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.

2.1 OSPF Router ID

OSPF Router ID

OSPF Reference Topology

The figure shows the topology used for configuring OSPFv2 in this module. The routers in the topology have a starting configuration, including interface addresses. There is currently no static routing or dynamic routing configured on any of the routers. All interfaces on R1, R2, and R3 (except the loopback 1 on R2) are within the OSPF backbone area. The ISP router is used as the gateway to the internet of the routing domain.

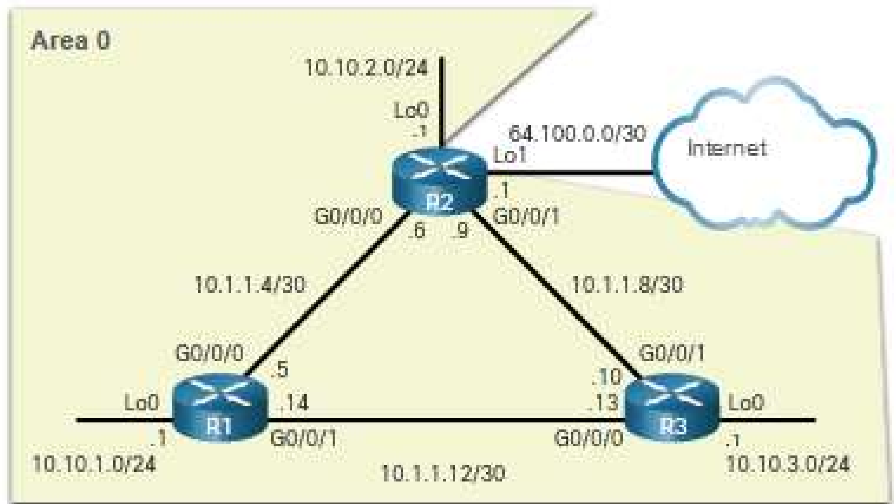

OSPF Router ID
Router Configuration Mode for OSPF

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.

```
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)#
```

OSPF Router ID

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 multiaccess 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).

OSPF Router ID

Router ID Order of Precedence

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

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 of configured loopback interfaces.
3. The router chooses the highest active IPv4 address of any of its physical interfaces.

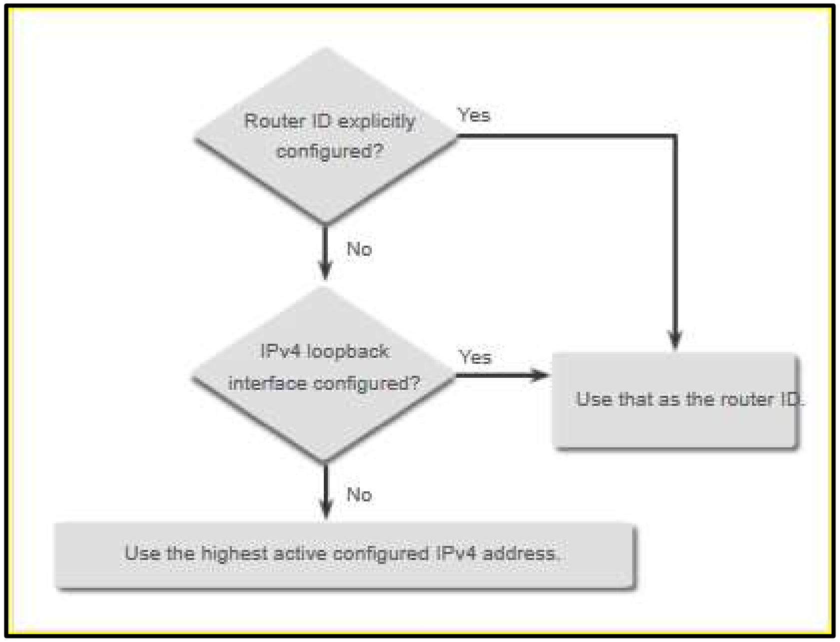

OSPF Router ID

Configure a Loopback Interface as the Router ID

Instead of relying on physical interface, the router ID can 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.

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#
```

OSPF Router ID

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#
```

OSPF Router ID

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 or after 30 min.
- 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 Point-to-Point OSPF Networks

Point-to-Point OSPF Networks

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

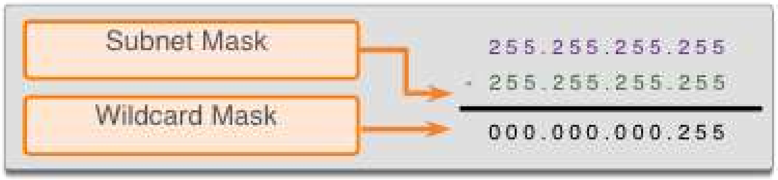

Calculating a Wildcard Mask for /26

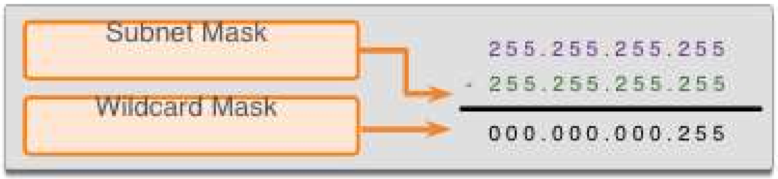

Point-to-Point OSPF Networks

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)#
```

Point-to-Point OSPF Networks

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

```
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
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
R1(config-if)#
```

Point-to-Point OSPF Networks

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.

Point-to-Point OSPF Networks

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#
```

Point-to-Point OSPF Networks
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
```


Point-to-Point OSPF Networks

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
```

Point-to-Point OSPF Networks

Packet Tracer - Point-to-Point Single-Area OSPFv2 Configuration

In this Packet Tracer activity, you will do the following:

- Explicitly configure router IDs.
- Configure the **network** command on R1 using wildcard mask based on the subnet mask.
- Configure the **network** command on R2 using a quad-zero wildcard mask.
- Configure the **ip ospf** interface command on R3.
- Configure passive interfaces.
- Verify OSPF operation using the **show ip protocols** and **show ip route** commands.

2.3 Multiaccess OSPF Networks

Multiaccess OSPF Networks

OPSF 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.

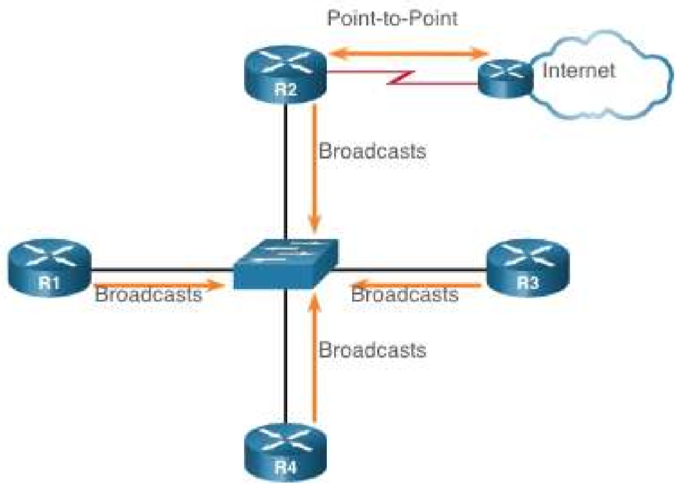

Multiaccess OSPF Networks

OPSF 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.

OPSF 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.

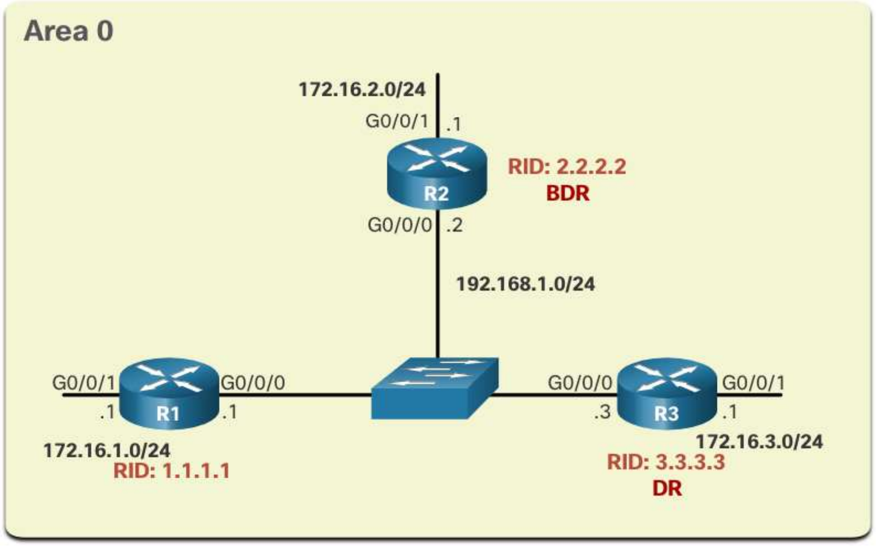

Multiaccess OSPF Networks

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#
```

Multiaccess OSPF Networks

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)

```
R3# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.1.3/24, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 1
  (output omitted)
  Transmit Delay is 1 sec, State DR, 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 2.2.2.2 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
R3#
```

Multiaccess OSPF Networks
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#
```

Multiaccess OSPF Networks

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.

Multiaccess OSPF Networks
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.

Multiaccess OSPF Networks

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
```


Multiaccess OSPF Networks

Packet Tracer - Determine the DR and BDR

In this activity, you will complete the following:

- Examine DR and BDR roles and watch the roles change when there is a change in the network.
- Modify the priority to control the roles and force a new election.
- Verify routers are filling the desired roles