OSPF (Open Shortest Path First) Basics for CCNA

OSPF Introduction and Course Overview

- OSPF (Open Shortest Path First) is the only dynamic routing protocol explicitly listed in the CCNA exam topics.
- OSPF is CCNA exam topic 3.4, "Configure and verify single area OSPFv2."
- Sub-topics for OSPF on the CCNA include:
  - Neighbor adjacencies
  - Point-to-point
  - Broadcast
  - Router ID
- This video provides a high-level introduction to OSPF concepts and basic configuration.
- Later lectures will go into the deeper details of how OSPF functions.

Lecture Topics

- Basic operations of OSPF (quick introduction).
- OSPF areas, a unique feature that splits larger networks into smaller sections.
- Basic OSPF configurations.

OSPF Protocol Type

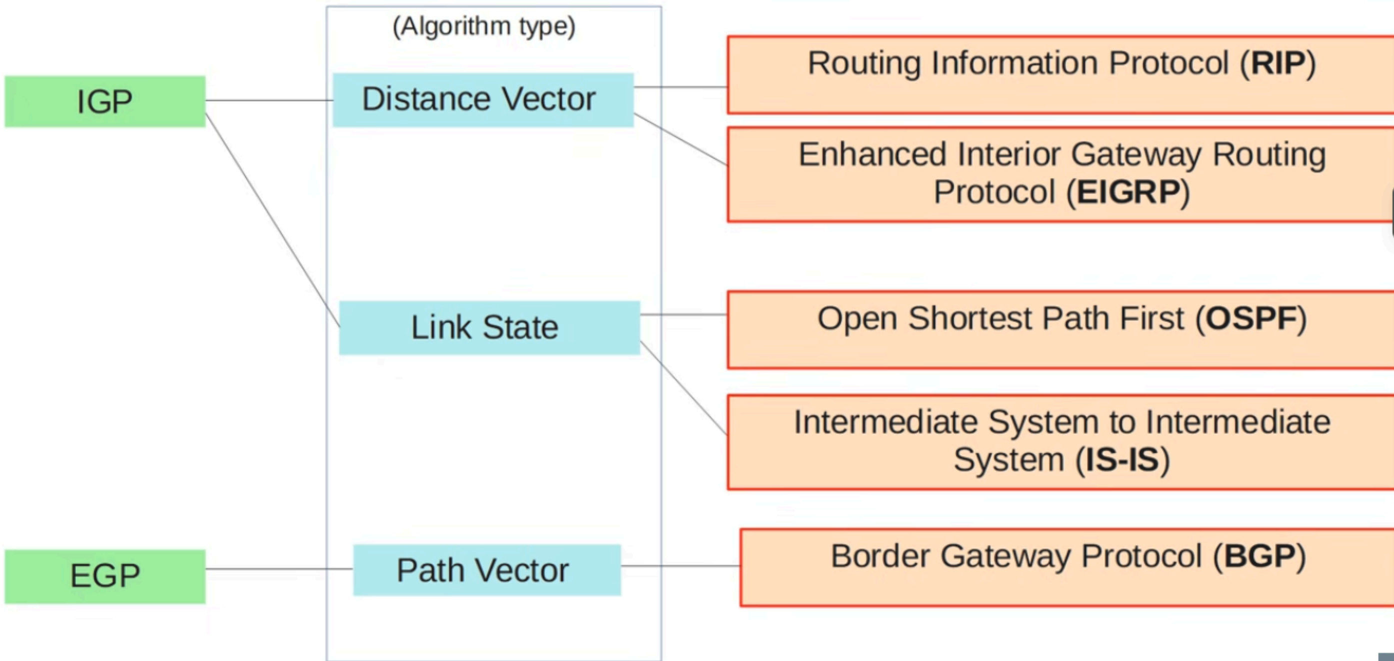

- OSPF is a Link State dynamic routing protocol.
- It functions differently than the distance vector routing protocols RIP and EIGRP.

Distance Vector Protocol Review

- Distance vector protocols use "routing by rumor."
- Each router shares information about the routes they know and their metric cost to reach each destination.
- Routers do not have a complete map of the network.
- They use the information their neighboring routers tell them to determine the best route to each destination.

Link State Protocol Review

- When using a link state routing protocol, every router creates a "connectivity map" of the network.

- Each router advertises information about its interfaces (its connected networks) to its neighbors.
- These advertisements are passed along to other routers, until all routers in the network develop the same map of the network.
- **Important:** All routers have the same complete map of the network.
- Each router independently uses this map to calculate the best routes to each destination.
- Link state protocols use more resources on the router because more information is shared; it's more demanding on the router.
- Link state protocols tend to be faster in reacting to changes in the network than distance vector protocols.

OSPF Overview

- OSPF stands for 'Open Shortest Path First'.
- OSPF uses the 'shortest path first' algorithm, created by Dutch computer scientist Edsger Dijkstra.
- Another name for the algorithm is 'Dijkstra's algorithm' — this could be an exam question.

OSPF Versions

- There are three versions of OSPF:
  - **Version 1:** Released in 1989. It is old and not in use anymore.
  - **Version 2:** Released in 1989. Typically used in IPv4 networks. This is the version listed in the CCNA exam topics.
  - **Version 3:** Developed for IPv6. It can also be used for IPv4, but version 2 is more common in IPv4 networks.

OSPF Key Terms

- Routers store information about the network in **LSAs (Link State Advertisements)**, which are organized in a structure called the **LSDB (Link State Database)**.

- Routers will **flood** LSAs until all routers in the OSPF **area** develop the same map of the network, meaning the same LSDB.
- In OSPF, "flood" means the router sends the LSAs to all of its OSPF neighbors.
- OSPF areas are a unique aspect of OSPF.

LSAs and LSDB Process Example

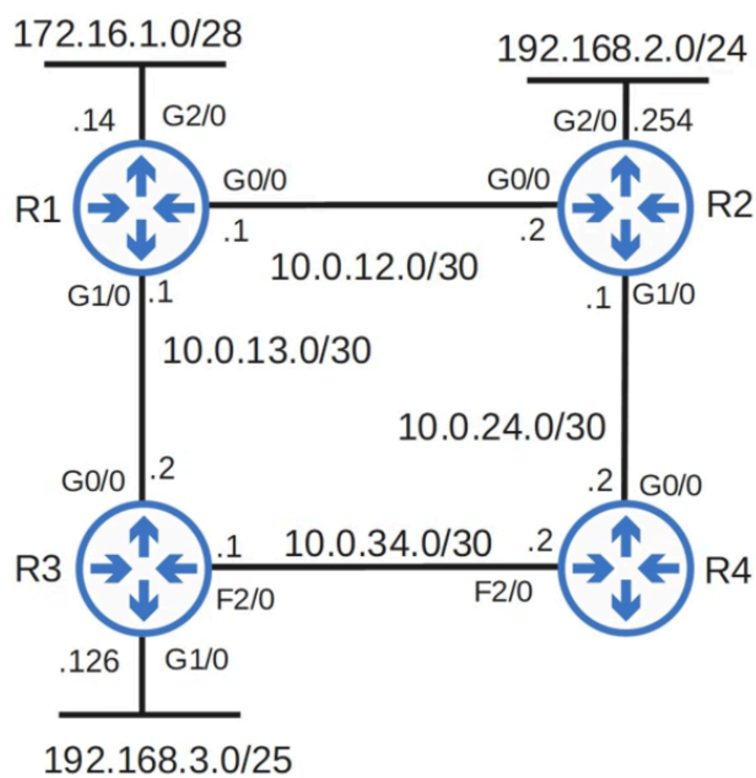

- Consider a network of four routers running OSPF, all OSPF neighbors with the same LSDB.
- When OSPF is enabled on a new interface (e.g., R4's G1/0), R4 creates an LSA to inform neighbors about the new network segment.

- Basic information included in an LSA:

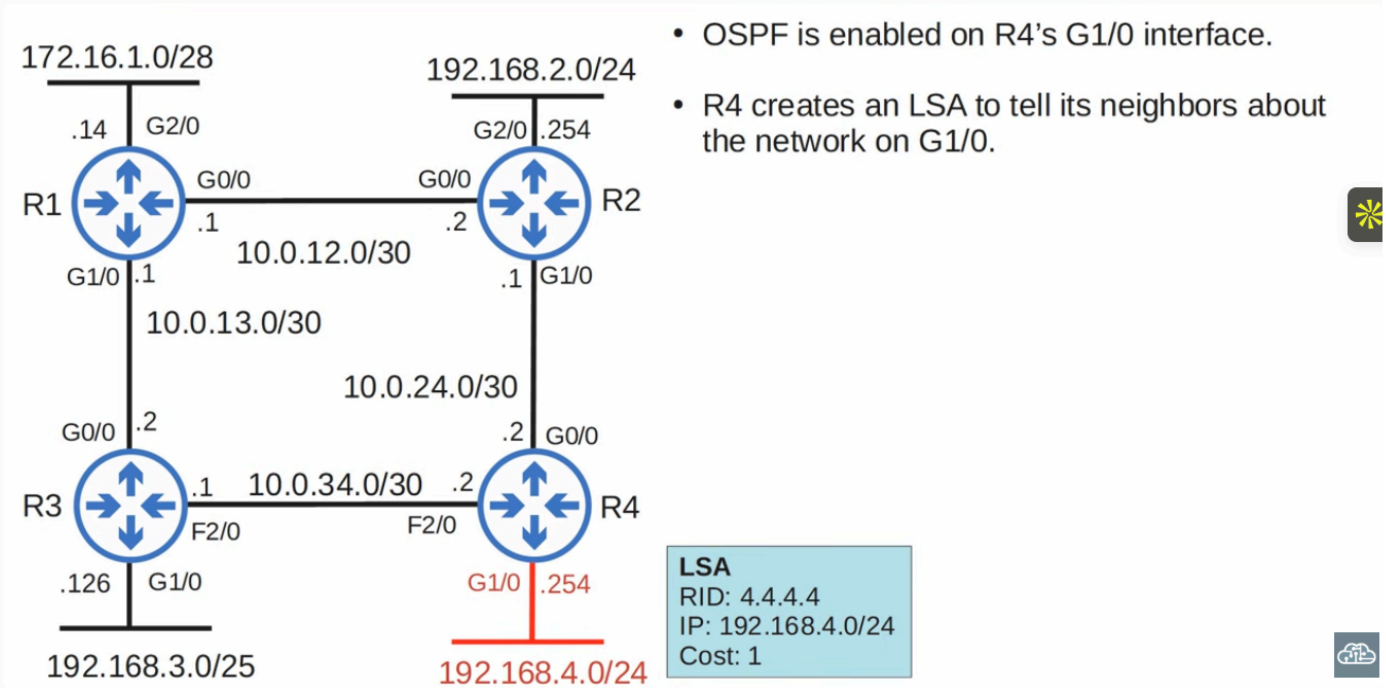

- **RID (Router ID):** e.g., 4.4.4.4.
- **Network:** The network on the advertised interface (e.g., G1/0).
- **Cost:** OSPF's metric (e.g., a GigabitEthernet interface has a cost of 1).

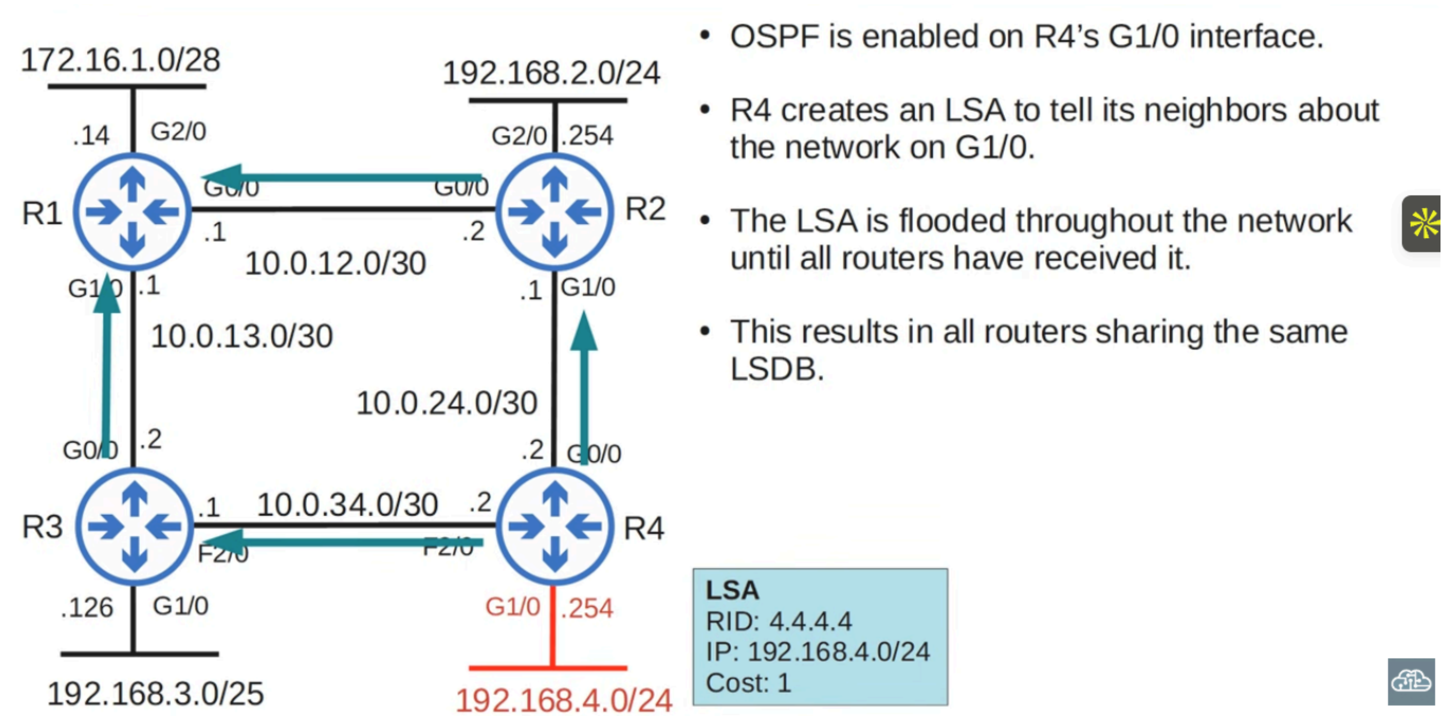

- The LSA is flooded throughout the network until all routers receive a copy.

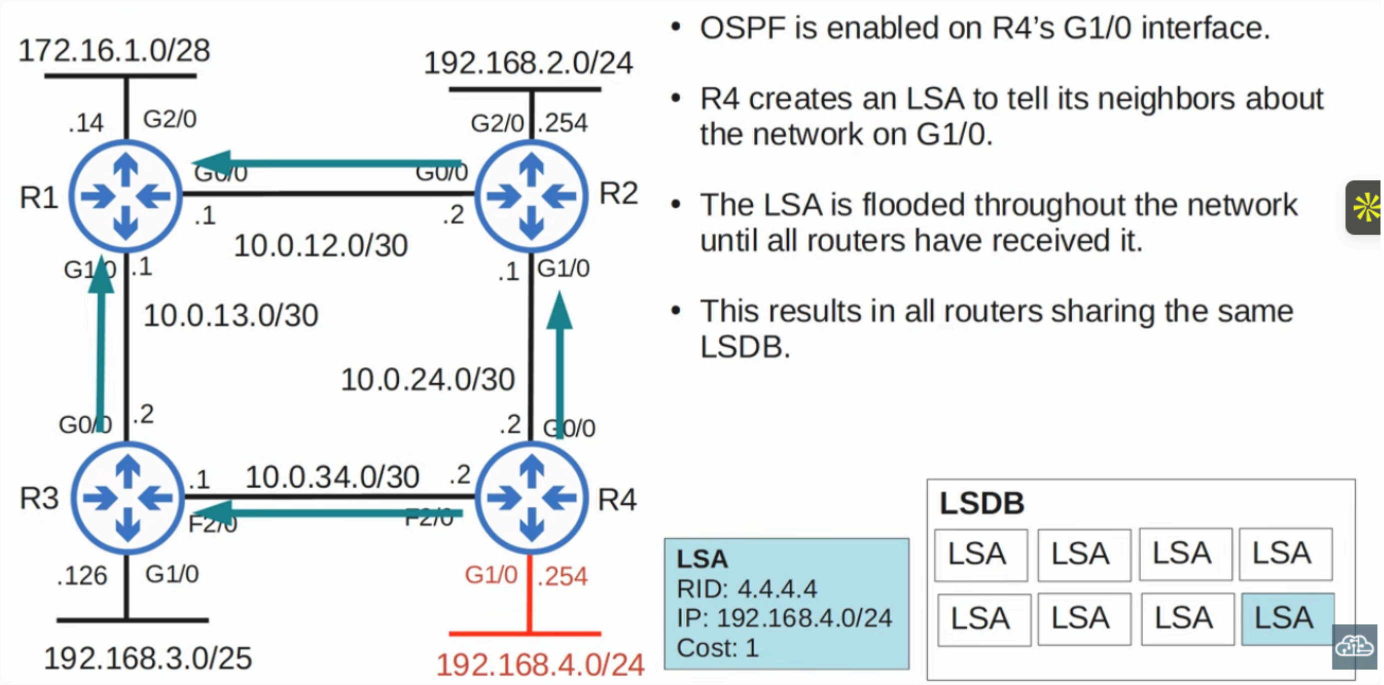

- This results in all routers in the OSPF area having the same LSDB, containing LSAs for all of the different links in the network.
- Now that OSPF has been activated on R4's G1/0 interface, that new LSA is added to the LSDB.
- The LSDB is identical for all routers in the OSPF area.

- Each router then uses the SPF algorithm (Dijkstra's algorithm) to calculate its best route to 192.168.4.0/24.

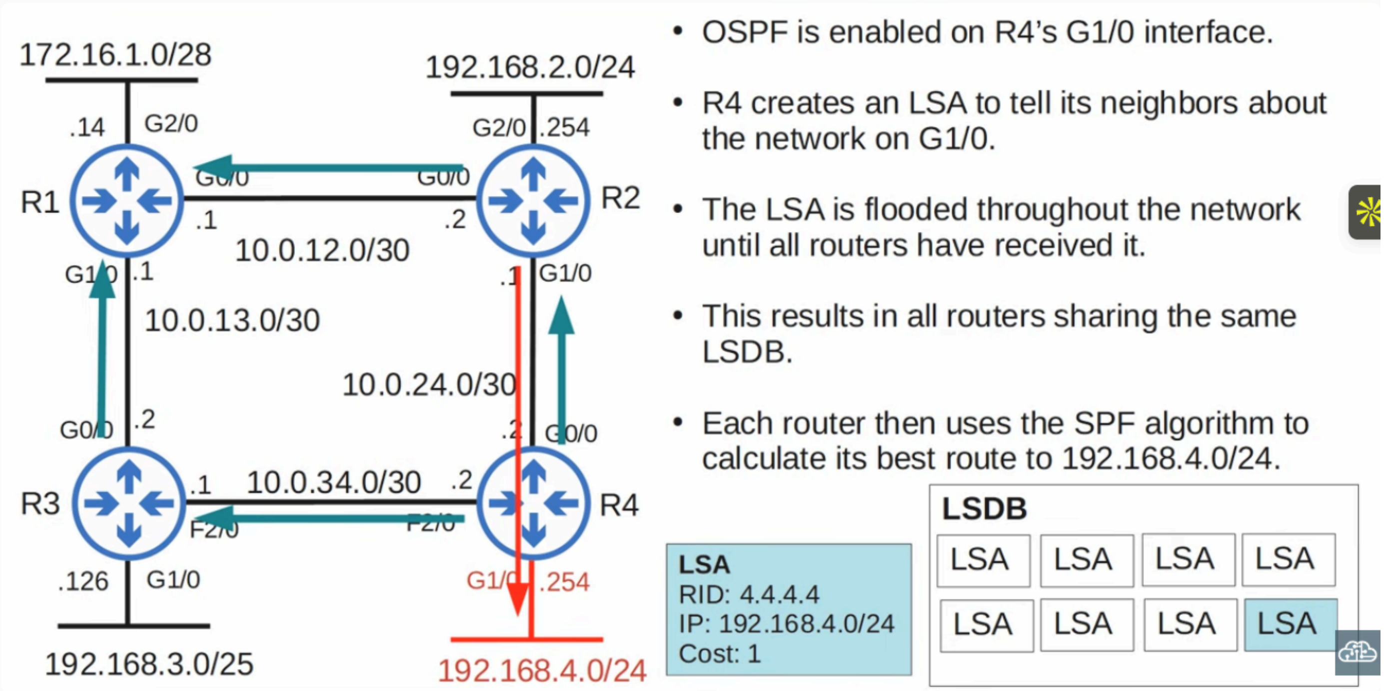

- Each router has a complete map of the network. For example, R2 can calculate that sending traffic via G1/0 is the best route to 192.168.4.0/24, just like looking at the network diagram.
- Each individual LSA has an aging timer, which is 30 minutes by default. The LSA will be flooded again after the timer expires, once every 30 minutes by default.

OSPF Process Summary

- Step 1: Become neighbors with other routers connected to the same segment.
  - *Example:* In the previous network, R4 was OSPF neighbors with R2 and R3.
- Step 2: Exchange LSAs with neighbor routers.
- Step 3: Each router independently calculates its best routes to each destination and inserts them into the routing table.

OSPF Areas

- OSPF uses areas to divide up the network. This is a fundamental concept of OSPF that isn't part of RIP or EIGRP.
- Small networks can be single-area without any negative effects on network performance.
    - *Example:* A network with four routers is a small network where a single OSPF area will not degrade performance.
- In larger networks, a single-area design can have negative effects.
    - *Example:* A network with 500 routers and over 1000 subnets should not use a single OSPF area.
- **Negative effects of a single-area design in a large network:**
    - The SPF algorithm takes more time to calculate routes.
    - It requires exponentially more processing power on each router.
    - Sharing a single, huge link state database takes up more memory on the routers.
    - Every small change on the network (e.g., a new interface being activated) would cause LSAs to be flooded to all 500 routers, and all of those routers would have to do the SPF calculation again.
- Dividing a large OSPF network into several smaller areas can avoid these negative effects.
- The CCNA exam topics only mention single area OSPF, so only a brief overview of OSPF areas is needed.

Backbone Area

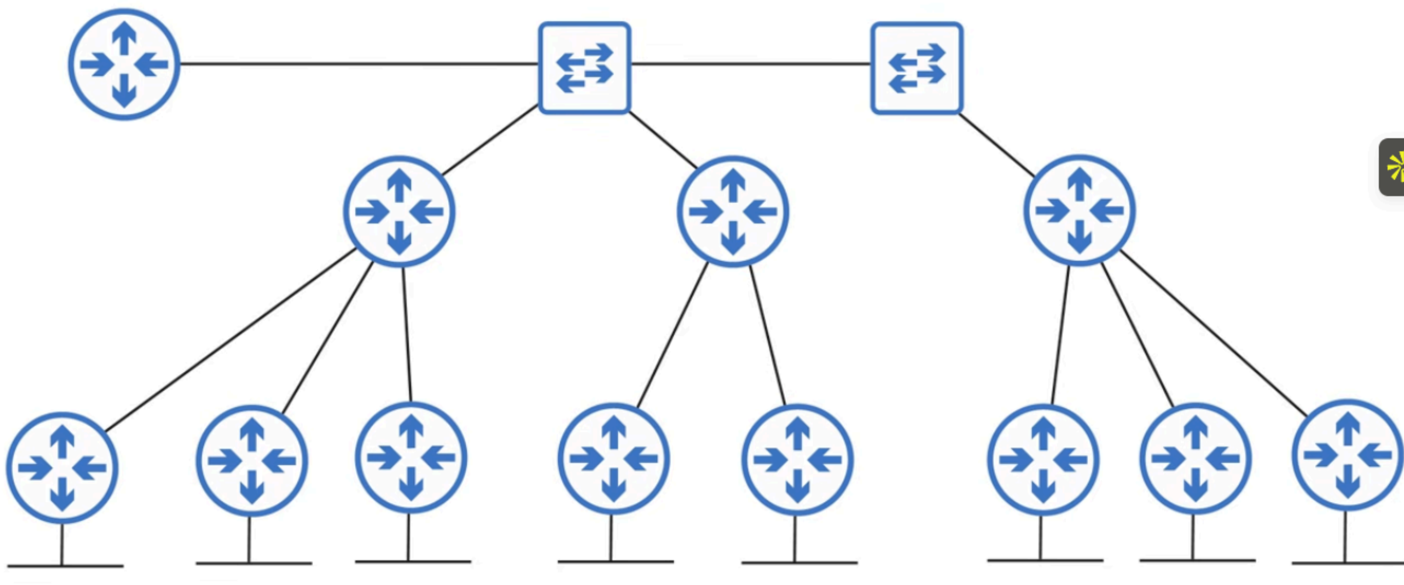

- A larger network can be made into one large single-area network where all interfaces on all routers are assigned to area 0, also known as the **backbone area**.

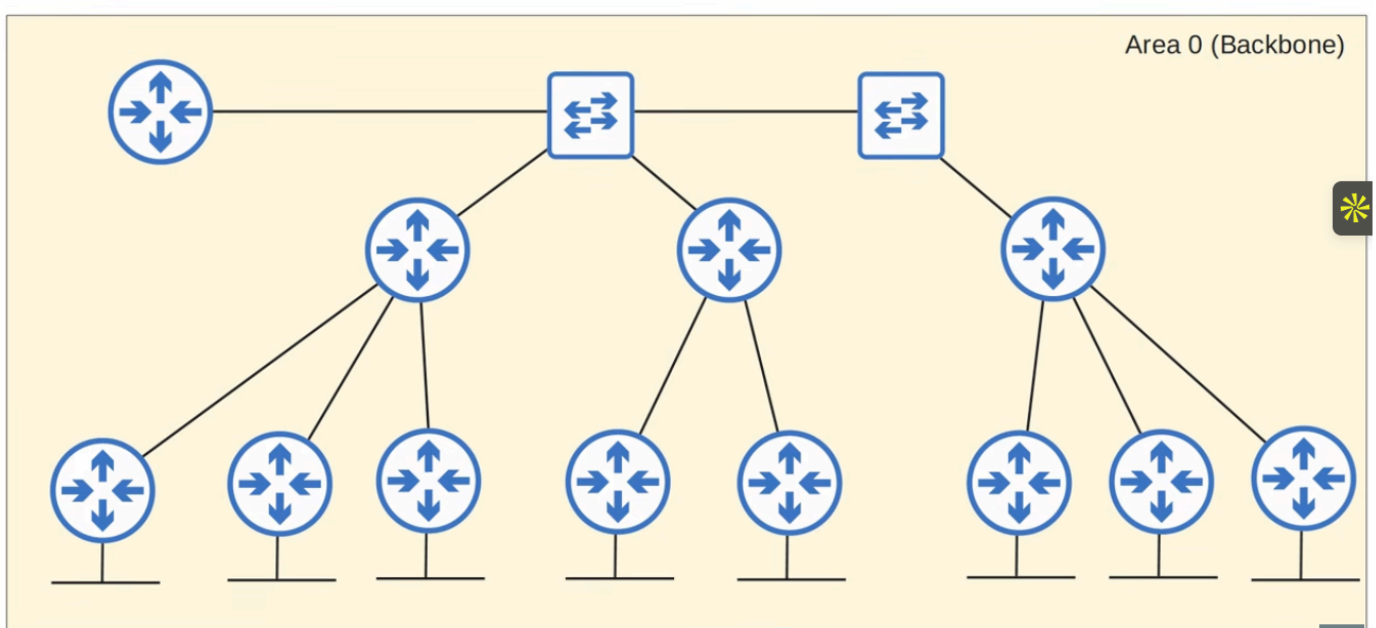

- Area 0 is of special importance in OSPF.

- An **area** is a set of routers and links that share the same LSDB (Link State Database).

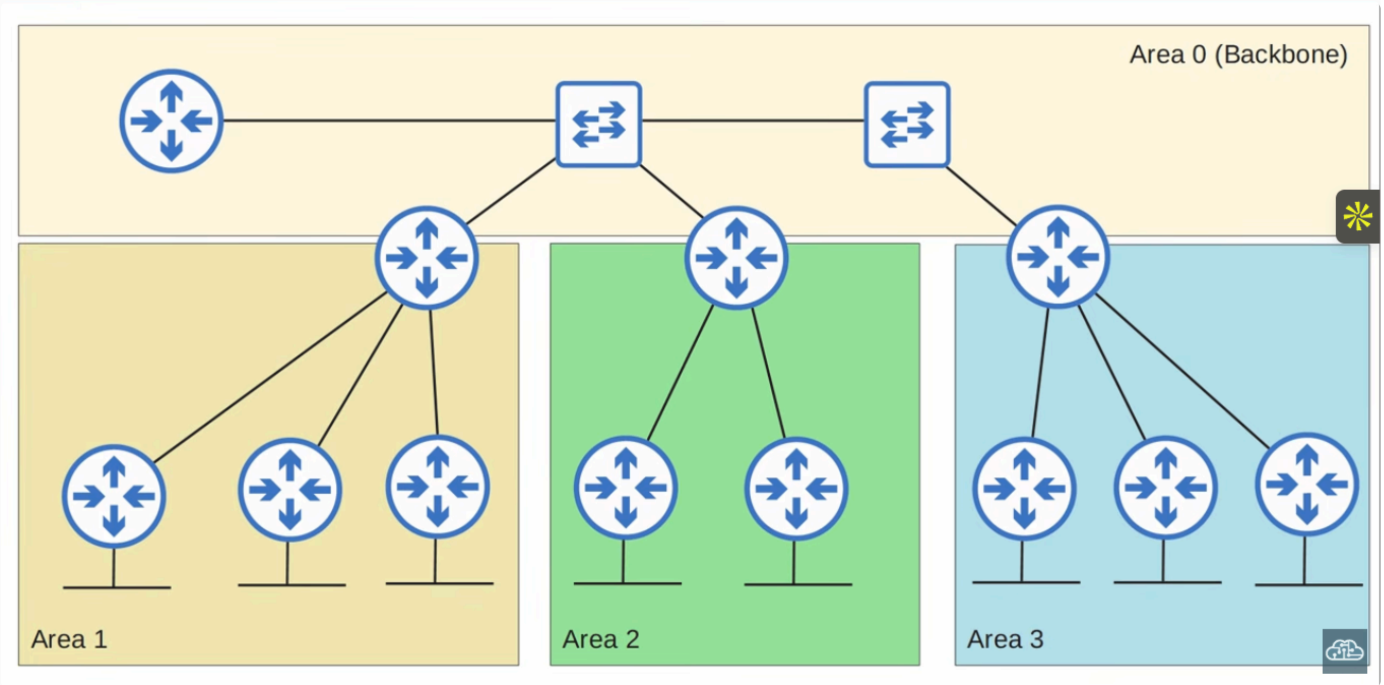

  - *Example:* In a divided network diagram, there are four areas: Area 0, Area 1, Area 2, and Area 3.
- Each area maintains a unique LSDB.

Backbone Area Rule

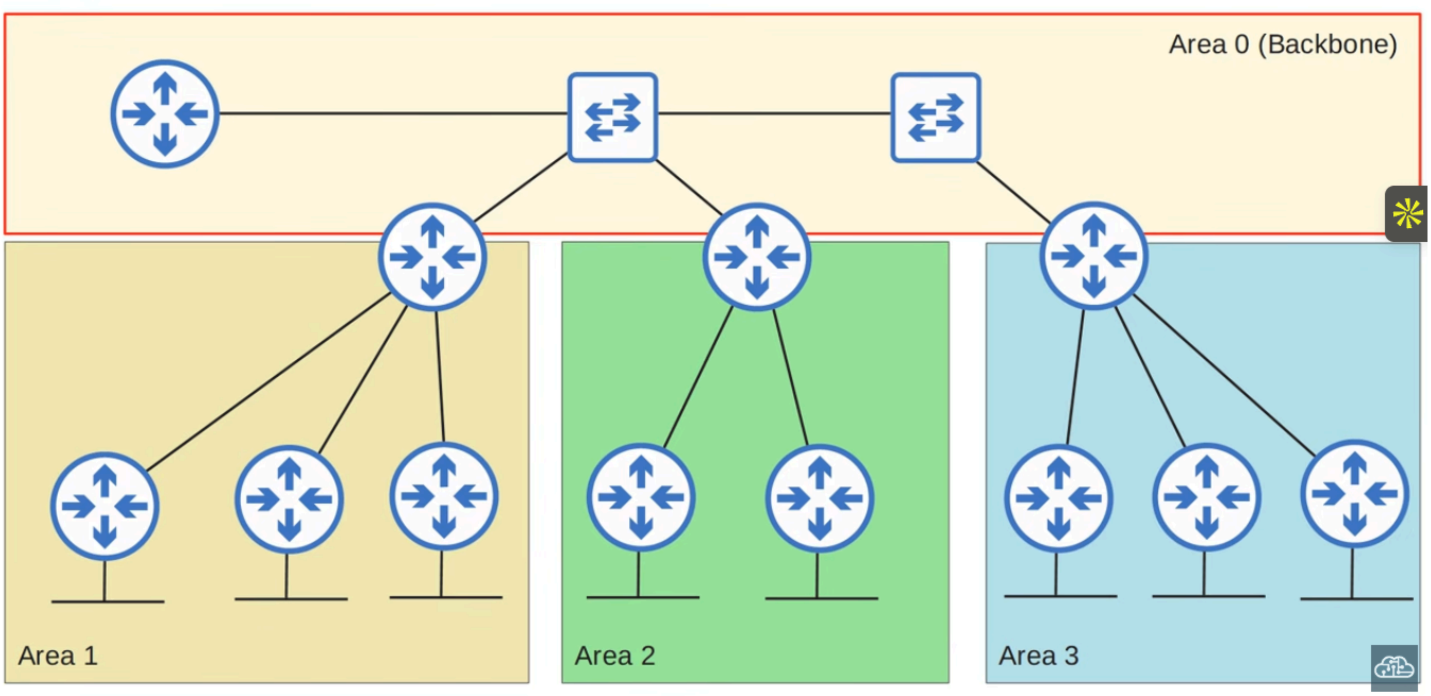

- The backbone area (area 0) is a special area that all other areas must connect to.
- *Example:* In the network diagram, area 1, area 2, and area 3 all connect to area 0, the backbone area.

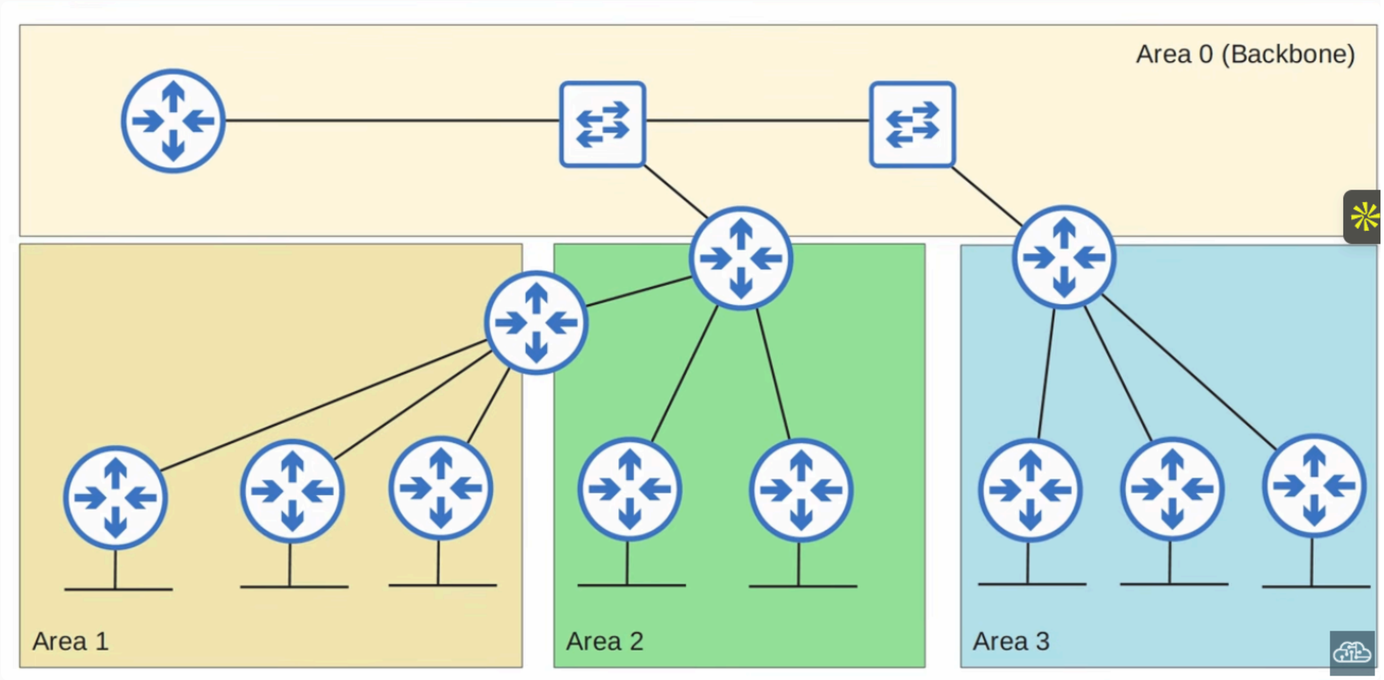

- A design where area 1 is not connected to area 0 (the backbone area) and is only connected to area 2 is not allowed in OSPF.

Internal Routers

- Routers with all interfaces in the same area are called **internal routers**.

- *Example:* In the diagram:

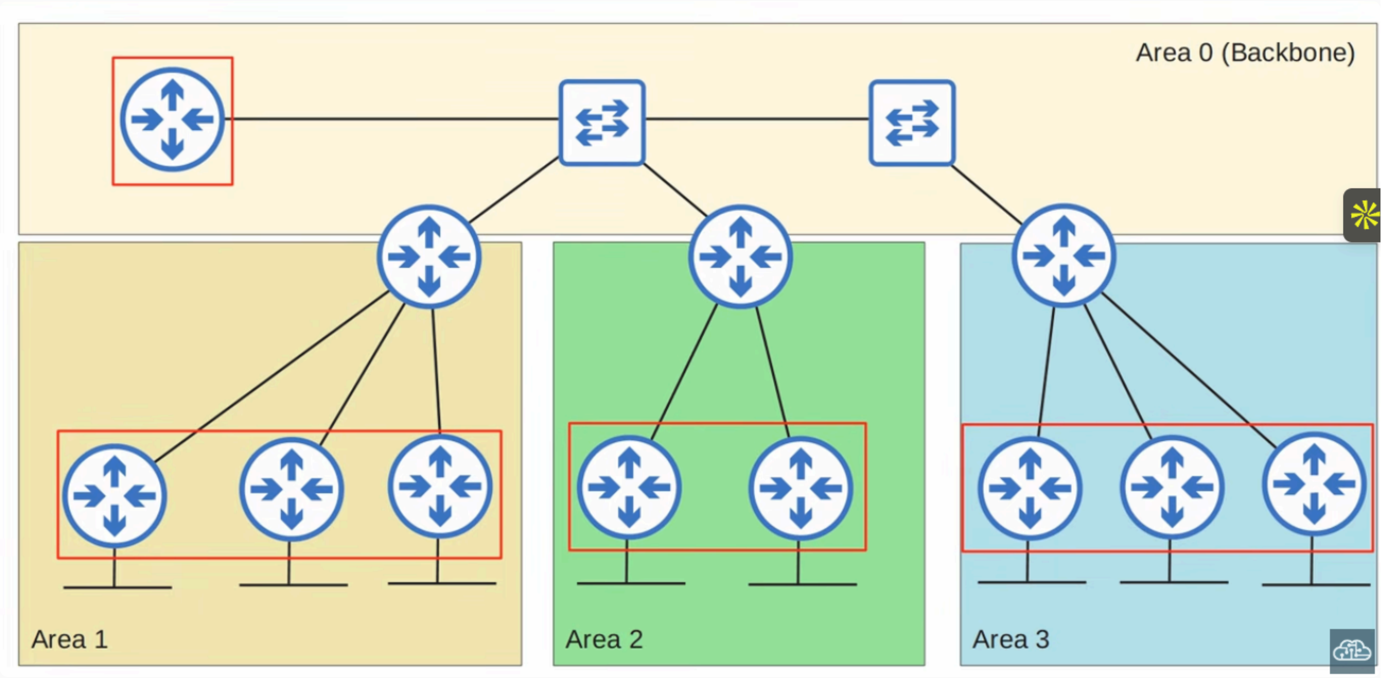

  - A router with all interfaces in area 0 is internal to area 0.
  - Routers with all interfaces in area 1 are internal routers in area 1.
  - Same for routers in area 2 and routers in area 3.
- Internal routers have all of their interfaces in the same OSPF area.

Area Border Routers (ABRs)

- Routers with interfaces in multiple areas are called **area border routers** (ABRs), because they are the border between different OSPF areas.

- *Example:* In the diagram:

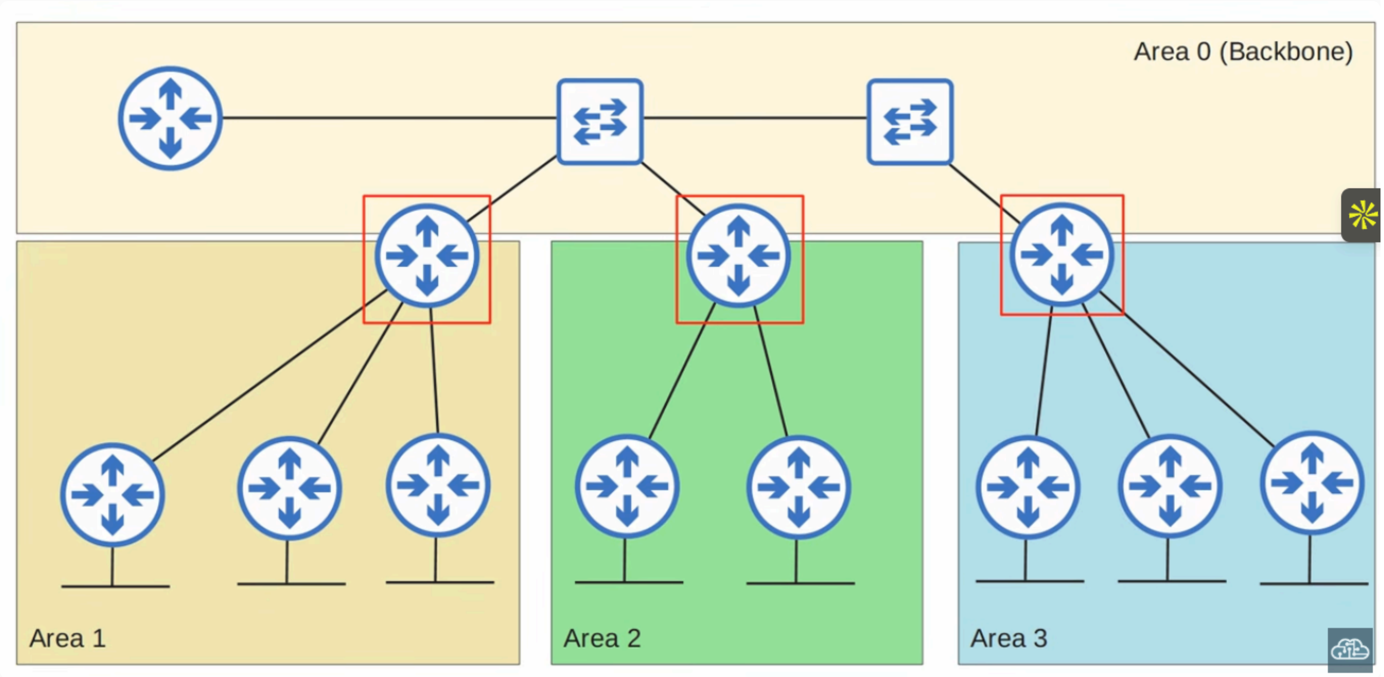

  - The router connected to area 0 and area 1 is an ABR.
  - The router connected to area 0 and area 2 is also an ABR.
  - The router connected to area 0 and area 3 is an ABR.
- ABRs maintain a separate LSDB for each area they are connected to.
- It is recommended that you connect an ABR to a maximum of 2 areas. Connecting an ABR to 3+ areas can overburden the router.
  - *Example:* A design with each ABR only connected to 2 areas is good OSPF network design.

Backbone Routers

- Routers connected to the backbone area (area 0) are called **backbone routers**.
  - *Note:* This includes area border routers (ABRs).

- *Example:* In the diagram:

  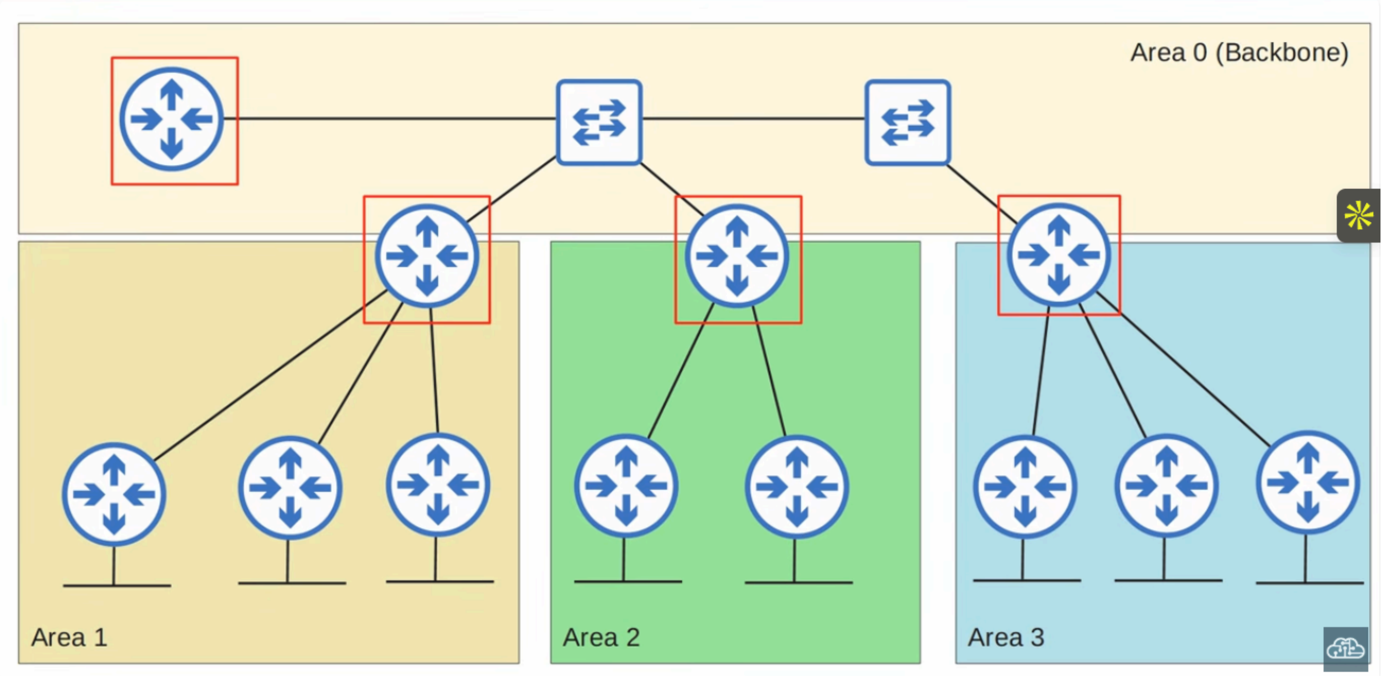

  - A router connected only to area 0 is a backbone router. It is also an internal router internal to area 0.
  - Routers connected to area 0 and another area are backbone routers on top of being ABRs.

Intra-area and Interarea Routes

- An **intra-area route** is a route to a destination inside the same OSPF area.

  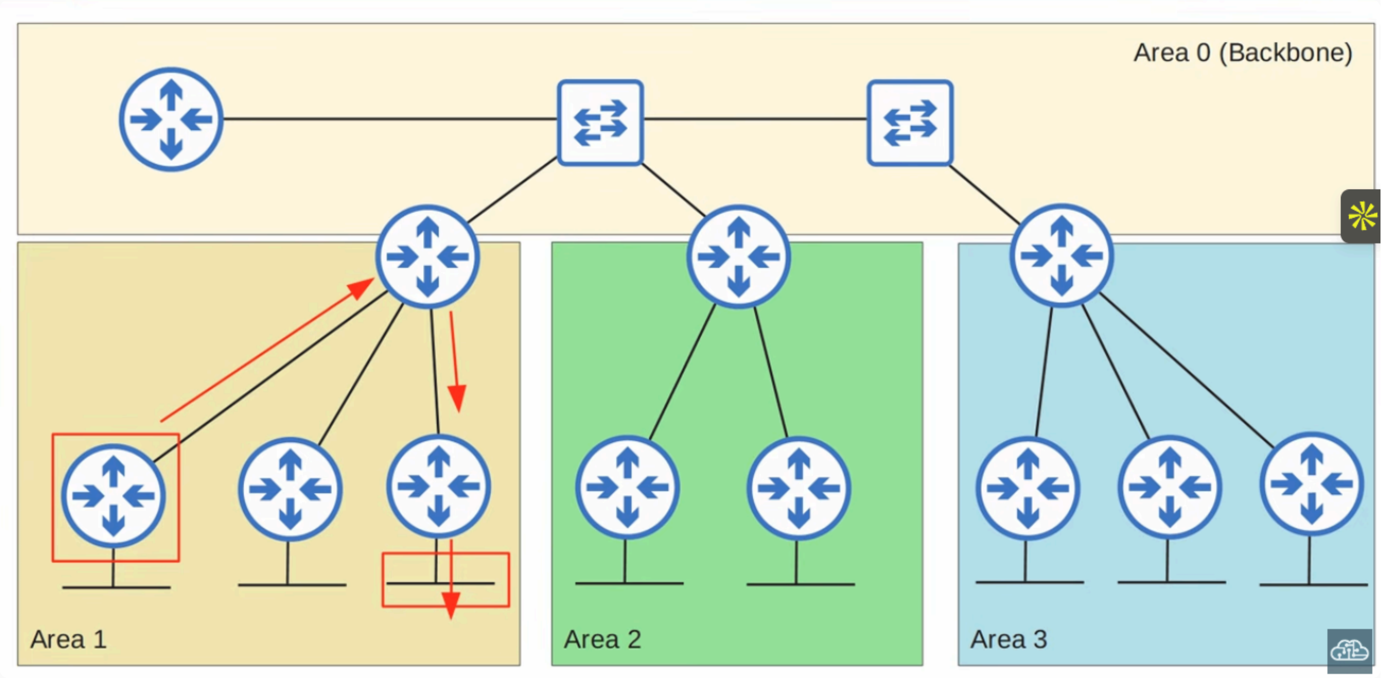

  - *Example:* If a router in area 1 learns a route to a subnet that is also in area 1, it is an intra-area route.

- An **interarea route** is a route to a destination in a different OSPF area.

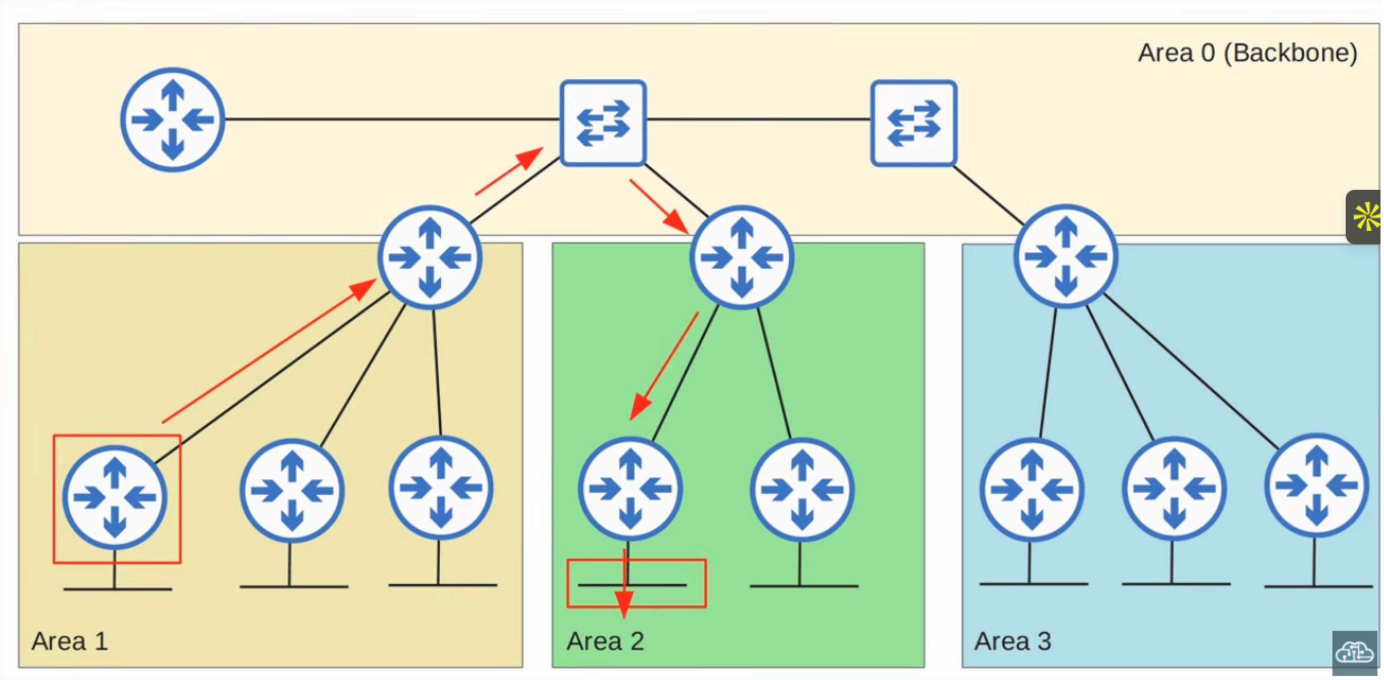

  - *Example:* If a router in area 1 learns a route to a subnet in area 2, it is an interarea route. The router and the destination are in two different OSPF areas.

Key OSPF area terms to remember

- **Area:** A set of routers and links that share the same LSDB (Link State Database).
- **Backbone area:** Area 0, which is a special area that all other areas must connect to.
- **Internal router:** A router with all of its interfaces in the same OSPF area.
- **Area border router (ABR):** A router with interfaces in multiple OSPF areas. It serves as the border between different areas.
- **Backbone router:** A router connected to the backbone area (area 0). This category includes both internal routers of area 0 and ABRs connected to area 0.
- **Intra-area route:** A route to a destination inside the same OSPF area as the router.
- **Interarea route:** A route to a destination in a different OSPF area than the router.

OSPF Area Rules

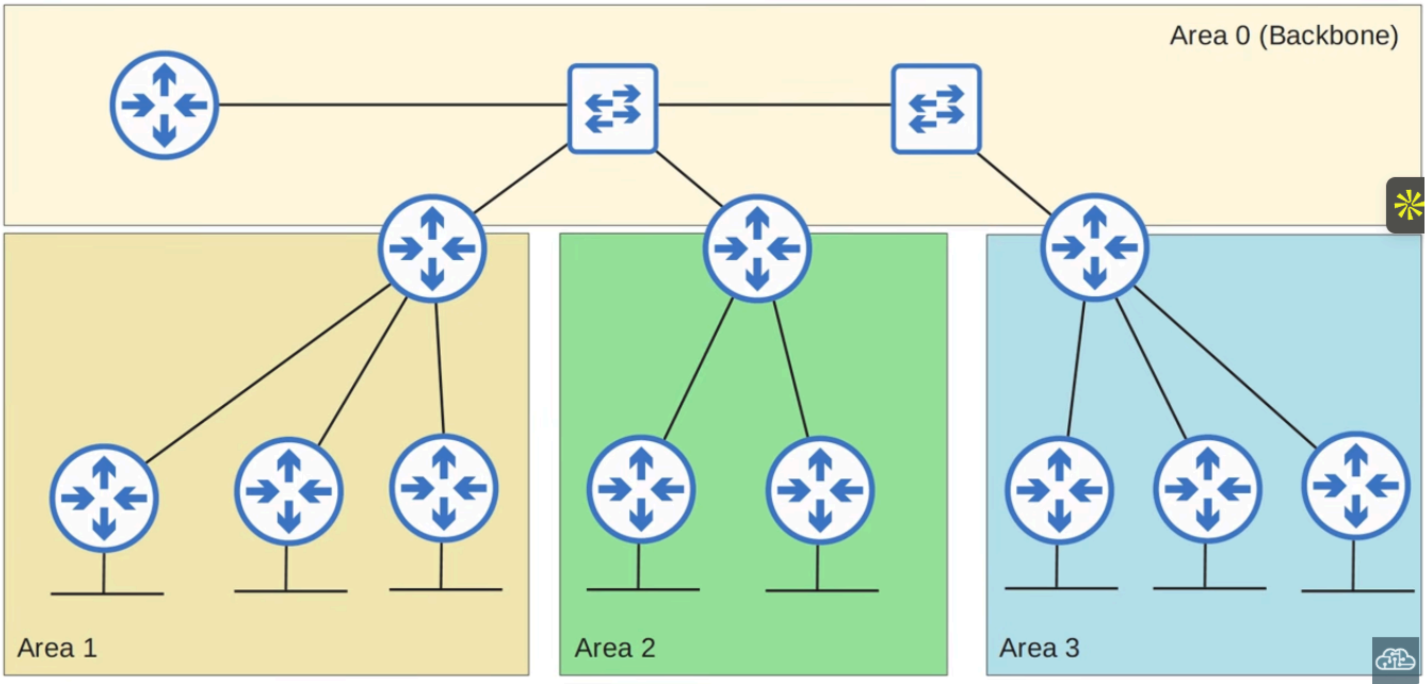

- OSPF areas should be ‘contiguous’, meaning each individual area should be connected and not divided up.
    - *Example:* A network diagram where all areas are connected satisfies this rule.

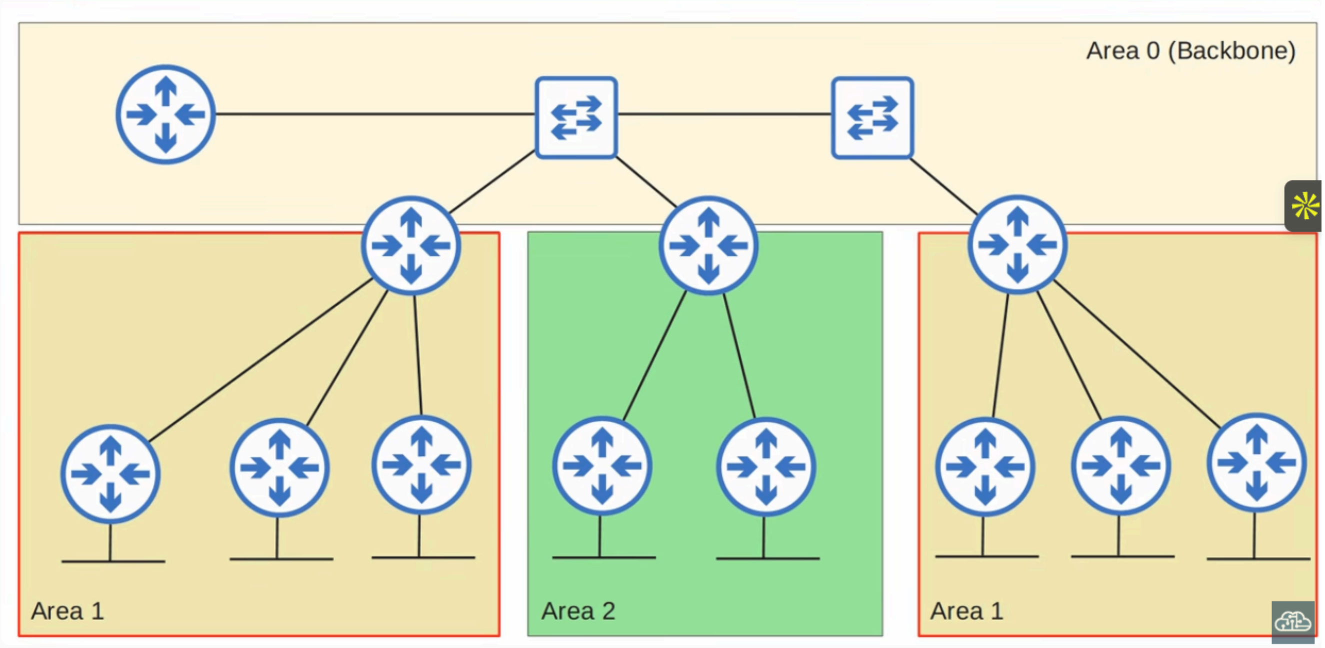

    - *Example:* If area 1 is split up—half here and half there—it is non-contiguous. This kind of network design is not allowed in OSPF and will cause problems.
    - *Note:* To fix a non-contiguous area, make the separate section a distinct area (e.g., turn the separated right section into area 3) so all areas are contiguous.

- All OSPF areas must have at least one ABR connected to the backbone area.

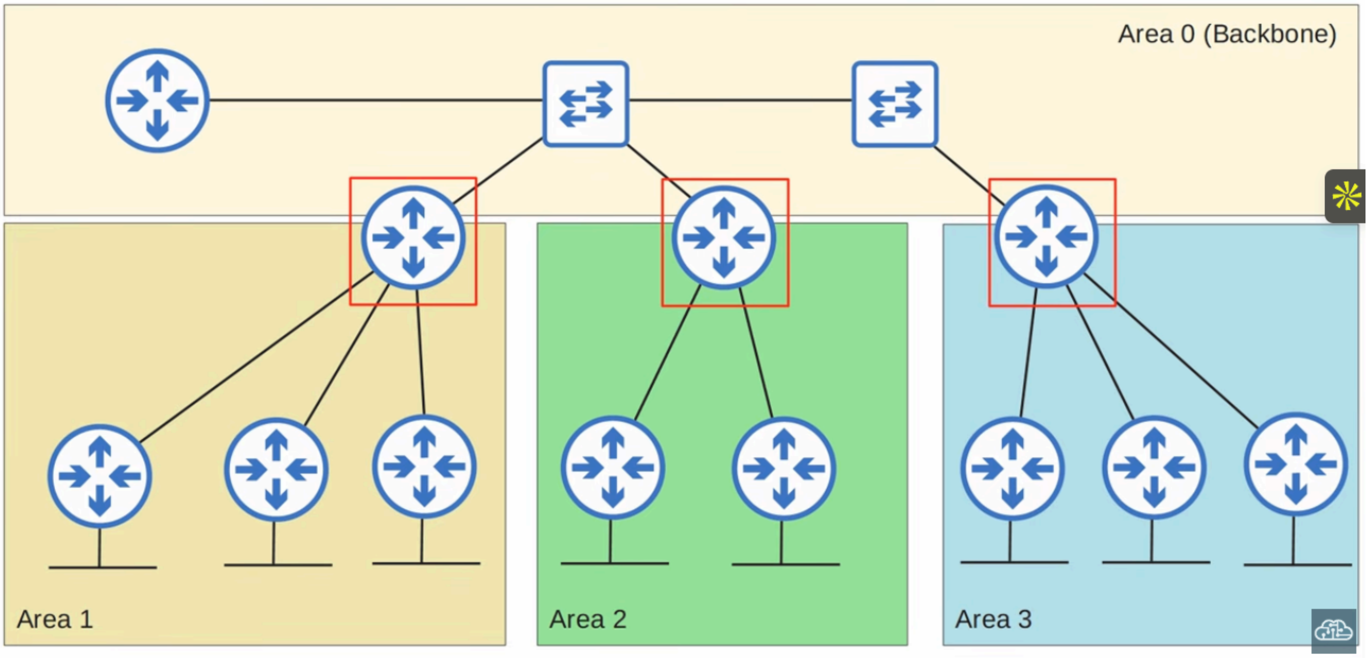

  - *Example:* In the network diagram, area 1, area 2, and area 3 each have an ABR connected to both their respective area and area 0. This is correct OSPF network design.

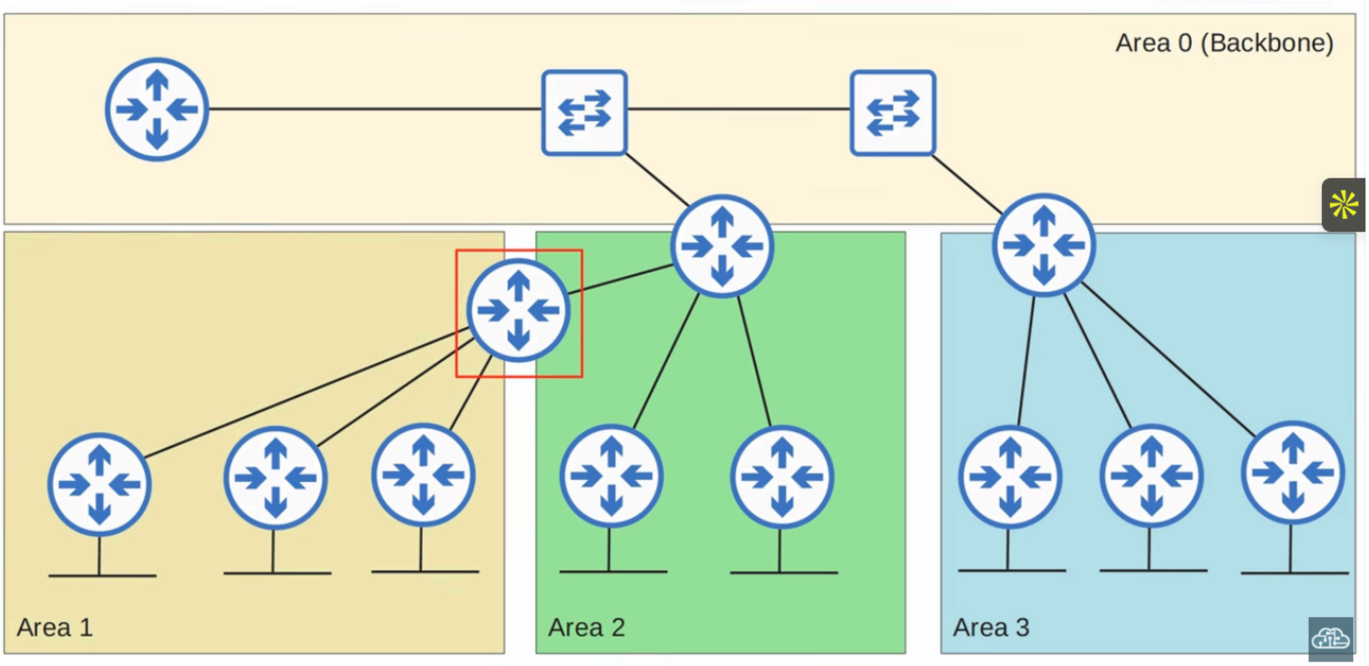

  - *Example:* A network where area 1 is not connected to area 0 is not correct OSPF design and will cause problems because area 1 lacks an ABR connected to the backbone area.

- OSPF interfaces in the same subnet must be in the same area. If they are not in the same area, they will not become OSPF neighbors and will not exchange information.

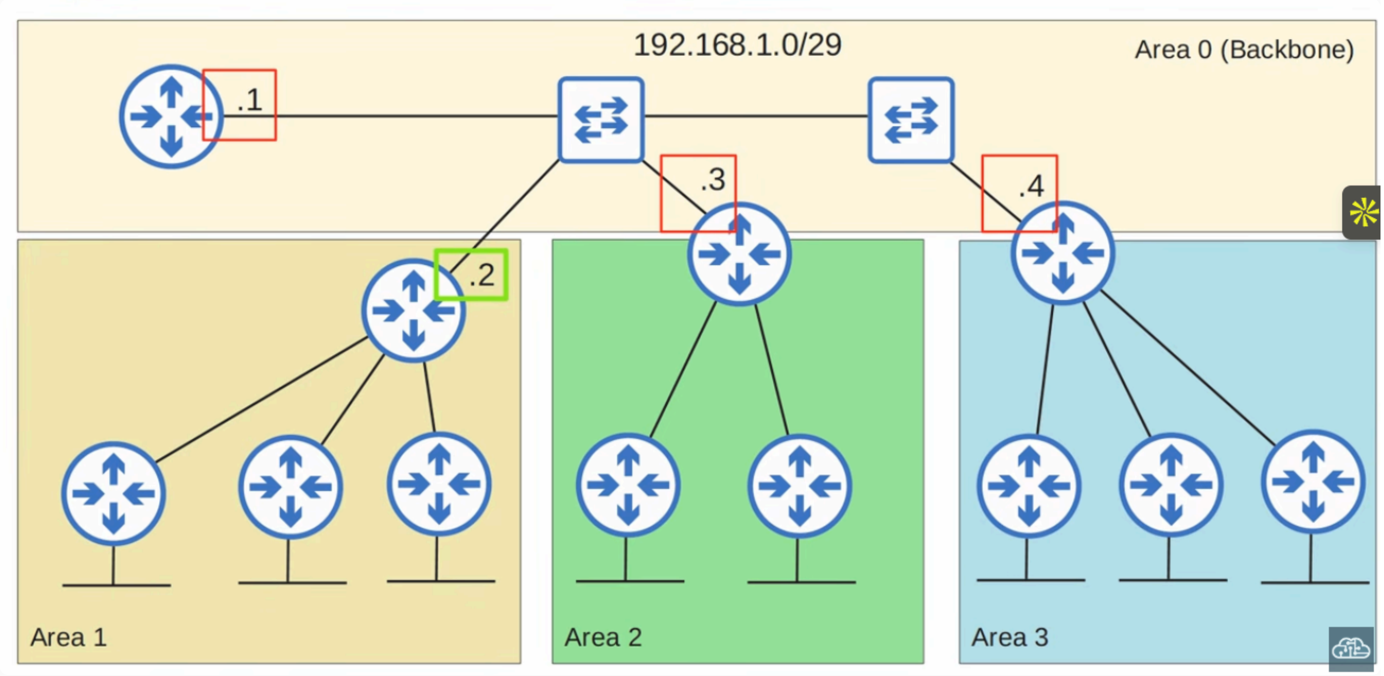

  - *Example:* Three routers have interfaces in area 0 in the 192.168.1.0/29 subnet. A fourth router has an interface in the same 192.168.1.0/29 subnet but in area 1. Even though all interfaces are in the same subnet with OSPF enabled, the area 1 router will not become OSPF neighbors with the area 0 routers.

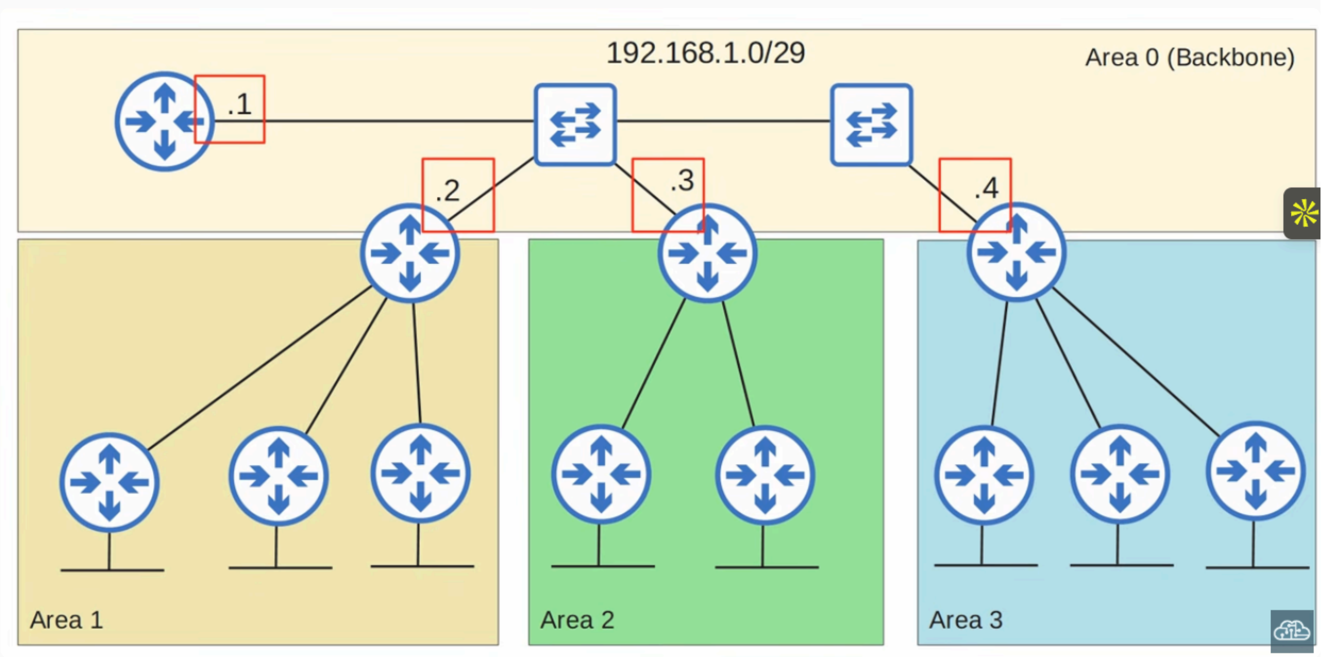

- *Example:* If the area 1 ABR's interface in the 192.168.1.0/29 subnet is correctly configured in area 0, all four routers will become OSPF neighbors.

OSPF Configuration

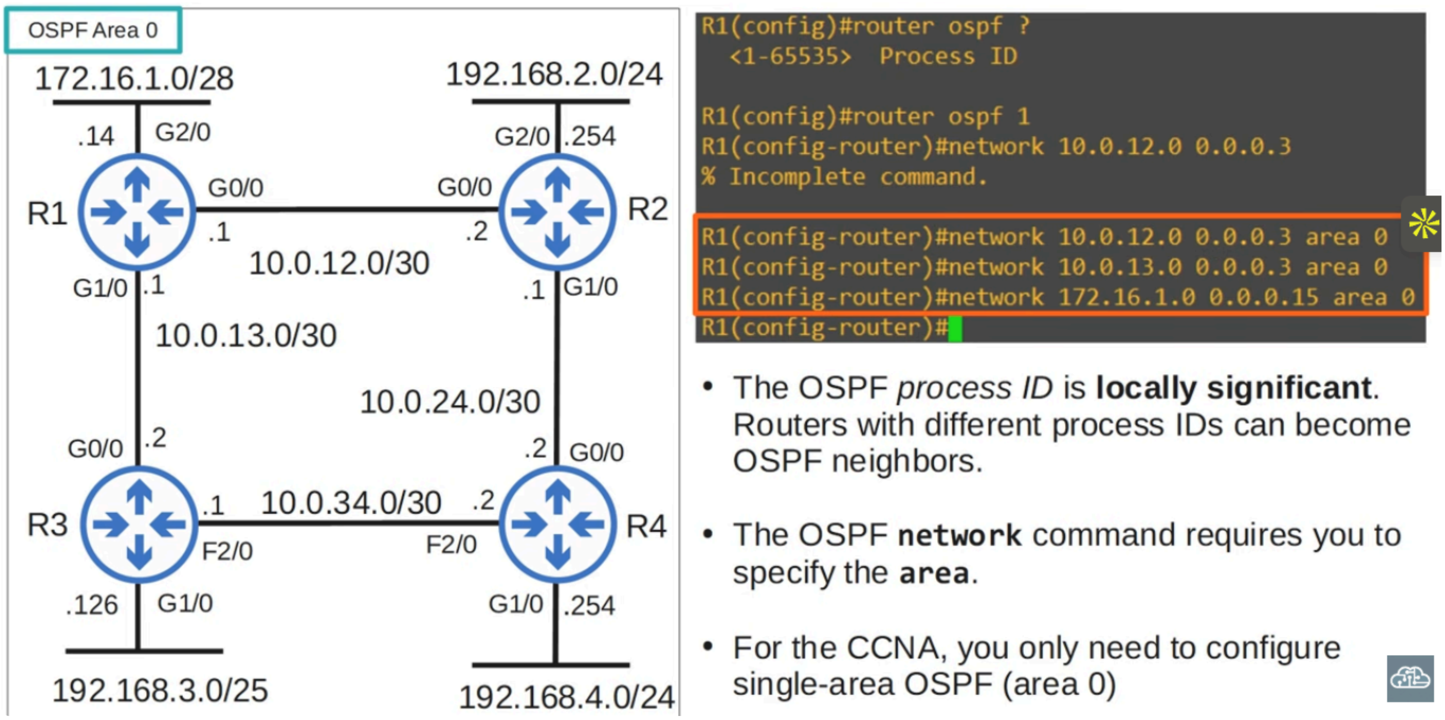

- For the CCNA, you only need to configure single-area OSPF. Best practice is to use area 0.
- To enter OSPF configuration mode, use the `ROUTER OSPF` command followed by a process ID.
- The OSPF process ID is locally significant. Routers with different process IDs can become OSPF neighbors (unlike EIGRP AS numbers which must match).
- The OSPF network command requires a wildcard mask and an area number.
- The `NETWORK` command tells OSPF to look for interfaces with an IP address contained in the specified range, and then activate OSPF on the interface in the specified area.
    - *Example:* The network command `10.0.12.0` with a `/28` wildcard mask, in area 0, will activate OSPF on R1's G0/0 interface (IP 10.0.12.1) because that IP is contained in the range.

- When OSPF is activated on an interface, the router will try to become OSPF neighbors with other OSPF-activated neighbor routers on that segment.
    - *Example:* R1 will become OSPF neighbors with R2 and R3.

```
R1(config)# router ospf 1
R1(config-router)# network 10.0.12.0 0.0.0.15 area 0
R1(config-router)# network 10.0.13.0 0.0.0.15 area 0
R1(config-router)# network 10.0.14.0 0.0.0.15 area 0
R1(config-router)# network 172.16.1.0 0.0.0.15 area 0
```

- *Note:* The `NETWORK` command simply tells the router which interfaces to activate OSPF on; it does not tell the router to "advertise these networks".

More OSPF Commands

Passive Interface Command

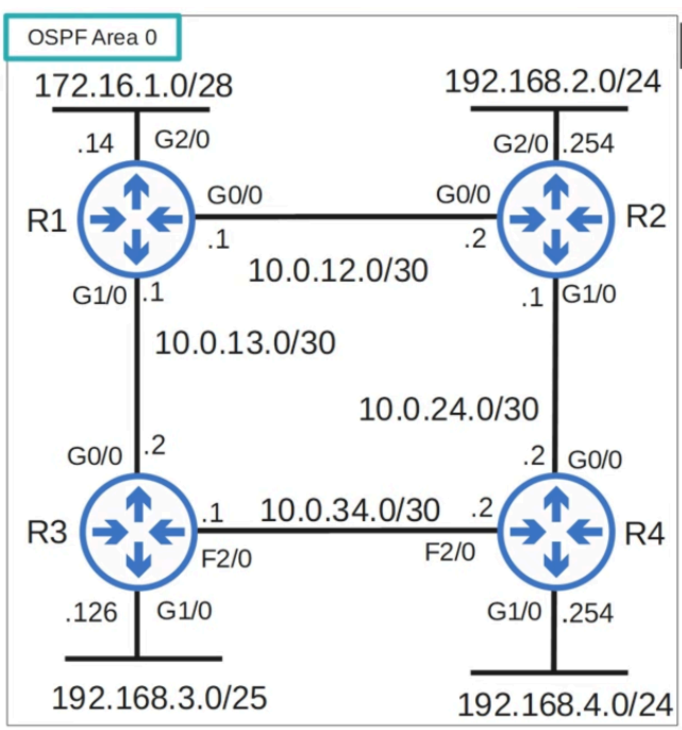

`R1(config-router)#passive-interface g2/0`

- You already know this command from RIP and EIGRP.
- The **passive-interface** command tells the rout to stop sending OSPF ‘hello’ messages out of the interface.
- However, the router will continue to send LSAs informing it’s neighbors about the subnet configured on the interface.
- You should always use this command on interfaces which don’t have any OSPF neighbors.

- The `passive-interface` command tells the router to stop sending OSPF 'hello' messages out of the specified interface.
- OSPF uses hello messages to tell other routers about itself.
- The router will continue to send LSAs informing its neighbors about the subnet configured on the interface.
- *Example:* Although R1 won't send hellos out of G2/0 and try to find OSPF neighbors, it will still tell its other neighbors about the 172.16.1.0/28 network.
- You should always use this command on interfaces which don't have any OSPF neighbors. It's a waste to continuously send hello messages out of an interface with no other routers connected to it.
- The `passive-interface` command is basically the same as for RIP and EIGRP.

```
R1(config)# router ospf 1
R1(config-router)# passive-interface GigabitEthernet2/0
```

Advertise a Default Route into OSPF

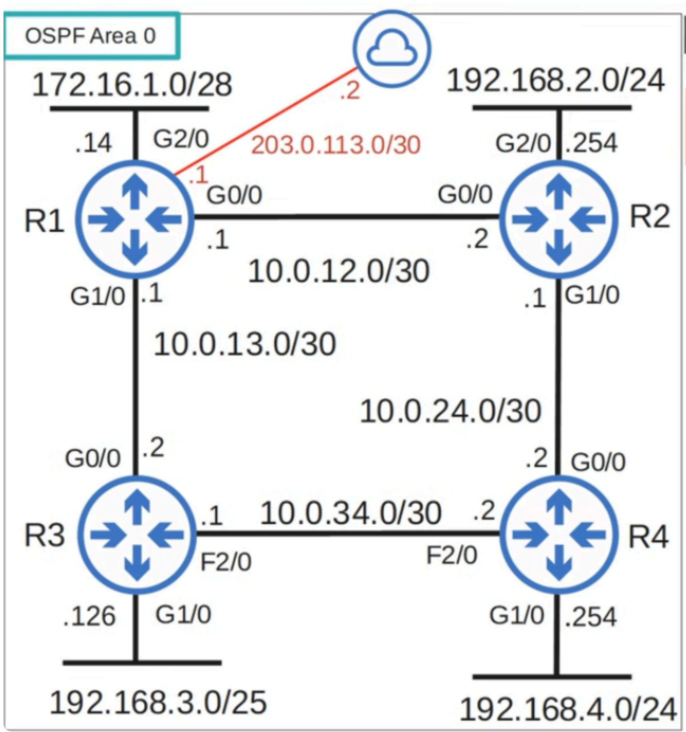

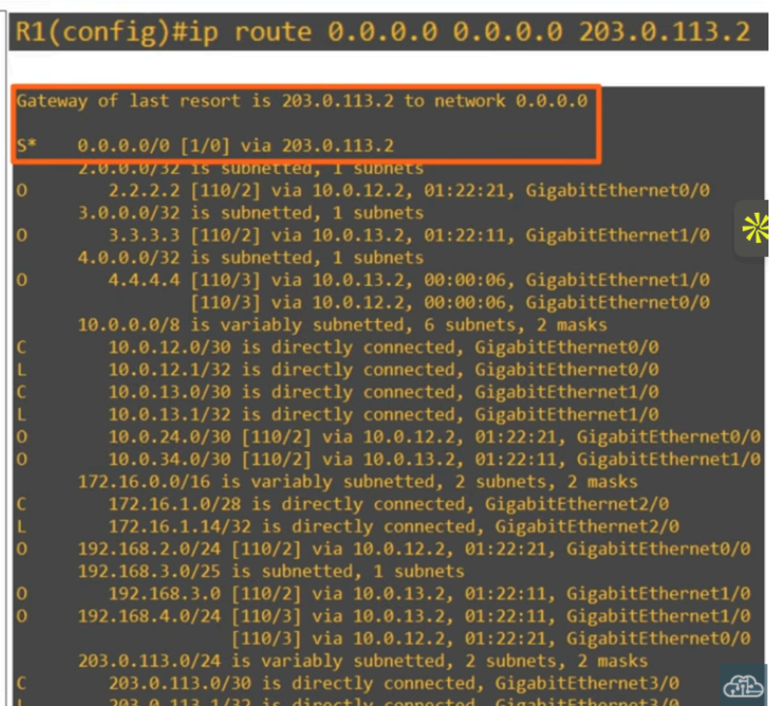

```
R1(config)#ip route 0.0.0.0 0.0.0.0 203.0.113.2

Gateway of last resort is 203.0.113.2 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 203.0.113.2
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/2] via 10.0.12.2, 01:22:21, GigabitEthernet0/0
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/2] via 10.0.13.2, 01:22:11, GigabitEthernet1/0
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/3] via 10.0.13.2, 00:00:06, GigabitEthernet1/0
                 [110/3] via 10.0.12.2, 00:00:06, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C        10.0.12.0/30 is directly connected, GigabitEthernet0/0
L        10.0.12.1/32 is directly connected, GigabitEthernet0/0
C        10.0.13.0/30 is directly connected, GigabitEthernet1/0
L        10.0.13.1/32 is directly connected, GigabitEthernet1/0
O        10.0.24.0/30 [110/2] via 10.0.12.2, 01:22:21, GigabitEthernet0/0
O        10.0.34.0/30 [110/2] via 10.0.13.2, 01:22:11, GigabitEthernet1/0
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.16.1.0/28 is directly connected, GigabitEthernet2/0
L        172.16.1.14/32 is directly connected, GigabitEthernet2/0
O     192.168.2.0/24 [110/2] via 10.0.12.2, 01:22:21, GigabitEthernet0/0
      192.168.3.0/25 is subnetted, 1 subnets
O        192.168.3.0 [110/2] via 10.0.13.2, 01:22:11, GigabitEthernet1/0
O     192.168.4.0/24 [110/3] via 10.0.13.2, 01:22:11, GigabitEthernet1/0
                     [110/3] via 10.0.12.2, 01:22:21, GigabitEthernet0/0
      203.0.113.0/24 is variably subnetted, 2 subnets, 2 masks
C        203.0.113.0/30 is directly connected, GigabitEthernet3/0
L        203.0.113.1/32 is directly connected, GigabitEthernet3/0
```

- The professor configured the static default route (`ip route 0.0.0.0 0.0.0.0 203.0.113.2`) because R1 has a direct **Internet connection** to an ISP with the IP address 203.0.113.2.

  He did this so that R1 would have a route to external networks, which he then advertised to the rest of the OSPF network using the `default-information originate` command.

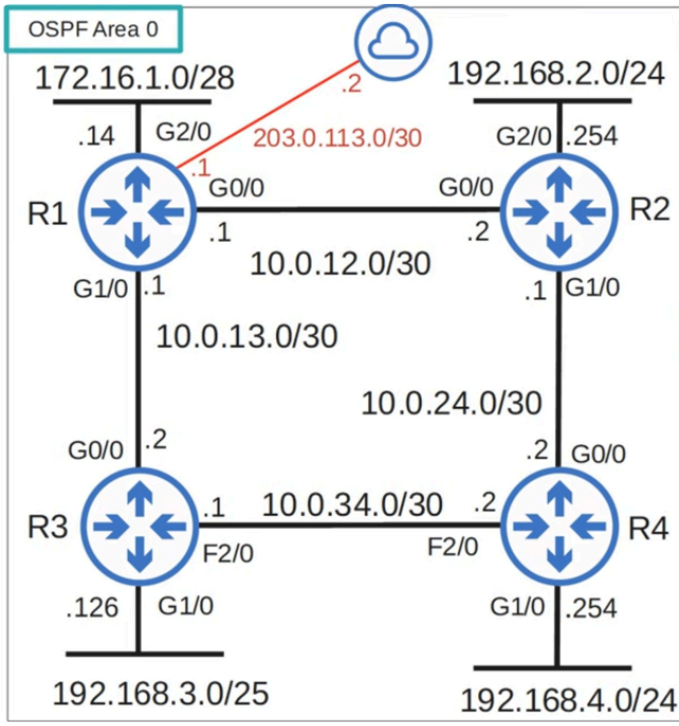

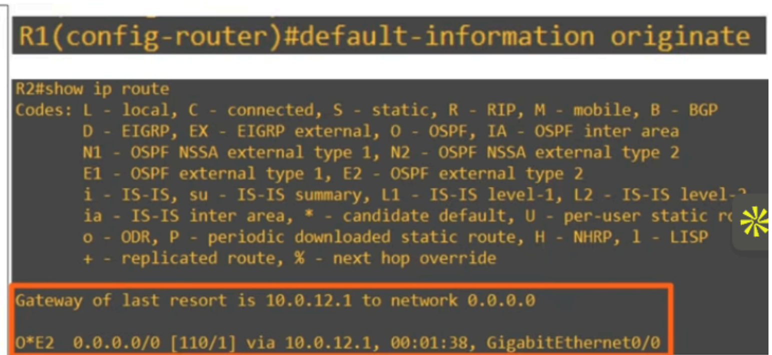

- To advertise a default route into OSPF, use the `DEFAULT-INFORMATION ORIGINATE` command.
- This will cause the router to create a new LSA and flood it.
- *Example:* R1 has an Internet connection and a default route pointing to the ISP's IP address (203.0.113.2). By using the `DEFAULT-INFORMATION ORIGINATE` command, R2, R3, and R4 will add a default route via R1 to their routing tables.

```
R1(config)# ip route 0.0.0.0 0.0.0.0 203.0.113.2
R1(config)# router ospf 1
R1(config-router)# default-information originate
```

Verification and Router ID Configuration

- The `SHOW IP PROTOCOLS` command displays the configured OSPF process ID.

Verification:

```
R1# show ip protocols
```

(Output snippet)

```
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 172.16.1.14
  It is an autonomous system boundary router
 Redistributing External Routes from,
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.12.0 0.0.0.3 area 0
    10.0.13.0 0.0.0.3 area 0
    172.16.1.0 0.0.0.15 area 0
  Passive Interface(s):
    GigabitEthernet2/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    4.4.4.4              110      00:00:08
    2.2.2.2              110      00:01:07
    3.3.3.3              110      00:01:07
    192.168.4.254        110      00:02:29
  Distance: (default is 110)
```

```
Routing Protocol is "ospf 1"
```

- "1" is the process ID. OSPF also uses a router ID.
- The OSPF router ID is determined in the same order of priority as EIGRP:
    1. Manually configured router ID.
    2. Highest IP address on a loopback interface.
    3. Highest IP address on a physical interface.
- *Example:* R1's router ID is 172.16.1.14 because no router ID is manually configured and there are no loopback interfaces.

Manually Configuring the Router ID

- The OSPF router ID is manually configured using the `ROUTER-ID` command in OSPF configuration mode.
- *Note:* In EIGRP, the command is `EIGRP ROUTER-ID`, but in OSPF it is just `ROUTER-ID`.

```
R1(config)# router ospf 1
R1(config-router)# router-id 1.1.1.1
```

```
R1(config-router)#router-id ?
  A.B.C.D  OSPF router-id in IP address format

R1(config-router)#router-id 1.1.1.1
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect
```

```
R1#clear ip ospf process
Reset ALL OSPF processes? [no]: yes
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
```

- After changing the router ID, the router displays a message indicating that a reload or the `clear ip ospf process` command is required for the change to take effect.
- The `CLEAR IP OSPF PROCESS` command is used in privileged exec mode to reset the OSPF process.
- *Note:* This is a bad idea in a real network because the router will lose all OSPF routes for a short time and won't be able to forward traffic to those destinations.
- *Note:* When `[no]` appears after a command, it means 'no' is the default choice.

```
R1# clear ip ospf process
```

ASBR and Load Balancing

```
R1#sh 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
  It is an autonomous system boundary router
 Redistributing External Routes from,
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.12.0 0.0.0.3 area 0
    10.0.13.0 0.0.0.3 area 0
    172.16.1.0 0.0.0.15 area 0
  Passive Interface(s):
    GigabitEthernet2/0
  Routing Information Sources:
    Gateway          Distance      Last Update
    2.2.2.2               110      00:01:40
    3.3.3.3               110      00:01:40
    4.4.4.4               110      00:01:40
  Distance: (default is 110)
```

- The `SHOW IP PROTOCOLS` output states, "It is an autonomous system boundary router".
- An **autonomous system boundary router (ASBR)** is an OSPF router that connects the OSPF network to an external network.
- *Example:* R1 is connected to the Internet. By using the `DEFAULT-INFORMATION ORIGINATE` command, R1 becomes an ASBR.
- OSPF supports ECMP (Equal-Cost Multi-Path) load-balancing over a maximum of 4 paths by default (unlike EIGRP, which supports unequal-cost load-balancing).
- The `MAXIMUM-PATHS` command is used in OSPF configuration mode to change the maximum number of load-balancing paths.

Verification:

```
R1# show ip protocols
```

(Output snippet)

```
Maximum path: 8
```

```
R1(config)# router ospf 1
R1(config-router)# maximum-paths 8
```

- The "routing for networks" section in the `show ip protocols` output shows the network commands used, which only determine which interfaces OSPF is activated on.
- The output also shows the configured passive interface and R1's neighbors with their router IDs.
- OSPF's default Administrative Distance (AD) is 110.
- The `DISTANCE` command is used in OSPF configuration mode to change the AD.
- *Example:* The AD was changed to 85 so OSPF routes are preferred over EIGRP routes on this router.

```
R1(config)# router ospf 1
R1(config-router)# distance 85
```