

CCNA 3 v7.0 Curriculum: Module 2 – Single-Area OSPFv2 Configuration

 itexamanswers.net/ccna-3-v7-0-curriculum-module-2-single-area-ospfv2-configuration.html

April 7, 2020

2.0. Introduction

2.0.1. Why should I take this module?

Welcome to Single-Area OSPFv2 Configuration!

Now that you know about single-area OSPFv2, you can probably think of all the ways it could benefit your own network. As a link-state protocol, OSPF is designed to not only find the fastest available route, it is designed to create fast, available routes. If you prefer a bit more control over some areas of your network, OSPF gives you several ways to manually override the DR election process and create your own preferred routes. With OSPF, your network can combine the automated processes with your own choices to make a network that you could troubleshoot in your sleep! You know you want to learn how to do this!

2.0.2. What will I learn to do in this module?

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

2.1.1. OSPF Reference Topology

To get you started, this topic discusses the foundation on which OSPF bases its entire process, the OSPF router ID.

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.

Note: In this topology the loopback interface is used to simulate the WAN link to the Internet and a LAN connected to each router. This is done to allow this topology to be duplicated for demonstration purposes on routers that only have two Gigabit Ethernet interfaces.



2.1.2. Router Configuration Mode for OSPF

OSPFv2 is enabled using the **router ospf process-id** global configuration mode command, as shown in the command window for R1. 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, which means that it does not have to be the same value on the other OSPF routers to establish adjacencies with those neighbors. It is considered best practice to use the same *process-id* on all OSPF routers.

After entering the **router ospf process-id** command, the router enters router configuration mode, as indicated by the **R1(config-router)#** prompt. Enter a question mark (?), to view all the commands available in this mode. The list of commands shown here has been altered to display only the commands that are relevant to this module.

```

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)#

```

2.1.3. Router IDs

An OSPF router ID is a 32-bit value, represented as an IPv4 address. The router ID is used to uniquely identify an OSPF router. All OSPF packets include the router ID of the originating router. Every router requires a router ID to participate in an OSPF domain. The router ID 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:

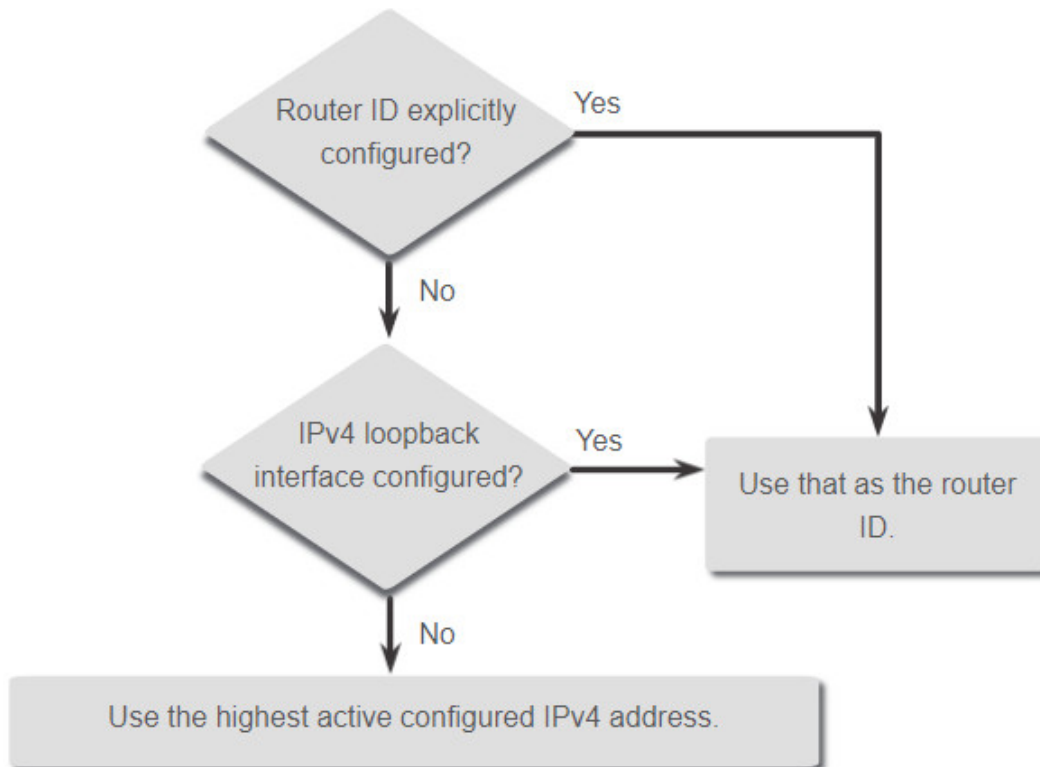
- **Participate in the synchronization of OSPF databases** – During the Exchange State, the router with the highest router ID will send their database descriptor (DBD) packets first.
- **Participate in the election of the designated router (DR)** – In a multiaccess LAN environment, the router with the highest router ID is elected the DR. The routing device with the second highest router ID is elected the backup designated router (BDR).

Note: The DR and BDR election process is discussed in more detail later in this module.

2.1.4. Router ID Order of Precedence

But how does the router determine the router ID? As illustrated in the figure, 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. The *rid* value is any 32-bit value expressed as an IPv4 address. This is the recommended method to assign a router ID.
2. If the router ID is not explicitly configured, the router chooses the highest IPv4 address of any of configured loopback interfaces. This is the next best alternative to assigning a router ID.
3. If no loopback interfaces are configured, then the router chooses the highest active IPv4 address of any of its physical interfaces. This is the least recommended method because it makes it more difficult for administrators to distinguish between specific routers.



2.1.5. Configure a Loopback Interface as the Router ID

In the reference topology, only the physical interfaces are configured and active. The loopback interfaces have not been configured. When OSPF routing is enabled on the router, the routers would pick the following highest active configured IPv4 address as the router ID.

- R1: 10.1.1.14 (G0/0/1)
- R2: 10.1.1.9 (G0/0/1)
- R3: 10.1.1.13 (G0/0/0)

Note: OSPF does not need to be enabled on an interface for that interface to be chosen 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.

The example shows how to configure a loopback interface on R1. Assuming the router ID was not explicitly configured or previously learned, R1 will use IPv4 address 1.1.1.1 as its router ID. Assume R1 has not yet learned a router ID.

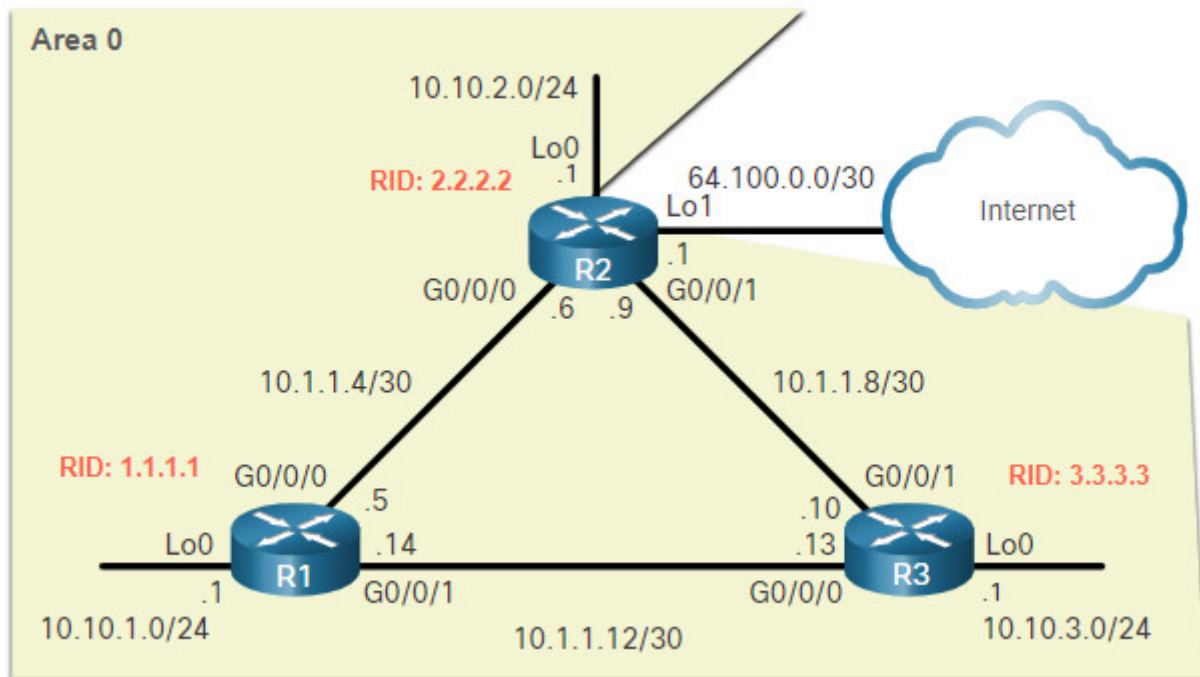
```

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#
  
```

2.1.6. Explicitly Configure a Router ID

In the figure, the topology has been updated to show the router ID for each router:

- 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#
```

2.1.7. 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.

In example for R1, the configured router ID has been removed and the router reloaded. Notice that the current router ID is 10.10.1.1, which is the Loopback 0 IPv4 address. The router ID should be 1.1.1.1. Therefore, R1 is configured with the command **router-id 1.1.1.1**.

Notice how an informational message appears stating that the OSPF process must be cleared or that the router must be reloaded. The reason is because R1 already has adjacencies with other neighbors using the router ID 10.10.1.1. Those adjacencies must be renegotiated using the new router ID 1.1.1.1. Use the **clear ip ospf process** command to reset the adjacencies. You can then verify that R1 is using the new router ID command with the **show ip protocols** command piped to display only the router ID section.

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 01:09:46.975: %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 01:09:46.981: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/1 from LOADING to FULL, Loading Done
*Jun  6 01:09:46.981: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R1# show ip protocols | include Router ID
  Router ID 1.1.1.1
R1#
```

Note: The **router-id** command is the preferred method. However, some older versions of the IOS do not recognize the **router-id** command; therefore, the best way to set the router ID on those routers is by using a loopback interface.

2.1.8. Syntax Checker – Configure R2 and R3 Router IDs

Use the Syntax Checker to configure R2 and R3 with router IDs.

You are currently logged into R2:

- Enter OSPF router configuration mode using process ID 10
- Assign a router ID of 2.2.2.2
- Return to privileged EXEC mode

```
R2(config)#router ospf 10
R2(config-router)#router-id 2.2.2.2
R2(config-router)#end
R2#
```

*Mar 25 20:03:56.863: %SYS-5-CONFIG_I: Configured from console by console
Enter the **show ip protocols** command to verify the router ID.

```
R2#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 2.2.2.2
  Number of areas in this router is 0. 0 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance         Last Update
  Distance: (default is 110)
```

R2#

You are now on R3:

- Enter OSPF router configuration mode using process ID 10
- Assign a router ID of 3.3.3.3
- Return to privileged EXEC mode.

```
R3(config)#router ospf 10
R3(config-router)#router-id 3.3.3.3
R3(config-router)#end
R3#
```

*Mar 25 20:11:05.415: %SYS-5-CONFIG_I: Configured from console by console
Enter the **show ip protocols** command to verify the router ID.

```
R3#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 3.3.3.3
  Number of areas in this router is 0. 0 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance         Last Update
  Distance: (default is 110)
```

R3#

You have successfully assigned the router IDs to R2 and R3.

2.2.1. The network Command Syntax

One type of network that uses OSPF is the point-to-point network. 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, as we will see later.

Both commands are used to determine which interfaces participate in the routing process for an OSPFv2 area. 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 the network address in the **network** 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.


2.2.2. The Wildcard Mask

The wildcard mask is typically the inverse of the subnet mask configured on that interface. In a subnet mask, binary 1 is equal to a match and binary 0 is not a match. In a wildcard mask, the reverse is true, as shown in here:


- **Wildcard mask bit 0** – Matches the corresponding bit value in the address.
- **Wildcard mask bit 1** – Ignores the corresponding bit value in the address.

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

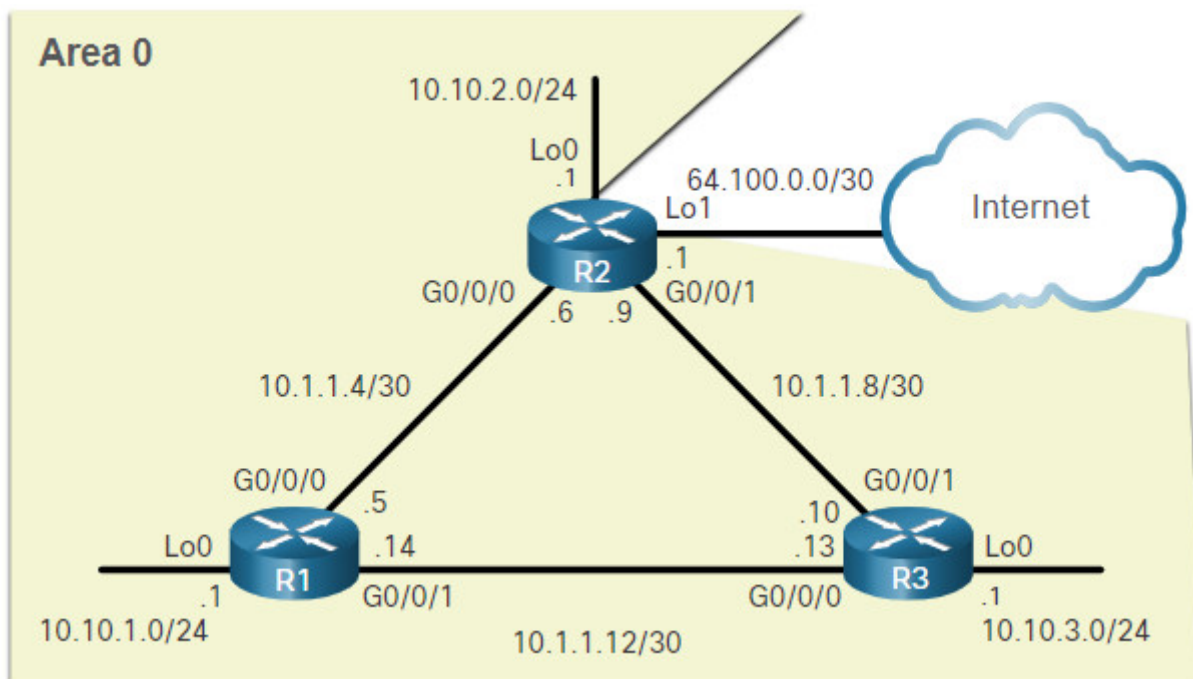


Calculating a Wildcard Mask for /26



2.2.4. 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. The figure shows the reference topology.



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.

```
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)#
```

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.

As an alternative, the second example shows how 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. As a result, the OSPFv2 process will advertise the network that is on this interface (10.1.1.4/30).

```
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)#
```

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.

2.2.5. Syntax Checker – Configure R2 and R3 Using the network Command

Use the Syntax Checker to advertise the networks connected to R2 and R3.

Note: While completing the syntax checker, observe the informational messages describing the adjacency between R1 (1.1.1.1) and R2 (2.2.2.2). The IPv4 addressing scheme used for the router ID makes it easy to identify the neighbor.

You are currently logged into R2:

- Enter OSPF router configuration mode using process ID 10
- Configure the R2 router ID of 2.2.2.2

```
R2(config)#router ospf 10
```

```
R2(config-router)#router-id 2.2.2.2
```

```
R2(config-router)#
```

Advertise the networks connected to R2 with the appropriate wildcard mask using area 0. Configure the networks in the following order:

1. 10.10.2.0/24
2. 10.1.1.4/30
3. 10.1.1.8/30

```
R2(config-router)#network 10.10.2.0 0.0.0.255 area 0
```

```
R2(config-router)#network 10.1.1.4 0.0.0.3 area 0
```

```
R2(config-router)#network 10.1.1.8 0.0.0.3 area 0
```

```
R2(config-router)#
```

```
\*Mar 25 21:19:21.938: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on  
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
```

You are now configuring R3:

- Enter OSPF router configuration mode using process ID 10
- Configure the R3 router ID.
- Use the network statement to enable OSPF based on the interface address and quad zero wildcard mask for area 0.
- Return to privileged EXEC mode when complete

```
R3(config)#router ospf 10
```

```
R3(config-router)#router-id 3.3.3.3
```

```
R3(config-router)#
```

Use the network statement to enable OSPF based on the interface address and quad zero wildcard mask for area 0. Configure the interfaces in the following order:

1. 10.10.3.1
2. 10.1.1.10
3. 10.1.1.13

```
R3(config-router)#network 10.10.3.1 0.0.0.0 area 0
```

```
R3(config-router)#network 10.1.1.10 0.0.0.0 area 0
```

```
R3(config-router)#network 10.1.1.13 0.0.0.0 area 0
```

```
R3(config-router)#
```

```
\*Mar 26 14:00:55.183: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on  
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
```

```
\*Mar 26 14:00:55.243: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on  
GigabitEthernet0/0/1 from LOADING to FULL, Loading Done
```

```
R3#
```

You successfully advertised the OSPF networks on R2 and R3.

2.2.6. Configure OSPF Using the ip ospf Command

You can also configure OSPF directly on the interface instead of using the **network** 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
```

For R1, remove the network commands by using the **no** form of the **network** commands. And then go to each interface and configure the **ip ospf** command, as shown in the command window.

```
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)#
```

2.2.7. Syntax Checker – Configure R2 and R3 Using the ip ospf Command

Use the Syntax Checker to advertise the networks by configuring the interfaces for OSPF on R2 and R3.

You are currently logged into R2. The network commands are already removed. Configure OSPF routing using process ID 10, in area 0, on each interface, in that order. Use the following shortened interface names:

```
- lo0
- g0/0/0
- g0/0/1
R2(config)#interface lo0
R2(config-if)#ip ospf 10 area 0
R2(config-if)#interface g0/0/0
R2(config-if)#ip ospf 10 area 0
R2(config-if)#interface g0/0/1
R2(config-if)#ip ospf 10 area 0
\*Mar 25 21:19:21.938: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
```

You are now logged into R3. The network commands are already removed. Configure OSPF routing using process ID 10, in area 0, on each interface, in that order. Use the following shortened interface names:

```
- lo0
- g0/0/0
- g0/0/1
R3(config)#interface lo0
R2(config-if)#ip ospf 10 area 0
R2(config-if)#interface g0/0/0
R2(config-if)#ip ospf 10 area 0
R2(config-if)#interface g0/0/1
R2(config-if)#ip ospf 10 area 0
\*Mar 26 14:00:55.183: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
\*Mar 26 14:00:55.243: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/1 from LOADING to FULL, Loading Done
R3(config-router)#
You successfully configured the interfaces to advertise the OSPF networks.
```


2.2.8. Passive Interface

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

Refer to the topology in the figure. OSPFv2 messages are forwarded out the three loopback interfaces even though no OSPFv2 neighbor exists on these simulated LANs. In a production network, these loopbacks would be physical interfaces to networks with users and traffic. Sending out unneeded messages on a LAN affects the network in three ways, as follows:

- **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.



2.2.9. 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 configuration example identifies the R1 Loopback 0/0/0 interface as passive.

The **show ip protocols** command is then used to verify that the Loopback 0 interface is listed as passive. The interface is still listed under the heading, "Routing on Interfaces Configured Explicitly (Area 0)", which means that this network is still included as a route entry in OSPFv2 updates that are sent to R2 and R3.

```

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#

```

2.2.10. Syntax Checker – Configure R2 and R3 Passive Interfaces

Use the Syntax Checker to configure the Loopback interfaces on R2 as a passive. As an alternative, all interfaces can be made passive using the **passive-interface default** command. Interfaces that should not be passive can be re-enabled using the **no passive-interface** command. Configure R3 with the **passive-interface default** command and then re-enable the Gigabit Ethernet interfaces.

You are currently logged into R2.

- Enter OSPF router configuration mode using process ID 10.
- Configure the Loopback interface as passive using the shortened interface name lo0.
- Return to privileged EXEC mode.
- Verify the OSPF settings with the show ip protocols command.

```

R2(config)#router ospf 10
R2(config-router)#passive-interface lo0
R2(config-router)#end
\*May 23 20:27:20.718: %SYS-5-CONFIG\I: Configured from console by console
R2#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 2.2.2.2
  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          02:07:48
    1.1.1.1          110          02:34:53
  Distance: (default is 110)
R2#
R3(config)#router ospf 10
R3(config-router)#passive-interface default
\*Jun  5 23:06:46.668: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
\*Jun  5 23:06:46.669: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
R3(config-router)#no passive-interface g0/0/0
\*Jun  5 23:07:07.746: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R3(config-router)#no passive-interface g0/0/1
\*Jun  5 23:07:17.841: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/1 from LOADING to FULL, Loading Done
R3(config-router)#end
\* Jun  5 23:07:35.732: %SYS-5-CONFIG\I: Configured from console by console
R3#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 3.3.3.3
  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):
    Serial0/1/0

```



```

Serial0/1/1
Loopback0
Routing Information Sources:
  Gateway         Distance      Last Update
  1.1.1.1         110          00:00:59
  2.2.2.2         110          00:00:48
Distance: (default is 110)
R3#

```

You successfully configured passive interfaces for R2 and R3.

2.2.11. 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, as shown in the example for G0/0/0 of R1.

```

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
    Hello due in 00:00:08
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2 (Designated Router)
  Suppress hello for 0 neighbor(s)
R1#

```

R1 is the BDR and R2 is the DR. 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. 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. The example below shows this configuration for R1. The OSPF neighbor adjacency status will go down for a few milliseconds.

```

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)# interface GigabitEthernet 0/0/1
R1(config-if)# ip ospf network point-to-point
*Jun  6 00:44:45.532: %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 00:44:45.535: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/1 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
  0                  1          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
    Hello due in 00:00:04
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
R1#

```

Notice that the Gigabit Ethernet 0/0/0 interface now lists the network type as POINT_TO_POINT and that there is no DR or BDR on the link.

2.2.12. Loopbacks and Point-to-Point Networks

We use loopbacks to provide additional interfaces for a variety of purposes. In this case, we are using loopbacks to simulate more networks than the equipment can support. By default, loopback interfaces are advertised as /32 host routes. For example, R1 would advertise the 10.10.1.0/24 network as 10.10.1.1/32 to R2 and R3.

```

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

```

To simulate a real LAN, the Loopback 0 interface is configured as a point-to-point network so that R1 will advertise the full 10.10.1.0/24 network to R2 and R3.

```

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

```

Now R2 receives the more accurate, simulated LAN network address of 10.10.1.0/24.

```
R2# show ip route | include 10.10.1
0      10.10.1.0/24 [110/2] via 10.1.1.5, 00:00:30, GigabitEthernet0/0/0
```

Note: At the time of this writing, Packet Tracer does not support the **ip ospf network point-to-point** command on Gigabit Ethernet interfaces. However, it is supported on Loopback interfaces.

2.2.13. Packet Tracer – Point-to-Point Single-Area OSPFv2 Configuration

In this Packet Tracer activity, you will configure the single-area OSPFv2 with 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.

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

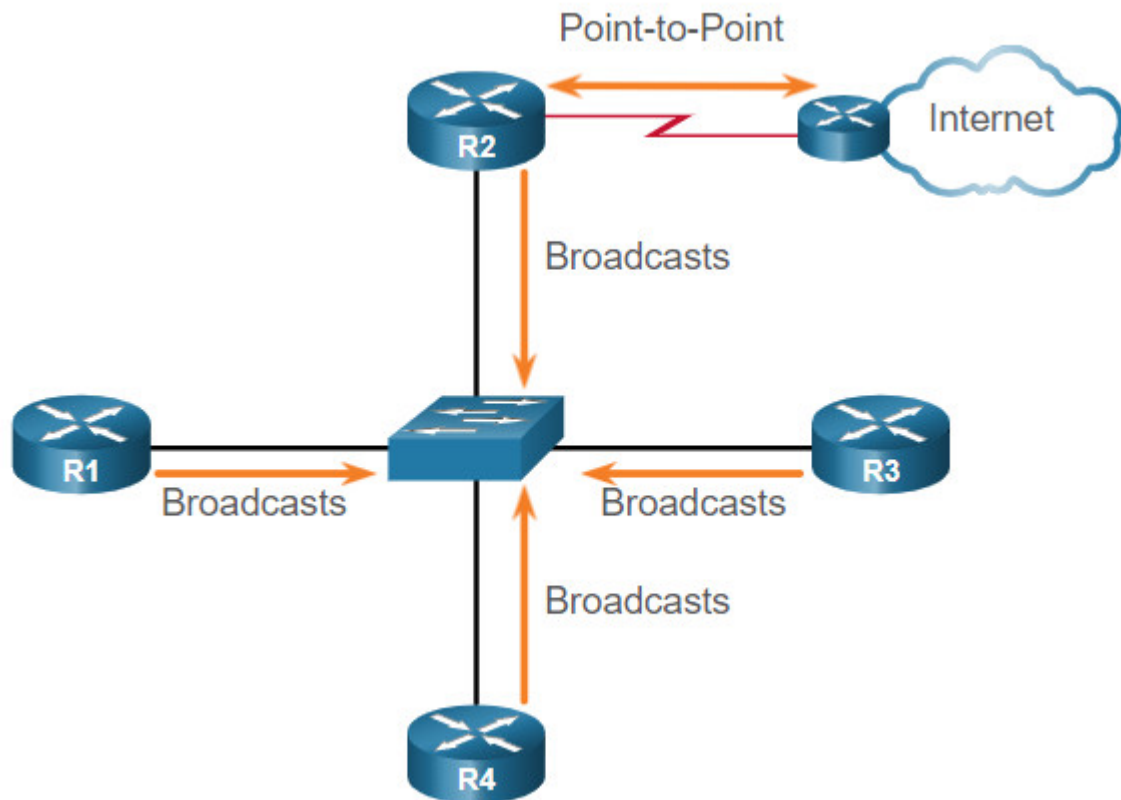
2.3. Multiaccess OSPF Networks

2.3.1. 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 may include additional processes depending on the type of network. The previous topology used point-to-point links between the routers. However, routers can be connected to the same switch to form a multiaccess network, as shown in the figure. Ethernet LANs are the most common example of broadcast multiaccess networks. In broadcast networks, all devices on the network see all broadcast and multicast frames.

OSPF Multiaccess Network



2.3.2. OSPF Designated Router


Recall that, in multiaccess networks, OSPF elects a DR and BDR as a solution to manage the number of adjacencies and the flooding of link-state advertisements (LSAs). 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.

In the figure, R1, R3, and R4 are DROTHERs. Click play to see the animation of R2 acting as DR. Notice that only the DR and the BDR process the LSA sent by R1 using the multicast IPv4 address 224.0.0.6. The DR then sends out the LSA to all OSPF routers using the multicast IPv4 address 224.0.0.5.

Role of the DR




2.3.3. 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. Each router is configured with the indicated IPv4 address on the Gigabit Ethernet 0/0/0 interface.

Because the routers are connected over a common multiaccess network, OSPF has automatically elected a DR and BDR. In this example, 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.

Area 0



2.3.4. Verify OSPF Router Roles

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

Click each button see the output for the show ip ospf interface command on each router.

- [Untitled](#)
- [R2 BDR](#)
- [R3 DR](#)

R1 DROTHER

The output generated by R1 confirms that the following:

1. R1 is not the DR or BDR, but is a DROTHER with a default priority of 1. (Line 7)
2. 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)
3. 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
  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 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
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:07
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 1
  Last flood scan time is 0 msec, maximum is 1 msec
  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#

```

2.3.5. Verify DR/BDR Adjacencies

To verify the OSPFv2 adjacencies, use the **show ip ospf neighbor** command, as shown in the example for R1. 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. For examples, DROTHERs will form a 2-WAY neighbor adjacency with any DROTHERs that join the network. When this happens, the neighbor state displays as 2-WAY/DROTHER.

Click each button see the output for the show ip ospf neighbor command on each router.

- [R1 Adjacencies](#)
- [R2 Adjacencies](#)
- [R3 Adjacencies](#)

R1 Adjacencies

The output generated by R1 confirms that R1 has adjacencies with the following routers:

- R2 with router ID 2.2.2.2 is in a Full state and the role of R2 is BDR.
- R3 with router ID 3.3.3.3 is in a Full state and the role of R3 is DR.

```
R1# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
2.2.2.2          1    FULL/BDR        00:00:31   192.168.1.2   GigabitEthernet0/0/0
3.3.3.3          1    FULL/DR         00:00:39   192.168.1.3   GigabitEthernet0/0/0
R1#
```

2.3.6. Default DR/BDR Election Process

How do the DR and BDR get elected? The OSPF DR and BDR election decision is based on the following criteria, in sequential order:

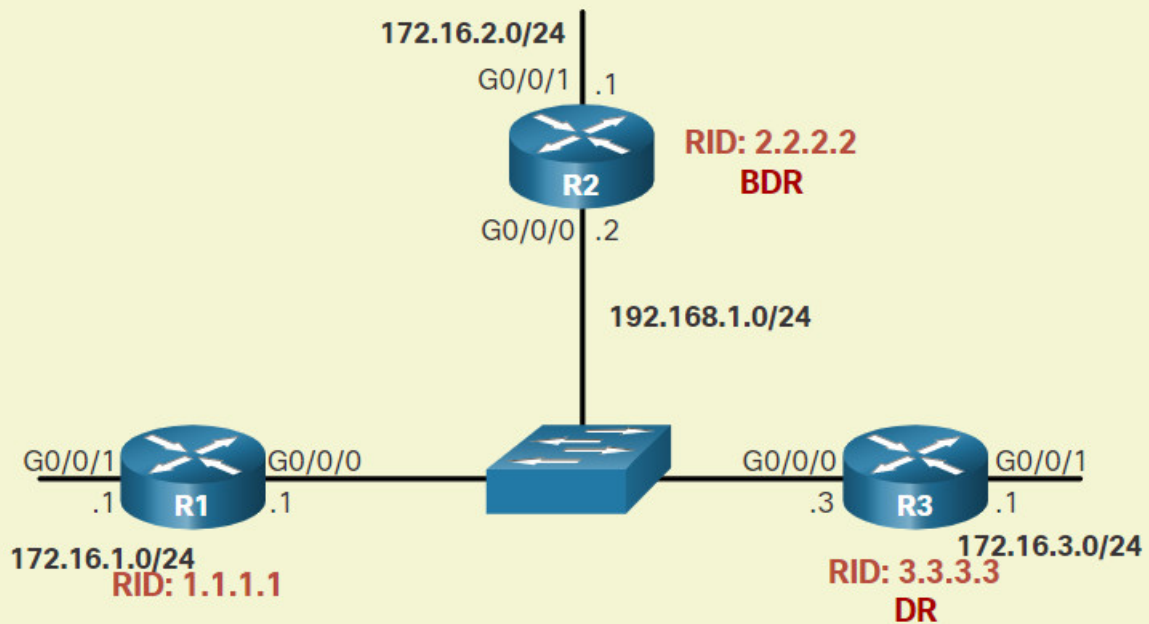
1. The routers in the network elect the router with the highest interface priority as the DR. The router with the second highest interface priority is elected as the BDR. The priority can be configured to be any number between 0 – 255. If the interface priority value is set to 0, that interface cannot be elected as DR nor BDR. The default priority of multiaccess broadcast interfaces is 1. Therefore, unless otherwise configured, all routers have an equal priority value and must rely on another tie breaking method during the DR/BDR election.
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.

Recall that the router ID is determined in one of the following three ways:

1. The router ID can be manually configured.
2. If no router IDs are configured, the router ID is determined by the highest loopback IPv4 address.
3. If no loopback interfaces are configured, the router ID is determined by the highest active IPv4 address.

OSPFv2 Multiaccess Reference Topology

Area 0



In the figure, all Ethernet router interfaces have a default priority of 1. As a result, based on the selection criteria listed above, the OSPF router ID is used to elect the DR and BDR. R3 with the highest router ID becomes the DR; and R2, with the second highest router ID, becomes the BDR.

The DR and BDR election process takes place as soon as the first router with an OSPF-enabled interface is active on the multiaccess network. This can happen when the preconfigured OSPF routers are powered on, or when OSPF is activated on the interface. The election process only takes a few seconds. If all of the routers on the multiaccess network have not finished booting, it is possible that a router with a lower router ID becomes the DR.

OSPF DR and BDR elections are not pre-emptive. If a new router with a higher priority or higher router ID is added to the network after the DR and BDR election, the newly added router does not take over the DR or the BDR role. This is because those roles have already been assigned. The addition of a new router does not initiate a new election process.

2.3.7. DR Failure and Recovery

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

1. The DR fails.
2. The OSPF process on the DR fails or is stopped.
3. The multiaccess interface on the DR fails or is shutdown.

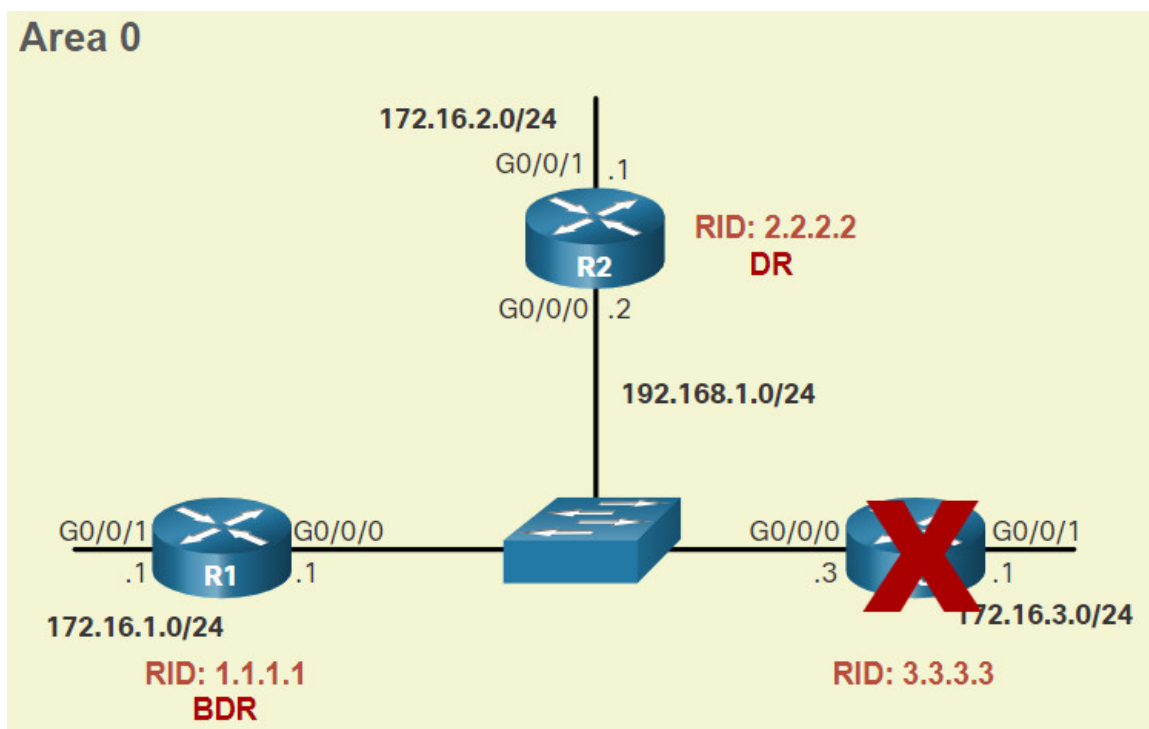
If the DR fails, the BDR is automatically promoted to DR. This is the case even if another DROTHER with a higher priority or router ID is added to the network after the initial DR/BDR election. However, after a BDR is promoted to DR, a new BDR election occurs and the DROTHER with the highest priority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

Click each button for an illustration of various scenarios relating to the DR and BDR election process.

- [R3 Fails](#)
- [R3 Re-Joins Network](#)
- [R4 Joins Network](#)
- [R2 Fails](#)

R3 Fails

In this scenario, the current DR (R3) fails. Therefore, the pre-elected BDR (R2) assumes the role of DR. Subsequently, an election is held to choose a new BDR. Because R1 is the only DROTHER, it is elected as the BDR.



2.3.8. 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. It is possible to configure the router ID to manipulate the DR/BDR election. However, this process only works if there is a stringent plan for setting the router ID on all routers. Configuring the router ID can help control this. However, in large networks this can be cumbersome.

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.

2.3.9. Configure OSPF Priority

In the topology, the **ip ospf priority** command will be used to change the DR and BDR as follows:

- R1 should be the DR and will be configured with a priority of 255.
- R2 should be the BDR and will be left with the default priority of 1.
- R3 should never be a DR or BDR and will be configured with a priority of 0.

Change the R1 G0/0/0 interface priority from 1 to 255.

```
R1(config)# interface GigabitEthernet 0/0/0
R1(config-if)# ip ospf priority 255
R1(config-if)# end
R1#
```

Change the R3 G0/0/0 interface priority from 1 to 0.

```
R3(config)# interface GigabitEthernet 0/0/0
R3(config-if)# ip ospf priority 0
R3(config-if)# end
R3#
```

The following example, shows how to clear the OSPF process on R1. The **clear ip ospf process** command also must be entered on R2 and R3 (not shown). Notice the OSPF state information that is generated.

```
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
*Jun  5 03:47:41.563: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
*Jun  5 03:47:41.569: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
*Jun  5 03:47:41.569: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
```

The output from the **show ip ospf interface g0/0/0** command on R1 confirms that R1 is now the DR with a priority of 255 and identifies the new neighbor adjacencies of R1.

```

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
  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 DR, Priority 255
  Designated Router (ID) 1.1.1.1, Interface address 192.168.1.1
  Backup Designated router (ID) 2.2.2.2, Interface address 192.168.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 1 msec
  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
  Suppress hello for 0 neighbor(s)
R1#

```

2.3.10. Syntax Checker – Configure OSPF Priority

Use the Syntax Checker to configure a different OSPF priority scenario for R1, R2, and R3.

Use the ip ospf priority command to change the DR and BDR as follows:

- R1 should be the BDR and will be configured with a priority of 10.
- R2 should never be a DR or BDR and will be configured with a priority of 0.
- R3 should be the DR and will be left with the default priority of 100.

On all routers, use g0/0/0 for the interface name.

You are logged into R1 in global configuration mode. Configure R1 with a priority of 10.

```
R1(config)#interface g0/0/0
R1(config-if)#ip ospf priority 10
R1(config-if)#
```

You are now logged into R2 in global configuration mode. Configure R2 with a priority of 0.

```
R2(config)#interface g0/0/0
R2(config-if)#ip ospf priority 0
R2(config-if)#
```

You are now logged into R3 in global configuration mode. Configure R3 with a priority of 100.

```
R3(config)#interface g0/0/0
R3(config-if)#ip ospf priority 100
R3(config-if)#
```

You are still logged into R3 in interface configuration mode. Return to privileged EXEC mode. Because R3 is to be the DR, restart the OSPF process on it first.

```
R3(config-if)#end
R3#clear ip ospf process
Reset ALL OSPF processes? \[no\]:y
\*Jun  5 05:29:35.231: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
\*Jun  5 05:29:35.231: %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  5 05:29:35.235: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
\*Jun  5 05:29:44.563: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R3#
```

You are now logged into R1 in privileged EXEC mode. Because R1 is to be the DR, restart the OSPF process on it next.

```
R1#clear ip ospf process
Reset ALL OSPF processes? \[no\]:y
\*Jun  5 05:27:20.691: %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  5 05:27:20.691: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
\*Jun  5 05:27:21.695: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
\*Jun  5 05:27:20.951: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R1#
```

You are now logged into R2 in privileged EXEC mode. R2 is to be DROTHER. Restart the OSPF process.

```
R2#clear ip ospf process
```

```
Reset ALL OSPF processes? \[no\]:y
\*Jun  5 15:37:08.978: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from 2WAY to DOWN, Neighbor Down: Interface down or detached
\*Jun  5 15:37:08.978: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
\*Jun  5 15:37:08.983: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
\*Jun  5 15:37:19.477: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R2#
You are now logged into R1. Use the show ip ospf interface g0/0/0 command to
verify the R1 is the BDR.
```

```
R1#show ip ospf interface g0/0/0
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
  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 10
  Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
  Backup Designated router (ID) 1.1.1.1, Interface address 192.168.1.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:04
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 2
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 2.2.2.2
    Adjacent with neighbor 3.3.3.3 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

```
R1#
You are now logged into R2. Use the show ip ospf interface g0/0/0 command to
verify the R2 is a DROTHER.
```

```
R2#show ip ospf interface g0/0/0
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
  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 DROTHER, Priority 0
  Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
  Backup Designated router (ID) 1.1.1.1, Interface address 192.168.1.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:03
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
```



```
IETF NSF helper support enabled
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
  Adjacent with neighbor 3.3.3.3 (Designated Router)
Suppress hello for 0 neighbor(s)
```

R2#

You are now logged into R3. Use the show ip ospf interface g0/0/0 command to verify the R3 is the DR.

```
R3#show ip ospf interface g0/0/0
```

```
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
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 DR, Priority 100
Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
Backup Designated router (ID) 1.1.1.1, Interface address 192.168.1.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:00
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 1/1/1, flood queue length 0
Next 0x0(0)/0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 3
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
  Adjacent with neighbor 2.2.2.2
Suppress hello for 0 neighbor(s)
```

R3#

You successfully changed the OSPF priority.

2.3.11. 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.

2.3.11 Packet Tracer – Determine the DR and BDR

2.4. Modify Single-Area OSPFv2

2.4.1. Cisco OSPF Cost Metric

Recall that a routing protocol uses a metric to determine the best path of a packet across a network. A metric gives indication of the overhead that is required to send packets across a certain interface. OSPF uses cost as a metric. A lower cost indicates a better path than a higher cost.

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 108 (100,000,000); therefore, the formula is:

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

Refer to the table for a breakdown of the cost calculation. 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.

Default Cisco OSPF Costs

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

2.4.2. 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.

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.

```
Router(config-router)# auto-cost reference-bandwidth Mbps
```

This command must be configured on every router in the OSPF domain. Notice 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.

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).

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

Use the `show ip ospf interface g0/0/0` command to verify the current OSPFv2 cost assigned to the R1 GigabitEthernet 0/0/0 interface. Notice how it displays a cost of 1. Then, after adjusting the reference bandwidth, the cost is now 10. This will allow for scaling to 10 Gigabit Ethernet interfaces in the future without having adjust the reference bandwidth again.

Note: The `auto-cost reference-bandwidth` command must be configured consistently on all routers in the OSPF domain to ensure accurate route calculations.

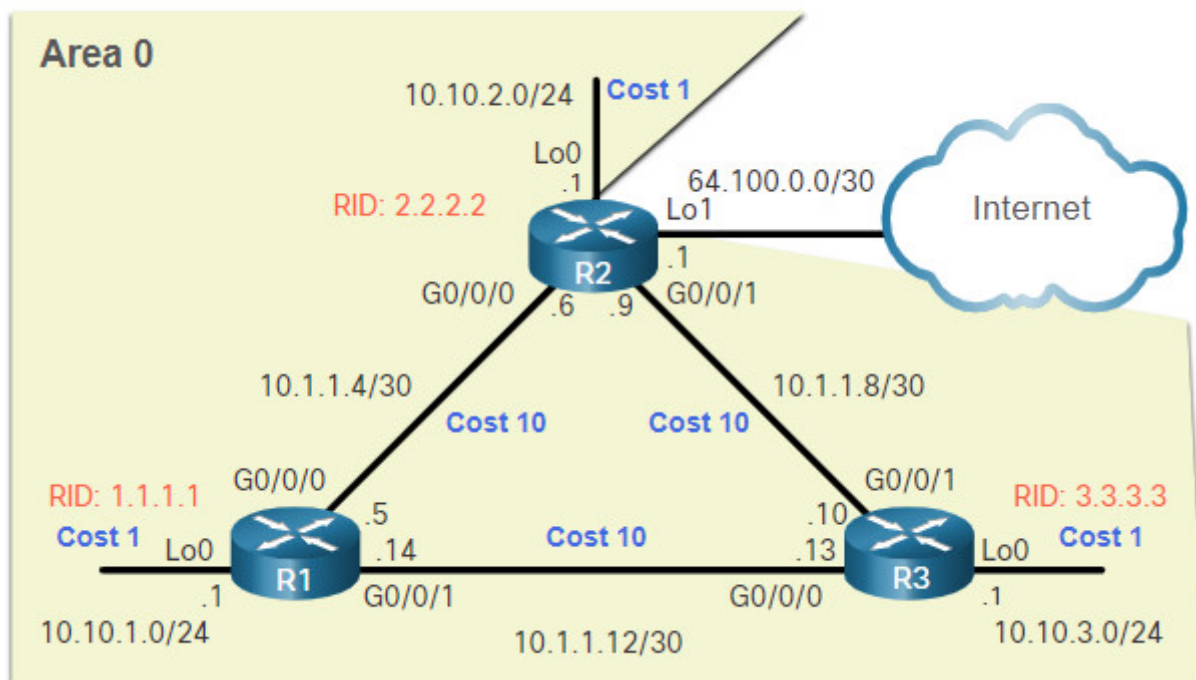
```

R1# show ip ospf interface gigabitethernet0/0/0
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
(output omitted)
R1# config t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router ospf 10
R1(config-router)# auto-cost reference-bandwidth 10000
% OSPF: Reference bandwidth is changed.
      Please ensure reference bandwidth is consistent across all routers.
R1(config-router)# do show ip ospf interface gigabitethernet0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet address is 172.16.1.1/24, Area 0
  Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
(output omitted)

```

2.4.3. OSPF Accumulates Costs

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, as shown in the figure.



Therefore, we 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. $10 + 1 = 11$.

The routing table of R1 in Figure 2 confirms that the metric to reach the R2 LAN is a cost of 11.

```

R1# show ip route | include 10.10.2.0
0       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#

```

2.4.4. Manually Set OSPF Cost Value

OSPF cost values can be manipulated to influence the route chosen by OSPF. For example, in the current configuration, R1 is load balancing to the 10.1.1.8/30 network. It will send some traffic to R2 and some traffic to R3. You can see this in the routing table.

```

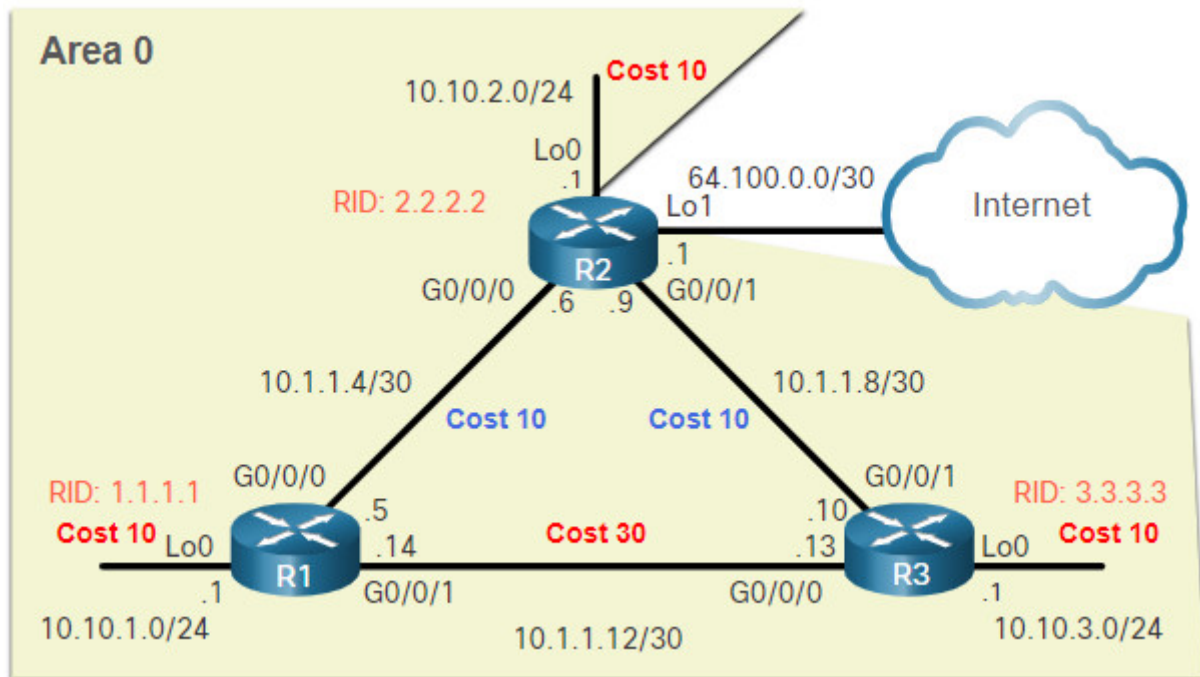
R1# show ip route ospf | begin 10
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
0       10.1.1.8/30 [110/20] via 10.1.1.13, 00:54:50, GigabitEthernet0/0/1
          [110/20] via 10.1.1.6, 00:55:14, GigabitEthernet0/0/0
(output omitted)
R1#

```

The administrator may want traffic to go through R2 and use R3 as a backup route in case the link between R1 and R2 goes down.

Another reason to change the cost value is because other vendors may calculate OSPF in a different manner. By manipulating the cost value, the administrator can make sure the route costs shared between OSPF multivendor routers are accurately reflected in routing tables.

To change the cost value reported by the local OSPF router to other OSPF routers, use the interface configuration command **ip ospf cost *value***. In the figure, we need to change cost of the loopback interfaces to 10 to simulate Gigabit Ethernet speeds. In addition, we will change the cost of the link between R2 and R3 to 30 so that this link is used as a backup link.



The following example is the configuration for R1.

```
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#
```

Assuming OSPF costs for R2 and R3 have been configured to match the topology in the above figure, the OSPF routes for R1 would have the following cost values. Notice that R1 is no longer load balancing to the 10.1.1.8/30 network. In fact, all routes go through R2 as desired by the network administrator.

```
R1# show ip route ospf | begin 10
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
0      10.1.1.8/30 [110/20] via 10.1.1.6, 01:18:25, GigabitEthernet0/0/0
0      10.10.2.0/24 [110/20] via 10.1.1.6, 00:04:31, GigabitEthernet0/0/0
0      10.10.3.0/24 [110/30] via 10.1.1.6, 00:03:21, GigabitEthernet0/0/0
R1#
```

Note: Although using the `ip ospf cost` command is the recommended method to manipulate the OSPF cost values, an administrator could also do this by using the interface configuration `bandwidth kbps` command. However, that would only work if all the routers are Cisco routers.

2.4.5. Test Failover to Backup Route

What happens if the link between R1 and R2 goes down? We 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(config)# interface g0/0/0
R1(config-if)# shutdown
*Jun  7 03:41:34.866: %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  7 03:41:36.865: %LINK-5-CHANGED: Interface GigabitEthernet0/0/0, changed
state to administratively down
*Jun  7 03:41:37.865: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/0/0, changed state to down
R1(config-if)# end
R1# show ip route ospf | begin 10
    10.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
0       10.1.1.4/30 [110/50] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
0       10.1.1.8/30 [110/40] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
0       10.10.2.0/24 [110/50] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
0       10.10.3.0/24 [110/40] via 10.1.1.13, 00:00:14, GigabitEthernet0/0/1
R1#
```

2.4.6. Syntax Checker – Modify the Cost Values for R2 and R3

Use the Syntax Checker to modify the cost values for R2 and R3.

You are logged into R2.

- Use the interface name lo0 to set the cost value for the loopback interface to 10.
- Return to global configuration mode.
- Verify the routing table entries with show ip route ospf.

```
R2(config)#interface lo0
R2(config-if)#ip ospf cost 10
R2(config-if)#end
R2#show ip route ospf
(output omitted)
Gateway of last resort is not set
    10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
0       10.1.1.12/30 \[110/20\] via 10.1.1.10, 00:01:13, GigabitEthernet0/0/1
0       10.10.1.0/24 \[110/20\] via 10.1.1.5, 00:00:32, GigabitEthernet0/0/0
0       10.10.3.0/24 \[110/11\] via 10.1.1.10, 00:00:49, GigabitEthernet0/0/1
R2#
```

The cost value to the R3 LAN is 11 because the loopback is still reporting a default cost of 1.

You are now logged into R3.

- Use the interface name lo0 and set the cost value for the loopback interface to 10.
- Use the interface name g0/0/0 and set the cost value of the link to R1 as 30.
- Return to privileged EXEC mode.
- Verify the routing table entries with show ip route ospf.

```
R3(config)#interface lo0
R3(config-if)#ip ospf cost 10
R3(config-if)#interface g0/0/0
R3(config-if)#ip ospf cost 30
R3(config-if)#end
R3#show ip route ospf
(output omitted)
Gateway of last resort is not set
    10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
0       10.1.1.4/30 \[110/20\] via 10.1.1.9, 01:48:54, GigabitEthernet0/0/1
0       10.10.1.0/24 \[110/30\] via 10.1.1.9, 00:00:06, GigabitEthernet0/0/1
0       10.10.2.0/24 \[110/20\] via 10.1.1.9, 00:35:24, GigabitEthernet0/0/1
R3#
```

You have successfully modified the OSPF cost values for R2 and R3.

2.4.7. Hello Packet Intervals

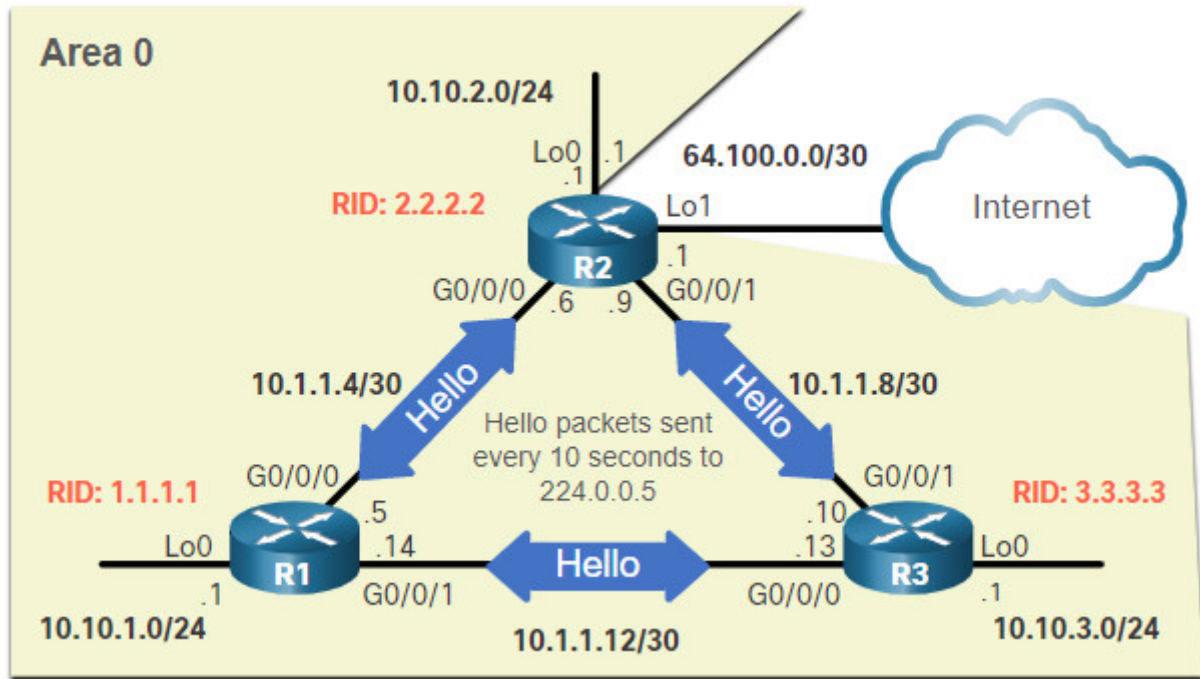
As shown in the figure, OSPFv2 Hello packets are transmitted to multicast address 224.0.0.5 (all OSPF routers) every 10 seconds. This is the default timer value on multiaccess and point-to-point networks.

Note: Hello packets are not sent on the simulated the LAN interfaces because those interfaces were set to passive by using the router configuration **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.

Note: On non-broadcast multiaccess (NBMA) networks, the default Hello interval is 30 seconds and the default dead interval is 120 seconds. NBMA networks are beyond the scope of this module.



2.4.8. 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, as shown in the example. 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
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
R1#

```

Use the **show ip ospf neighbor** command to see the Dead Time counting down from 40 seconds, as shown in the following example. 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#

```

2.4.9. 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.

In the example, the Hello interval for the link between R1 and R2 is changed to 5 seconds. Immediately after changing the Hello interval, 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.

As displayed by the highlighted OSPFv2 adjacency message, when the Dead Timer on R1 expires, R1 and R2 lose adjacency. The reason is because the 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. Notice that the only neighbor listed is the 3.3.3.3 (R3) router and that R1 is no longer adjacent with the 2.2.2.2 (R2) neighbor.

```
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#
```

To restore adjacency between R1 and R2, the R2 Gigabit Ethernet 0/0/0 interface Hello interval is set to 5 seconds, as shown in the following example. Almost immediately, the IOS displays a message that adjacency has been established with a state of FULL. Verify the interface intervals using the **show ip ospf interface** command. Notice that the Hello time is 5 seconds and that the Dead Time was automatically set to 20 seconds instead of the default 40 seconds.

```
R2(config)# interface g0/0/0
R2(config-if)# ip ospf hello-interval 5
*Jun  7 15:08:30.211: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R2(config-if)# end
R2# show ip ospf interface g0/0/0 | include Timer
  Timer intervals configured, Hello 5, Dead 20, Wait 20, Retransmit 5
R2# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address      Interface
3.3.3.3          0    FULL/ -         00:00:38   10.1.1.10    GigabitEthernet0/0/1
1.1.1.1          0    FULL/ -         00:00:17   10.1.1.5     GigabitEthernet0/0/0
R2#
```

2.4.10. Syntax Checker – Modifying Hello and Dead Intervals on R3

The Hello and Dead intervals are set to 5 and 20, respectively, on R1 and R2. Use the Syntax Checker to modify the Hello and Dead intervals on R3 and verify adjacencies are re-established with R1 and R2.

You are logged into R3. Enter the show ip ospf neighbor command to see that no adjacencies currently exist with R1 and R2.

```
R3#show ip ospf neighbor
R3#
```

No results are returned for the adjacent neighbors.
You are now in global configuration mode.

- Use g0/0/0 as the interface name and configure the Hello interval to match R1.
- Use g0/0/1 as the interface name and configure the Hello interval to match R2.
- Return to privileged EXEC mode.
- Verify neighbor adjacencies are re-established with the show ip ospf neighbor command.

```
R3(config)#interface g0/0/0
R3(config-if)#ip ospf hello-interval 5
\*Jun  7 05:11:34.423: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
GigabitEthernet0/0/0 from LOADING to FULL, Loading Done
R3(config-if)#interface g0/0/1
R3(config-if)#ip ospf hello-interval 5
\*Jun  7 05:11:43.081: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
GigabitEthernet0/0/1 from LOADING to FULL, Loading Done
R3(config-if)#end
R3#show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
2.2.2.2          0    FULL/ -         00:00:19    10.1.1.9       GigabitEthernet0/0/1
1.1.1.1          0    FULL/ -         00:00:19    10.1.1.14      GigabitEthernet0/0/0
```

```
R3#
```

You have successfully modified the OSPF Hello and Dead intervals on R3.

2.4.11. Packet Tracer – Modify Single-Area OSPFv2

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

- Adjust the reference bandwidth to account for gigabit and faster speeds
- Modify the OSPF cost value
- Modify the OSPF Hello timers
- Verify the modifications are accurately reflected in the routers.


2.4.11 Packet Tracer – Modify Single-Area OSPFv2

2.5. Default Route Propagation

2.5.1. Propagate a Default Static Route in OSPFv2

Your network users will need to send packets out of your network to non-OSPF networks, such as the internet. This is where you will need to have a default static route that they can use. In the topology in the figure, R2 is connected to the internet and should propagate a default route to R1 and R3. The router connected to the internet is

sometimes called the edge router or the gateway router. However, in OSPF terminology, the router located between an OSPF routing domain and a non-OSPF network is called the autonomous system boundary router (ASBR).



All that is required for R2 to reach the internet is a default static route to the service provider.

Note: In this example, a loopback interface with IPv4 address 64.100.0.1 is used to simulate the connection to the service provider.

To propagate a default route, the edge router (R2) 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 following example, R2 is configured with a loopback to simulate a connection to the internet. Then 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#

```

2.5.2. 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 and R3 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.

Click each button see the output for the show ip route command on each router.

- [R2 Routing Table](#)
- [R1 Routing Table](#)
- [R3 Routing Table](#)

R2 Routing Table

```

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
C     10.1.1.4/30 is directly connected, GigabitEthernet0/0/0
L     10.1.1.6/32 is directly connected, GigabitEthernet0/0/0
C     10.1.1.8/30 is directly connected, GigabitEthernet0/0/1
L     10.1.1.9/32 is directly connected, GigabitEthernet0/0/1
O     10.1.1.12/30 [110/40] via 10.1.1.10, 00:48:42, GigabitEthernet0/0/1
           [110/40] via 10.1.1.5, 00:59:30, GigabitEthernet0/0/0
O     10.10.1.0/24 [110/20] via 10.1.1.5, 00:59:30, GigabitEthernet0/0/0
C     10.10.2.0/24 is directly connected, Loopback0
L     10.10.2.1/32 is directly connected, Loopback0
O     10.10.3.0/24 [110/20] via 10.1.1.10, 00:48:42, GigabitEthernet0/0/1
     64.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C     64.100.0.0/30 is directly connected, Loopback1
L     64.100.0.1/32 is directly connected, Loopback1
R2#

```

2.5.3. Packet Tracer – Propagate a Default Route in OSPFv2

In this activity, you will configure an IPv4 default route to the internet and propagate that default route to other OSPF routers. You will then verify the default route is in downstream routing tables and that hosts can now access a web server on the internet.

2.5.3 Packet Tracer – Propagate a Default Route in OSPFv2

2.6.1. Verify OSPF Neighbors

If you have configured single-area OSPFv2, you will need to verify your configurations. This topic details the many commands that you can use to verify OSPF.

As you know, the following two commands are particularly useful for verifying routing:

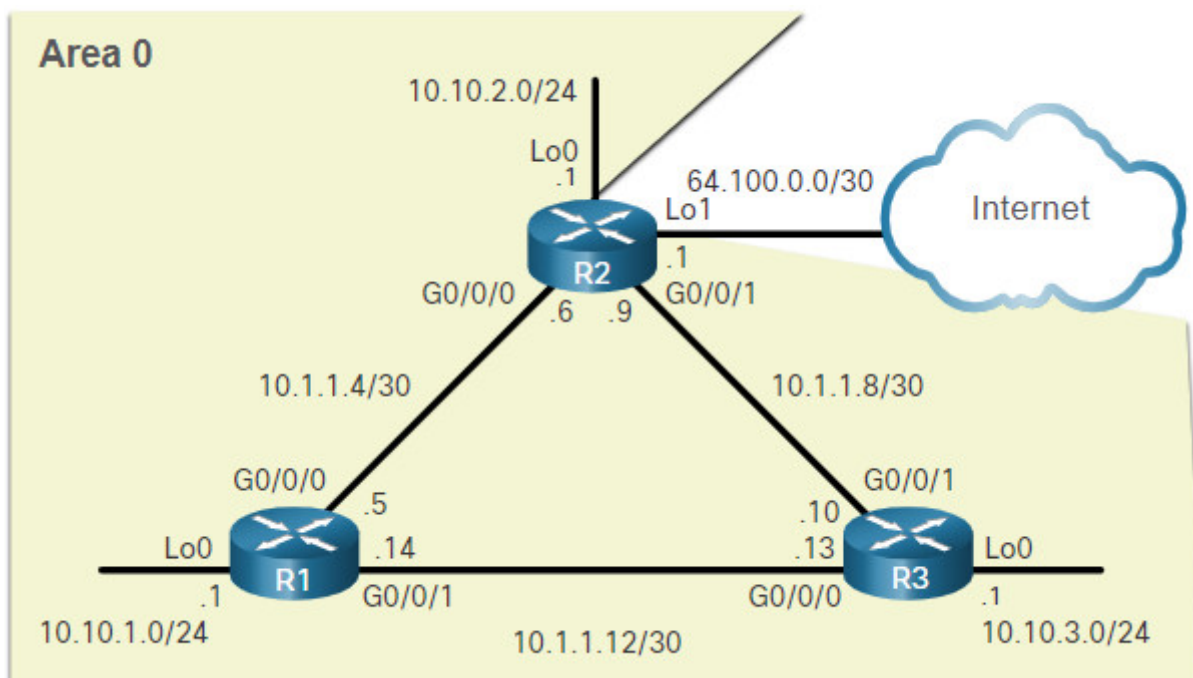
- **show ip interface brief** – This verifies that the desired interfaces are active with correct IP addressing.
- **show ip route**– This verifies that the routing table contains all the expected routes.

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**

The figure shows the OSPF reference topology used to demonstrate these commands.

OSPFv2 Reference Topology



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.

If two routers do not establish adjacency, link-state information is not exchanged. Incomplete LSDBs can cause inaccurate SPF trees and routing tables. Routes to destination networks may not exist, or may not be the most optimum path.

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/ -         00:00:19   10.1.1.13   GigabitEthernet0/0/1
2.2.2.2          0    FULL/ -         00:00:18   10.1.1.6    GigabitEthernet0/0/0
R1#
```

For each neighbor, this command displays the following:

- **Neighbor ID** – This is the router ID of the neighboring router.
- **Pri** – This is the OSPFv2 priority of the interface. This value is used in the DR and BDR election.
- **State** – This is the OSPFv2 state of the interface. FULL state means that the router and its neighbor have identical OSPFv2 LSDBs. On multiaccess networks, such as Ethernet, two routers that are adjacent may have their states displayed as 2WAY. The dash indicates that no DR or BDR is required because of the network type.
- **Dead Time** – This is the amount of time remaining that the router waits to receive an OSPFv2 Hello packet from the neighbor before declaring the neighbor down. This value is reset when the interface receives a Hello packet.
- **Address** – This is the IPv4 address of the interface of the neighbor to which this router is directly connected.
- **Interface** – This is the interface on which this router has formed adjacency with the neighbor.

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.

2.6.2. Verify OSPF Protocol Settings

The **show ip protocols** command is a quick way to verify vital OSPF configuration information, as shown in the following 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#
```

2.6.3. 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 following 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
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
Supports NSSA (compatible with RFC 3101)
Supports Database Exchange Summary List Optimization (RFC 5243)
Event-log enabled, Maximum number of events: 1000, Mode: cyclic
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
EXCHANGE/LOADING adjacency limit: initial 300, process maximum 300
Number of external LSA 1. Checksum Sum 0x00A1FF
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Number of areas transit capable is 0
External flood list length 0
IETF NSF helper support enabled
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#

2.6.4. 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, as shown in the following output for Gigabit Ethernet 0/0/0. 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
  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 5, Dead 20, Wait 20, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
R1#

```

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

- 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#

```

2.6.5. Syntax Checker – Verify Single-Area OSPFv2

Use the Syntax Checker to verify single-area OSPFv2 configuration on R2 and R3.

OSPF Reference Topology



Display a summary of the status of IPv4 interfaces on R2.

```
R2#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
GigabitEthernet0/0/0	10.1.1.6	YES	NVRAM	up	up
GigabitEthernet0/0/1	10.1.1.9	YES	NVRAM	up	up
Serial0/1/0	unassigned	YES	NVRAM	administratively down	down
Serial0/1/1	unassigned	YES	NVRAM	administratively down	down
GigabitEthernet0	unassigned	YES	NVRAM	administratively down	down
Loopback0	10.10.2.1	YES	NVRAM	up	up
Loopback1	64.100.0.1	YES	NVRAM	up	up

```
R2#
```

Display the OSPF routes installed in the routing table on R2.

```
R2#show ip route ospf
```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks

```
0      10.1.1.12/30 \[110/40\] via 10.1.1.10, 00:45:28, GigabitEthernet0/0/1
      \[110/40\] via 10.1.1.5, 00:45:49, GigabitEthernet0/0/0
0      10.10.1.0/24 \[110/20\] via 10.1.1.5, 00:45:49, GigabitEthernet0/0/0
0      10.10.3.0/24 \[110/20\] via 10.1.1.10, 00:45:28, GigabitEthernet0/0/1
```

```
R2#
```

Display the OSPF neighbor table for R2.

```
R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
3.3.3.3	0	FULL/ -	00:00:16	10.1.1.10	GigabitEthernet0/0/1
1.1.1.1	0	FULL/ -	00:00:19	10.1.1.5	GigabitEthernet0/0/0

```
R2#
```

Verify the protocol settings on R2.

```
R2#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 2.2.2.2

It is an autonomous system boundary router

Redistributing External Routes from,

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          00:46:14
  1.1.1.1          110          00:46:36
Distance: (default is 110)
```

R2#

Verify the OSPF process on R2.

R2#show ip ospf

```
Routing Process "ospf 10" with ID 2.2.2.2
  Start time: 00:01:54.811, Time elapsed: 00:48:04.766
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  Supports Link-local Signaling (LLS)
  Supports area transit capability
  Supports NSSA (compatible with RFC 3101)
  Event-log enabled, Maximum number of events: 1000, Mode: cyclic
  It is an autonomous system boundary router
  Redistributing External Routes from,
  Router is not originating router-LSAs with maximum metric
  Initial SPF schedule delay 5000 msec
  Minimum hold time between two consecutive SPF's 10000 msec
  Maximum wait time between two consecutive SPF's 10000 msec
  Incremental-SPF disabled
  Minimum LSA interval 5 secs
  Minimum LSA arrival 1000 msec
  LSA group pacing timer 240 secs
  Interface flood pacing timer 33 msec
  Retransmission pacing timer 66 msec
  Number of external LSA 1. Checksum Sum 0x009F01
  Number of opaque AS LSA 0. Checksum Sum 0x000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Number of areas transit capable is 0
  External flood list length 0
  IETF NSF helper support enabled
  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 00:47:04.655 ago
    SPF algorithm executed 5 times
    Area ranges are
    Number of LSA 3. Checksum Sum 0x00E181
    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
```

R2#

Use the interface name g0/0/0 to verify the OSPF interface settings on R2.

```

R2#show ip ospf interface g0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.6/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 2.2.2.2, 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 5, Dead 20, Wait 20, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1
  Suppress hello for 0 neighbor(s)

```

R2#

Display summary information for all the OSPF interfaces on R2.

```
R2#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	10	0	10.10.2.1/24	10	P2P	0/0	
Gi0/0/1	10	0	10.1.1.9/30	10	P2P	1/1	
Gi0/0/0	10	0	10.1.1.6/30	10	P2P	1/1	

R2#

You are now logged into R3. Display a summary of the status of IPv4 interfaces on R3.

```
R3#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
GigabitEthernet0/0/0	10.1.1.13	YES	NVRAM	up	up
GigabitEthernet0/0/1	10.1.1.10	YES	NVRAM	up	up
Serial0/1/0	unassigned	YES	NVRAM	administratively down	down
Serial0/1/1	unassigned	YES	NVRAM	administratively down	down
GigabitEthernet0	unassigned	YES	NVRAM	administratively down	down
Loopback0	10.10.3.1	YES	NVRAM	up	up

R3#

Display the OSPF routes installed in the routing table on R3.

```
R3#show ip route ospf
```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 a - application route
 + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is 10.1.1.9 to network 0.0.0.0

```

O*E2 0.0.0.0/0 \[110/1\] via 10.1.1.9, 00:49:56, GigabitEthernet0/0/1

```

```
10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
0      10.1.1.4/30 \[110/20\] via 10.1.1.9, 00:49:56, GigabitEthernet0/0/1
0      10.10.1.0/24 \[110/30\] via 10.1.1.9, 00:49:56, GigabitEthernet0/0/1
0      10.10.2.0/24 \[110/20\] via 10.1.1.9, 00:49:56, GigabitEthernet0/0/1
R3#
```

Display the neighbor table for R3.

```
R3#show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address      Interface
2.2.2.2          0    FULL/ -         00:00:18   10.1.1.9     GigabitEthernet0/0/1
1.1.1.1          0    FULL/ -         00:00:19   10.1.1.14    GigabitEthernet0/0/0
R3#
```

Verify the protocol settings on R3.

```
R3#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 3.3.3.3
  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
    1.1.1.1          110           00:50:44
    2.2.2.2          110           00:50:44
  Distance: (default is 110)
```

R3#

Verify the OSPF process on R3.

```
R3#show ip ospf
Routing Process "ospf 10" with ID 3.3.3.3
  Start time: 00:01:38.093, Time elapsed: 00:52:05.897
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  Supports Link-local Signaling (LLS)
  Supports area transit capability
  Supports NSSA (compatible with RFC 3101)
  Supports Database Exchange Summary List Optimization (RFC 5243)
  Event-log enabled, Maximum number of events: 1000, Mode: cyclic
  Router is not originating router-LSAs with maximum metric
  Initial SPF schedule delay 5000 msec
  Minimum hold time between two consecutive SPFs 10000 msec
  Maximum wait time between two consecutive SPFs 10000 msec
  Incremental-SPF disabled
  Minimum LSA interval 5 secs
  Minimum LSA arrival 1000 msec
  LSA group pacing timer 240 secs
```



```

Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
EXCHANGE/LOADING adjacency limit: initial 300, process maximum 300
Number of external LSA 1. Checksum Sum 0x009F01
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Number of areas transit capable is 0
External flood list length 0
IETF NSF helper support enabled
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 00:51:04.059 ago
    SPF algorithm executed 4 times
    Area ranges are
    Number of LSA 3. Checksum Sum 0x00E181
    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

```

R3#

Use the interface name g0/0/0 to verify the OSPF interface settings on R3.

```

R3#show ip ospf interface g0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 10.1.1.13/30, Area 0, Attached via Interface Enable
  Process ID 10, Router ID 3.3.3.3, Network Type POINT_TO_POINT, Cost: 30
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                30        no           no           Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 5, Dead 20, Wait 20, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1
  Suppress hello for 0 neighbor(s)

```

R3#

Display summary information for all the OSPF interfaces on R3.

```

R3#show ip ospf interface brief

```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	10	0	10.10.3.1/24	10	P2P	0/0	
Gi0/0/1	10	0	10.1.1.10/30	10	P2P	1/1	

Gi0/0/0 10 0 10.1.1.13/30 30 P2P 1/1

R3#

You have successfully verified single-area OSPFv2 on R2 and R3.

2.6.6. Packet Tracer – Verify Single-Area OSPFv2

In this Packet Tracer activity, you will use a variety of commands to verify the single-area OSPFv2 configuration.

2.6.6 Packet Tracer – Verify Single-Area OSPFv2

2.7. Module Practice and Quiz

2.7.1. Packet Tracer – Single-Area OSPFv2 Configuration

You are helping a network engineer test an OSPF set up by building the network in the lab where you work. You have interconnected the devices and configured the interfaces and have connectivity within the local LANs. Your job is to complete the OSPF configuration according to the requirements left by the engineer.

In this Packet Tracer activity, use the information provided and the list of requirements to configure the test network. When the task has been successfully completed, all hosts should be able to ping the internet server.

2.7.1 Packet Tracer – Single-Area OSPFv2 Configuration

2.7.2. Lab – Single-Area OSPFv2 Configuration

Skills Practice Opportunity

In this lab, you will complete the following objectives:

- Part 1: Build the network and configure basic device settings
- Part 2: Configure and verify single-area OSPFv2 for basic operation
- Part 3: Optimize and verify the single-area OSPFv2 configuration

You can practice these skills using the Packet Tracer or lab equipment, if available.

Packet Tracer – Physical Mode (PTPM)

2.7.2 Packet Tracer – Configure Single-Area OSPFv2 – Physical Mode

Lab Equipment

2.7.2 Lab – Configure Single-Area OSPFv2

2.7.3 Packet Tracer – Multiarea OSPF Exploration

This Packet Tracer exploration activity is divided into three parts and uses Packet Tracer Physical Mode (PTPM):

Part 1: The Beginning

Casual Recording Company, based in Sao Paulo, Brazil, provides self-serve, mini-recording studios around the city so that anyone can rent a timeslot and record their songs by themselves. CRC began with a single area OSPF network located in one building. This idea was very popular and consequently the business has grown, causing the company to expand into a branch office in a second building at the far end of the city. They continued to use single-area OSPF. You will evaluate the impact of the expanded network.

2.7.3 Packet Tracer – Multiarea OSPF Exploration – Physical Mode (Part 1)

Part 2: Business is Booming

The IT department in CRC decided to migrate to a multiarea OSPF network. You will evaluate the impact and benefits derived from the change to determine if it was a good decision or not.

2.7.3 Packet Tracer – OSPF Multiarea Exploration – Physical Mode (Part 2)

Part 3: CRC Expansion Continues

CRC has continued to grow and will open a new branch office in Montevideo, Uruguay. You will configure the Area Border Router for the new area and physically connect the branch office network to the corporate headquarters network through the internet.

2.7.3 Packet Tracer – OSPF Multiarea Exploration – Physical Mode (Part 3)

2.7.4 What did I learn in this module?

OSPF Router ID

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 the following preferential order:

1. The router ID is explicitly configured using the OSPF **router-id rid** router configuration mode command. The *rid* value is any 32-bit value expressed as an IPv4 address.
2. If the router ID is not explicitly configured, the router chooses the highest IPv4 address of any of configured loopback interfaces.
3. If no loopback interfaces are configured, then the router chooses the highest active IPv4 address of any of its physical interfaces.

The router ID can be assigned to a loopback interface. The IPv4 address for this type of loopback interface should be configured using a 32-bit subnet mask (255.255.255.255), creating a host route. A 32-bit host route would not get advertised as a route to other OSPF routers. 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. Use the **clear ip ospf process** command to reset the adjacencies. You can then verify that R1 is using the new router ID command with the **show ip protocols** command piped to display only the router ID section.

Point-to-Point OSPF Networks

The network command is used to determine which interfaces participate in the routing process for an OSPFv2 area. 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. In a wildcard mask:

- **Wildcard mask bit 0** – Matches the corresponding bit value in the address
- **Wildcard mask bit 1** – Ignores the corresponding bit value in the address

Within routing configuration mode, there are two ways to identify the interfaces that will participate in the OSPFv2 routing process. One way is when 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. The other way is OSPFv2 can be enabled by specifying the exact interface IPv4 address using a quad zero wildcard mask. 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**. Sending out unneeded messages on a LAN affects the network through inefficient use of bandwidth and resources, and creates an increased security risk. 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 **show ip protocols** command is then used to verify that the Loopback 0 interface is listed as passive. 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. Use loopbacks to simulate more networks than the equipment can support. 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

Routers can be connected to the same switch to form a multiaccess network. Ethernet LANs are the most common example of broadcast multiaccess networks. In broadcast networks, all devices on the network see all broadcast and multicast frames. 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. To verify the OSPFv2 adjacencies, use the **show ip ospf neighbor** command. The state of neighbors in multiaccess networks can be:

- **FULL/DROTHER** – This is a DR or BDR router that is fully adjacent with a non-DR or BDR router.
- **FULL/DR** – The router is fully adjacent with the indicated DR neighbor.
- **FULL/BDR** – The router is fully adjacent with the indicated BDR neighbor.
- **2-WAY/DROTHER** – The non-DR or BDR router has a neighbor relationship with another non-DR or BDR router.

The OSPF DR and BDR election decision 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 elected as 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. Therefore, unless otherwise configured, all routers have an equal priority value and must rely on another tie breaking method during the DR/BDR election.
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.

OSPF DR and BDR elections are not pre-emptive. 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. 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.

Modify Single-Area OSPFv2

OSPF uses cost as a metric. A lower cost indicates a better path than a higher cost. 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: $\text{Cost} = \text{reference bandwidth} / \text{interface bandwidth}$. 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. To adjust the reference bandwidth, use the **auto-cost reference-bandwidth Mbps** router configuration 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 seconds** and **ip ospf dead-interval seconds**.

Default Route Propagation

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. This instructs the ASBR to be the source of the default route information and propagate the default static route in OSPF updates. Verify the default route settings on the ASBR using the **show ip route** command.

Verify Single-Area OSPFv2

The following two commands are used to verify routing:

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

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. For each neighbor, this command displays:

- **Neighbor ID** – The router ID of the neighboring router.
- **Pri** – The OSPFv2 priority of the interface. This value is used in the DR and BDR election.
- **State** – The OSPFv2 state of the interface. FULL state means that the router and its neighbor have identical OSPFv2 LSDBs. On multiaccess networks, such as Ethernet, two routers that are adjacent may have their states displayed as 2WAY. The dash indicates that no DR or BDR is required because of the network type.

- **Dead Time** – The amount of time remaining that the router waits to receive an OSPFv2 Hello packet from the neighbor before declaring the neighbor down. This value is reset when the interface receives a Hello packet.
- **Address** – The IPv4 address of the neighbor's interface to which this router is directly connected.
- **Interface** – The interface on which this router has formed adjacency with the neighbor.

The **show ip protocols** command is a quick way to verify vital OSPF configuration information such as 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. Use the **show ip ospf** command to examine the OSPFv2 process ID and router ID. This command displays the OSPFv2 area information and the last time the SPF algorithm was executed. The **show ip ospf interface** command provides a detailed list for every OSPFv2-enabled interface. Specify an interface for just one interface to display the process ID, the local router ID, the type of network, OSPF cost, DR and BDR information on multiaccess links, and adjacent neighbors.

2.7.5 Module Quiz – Single-Area OSPFv2 Configuration

Post navigation

[← Previous Article](#)

[CCNA 3 v7.0 Curriculum: Module 1 – Single-Area OSPFv2 Concepts](#)

[Next Article →](#)

[CCNA 3 v7.0 Curriculum: Module 3 – Network Security Concepts](#)