


OSPF Lecture Overview

- **OSPF Network Types:** Refers to the different kinds of connections between OSPF routers and how those connections influence OSPF's behavior.
- **OSPF Neighbor and Adjacency Requirements:** Examines the specific requirements for a successful neighbor adjacency (covered after the general process in Day 2).
- **LSA Types (Link State Advertisement):** An overview of a few specific LSA types required for the CCNA.
- **Quiz:** A final quiz followed by a bonus question from Boson ExSim.


Loopback Interfaces



- A loopback interface is a virtual interface in the router.
- It is always up/up (unless you manually shut it down)
- It is not dependent on a physical interface.
- So, it provides a consistent IP address that can be used to reach/identify the router.


- A loopback interface is a virtual interface in the router.
- It is always up/up (unless you manually shut it down).
- The status of a loopback interface is not dependent on a physical interface.

- Physical interfaces can have hardware problems and fail, but that can't happen to a loopback interface unless the router itself fails.
- It provides a consistent IP address that can be used to reach and identify the router.




- A loopback interface is a virtual interface in the router.
- It is always up/up (unless you manually shut it down)
- It is not dependent on a physical interface.
- So, it provides a consistent IP address that can be used to reach/identify the router.

- Example (Above):** R1 has no loopback interface. R4 receives a packet destined for R1 at 10.0.13.1, the IP address of its G1/0 interface. It forwards it to R1.



- A loopback interface is a virtual interface in the router.
- It is always up/up (unless you manually shut it down)
- It is not dependent on a physical interface.
- So, it provides a consistent IP address that can be used to reach/identify the router.

- If R1's G1/0 interface goes down, R4 will not be able to send the packet to R1.



- A loopback interface is a virtual interface in the router.
- It is always up/up (unless you manually shut it down)
- It is not dependent on a physical interface.
- So, it provides a consistent IP address that can be used to reach/identify the router.


- **Example (Above):** R1 has a loopback interface, 1.1.1.1, which is used to identify R1 instead of 10.0.13.1. Even if a physical interface fails, when R4 receives a packet destined for R1's loopback interface, it will still be able to send the packet to R1.
- Configuring a loopback interface on a router is a good idea because it provides an interface with an IP address that is always up, and can consistently be used to identify and reach the router.

OSPF Network Types

- The OSPF network type refers to the type of connection between OSPF neighbors, and the type influences how OSPF behaves in some ways.
- The most common type of connection in modern networks is Ethernet.


- There are three main OSPF network types:
 - **Broadcast:** Enabled by default on Ethernet and FDDI interfaces.
 - FDDI is an old technology, but for the exam you might want to remember that the OSPF broadcast network type is used by default on it.
 - **Point-to-point:** Enabled by default on PPP and HDLC interfaces.
 - **Non-broadcast:** Enabled by default on Frame Relay and X.25 interfaces.

Characteristics of the Broadcast Network Type




- **Neighbor Discovery:** Routers dynamically discover neighbors by sending and listening for OSPF Hello messages using the multicast address 224.0.0.5.
 - Note: Not all network types dynamically discover neighbors. For example, with the 'Non-broadcast' network type, you must manually configure neighbors.
- **DR/BDR Election:** A DR (designated router) and BDR (backup designated router) must be elected on each subnet.
 - If there are no OSPF neighbors on an interface (like the G1/0 interface of R1, R3, R4, and R5), there is just a DR, no BDR.
 - Routers that are not the DR or BDR for the subnet become a 'DROther'.


- Example (Subnets needing a DR):



- Subnets with no OSPF neighbors: Each router becomes the DR for its own subnet. **(Above)**




- 192.168.1.0/30 subnet (R1 and R2): Let's say R2 is the DR. Then R1 becomes the BDR for the segment. **(Above)**



- 192.168.2.0/29 subnet (R2, R3, R4, R5): For example, R5 might be the DR, R4 the BDR, and then R2 and R3 become DROthers. **(Above)**

DR and BDR Election Process

- There is an order of priority for the election:
 - The router with the highest OSPF interface priority in the subnet becomes the DR.
 - However, all interfaces have the same priority by default, so if the priorities are equal, the routers compare their OSPF router IDs.
 - The router with the highest OSPF router ID wins.
- 'First place' in the election becomes the DR, and 'second place' becomes the BDR.
- The default OSPF interface priority is 1 on all interfaces, so by default the router with the highest router ID will become the DR for the segment.



- Example (Above) (R5 is DR):** Output from `show ip ospf interface g0/0`
 - on R5 shows its router ID (from its loopback interface)
 - state is DR, and priority is the default of 1.
 - R5 has the highest router ID on the 192.168.2.0/29 subnet, so it became the DR.
 - The output also lists R5 as the DR and R4 as the BDR for the segment.



```

R2#show ip ospf int g0/0
GigabitEthernet0/0 is up, line protocol is up
 Internet Address 192.168.2.1/29, Area 0, Attached via Network Statement
 Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
      0              1          no           no           Base
Transmit Delay is 1 sec, State DROTHER, Priority 1
Designated Router (ID) 5.5.5.5, Interface address 192.168.2.4
Backup Designated router (ID) 4.4.4.4, Interface address 192.168.2.3

```

- **Example (R2 is DROther) (Above):** The same command output on R2 shows its state as DROTHER.

Changing OSPF Interface Priority



```


R2(config)#int g0/0
R2(config-if)#ip ospf priority ?
<0-255> Priority
R2(config-if)#ip ospf priority 255

```

If you set the OSPF interface priority to 0, the router CANNOT be the DR/BDR for the subnet.

- The command to change the OSPF priority of an interface is **IP OSPF PRIORITY**, with a range of 0 to 255. **(Above)**
 - Note: If you set the OSPF interface priority to 0, the router CANNOT be the DR/BDR for the subnet, no matter what.

```
R2(config-if)# ip ospf priority 255
```



```
R2(config)#int g0/0
R2(config-if)#ip ospf priority ?
<0-255> Priority
R2(config-if)#ip ospf priority 255
```

```
R2#show ip ospf int g0/0
GigabitEthernet0/0 is up, line protocol is up
Internet Address 192.168.2.1/29, Area 0, Attached via Network Statement
Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
0 1 no no Base
Transmit Delay is 1 sec, State DROTHER, Priority 255
Designated Router (ID) 5.5.5.5, Interface address 192.168.2.4
Backup Designated router (ID) 4.4.4.4, Interface address 192.168.2.3
```


- **Non-Preemptive Election (Above):** The DR/BDR election is 'non-preemptive'. This means that once the DR/BDR are selected, they will keep their role until OSPF is reset, the interface fails, is shut down, etc. Increasing the priority on another router will not cause it to take over the DR role automatically.
 - Example: R2's priority was changed to 255, but its state remained DROTHER because the election is non-preemptive.

DR/BDR Failure and Adjacency States


DR Failure

- When the DR goes down, the BDR automatically becomes the new DR. An election is then held for the next BDR.


- **Example:** The OSPF process on R5 (the DR) was reset.



- R5's neighbors went DOWN. (**Above**)



- R2 and R4 returned to the FULL state, but not R3. **(Above)**




- Using **SHOW IP OSPF NEIGHBOR** on R5 shows that R4 became the DR, not R2. **(Above)**
- R4, which was the BDR, instantly stepped up to be the new DR.
- An election was then held between the other routers for the next BDR. R2 has the highest priority (255), so it became the new BDR.

DROther Adjacencies

- DROthers will only move to the FULL state with the DR and BDR of the subnet. The neighbor state with other DROthers will be 2-way. This gives us a hint to the purpose of the DR and BDR.
- Example:**
 - After R5's OSPF process was reset, R5 became a DROther.
 - R3 is also a DROther.
 - The neighbor state between R5 and R3 was stable in the 2-way state.

DR/BDR Purpose and LSA Exchange




- LSA Exchange:**


- In the broadcast network type, routers will only form a full OSPF adjacency with the DR and BDR of the segment.
- Therefore, routers only exchange LSAs with the DR and BDR.
- DROthers will not exchange LSAs with each other.
- In the 2-way state, routers haven't shared LSAs with each other yet.

- **Purpose:**

- All routers will still have the same LSDB, but this reduces the amount of LSAs flooding the network.




- **Example (Above):** If 6 routers are connected to the same segment and they all share LSAs with each other, there will be a lot of LSAs flooding and congesting the network.



- **Example (Above):** If routers only exchange LSAs with the DR and BDR, the number of LSAs flooding around the network is reduced.
- **Multicast Addresses:**
 - When routers need to send messages to the DR and BDR, they use multicast address 224.0.0.6.
 - This is different than the OSPF 'all routers' multicast address of 224.0.0.5.

OSPF Neighbor Process Review




- When routers become OSPF neighbors, they have completed the initial steps of the neighbor process.
- Connections between two DROthers stop at this point (2-way state).
- Routers will only continue on to exchange LSAs and form a full OSPF adjacency with the DR and BDR.
- The DR and BDR form a full adjacency with ALL routers in the subnet, including the DROthers.
- DROthers will form a FULL adjacency only with the DR/BDR.

Verifying Adjacencies

- Verification:

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



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

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs F/C
Gi0/0	1	0	192.168.2.2/29	1	DROTH	2/3
Gi1/0	1	0	10.0.3.1/24	1	DR	0/0

- On R3 (a DROther), the neighbor count on its G0/0 interface shows:
 - F** indicates the number of full adjacencies.
 - C** indicates the total count of neighbors.
 - R3 has two Full adjacencies (with R2 and R4) but three total neighbors (R2, R4, and R5).

- Verification:

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


```

R3#show ip ospf interface g0/0
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 192.168.2.2/29, Area 0, Attached via Network Statement
  Process ID 1, 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 DROTHER, Priority 1
  Designated Router (ID) 4.4.4.4, Interface address 192.168.2.3
  Backup Designated router (ID) 2.2.2.2, Interface address 192.168.2.1
  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 0, maximum is 1
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 3, Adjacent neighbor count is 2
    Adjacent with neighbor 2.2.2.2  (Backup Designated Router)
    Adjacent with neighbor 4.4.4.4  (Designated Router)
  Suppress hello for 0 neighbor(s)

```

- This command gives the same information:
 - 'Neighbor Count is 3' is R3's total number of OSPF neighbors.
 - 'Adjacent neighbor count is 2' are the neighbors R3 has formed a Full adjacency with.
 - The two adjacent neighbors listed are R2 (the BDR) and R4 (the DR).

Point-to-Point Network Type




- This network type is enabled on serial interfaces using the PPP or HDLC encapsulations by default.
- PPP and HDLC are both Layer 2 encapsulations, similar to Ethernet, except they are used on serial connections.
- Same as the Broadcast network type, routers dynamically discover neighbors by sending/listening for OSPF Hello messages using multicast address 224.0.0.5.
- **Difference:** A DR and BDR are NOT elected.
 - Why? As the network type name implies, these encapsulations are used for 'point-to-point' connections between two routers. Therefore there is no point in electing a DR and BDR.
 - The two routers will form a Full adjacency with each other, without the need to elect a DR and BDR.

Serial Connections Overview



- Serial connections are an old technology which isn't very common anymore. Ethernet is much more dominant.
- Serial connections have been removed from exam topics except for the OSPF 'point-to-point' network type.
- Although you won't be tested directly on knowledge of serial interfaces, you should have a basic understanding of them.

Serial Interface Basics




- One side of a serial connection functions as **DCE** (Data Communications Equipment).
- The other side functions as **DTE** (Data Terminal Equipment).
- The DCE side needs to specify the **clock rate**, which is the speed of the connection.

Configuring R1's Serial Interface (DCE Side)

- The command to specify the speed is **clock rate**, followed by standard values in bits per second.
- R1 has the DCE side of the cable and needs to tell R2 what speed the connection will operate at.
- A clock rate of 64000 bits per second (64 kbps) was configured, along with an IP address and **no shutdown**.
- **Important Point:** Ethernet interfaces use the **SPEED** command to configure the operating speed. Serial interfaces use the **CLOCK RATE** command.

Example: R1 Configuration



```
R1(config)#interface s2/0
R1(config-if)#clock rate ?
    With the exception of the following standard values not subject to rounding,

    1200 2400 4800 9600 14400 19200 28800 38400
    56000 64000 128000 2015232

    accepted clockrates will be bestfitted (rounded) to the nearest value
    supportable by the hardware.

    <246-8064000>    DCE clock rate (bits per second)

R1(config-if)#clock rate 64000
R1(config-if)#ip address 192.168.1.1 255.255.255.0
R1(config-if)#no shut
```

- One side of a serial connection functions as DCE (Data Communications Equipment)
- The other side functions as DTE (Data Terminal Equipment)
- The DCE side needs to specify the *clock rate* (speed) of the connection

```

R1(config)# interface s2/0
R1(config-if)# clock rate 64000
R1(config-if)# ip address <ip_address> <subnet_mask>
R1(config-if)# no shutdown

```

HDLC Encapsulation



```

R1#show interface s2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC   crc 16, loopback not set

```

- The default encapsulation on a serial interface is HDLC.
**actually cHDLC (Cisco HDLC)

cHDLC frame structure [\[edit \]](#)


The following table describes the structure of a cHDLC frame on the wire. [citation needed]

Address	Control	Protocol Code	Information	Frame Check Sequence (FCS)	Flag
8 bits	8 bits	16 bits	Variable length, 0 or more bits, in multiples of 8	16 bits	8 bits

- The Address field is used to specify the type of packet contained in the cHDLC frame; 0x0F for Unicast and 0x8F for Broadcast packets.
- The Control field is always set to zero (0x00).
- The Protocol Code field is used to specify the protocol type encapsulated within the cHDLC frame (e.g. 0x0800 for [Internet Protocol](#)).

- After checking the interface with **SHOW INTERFACE S2/0**, the encapsulation is HDLC.
- On Cisco routers, the default encapsulation on a serial interface is HDLC (actually Cisco's version called 'cHDLC', but it displays as 'HDLC' in the CLI).
- This is a Layer 2 encapsulation, like Ethernet, except it's used on serial connections.
- Example:** The structure of an HDLC frame (from Wikipedia) shows there is no MAC address field. So, MAC addresses actually aren't used.

Changing Encapsulation to PPP



```
R1#show interface s2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC, crc 16, loopback not set
```

```
R1(config)#int s2/0
R1(config-if)#encapsulation ppp
R1(config-if)#do show interface s2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation PPP, LCP Open
```

- The default encapsulation on a serial interface is HDLC.
**actually cHDLC (Cisco HDLC)

- To configure the router to use PPP encapsulation instead, use the **ENCAPSULATION PPP** command on the interface.
- Note: If you change the encapsulation, it must match on both ends or the interface will go down. If they use two different encapsulations it's like they are speaking two different languages, they won't be able to communicate.

Example: R1 Verification

```
R1# show interface s2/0
```




```
R1#show interface s2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC, crc 16, loopback not set
```


```
R1(config)#int s2/0
R1(config-if)#encapsulation ppp
R1(config-if)#do show interface s2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation PPP, LCP Open
```

- The default encapsulation on a serial interface is HDLC.
**actually cHDLC (Cisco HDLC)
- If you change the encapsulation, it must match on both ends or the interface will go down.

- The encapsulation has changed to PPP.

Final Configurations

- Example: R1 Final Configuration




```
R1#show running-config interface s2/0
Building configuration...

Current configuration : 126 bytes
!
interface Serial2/0
ip address 192.168.1.1 255.255.255.0
encapsulation ppp
serial restart-delay 0
clock rate 64000
end
```

- The **SERIAL RESTART-DELAY 0** command was there by default. The clock rate, encapsulation, and IP address were configured.

```
R1(config)# interface s2/0
R1(config-if)# clock rate 64000
R1(config-if)# encapsulation ppp
R1(config-if)# ip address <ip_address> <subnet_mask>
R1(config-if)# no shutdown
```

- **Example: R2 Configuration (DTE Side)**



```
R1#show running-config interface s2/0
Building configuration...

Current configuration : 126 bytes
!
interface Serial2/0
 ip address 192.168.1.1 255.255.255.0
 encapsulation ppp
 serial restart-delay 0
 clock rate 64000
end
```

```
R2#show running-config interface s2/0
Building configuration...


Current configuration : 110 bytes
!
interface Serial2/0
 ip address 192.168.1.2 255.255.255.252
 encapsulation ppp
 serial restart-delay 0
end
```

- R2 is the DTE end, so no **CLOCK RATE** command is needed.

```
R2(config)# interface s2/0
R2(config-if)# encapsulation ppp
R2(config-if)# ip address <ip_address> <subnet_mask>
R2(config-if)# no shutdown
```


Verifying DCE/DTE IN Packet Tracer

- The command to view which side is DCE and which is DTE is **SHOW CONTROLLERS**, followed by the interface ID.
- Verification:**



```
sw# show controllers s2/0
```

- Output shows R1 is the DCE side with the configured clock rate of 64,000 bits.

```
R2# show controllers s2/0
```



```

R1#
R1#show con
R1#show controllers s2/0
Interface Serial2/0
Hardware is PowerQUICC MPC860
DCE V.35, clock rate 64000
idb at 0x81081AC4, driver data structure at 0x81084AC0
SCC Registers:
General [GSMR]=0x2:0x00000000, Protocol-specific [PSMR]=0x8
Events [SCCE]=0x0000, Mask [SCCM]=0x0000, Status [SCCS]=0x00
Transmit on Demand [TODR]=0x0, Data Sync [DSR]=0x7E7E
Interrupt Registers:
Config [CICR]=0x00367F80, Pending [CIPR]=0x0000C000
Mask [CIMR]=0x00200000, In-srv [CISR]=0x00000000
Command register [CR]=0x580
Ctrl+F6 to exit CLI focus

```

```

R2#
R2#show controllers s2/0
Interface Serial2/0
Hardware is PowerQUICC MPC860
DTE V.35 TX and RX clocks detected
idb at 0x81081AC4, driver data structure at 0x81084AC0
SCC Registers:
General [GSMR]=0x2:0x00000000, Protocol-specific [PSMR]=0x8
Events [SCCE]=0x0000, Mask [SCCM]=0x0000, Status [SCCS]=0x00
Transmit on Demand [TODR]=0x0, Data Sync [DSR]=0x7E7E
Interrupt Registers:
Config [CICR]=0x00367F80, Pending [CIPR]=0x0000C000
Mask [CIMR]=0x00200000, In-srv [CISR]=0x00000000
Command register [CR]=0x580
Port 2 [PADIR]=0x1030 [PADAR]=0xFFFF
Ctrl+F6 to exit CLI focus

```

- Output shows R2 is the DTE side and has detected the Tx (transmit) and Rx (receive) clocks from R1.


Serial Connections Summary

- The default encapsulation on a serial interface is HDLC.
- You can configure them to use PPP encapsulation instead with this command: **ENCAPSULATION PPP**.
- If you change the encapsulation on one side, remember to change it on the other too!
- One side of the connection is DCE and the other is DTE.
- You can identify which side is DCE and which is DTE with this command: **SHOW CONTROLLERS**, followed by the interface ID.
- Configure the clock rate, the speed of the connection, on the DCE side with this command: **CLOCK RATE**, followed by the clock rate in bits per second.

OSPF Point-to-Point Verification

- **Verification:**

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



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

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	0	FULL/ -	00:00:31	192.168.1.1	Serial2/0
3.3.3.3	1	2WAY/DROTHER	00:00:39	192.168.2.2	GigabitEthernet0/0
4.4.4.4	1	FULL/DR	00:00:38	192.168.2.3	GigabitEthernet0/0
5.5.5.5	1	FULL/BDR	00:00:31	192.168.2.4	GigabitEthernet0/0


- R2 has a full adjacency with R1, but instead of DR, BDR, or DROTHER, a dash is displayed.
- This is because the point-to-point network type doesn't use DRs or BDRs.

Configuring OSPF Network Type

```
R1(config-if)#ip ospf network ?  
broadcast          Specify OSPF broadcast multi-access network  
non-broadcast       Specify OSPF NBMA network  
point-to-multipoint Specify OSPF point-to-multipoint network  
point-to-point      Specify OSPF point-to-point network
```

- You can manually configure the network type an interface uses with the command **IP OSPF NETWORK**, followed by the network type.
- There is also a 'point-to-multipoint' network type. It's more of a 'sub-type'. You don't need to learn about it for the CCNA.

Reasons to Change OSPF Network Type



- **Example (Above):** If two routers are directly connected with an Ethernet link (like R1 and R2 in the diagram), there is no need for a DR/BDR. You can configure the point-to-point network type in this case, although you don't have to.
- **Compatibility Note:** Not all network types work on all link types.
 - **Example:** A serial link cannot use the broadcast network type. This is because serial links don't support Layer 2 broadcast frames, which is necessary for the broadcast network type.

OSPF Network Types Quick Review

Broadcast	Point-to-point
Default on Ethernet , FDDI interfaces	Default on HDLC , PPP (serial) interfaces
DR/DBR elected	No DR/BDR
Neighbors dynamically discovered	Neighbors dynamically discovered
Default timers: Hello 10, Dead 40	Default timers: Hello 10, Dead 40

- Point-to-point networks use the same default timers as broadcast networks.
- The default Hello timer is 10 seconds and the default Dead timer is 40 seconds.


- For reference only (not required for the exam): The non-broadcast network type uses a default Hello timer of 30 seconds and Dead timer of 120 seconds.

OSPF Neighbor Requirements

- Usually, routers will become OSPF neighbors without issue, but there are a few problems that can occur.

Rule 1. Area Number Must Match

- The area number must match for two routers to become OSPF neighbors.




```
R1#show running-config | section ospf
router ospf 1
network 192.168.1.0 0.0.0.3 area 0
```

```
R2#show running-config | section ospf
router ospf 1
network 192.168.1.0 0.0.0.3 area 1
```

```
R1#show ip ospf neighbor
R1#
R1#
```

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

- **Example:** A small topology of two routers (R1 and R2).
 - R1's OSPF config: OSPF is enabled on G0/0 in area 0.
 - R2's OSPF config: OSPF is enabled on G0/0 in area 1.
 - When `SHOW IP OSPF NEIGHBOR` is entered on both devices, they have no OSPF neighbors.




- The network command on R2 was changed to use area 0.
- As a result, they successfully became OSPF neighbors.

For two routers to become OSPF neighbors, they must be in the same area.

Rule 2. Interfaces Must Be in the Same Subnet

- The interfaces must be in the same subnet to become OSPF neighbors.

- **Example:** R1 and R2's G0/0 interfaces are in different subnets.



```

R1#show running-config | section ospf
router ospf 1
 network 192.168.1.0 0.0.0.3 area 0
R1#
  
```

```

R2#show running-config | section ospf
router ospf 1
 network 192.168.2.0 0.0.0.3 area 0
R2#
  
```


```

R1#show ip ospf neighbor
R1#
R1#
  
```

```

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

- OSPF was activated on both of the interfaces.
- When checking for OSPF neighbors, they don't come up.



```

R1#show running-config | section ospf
router ospf 1
 network 192.168.1.0 0.0.0.3 area 0
R1#
  
```

```

R2#show running-config | section ospf
router ospf 1
 network 192.168.1.0 0.0.0.3 area 0
R2#
  
```

```

R1#show ip ospf neighbor
  
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.2	1	FULL/BDR	00:00:34	192.168.1.2	GigabitEthernet0/0

```

R2#show ip ospf neighbor
  
```


Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.1	1	FULL/DR	00:00:39	192.168.1.1	GigabitEthernet0/0

- R2's interface was configured in the same subnet as R1, and the network command was edited so that OSPF is once again activated on G0/0.

- Now R1 and R2 are OSPF neighbors again.

Rule 3. OSPF Process Must Not Be Shutdown

- The OSPF process must not be shutdown.
- You can 'shutdown' the OSPF process on the router like an interface. This disables OSPF operation, without removing the OSPF configurations.



```
R2(config)#router ospf 1
R2(config-router)#shutdown
R2(config-router)#
*Aug 23 03:43:31.719: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Int
R2(config-router)#do show ip ospf neighbor
R2(config-router)#
```

```
R2(config-router)#no shutdown
R2(config-router)#
*Aug 23 03:49:52.931: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-router)#do show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.1	1	FULL/DR	00:00:38	192.168.1.1	GigabitEthernet0/0


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

- **Example:** From OSPF configuration mode on R2, the **SHUTDOWN** command was used.
 - A message is displayed indicating that its neighbor R1 has gone from FULL to DOWN.
 - No neighbors are displayed in **SHOW IP OSPF NEIGHBOR**.
 - OSPF was then enabled again with **NO SHUTDOWN**.
 - A message indicates the neighbor is back up, and it can be seen in **SHOW IP OSPF NEIGHBOR**.
- This won't be a problem unless you manually shutdown the OSPF process, so usually it's not a problem.

Rule 4. OSPF Router IDs Must Be Unique

- The OSPF router IDs must be unique.

- **Example:**




```

R2(config-router)#router-id 192.168.1.1
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect
R2(config-router)#end
R2#clear ip
*Aug 23 03:57:58.835: %SYS-5-CONFIG_I: Configured from console by console
R2#clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2#
*Aug 23 03:58:04.055: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or d
R2#
*Aug 23 03:58:06.495: %OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 192.168.1.1 from 192.168.1.1 on interface GigabitEthernet0/0
R2#show ip ospf neighbor
R2#

```

- R1 and R2 have no manually configured router IDs or loopback interfaces.
- Each router selected its G0/0 IP address as its router ID: 192.168.1.1 for R1 and 192.168.1.2 for R2.
- 192.168.1.1 was configured as R2's router ID, making it the same as R1's.
- **CLEAR IP OSPF PROCESS** was used for the new router ID to take effect.
- The neighbor went down immediately because the OSPF process was cleared.
- Instead of coming back up, a message was displayed: 'OSPF detected duplicate router-id 192.168.1.1 from 192.168.1.1 on interface GigabitEthernet0/0'.
- The neighbor stayed down.

- **Fix:** The manually configured router ID was removed with **NO ROUTER-ID**.




```

R2(config-router)#router-id 192.168.1.1
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect
R2(config-router)#end
R2#clear ip
*Aug 23 03:57:58.835: %SYS-5-CONFIG_I: Configured from console by console
R2#clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2#
*Aug 23 03:58:04.055: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or d
R2#
*Aug 23 03:58:06.495: %OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 192.168.1.1 from 192.168.1.1 on interface GigabitEthernet0/0
R2#show ip ospf neighbor
R2#
R2(config-router)#no router-id
R2(config-router)#
*Aug 23 04:10:10.207: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-router)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    1     FULL/DR         00:00:35   192.168.1.1   GigabitEthernet0/0
R2(config-router)#
  
```

- Note: You don't have to specify the router ID when removing the command. **NO ROUTER-ID** has the same effect as **NO ROUTER-ID 192.168.1.1**.
- Without having to reset the OSPF process, R2's router ID returned to 192.168.1.2 and the neighbor came up again.
- The router was free to change R2's router ID because R2 had no other OSPF neighbors at the time.
- Watch out for duplicate router IDs.

Rule 5. Hello and Dead Timers Must Match



```


R2(config-if)#ip ospf hello-interval ?
<1-65535> Seconds

R2(config-if)#ip ospf hello-interval 5
R2(config-if)#ip ospf dead-interval ?
<1-65535> Seconds
minimal Set to 1 second

R2(config-if)#ip ospf dead-interval 20
R2(config-if)#
*Aug 23 04:29:30.623: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Dead timer expired
R2(config-if)#do show ip ospf neighbor
R2(config-if)#
  
```

- The Hello and Dead timers must match.
- In both Broadcast and Point-to-Point network types, the defaults are 10 seconds (Hello) and 40 seconds (Dead).
- The Hello and Dead timers are configured directly on the interface.
- The command to configure the Hello timer is **IP OSPF HELLO-INTERVAL**, followed by the interval in seconds.
- The command to configure the Dead timer is **IP OSPF DEAD-INTERVAL**, followed by the interval in seconds.
- Note: If you change the Hello timer and/or the Dead timer and they don't match on both sides, the neighbor will go down.

```
R2(config-if)# ip ospf hello-interval 5
R2(config-if)# ip ospf dead-interval 20
```



```

R2(config-if)#ip ospf hello-interval ?
<1-65535> Seconds

R2(config-if)#ip ospf hello-interval 5
R2(config-if)#ip ospf dead-interval ?
<1-65535> Seconds
minimal Set to 1 second

R2(config-if)#ip ospf dead-interval 20
R2(config-if)#
*Aug 23 04:29:30.623: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Dead timer expired
R2(config-if)#do show ip ospf neighbor
R2(config-if)#

R2(config-if)#no ip ospf hello-interval
R2(config-if)#no ip ospf dead-interval
R2(config-if)#
*Aug 23 04:31:32.727: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-if)#do show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
192.168.1.1      1     FULL/BDR        00:00:35   192.168.1.1   GigabitEthernet0/0
R2(config-if)#
  
```


Fix: Use **NO IP OSPF HELLO-INTERVAL** and **NO IP OSPF DEAD-INTERVAL** to return the timers to their default values.

- Note: You don't have to specify the values when removing the commands (e.g., `NO IP OSPF HELLO-INTERVAL 5` is not required).
- Once the Hello and Dead timers match again, the neighbor is back up.
- Remember to check the Hello and Dead intervals if you're having trouble with OSPF adjacencies.

Rule 6. Authentication Settings Must Match


- You can configure an OSPF password, and then the router will only become neighbors with routers that have a matching OSPF password.
- The OSPF password is configured directly on the interface with `IP OSPF AUTHENTICATION-KEY`, followed by the password.
- Note: This command configures the password but does not enable OSPF authentication on the interface.
- To enable OSPF authentication on the interface, use `IP OSPF AUTHENTICATION`.

Example:



```
R2(config-if)#ip ospf authentication-key jeremy
R2(config-if)#ip ospf authentication
R2(config-if)#
*Aug 23 04:56:28.435: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: D
R2(config-if)#do show ip ospf neighbor
R2(config-if)#
```

- On R2, the OSPF password was configured as 'jeremy' and authentication was enabled.
- The neighbor (R1) went down because OSPF authentication was not configured on R1.



```

R2(config-if)#ip ospf authentication-key jeremy
R2(config-if)#ip ospf authentication
R2(config-if)#
*Aug 23 04:56:28.435: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: D
R2(config-if)#do show ip ospf neighbor
R2(config-if)#
  
```

```

R2(config-if)#no ip ospf authentication
R2(config-if)#no ip ospf authentication-key jeremy
*Aug 23 04:59:37.315: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-if)#do show ip ospf neighbor


Neighbor ID      Pri   State           Dead Time   Address        Interface
192.168.1.1      1     FULL/BDR        00:00:34    192.168.1.1    GigabitEthernet0/0
R2(config-if)#
  
```

- **Fix:** The authentication settings were removed from R2, and the neighbor came back up.

Rule 7. IP MTU Settings Must Match

- The IP MTU settings on the interfaces must match.
- The IP MTU is the maximum size of an IP packet that will be sent out of the interface. The default is usually 1500 bytes, but it can be configured.
- Note: Even if the settings don't match, the neighbors can become OSPF neighbors, but OSPF won't function properly.
- The command to configure the IP MTU of an interface is **IP MTU**, followed by the MTU in bytes.

Example:




```

R2(config-if)#ip mtu ?
<68-1500> MTU (bytes)
R2(config-if)#ip mtu 1400
R2(config-if)#do show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
192.168.1.1      1     FULL/BDR        00:00:34    192.168.1.1    GigabitEthernet0/0
R2(config-if)#do clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2(config-if)#
*Aug 23 05:16:07.474: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
R2(config-if)#do show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
192.168.1.1      1     EXSTART/DR      00:00:38    192.168.1.1    GigabitEthernet0/0
  
```

- The IP MTU was changed to 1400 on R2's G0/0, so it doesn't match with the default 1500 on R1's G0/0.
- R1 and R2 actually remained neighbors.
- The OSPF process on R2 was reset.



```
R2(config-if)#ip mtu ?
<68-1500> MTU (bytes)


R2(config-if)#ip mtu 1400
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    1     FULL/BDR        00:00:34    192.168.1.1    GigabitEthernet0/0
R2(config-if)#do clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2(config-if)#
*Aug 23 05:16:07.474: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    1     EXSTART/DR      00:00:38    192.168.1.1    GigabitEthernet0/0

*Aug 23 05:21:12.946: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from EXSTART to DOWN, Neighbor Down: Too many retransmissions
R2(config-if)#
*Aug 23 05:22:12.946: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from DOWN to DOWN, Neighbor Down: Ignore timer expired
```

- The neighbors were stuck in the EXSTART state and did not reach the FULL state. Messages were displayed and kept repeating, indicating OSPF wasn't functioning properly.



```
R2(config-if)#ip mtu ?
<68-1500> MTU (bytes)

R2(config-if)#ip mtu 1400
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    1     FULL/BDR        00:00:34    192.168.1.1    GigabitEthernet0/0
R2(config-if)#do clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2(config-if)#
*Aug 23 05:16:07.474: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    1     EXSTART/DR      00:00:38    192.168.1.1    GigabitEthernet0/0

*Aug 23 05:21:12.946: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from EXSTART to DOWN, Neighbor Down: Too many retransmissions
R2(config-if)#
*Aug 23 05:22:12.946: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from DOWN to DOWN, Neighbor Down: Ignore timer expired


R2(config-if)#no ip mtu
R2(config-if)#
*Aug 23 05:25:49.362: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
```

- **NO IP MTU** was used to return the MTU to the default value, and then R1 and R2 reached the FULL state.
- If your OSPF neighbors are having trouble reaching the FULL state, make sure to check the MTU settings.

Rule 8. OSPF Network Type Must Match

- The OSPF network type must match between neighbors.

Example:




```

R2(config)#interface l0
R2(config-if)#
*Aug 23 05:52:53.898: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
*Aug 23 05:52:54.898: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed state to up
R2(config-if)#ip address 2.2.2.2 255.255.255.255
R2(config-if)#router ospf 1
R2(config-router)#network 2.2.2.2 0.0.0.0 area 0
R2(config-router)#interface g0/0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#
*Aug 23 05:53:34.818: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface
*Aug 23 05:53:34.914: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    0     FULL/ -         00:00:36    192.168.1.1    GigabitEthernet0/0
R2(config-if)#
  
```

- A loopback interface was configured on R2 with an IP address, and advertised it to R1.




```

R2(config)#interface l0
R2(config-if)#
*Aug 23 05:52:53.898: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
*Aug 23 05:52:54.898: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed state to up
R2(config-if)#ip address 2.2.2.2 255.255.255.255
R2(config-if)#router ospf 1
R2(config-router)#network 2.2.2.2 0.0.0.0 area 0
R2(config-router)#interface g0/0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#
*Aug 23 05:53:34.818: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interfac
*Aug 23 05:53:34.914: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    0     FULL/ -         00:00:36    192.168.1.1    GigabitEthernet0/0
R2(config-if)#
  
```

- The network type on R2's G0/0 was changed to point-to-point.
- R1's G0/0 is still using the default broadcast network type.




```

R2(config)#interface lo
R2(config-if)#
*Aug 23 05:52:53.898: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
*Aug 23 05:52:54.898: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed state to up
R2(config-if)#ip address 2.2.2.2 255.255.255.255
R2(config-if)#router ospf 1
R2(config-router)#network 2.2.2.2 0.0.0.0 area 0
R2(config-router)#interface g0/0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#
*Aug 23 05:53:34.818: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or forced down
*Aug 23 05:53:34.914: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on GigabitEthernet0/0 from LOADING to FULL, Loading Done
R2(config-if)#do show ip ospf neighbor

Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.1.1    0     FULL/ -         00:00:36    192.168.1.1    GigabitEthernet0/0
R2(config-if)#
  
```

- A message is displayed saying the neighbor went down, but then it went right back to the FULL state.
- R1 even displays as FULL in **SHOW IP OSPF NEIGHBOR**.
- **Verification on R1:**

```
R1# show ip ospf neighbor
```



```
R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.2	1	FULL/BDR	00:00:31	192.168.1.2	GigabitEthernet0/0

```
R1#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
 + - replicated route, % - next hop override

Gateway of last resort is not set

192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks


C 192.168.1.0/30 is directly connected, GigabitEthernet0/0

L 192.168.1.1/32 is directly connected, GigabitEthernet0/0

```
R1#
```

- R2 appears in the neighbor table in the FULL state, so you might think OSPF is working fine.
- **Verification:**

```
R1# show ip route
```

```
R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.2	1	FULL/BDR	00:00:31	192.168.1.2	GigabitEthernet0/0

```
R1#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
 + - replicated route, % - next hop override

Gateway of last resort is not set

```

192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/30 is directly connected, GigabitEthernet0/0
L    192.168.1.1/32 is directly connected, GigabitEthernet0/0
  
```

```
R1#
```

- But R2's loopback address should be in R1's routing table, but it's not.
- This is what will happen when the OSPF network type doesn't match.
- It can be tricky to troubleshoot because the neighbor state is Full, so it can seem like everything is working fine.
- Make sure to watch out for that when troubleshooting OSPF.

LSA Types Overview

- The OSPF LSDB is made up of LSAs.
- All routers in the same OSPF area share the same LSDB.
- There are 11 types of LSA, but there are only 3 you should know for the CCNA: Type 1, Type 2, and Type 5.

Type 1 - Router LSA

- Every router running OSPF generates this type of LSA.
- The router LSA identifies the router using its router ID.
- The router LSA lists networks attached to the router's OSPF-activated interfaces.


Type 2 - Network LSA

- It is generated by the DR of each 'multi-access' network.
- An example of a multi-access network is an Ethernet network using the broadcast network type.
- This type of LSA lists the routers which are attached to the multi-access network.

Type 5 - AS-External LSA


- This type of LSA is generated by ASBRs to describe routes to destinations outside of the autonomous system (the OSPF domain).

Verifying LSAs



Verification:

```
R1# show ip ospf database
```



```
R1#show ip ospf database
```

OSPF Router with ID (1.1.1.1) (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	1396	0x80000002	0x00FE8D	4
2.2.2.2	2.2.2.2	932	0x80000005	0x00753F	4
3.3.3.3	3.3.3.3	974	0x80000004	0x00AD70	2
4.4.4.4	4.4.4.4	975	0x80000005	0x004CC2	2
5.5.5.5	5.5.5.5	976	0x80000004	0x00D212	3

Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.2.3	4.4.4.4	932	0x80000002	0x00740D

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	4.4.4.4	273	0x80000002	0x00C0E0	1

```
R1#
```

- It doesn't matter which router the command is used on, the output will be the same because all routers in the area have the same LSDB.
- Each router has generated a Type 1 Router LSA identifying itself. Each of these router LSAs contains information about the networks each router is connected to.

```
R1#show ip ospf database
```

```
OSPF Router with ID (1.1.1.1) (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	1396	0x80000002	0x00FE8D	4
2.2.2.2	2.2.2.2	932	0x80000005	0x00753F	4
3.3.3.3	3.3.3.3	974	0x80000004	0x00AD70	2
4.4.4.4	4.4.4.4	975	0x80000005	0x004CC2	2
5.5.5.5	5.5.5.5	976	0x80000004	0x00D212	3

```
Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.2.3	4.4.4.4	932	0x80000002	0x00740D

```
Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	4.4.4.4	273	0x80000002	0x00C0E0	1

```
R1#
```

- There is only one Type 2 Network LSA, generated by R4 for the 192.168.2.0/29 subnet.
 - Even though R1, R3, and R5 are DRs on their G1/0 interfaces, no other routers are connected to the interfaces so no Type 2 LSA is generated.

```
R1#show ip ospf database
```

```
OSPF Router with ID (1.1.1.1) (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	1396	0x80000002	0x00FE8D	4
2.2.2.2	2.2.2.2	932	0x80000005	0x00753F	4
3.3.3.3	3.3.3.3	974	0x80000004	0x00AD70	2
4.4.4.4	4.4.4.4	975	0x80000005	0x004CC2	2
5.5.5.5	5.5.5.5	976	0x80000004	0x00D212	3

```
Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.2.3	4.4.4.4	932	0x80000002	0x00740D

```
Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	4.4.4.4	273	0x80000002	0x00C0E0	1

```
R1#
```

- A Type 5 AS-External LSA is generated by R4. It is sharing its default route to the Internet with the other routers.

Lecture Review and Summary

- **OSPF Network Types:**
 - Focused on the two needed for the CCNA: Broadcast and Point-to-Point.
 - Ethernet connections are dominant, so the Broadcast network type is most commonly used.
 - Be familiar with the Point-to-Point type and the basics of serial interfaces.
- **OSPF Neighbor and Adjacency Requirements:**
 - You don't have to memorize the list, but you must be able to identify them and fix any problems.

- **OSPF LSA Types:**

- Introduced three basic types: Type 1 (Router LSA), Type 2 (Network LSA), and Type 5 (AS-External LSA).