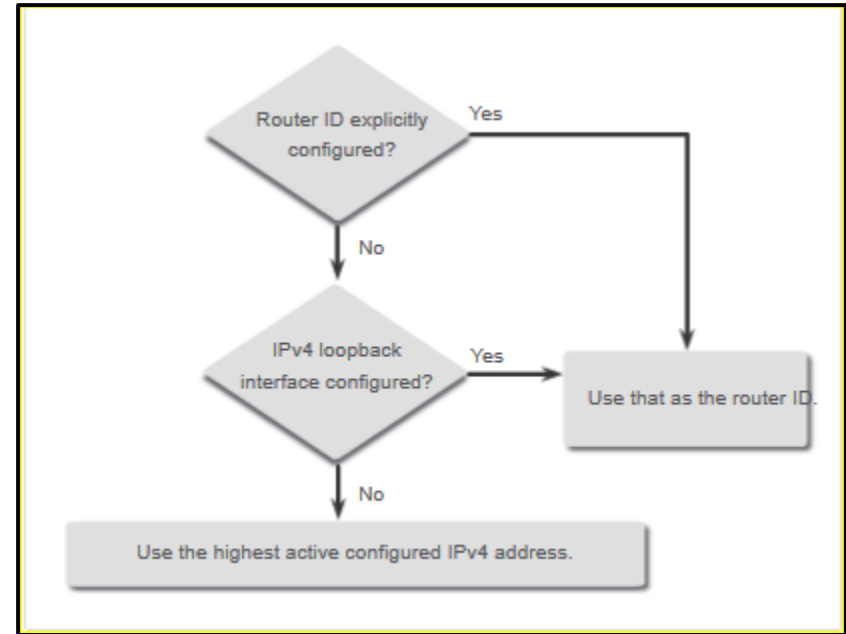
Router IDs

- An OSPF router ID is a 32-bit value, represented as an IPv4 address. It is used to uniquely identify an OSPF router, and all OSPF packets include the router ID of the originating router.
- Every router requires a router ID to participate in an OSPF domain. It can be defined by an administrator or automatically assigned by the router. The router ID is used by an OSPF-enabled router to do the following:
 - **Participate in the synchronization of OSPF databases** – During the Exchange State, the router with the highest router ID will send their database descriptor (DBD) packets first.
 - **Participate in the election of the designated router (DR)** - In a multi-access LAN environment, the router with the highest router ID is elected the DR. The routing device with the second highest router ID is elected the backup designated router (BDR).

Router ID Order of Precedence

Cisco routers derive the router ID based on one of three criteria, in the following order of precedence:

1. The router ID is explicitly configured using the OSPF **router-id** *rid* router configuration mode command. This is the recommended method to assign a router ID.
2. The router chooses the highest IPv4 address of any (**active!**) configured loopback interfaces.
3. The router chooses the highest **active** IPv4 address of any of its physical interfaces.



Configure a Loopback Interface as the Router ID

Instead of relying on physical interface, the router ID ~~can~~ should be assigned to a loopback interface. Typically, the IPv4 address for this type of loopback interface should be configured using a 32-bit subnet mask (255.255.255.255). This effectively creates a host route.

~~A 32-bit host route would not get advertised as a route to other OSPF routers.~~
(That's pure nonsense; wait until NET3012! And also see slide 24.)

OSPF does not need to be enabled on an interface for that interface to be chosen as the router ID.

```
R1(config-if)# interface Loopback 1
R1(config-if)# ip address 1.1.1.1 255.255.255.255
R1(config-if)# end
R1# show ip protocols | include Router ID
  Router ID 1.1.1.1
R1#
```

REAL Example of the Address/Router ID on a Loopback Interface

```
C:\>tracert netlab.algonquincollege.com
```

```
Tracing route to netlab.algonquincollege.com [205.211.76.254]
over a maximum of 30 hops:
```

Traceroute run on Bell residential fiber network.

1	<1 ms	<1 ms	<1 ms	Tomato-GW.dsl.teksavvy.com [192.168.0.1]
2	2 ms	1 ms	1 ms	mynetwork.home [192.168.2.1]
3	2 ms	3 ms	2 ms	10.11.12.169 ← What do you notice about this IP?
4	*	*	*	Request timed out.
5	*	*	*	Request timed out.
6	6 ms	5 ms	5 ms	cr02-mtrlpq02ho6-bundle-ether10.net.bell.ca [142.124.127.124]
7	5 ms	3 ms	4 ms	bx1-montrealgz_et-0-0-5.net.bell.ca [64.230.26.135]
8	5 ms	4 ms	4 ms	209.148.215.253
9	10 ms	8 ms	9 ms	209.148.233.129
10	9 ms	8 ms	8 ms	69.63.248.94
11	9 ms	8 ms	9 ms	209.148.253.178
12	9 ms	9 ms	9 ms	207.107.79.174
13	8 ms	8 ms	8 ms	aip76-254.algonquincollege.com [205.211.76.254]

Trace complete.

Explicitly Configure a Router ID

In our reference topology the router ID for each router is assigned as follows:

- R1 uses router ID 1.1.1.1
- R2 uses router ID 2.2.2.2
- R3 uses router ID 3.3.3.3
- Use the **router-id** *rid* router configuration mode command to manually assign a router ID. In the example, the router ID 1.1.1.1 is assigned to R1. Use the **show ip protocols** command to verify the router ID.

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# end
*May 23 19:33:42.689: %SYS-5-CONFIG_I: Configured from console by console
R1# show ip protocols | include Router ID
    Router ID 1.1.1.1
R1#
```

Modify a Router ID

- After a router selects a router ID, an active OSPF router does not allow the router ID to be changed until the router is reloaded or the OSPF process is reset.
- Clearing the OSPF process is the preferred method to reset the router ID.

```
R1# show ip protocols | include Router ID
Router ID 10.10.1.1
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect
R1(config-router)# end
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
*Jun 6 01:09:46.975: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on GigabitEthernet0/0/1 from FULL to
DOWN, Neighbor Down: Interface down or detached
*Jun 6 01:09:46.981: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on GigabitEthernet0/0/1 from LOADING
to FULL, Loading Done *
R1# show ip protocols | include Router ID
Router ID 1.1.1.1
R1#
```

2.2 Activate Interfaces for OSPF

(Whether point-to-point OR multi-access!!)

The HARD way: slide 17

The EASY way: slide 19

The network Command Syntax

- You can specify the interfaces that belong to a point-to-point network by configuring the **network** command. You can also configure OSPF directly on the interface with the **ip ospf** command.
- The basic syntax for the **network** command is as follows:

```
Router(config-router)# network network-address wildcard-mask area area-id
```

- The *network-address wildcard-mask* syntax is used to enable OSPF on interfaces. Any interfaces on a router that match this part of the command are enabled to send and receive OSPF packets.
- The **area area-id** syntax refers to the OSPF area. When configuring single-area OSPFv2, the **network** command must be configured with the same *area-id* value on all routers. Although any area ID can be used, it is good practice to use an area ID of 0 with single-area OSPFv2. This convention makes it easier if the network is later altered to support multiarea OSPFv2.

Point-to-Point OSPF Networks


The Wildcard Mask

- The wildcard mask is typically the inverse of the subnet mask configured on that interface.
- The easiest method for calculating a wildcard mask is to subtract the network subnet mask from 255.255.255.255, as shown for /24 and /26 subnet masks in the figure.

Calculating a Wildcard Mask for /24



Calculating a Wildcard Mask for /26



Note!! One quick check on your math: since **all** subnet masks are even (except for .0 and .255), all wildcard masks are odd numbered.

Configure OSPF Using the network Command

Within routing configuration mode, there are two ways to identify the interfaces that will participate in the OSPFv2 routing process.

- In the first example, the wildcard mask identifies the interface based on the network addresses. Any active interface that is configured with an IPv4 address belonging to that network will participate in the OSPFv2 routing process.
- **Note:** Some IOS versions allow the subnet mask to be entered instead of the wildcard mask. The IOS then converts the subnet mask to the wildcard mask format.

```
R1 (config)# router ospf 10
R1 (config-router)# network 10.10.1.0 0.0.0.255 area 0
R1 (config-router)# network 10.1.1.4 0.0.0.3 area 0
R1 (config-router)# network 10.1.1.12 0.0.0.3 area 0
R1 (config-router)#
```

Configure OSPF Using the network Command (Cont.)

- As an alternative, OSPFv2 can be enabled by specifying the exact interface IPv4 address using a quad zero wildcard mask. Entering **network 10.1.1.5 0.0.0.0 area 0** on R1 tells the router to enable interface Gigabit Ethernet 0/0/0 for the routing process.
- The advantage of specifying the interface is that the wildcard mask calculation is not necessary. Notice that in all cases, the **area** argument specifies area 0.

```
R1(config)# router ospf 10
R1(config-router)# network 10.10.1.1 0.0.0.0 area 0
R1(config-router)# network 10.1.1.5 0.0.0.0 area 0
R1(config-router)# network 10.1.1.14 0.0.0.0 area 0
R1(config-router)#
```

Configure OSPF Using the ip ospf Command

To configure OSPF directly on the interface, use the **ip ospf** interface configuration mode command. The syntax is as follows:

```
Router(config-if)# ip ospf process-id area area-id
```

Remove the network commands using the **no** form of the command. Then go to each interface and configure the **ip ospf** command

```
! Undo the HARD way of configuring OSPF
R1(config)# router ospf 10
R1(config-router)# no network 10.10.1.1 0.0.0.0 area 0
R1(config-router)# no network 10.1.1.5 0.0.0.0 area 0
R1(config-router)# no network 10.1.1.14 0.0.0.0 area 0

! Configure the EASY way (no wildcard needed; same as IPv6!)
R1(config-router)# interface GigabitEthernet 0/0/0
R1(config-if)# ip ospf 10 area 0
R1(config-if)# interface GigabitEthernet 0/0/1
R1(config-if)# ip ospf 10 area 0
R1(config-if)# interface Loopback 0
R1(config-if)# ip ospf 10 area 0
```

Passive Interface

By default, OSPF messages are forwarded out all OSPF-enabled interfaces. However, these messages only need to be sent out interfaces that are connecting to other OSPF-enabled routers.

Sending out unneeded messages on a LAN affects the network in three ways:

- **Inefficient Use of Bandwidth** - Available bandwidth is consumed transporting unnecessary messages.
- **Inefficient Use of Resources** - All devices on the LAN must process and eventually discard the message.
- **Increased Security Risk** - Without additional OSPF security configurations, OSPF messages can be intercepted with packet sniffing software. Routing updates can be modified and sent back to the router, corrupting the routing table with false metrics that misdirect traffic.

Configure Passive Interfaces

- Use the **passive-interface** router configuration mode command to prevent the transmission of routing messages through a router interface, but still allow that network to be advertised to other routers.
- The **show ip protocols** command is then used to verify that the interface is listed as passive.

```
R1(config)# router ospf 10
R1(config-router)# passive-interface loopback 0
R1(config-router)# end
R1#
*May 23 20:24:39.309: %SYS-5-CONFIG_I: Configured from console by console
R1# show ip protocols
*** IP Routing is NSF aware ***
(output omitted)
Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
    GigabitEthernet0/0/1
    GigabitEthernet0/0/0
  Passive Interface(s):
    Loopback0
  Routing Information Sources:
    Gateway         Distance       Last Update
    3.3.3.3          110            01:01:48
    2.2.2.2          110            01:01:38
  Distance: (default is 110)
R1#
```

OSPF Point-to-Point Networks

By default, Cisco routers elect a DR and BDR on Ethernet interfaces, even if there is only one other device on the link. You can verify this with the **show ip ospf interface** command. The DR/ BDR election process is unnecessary as there can only be two routers on the point-to-point network between R1 and R2. Notice in the output that the router has designated the network type as BROADCAST.

```
R1# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.5/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, 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) 2.2.2.2, Interface address 10.1.1.6
  Backup Designated router (ID) 1.1.1.1, Interface address 10.1.1.5
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
```

OSPF Point-to-Point Networks (Cont.)

To change this to a point-to-point network, use the interface configuration command **ip ospf network point-to-point** on all interfaces where you want to disable the DR/BDR election process.

```
R1(config)# interface GigabitEthernet 0/0/0
R1(config-if)# ip ospf network point-to-point
*Jun 6 00:44:05.208: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on GigabitEthernet0/0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Jun 6 00:44:05.211: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on GigabitEthernet0/0/0 from
LOADING to FULL, Loading Done
R1(config-if)# end
R1# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.5/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 1

Topology-MTID      Cost      Disabled      Shutdown      Topology Name
```

Loopbacks and Point-to-Point Networks

- Use loopbacks to provide additional interfaces for a variety of purposes. By default, loopback interfaces are advertised as /32 host routes.
- To simulate a real LAN, the loopback interface can be configured as a point-to-point network to advertise the full network.
- What R2 sees when R1 advertises the loopback interface as-is:

```
R2# show ip route | include 10.10.1
O    10.10.1.1/32 [110/2] via 10.1.1.5, 00:03:05, GigabitEthernet0/0/0
```

- Configuration change at R1:

```
R1(config-if)# interface Loopback 0
R1(config-if)# ip ospf network point-to-point
```

- Result at R2:

```
R2# show ip route | include 10.10.1
O    10.10.1.0/24 [110/2] via 10.1.1.5, 00:03:05, GigabitEthernet0/0/0
```


2.3 Multi-access Links in OSPF

Summary of next 12 slides:


```
# ip ospf priority xx
```

OSPF Network Types

Another type of network that uses OSPF is the multiaccess OSPF network.

Multiaccess OSPF networks are unique in that one router controls the distribution of LSAs.

The router that is elected for this role should be determined by the network administrator through proper configuration.




OSPF Designated Router

- In multiaccess networks, OSPF elects a DR and BDR. The DR is responsible for collecting and distributing LSAs sent and received. The DR uses the multicast IPv4 address 224.0.0.5 which is meant for all OSPF routers.
- A BDR is also elected in case the DR fails. The BDR listens passively and maintains a relationship with all the routers. If the DR stops producing Hello packets, the BDR promotes itself and assumes the role of DR.
- All other routers become a DROTHER (a router that is neither the DR nor the BDR). DROTHERs use the multiaccess address 224.0.0.6 (all designated routers) to send OSPF packets to the DR and BDR. Only the DR and BDR listen for 224.0.0.6.

OSPF Multiaccess Reference Topology

- In the multiaccess topology shown in the figure, there are three routers interconnected over a common Ethernet multiaccess network, 192.168.1.0/24.
- Because the routers are connected over a common multiaccess network, OSPF has automatically elected a DR and BDR. R3 has been elected as the DR because its router ID is 3.3.3.3, which is the highest in this network. R2 is the BDR because it has the second highest router ID in the network.



Verify OSPF Router Roles

To verify the roles of the OSPFv2 router, use the **show ip ospf interface** command.

The output generated by R1 confirms that the following:

- R1 is not the DR or BDR, but is a DROTHER with a default priority of 1. (Line 7)
- The DR is R3 with router ID 3.3.3.3 at IPv4 address 192.168.1.3, while the BDR is R2 with router ID 2.2.2.2 at IPv4 address 192.168.1.2. (Lines 8 and 9)
- R1 has two adjacencies: one with the BDR and one with the DR. (Lines 20-22)

```
R1# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.1.1/24, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
  (output omitted)
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
  Backup Designated router (ID) 2.2.2.2, Interface address 192.168.1.2
  (output omitted)
  Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 2.2.2.2 (Backup Designated Router)
  Adjacent with neighbor 3.3.3.3 (Designated Router)
  Suppress hello for 0 neighbor(s)
R1#
```

Verify OSPF Router Roles (Cont.)

The output generated by R2 confirms that:

- R2 is the BDR with a default priority of 1. (Line 7)
- The DR is R3 with router ID 3.3.3.3 at IPv4 address 192.168.1.3, while the BDR is R2 with router ID 2.2.2.2 at IPv4 address 192.168.1.2. (Lines 8 and 9)
- R2 has two adjacencies; one with a neighbor with router ID 1.1.1.1 (R1) and the other with the DR. (Lines 20-22)

```
R2# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.1.2/24, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  (output omitted)
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
  Backup Designated Router (ID) 2.2.2.2, Interface address 192.168.1.2
  (output omitted)
  Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 1.1.1.1
  Adjacent with neighbor 3.3.3.3 (Designated Router)
  Suppress hello for 0 neighbor(s)
R2#
```

Verify OSPF Router Roles (Cont.)

The output generated by R3 confirms that:

- R3 is the DR with a default priority of 1. (Line 7)
- The DR is R3 with router ID 3.3.3.3 at IPv4 address 192.168.1.3, while the BDR is R2 with router ID 2.2.2.2 at IPv4 address 192.168.1.2. (Lines 8 and 9)
- R3 has two adjacencies: one with a neighbor with router ID 1.1.1.1 (R1) and the other with the BDR. (Lines 20-22)

```
R1# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.1.1/24, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
  (output omitted)
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
  Backup Designated router (ID) 2.2.2.2, Interface address 192.168.1.2
  (output omitted)
  Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 2.2.2.2 (Backup Designated Router)
  Adjacent with neighbor 3.3.3.3 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

```
R1#
```

Verify DR/BDR Adjacencies

To verify the OSPFv2 adjacencies, use the **show ip ospf neighbor** command. The state of neighbors in multiaccess networks can be as follows:

- **FULL/DROTHER** - This is a DR or BDR router that is fully adjacent with a non-DR or BDR router. These two neighbors can exchange Hello packets, updates, queries, replies, and acknowledgments.
- **FULL/DR** - The router is fully adjacent with the indicated DR neighbor. These two neighbors can exchange Hello packets, updates, queries, replies, and acknowledgments.
- **FULL/BDR** - The router is fully adjacent with the indicated BDR neighbor. These two neighbors can exchange Hello packets, updates, queries, replies, and acknowledgments.
- **2-WAY/DROTHER** - The non-DR or BDR router has a neighbor relationship with another non-DR or BDR router. These two neighbors exchange Hello packets.

The normal state for an OSPF router is usually FULL. If a router is stuck in another state, it is an indication that there are problems in forming adjacencies. The only exception to this is the 2-WAY state, which is normal in a multiaccess broadcast network.

Verify DR/BDR Adjacencies (Cont.)

The output generated by R2 confirms that R2 has adjacencies with the following routers:

- R1 with router ID 1.1.1.1 is in a Full state and R1 is neither the DR nor BDR.
- R3 with router ID 3.3.3.3 is in a Full state and the role of R3 is DR.

```
R2# show ip ospf neighbor
Neighbor ID    Pri  State           Dead Time   Address        Interface
1.1.1.1        1    FULL/DROTHER    00:00:31   192.168.1.1   GigabitEthernet0/0/0
3.3.3.3        1    FULL/DR         00:00:34   192.168.1.3   GigabitEthernet0/0/0 R2#
```

Default DR/BDR Election Process

The OSPF DR and BDR election is based on the following criteria, in sequential order:

1. The routers in the network elect the router with the highest interface priority as the DR. The router with the second highest interface priority is becomes the BDR.
 - The priority can be configured to be any number between 0 – 255.
 - If the interface priority value is set to 0, that interface cannot be elected as DR nor BDR.
 - The default priority of multiaccess broadcast interfaces is 1.
2. If the interface priorities are equal, then the router with the highest router ID is elected the DR. The router with the second highest router ID is the BDR.
 - The election process takes place when the first router with an OSPF-enabled interface is active on the network. If all of the routers on the network have not finished booting, it is possible that a router with a lower router ID becomes the DR.
 - The addition of a new router does not initiate a new election process.

DR Failure and Recovery

After the DR is elected, it remains the DR until one of the following events occurs:

- The DR fails.
- The OSPF process on the DR fails or is stopped.
- The multiaccess interface on the DR fails or is shutdown.

If the DR fails, the BDR is automatically promoted to DR. This is the case even if another DROTHER with a higher priority or router ID is added to the network after the initial DR/BDR election. However, after a BDR is promoted to DR, a new BDR election occurs and the DROTHER with the highest priority or router ID is elected as the new BDR.

The ip ospf priority Command

- If the interface priorities are equal on all routers, the router with the highest router ID is elected the DR.
- Instead of relying on the router ID, it is better to control the election by setting interface priorities. This also allows a router to be the DR in one network and a DROTHER in another.
- To set the priority of an interface, use the command **ip ospf priority *value***, where *value* is 0 to 255.
 - A value of 0 does not become a DR or a BDR.
 - A value of 1 to 255 on the interface makes it more likely that the router becomes the DR or the BDR.

Configure OSPF Priority

The example shows the commands being used to change the R1 G0/0/0 interface priority from 1 to 255 and then reset the OSPF process.

```
R1(config)# interface GigabitEthernet 0/0/0
R1(config-if)# ip ospf priority 255
R1(config-if)# end
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R1# *Jun 5 03:47:41.563: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on GigabitEthernet0/0/0
from FULL to DOWN, Neighbor Down: Interface down or detached
```

2.4 Modify Single-Area OSPFv2

Cisco OSPF Cost Metric

- Routing protocols use a metric to determine the best path of a packet across a network. OSPF uses cost as a metric. A lower cost indicates a better path.
- The Cisco cost of an interface is inversely proportional to the bandwidth of the interface. Therefore, a higher bandwidth indicates a lower cost. The formula used to calculate the OSPF cost is:

$$\text{Cost} = \text{reference bandwidth} / \text{interface bandwidth}$$

- The default reference bandwidth is 10^8 (100,000,000); therefore, the formula is:

$$\text{Cost} = 100,000,000 \text{ bps} / \text{interface bandwidth in bps}$$

- Because the OSPF cost value must be an integer, FastEthernet, Gigabit Ethernet, and 10 GigE interfaces share the same cost. To correct this situation, you can:
 - Adjust the reference bandwidth with the **auto-cost reference-bandwidth** command on each OSPF router.
 - Manually set the OSPF cost value with the **ip ospf cost** command on necessary interfaces.

Cisco OSPF Cost Metric (Cont.)

Refer to the table for a breakdown of the cost calculation

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	100,000,000	÷	10,000,000,000	0.01 = 1
Gigabit Ethernet 1 Gbps	100,000,000	÷	1,000,000,000	0.1 = 1
Fast Ethernet 100 Mbps	100,000,000	÷	100,000,000	1
Ethernet 10 Mbps	100,000,000	÷	10,000,000	1

Same Costs due to reference bandwidth

Adjust the Reference Bandwidth

- The cost value must be an integer. If something less than an integer is calculated, OSPF rounds up to the nearest integer. Therefore, the OSPF cost assigned to a Gigabit Ethernet interface with the default reference bandwidth of 100,000,000 bps would equal 1, because the nearest integer for 0.1 is 0 instead of 1.

$$\text{Cost} = 100,000,000 \text{ bps} / 1,000,000,000 = 1$$

- For this reason, all interfaces faster than Fast Ethernet will have the same cost value of 1 as a Fast Ethernet interface.
- To assist OSPF in making the correct path determination, the reference bandwidth must be changed to a higher value to accommodate networks with links faster than 100 Mbps.

Adjust the Reference Bandwidth (Cont.)

- Changing the reference bandwidth does not actually affect the bandwidth capacity on the link; rather, it simply affects the calculation used to determine the metric.
- To adjust the reference bandwidth, use the **auto-cost reference-bandwidth *Mbps*** router configuration command.
 - This command must be configured on every router in the OSPF domain.
 - Notice in the command that the value is expressed in Mbps; therefore, to adjust the costs for Gigabit Ethernet, use the command **auto-cost reference-bandwidth 1000**. For 10 Gigabit Ethernet, use the command **auto-cost reference-bandwidth 10000**.
 - To return to the default reference bandwidth, use the **auto-cost reference-bandwidth 100** command.
- Another option is to change the cost on one specific interface using the **ip ospf cost *cost*** command.

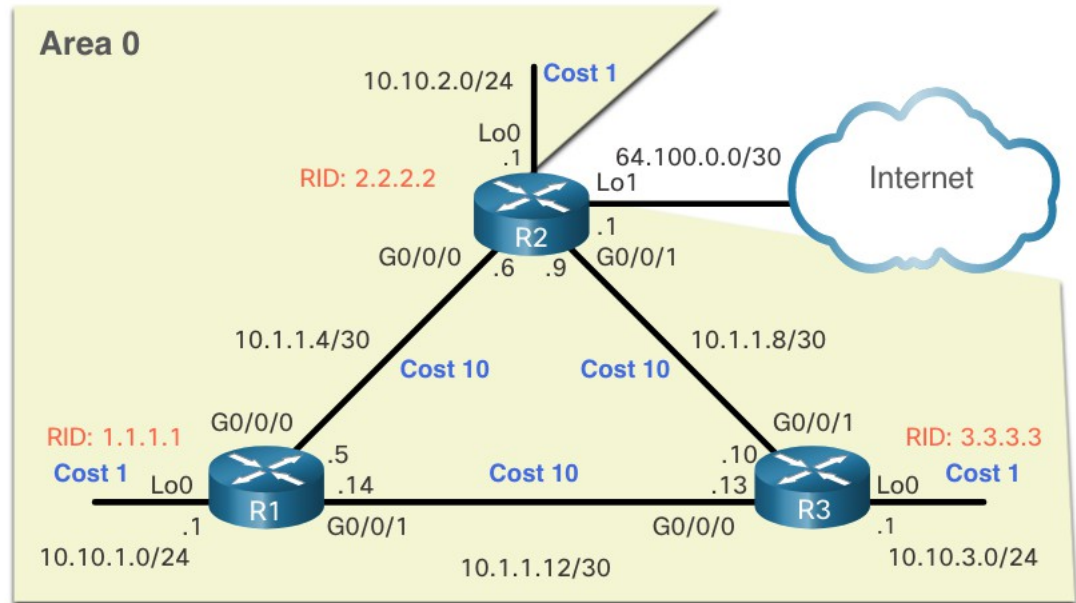
Adjust the Reference Bandwidth (Cont.)

- Whichever method is used, it is important to apply the configuration to all routers in the OSPF routing domain.
- The table shows the OSPF cost if the reference bandwidth is adjusted to accommodate 10 Gigabit Ethernet links. The reference bandwidth should be adjusted anytime there are links faster than FastEthernet (100 Mbps).
- Use the **show ip ospf interface** command to verify the current OSPFv2 cost assigned to the interface.

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	10,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	10,000,000,000	÷	1,000,000,000	10
Fast Ethernet 100 Mbps	10,000,000,000	÷	100,000,000	100
Ethernet 10 Mbps	10,000,000,000	÷	10,000,000	1000


OSPF Accumulates Cost

- The cost of an OSPF route is the accumulated value from one router to the destination network.
- Assuming the **auto-cost reference-bandwidth 10000** command has been configured on all three routers, the cost of the links between each router is now 10. The loopback interfaces have a default cost of 1.



OSPF Accumulates Cost (Cont.)

- You can calculate the cost for each router to reach each network.
- For example, the total cost for R1 to reach the 10.10.2.0/24 network is 11. This is because the link to R2 cost = 10 and the loopback default cost = 1. $10 + 1 = 11$.
- You can verify this with the **show ip route** command.



OSPF Accumulates Cost (Cont.)

Verifying the accumulated cost for the path to the 10.10.2.0/24 network:

```
R1# show ip route | include 10.10.2.0
O      10.10.2.0/24 [110/11] via 10.1.1.6, 01:05:02, GigabitEthernet0/0/0
R1# show ip route 10.10.2.0
Routing entry for 10.10.2.0/24
  Known via "ospf 10", distance 110, metric 11, type intra area
  Last update from 10.1.1.6 on GigabitEthernet0/0/0, 01:05:13 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 2.2.2.2, 01:05:13 ago, via GigabitEthernet0/0/0
    Route metric is 11, traffic share count is 1
R1#
```

Manually Set OSPF Cost Value

Reasons to manually set the cost value include:

- The Administrator may want to influence path selection within OSPF, causing different paths to be selected than what normally would given default costs and cost accumulation.
- Connections to equipment from other vendors who use a different formula to calculate OSPF cost.

To change the cost value reported by the local OSPF router to other OSPF routers, use the interface configuration command **ip ospf cost** *value*.

```
R1(config)# interface g0/0/1
R1(config-if)# ip ospf cost 30
R1(config-if)# interface lo0
R1(config-if)# ip ospf cost 10
R1(config-if)# end
R1#
```

Test Failover to Backup Route

What happens if the link between R1 and R2 goes down? You can simulate that by shutting down the Gigabit Ethernet 0/0/0 interface and verifying the routing table is updated to use R3 as the next-hop router. Notice that R1 can now reach the 10.1.1.4/30 network through R3 with a cost value of 50.

```
R1# show ip route ospf | begin 10
  10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
O       10.1.1.4/30 [110/50] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
O       10.1.1.8/30 [110/40] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
O       10.10.2.0/24 [110/50] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
O       10.10.3.0/24 [110/40] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
R1#
```


Hello Packet Intervals

- OSPFv2 Hello packets are transmitted to multicast address 224.0.0.5 (all OSPF routers) every 10 seconds. This is the default timer value on multiaccess and point-to-point networks.

Note: Hello packets are not sent on interfaces set to passive by the **passive-interface** command.

- The Dead interval is the period that the router waits to receive a Hello packet before declaring the neighbor down. If the Dead interval expires before the routers receive a Hello packet, OSPF removes that neighbor from its link-state database (LSDB). The router floods the LSDB with information about the down neighbor out all OSPF-enabled interfaces. Cisco uses a default of 4 times the Hello interval. This is 40 seconds on multiaccess and point-to-point networks.

Verify Hello and Dead Intervals

- The OSPF Hello and Dead intervals are configurable on a per-interface basis.
- The OSPF intervals must match or a neighbor adjacency does not occur.
- To verify the currently configured OSPFv2 interface intervals, use the **show ip ospf interface** command. The Gigabit Ethernet 0/0/0 Hello and Dead intervals are set to the default 10 seconds and 40 seconds respectively.

```
R1# show ip ospf interface g0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.5/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 10
  Topology-MTID      Cost      Disabled  Shutdown  Topology Name
    0                10        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
    oob-resync timeout 40
(output omitted)
```

Verify Hello and Dead Intervals (Cont.)

Use the **show ip ospf neighbor** command to see the Dead Time counting down from 40 seconds. By default, this value is refreshed every 10 seconds when R1 receives a Hello from the neighbor.

```
R1# show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time      Address        Interface
3.3.3.3        0     FULL/ -         00:00:35      10.1.1.13      GigabitEthernet0/0/1
2.2.2.2        0     FULL/ -         00:00:31      10.1.1.6       GigabitEthernet0/0/0
R1#
```

Modify OSPFv2 Intervals

- It may be desirable to change the OSPF timers so that routers detect network failures in less time. Doing this increases traffic, but sometimes the need for quick convergence is more important than the extra traffic it creates.

Note: The default Hello and Dead intervals are based on best practices and should only be altered in rare situations.

- OSPFv2 Hello and Dead intervals can be modified manually using the following interface configuration mode commands:

```
Router(config-if)# ip ospf hello-interval seconds
Router(config-if)# ip ospf dead-interval seconds
```

- Use the **no ip ospf hello-interval** and **no ip ospf dead-interval** commands to reset the intervals to their default.

Modify OSPFv2 Intervals (Cont.)

- In the example, the Hello interval for the link between R1 and R2 is changed to 5 seconds. The Cisco IOS automatically modifies the Dead interval to four times the Hello interval. However, you can document the new Dead interval in the configuration by manually setting it to 20 seconds, as shown.
- When the Dead Timer on R1 expires, R1 and R2 lose adjacency. R1 and R2 must be configured with the same Hello interval. Use the **show ip ospf neighbor** command on R1 to verify the neighbor adjacencies.

```
R1(config)# interface g0/0/0
R1(config-if)# ip ospf hello-interval 5
R1(config-if)# ip ospf dead-interval 20
R1(config-if)#
*Jun 7 04:56:07.571: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on GigabitEthernet0/0/0
from FULL to DOWN, Neighbor Down: Dead timer expired
R1(config-if)# end
R1# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time           Address             Interface
3.3.3.3          0     FULL/ -         00:00:37           10.1.1.13          GigabitEthernet0/0/1
R1#
```

2.5 Default Route Propagation

Propagate a Default Static Route in OSPFv2

To propagate a default route, the edge router must be configured with the following:

- A default static route using the **ip route 0.0.0.0 0.0.0.0 [next-hop-address | exit-intf]** command.
- The **default-information originate** router configuration command. This instructs R2 to be the source of the default route information and propagate the default static route in OSPF updates.

In the example, R2 is configured with a loopback to simulate a connection to the internet. A default route is configured and propagated to all other OSPF routers in the routing domain.

Note: When configuring static routes, best practice is to use the next-hop IP address. However, when simulating a connection to the internet, there is no next-hop IP address. Therefore, we use the *exit-intf* argument.

```
R2(config)# interface lo1
R2(config-if)# ip address 64.100.0.1 255.255.255.252
R2(config-if)# exit
R2(config)# ip route 0.0.0.0 0.0.0.0 loopback 1
%Default route without gateway, if not a point-to-point interface, may impact performance
R2(config)# router ospf 10
R2(config-router)# default-information originate
R2(config-router)# end
R2#
```

Verify the Propagated Default Route

- You can verify the default route settings on R2 using the **show ip route** command. You can also verify that R1 and R3 received a default route.
- Notice that the route source on R1 is **O*E2**, signifying that it was learned using OSPFv2. The asterisk identifies this as a good candidate for the default route. The E2 designation identifies that it is an external route. The meaning of E1 and E2 is beyond the scope of this module.

```
R2# show ip route | begin Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
S*    0.0.0.0/0 is directly connected, Loopback1
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
(output omitted)
```

```
R1# show ip route | begin Gateway
Gateway of last resort is 10.1.1.6 to network 0.0.0.0
O*E2  0.0.0.0/0 [110/1] via 10.1.1.6, 00:11:08, GigabitEthernet0/0/0
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
(output omitted)
```


2.6 Verify Single-Area OSPFv2

Verify OSPF Neighbors

After configuring single-area OSPFv2, you will need to verify your configurations. The following two commands are particularly useful for verifying routing:

- **show ip interface brief** - This verifies that the desired interfaces are active with correct IP addressing.
- **show ip route**- This verifies that the routing table contains all the expected routes.

Additional commands for determining that OSPF is operating as expected include the following:

- **show ip ospf neighbor**
- **show ip protocols**
- **show ip ospf**
- **show ip ospf interface**

Verify OSPF Neighbors (Cont.)

- Use the **show ip ospf neighbor** command to verify that the router has formed an adjacency with its neighboring routers. If the router ID of the neighboring router is not displayed, or if it does not show as being in a state of FULL, the two routers have not formed an OSPFv2 adjacency.

Note: A non-DR or BDR router that has a neighbor relationship with another non-DR or BDR router will display a two-way adjacency instead of full.

- The following command output displays the neighbor table of R1.

```
R1# show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time           Address             Interface
3.3.3.3        0     FULL/ -         00:00:35           10.1.1.13          GigabitEthernet0/0/1
2.2.2.2        0     FULL/ -         00:00:31           10.1.1.6           GigabitEthernet0/0/0
R1#
```

Verify OSPF Neighbors (Cont.)

Two routers may not form an OSPFv2 adjacency if the following occurs:

- The subnet masks do not match, causing the routers to be on separate networks.
- The OSPFv2 Hello or Dead Timers do not match.
- The OSPFv2 Network Types do not match.
- There is a missing or incorrect OSPFv2 network command.

Verify OSPF Protocol Settings

The **show ip protocols** command is a quick way to verify vital OSPF configuration information, as shown in the command output. This includes the OSPFv2 process ID, the router ID, interfaces explicitly configured to advertise OSPF routes, the neighbors the router is receiving updates from, and the default administrative distance, which is 110 for OSPF.

```
R1# show ip protocols
*** IP Routing is NSF aware ***
(output omitted)
Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
    GigabitEthernet0/0/1
    GigabitEthernet0/0/0
  Routing Information Sources:
    Gateway         Distance         Last Update
    3.3.3.3          110              00:09:30
    2.2.2.2          110              00:09:58
  Distance: (default is 110)
R1#
```

Verify OSPF Process Information

The **show ip ospf** command can also be used to examine the OSPFv2 process ID and router ID, as shown in the command output. This command displays the OSPFv2 area information and the last time the SPF algorithm was executed.

```
R1# show ip ospf
Routing Process "ospf 10" with ID 1.1.1.1
Start time: 00:01:47.390, Time elapsed: 00:12:32.320
(output omitted)
Cisco NSF helper support enabled
Reference bandwidth unit is 10000 mbps
  Area BACKBONE (0)
    Number of interfaces in this area is 3
    Area has no authentication
    SPF algorithm last executed 00:11:31.231 ago
    SPF algorithm executed 4 times
    Area ranges are
    Number of LSA 3. Checksum Sum 0x00E77E
    Number of opaque link LSA 0. Checksum Sum 0x000000
    Number of DCbitless LSA 0 Number of indication LSA 0
    Number of DoNotAge LSA 0 Flood list length 0

R1#
```

Verify OSPF Interface Settings

The **show ip ospf interface** command provides a detailed list for every OSPFv2-enabled interface. Specify an interface to display the settings of just that interface. This command shows the process ID, the local router ID, the type of network, OSPF cost, DR and BDR information on multiaccess links (not shown), and adjacent neighbors.

```
R1# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.5/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 10

<output omitted>

  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
R1#
```

Verify OSPF Interface Settings (Cont.)

To get a quick summary of OSPFv2-enabled interfaces, use the **show ip ospf interface brief** command, as shown in the command output. This command is useful for seeing important information including:

- Interfaces are participating in OSPF
- Networks that are being advertised (IP Address/Mask)
- Cost of each link
- Network state
- Number of neighbors on each link

```
R1# show ip ospf interface brief
Interface      PID  Area      IP Address/Mask      Cost      State      Nbrs F/C
Lo0            10   0         10.10.1.1/24        10        P2P        0/0
Gi0/0/1       10   0         10.1.1.14/30        30        P2P        1/1
Gi0/0/0       10   0         10.1.1.5/30         10        P2P        1/1
R1#
```


Module Summary

What Did I Learn In This Module?

- OSPFv2 is enabled using the **router ospf process-id** global configuration mode command. The process-id value represents a number between 1 and 65,535 and is selected by the network administrator.
- An OSPF router ID is a 32-bit value, represented as an IPv4 address. The router ID is used by an OSPF-enabled router to synchronize OSPF databases and participate in the election of the DR and BDR.
- Cisco routers derive the router ID based on one of three criteria, in this order: 1) Router ID is explicitly configured using the OSPF **router-id rid** router configuration mode command, 2) the router chooses the highest IPv4 address of any of configured loopback interfaces or 3) the router chooses the highest active IPv4 address of any of its physical interfaces.
- The basic syntax for the **network** command is **network network-address wildcard-mask area area-id**. Any interfaces on a router that match the network address in the **network** command can send and receive OSPF packets.
- When configuring single-area OSPFv2, the **network** command must be configured with the same area-id value on all routers. The wildcard mask is typically the inverse of the subnet mask configured on that interface, but could also be a quad zero wildcard mask, which would specify the exact interface.

What Did I Learn In This Module? (Cont.)

- To configure OSPF directly on the interface, use the **ip ospf interface** configuration mode command. The syntax is **ip ospf process-id area area-id**.
- Use the **passive-interface router** configuration mode command to stop transmitting routing messages through a router interface, but still allow that network to be advertised to other routers.
- The DR/ BDR election process is unnecessary as there can only be two routers on the point-to-point network between R1 and R2. Use the interface configuration command **ip ospf network point-to-point** on all interfaces where you want to disable the DR/BDR election process.
- By default, loopback interfaces are advertised as /32 host routes. To simulate a real LAN, the Loopback 0 interface is configured as a point-to-point network.
- OSPF Network Types
- The DR is responsible for collecting and distributing LSAs . The DR uses the multicast IPv4 address 224.0.0.5 which is meant for all OSPF routers. If the DR stops producing Hello packets, the BDR promotes itself and assumes the role of DR. All other routers become a DROTHER.
- DROTHERs use the multiaccess address 224.0.0.6 (all designated routers) to send OSPF packets to the DR and BDR. Only the DR and BDR listen for 224.0.0.6.
- To verify the roles of the OSPFv2 router, use the **show ip ospf interface** command.

What Did I Learn In This Module? (Cont.)

- To verify the OSPFv2 adjacencies, use the **show ip ospf neighbor** command. The state of neighbors in multiaccess networks can be: FULL/DROTHER, FULL/DR, FULL/BDR, or 2-WAY/DROTHER.
- The OSPF DR and BDR election decision is based on the router with the highest interface priority as the DR. The router with the second highest interface priority is elected as the BDR. If the interface priorities are equal, then the router with the highest router ID is elected the DR. The router with the second highest router ID is the BDR.
- The interface priority can be configured to be any number between 0 – 255. If the interface priority value is set to 0, that interface cannot be elected as DR nor BDR. The default priority of multiaccess broadcast interfaces is 1.
- OSPF DR and BDR elections are not pre-emptive. If the DR fails, the BDR is automatically promoted to DR.
- To set the priority of an interface, use the command **ip ospf priority value**, where value is 0 to 255. If the value is 0, the router will not become a DR or BDR. If the value is 1 to 255, then the router with the higher priority value will more likely become the DR or BDR on the interface.
- OSPF uses cost as a metric. A lower cost indicates a better path than a higher cost.
- The formula used to calculate the OSPF cost is: $\text{Cost} = \text{reference bandwidth} / \text{interface bandwidth}$.

What Did I Learn In This Module? (Cont.)

- Because the OSPF cost value must be an integer, FastEthernet, Gigabit Ethernet, and 10 GigE interfaces share the same cost. To correct this situation, you can adjust the reference bandwidth with the **auto-cost reference-bandwidth** command on each OSPF router, or manually set the OSPF cost value with the **ip ospf cost** command.
- The cost of an OSPF route is the accumulated value from one router to the destination network. OSPF cost values can be manipulated to influence the route chosen by OSPF. To change the cost value report by the local OSPF router to other OSPF routers, use the interface configuration command **ip ospf cost value**.
- If the Dead interval expires before the routers receive a Hello packet, OSPF removes that neighbor from its link-state database (LSDB). The router floods the LSDB with information about the down neighbor out all OSPF-enabled interfaces.
- Cisco uses a default of 4 times the Hello interval or 40 seconds on multiaccess and point-to-point networks. To verify the OSPFv2 interface intervals, use the **show ip ospf interface** command.
- OSPFv2 Hello and Dead intervals can be modified manually using the following interface configuration mode commands: **ip ospf hello-interval** and **ip ospf dead-interval**.

What Did I Learn In This Module? (Cont.)

- In OSPF terminology, the router located between an OSPF routing domain and a non-OSPF network is called the ASBR. To propagate a default route, the ASBR must be configured with a default static route using the **ip route 0.0.0.0 0.0.0.0 [next-hop-address | exit-intf]** command, and the **default-information originate router** configuration command.
- Verify the default route settings on the ASBR using the **show ip route** command.
- Additional commands for determining that OSPF is operating as expected include: **show ip ospf neighbor**, **show ip protocols**, **show ip ospf**, and **show ip ospf interface**.
- Use the **show ip ospf neighbor** command to verify that the router has formed an adjacency with its neighboring routers.