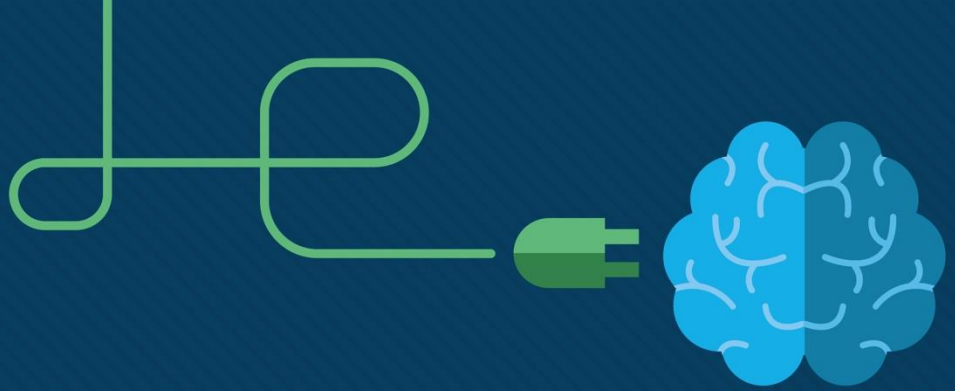

Chapter 10: OSPF Tuning and Troubleshooting

CCNA Routing and Switching

Scaling Networks v6.0

Chapter 10 - Sections & Objectives

- 10.1 Advanced Single-Area OSPF Configurations
  - Configure the OSPF interface priority to influence the DR/BDR election.
  - Configure OSPF to propagate a default route.
  - Configure OSPF interface settings to improve network performance.

- 10.2 Troubleshooting Single-Area OSPF Implementations
  - Explain the process and tools used to troubleshoot a single-area OSPF network.
  - Troubleshoot missing route entries in the single-area OSPFv2 routing table.
  - Troubleshoot missing route entries in a single-area OSPFv3 routing table.
  - Troubleshoot missing route entries in multiarea OSPFv2 and OSPFv3 routing tables.

10.1 Advanced Single-Area OSPF Configurations

OSPF Network Types

Point-to-point	Broadcast multiaccess	Nonbroadcast multiaccess (NBMA)	Point-to-multipoint	Virtual links
• Two routers interconnected over a common link. • No other routers are on the link. • Common configuration in WAN links.	• Multiple routers interconnected over an Ethernet network. • Ethernet LANs are the most common example of broadcast multiaccess networks.	• Multiple routers interconnected in a network that does not allow broadcasts. • The Frame Relay WAN protocol is an example NBMA network.	• Multiple routers interconnected in a hub-and-spoke topology over an NBMA network. • Often used to connect branch sites (spokes) to a central site (hub).	• Special OSPF network used to interconnect distant OSPF areas to the backbone area.
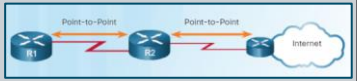	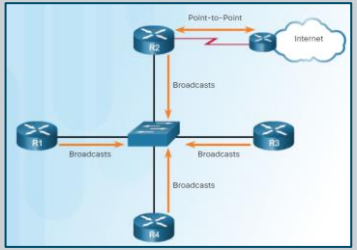	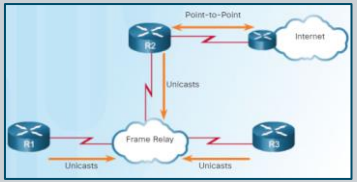	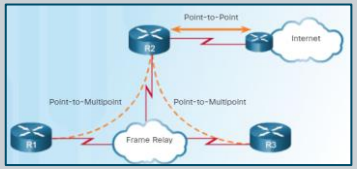	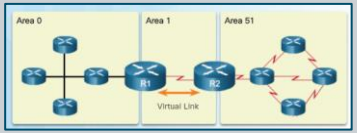

Challenges in Multiaccess Networks

- Multiaccess networks create two challenges regarding the flooding of OSPF LSAs.

<table>
<tr><th>OSPF Challenges</th><th>Description</th><th></th></tr>
<tr><td>Creation of multiple adjacencies</td><td>• Ethernet networks could potentially interconnect many OSPF routers creating numerous adjacencies with every router.
• Use the n (n-1) / 2 formula to calculate the number of adjacencies required for any number of routers (i.e., n) on a multiaccess network.</td><td>
<table>
<tr><th>Routers
n</th><th>Adjacencies
n (n-1)/2</th></tr>
<tr><td>4</td><td>6</td></tr>
<tr><td>5</td><td>10</td></tr>
<tr><td>10</td><td>45</td></tr>
<tr><td>20</td><td>190</td></tr>
<tr><td>50</td><td>1225</td></tr>
</table>
</td></tr>
<tr><td>Extensive flooding of LSAs</td><td colspan="2">• Link-state routers flood their link-state packets when OSPF is initialized, or when there is a change in the topology.
• This flooding can become excessive.</td></tr>
</table>

OSPF Designated Router

- On multiaccess networks, OSPF elects a DR to be the collection and distribution point for LSAs sent and received.
  - A BDR is also elected in case the DR fails. If the DR stops producing Hello packets, the BDR promotes itself and assumes the role of DR.
  - All other non-DR or BDR routers become DROTHER (a router that is neither the DR nor the BDR) and DROTHERs only form full adjacencies with the DR and BDR in the network.
  - Instead of flooding LSAs to all routers in the network, DROTHERs only send their LSAs to the DR and BDR using the multicast address 224.0.0.6 (all DR routers).

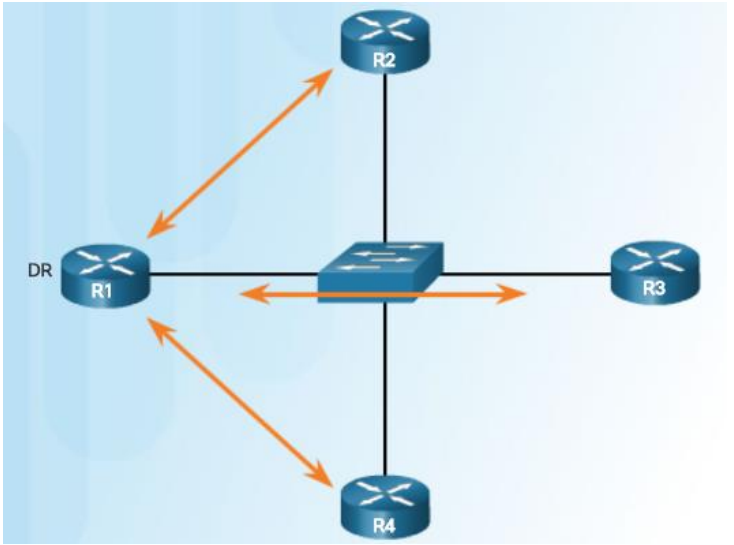

Verifying DR/BDR Roles

- OSPF has automatically elected a DR and BDR.

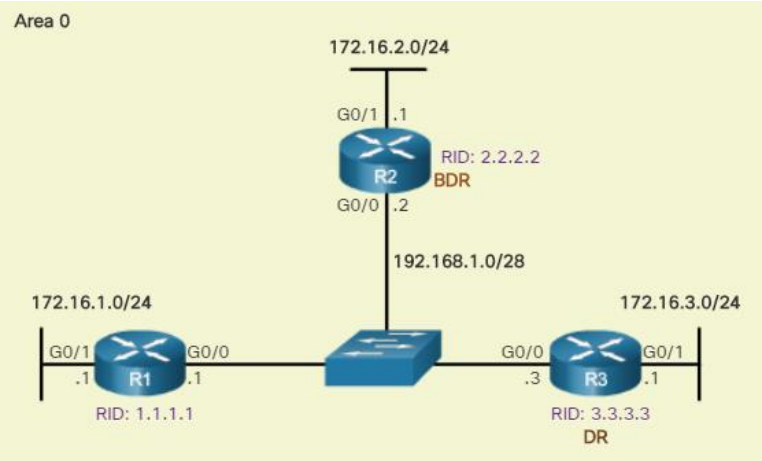

- R3 is the DR because of its higher router ID.
- R2 is the BDR because of its 2nd highest router ID.
- R1 is a DROTHER.

Verify the roles of the OSPFv2 router using the **show ip ospf interface** command.

```
R3# show ip ospf interface GigabitEthernet 0/0
GigabitEthernet0/0 is up, line protocol is up
 Internet Address 192.168.1.3/28,Area 0,Attached via Network Statement
 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
 Transmit Delay is 1 sec, State DR Priority 1
 Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
 Backup Designated router (ID) 2.2.2.2, Interface address 192.168.1.2
 Timer intervals configured, Hello 10, Dead 40, Wait 40, 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 3, 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
   Adjacent with neighbor 2.2.2.2  (Backup Designated Router)
 Suppress hello for 0 neighbor(s)
R3#
```

Note: For the equivalent OSPFv3 command, simply substitute **ip** with **ipv6**.

Verifying DR/BDR Adjacencies

- Verify OSPFv2 adjacencies using **show ip ospf neighbor**.

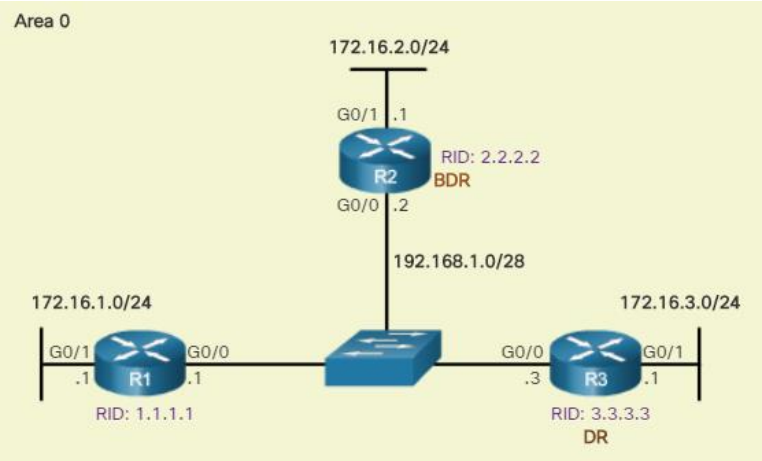


```
R1# show ip ospf neighbor

Neighbor ID Pri State       Dead Time  Address     Interface
2.2.2.2     1  FULL/BDR     00:00:36  192.168.1.2 GigabitEthernet0/0
3.3.3.3     1  FULL/DR      0:00:35   192.168.1.3 GigabitEthernet0/0
R1#
```

```
R2# show ip ospf neighbor

Neighbor ID Pri State       Dead Time  Address     Interface
1.1.1.1     1  FULL/DROTHER 00:00:31  192.168.1.1 GigabitEthernet0/0
3.3.3.3     1  FULL/DR      00:00:39  192.168.1.3 GigabitEthernet0/0
R2#
```

```
R3# show ip ospf neighbor

Neighbor ID Pri State       Dead Time  Address     Interface
1.1.1.1     1  FULL/DROTHER 00:00:34  192.168.1.1 GigabitEthernet0/0
2.2.2.2     1  FULL/BDR     00:00:39  192.168.1.2 GigabitEthernet0/0
R3#
```

- Routers can be in the following states:
  - FULL/DROTHER
  - FULL/DR
  - FULL/BDR
  - 2-WAY/DROTHER

Note: For the equivalent OSPFv3 command, simply substitute **ip** with **ipv6**.

Default DR/BDR Election Process

- 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 but the default priority is 1.
2. If the interface priorities are equal, then the router with the highest router ID is elected the DR.
   - The router with the second highest router ID is the BDR.

- Recall that the router ID is determined in one of three ways:
  - The router ID can be manually configured.
  - If no router IDs are configured, the router ID is determined by the highest loopback IPv4 address.
  - If no loopback interfaces are configured, the router ID is determined by the highest active IPv4 address.

Note: In an IPv6 network, if there are no IPv4 addresses configured on the router, then the router ID must be manually configured with the **router-id** *rid* router configuration command; otherwise, OSPFv3 does not start.

DR/BDR Election Process

- After the DR is elected, it remains the DR until one of the following events occurs:
  - The DR fails.
  - The OSPF process on the DR fails or is stopped.
  - The multiaccess interface on the DR fails or is shutdown.

- OSPF DR and BDR elections are not pre-emptive.
  - If a new router with a higher priority is added to the network after the DR election, the newly added router does not take over the DR or the BDR role because those roles have already been assigned.
  - If the DR fails, the BDR is automatically promoted to DR even if another DROTHER with a higher priority or router ID is added to the network after the initial DR/BDR election.
  - After a BDR is promoted to DR, a new BDR election occurs and the DROTHER with the higher priority or router ID is elected as the new BDR.

The OSPF Priority

- To control the DR and BDR election, the priority of an interface can be configured using:
  - **ip ospf priority** *value* - OSPFv2 interface command
  - **ipv6 ospf priority** *value* - OSPFv3 interface command

- The *value* can be:
  - **0** - Does not become a DR or BDR.
  - **1 – 255** - The higher the priority value, the more likely the router becomes the DR or BDR on the interface.

Changing the OSPF Priority

- Changing the priority value on an interface from 1 to a higher value would enable the router to become a DR or BDR router during the next election.
  - Priority changes do not automatically take effect because the DR and BDR are already elected.

- To force an election, use one of the following methods:
  - Shutdown the router interfaces and then re-enable them starting with the desired DR, then the desired BDR, and then all other routers.
  - Reset the OSPF process using the **clear ip ospf process** privileged EXEC mode command on all routers.

Default Route Propagation

- An OSPF ASBR router (aka, edge, entrance, or the gateway router) connects to the Internet and can be configured to propagate a default route to other routers in the OSPF routing domain.

- To propagate a default route, R2 is configured with:
  - A default static route.
    - **ip route 0.0.0.0 0.0.0.0** {*ip-address* **|** *exit-intf*} command.
  - The **default-information originate** router config mode command to propagate the default route in OSPF updates.

```
R2(config)# ip route 0.0.0.0 0.0.0.0 209.165.200.226
R2(config)#
R2(config)# router ospf 10
R2(config-router)# default-information originate
R2(config-router)# end
R2#
```

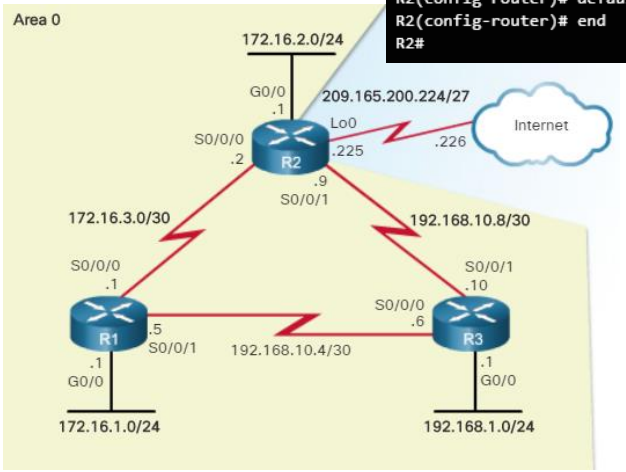

Verifying the Propagated Default Route

- Use the **show ip route** command to verify the default route settings..

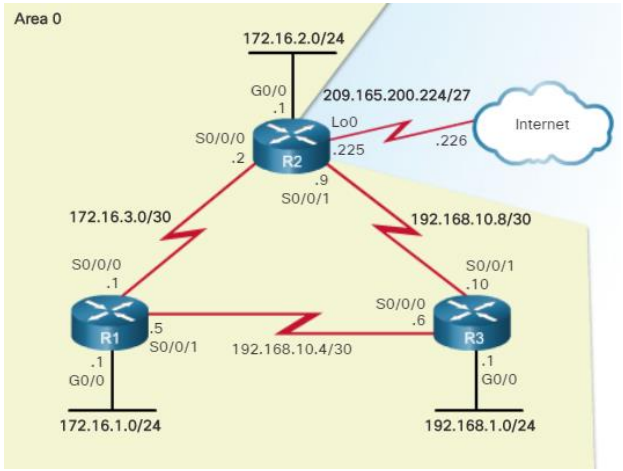


```
R2# show ip route | begin Gateway

Gateway of last resort is 209.165.200.226 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.226, Loopback0
   172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
O  172.16.1.0/24 [110/65] via 172.16.3.1, 00:01:44,
   Serial0/0/0
C  172.16.2.0/24 is directly connected, GigabitEthernet0/0
L  172.16.2.1/32 is directly connected, GigabitEthernet0/0
C  172.16.3.0/30 is directly connected, Serial0/0/0
L  172.16.3.2/32 is directly connected, Serial0/0/0
O  192.168.1.0/24 [110/65] via 192.168.10.10, 00:01:12,
   Serial0/0/1
   192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
O  192.168.10.4/30 [110/128] via 192.168.10.10, 00:01:12,
   Serial0/0/1
           [110/128] via 172.16.3.1, 00:01:12, Serial0/0/0
C  192.168.10.8/30 is directly connected, Serial0/0/1
L  192.168.10.9/32 is directly connected, Serial0/0/1
   209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C  209.165.200.224/30 is directly connected, Loopback0
L  209.165.200.225/32 is directly connected, Loopback0
R2#
```

Propagating a Default Static Route in OSPFv3

- To propagate a default route, the edge router (R2) must be configured with:
  - A default static route using the **ipv6 route ::/0 {***ipv6-address* **|** *exit-intf***}** command.
  - The **default-information originate** router configuration mode command.

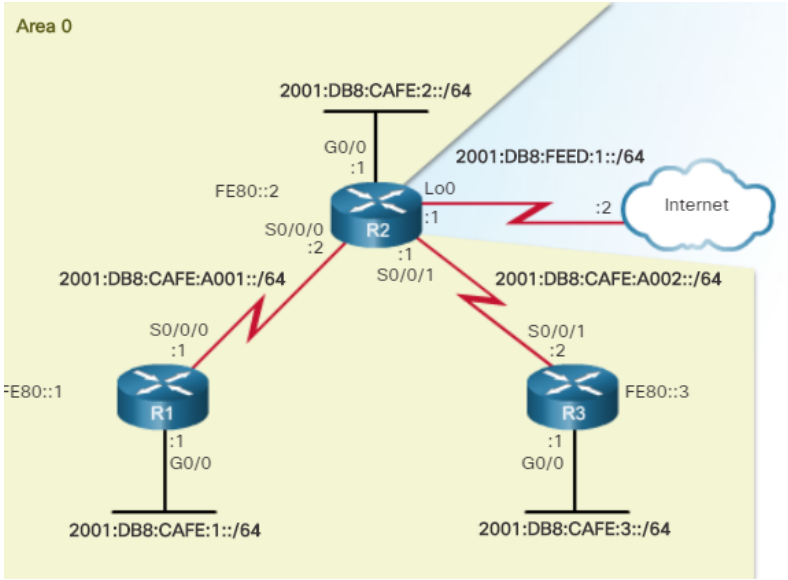


```
R2(config)# ipv6 route ::/0 2001:DB8:FEED:1::2
R2(config)#
R2(config)# ipv6 router ospf 10
R2(config-rtr)# default-information originate
R2(config-rtr)# end
R2#
*Apr 10 11:36:21.995: %SYS-5-CONFIG_I: Configured from console by console
R2#
```

Verifying the Propagated IPv6 Default Route

- Verify the default static route setting on R2 using the **show ipv6 route static** command.

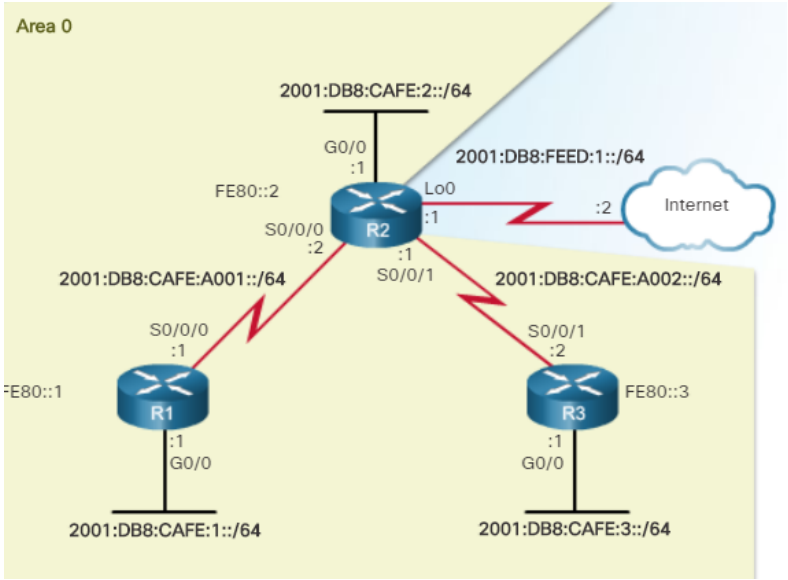


```
R2# show ipv6 route static
IPv6 Routing Table - default - 12 entries
Codes:C -Connected, L - Local, S - Static, U - Per-user Static route
      B -BGP, R - RIP, H - NHRP, I1 - ISIS L1
      I2 -ISIS L2, IA - ISIS interarea, IS-ISIS summary,D-EIGRP
      EX -EIGRP external, ND-ND Default,NDp-ND Prefix,DCE-Destination
      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
S   ::/0 [1/0]
     via 2001:DB8:FEED:1::2, Loopback0
R2#
```

OSPF Hello and Dead Intervals

- The OSPF Hello and Dead intervals used between two adjacent peers must match or a neighbor adjacency does not occur.
  - The OSPF Hello and Dead intervals are configurable on a per-interface basis.
  - The Serial 0/0/0 Hello and Dead intervals are set to the default 10 seconds and 40 seconds respectively.

- To verify the currently configured OSPFv2 interface intervals, use the **show ip ospf interface** command
- Use the **show ip ospf neighbor** command to verify that a router is adjacent with other routers.

```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 172.16.3.1/30, Area 0, Attached via
Network Statement
  Process ID 10, Router ID 1.1.1.1, Network Type
POINT_TO_POINT, Cost: 64
 Topology-MTID  Cost  Disabled    Shutdown    Topology Name
      0          64      no          no             Base
  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 00:00:03
  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 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
R1#
R1# show ip ospf neighbor

Neighbor ID  Pri  State   Dead Time  Address      Interface
3.3.3.3      0    FULL/-  00:00:35   192.168.10.6 Serial0/0/1
2.2.2.2      0    FULL/-  00:00:33   172.16.3.2   Serial0/0/0
R1#
```

Modifying OSPFv2 Intervals

- OSPFv2 Hello and Dead intervals can be modified using the interface configuration mode commands:
  - **ip ospf hello-interval** *seconds*
  - **ip ospf dead-interval** *seconds*

- Use the **no ip ospf hello-interval** and **no ip ospf dead-interval** interface configuration commands to reset the intervals to their default.

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip ospf hello-interval 5
R1(config-if)# ip ospf dead-interval 20
R1(config-if)# end
R1#
R1#
*Apr  7 17:28:21.529: %OSPF-5-ADJCHG: Process 10, Nbr
2.2.2.2 on Serial0/0/0 from FULL to DOWN, Neighbor Down: Dead timer expired
R1#
```

```
R2(config)# interface serial 0/0/0
R2(config-if)# ip ospf hello-interval 5
R2(config-if)#
*Apr  7 17:41:49.001: %OSPF-5-ADJCHG: Process 10, Nbr
1.1.1.1 on Serial0/0/0 from LOADING to FULL, Loading Done
R2(config-if)# end
R2#
R2# show ip ospf interface s0/0/0 | include Timer
  Timer intervals configured, Hello 5, Dead 20, Wait 20,
Retransmit 5
R2#
R2# show ip ospf neighbor

Neighbor ID   Pri  State  Dead Time  Address       Interface
3.3.3.3       0    FULL/- 00:00:35   192.168.10.10 Serial0/0/1
1.1.1.1       0    FULL/- 00:00:17   172.16.3.1    Serial0/0/0
R2#
```

Modifying OSPFv3 Intervals

- OSPFv3 Hello and Dead intervals can be modified using the interface configuration mode commands:
  - **ipv6 ospf hello-interval** *seconds*
  - **ipv6 ospf dead-interval** *seconds*

- Use the **no ipv6 ospf hello-interval** and **no ipv6 ospf dead-interval** interface configuration commands to reset the intervals to their default.

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 ospf hello-interval 5
R1(config-if)# ipv6 ospf dead-interval 20
R1(config-if)# end
R1#
*Apr 10 15:03:51.175: %OSPFv3-5-ADJCHG: Process 10, Nbr
2.2.2.2 on Serial0/0/0 from FULL to DOWN, Neighbor Down:
Dead timer expired
R1#
```

```
R2(config)# interface serial 0/0/0
R2(config-if)# ipv6 ospf hello-interval 5
R2(config-if)#
*Apr 10 15:07:28.815: %OSPFv3-5-ADJCHG: Process 10, Nbr
1.1.1.1 on Serial0/0/0 from LOADING to FULL, Loading Done
R2(config-if)# end
R2#
R2# show ipv6 ospf interface s0/0/0 | include Timer
  Timer intervals configured, Hello 5, Dead 20, Wait 20,
Retransmit 5
R2#
R2# show ipv6 ospf neighbor

          OSPFv3 Router with ID (2.2.2.2) (Process ID 10)

Neighbor ID Pri State  Dead Time Interface ID   Interface
3.3.3.3      0  FULL/- 00:00:38  7            Serial0/0/1
1.1.1.1      0  FULL/- 00:00:19  6            Serial0/0/0
R2#
```

10.2 Troubleshooting Single-Area OSPF Implementations

Components of Troubleshooting Single-Area OSPF

- OSPF is a popular routing protocol in large enterprise networks.

- Troubleshooting problems related to the exchange of routing information is one of the most essential skills for a network administrator.

OSPF Adjacencies Will Not Form If

- The interfaces are not on the same network.
- OSPF network types do not match.
- OSPF Hello or Dead Timers do not match.
- Interface to neighbor is incorrectly configured as passive.
- There is a missing or incorrect OSPF **`network`** command.
- Authentication is misconfigured.
- Each interface must be properly addressed and in the "up and up" condition.

OSPF States

- To troubleshoot OSPF, it is important to understand how OSPF routers traverse different OSPF states when adjacencies are being established.

- When troubleshooting OSPF neighbors, be aware that the FULL or 2WAY states are normal.
  - All other states are transitory and the router should not remain in those states for extended periods of time.

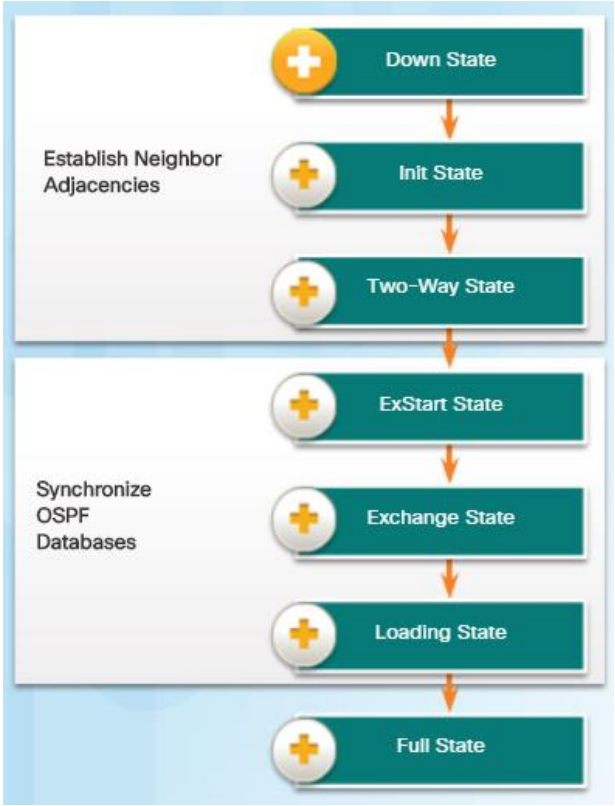

OSPF Troubleshooting Commands

- Common OSPFv2 troubleshooting commands include:
  - **show ip protocols** - Used to verify vital OSPFv2 configuration information.
  - **show ip ospf neighbor** - Used to verify that the router has formed an OSPFv2 adjacency with its neighboring routers.
  - **show ip ospf interface** - Used to display the OSPFv2 parameters configured on an interface.
  - **show ip ospf** - Used to examine the OSPFv2 process ID and router ID.
  - **show ip route ospf** - Used to display only the OSPFv2 learned routes in the IPv4 routing table. T
  - **clear ip ospf** [*process-id*] **process** - Used to reset the OSPFv2 neighbor adjacencies

```
R1# show ip protocols
*** IP Routing is NSF aware ***

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:
   172.16.1.1 0.0.0.0 area 0
   172.16.3.1 0.0.0.0 area 0
   192.168.10.5 0.0.0.0 area 0
 Passive Interface(s):
   GigabitEthernet0/0
 Routing Information Sources:
   Gateway         Distance      Last Update
   3.3.3.3              110      00:08:35
   2.2.2.2              110      00:08:35
 Distance: (default is 110)

R1# show ip ospf neighbor

Neighbor ID Pri State        Dead Time Address     Interface
2.2.2.2       1 FULL/BDR     00:00:30  192.168.1.2 GigabitEthernet0/0
3.3.3.3       0 FULL/DROTHER 00:00:38  192.168.1.3 GigabitEthernet0/0
R1#
```

Note: For the equivalent OSPFv3 command, simply substitute **ip** with **ipv6**.

Components of Troubleshooting OSPF

OSPF problems usually relate to:

- Neighbor adjacencies
- Missing routes
- Path selection

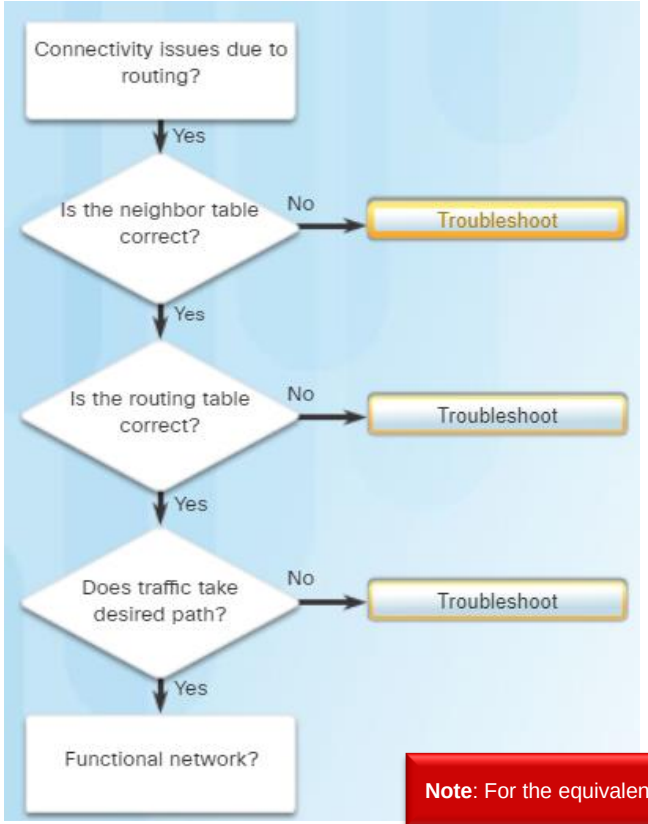

Use:
- **show ip ospf neighbor**
- **show ip interface brief**
- **show ip ospf interface**

Use:
- **show ip protocols**
- **show ip route ospf**

Use:
- **show ip route ospf**
- **show ip ospf interface**

Note: For the equivalent OSPFv3 command, simply substitute **ip** with **ipv6**.

Troubleshooting Neighbor Issues

- When troubleshooting neighbor issues:
  - Verify the routing table using the **show ip route ospf** command.
  - Verify that interfaces are active using the **show ip interface brief** command.
  - Verify active OSPF interfaces using the **show ip ospf interface** command.
  - Verify the OSPFv2 settings using the **show ip protocols** command.

```
R1# show ip protocols
*** IP Routing is NSF aware ***

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:
    172.16.1.1 0.0.0.0 area 0
    172.16.3.1 0.0.0.0 area 0
  Passive Interface(s):
    GigabitEthernet0/0
    Serial0/0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    3.3.3.3              110      00:50:03
    2.2.2.2              110      04:27:25
  Distance: (default is 110)

R1#
```

- Recall that the **passive-interface** command stops both outgoing and incoming routing updates and for that reason, routers will not become neighbors.

```
R1(config)# router ospf 10
R1(config-router)# no passive-interface s0/0/0
R1(config-router)#
*Apr  9 13:14:15.454: %OSPF-5-ADJCHG: Process 10, Nbr
2.2.2.2 on Serial0/0/0 from LOADING to FULL, Loading Done
R1(config-router)# end
R1#
```

Troubleshooting OSPFv2 Routing Table Issues

- When troubleshooting routing table issues:
  - Verify the routing table using the **show ip route ospf** command.
  - Verify the OSPFv2 settings using the **show ip protocols** command.
  - Verify the OSPF configuration using the **show running-config | section router ospf** command.

```
R3# show running-config | section router ospf
router ospf 10
 router-id 3.3.3.3
 passive-interface default
 no passive-interface Serial0/0/1
 network 192.168.10.8 0.0.0.3 area 0
R3#
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router ospf 10
R3(config-router)# network 192.168.1.0 0.0.0.255 area 0
R3(config-router)# end
R3#
*Apr 10 11:03:11.115: %SYS-5-CONFIG_I: Configured from
console by console
R3#
```

OSPFv3 Troubleshooting Commands

- Common OSPFv3 troubleshooting commands include:
  - **show ipv6 protocols** - Used to verify vital OSPFv3 configuration information.
  - **show ipv6 ospf neighbor** - Used to verify that the router has formed an OSPFv3 adjacency with its neighboring routers.
  - **show ipv6 ospf interface** - Used to display the OSPFv3 parameters configured on an interface.
  - **show ipv6 ospf** - Used to examine the OSPFv3 process ID and router ID.
  - **show ipv6 route ospf** - Used to display only the OSPFv3 learned routes in the IPv4 routing table. T
  - **clear ipv6 ospf** [*process-id*] **process** - Used to reset the OSPFv3 neighbor adjacencies
- .

```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#
R1# show ipv6 ospf neighbor

Neighbor ID   Pri  State   Dead Time  Interface ID  Interface
2.2.2.2         0  FULL/-  00:00:33       7         Serial0/0/0
R1#
```

Troubleshooting OSPFv3

- In this example, R1 is not receiving the R3 LAN OSPFv3 route (2001:DB8:CAFE:3::/64).
- Verifying the R3 routing protocol settings reveals that R3 is not enabled on the G0/0 R3 interface.
- Enable OSPFv3 on the R3 Gigabit Ethernet 0/0 interface.
- The R3 LAN is now in the routing table of R1.

```
R3# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 3.3.3.3
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
  Serial0/0/1
  Redistribution:
    None
R3#
```

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# interface g0/0
R3(config-if)# ipv6 ospf 10 area 0
R3(config-if)# end
R3#
```

Multiarea OSPF Troubleshooting Skills

- Before you can begin to diagnose and resolve problems related to a multiarea OSPF implementation, you must be able to do the following:
  - Understand the processes OSPF uses to distribute, store, and select routing information.
  - Understand how OSPF information flows within and between areas.
  - Use Cisco IOS commands to gather and interpret the information necessary to troubleshoot OSPF operation.

Multiarea OSPF Troubleshooting Data Structures

- OSPF stores routing information in four main data structures:

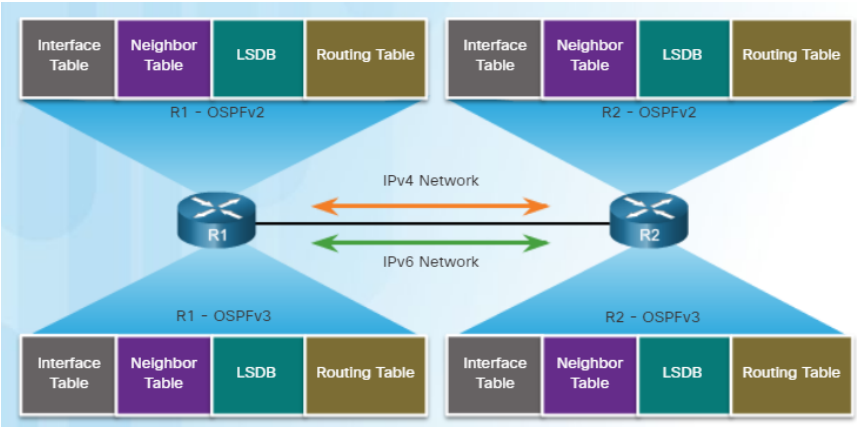

OSPF Data Structures	Description
Interface table	• Includes a list of all active OSPF interfaces. • Type 1 LSAs include the subnets associated with each active interface.
Neighbor table	• Used to manage neighbor adjacencies through hello timers and dead timers. • Neighbor entries are added and refreshed when a hello is received. • Neighbors are removed when dead timer expires.
Link-state database (LSDB)	• This is the primary data structure used by OSPF to store network topology information. • It includes full topological information about each area that the OSPF router is connected to and any paths that are available to reach other networks.
Routing table	• After the SPF algorithm is calculated, the best routes are offered to the routing table

10.3 Chapter Summary

Chapter 10: OSPF Tuning and Troubleshooting

- OSPF defines five network types: point-to-point, broadcast multiaccess, nonbroadcast multiaccess, point-to-multipoint, and virtual links.
- Multiaccess networks can create two challenges for OSPF regarding the flooding of LSAs: creation of multiple adjacencies and extensive flooding of LSAs. The solution to managing the number of adjacencies and the flooding of LSAs on a multiaccess network is the DR and BDR. If the DR stops producing Hellos, the BDR promotes itself and assumes the role of DR.
- The routers in the network elect the router with the highest interface priority as DR. The router with the second highest interface priority is elected the BDR. The higher the priority, the likelier the router will be selected as the DR. If set to 0, the router is not capable of becoming the DR. 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. If the interface priorities are equal, then the router with the highest router ID is elected the DR. The router with the second highest router ID is the BDR. The addition of a new router does not initiate a new election process.

Chapter 10: OSPF Tuning and Troubleshooting (Cont.)

- To propagate a default route in OSPF, the router must be configured with a default static route and the **default-information originate** command must be added to the configuration. Verify routes with the **show ip route** or **show ipv6 route** command.
- 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. To adjust the reference bandwidth, use the **auto-cost reference-bandwidth** *Mbps* router configuration mode command. To adjust the interface bandwidth, use the **bandwidth** *kilobits* interface configuration mode command. The cost can be manually configured on an interface using the **ip ospf cost** *value* interface configuration mode command.
- The OSPF Hello and Dead intervals must match or a neighbor adjacency does not occur. To modify these intervals, use the following interface commands:
  - **ip ospf hello-interval** *seconds*
  - **ip ospf dead-interval** *seconds*
  - **ipv6 ospf hello-interval** *seconds*
  - **ipv6 ospf dead-interval** *seconds*

Chapter 10: OSPF Tuning and Troubleshooting (Cont.)

- When troubleshooting OSPF neighbors, be aware that the FULL or 2WAY states are normal. The following commands summarize OSPFv2 troubleshooting:
  - **show ip protocols**
  - **show ip ospf neighbor**
  - **show ip ospf interface**
  - **show ip ospf**
  - **show ip route ospf**
  - **clear ip ospf** [*process-id*] **process**
- Troubleshooting OSPFv3 is similar to OSPFv2. The following commands are the equivalent commands used with OSPFv3: **show ipv6 protocols**, **show ipv6 ospf neighbor**, **show ipv6 ospf interface**, **show ipv6 ospf**, **show ipv6 route ospf**, and **clear ipv6 ospf** [*process-id*] **process**.


CISCO