OSPF PART 2

Introduction

- OSPF is item 3.4 on the official CCNA exam topics list
- Assumes knowledge of basic single-area OSPFv2 configuration and OSPF router ID from Part 1

Topics Covered in This Lecture

- OSPF's metric (cost)
- How routers become OSPF neighbors (the process)
- Additional OSPF configurations

OSPF's Metric (Cost)

- OSPF's metric is called **cost**
- Automatically calculated based on the bandwidth (speed) of the interface
- You can also manually configure the cost of each interface (covered later)

OSPF Cost Formula

- Cost = Reference Bandwidth ÷ Interface Bandwidth
- **Default reference bandwidth**: 100 megabits per second

Default Cost Examples

Interface Type	Speed	Calculation	Cost
Ethernet	10 Mbps	100 ÷ 10	10
FastEthernet	100 Mbps	100 ÷ 100	1
Gigabit Ethernet	1000 Mbps	100 ÷ 1000 = 0.1	1
10Gig Ethernet	10000 Mbps	100 ÷ 10000 = 0.01	1

- In OSPF, all values less than 1 are converted to 1
- Therefore FastEthernet, Gigabit Ethernet, and 10Gig Ethernet all have a cost of 1 by default

Verification: Checking OSPF Cost

```
R3# show ip ospf interface FastEthernet 2/0
```

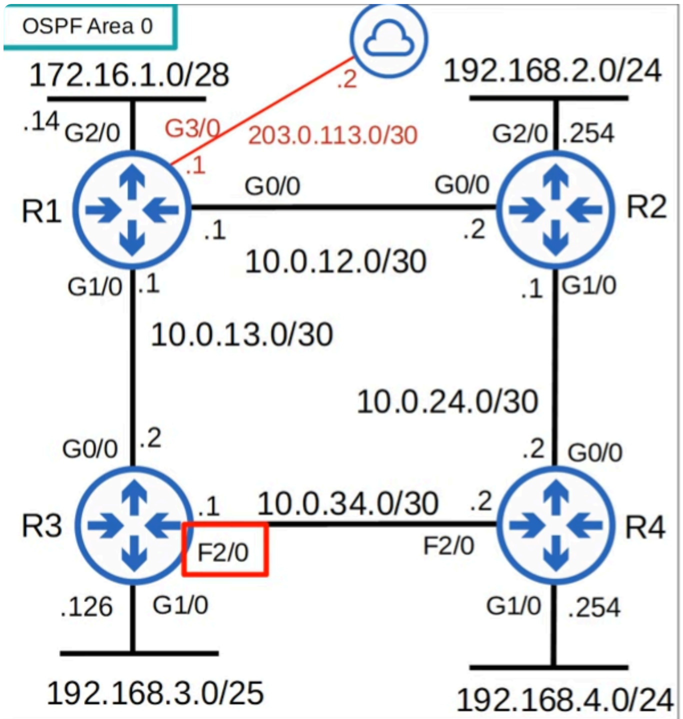

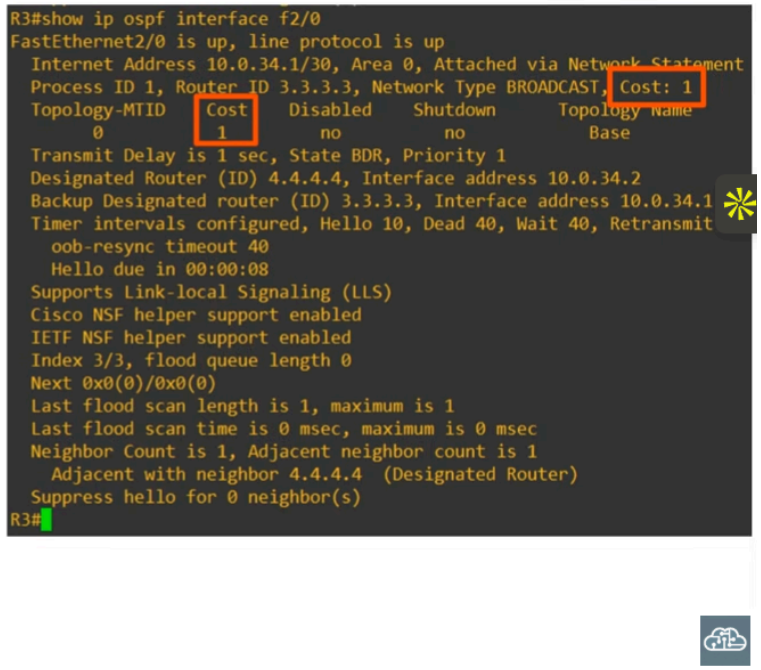

```
R3#show ip ospf interface f2/0
FastEthernet2/0 is up, line protocol is up
  Internet Address 10.0.34.1/30, Area 0, Attached via Network Statement
  Process ID 1, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 4.4.4.4, Interface address 10.0.34.2
  Backup Designated router (ID) 3.3.3.3, Interface address 10.0.34.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit
    oob-resync timeout 40
    Hello due in 00:00:08
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 3/3, 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 4.4.4.4  (Designated Router)
  Suppress hello for 0 neighbor(s)
R3#
```

- Cost is displayed in two places in the output
- Default cost of FastEthernet interface is 1

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

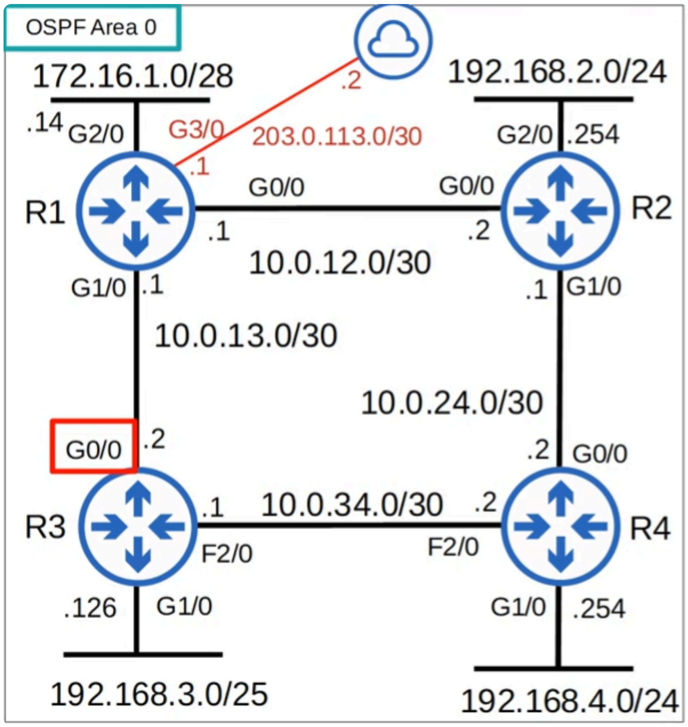

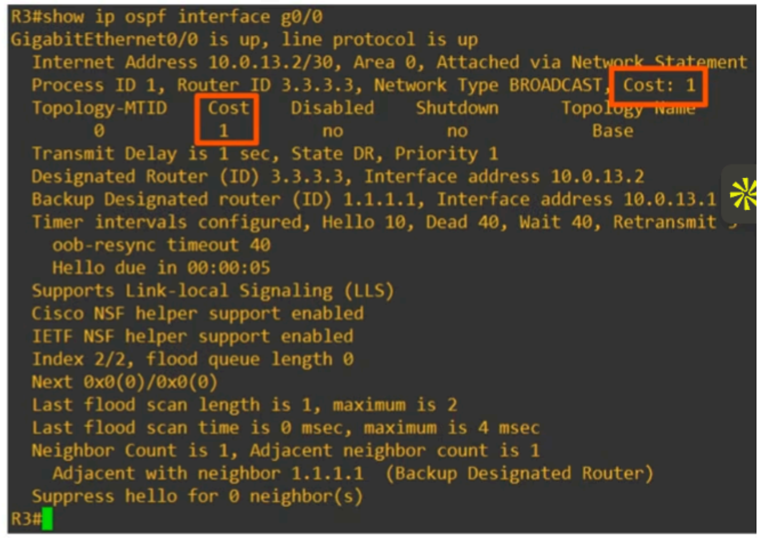

```
R3#show ip ospf interface g0/0
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 10.0.13.2/30, Area 0, Attached via Network Statement
  Process ID 1, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 3.3.3.3, Interface address 10.0.13.2
  Backup Designated router (ID) 1.1.1.1, Interface address 10.0.13.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:05
  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 2
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
R3#
```

- GigabitEthernet interface also has a cost of 1 by default

Changing the Reference Bandwidth

- The default situation is not ideal — you can and should change the reference bandwidth

```
R1(config-router)# auto-cost reference-bandwidth <value in Mbps>
```

- Reference bandwidth is configured in **megabits per second**

Example: Configuring Reference Bandwidth

```
R1(config-router)# auto-cost reference-bandwidth 100000
```

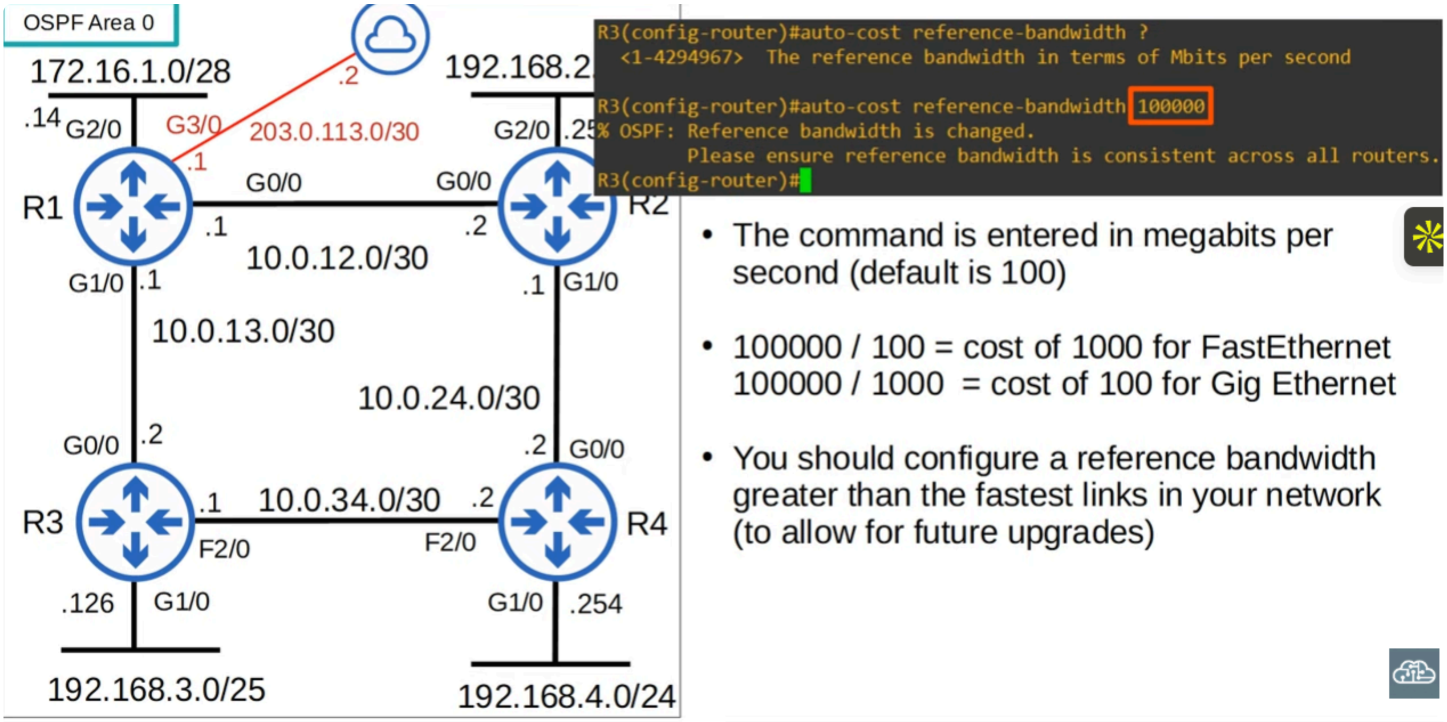

- Configured value: 100,000 (in megabits per second)
- FastEthernet cost: 100,000 ÷ 100 = **1000**
- Gigabit Ethernet cost: 100,000 ÷ 1000 = **100**

Why Configure a Large Reference Bandwidth

- Configure a reference bandwidth greater than the fastest links in your network to allow for future upgrades
- 100,000 Mbps = 100 Gbps, which is 100 times faster than the fastest interfaces in the example network (regular Gigabit Ethernet)

Important Warning

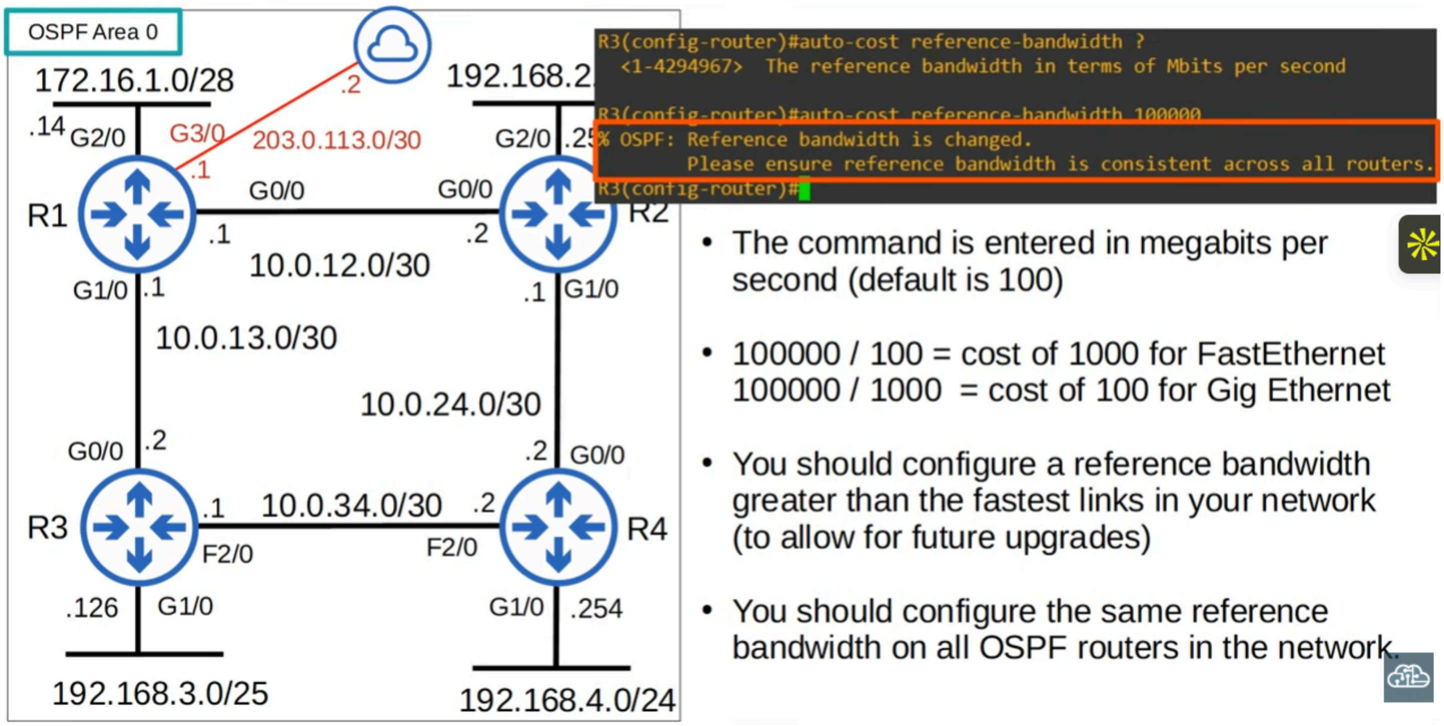

- A message is displayed when you configure the reference bandwidth: *"Please ensure reference bandwidth is consistent across all routers"*
- You should configure the same reference bandwidth on all OSPF routers in the network to provide consistent cost values for each interface bandwidth

---

Calculating OSPF Cost to a Destination

- The OSPF cost to a destination is the total cost of the 'outgoing' or 'exit' interfaces (just like spanning tree cost)
- All routers in the example have been set to a reference bandwidth of 100 Gbps (100,000 Mbps), making Gigabit Ethernet interfaces have a cost of 100

Example 1: R1's cost to reach 192.168.4.0/24

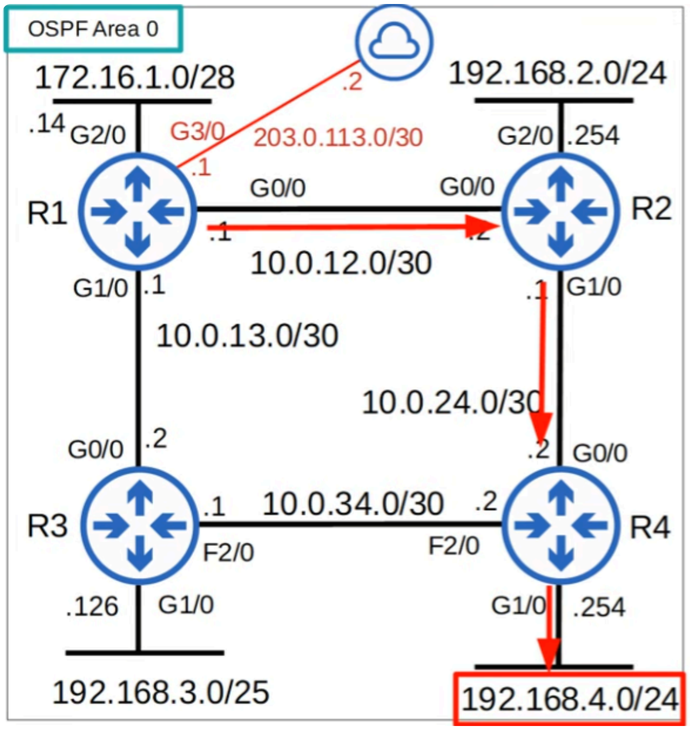

- The OSPF cost to a destination is the total cost of the 'outgoing/exit interfaces'
- For example, R1's cost to reach 192.168.4.0/24 is :
  100 (R1 G0/0) + 100 (R2 G1/0) + 100 (R4 G1/0) = 300

- Packet exits: R1's G0/0 → R2's G1/0 → R4's G1/0
- Calculation: 100 + 100 + 100 = **300**

Example 2: R1's cost to reach 2.2.2.2 (R2's Loopback0)

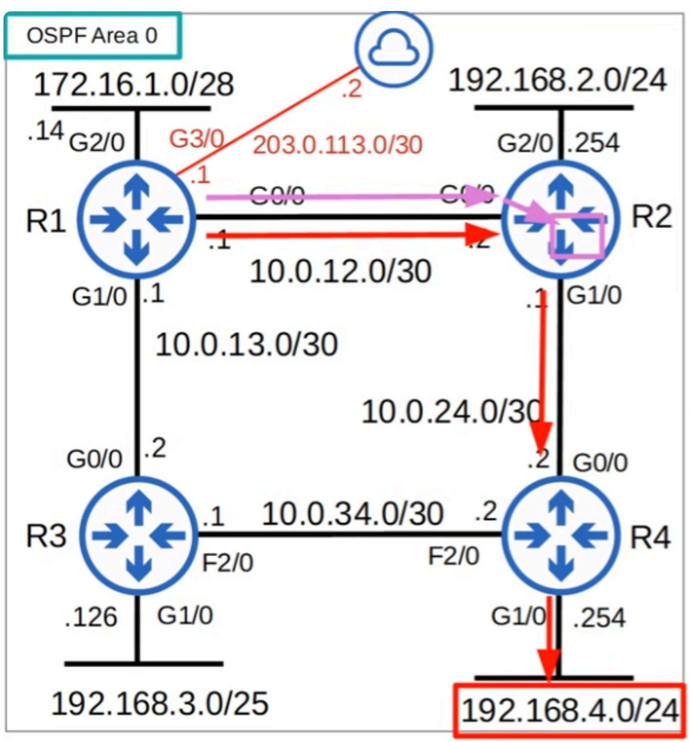

- The OSPF cost to a destination is the total cost of the 'outgoing/exit interfaces'
- For example, R1's cost to reach 192.168.4.0/24 is :
  100 (R1 G0/0) + 100 (R2 G1/0) + 100 (R4 G1/0) = 300
- Loopback interfaces have a cost of 1
- What is R1's cost to reach 2.2.2.2 (R2's loopback0 interface)?
- 100 (R1 G1/0) + 1 (R2 L0) = 101

- Packet exits: R1's G0/0 → R2's Loopback0

- Loopback interfaces have a cost of 1
- Calculation: 100 + 1 = **101**

Verification: R1's Routing Table (Before vs. After Reference Bandwidth Change)

Before (default reference bandwidth of 100 Mbps)

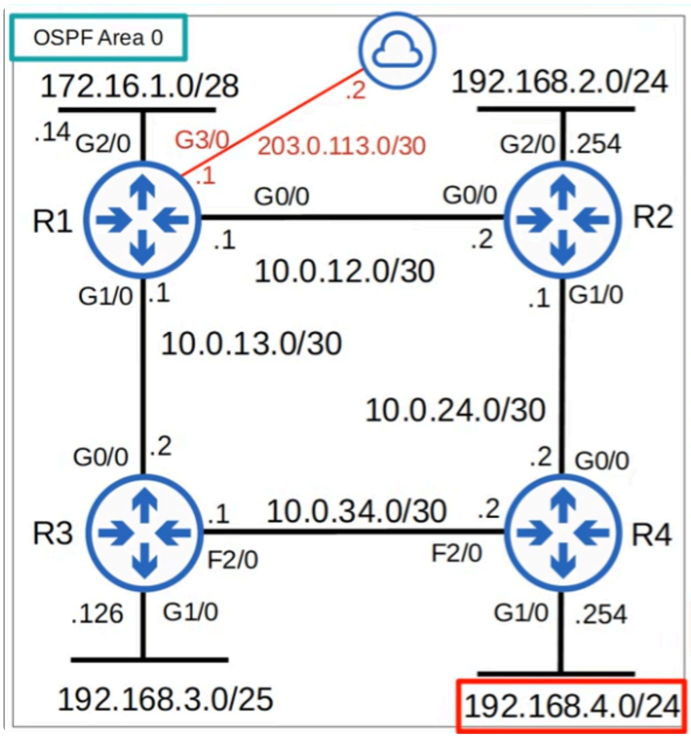

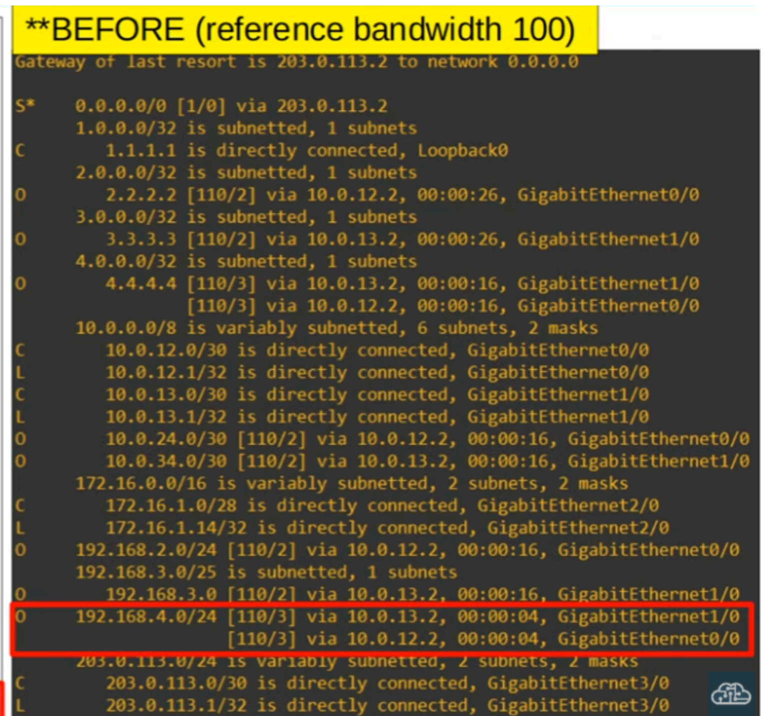

- R1 has two routes to 192.168.4.0 (one via R2, one via R3)
- The slower FastEthernet connection (R3 to R4) has the same cost of 1 as the Gigabit Ethernet interfaces, so both paths are equal

After (reference bandwidth changed to 100,000 Mbps on all routers)

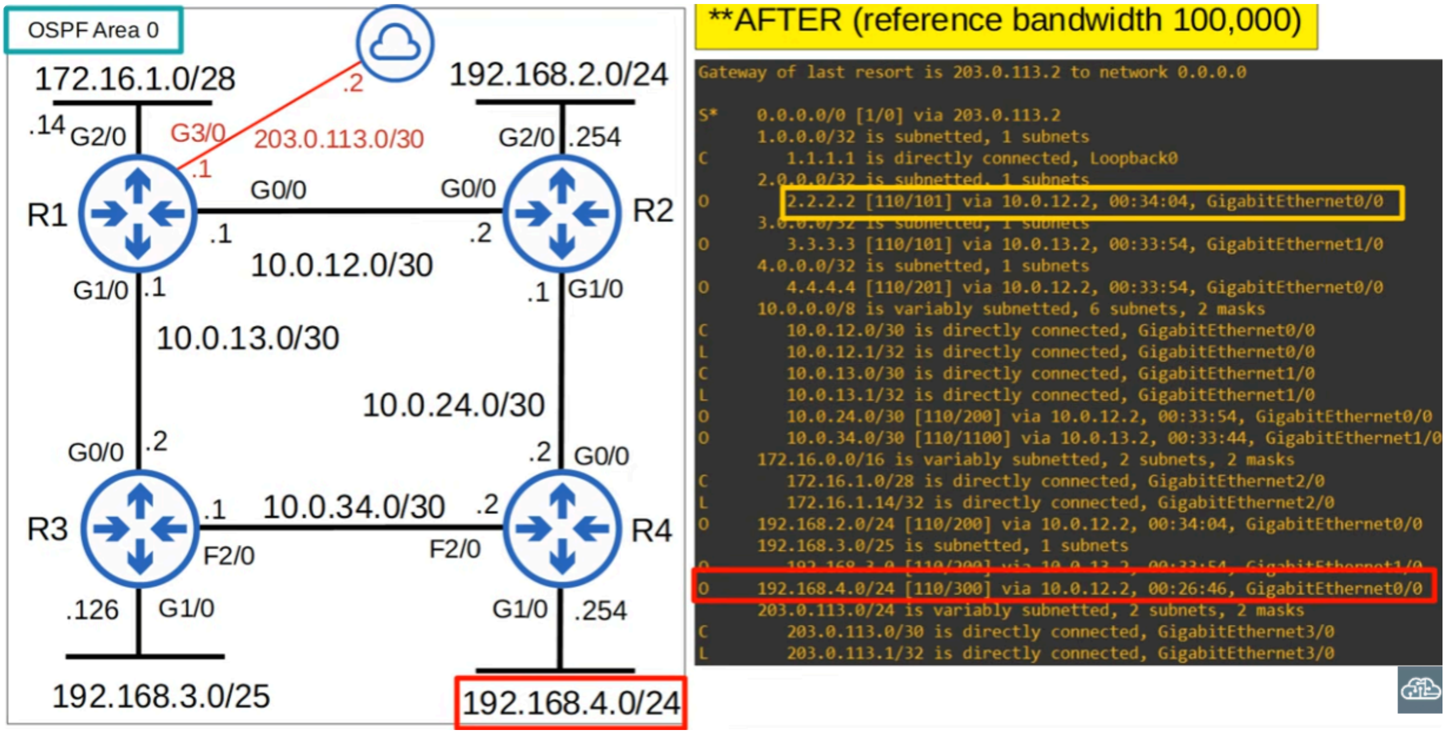

- R1 only has one route to 192.168.4.0
- Cost to 192.168.4.0 is 300 (as calculated above)
- Cost to 2.2.2.2 is 101 (as calculated above)

Manually Configuring OSPF Cost

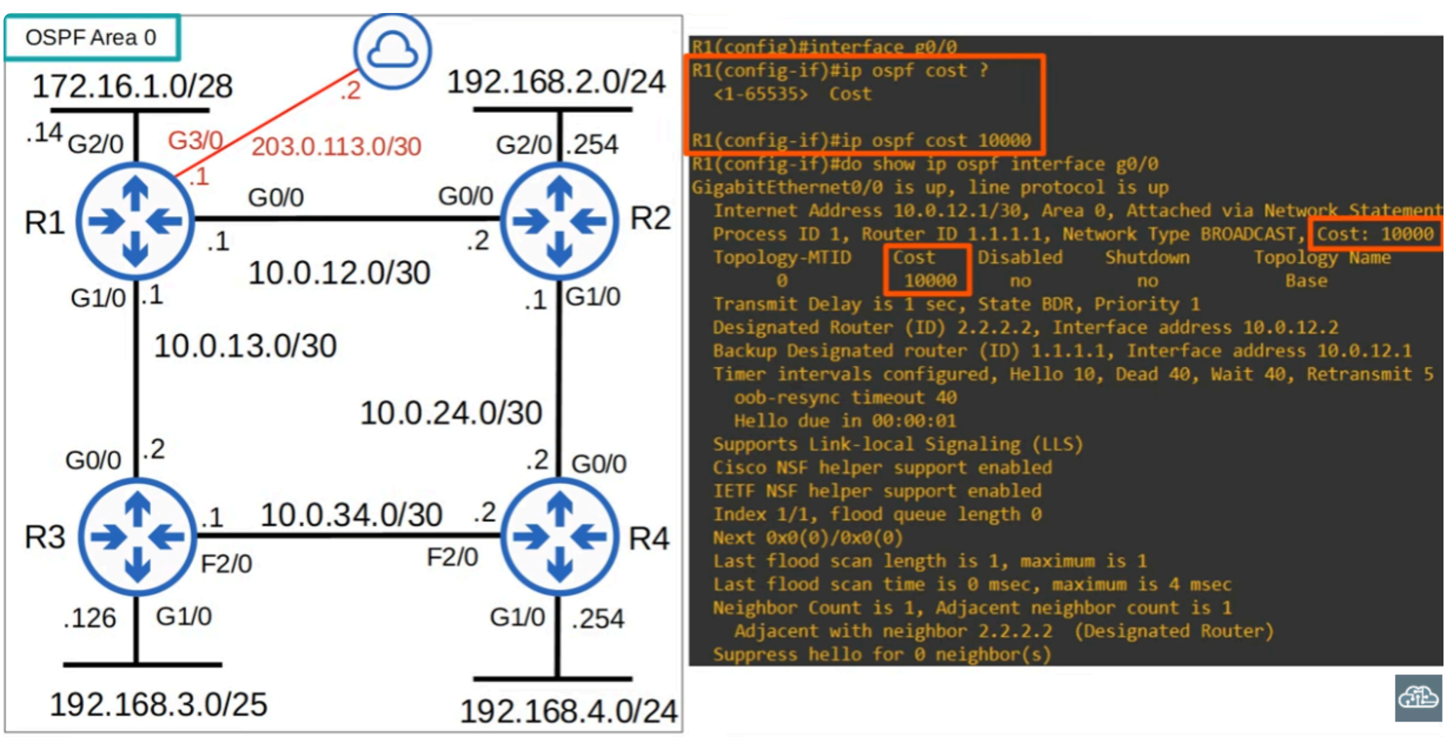

- You can manually configure the OSPF cost of an interface
- This manually configured cost takes priority over the auto-calculated cost

```
R1(config-if)# ip ospf cost <value>
```

- **Example:** Configured R1's G0/0 interface with a cost of 10,000
- The cost is now 10,000 instead of the auto-calculated 100

Changing Interface Bandwidth to Affect Cost

- Formula review: Cost = Reference Bandwidth ÷ Interface Bandwidth
- You can change the interface bandwidth to alter the cost calculation

Speed vs. Bandwidth

- **Bandwidth** and **Speed** are different
- Changing the interface **bandwidth** does NOT change the actual speed at which the interface operates
- The **bandwidth** value is just a number used in calculations (OSPF cost, EIGRP metric, etc.)
- To change the actual physical operating speed, use the **speed** command

Example:

- If you change a Gigabit Ethernet interface's bandwidth to 100 Mbps, it will still physically operate at 1 Gbps
- However, OSPF will calculate cost using 100 Mbps
- Because the bandwidth value is used in other calculations (not just OSPF cost), it is **not recommended** to change it to alter OSPF cost
- **Recommended method:** Change the reference bandwidth, then use the `ip ospf cost` command for individual interfaces if needed

The Bandwidth Command

- Configured in **kilobits per second** (different from the reference bandwidth, which is in megabits per second)

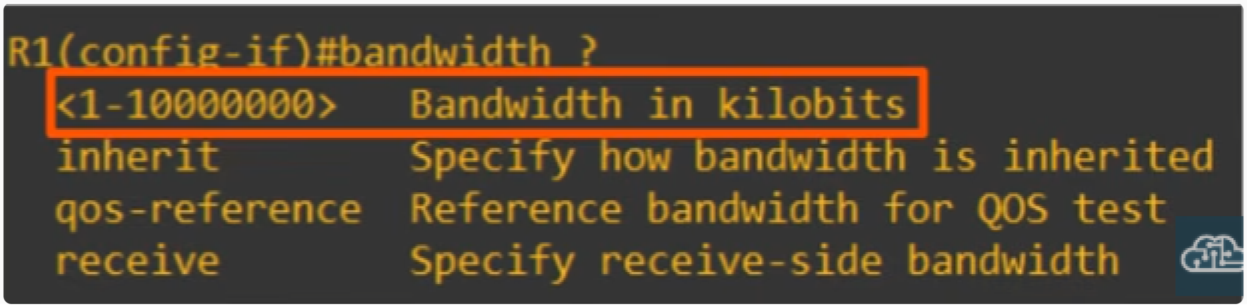


```
R1(config-if)# bandwidth <value in kbps>
```

- **Tip:** Always use the question mark (`?`) to check the units a command expects before entering it (e.g., time commands may use seconds or minutes; speed commands may use kilobits or megabits)

Three Ways to Modify OSPF Cost

1. **Change the reference bandwidth (Recommended)**

   ```
   R1(config-router)# auto-cost reference-bandwidth <value in Mbps>
   ```

2. **Manually configure the OSPF cost on the interface**

   ```
   R1(config-if)# ip ospf cost <value>
   ```

3. **Change the interface bandwidth (Not recommended)**

```
R1(config-if)# bandwidth <value in kbps>
```

Verification Commands

```
R1# show ip ospf interface brief
```

```
R3#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               3.3.3.3/32         1     LOOP  0/0
Gi1/0        1     0               192.168.3.126/25   100   DR    0/0
Fa2/0        1     0               10.0.34.1/30       1000  BDR   1/1
Gi0/0        1     0               10.0.13.2/30       100   DR    1/1
```

- Provides a convenient overview of the OSPF cost for each OSPF-enabled interface on the router.

OSPF Neighbors

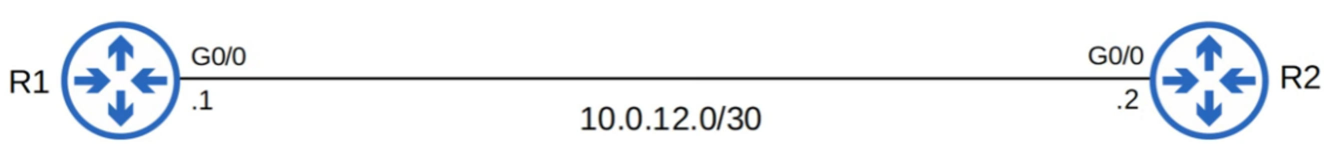

- Making sure routers successfully become OSPF neighbors is the main task in configuring and troubleshooting OSPF
- Once routers become neighbors, they automatically share network information and calculate routes

- If routers cannot become OSPF neighbors, OSPF cannot operate at all

OSPF Hello Messages

- When OSPF is activated on an interface, the router starts sending OSPF Hello messages out of the interface at regular intervals (determined by the hello timer)
- Hello messages are used to introduce the router to potential OSPF neighbors
- By exchanging Hello messages, routers check that they are compatible to become OSPF neighbors, and then negotiate their neighbor relationship

Key OSPF Details

- **Default Hello Timer:** 10 seconds on an Ethernet connection
- **Multicast Address:** OSPF Hello messages are multicast to 224.0.0.5, which is the multicast address for all OSPF routers
    - Note: RIP's multicast address is 224.0.0.9
    - Note: EIGRP's multicast address is 224.0.0.10
- **IP Protocol Number:** OSPF messages are encapsulated in an IP header, and the 'protocol' field of the IP header has a value of 89 to indicate OSPF

OSPF Neighbor States

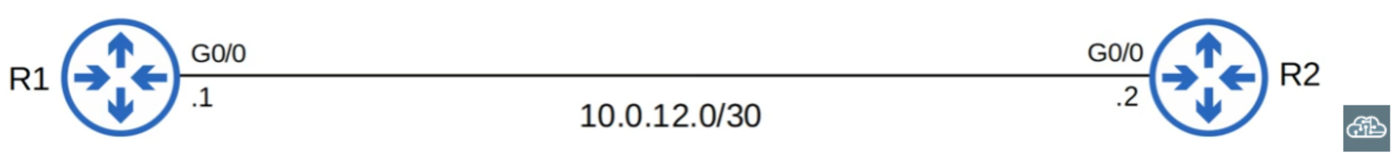

- For OSPF routers to become neighbors, they must go through various neighbor states.

Scenario: OSPF is already activated on R2's G0/0 interface. Then, OSPF is activated on R1's G0/0 interface.

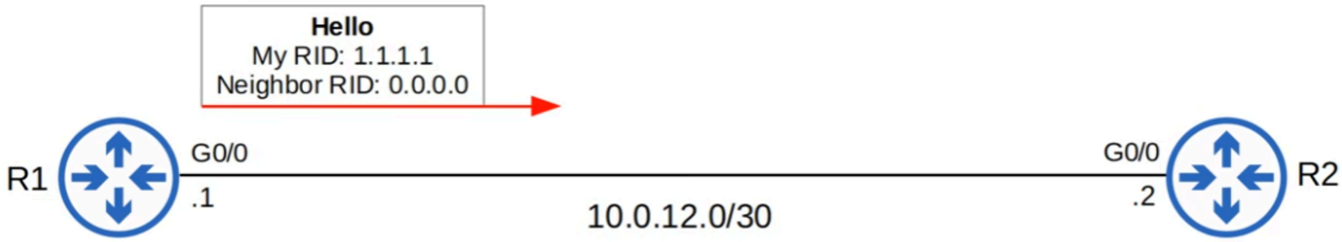

- R1 sends an OSPF Hello message to 224.0.0.5.
- Two important fields in the Hello message are the **Router ID** and the **Neighbor's Router ID**.
- R1 does not know about R2 yet, so the neighbor Router ID field is 0.0.0.0.

The 'Down' State

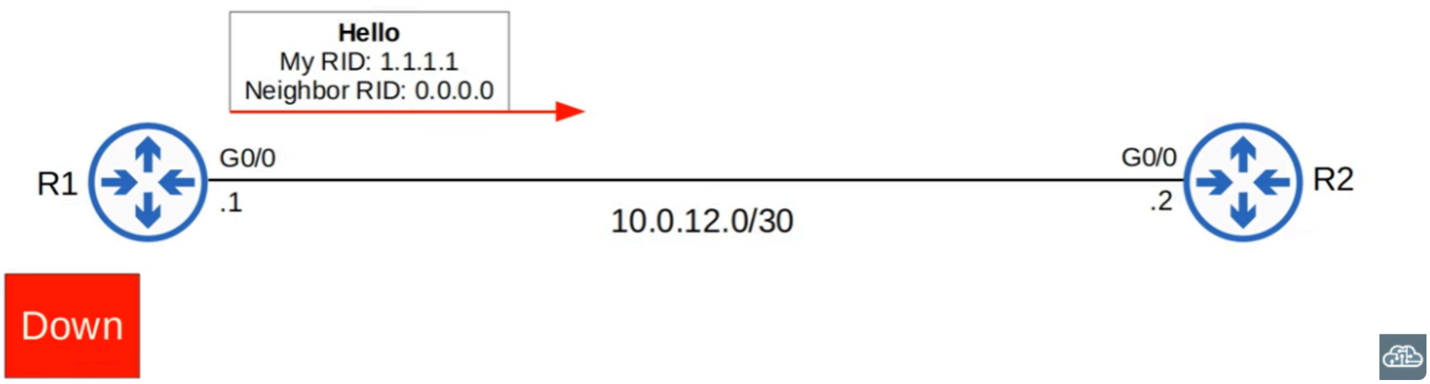

- R1 does not know about any OSPF neighbors yet, so the current neighbor state is **Down**.
- This is the first OSPF neighbor state.

The 'Init' State

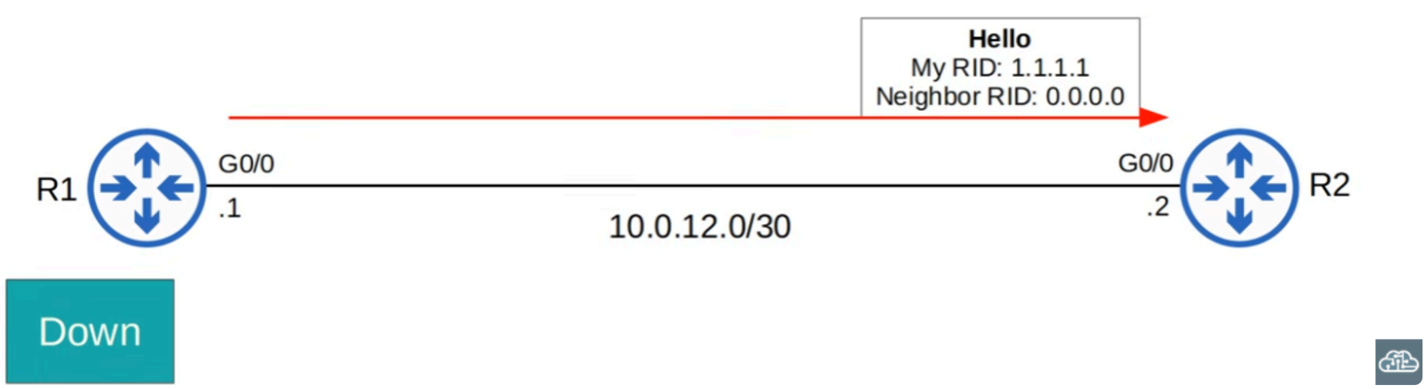

- When R2 receives the Hello packet, it adds an entry for R1 to its OSPF neighbor table.

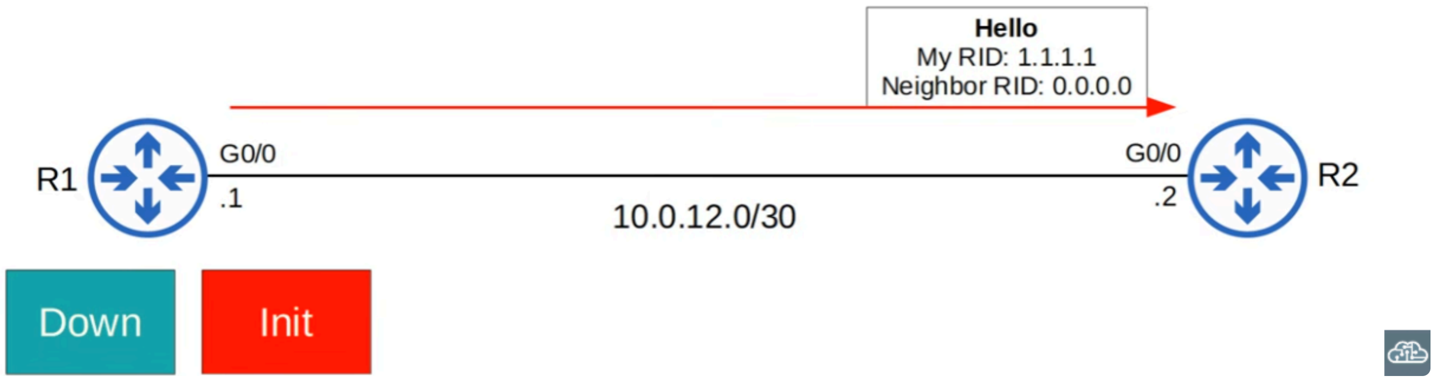

- In R2's neighbor table, the relationship with R1 is now in the **Init** state.
- R1 still does not know about R2, so it will have no entries in its OSPF neighbor table.
- The **Init** state means a Hello packet was received, but the receiver's own Router ID is not in that Hello packet.
- **Example:** R2's Router ID is 2.2.2.2, but the neighbor Router ID in R1's Hello packet is 0.0.0.0.

The '2-Way' State

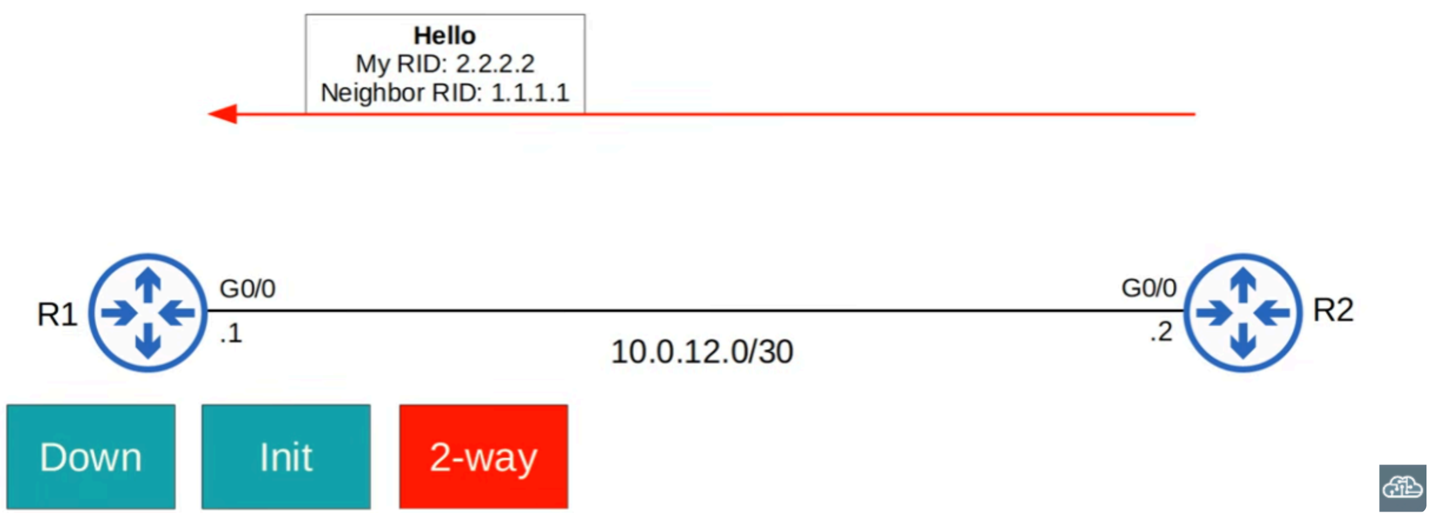

- **Transition:** R2 sends a Hello packet containing the RID of both routers. R1 then inserts R2 into its OSPF neighbor table in the 2-way state.

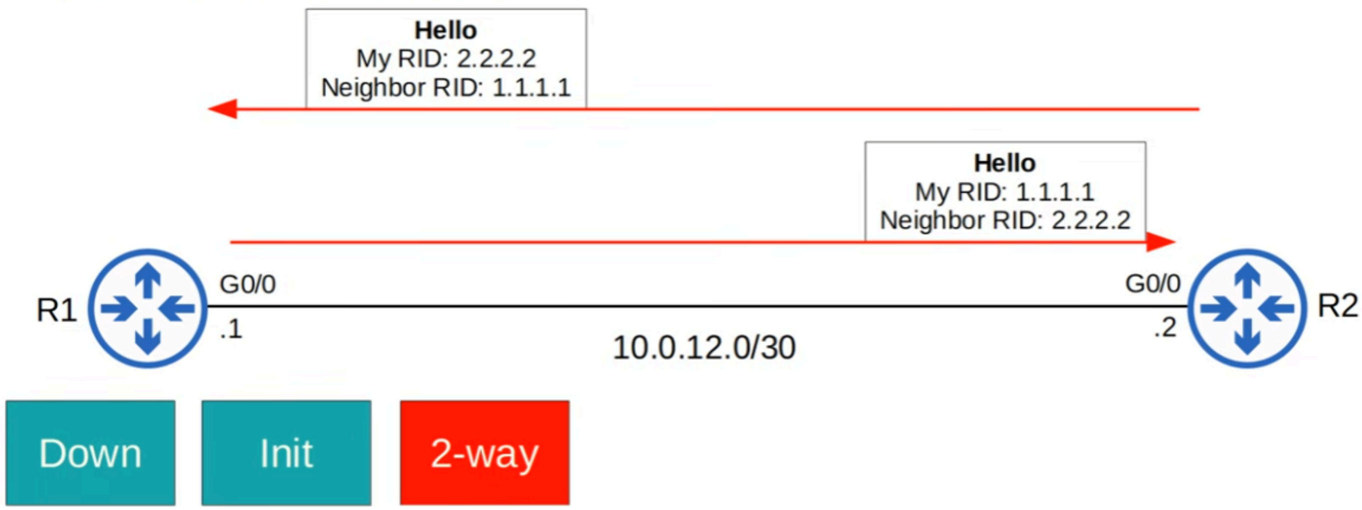

- R1 responds with another Hello message containing R2's RID.
- **Status:** Both routers are now in the 2-way state.
- **Meaning:** This state indicates the router has received a Hello packet containing its own RID. Reaching this state means all conditions are met for the routers to become OSPF neighbors.
- **Action:** The routers are now ready to share LSAs to build a common LSDB.
- **Troubleshooting Note:** If routers fail to reach this 2-way state, you must troubleshoot to find the cause.
- **Additional Info:** In some network types, a **DR (Designated Router)** and **BDR (Backup Designated Router)** are elected at this point (covered in Day 28).
- **Conclusion:** At this point, the routers are already OSPF neighbors. Over the next few states, they will share LSAs and form a full OSPF adjacency.

The 'Exstart' State

- After the 2-way state, the two routers prepare to exchange information about their LSDB

- Before exchanging, they must choose which one will start the exchange
- They decide which will be the **Master router** and which will be the **Slave router**
- Note: These are different than the DR and BDR mentioned previously; this Master/Slave relationship is only needed for this initial exchange of LSDB information
- The router with the **higher RID** becomes the Master and initiates the exchange
- The router with the **lower RID** becomes the Slave
- **Example:** R2 has the higher RID, so R2 will be the Master and R1 will be the Slave
- To decide the Master and Slave, they exchange **DBD (Database Description) packets**

- R1 sends a DBD packet claiming to be the Master.

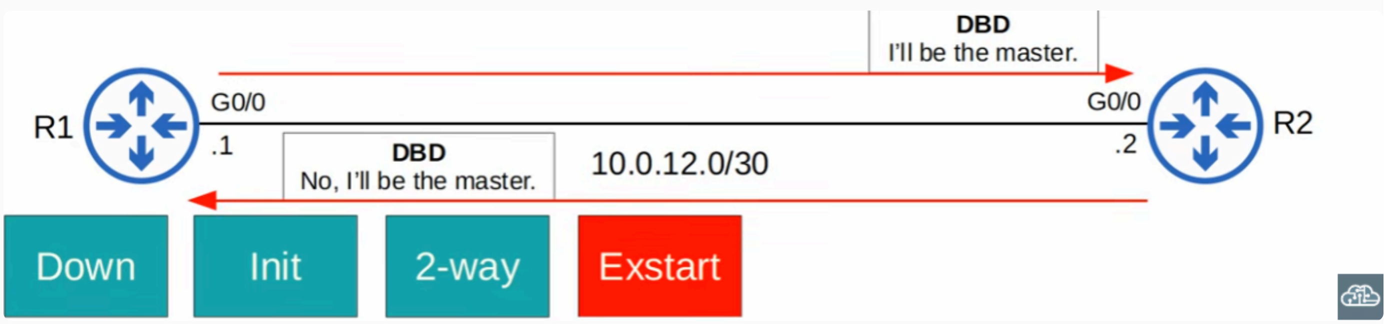

- However, R2 corrects R1 — R2 has the higher Router ID and says it will be the Master
- The Exstart state is just to prepare for the next state

The 'Exchange' State

- The routers exchange DBDs (Database Description packets) which contain a list of the LSAs in their LSDB
- These DBDs do not include detailed information about the LSAs, just basic information telling their neighbor what LSAs they have
- Basically the routers are telling each other, "I have these LSAs", but they aren't actually sending the LSAs yet
- The routers compare the information in the DBD they received to the information in their own LSDB to determine which LSAs they must receive from their neighbor
- After exchanging DBDs, they move to the next state

The 'Loading' State

- In the Loading state, routers send **Link State Request (LSR)** messages to request that their neighbors send them any LSAs they don't have
- In the Exchange state they exchanged DBD packets, so they know which LSAs their neighbors are holding
- These LSRs are used to request any missing LSAs to make sure each router has the same LSAs

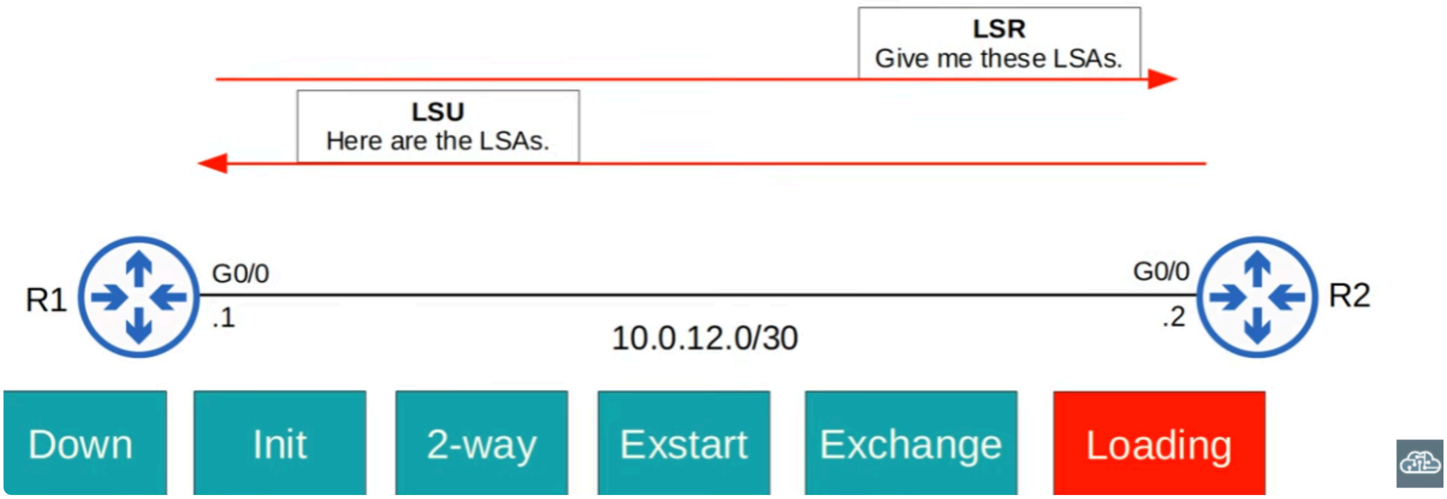

- **Example:** R2 will also send R1 LSRs for any missing LSAs (shown from one side of the exchange only)
- Then the LSAs themselves are sent in **Link State Update (LSU)** messages

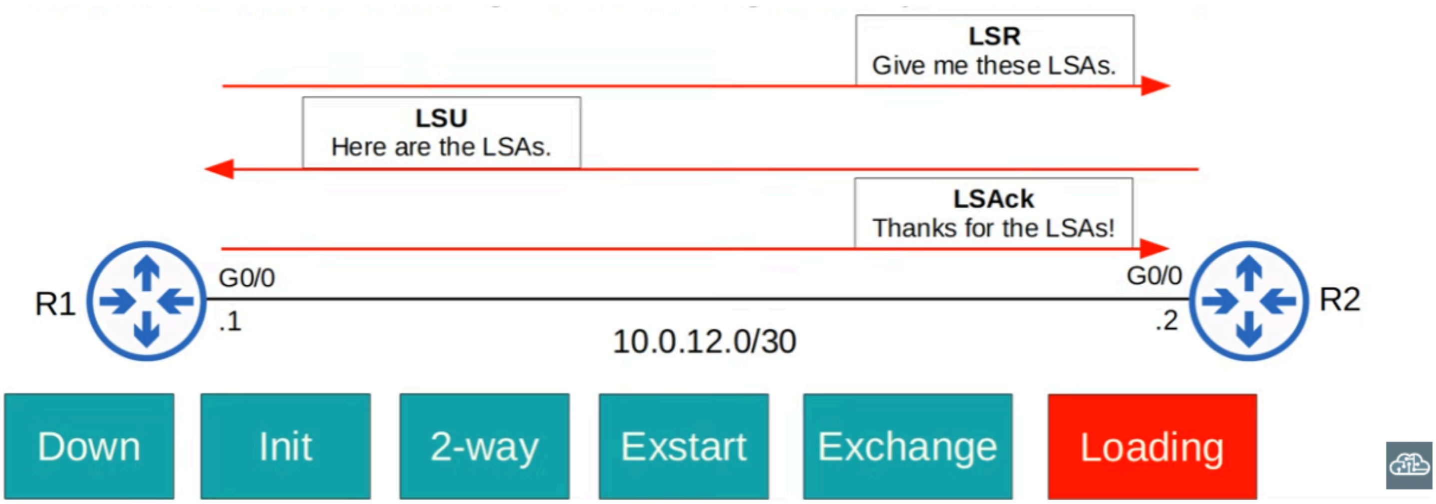

- **Example:** R2 sends R1 the requested LSAs in an LSU; R1 will also do the same for R2
- Finally, the routers send **LSAck** messages to acknowledge that they received the LSAs
- The Loading state is now complete, and the routers have the same LSDB

The 'Full' State

- In the Full state, the routers have a full OSPF adjacency and identical LSDBs.
- To maintain the adjacency, routers continue to send and listen for **Hello packets** once every 10 seconds (default).
- Another timer called the **'Dead' timer** is used to maintain the adjacency.
  - **Default Dead Timer:** 40 seconds.
  - Every time a Hello packet is received, the Dead timer is reset.
  - If the Dead timer counts down to 0 and no Hello message is received, the neighbor is removed.
- If neighbors remain up, routers will continue to share LSAs as the network changes to ensure each router has a complete and accurate map of the network.
- This is the main advantage of dynamic routing protocols: routers automatically react to changes and add, remove, or change routes as necessary.

OSPF Neighbor States Summary

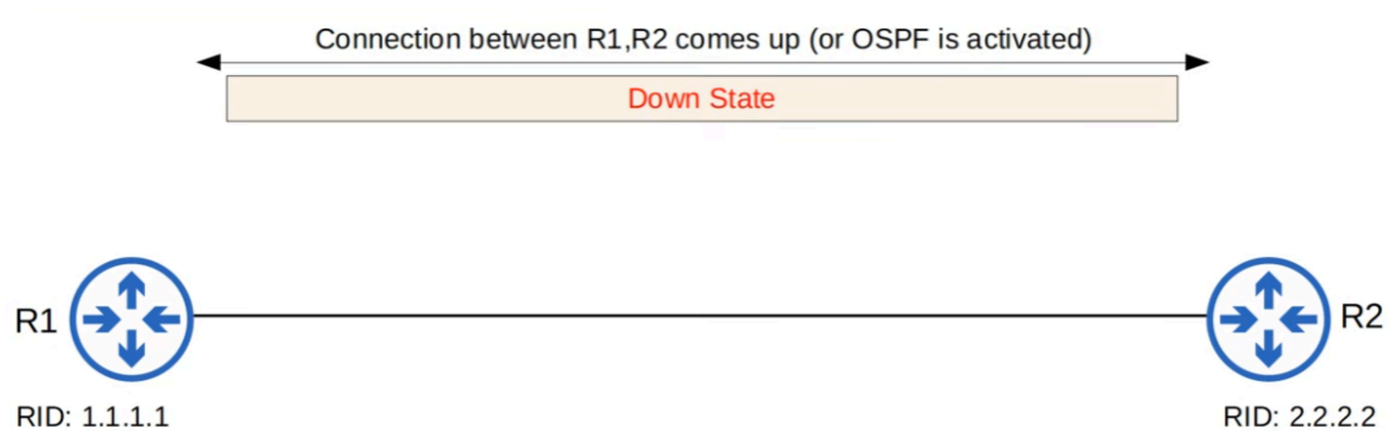

- **Down State:** The connection between R1 and R2 comes up (or OSPF is activated). R1 and R2 don't know each other yet, but they send Hello packets. Assuming R1 sends the first Hello packet.

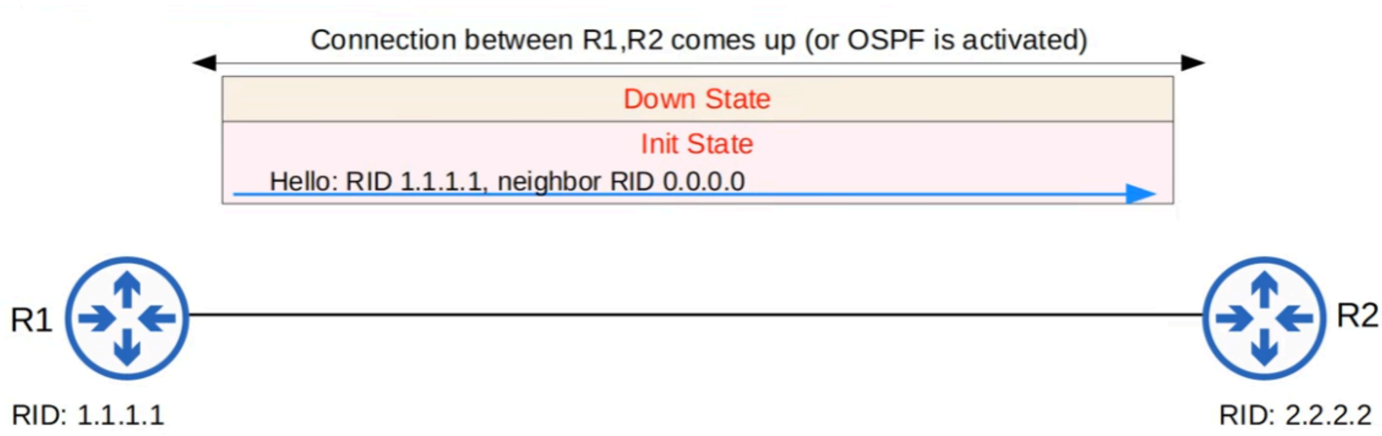

- **Init State:** R2 receives the first Hello packet from R1, but R2's own router ID isn't in the packet yet.

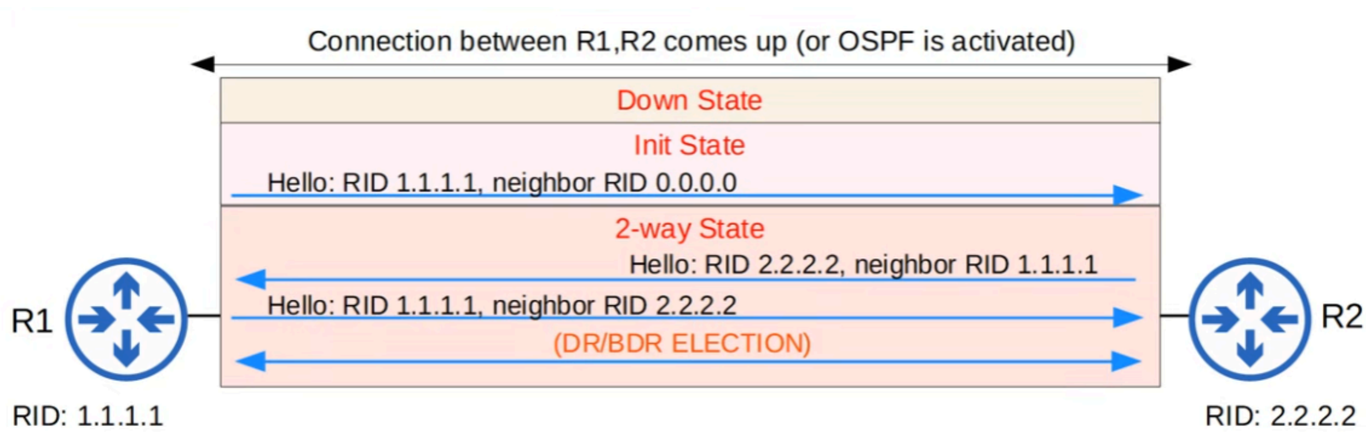

- **2-Way State:** The routers exchange more Hello packets. R1's router ID is included in R2's Hellos, and R2's router ID is included in R1's Hellos. In some OSPF connections, an election for the Designated Router (DR) and Backup Designated Router (BDR) takes place here.

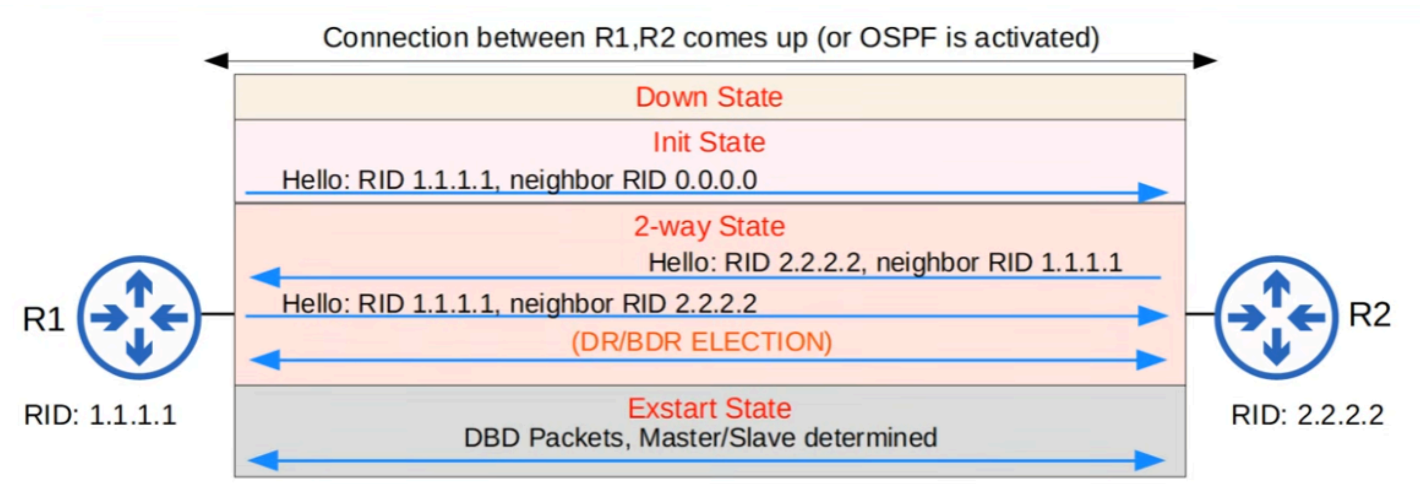

- **Exstart State:** The routers exchange DBD packets to determine which will be the Master and which will be the Slave. The Master is the router that starts the DBD exchange in the next state.

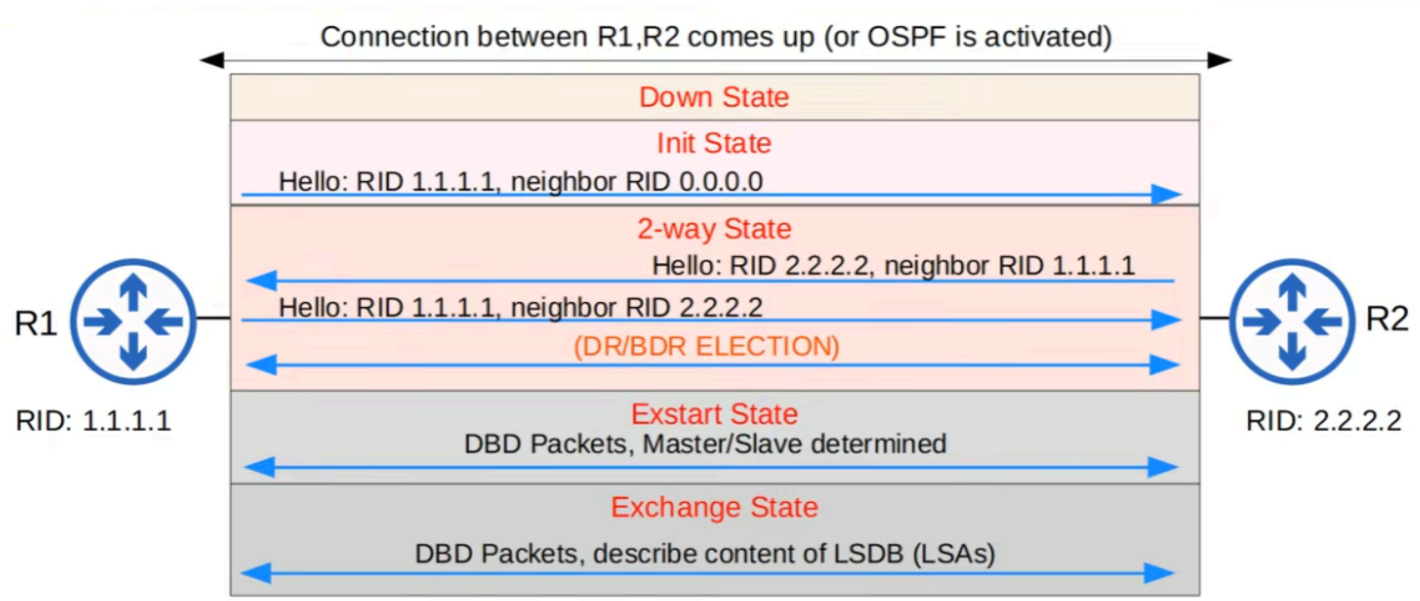

- **Exchange State:** The routers exchange DBD packets to tell each other about the contents of their LSDB.

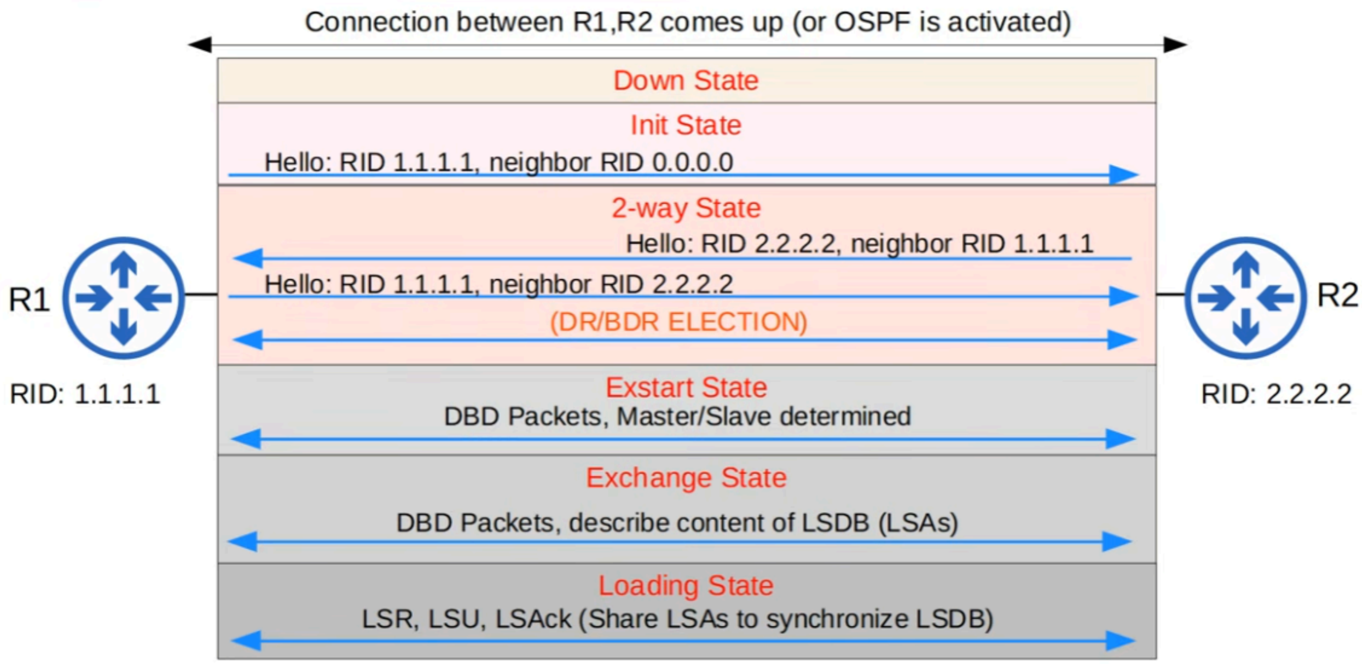

- **Loading State:** The routers use LSRs (Link State Requests) to request LSAs from each other. The LSAs are sent in LSU (Link State Update) packets. Finally, LSAck packets are sent to acknowledge that they received the LSAs.

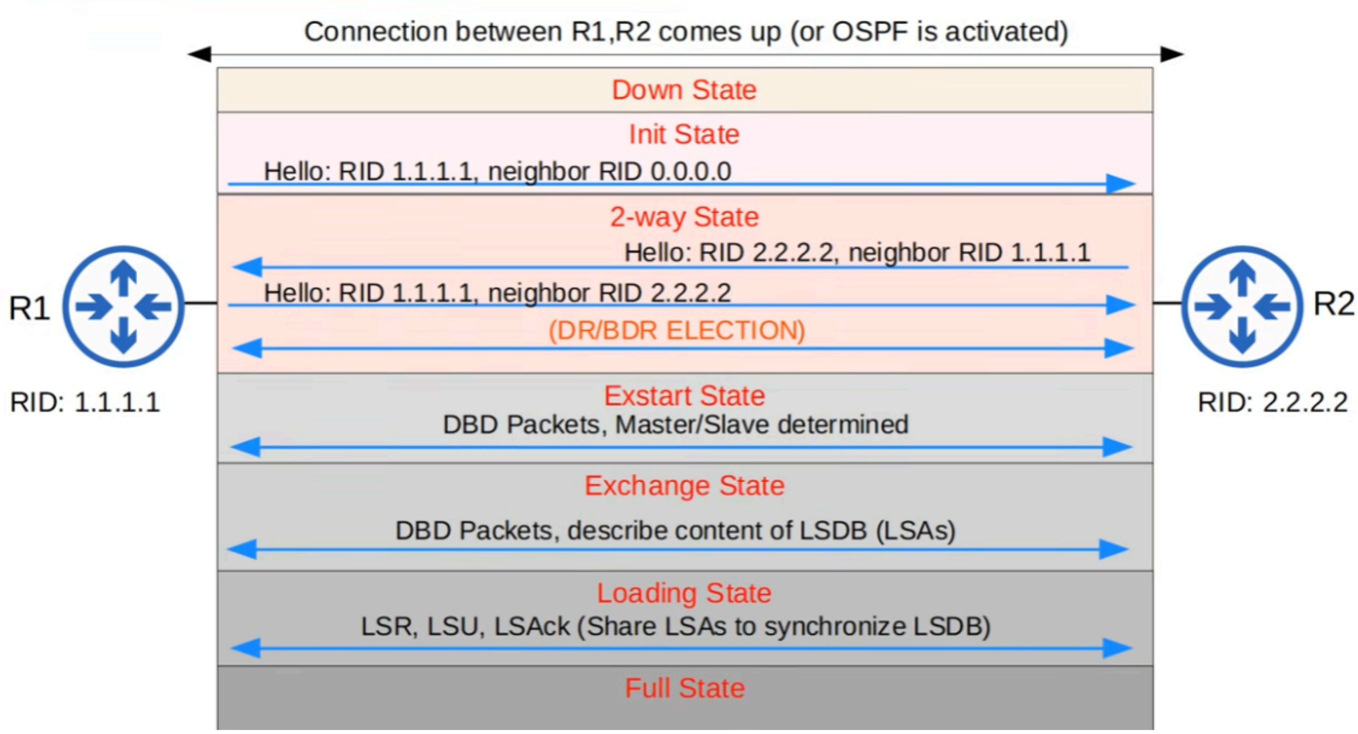

- **Full State:** The routers reach the Full state and have a full OSPF adjacency.

OSPF Process Phases

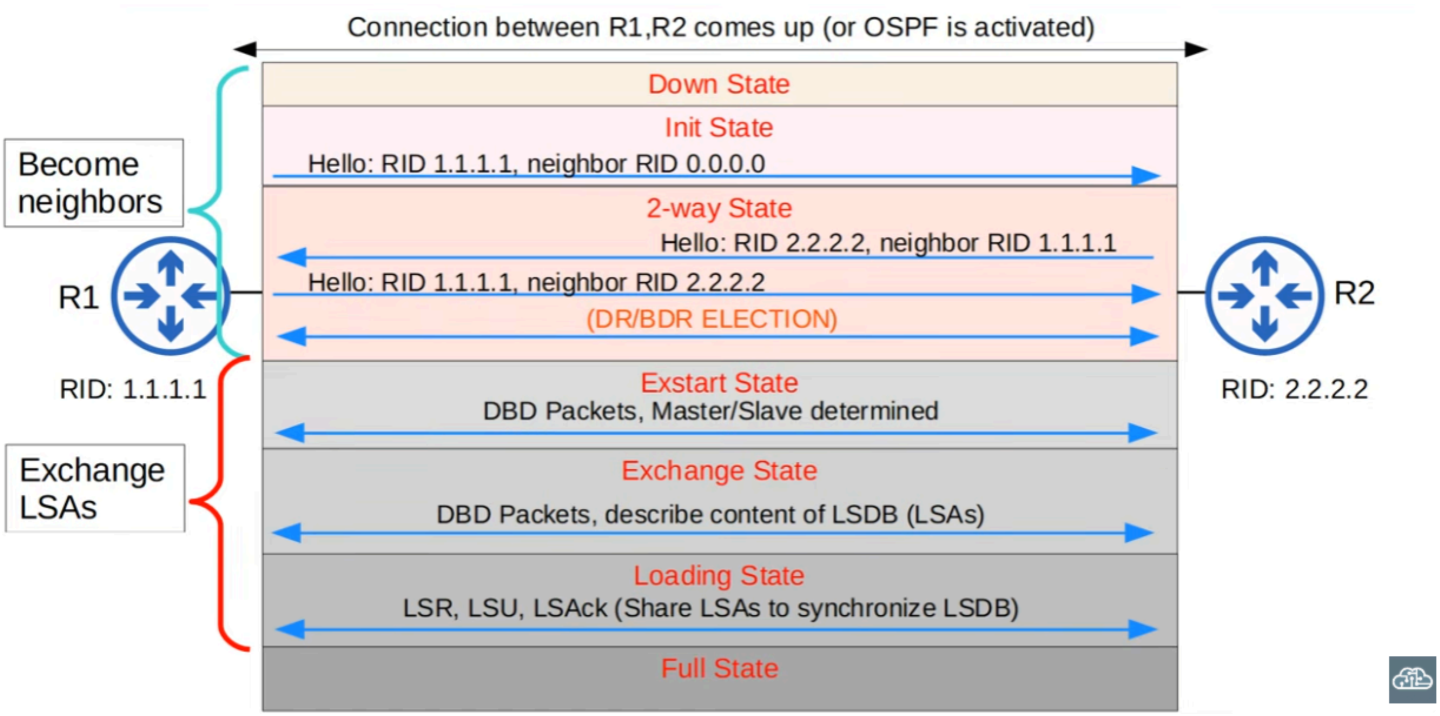

- **1. Become Neighbors:** Involves the Down, Init, and 2-Way states.
- **2. Exchange LSAs:** Involves the Exstart, Exchange, and Loading states to synchronize the LSDB.
- **3. Calculate Best Routes:** The routers use the OSPF metric (cost) to calculate the best route to each destination.

OSPF Message Types Summary

Type	Name	Purpose
1	**Hello**	Neighbor discovery and maintenance.
2	**Database Description (DBD)**	Summary of the LSDB of the router. Used to check if the LSDB of each router is the same.
3	**Link-State Request (LSR)**	Requests specific LSAs from the neighbor.
4	**Link-State Update (LSU)**	Sends specific LSAs to the neighbor.
5	**Link-State Acknowledgement (LSAck)**	Used to acknowledge that the router received a message.

#	Message Type	Purpose
1	Hello	Discover and maintain neighbors
2	DBD (Database Description)	Summarize LSDB contents (list of LSAs)
3	LSR (Link State Request)	Request specific LSAs from a neighbor
4	LSU (Link State Update)	Send requested LSAs to a neighbor
5	LSAck (Link State Acknowledgment)	Acknowledge receipt of LSAs

Verification: OSPF Show Commands

SHOW IP OSPF NEIGHBOR

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

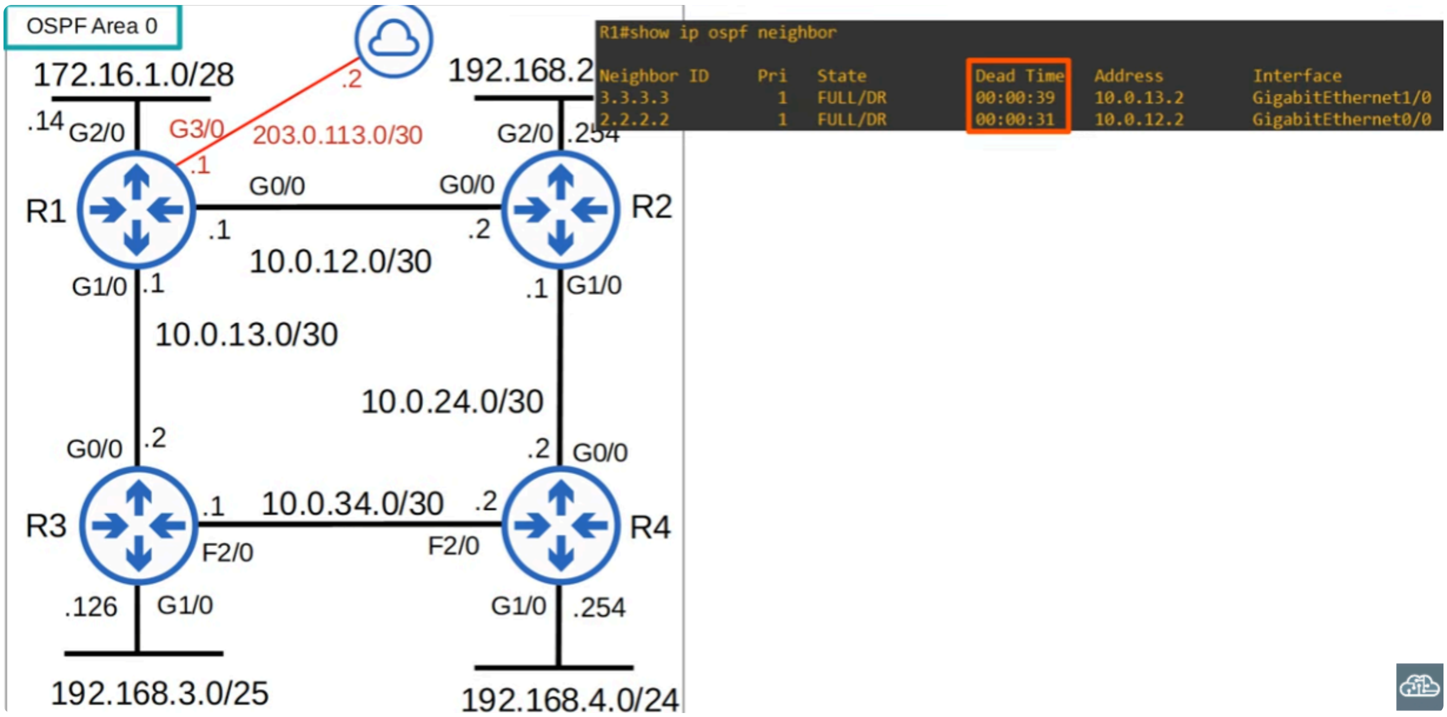

- Entered on R1
- Note the **Full** state with both neighbors R2 and R3
- Both R2 and R3 are DRs (covered in next video)
- **Dead Time:** Counts down from 40, but resets as soon as R1 receives a Hello packet from the neighbor
- Assuming a Hello packet is received every 10 seconds, it should count down to 30, reset to 40, count down to 30, reset to 40, etc.

SHOW IP OSPF INTERFACE

```
R1# show ip ospf interface GigabitEthernet 0/0
```

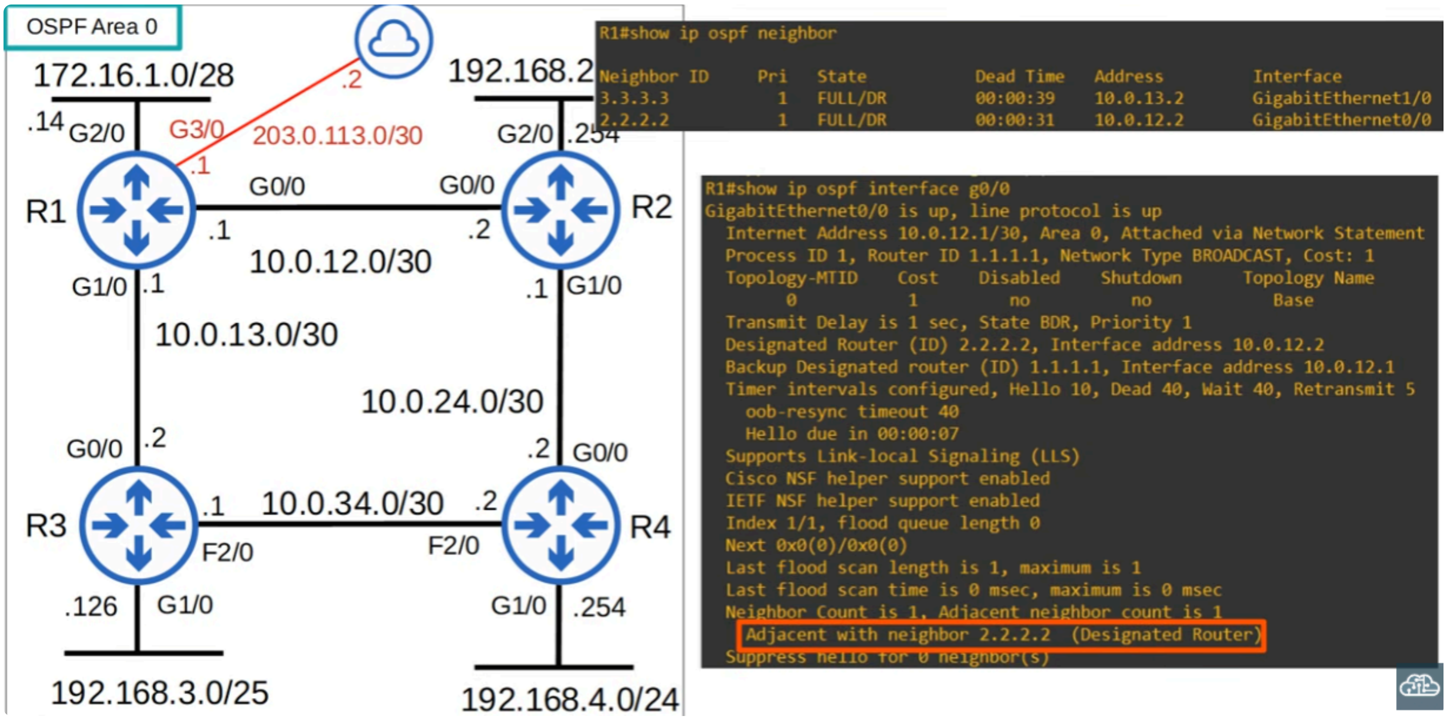

- Looking at R1's G0/0 interface
- Default Hello and Dead timers of 10 and 40 are displayed
- "Hello due in 7 seconds" means R1 will send a Hello message out of this interface in 7 seconds (sent once every 10 seconds)
- Neighbor count is 1, adjacent neighbor count is 1
- R1 has only 1 neighbor connected to its G0/0 interface, that's R2
- The difference between a neighbor and an adjacent neighbor will be covered in the next video
- Adjacent neighbor 2.2.2.2 is the designated router (covered in Day 28)

Additional OSPF Configurations

Enabling OSPF Directly on an Interface

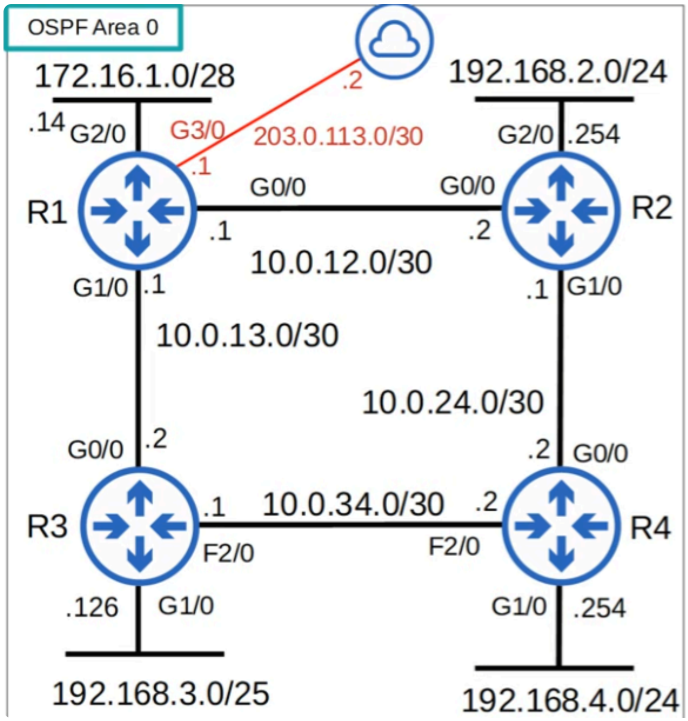


```
R1(config)#int g0/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int g1/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int g2/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int l0
R1(config-if)#ip ospf 1 area 0
```

- You can activate OSPF directly on an interface with this command:
  R1(config-if)#**ip ospf** *process-id* **area** *area*

- The **network** command tells the router which interfaces to activate the routing protocol on (same for RIP, EIGRP, and OSPF)
- You can enable OSPF directly on an interface without using the network command
- **Example:** R1 has no OSPF configurations on it yet

```
R1(config-if)# ip ospf <process-id> area <area-id>
```

- This is done from interface configuration mode
- OSPF is now enabled on those interfaces without entering OSPF configuration mode

Alternative Method to Configure Passive Interfaces

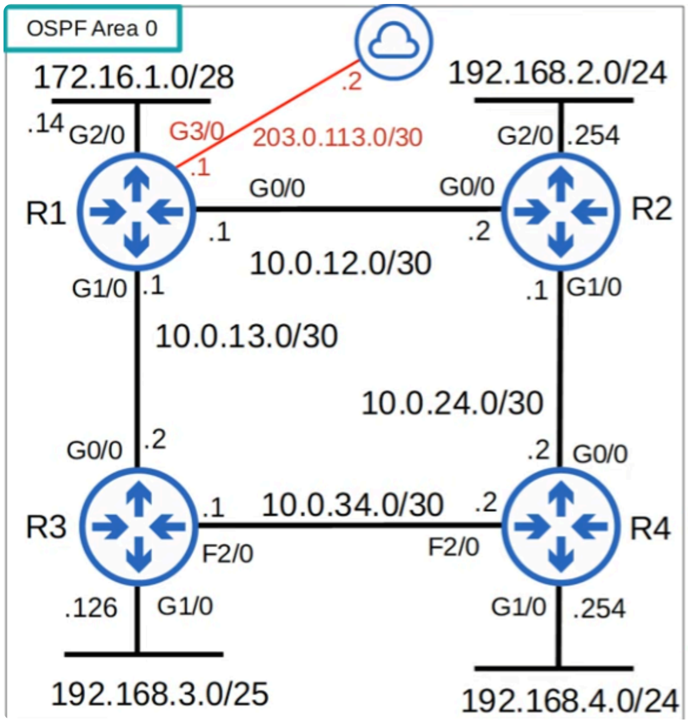


```
R1(config)#int g0/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int g1/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int g2/0
R1(config-if)#ip ospf 1 area 0
R1(config-if)#int l0
R1(config-if)#ip ospf 1 area 0
```

- You can activate OSPF directly on an interface with this command:
  R1(config-if)#**ip ospf** *process-id* **area** *area*

```
R1(config-if)#router ospf 1
R1(config-router)#passive-interface default
R1(config-router)#no passive-interface g0/0
R1(config-router)#no passive-interface g1/0
```

- Configure ALL interfaces as OSPF passive interfaces:
  R1(config-router)#**passive-interface default**
- Then configure specific interfaces as active:
  R1(config-router)#**no passive-interface** *int-id*

- You can configure all of the router's interfaces to be OSPF passive interfaces by default with the **passive-interface default** command
- Then, use **no passive-interface** to remove passive only from specific interfaces

```
R1(config-router)# passive-interface default
R1(config-router)# no passive-interface <interface>
```

- This is simply another way to configure passive interfaces
- Depending on the number of passive interfaces you need to configure, this method might be faster, or the normal method might be faster
- Either way, the effect is the same

Verification: SHOW IP PROTOCOLS Output Difference

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

Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
    GigabitEthernet1/0
    GigabitEthernet0/0
    GigabitEthernet2/0
  Passive Interface(s):
    Ethernet0/0
    GigabitEthernet2/0
    GigabitEthernet3/0
    Loopback0
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    2.2.2.2              110      00:09:53
    Gateway         Distance      Last Update
    3.3.3.3              110      00:09:54
    4.4.4.4              110      00:09:54
  Distance: (default is 110)
```

- If you configure OSPF directly on the interfaces, the output from `show ip protocols` will be slightly different
- The 'routing for networks' section is empty
- Instead, the interfaces you activated OSPF on are displayed under 'routing on interfaces configured explicitly'
- The rest of the output is the same

Review

OSPF Metric (Cost)

- Automatically calculated: Reference Bandwidth ÷ Interface Bandwidth
- If the result is less than 1, it is converted to 1
- Default reference bandwidth is 100 Mbps, so anything equal to or faster than 100 Mbps has an equal cost of 1
- Three ways to modify cost:
  - `auto-cost reference-bandwidth` command
  - `ip ospf cost` command
  - `bandwidth` command (not recommended)
- A route's metric is the total cost of the outgoing interfaces in the route

OSPF Neighbor States

- Summary diagram provided
- This is probably the most difficult section of the lecture
- Recommended to watch multiple times and do a Google search for "ospf neighbor states" to learn more

Additional OSPF Configurations

- Instead of using the `network` command, you can activate OSPF directly on an interface with this command:

  ```
  R1(config-if)# ip ospf <process-id> area <area-id>
  ```

- As an alternative method of configuring passive interfaces, you can configure all interfaces as passive with the `passive-interface default` command, then make specific interfaces active with `no passive-interface`