CCNA 200-301, Volume I

Chapter 20
Implementing OSPF

Objectives

- Implementing Single-Area OSPFv2
- OSPFv2 Interface Configuration
- Additional OSPFv2 Features

Implementing Single-Area OSPF

- Step 1. Use the **router ospf *process-id*** global command to enter OSPF configuration mode for a particular OSPF process.
- Step 2. (Optional) Configure the OSPF router ID by doing the following:
  - A. Use the **router-id *id-value*** router subcommand to define the router ID.
  - B. Use the **interface loopback *number*** global command, along with an **ip address *address mask*** command, to configure an IP address on a loopback interface (chooses the highest IP address of all working loopbacks).
  - C. Rely on an interface IP address (chooses the highest IP address of all working nonloopbacks).
- Step 3. Use one or more **network *ip-address wildcard-mask* area *area-id*** router subcommands to enable OSPFv2 on any interfaces matched by the configured address and mask, enabling OSPF on the interface for the listed area.
- Step 4. (Optional) Use the **passive-interface *type number*** router subcommand to configure any OSPF interfaces as passive if no neighbors can or should be discovered on the interface.

Implementing Single-Area OSPF

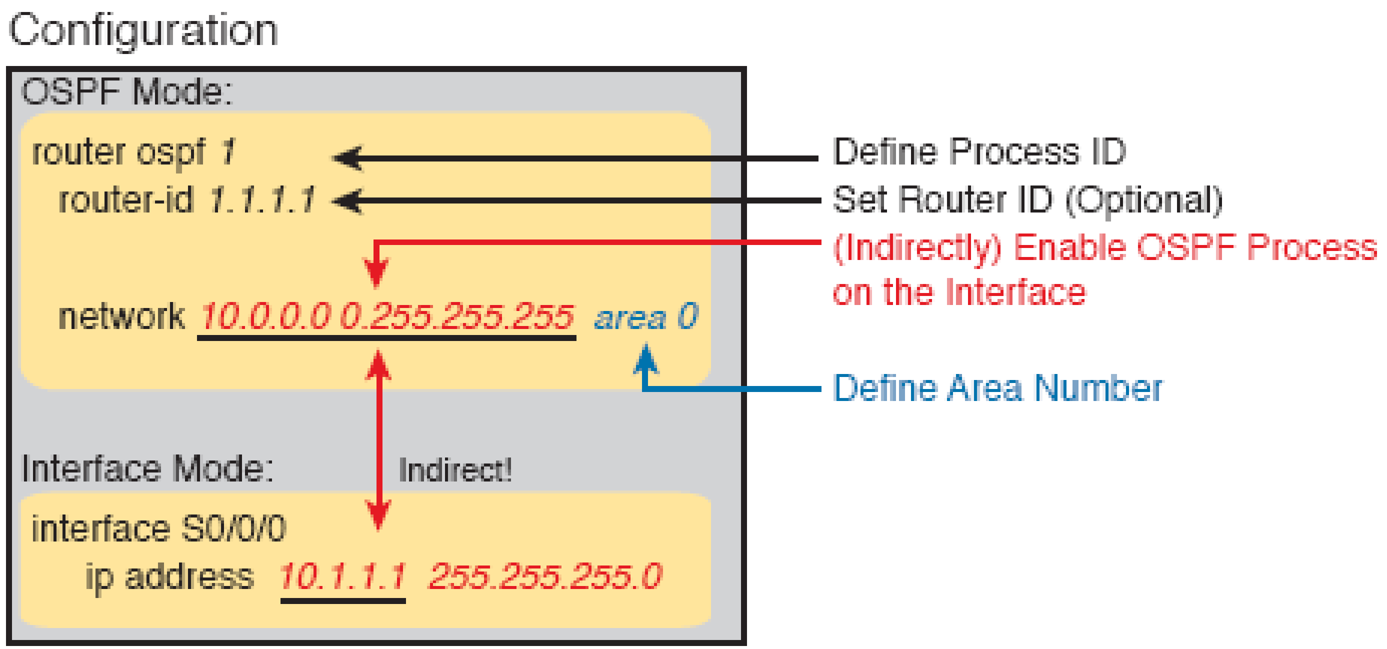

Sample Network for OSPF Single-Area Configuration

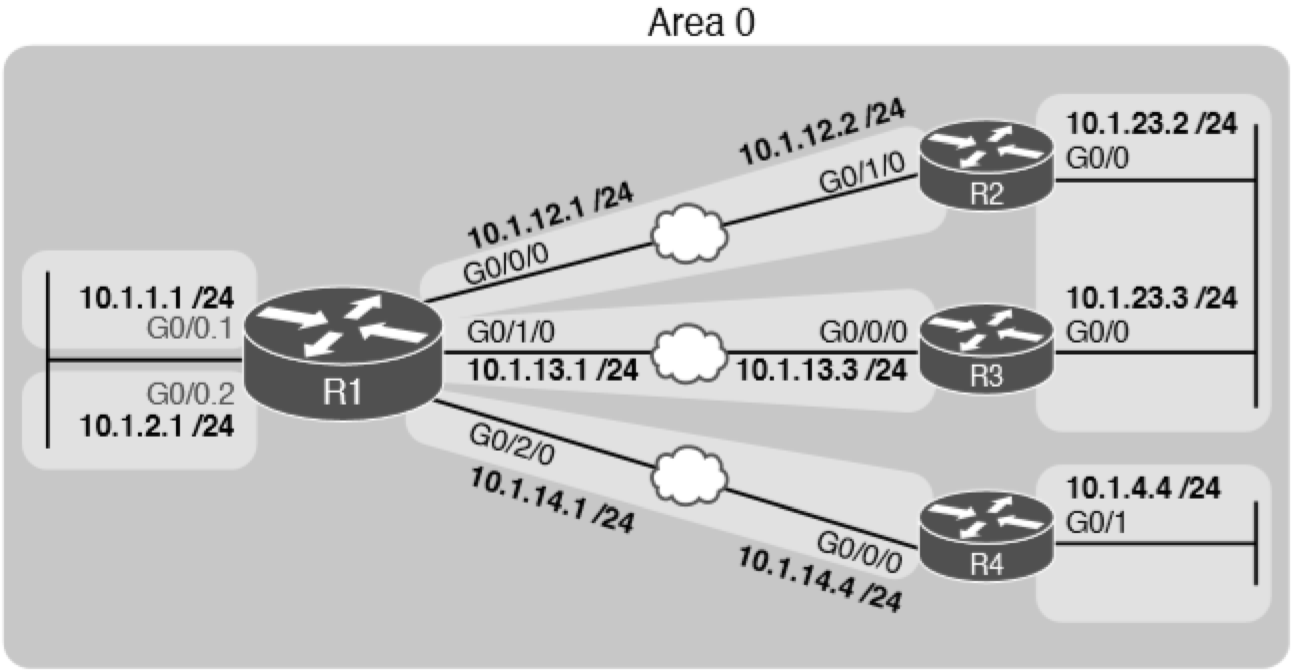

IPv4 Address Configuration on R1 (Including VLAN Trunking)

```
interface GigabitEthernet0/0.1
 encapsulation dot1q 1 native
 ip address 10.1.1.1 255.255.255.0
!
interface GigabitEthernet0/0.2
 encapsulation dot1q 2
 ip address 10.1.2.1 255.255.255.0
!
interface GigabitEthernet0/0/0
 ip address 10.1.12.1 255.255.255.0
!
interface GigabitEthernet0/1/0
 ip address 10.1.13.1 255.255.255.0
!
interface GigabitEthernet0/2/0
 ip address 10.1.14.1 255.255.255.0
```

Implementing Single-Area OSPF

```
router ospf 1
 network 10.0.0.0 0.255.255.255 area 0
```

- The **router ospf 1** global command puts the users into OSPF configuration mode and sets the OSPF *process-id* to 1.
- This *process-id* just needs to be unique on the local router and be between 1 and 65,535.
- The OSPF **network** command tells the router to find its local interfaces that match the first two parameters in the command.
- Those interfaces then discover neighbors, create neighbor relationships and assign the interface to the area listed.

Matching with the OSPF network Command

- The wildcard mask gives the local router its rules for matching its own interfaces:
  - Wildcard 0.0.0.0: Compare all 4 octets. In other words, the numbers must exactly match.
  - Wildcard 0.0.0.255: Compare the first 3 octets only. Ignore the last octet when comparing the numbers.
  - Wildcard 0.0.255.255: Compare the first 2 octets only. Ignore the last 2 octets when comparing the numbers.
  - Wildcard 0.255.255.255: Compare the first octet only. Ignore the last 3 octets when comparing the numbers.
  - Wildcard 255.255.255.255: Compare nothing—this wildcard mask means that all addresses will match the network command.

Matching with the OSPF network Command

Command	Logic in Command	Matched Interfaces
network 10.1.0.0 0.0.255.255	Match interface IP addresses that begin with 10.1	G0/0.341 G0/0.342 S0/0/0
network 10.0.0.0 0.255.255.255	Match interface IP addresses that begin with 10	G0/0.341 G0/0.342 S0/0/0
network 0.0.0.0 255.255.255.255	Match all interface IP addresses	G0/0.341 G0/0.342 S0/0/0
network 10.1.13.0 0.0.0.255	Match interface IP addresses that begin with 10.1.13	S0/0/0
network 10.1.3.1 0.0.0.0	Match one IP address: 10.1.3.1	G0/0.341

Matching with the OSPF network Command

```
! R1 configuration next - one network command enables OSPF
! on all three interfaces
router ospf 1
 network 10.1.0.0 0.0.255.255 area 0
```

```
! R2 configuration next - One network command per interface
router ospf 1
 network 10.1.12.2 0.0.0.0 area 0
 network 10.1.24.2 0.0.0.0 area 0
 network 10.1.2.2 0.0.0.0 area 0
```

OSPF Verification Commands

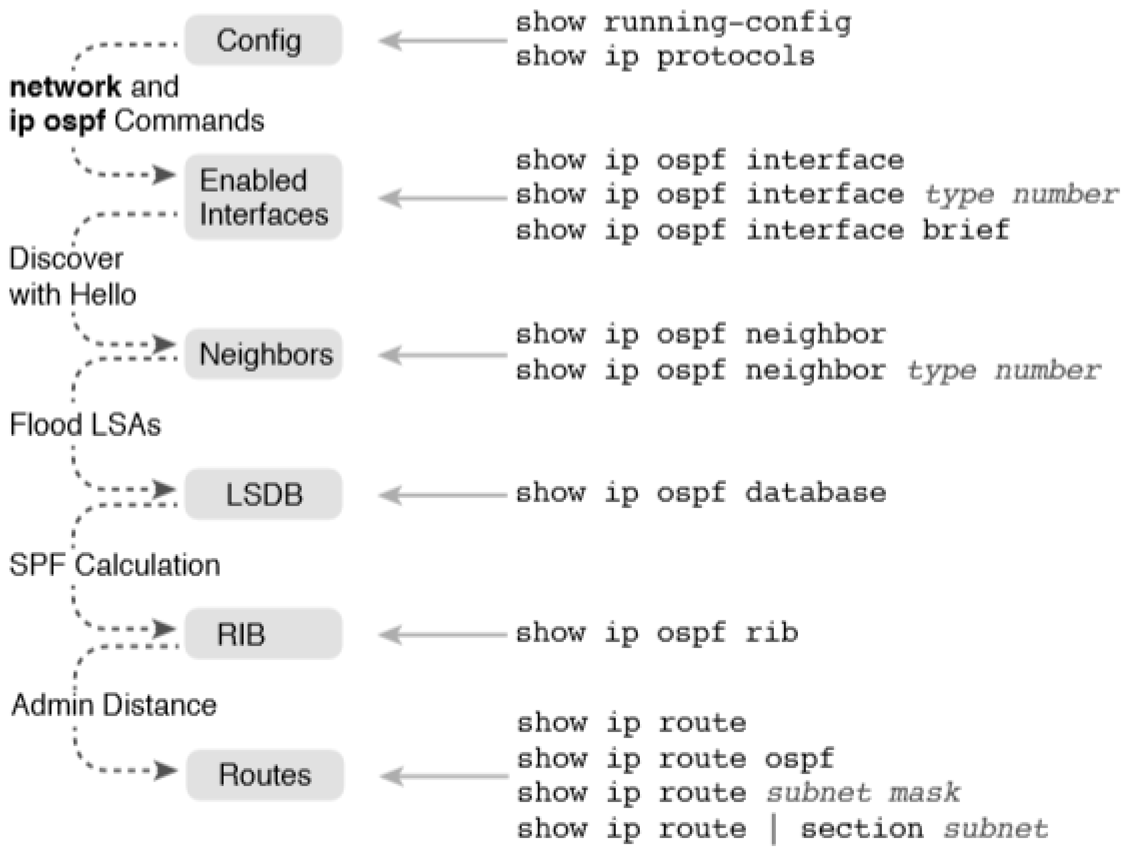

Verifying OSPFv2 Single Area

```
R3# show ip ospf neighbor

Neighbor ID     Pri   State       Dead Time   Address         Interface
1.1.1.1           0   FULL/  -    00:00:33    10.1.13.1       Serial0/0/0
10.1.24.4         1   FULL/DR     00:00:35    10.1.3.130      GigabitEthernet0/0.342
10.1.24.4         1   FULL/DR     00:00:36    10.1.3.4        GigabitEthernet0/0.341
```

- Interface: This is the local router's interface connected to the neighbor. For example, the first neighbor in the list is reachable through R3's S0/0/0 interface.
- Address: This is the neighbor's IP address on that link. Again, for this first neighbor, the neighbor, which is R1, uses IP address 10.1.13.1.
- State: While many possible states exist, for the details discussed in this chapter, FULL is the correct and fully working state in this case.
- Neighbor ID: This is the router ID of the neighbor.

Verifying OSPFv2 Single Area

```
R3# show ip ospf database

            OSPF Router with ID (10.1.13.3) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         498         0x80000006 0x002294 6
2.2.2.2         2.2.2.2         497         0x80000004 0x00E8C6 5
10.1.13.3       10.1.13.3       450         0x80000003 0x001043 4
10.1.24.4       10.1.24.4       451         0x80000003 0x009D7E 4

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.3.4        10.1.24.4       451         0x80000001 0x0045F8
10.1.3.130      10.1.24.4       451         0x80000001 0x00546B
```

- When OSPF is working correctly in the internetwork with a single-area design, all of the routers will have the same LSDB contents.
- For the purposes of the CCNA knowledge of the specifics about the output of this command are not required.

IPv4 Routes Added by OSPF

```
R4# show ip route
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
! Additional legend lines omitted for brevity

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
O        10.1.1.0/24 [110/2] via 10.1.14.1, 00:27:24, GigabitEthernet0/0/0
O        10.1.2.0/24 [110/2] via 10.1.14.1, 00:27:24, GigabitEthernet0/0/0
C        10.1.4.0/24 is directly connected, Vlan4
L        10.1.4.4/32 is directly connected, Vlan4
O        10.1.12.0/24 [110/2] via 10.1.14.1, 00:27:24, GigabitEthernet0/0/0
O        10.1.13.0/24 [110/2] via 10.1.14.1, 00:25:15, GigabitEthernet0/0/0
C        10.1.14.0/24 is directly connected, GigabitEthernet0/0/0
L        10.1.14.4/32 is directly connected, GigabitEthernet0/0/0
O        10.1.23.0/24 [110/3] via 10.1.14.1, 00:27:24, GigabitEthernet0/0/0
```

Router R3 Configuration and the *show ip protocols* Command

```
! First, a reminder of R3's configuration per Example 20-3:
router ospf 1
 network 10.1.13.3 0.0.0.0 area 0
 network 10.1.23.3 0.0.0.0 area 0
!
! The output from router R3:
R3# show ip protocols
*** IP Routing is NSF aware ***

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 3.3.3.3
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.13.3 0.0.0.0 area 0
    10.1.23.3 0.0.0.0 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    1.1.1.1              110      02:05:26
    4.4.4.4              110      02:05:26
    2.2.2.2              110      01:51:16
  Distance: (default is 110)
```

show ip ospf interface brief Command

```
R1# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
G10/0/0      1     0               10.1.12.1/24       1     BDR   1/1
G10/1/0      1     0               10.1.13.1/24       1     BDR   1/1
G10/2/0      1     0               10.1.14.1/24       1     DR    1/1
G10/0.2      1     0               10.1.2.1/24        1     DR    0/0
G10/0.1      1     0               10.1.1.1/24        1     DR    0/0
```

Configuring the OSPF Router ID

- OSPF speaking routers must have a router-ID (RID) for proper operation.
- By default, routers will choose an interface IP address to use as the RID.
- Typically, engineers prefer to choose each router's RID.
- A Cisco router uses the following process when the router reloads and brings up the OSPF process to find the RID:

1. If the **router-id *rid*** OSPF subcommand is configured, this value is used as the RID.
2. If any loopback interfaces have an IP address configured, and the interface has an interface status of up, the router picks the highest numeric IP address among these loopback interfaces.
3. The router picks the highest numeric IP address from all other interfaces whose interface status code (first status code) is up. (In other words, an interface in up/down state will be included by OSPF when choosing its router ID.)

OSPF Router ID Configuration Examples

```
! R1 Configuration first
router ospf 1
 router-id 1.1.1.1
 network 10.1.0.0 0.0.255.255 area 0
```

```
! R2 Configuration next
!
interface Loopback2
 ip address 2.2.2.2 255.255.255.255
```

```
R1# show ip ospf
 Routing Process "ospf 1" with ID 1.1.1.1
! lines omitted for brevity
```

Implementing Multiarea OSPF

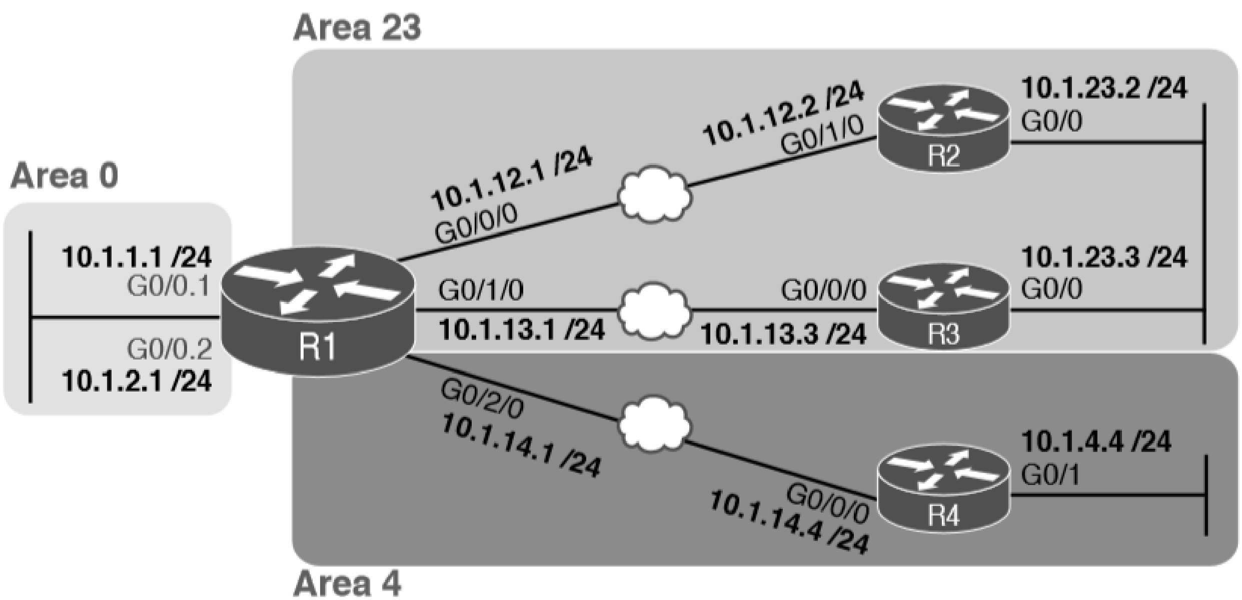


```
router ospf 1
 network 10.1.1.1 0.0.0.0 area 0
 network 10.1.2.1 0.0.0.0 area 0
 network 10.1.12.1 0.0.0.0 area 23
 network 10.1.13.1 0.0.0.0 area 23
 network 10.1.14.1 0.0.0.0 area 4
```

OSPF Single-Area Configuration

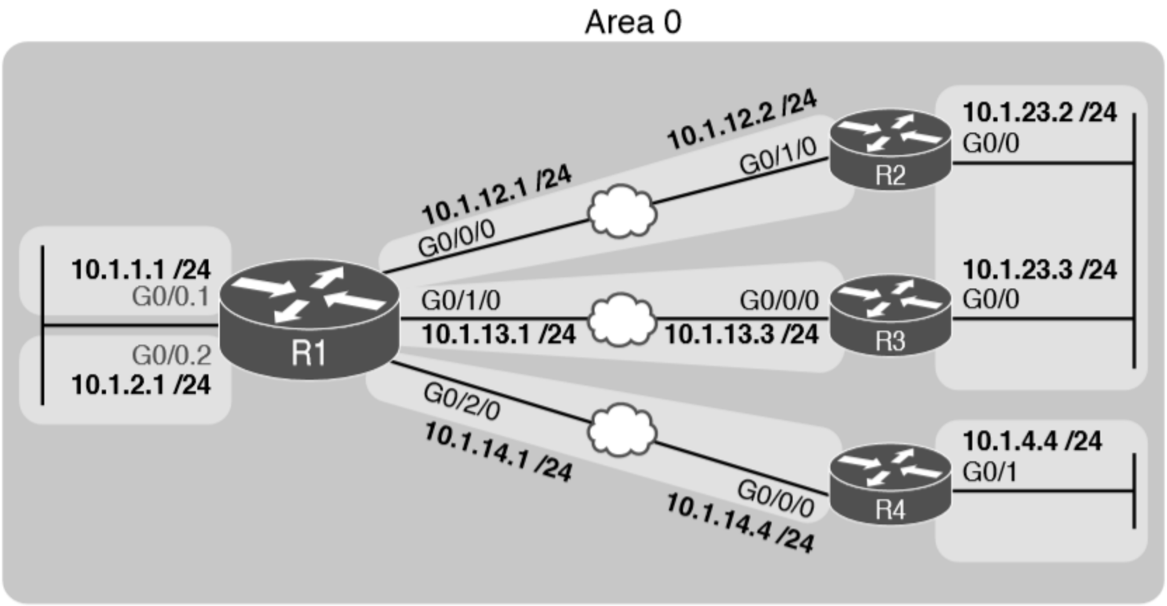

OSPF Single-Area Configuration

```
R1# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router ospf 1
R1(config-router)# no network 10.0.0.0 0.255.255.255 area 0
R1(config-router)#
*Apr  8 19:35:24.994: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on GigabitEthernet0/0/0
  from FULL to DOWN, Neighbor Down: Interface down or detached
*Apr  8 19:35:24.994: %OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on GigabitEthernet0/1/0
  from FULL to DOWN, Neighbor Down: Interface down or detached
*Apr  8 19:35:24.994: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on GigabitEthernet0/2/0
  from FULL to DOWN, Neighbor Down: Interface down or detached
R1(config-router)# interface g0/0.1
R1(config-subif)# ip ospf 1 area 0
R1(config-subif)# interface g0/0.2
R1(config-subif)# ip ospf 1 area 0
R1(config-subif)# interface g0/0/0
R1(config-if)# ip ospf 1 area 0
R1(config-if)#
*Apr  8 19:35:52.970: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on GigabitEthernet0/0/0
  from LOADING to FULL, Loading Done
R1(config-if)# interface g0/1/0
R1(config-if)# ip ospf 1 area 0
R1(config-if)#
*Apr  8 19:36:13.362: %OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on GigabitEthernet0/1/0
  from LOADING to FULL, Loading Done
R1(config-if)# interface g0/2/0
R1(config-if)# ip ospf 1 area 0
R1(config-if)#
*Apr  8 19:37:05.398: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on GigabitEthernet0/2/0
  from LOADING to FULL, Loading Done
R1(config-if)#
```

OSPF Passive Interfaces

- An OSPF interface configured as passive will do the following:
  - Quit sending OSPF Hellos on the interface.
  - Ignore received Hellos on the interface.
  - Do not form neighbor relationships over the interface.
- OSPF does not form neighbor relationships over interfaces configured as passive, but does still advertise about the subnet connected to that interface.

OSPF Passive Interfaces

```
! First, make each subinterface passive directly
router ospf 1
 passive-interface GigabitEthernet0/0.11
 passive-interface GigabitEthernet0/0.12

! Or, change the default to passive, and make the other interfaces
! not be passive

router ospf 1
 passive-interface default
 no passive-interface serial0/0/0
 no passive-interface serial0/0/1
```

- To configure an interface as passive there are two options:
  - Adding the **passive-interface *type number*** command under router configuration mode.
  - Configuring interfaces as passive by default with the **passive-interface default** command under global configuration mode and using the **no passive-interface *type number*** command to disable passive of specific interfaces.

OSPF Passive Interfaces

```
R1# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Gi0/0.12     1     0               10.1.1.129/25      1     DR    0/0
Gi0/0.11     1     0               10.1.1.1/25        1     DR    0/0
Se0/0/0      1     0               10.1.12.1/25       64    P2P   0/0
Se0/0/1      1     0               10.1.13.1/25       64    P2P   0/0

R1# show ip ospf interface g0/0.11
GigabitEthernet0/0.11 is up, line protocol is up
  Internet Address 10.1.1.1/25, Area 0, Attached via Network Statement
  Process ID 1, Router ID 10.1.1.129, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.1.1.129, Interface address 10.1.1.1
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    No Hellos (Passive interface)
! Lines omitted for brevity
```

- OSPF makes it interesting to figure out which interfaces are passive:
  - The **show ip ospf interface brief** command lists all interfaces on which OSPF is enabled *including passive interfaces.*
  - The **show ip ospf interface** command lists a single line that mentions that the interface is passive.

OSPF Default Routes

- The most classic case for using a routing protocol to advertise a default route has to do with an enterprise's connection to the Internet.

- The enterprise engineer uses the following design goals:

  - All routers learn specific routes for subnets inside the company; a default route is not needed when forwarding packets to these destinations.

  - One router connects to the Internet, and it has a default route that points toward the Internet.

  - All routers should dynamically learn a default route, used for all traffic going to the Internet, so that all packets destined to locations in the Internet go to the one router connected to the Internet.

OSPF Default Routes

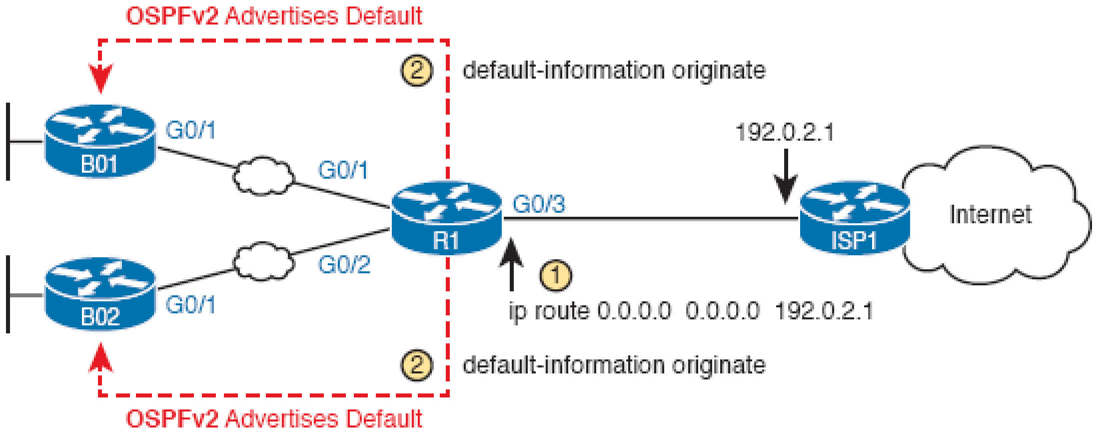

- R1 has a static default route with a next-hop address of the ISP router.
- R1 uses the OSPF **default-information originate** command to advertise a default route using OSPF to B1 and B2.
- R1 needs a default route, either defined as a static default route, learned from the ISP with DHCP or learned from the ISP with a routing protocol like eBGP unless the **always** parameter is used.

OSPF Default Routes

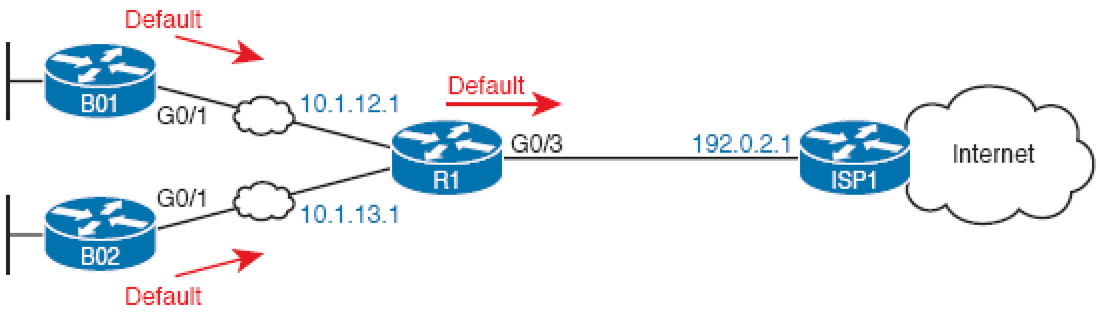

- The branch routers then place the learned OSPF default route into their routing tables.

OSPF Default Routes

```
! The next command is from Router R1. Note the static code for the default route
R1# show ip route static
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
! Rest of the legend omitted for brevity

Gateway of last resort is 192.0.2.1 to network 0.0.0.0

S*    0.0.0.0/0 [254/0] via 192.0.2.1
```

```
! The next command is from router B01; notice the External route code for the default
B01# 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
! Rest of the legend omitted for brevity

Gateway of last resort is 10.1.12.1 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 10.1.12.1, 00:20:51, GigabitEthernet0/1
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O        10.1.3.0/24 [110/3] via 10.1.12.1, 00:20:51, GigabitEthernet0/1
O        10.1.13.0/24 [110/2] via 10.1.12.1, 00:20:51, GigabitEthernet0/1
```

- R1 uses DHCP to learn its IP address on its Gi0/3 interface from the ISP.
- R1 then creates a static default route with the ISP router’s IP address of 192.0.2.1 as the next-hop address.

OSPF Metrics (Cost)

- Cisco routers allow three different ways to change the OSPF interface cost:
  - Directly, using the interface subcommand ip ospf cost x.
  - Using the default calculation per interface, and changing the interface bandwidth setting, which changes the calculated value.
  - Using the default calculation per interface, and changing the OSPF reference bandwidth setting, which changes the calculated value.
- IOS uses the Reference_bandwidth / Interface_bandwidth formula to choose an interface's cost.
- With this formula the following sequence of logic happens:

1. A higher interface bandwidth—that is, a faster bandwidth—results in a lower number in the calculation.
2. A lower number in the calculation gives the interface a lower cost.
3. An interface with a lower cost is more likely to be used by OSPF when calculating the best routes.

OSPF Metrics (Cost)

```
R1# show ip ospf interface brief
Interface     PID   Area          IP Address/Mask    Cost  State Nbrs F/C
Gi0/0.12      1     0             10.1.2.1/24        1     DR    0/0
Gi0/0.11      1     0             10.1.1.1/24        1     DR    0/0
Gi0/1         1     4             10.1.14.1/24       1     BDR   1/1
Se0/0/1       1     23            10.1.13.1/24       64    P2P   1/1
Se0/0/0       1     23            10.1.12.1/24       64    P2P   1/1
```

OSPF Metrics (Cost)

Interface	Interface Default Bandwidth (Kbps)	Formula (Kbps)	OSPF Cost
Serial	1544 Kbps	100,000 / 1544	64
Ethernet	10,000 Kbps	100,000 / 10,000	10
Fast Ethernet	100,000 Kbps	100,000/100,000	1
Gigabit Ethernet	1,000,000 Kbps	100,000/1,000,000	1
10 Gigabit Ethernet	10,000,000 Kbps	100,000/10,000,000	1
100 Gigabit Ethernet	100,000,000 Kbps	100,000/100,000,000	1

OSPF Load Balancing

- Typically when SPF calculates the metric of several routes, one route will have the lowest metric, so OSPF puts it into the routing table.
- When metrics tie for multiple routes to the same subnet, the router can put multiple equal-cost routes into the routing table. (By default, up to 4.)
- This default can be altered with the **maximum-paths *number*** command.
- Routers can load balance in one of two ways:
  - On a per-packet basis, where each new packet it forwarded out the next path in round robin.
  - On a per-destination bases, where packets are forwarded using the same path for each specific IP destination.