CCNPv7 ROUTE

Chapter 3 Lab 3-2, Multi-Area OSPFv2 and OSPFv3 with Stub Area

Topology

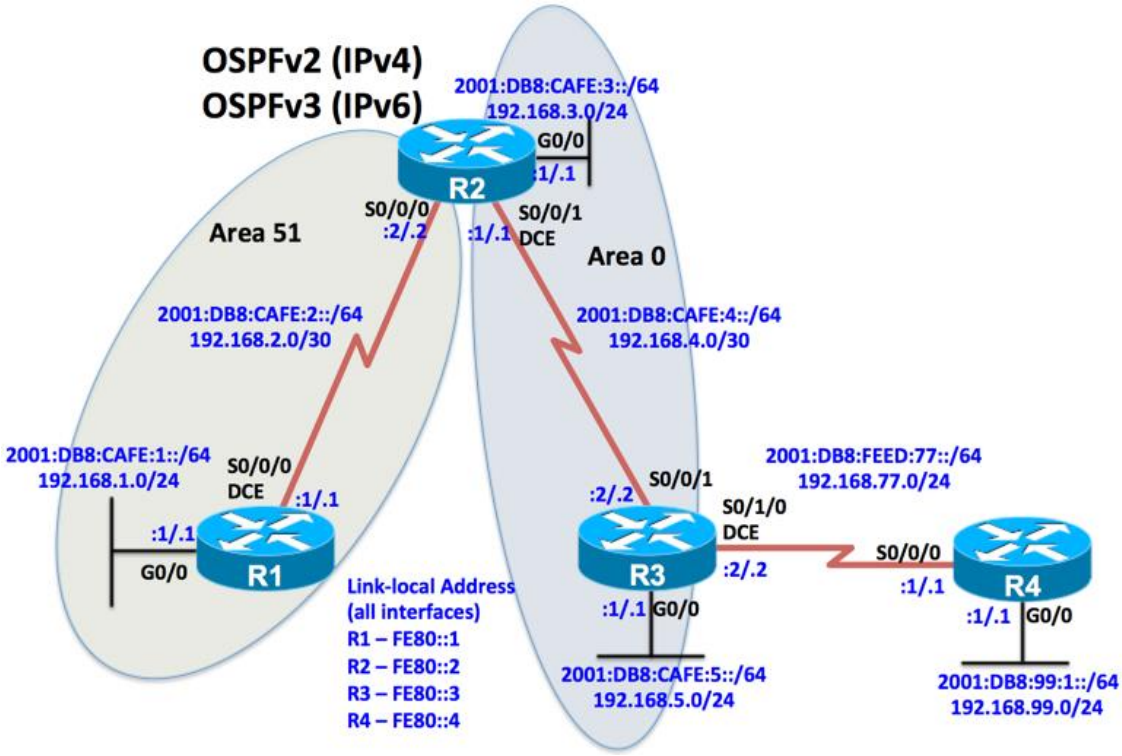

Objectives

- Configure multi-area OSPFv2 for IPv4.
- Configure multi-area OSPFv3 for IPv6
- Verify multi-area behavior.
- Configure stub and totally stubby areas for OSPFv2.
- Configure stub and totally stubby areas for OSPFv3.

Background

In this lab, you will configure the network with multi-area OSPFv2 routing for IPv4 and multi-area OSPFv3 routing for IPv6. For both OSPFv2 and OSPFv3, area 51 will be configured as a normal OSPF area, a stub area and then a totally stubby area.

Note: This lab uses Cisco 1941 routers with Cisco IOS Release 15.4 with IP Base. The switches are Cisco WS-C2960-24TT-L with Fast Ethernet interfaces, therefore the router will use routing metrics associated with a 100 Mb/s interface. 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

- 4 routers (Cisco IOS Release 15.2 or comparable)
- 4 switches (LAN interfaces)
- Serial and Ethernet cables

Step 0: Suggested starting configurations.

a. Apply the following configuration to each router along with the appropriate **hostname**. The **exec-timeout 0 0** command should only be used in a lab environment.

```
Router(config)# no ip domain-lookup
Router(config)# line con 0
Router(config-line)# logging synchronous
Router(config-line)# exec-timeout 0 0
```

Step 1: Configure the addressing and serial links.

a. Using the topology, configure the IPv4 and IPv6 addresses on the interfaces of each router.

```
R1(config)# interface GigabitEthernet0/0
R1(config-if)# ip address 192.168.1.1 255.255.255.0
R1(config-if)# ipv6 address FE80::1 link-local
R1(config-if)# ipv6 address 2001:DB8:CAFE:1::1/64
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)# interface Serial0/0/0
R1(config-if)# ip address 192.168.2.1 255.255.255.252
R1(config-if)# ipv6 address FE80::1 link-local
R1(config-if)# ipv6 address 2001:DB8:CAFE:2::1/64
R1(config-if)# clock rate 64000
R1(config-if)# no shutdown

R2(config)# interface GigabitEthernet0/0
R2(config-if)# ip address 192.168.3.1 255.255.255.0
R2(config-if)# ipv6 address FE80::2 link-local
R2(config-if)# ipv6 address 2001:DB8:CAFE:3::1/64
R2(config-if)# no shutdown
R2(config-if)# exit
R2(config)# interface Serial0/0/0
R2(config-if)# ip address 192.168.2.2 255.255.255.252
R2(config-if)# ipv6 address FE80::2 link-local
R2(config-if)# ipv6 address 2001:DB8:CAFE:2::2/64
R2(config-if)# no shutdown
R2(config-if)# exit
R2(config)# interface Serial0/0/1
R2(config-if)# ip address 192.168.4.1 255.255.255.252
R2(config-if)# ipv6 address FE80::2 link-local
R2(config-if)# ipv6 address 2001:DB8:CAFE:4::1/64
R2(config-if)# clock rate 64000
R2(config-if)# no shutdown

R3(config)# interface GigabitEthernet0/0
R3(config-if)# ip address 192.168.5.1 255.255.255.0
R3(config-if)# ipv6 address FE80::3 link-local
R3(config-if)# ipv6 address 2001:DB8:CAFE:5::1/64
R3(config-if)# no shutdown
R3(config-if)# exit
```

```
R3(config)# interface Serial0/0/1
R3(config-if)# ip address 192.168.4.2 255.255.255.252
R3(config-if)# ipv6 address FE80::3 link-local
R3(config-if)# ipv6 address 2001:DB8:CAFE:4::2/64
R3(config-if)# no shutdown
R3(config-if)# exit
R3(config)# interface Serial0/1/0
R3(config-if)# ip address 192.168.77.2 255.255.255.0
R3(config-if)# ipv6 address FE80::3 link-local
R3(config-if)# ipv6 address 2001:DB8:FEED:77::2/64
R3(config-if)# clock rate 64000
R3(config-if)# no shutdown
R3(config-if)#

R4(config)# interface Serial0/0/0
R4(config-if)# ip address 192.168.77.1 255.255.255.0
R4(config-if)# ipv6 address FE80::4 link-local
R4(config-if)# ipv6 address 2001:DB8:FEED:77::1/64
R4(config-if)# no shutdown
R4(config-if)# exit
R4(config)# interface gigabitethernet 0/0
R4(config-if)# ip address 192.168.99.1 255.255.255.0
R4(config-if)# ipv6 address 2001:db8:99:1::1/64
R4(config-if)# no shutdown
R4(config-if)# exit
R4(config)# ipv6 unicast-routing
R4(config)# ipv6 route 2001:DB8:CAFE::/48 2001:DB8:FEED:77::2
R4(config)# ip route 0.0.0.0 0.0.0.0 192.168.77.2
R4(config)#
```

b. Verify connectivity by pinging across each of the local networks connected to each router.

c. Issue the **show ip interface brief** and the **show ipv6 interface brief** command on each router. These commands display a brief listing of the interfaces, their status, and their IP addresses. Router R1 is shown as an example.

```
R1# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
Embedded-Service-Engine0/0 unassigned      YES unset  administratively down down
GigabitEthernet0/0         192.168.1.1     YES manual up                    up
GigabitEthernet0/1         unassigned      YES unset  administratively down down
Serial0/0/0                192.168.2.1     YES manual up                    up
Serial0/0/1                unassigned      YES unset  administratively down down
R1# show ipv6 interface brief
Em0/0                  [administratively down/down]
    unassigned
GigabitEthernet0/0     [up/up]
    FE80::1
    2001:DB8:CAFE:1::1
GigabitEthernet0/1     [administratively down/down]
    unassigned
Serial0/0/0            [up/up]
    FE80::1
    2001:DB8:CAFE:2::1
Serial0/0/1            [administratively down/down]
    unassigned
R1#
```

Step 2: Configure multi-area OSPFv2.

Create OSPFv2 process 1 on routers R1, R2 and R3. Configure the OSPF router ID on each router. Enable directly connected networks into the OSPF process using the **ip ospf** *process-id* **area** *area-id* interface command that is available with Cisco IOS version 12.3(11)T and later.

Note: The **show ip ospf** command should used to verify the OSPF router ID. If the OSPF router ID is using a 32-bit value other than the one specified by the **router-id** command, you can reset the router ID by using the **clear ip ospf** *pid* **process** command and re-verify using the command **show ip ospf**.

a. Configure R3 as an OSPFv2 router in area 0.

```
R3(config)# router ospf 1
R3(config-router)# router-id 3.3.3.3
R3(config-router)# exit
R3(config)# interface gigabitethernet 0/0
R3(config-if)# ip ospf 1 area 0
R3(config-if)# exit
R3(config)# interface serial 0/0/1
R3(config-if)# ip ospf 1 area 0
R3(config-if)#
```

Note: Another option is to use the OSPF **network** command in router configuration mode.

b. Configure R2 as an ABR router for area 0 and area 51. Interfaces S0/0/1 and G0/0 are in area 0, while interface S0/0/0 is in area 51.

```
R2(config)# router ospf 1
R2(config-router)# router-id 2.2.2.2
R2(config-router)# exit
R2(config)# interface serial 0/0/1
R2(config-if)# ip ospf 1 area 0
R2(config-if)# exit
R2(config)# interface gigabitethernet 0/0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# exit
R2(config)# interface serial 0/0/0
R2(config-if)# ip ospf 1 area 51
R2(config-if)#
```

What address on R2 is used to form the neighbor adjacency with R1? What type of IPv6 address is used to establish the adjacencies?

_____

c. Configure R1 as an internal OSPFv2 router in area 51.

```
R1(config)# router ospf 1
R1(config-router)# router-id 1.1.1.1
R1(config-router)# exit
R1(config)# interface serial 0/0/0
R1(config-if)# ip ospf 1 area 51
```

```
R1(config-if)# exit
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip ospf 1 area 51
R1(config-if)#
```

d. Verify that the routers have OSPFv2 neighbors using the **show ip ospf neighbors** command. The output for R2 is displayed.

```
R2# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:36    192.168.4.2     Serial0/0/1
1.1.1.1           0   FULL/  -        00:00:32    192.168.2.1     Serial0/0/0
R2#
```

d. Verify that router R3 can see all the IPv4 networks in the OSPFv2 routing domain using the **show ip route** command.

```
R3# 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
       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 not set

O IA  192.168.1.0/24 [110/129] via 192.168.4.1, 00:14:43, Serial0/0/1
      192.168.2.0/30 is subnetted, 1 subnets
O IA     192.168.2.0 [110/128] via 192.168.4.1, 00:20:16, Serial0/0/1
O     192.168.3.0/24 [110/65] via 192.168.4.1, 00:26:25, Serial0/0/1
      192.168.4.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.4.0/30 is directly connected, Serial0/0/1
L        192.168.4.2/32 is directly connected, Serial0/0/1
      192.168.5.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.5.0/24 is directly connected, GigabitEthernet0/0
L        192.168.5.1/32 is directly connected, GigabitEthernet0/0
      192.168.77.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.77.0/24 is directly connected, Serial0/1/0
L        192.168.77.2/32 is directly connected, Serial0/1/0
R3#
```

How many OSPFv2 intra-area routes area routes are in R3's IPv4 routing table? How many inter-area routes are in R3's IPv4 routing table?

_____

e. Issue the **show ip route** command on R2.

```
R2# 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
       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 not set

O     192.168.1.0/24 [110/65] via 192.168.2.1, 00:22:38, Serial0/0/0
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.2.0/30 is directly connected, Serial0/0/0
L        192.168.2.2/32 is directly connected, Serial0/0/0
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.3.0/24 is directly connected, GigabitEthernet0/0
L        192.168.3.1/32 is directly connected, GigabitEthernet0/0
      192.168.4.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.4.0/30 is directly connected, Serial0/0/1
L        192.168.4.1/32 is directly connected, Serial0/0/1
O     192.168.5.0/24 [110/65] via 192.168.4.2, 00:28:17, Serial0/0/1
R2#
```

Why doesn't R2 have any inter-area OSPFv2 routes in its routing table?

_____

_____

f. Configure an IPv4 default route on the ASBR R3 forwarding traffic to R4. Propagate the default routing into OSPFv2.

```
R3(config)# ip route 0.0.0.0 0.0.0.0 192.168.77.1
R3(config)# router ospf 1
R3(config-router)# default-information originate
R3(config-router)#
```

g. Issue the **show ip route static** command on R3 to verify the static route is in the IPv4 routing table.

```
R3# show ip route static
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 192.168.77.1 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 192.168.77.1
R3#
```

h. Configure an IPv4 static route on the ASBR, R3 for the 192.168.99.0/24 network on R4. Redistribute the static route into OSPFv2 using the **redistribute static subnets** command. The subnets parameter is used to include subnets and not just classful network addresses. The **redistribute** command is discussed in more detail in later chapters.

```
R3(config)# ip route 192.168.99.0 255.255.255.0 192.168.77.1
R3(config)# router ospf 1
R3(config-router)# redistribute static subnets
```

i. Issue the **show ip route ospf** command on R1 to verify that the default route and the redistributed static route are being advertised into the OSPFv2 domain.

```
R1# 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 192.168.2.2 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 192.168.2.2, 00:01:53, Serial0/0/0
O IA  192.168.3.0/24 [110/65] via 192.168.2.2, 00:06:09, Serial0/0/0
      192.168.4.0/30 is subnetted, 1 subnets
O IA     192.168.4.0 [110/128] via 192.168.2.2, 00:06:09, Serial0/0/0
O IA  192.168.5.0/24 [110/129] via 192.168.2.2, 00:06:09, Serial0/0/0
O E2  192.168.99.0/24 [110/20] via 192.168.2.2, 00:01:53, Serial0/0/0
R1#
```

What does the “E2” for the default route and the redistributed external route signify?

_____

_____

Step 3: Configure an OSPFv2 stub area.

a. Under the OSPFv2 process on R1 and R2, make area 51 a stub area using the **area** *area* **stub** command. The adjacency between the two routers might go down during the transition period, but it should come back up afterwards.

```
R1(config)# router ospf 1
R1(config-router)# area 51 stub

R2(config)# router ospf 1
R2(config-router)# area 51 stub
```

b. Confirm that both R1 and R2 are neighbors using the **show ip ospf neighbors** command.

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/  -        00:00:36    192.168.2.2     Serial0/0/0
R1#

R2# show ip ospf neighbor
```

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:37    192.168.4.2     Serial0/0/1
1.1.1.1           0   FULL/  -        00:00:38    192.168.2.1     Serial0/0/0
R2#
```

c. Issue the **show ip route ospf** command on R1. Notice that R1 still has a default route pointing toward R2 but with a different cost than it had prior to being configured in a stub area. This is not the default route propagated by the ASBR R3, but the default route injected by the ABR of the stub area. Also, R1 does not receive any external routes, so it no longer has the external network 192.168.99.0/24 in its routing table. Stub routers continue to receive inter-area routes from area 0.

```
R1# 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 192.168.2.2 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/65] via 192.168.2.2, 00:06:09, Serial0/0/0
O IA  192.168.3.0/24 [110/65] via 192.168.2.2, 00:06:09, Serial0/0/0
      192.168.4.0/30 is subnetted, 1 subnets
O IA     192.168.4.0 [110/128] via 192.168.2.2, 00:06:09, Serial0/0/0
O IA  192.168.5.0/24 [110/129] via 192.168.2.2, 00:06:09, Serial0/0/0
R1#
```

d. View the output of the **show ip ospf** command on ABR R2 to see what type each area is and the number of interfaces in each area.

```
R2# show ip ospf
 Routing Process "ospf 1" with ID 2.2.2.2
 Start time: 01:49:34.272, Time elapsed: 02:04:19.324
 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 area border router
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Incremental-SPF disabled
 Minimum LSA interval 5 secs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Number of external LSA 2. Checksum Sum 0x0174F7
 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 2. 1 normal 1 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 100 mbps
   Area BACKBONE(0)
       Number of interfaces in this area is 2
       Area has no authentication
       SPF algorithm last executed 00:23:27.416 ago
       SPF algorithm executed 20 times
       Area ranges are
       Number of LSA 6. Checksum Sum 0x0413D3
       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
   Area 51
       Number of interfaces in this area is 1
       It is a stub area
       Generates stub default route with cost 1
       Area has no authentication
       SPF algorithm last executed 00:23:17.416 ago
       SPF algorithm executed 4 times
       Area ranges are
       Number of LSA 6. Checksum Sum 0x02E70A
       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#
```

What are the advantages of having a router receive a default route rather than a more specific route?

_____

_____

Why do all routers in a stub area need to know that the area is a stub?

_____

_____

_____

Step 4: Configure a totally stubby area.

A modified version of a stubby area is a totally stubby area. A totally stubby area ABR only allows in a single, default route from the backbone, injected by the ABR. To configure a totally stubby area, you only need to change a command at the ABR, R2 in this scenario. Under the router OSPFv2 process, you will enter the **area 51 stub no-summary** command to replace the existing stub command for area 51. The **no-summary** option tells the router that this area will not receive summary (inter-area) routes.

a. To see how this works, issue the **show ip route ospf** command on R1. Notice the inter-area routes, in addition to the default route generated by R2.

```
R1# 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 192.168.2.2 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/65] via 192.168.2.2, 00:28:13, Serial0/0/0
O IA  192.168.3.0/24 [110/65] via 192.168.2.2, 00:28:13, Serial0/0/0
      192.168.4.0/30 is subnetted, 1 subnets
O IA     192.168.4.0 [110/128] via 192.168.2.2, 00:28:13, Serial0/0/0
O IA  192.168.5.0/24 [110/129] via 192.168.2.2, 00:28:13, Serial0/0/0
R1#
```

b. Look at the output of the **show ip ospf database** command on R2 to see which LSAs are in its OSPFv2 database.

```
R2# show ip ospf database

            OSPF Router with ID (2.2.2.2) (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         2231        0x80000002 0x00EECE 2
2.2.2.2         2.2.2.2         41          0x8000000D 0x00E63E 3
3.3.3.3         3.3.3.3         385         0x80000007 0x0071B1 3

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.1.0     1.1.1.1         2241        0x80000002 0x00B616
192.168.1.0     2.2.2.2         1838        0x80000001 0x001D6C
192.168.2.0     2.2.2.2         41          0x80000002 0x00F397

                Router Link States (Area 51)

Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         1847        0x8000000B 0x0043F8 3
2.2.2.2         2.2.2.2         1841        0x8000000A 0x009C16 2

                Summary Net Link States (Area 51)

Link ID         ADV Router      Age         Seq#       Checksum
0.0.0.0         2.2.2.2         41          0x80000002 0x0073C1
192.168.3.0     2.2.2.2         41          0x80000007 0x00962D
192.168.4.0     2.2.2.2         41          0x80000007 0x00F194
192.168.5.0     2.2.2.2         41          0x80000007 0x00037E

                Type-5 AS External Link States
```

```
Link ID         ADV Router      Age         Seq#       Checksum Tag
0.0.0.0         3.3.3.3         385         0x80000003 0x00DCC7 1
192.168.99.0    3.3.3.3         385         0x80000002 0x009432 0
R2#
```

c. Enter the **area 51 stub no-summary** command on R2 (the ABR) under the OSPF process.

```
R2(config)# router ospf 1
R2(config-router)# area 51 stub no-summary
```

d. Go back to R1 and issue the **show ip route ospf** command. Notice that it shows only one incoming route from the ABR R2. The default route is injected by the ABR R2. There are no inter-area OSPFv2 routes and no external OSPFv2 routes.

```
R1# 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 192.168.2.2 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/65] via 192.168.2.2, 00:01:14, Serial0/0/0
R1#
```

e. Examine the output of the **show ip ospf database** command to see which routes are in area 51. You may need to clear the OSPFv2 process to reset the entries in the OSPF LSDB.

```
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: yes
*Oct  8 03:56:06.802: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Oct  8 03:56:06.894: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0/0 from
LOADING to FULL, Loading Done
R1#
R1# show ip ospf database

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 51)

Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         7           0x8000000D 0x003FFA 3
2.2.2.2         2.2.2.2         284         0x8000000B 0x009A17 2

                Summary Net Link States (Area 51)

Link ID         ADV Router      Age         Seq#       Checksum
0.0.0.0         2.2.2.2         330         0x80000004 0x006FC3
R1#
```

What are the advantages of making an area totally stubby instead of a regular stub area? What are the disadvantages?

_____

_____

_____

_____

Why did only the ABR need to know that the area was totally stubby rather than all routers in the area?

_____

_____

_____

Step 5: Configure multi-area OSPFv3.

Traditional OSPFv3 implements OSPF routing for IPv6. In our dual-stack (IPv4/IPv6) environment we have previously configured OSPFv2 for routing IPv4 and now we will configure OSPFv3 for routing IPv6.

a. OSPFv3 messages are sourced from the router's IPv6 link-local address. Earlier in this lab, IPv6 GUA and link-local addresses were statically configured on each router's interface. The link-local addresses were configured to make these addresses more recognizable than being automatically created using EUI-64. Issue the **show ipv6 interface brief** command to verify the GUA and link-local addresses on the router's interfaces.

```
R1# show ipv6 interface brief
Em0/0                  [administratively down/down]
    unassigned
GigabitEthernet0/0     [up/up]
    FE80::1
    2001:DB8:CAFE:1::1
GigabitEthernet0/1     [administratively down/down]
    unassigned
Serial0/0/0            [up/up]
    FE80::1
    2001:DB8:CAFE:2::1
Serial0/0/1            [administratively down/down]
    unassigned
R1#
```

b. IPv6 routing is disabled by default. The Cisco IOS version used with the routers in this lab has IPv6 CEF enabled by default once IPv6 routing is enabled. To enable IPv6 routing, use the **ipv6 unicast-routing** command in global configuration mode. Use the **show ipv6 cef** command to verify whether IPv6 CEF is enabled. If you need to enable IPv6 CEF, use the **ipv6 cef** command. If IPv6 CEF is disabled you will see the an IOS message similar to "%IPv6 CEF not running:. Enter these commands on routers R1, R2 and R3. IPv6 routing on R4 has been enabled in Step 1.

```
R1(config)# ipv6 unicast-routing
R1(config)# end
R1# show ipv6 cef
```

```
::/0
  no route
::/127
  discard
2001:DB8:CAFE:1::/64
  attached to GigabitEthernet0/0
2001:DB8:CAFE:1::1/128
  receive for GigabitEthernet0/0
2001:DB8:CAFE:2::/64
  attached to Serial0/0/0
2001:DB8:CAFE:2::1/128
  receive for Serial0/0/0
FE80::/10
  receive for Null0
FF00::/8
  multicast
R1#

R2(config)# ipv6 unicast-routing

R3(config)# ipv6 unicast-routing
```

c. Configure the OSPFv3 process on each router. Similar to OSPFv2, the process ID does not have to match other routers to form neighbor adjacencies. Configure the 32-bit OSPFv3 router ID on each router. The OSPFv3 router ID uses the same process as OSPFv2 and is required if there are no IPv4 addresses configured on the router.

Note: The **show ipv6 ospf** command should used to verify the OSPF router ID. If the OSPFv3 router ID is uses a 32-bit value other than the one specified by the **router-id** command, you can reset the router ID by using the **clear ipv6 ospf** *pid* **process** command and re-verify using the command **show ipv6 ospf**.

```
R1(config)# ipv6 router ospf 2
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# exit
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ipv6 ospf 2 area 51
R1(config-if)# exit
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 ospf 2 area 51
R1(config-if)#


R2(config)# ipv6 router ospf 2
R2(config-rtr)# router-id 2.2.2.2
R2(config-rtr)# exit
R2(config)# interface serial 0/0/1
R2(config-if)# ipv6 ospf 2 area 0
R2(config-if)# exit
R2(config)# interface gigabitethernet 0/0
R2(config-if)# ipv6 ospf 2 area 0
```

```
R2(config-if)# exit
R2(config)# interface serial 0/0/0
R2(config-if)# ipv6 ospf 2 area 51
R2(config-if)#

R3(config)# ipv6 router ospf 2
R3(config-rtr)# router-id 3.3.3.3
R3(config-rtr)# exit
R3(config)# interface gigabitethernet 0/0
R3(config-if)# ipv6 ospf 2 area 0
R3(config-if)# exit
R3(config)# interface serial 0/0/1
R3(config-if)# ipv6 ospf 2 area 0
R3(config-if)#
```

d. Verify that you have OSPFv3 neighbors with the **show ipv6 ospf neighbor** command. The output for R2 is displayed.

```
R2# show ipv6 ospf neighbor

            OSPFv3 Router with ID (2.2.2.2) (Process ID 2)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           0   FULL/  -        00:00:36    6               Serial0/0/1
1.1.1.1           0   FULL/  -        00:00:34    6               Serial0/0/0
R2#
```

e. View the OSPF routes in the IPv6 routing table on all three routers with the **show ipv6 route ospf** command.

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OI  2001:DB8:CAFE:3::/64 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:4::/64 [110/128]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:5::/64 [110/129]
     via FE80::2, Serial0/0/0
R1#

R2# show ipv6 route ospf
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
```

```
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
O   2001:DB8:CAFE:1::/64 [110/65]
     via FE80::1, Serial0/0/0
O   2001:DB8:CAFE:5::/64 [110/65]
     via FE80::3, Serial0/0/1
R2#

R3# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OI  2001:DB8:CAFE:1::/64 [110/129]
     via FE80::2, Serial0/0/1
OI  2001:DB8:CAFE:2::/64 [110/128]
     via FE80::2, Serial0/0/1
O   2001:DB8:CAFE:3::/64 [110/65]
     via FE80::2, Serial0/0/1
R3#
```

f. Configure an IPv6 default route on the ASBR R3 forwarding traffic to R4. Propagate the default routing into OSPFv3.

```
R3(config)# ipv6 route ::/0 2001:db8:feed:77::1
R3(config)# ipv6 router ospf 2
R3(config-rtr)# default-information originate
R3(config-rtr)#
```

g. Configure an IPv6 static route on the ASBR R3 for the 2001:DB8:99:1::/64 prefix on R4. Redistribute the static route into OSPFv3.

```
R3(config)# ipv6 route 2001:db8:99:1::/64 2001:db8:feed:77::1
R3(config)# ipv6 router ospf 2
R3(config-rtr)# redistribute static
R3(config-rtr)#
```

h. Issue the **show ipv6 route static** command on R3 to verify both static routes is in the IPv6 routing table.

```
R3# show ipv6 route static
IPv6 Routing Table - default - 12 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
```

```
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
S   ::/0 [1/0]
     via 2001:DB8:FEED:77::1
S   2001:DB8:99:1::/64 [1/0]
     via 2001:DB8:FEED:77::1
R3#
```

i. Issue the **show ipv6 route ospf** command on R1 to verify that the default route and the redistributed static route are now being advertised into the OSPFv3 domain.

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OE2 ::/0 [110/1], tag 2
     via FE80::2, Serial0/0/0
OE2 2001:DB8:99:1::/64 [110/20]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:3::/64 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:4::/64 [110/128]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:5::/64 [110/129]
     via FE80::2, Serial0/0/0
R1#
```

Step 6: Configure an OSPFv3 stub area.

a. Configuring stub areas for OSPFv3 is similar to that for OSPFv2. The stub area functionality is the same for OSPFv2 and OSPFv3. Under the OSPFv3 process on R1 and R2, make area 51 a stub area using the **area** *area* **stub** command. The adjacency between the two routers might go down during the transition period, but it should come back up afterwards.

```
R1(config)# ipv6 router ospf 2
R1(config-rtr)# area 51 stub

R2(config)# ipv6 router ospf 2
R2(config-rtr)# area 51 stub
```

b. Confirm that both R1 and R2 are neighbors using the **show ipv6 ospf neighbors** command.

```
R1# show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 2)
```

```
Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
2.2.2.2           0   FULL/  -        00:00:36    5               Serial0/0/0
R1#

R2# show ipv6 ospf neighbor

            OSPFv3 Router with ID (2.2.2.2) (Process ID 2)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           0   FULL/  -        00:00:35    6               Serial0/0/1
1.1.1.1           0   FULL/  -        00:00:34    6               Serial0/0/0
R2#
```

c. To verify that the stub area functionality is the same in OSPFv3 as in OSPFv2 issue the **show ipv6 route ospf** command on R1. Similar to OSPFv2, notice that R1 still has a default route pointing toward R2 but with a different cost than it had prior to being configured in a stub area. Again, this is not the default route propagated by the ASBR R3, but the default route injected by the ABR of the stub area. R1 also does not receive any external routes, so it no longer has the 2001:DB8:99:1::/64 prefix in its routing table. Stub routers continue to receive inter-area routes.

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OI  ::/0 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:3::/64 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:4::/64 [110/128]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:5::/64 [110/129]
     via FE80::2, Serial0/0/0
R1#
```

Step 7: Configure a totally stubby area.

As mentioned earlier in the lab, a totally stubby area ABR only allows in a single, default route from the backbone, injected by the ABR. Configuring a totally stubby area, you only need to change a command at the ABR, R2 in this scenario. Similar commands used to configure a totally stubby area for the OSPFv2 process are used for OSPFv3.

a. First, issue the **show ipv6 route ospf** command on R1 to verify that inter-area routes, in addition to the default route are being sent by R2.

```
R1#show ipv6 route ospf
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
```

```
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OI  ::/0 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:3::/64 [110/65]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:4::/64 [110/128]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:5::/64 [110/129]
     via FE80::2, Serial0/0/0
R1#
```

b. Enter the **area 51 stub no-summary** command on R2 (the ABR) under the OSPFv3 process.

```
R2(config)# ipv6 router ospf 2
R2(config-rtr)# area 51 stub no-summary
```

c. On R1 and issue the **show ipv6 route ospf** command. Similar to OSPFv2, there is only one incoming route from the ABR R2. The default route is injected by the ABR R2. There are no inter-area OSPFv3 routes and no external OSPFv3 routes.

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       a - Application
OI  ::/0 [110/65]
     via FE80::2, Serial0/0/0
R1#
```

d. View the output of the **show ipv6 ospf** command on ABR R2 to see what type each area is and the number of interfaces in each area.

```
R2# show ipv6 ospf
 Routing Process "ospfv3 2" with ID 2.2.2.2
 Supports NSSA (compatible with RFC 3101)
 Event-log enabled, Maximum number of events: 1000, Mode: cyclic
 It is an area border router
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Minimum LSA interval 5 secs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Retransmission limit dc 24 non-dc 24
 Number of external LSA 2. Checksum Sum 0x00FD33
 Number of areas in this router is 2. 1 normal 1 stub 0 nssa
 Graceful restart helper support enabled
 Reference bandwidth unit is 100 mbps
 RFC1583 compatibility enabled
```

```
    Area BACKBONE(0)
        Number of interfaces in this area is 2
        SPF algorithm executed 7 times
        Number of LSA 9. Checksum Sum 0x0539E9
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 51
        Number of interfaces in this area is 1
        It is a stub area, no summary LSA in this area
        Generates stub default route with cost 1
        SPF algorithm executed 5 times
        Number of LSA 7. Checksum Sum 0x028798
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0

R2#
```

What is meant by the high-lighted output for Area 51?

_____

_____

_____

_____