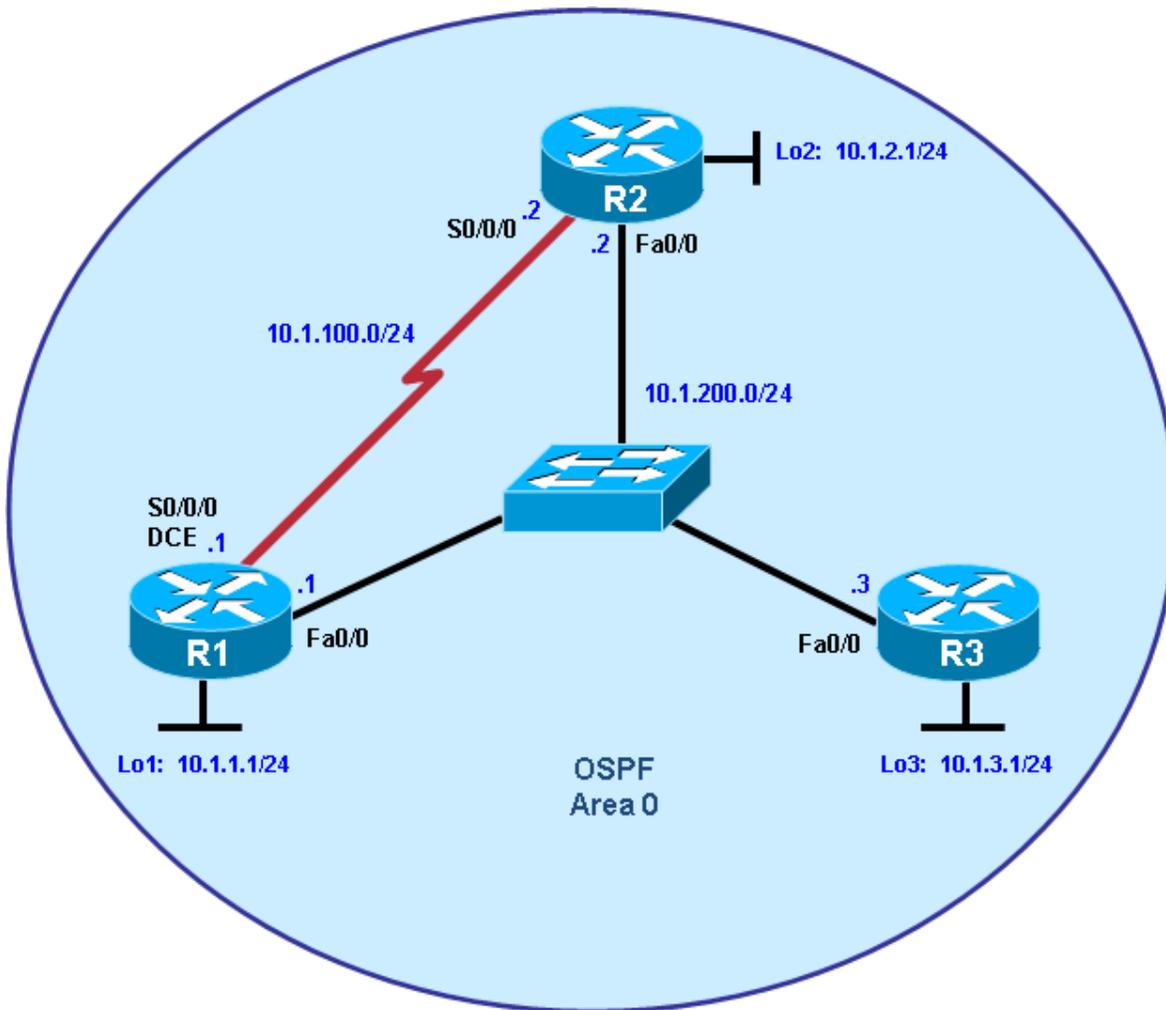


## Chapter 3 Lab 3-1, Single-Area OSPF Link Costs and Interface Priorities

### Topology



### Objectives

- Configure single-area OSPF on a router.
- Advertise loopback interfaces into OSPF.
- Verify OSPF adjacencies.
- Verify OSPF routing information exchange.
- Modify OSPF link costs.
- Change interface priorities.
- Utilize debugging commands for troubleshooting OSPF.

## Background

You are responsible for configuring the new network to connect your company's engineering, marketing, and accounting departments, represented by the loopback interfaces on each of the three routers. The physical devices have just been installed and connected by Fast Ethernet and serial cables. Configure OSPF to allow full connectivity between all departments.

**Note:** This lab uses Cisco 1841 routers with Cisco IOS Release 12.4(24)T1 and the Advanced IP Services image c1841-advipservicesk9-mz.124-24.T1.bin. The switch is a Cisco WS-C2960-24TT-L with the Cisco IOS image c2960-lanbasek9-mz.122-46.SE.bin. You can use other routers (such as a 2801 or 2811), switches (such as a 2950), and Cisco IOS Software versions if they have comparable capabilities and features. Depending on the router or switch model and Cisco IOS Software version, the commands available and output produced might vary from what is shown in this lab.

## Required Resources

- 3 routers (Cisco 1841 with Cisco IOS Release 12.4(24)T1 Advanced IP Services or comparable)
- 1 switch (Cisco 2960 with the Cisco IOS Release 12.2(46)SE C2960-LANBASEK9-M image or comparable)
- Serial and Ethernet cables

## Step 1: Configure addressing and loopbacks.

- a. Using the addressing scheme in the diagram, apply IP addresses to the Fast Ethernet interfaces on R1, R2, and R3. Create Loopback1 on R1, Loopback2 on R2, and Loopback3 on R3, and address them according to the diagram.

**Note:** Depending on the router models you have, you might need to add clock rates to the DCE end of each connection (newer equipment adds this automatically). Verify connectivity across each serial link.

```
R1# configure terminal
R1(config)# interface Loopback1
R1(config-if)# description Engineering Department
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# exit
R1(config)# interface FastEthernet0/0
R1(config-if)# ip address 10.1.200.1 255.255.255.0
R1(config-if)# no shutdown
```

```
R2# configure terminal
R2(config)# interface Loopback2
R2(config-if)# description Marketing Department
R2(config-if)# ip address 10.1.2.1 255.255.255.0
R2(config-if)# exit
R2(config)# interface FastEthernet0/0
R2(config-if)# ip address 10.1.200.2 255.255.255.0
R2(config-if)# no shutdown
```

```
R3# configure terminal
R3(config)# interface Loopback3
R3(config-if)# description Accounting Department
R3(config-if)# ip address 10.1.3.1 255.255.255.0
R3(config-if)# exit
R3(config)# interface FastEthernet0/0
R3(config-if)# ip address 10.1.200.3 255.255.255.0
R3(config-if)# no shutdown
```

Leave the switch in its default (blank) configuration. By default, all switch ports are in VLAN1 and are not administratively down.

- b. Configure the serial interfaces on R1 and R2 with the IP addresses shown in the diagram. Add the **clockrate** command where needed.

```
R1(config)# interface Serial 0/0/0
R1(config-if)# ip address 10.1.100.1 255.255.255.0
R1(config-if)# clockrate 64000
R1(config-if)# bandwidth 64
R1(config-if)# no shutdown
```

```
R2(config)# interface Serial 0/0/0
R2(config-if)# ip address 10.1.100.2 255.255.255.0
R2(config-if)# bandwidth 64
R2(config-if)# no shutdown
```

**Note:** The **bandwidth** command on the serial interfaces is used to match the actual bandwidth of the link. By default, OSPF calculates the cost of links based on the default interface bandwidth which may be either 128 or 1544 Kb/s, depending on the WIC type. In this case the **bandwidth 64** command is used because the real bandwidth of the serial interfaces is set to 64 Kbps. Refer to Step 5 for information on modifying OSPF link costs.

- c. Verify that the appropriate interfaces are up and that you can ping across each link.

### Step 2: Add physical interfaces to OSPF.

- a. Enter the OSPF configuration prompt using the **router ospf process\_number** command. The process number is a locally significant number that does not affect how OSPF works. For this lab, use process number 1 on all the routers.
- b. Add interfaces with the **network address wildcard\_mask area area** command. The address is an IP address. The mask is an inverse mask, similar to the kind used in an access list. The area is the OSPF area to put the interface. For this lab, use area 0, the backbone area, for all interfaces.

This command can be confusing at first. What it means is that any interface with an IP address that matches the address and wildcard mask combination in the network statement is added to the OSPF process in that area. The wildcard mask used in the **network** command has no influence on the actual IP subnet mask that is advertised with a network on an interface. The **network** command selects interfaces to be included into OSPF, but OSPF advertises the real subnet mask of the network attached to that interface (with the only exception being loopback interfaces).

For example, the command **network 10.1.200.1 0.0.0.0 area 0** adds the interface with the IP address of 10.1.200.1 and its network to the OSPF process into area 0. The wildcard mask of 0.0.0.0 means that all 32 bits of the IP address have to be an exact match. A 0 bit in the wildcard mask means that portion of the interface IP must match the address. A 1 bit means that the bit in the interface IP does not have to match that portion of the IP address.

The command **network 10.1.100.0 0.0.0.255 area 0** means that any interface whose IP address matches 10.1.100.0 for the first 3 octets will match the command and add it to area 0. The last octet is all 1s, because in the wildcard mask it is 255. This means that an interface with an IP of 10.1.100.1, 10.1.100.2, or 10.1.100.250 would match this address and wildcard combination and get added to OSPF.

Instead of using wildcard masks in the **network** command, it is possible to use subnet masks. The router converts the subnet masks to the wildcard format automatically. An easy way to calculate a wildcard

mask from the subnet mask is to subtract the octet value for each octet from 255. For example, a subnet mask of 255.255.255.252 (/30) becomes 0.0.0.3 to capture all interfaces on that subnet:

```
255.255.255.255
- 255.255.255.252
= 0. 0. 0. 3
```

**Note:** Another option for adding individual directly connected networks into the OSPF process is to use the **ip ospf process-id area area-id** interface command that is available with Cisco IOS version 12.3(11)T and later.

- c. Enter the commands on R1. Exit to privileged EXEC mode and type **debug ip ospf adj**. The **debug** command lets you watch OSPF neighbors come up and see neighbor relationships.

```
R1(config)# router ospf 1
R1(config-router)# network 10.1.100.0 0.0.0.255 area 0
R1(config-router)# network 10.1.200.0 0.0.0.255 area 0
R1(config-router)# end
R1#
```

```
R1# debug ip ospf adj
OSPF adjacency events debugging is on
```

- d. Add network statements to the other two routers.

```
R2(config)# router ospf 1
R2(config-router)# network 10.1.100.0 0.0.0.255 area 0
R2(config-router)# network 10.1.200.0 0.0.0.255 area 0
```

```
R3(config)# router ospf 1
R3(config-router)# network 10.1.200.0 0.0.0.255 area 0
```

- e. Observe the debug output on R1. When you are finished, turn off debugging on R1 with the **undebug all** command.
- f. What is the advantage of adding networks with a wildcard mask instead of using classful network addresses?

### Step 3: Use OSPF show commands.

- a. The **show ip protocols** command displays basic high-level routing protocol information. The output lists each OSPF process, the router ID, and which networks OSPF is routing for in each area. This information can be useful in debugging routing operations.

```
R1# show ip protocols
Routing Protocol is "ospf 1"
```

```
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Router ID 10.1.1.1
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Maximum path: 4
Routing for Networks:
 10.1.100.0 0.0.0.255 area 0
 10.1.200.1 0.0.0.0 area 0
Reference bandwidth unit is 100 mbps
Routing Information Sources:
  Gateway          Distance          Last Update
Distance: (default is 110)
```

- b. The **show ip ospf** command displays the OSPF process ID and router ID.

```
R1# show ip ospf
Routing Process "ospf 1" with ID 10.1.1.1
Start time: 00:17:44.612, Time elapsed: 00:10:51.408
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
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 0. Checksum Sum 0x000000
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
  Area BACKBONE(0)
    Number of interfaces in this area is 2
    Area has no authentication
    SPF algorithm last executed 00:03:21.132 ago
    SPF algorithm executed 5 times
    Area ranges are
    Number of LSA 4. Checksum Sum 0x021A30
    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
```

Notice the router ID listed in the output. The R1 ID is 10.1.1.1, even though you have not added this loopback into the OSPF process. The router chooses the router ID using the highest IP on a loopback interface when OSPF is configured. If an additional loopback interface with a higher IP address is added after OSPF is turned on, it does not become the router ID unless the router is reloaded, the OSPF configuration is removed and reentered, or the OSPF-level command **router-id** is used to modify the RID manually and the **clear ip ospf process** command is subsequently entered. If no loopback interfaces are

present on the router, the router selects the highest available IP address among interfaces that are activated using the **no shutdown** command. If no IP addresses are assigned to interfaces, the OSPF process does not start.

- c. The **show ip ospf neighbor** command displays important neighbor status, including the adjacency state, address, router ID, and connected interface.

R1# **show ip ospf neighbor**

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.2.1	1	FULL/BDR	00:00:36	10.1.200.2	FastEthernet0/0
10.1.3.1	1	FULL/DR	00:00:35	10.1.200.3	FastEthernet0/0
10.1.2.1	0	FULL/ -	00:00:36	10.1.100.2	Serial0/0/0

If you need more detail than the standard one-line summaries of neighbors, use the **show ip ospf neighbor detail** command. However, generally, the regular command gives you all that you need.

- d. The **show ip ospf interface interface\_type number** command shows interface timers and network types.

R1# **show ip ospf interface FastEthernet 0/0**

```
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.1.200.1/24, Area 0
  Process ID 1, Router ID 10.1.1.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 10.1.3.1, Interface address 10.1.200.3
  Backup Designated router (ID) 10.1.2.1, Interface address 10.1.200.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:09
  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 0 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 10.1.3.1 (Designated Router)
    Adjacent with neighbor 10.1.2.1
  Suppress hello for 0 neighbor(s)
```

- e. A variation of the previous command is the **show ip ospf interface brief** command, which displays each interface that is participating in the OSPF process on the router, the area it is in, its IP address, cost, state, and number of neighbors.

R1# **show ip ospf interface brief**

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs F/C
Fa0/0	1	0	10.1.200.1/24	1	DROTH	2/2
Se0/0/0	1	0	10.1.100.1/24	1	P2P	1/1

- f. The **show ip ospf database** command displays the various LSAs in the OSPF database, organized by area and type.

R1# **show ip ospf database**

OSPF Router with ID (10.1.1.1) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
---------	------------	-----	------	----------	------------

## CCNPv6 ROUTE

---

```
10.1.1.1      10.1.1.1      1782      0x80000002 0x001AC7 3
10.1.2.1      10.1.2.1      1783      0x80000001 0x001DC2 3
10.1.3.1      10.1.3.1      1720      0x80000002 0x00F077 1
```

Net Link States (Area 0)

```
Link ID      ADV Router    Age          Seq#          Checksum
10.1.200.1   10.1.1.1     1719        0x80000002 0x00EC3C
```

OSPF Router with ID (10.1.1.1) (Process ID 1)

### Step 4: Add loopback interfaces to OSPF.

- a. All three routers have loopback interfaces, but they are not yet advertised in the routing process. You can verify this with the **show ip route** command on the three routers.

R1# **show ip route**

```
Codes: 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
```

Gateway of last resort is not set

```
10.0.0.0/24 is subnetted, 3 subnets
C      10.1.1.0 is directly connected, Loopback1
C      10.1.100.0 is directly connected, Serial0/0/0
C      10.1.200.0 is directly connected, FastEthernet0/0
```

R2# **show ip route**

```
Codes: 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
```

Gateway of last resort is not set

```
10.0.0.0/24 is subnetted, 3 subnets
C      10.1.2.0 is directly connected, Loopback2
C      10.1.100.0 is directly connected, Serial0/0/0
C      10.1.200.0 is directly connected, FastEthernet0/0
```

R3# **show ip route**

```
Codes: 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

Gateway of last resort is not set

```

10.0.0.0/24 is subnetted, 3 subnets
C    10.1.3.0 is directly connected, Loopback3
O    10.1.100.0 [110/65] via 10.1.200.2, 00:06:39, FastEthernet0/0
      [110/65] via 10.1.200.1, 00:06:39, FastEthernet0/0
C    10.1.200.0 is directly connected, FastEthernet0/0
  
```

- b. For each router, the only loopback address displayed is the locally connected one. Add the loopbacks into the routing process for each router using the same **network** command previously used to add the physical interfaces.

```

R1(config)# router ospf 1
R1(config-router)# network 10.1.1.0 0.0.0.255 area 0
  
```

```

R2(config)# router ospf 1
R2(config-router)# network 10.1.2.0 0.0.0.255 area 0
  
```

```

R3(config)# router ospf 1
R3(config-router)# network 10.1.3.0 0.0.0.255 area 0
  
```

- c. Verify that these networks have been added to the routing table using the **show ip route** command.

```
R1# show ip route
```

```

Codes: 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
  
```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O    10.1.2.1/32 [110/2] via 10.1.200.2, 00:00:03, FastEthernet0/0
O    10.1.3.1/32 [110/2] via 10.1.200.3, 00:00:03, FastEthernet0/0
C    10.1.1.0/24 is directly connected, Loopback1
C    10.1.100.0/24 is directly connected, Serial0/0/0
C    10.1.200.0/24 is directly connected, FastEthernet0/0
  
```

Now you can see the loopbacks of the other routers, but their subnet mask is incorrect, because the default network type on loopback interfaces advertises them as /32 (host) routes. As you can see in the output of the **show ip ospf interface Lo1** command, the default OSPF network type for a loopback interface is LOOPBACK, causing the OSPF to advertise host routes instead of actual network masks.

```
R1# show ip ospf interface Lo1
```

```

Loopback1 is up, line protocol is up
Internet Address 10.1.1.1/24, Area 0
Process ID 1, Router ID 10.1.1.1, Network Type LOOPBACK, Cost: 1
Loopback interface is treated as a stub Host
  
```

**Note:** The OSPF network type of LOOPBACK is a Cisco-proprietary extension that is not configurable but that is present on loopback interfaces by default. In some applications such as MPLS, the possible discrepancy

between the real loopback interface mask and the advertised address/mask can lead to reachability or functionality issues, and care must be taken to either use /32 mask on loopbacks, or whenever a different mask is used, the OSPF network type must be changed to point-to-point.

- d. To change this default behavior use the **ip ospf network point-to-point** command in interface configuration mode for each loopback. After the routes propagate, you see the correct subnet masks associated with those loopback interfaces.

```
R1(config)# interface loopback1
R1(config-if)# ip ospf network point-to-point
```

```
R2(config)# interface loopback2
R2(config-if)# ip ospf network point-to-point
```

```
R3(config)# interface loopback3
R3(config-if)# ip ospf network point-to-point
```

```
R1# show ip route
```

```
Codes: 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
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 5 subnets
```

```
O    10.1.3.0 [110/2] via 10.1.200.3, 00:00:01, FastEthernet0/0
O    10.1.2.0 [110/2] via 10.1.200.2, 00:00:01, FastEthernet0/0
C    10.1.1.0 is directly connected, Loopback1
C    10.1.100.0 is directly connected, Serial0/0/0
C    10.1.200.0 is directly connected, FastEthernet0/0
```

- e. Use the following Tcl script to verify connectivity to all addresses in the topology.

```
R1# tclsh
```

```
foreach address {
10.1.1.1
10.1.2.1
10.1.3.1
10.1.100.1
10.1.100.2
10.1.200.1
10.1.200.2
10.1.200.3
} {
ping $address }
```

### Step 5: Modify OSPF link costs.

When you use the **show ip route** command on R1, you see that the most direct route to the R2 loopback is through its Ethernet connection. Next to this route is a pair in the form [administrative distance / metric]. The default administrative distance of OSPF on Cisco routers is 110. The metric depends on the link type. OSPF always chooses the route with the lowest metric, which is a sum of all link costs.

## CCNPv6 ROUTE

---

You can modify a single link cost by using the interface command **ip ospf cost cost**. Use this command on both ends of the link. In the following commands, the link cost of the Fast Ethernet connection between the three routers is changed to a cost of 50. Notice the change in the metrics in the routing table.

```
R1(config)# interface FastEthernet 0/0
R1(config-if)# ip ospf cost 50
```

```
R2(config)# interface FastEthernet 0/0
R2(config-if)# ip ospf cost 50
```

```
R3(config)# interface FastEthernet 0/0
R3(config-if)# ip ospf cost 50
```

```
R1# show ip route
```

```
Codes: 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
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 5 subnets
O       10.1.3.0 [110/51] via 10.1.200.3, 00:01:40, FastEthernet0/0
O       10.1.2.0 [110/51] via 10.1.200.2, 00:01:40, FastEthernet0/0
C       10.1.1.0 is directly connected, Loopback1
C       10.1.100.0 is directly connected, Serial0/0/0
C       10.1.200.0 is directly connected, FastEthernet0/0
```

For reference, here are some default link costs (taken from Cisco.com):

- 64-kb/s serial link: 1562
- T1 (1.544-Mb/s serial link): 64
- E1 (2.048-Mb/s serial link): 48
- Ethernet: 10
- Fast Ethernet: 1
- FDDI: 1
- X25: 5208
- ATM: 1

OSPF uses a reference bandwidth of 100 Mb/s for cost calculation. The formula to calculate the cost is the reference bandwidth divided by the interface bandwidth. For example, in the case of Ethernet, is the cost is  $100 \text{ Mb/s} / 10 \text{ Mb/s} = 10$ .

The above link costs do not include Gigabit Ethernet, which is significantly faster than Fast Ethernet, but would still have a cost of 1 using the default reference bandwidth of 100 Mb/s.

The cost calculation can be adjusted to account for network links that are faster than 100 Mb/s by using the **auto-cost reference-bandwidth** command to change the reference bandwidth. For example, to change the reference bandwidth to 1000 Mb/s (Gigabit Ethernet), use the following commands:

```
R1(config)# router ospf 1
R1(config-router)# auto-cost reference-bandwidth 1000
```

% OSPF: Reference bandwidth is changed.  
Please ensure reference bandwidth is consistent across all routers.

**Note:** If the `ip ospf cost cost` command is used on the interface, as is the case here, it overrides this formulated cost.

**Note:** The above example is for reference only and should not be entered on R1.

### Step 6: Modify interface priorities to control the DR and BDR election.

If you use the `show ip ospf neighbor detail` command on any of the routers, you see that for the Ethernet network, R3 is the DR (designated router) and R2 is the BDR (backup designated router). These designations are determined by the interface priority for all routers in that network, which you see in the `show` output.

The default priority is 1. If all the priorities are the same (which happens by default), the DR election is then based on router IDs. The highest router ID router becomes the DR, and the second highest becomes the BDR. All other routers become DROTHERs.

**Note:** If your routers do not have this exact behavior, it might be because of the order in which the routers came up. Sometimes a router does not leave the DR position unless its interface goes down and another router takes over. Your routers might not behave exactly like the example.

Use the `ip ospf priority number` interface command to change the OSPF priorities on R1 and R2 to make R1 the DR and R2 the BDR. After changing the priority on both interfaces, look at the output of the `show ip ospf neighbor detail` command. You can also see the change with the `show ip ospf neighbor` command, but it requires more interpretation because it comes up with states per neighbor, rather than stating the DR and BDR on a neighbor adjacency network.

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

R2(config)# interface FastEthernet 0/0
R2(config-if)# ip ospf priority 5

R1# show ip ospf neighbor detail
Neighbor 10.1.2.1, interface address 10.1.200.2
  In the area 0 via interface FastEthernet0/0
  Neighbor priority is 5, State is FULL, 12 state changes
  DR is 10.1.200.1 BDR is 10.1.200.2
  Options is 0x52
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:37
  Neighbor is up for 00:01:32
  Index 3/3, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
Neighbor 10.1.3.1, interface address 10.1.200.3
  In the area 0 via interface FastEthernet0/0
  Neighbor priority is 1, State is FULL, 12 state changes
  DR is 10.1.200.1 BDR is 10.1.200.2
  Options is 0x52
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:30
  Neighbor is up for 00:01:12
  Index 1/1, retransmission queue length 0, number of retransmission 3
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
```

```
Last retransmission scan time is 0 msec, maximum is 0 msec
Neighbor 10.1.2.1, interface address 10.1.100.2
In the area 0 via interface Serial0/0/0
Neighbor priority is 0, State is FULL, 12 state changes
DR is 0.0.0.0 BDR is 0.0.0.0
Options is 0x52
LLS Options is 0x1 (LR)
Dead timer due in 00:00:35
Neighbor is up for 00:01:44
Index 2/2, retransmission queue length 0, number of retransmission 2
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 2, maximum is 2
Last retransmission scan time is 0 msec, maximum is 0 msec
```

**Note:** To make a router take over as DR, use the **clear ip ospf process** command on all the routers after changing the priorities. Another method of demonstrating the election process and priorities is to shutdown and reactivate all ports on the switch simultaneously. The switch can be configured with **spanning-tree portfast default** and all ports can be shutdown and reactivated using the following commands.

```
interface range fa0/1 - 24
shutdown
no shutdown
```

What is the purpose of a DR in OSPF?

What is the purpose of a BDR in OSPF?

### Challenge: Topology Change

OSPF, like many link-state routing protocols, is reasonably fast when it comes to convergence. To test this, have R3 send a large number of pings to the R1 loopback. By default, the pings take the path from R3 to R1 over Fast Ethernet because it has the lowest total path cost.

- a. Check the path from R3 to R1 by performing a traceroute on R3 to the loopback of R1.

```
R3# traceroute 10.1.1.1
```

```
Type escape sequence to abort.  
Tracing the route to 10.1.1.1
```

```
 1 10.1.200.1 0 msec 0 msec *
```

**Note:** Read the next substep carefully before trying out the commands on routers.

- b. Initiate a ping from R3 to the R1 loopback with a high repeat number using the command **ping ip repeat number** command. While this ping is going on, shut down the R1 Fa0/0 interface.

```
R3# ping 10.1.1.1 repeat 10000
```

```
R1(config)# interface FastEthernet 0/0  
R1(config-if)# shutdown
```

Did you notice that some packets were dropped but then the pings started returning again?

How do you think OSPF convergence compares to other routing protocols, such as RIP? What about EIGRP?

---

## Router Interface Summary Table

Router Interface Summary				
Router Model	Ethernet Interface #1	Ethernet Interface #2	Serial Interface #1	Serial Interface #2
1700	Fast Ethernet 0 (FA0)	Fast Ethernet 1 (FA1)	Serial 0 (S0)	Serial 1 (S1)
1800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
2600	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0 (S0/0)	Serial 0/1 (S0/1)
2800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
<b>Note:</b> To find out how the router is configured, look at the interfaces to identify the type of router and how many interfaces the router has. Rather than list all combinations of configurations for each router class, this table includes identifiers for the possible combinations of Ethernet and serial interfaces in the device. The table does not include any other type of interface, even though a specific router might contain one. For example, for an ISDN BRI interface, the string in parenthesis is the legal abbreviation that can be used in Cisco IOS commands to represent the interface.				