OSPF Concepts

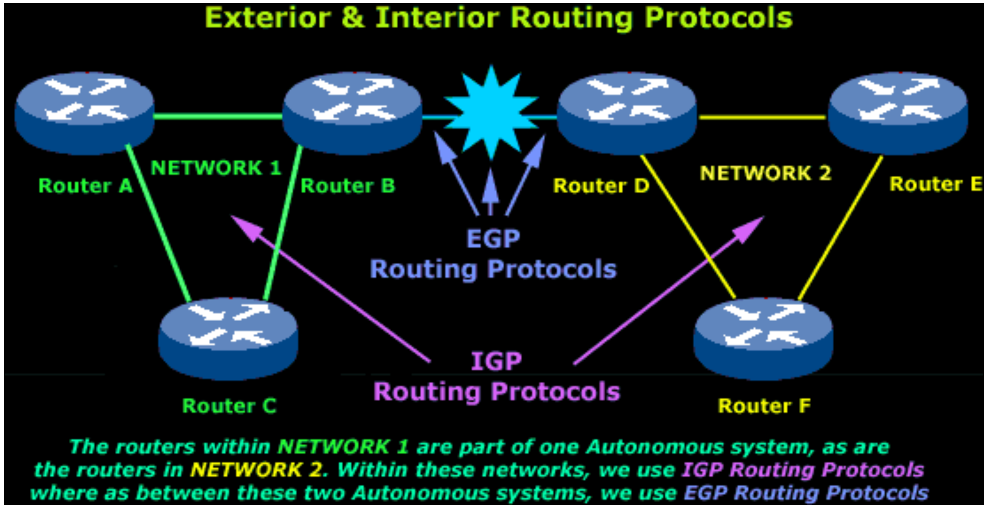

OSPF Overview

Introduction to OSPF

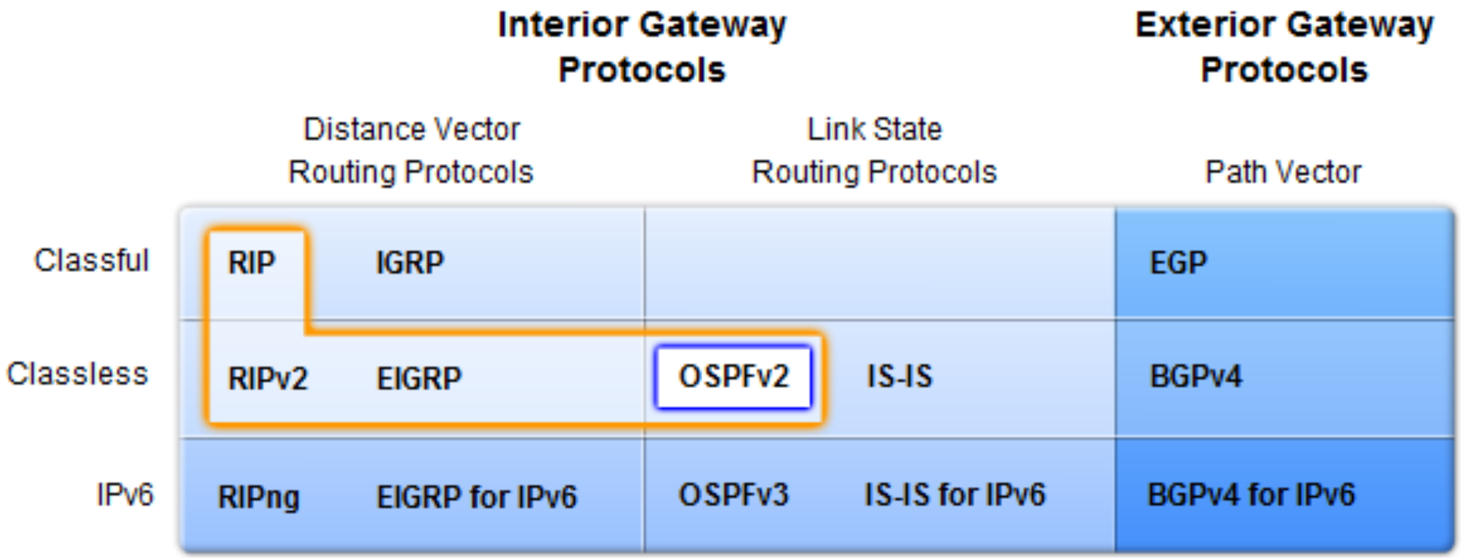

- OSPF is:
  - Classless
  - Link-state routing protocol
  - Uses areas for scalability
- RFC 2328 defines the OSPF metric as an arbitrary value called cost.
  - Cisco IOS software uses bandwidth to calculate the OSPF cost metric.

Background of OSPF

	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector Routing Protocols		Link State Routing Protocols		Path Vector
Classful	RIP	IGRP			EGP
Classless	RIPv2	EIGRP	OSPFv2	IS-IS	BGPv4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGPv4 for IPv6

- 1987 - Initial development by IETF OSPF Working Group.
- 1989 - OSPFv1 was published in RFC 1131.
- 1991 - OSPFv2 was introduced in RFC 1247 by John Moy.
- ISO was working IS-IS
- IETF chose OSPF as its recommended IGP (interior gateway protocol).
- In 1998 - OSPFv2 specification was updated in RFC 2328 and is the current RFC for OSPF.

Features of OSPF

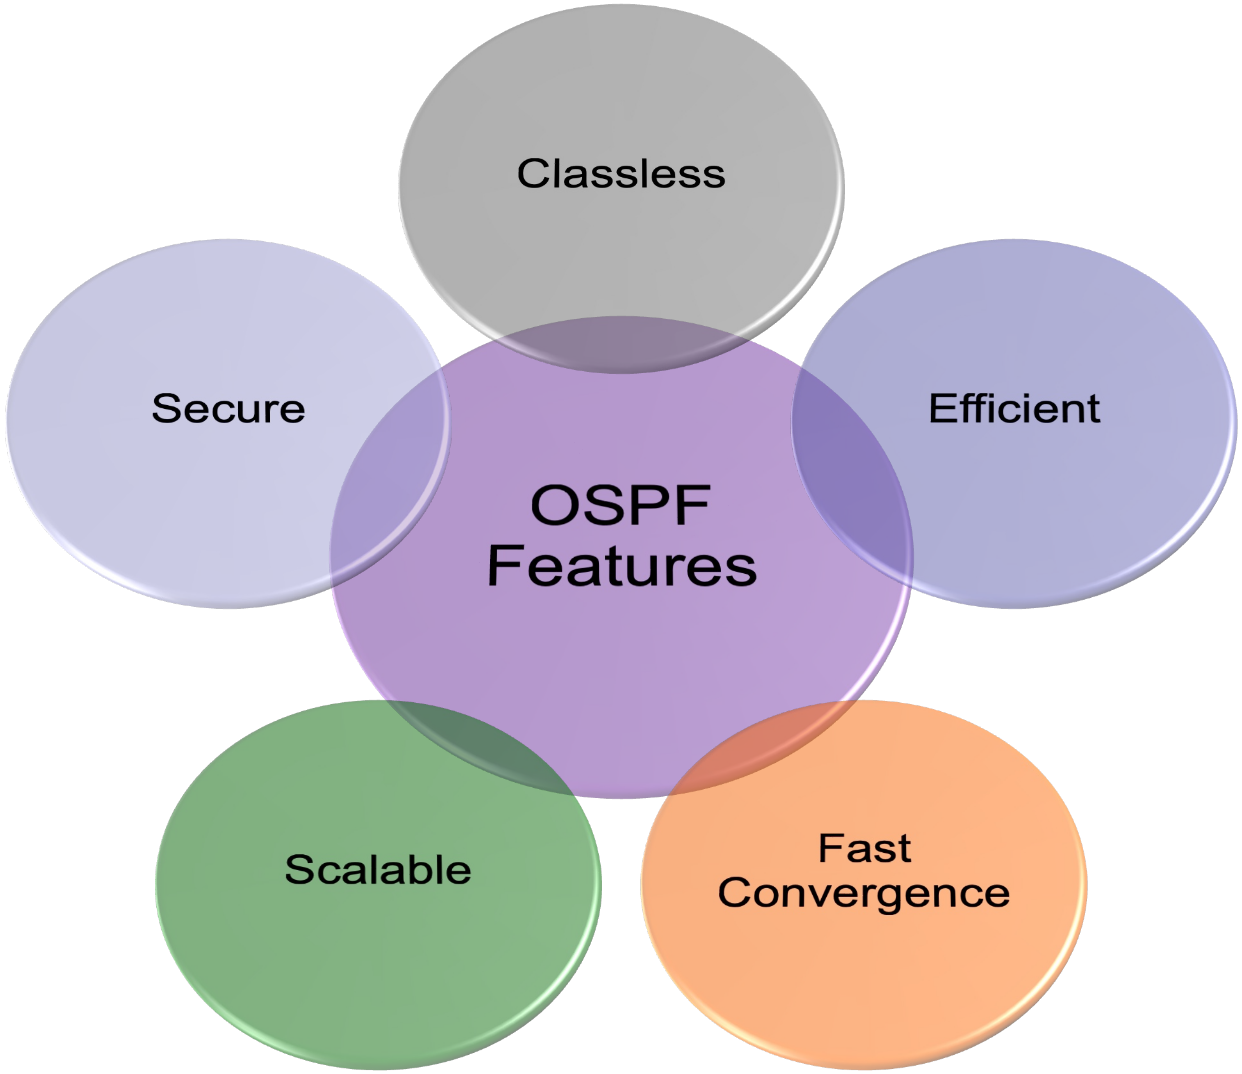

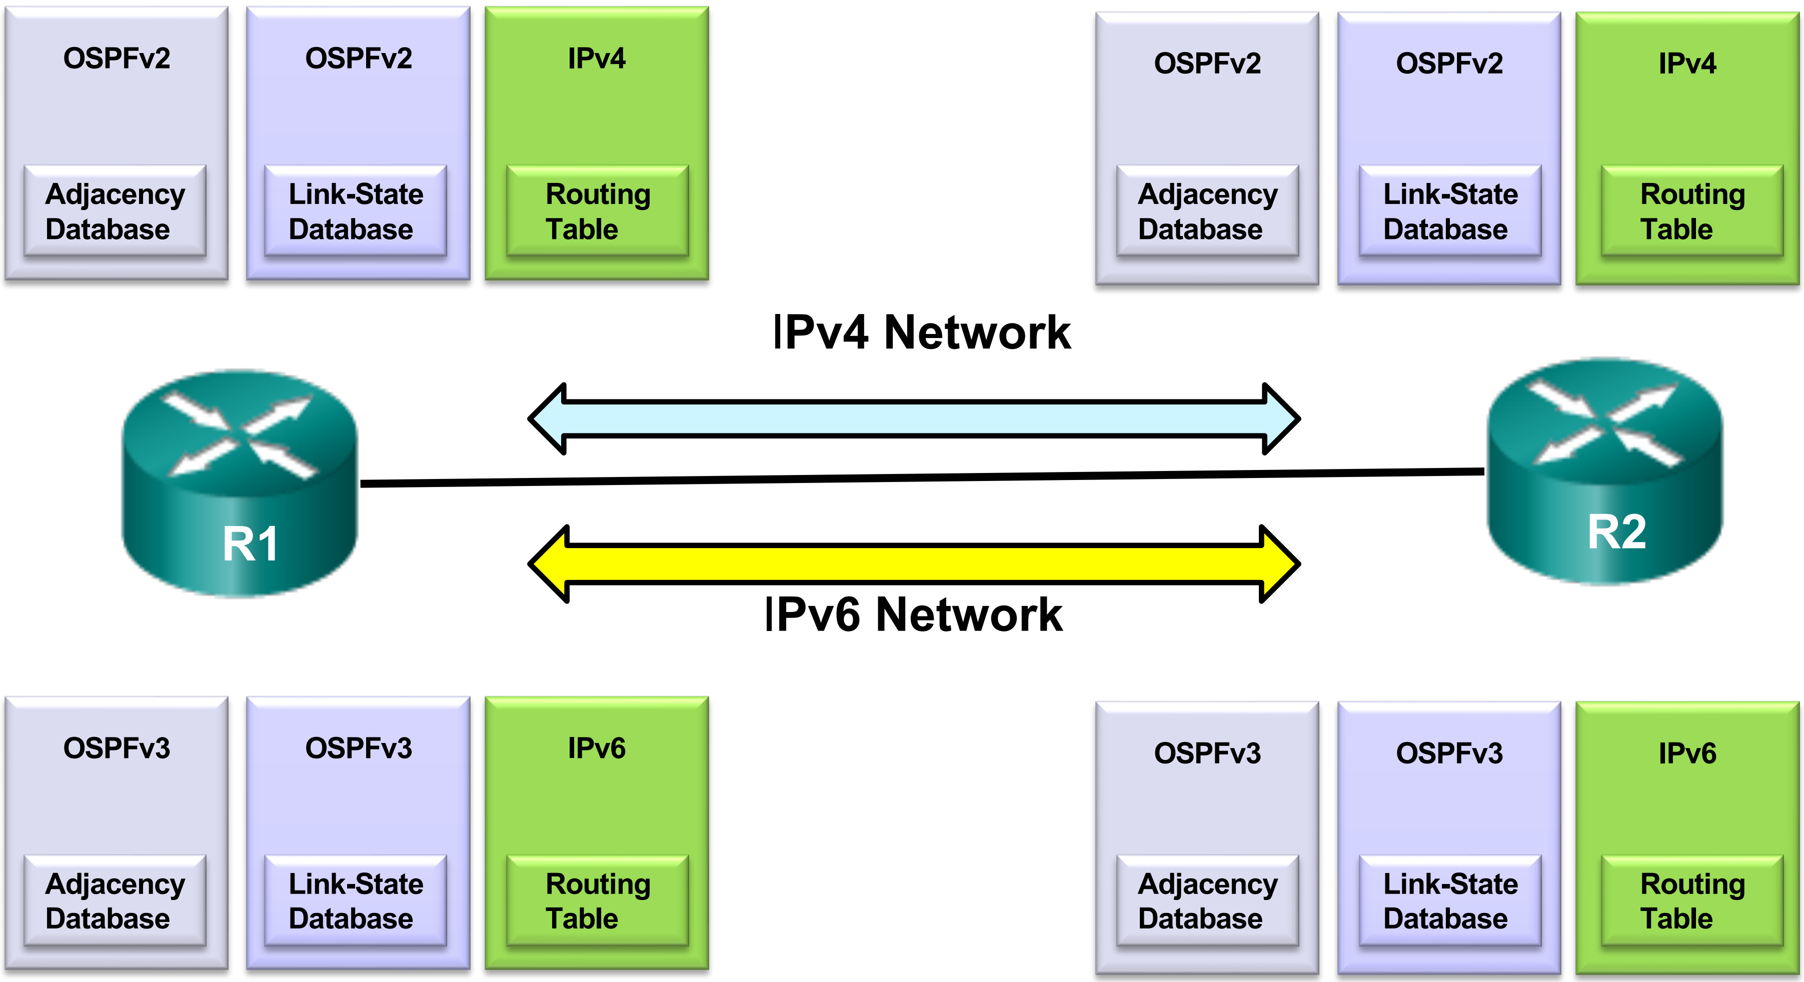

Note: OSPFv3 supports both IPv4 and IPv6 with the use of Address Families (beyond the scope of CCNA but in CCNP)

Link-State Routing Protocols

Shortest Path First (SPF) Algorithm

Distance Vector

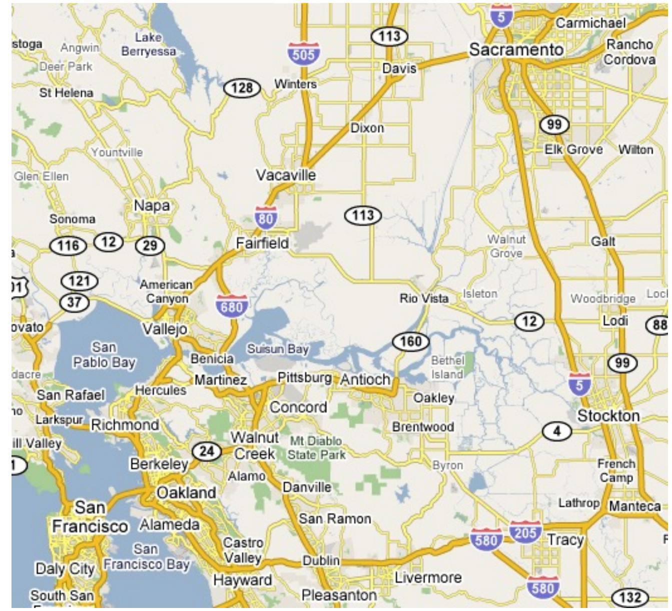

Link-State

- Link-state routing protocols (a.k.a. shortest path first protocols) are based on Edsger Dijkstra’s shortest path first (SPF) algorithm.

Shortest Path First (SPF) Algorithm

- Shortest path from R2 to the R3 LAN:
  - R2 to R1: 20
  - R1 to R3: 5
  - R3 - R3 LAN: 2

  Cost: 27

- Each router calculates the SPF algorithm and determines the cost from its own perspective. (more later)

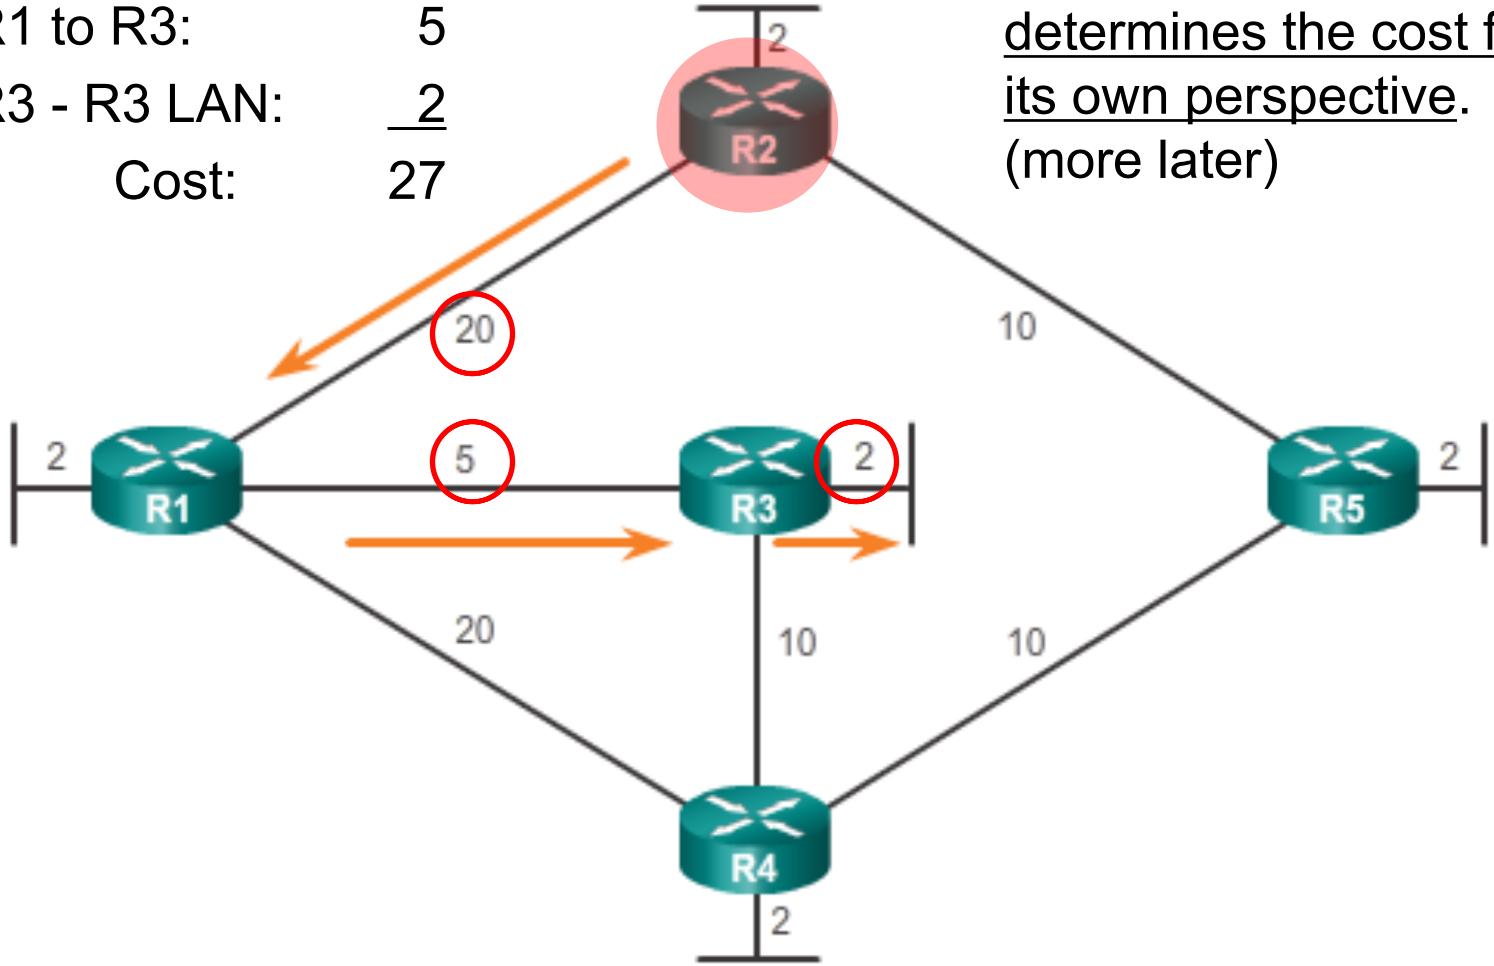

R1 SPF Tree

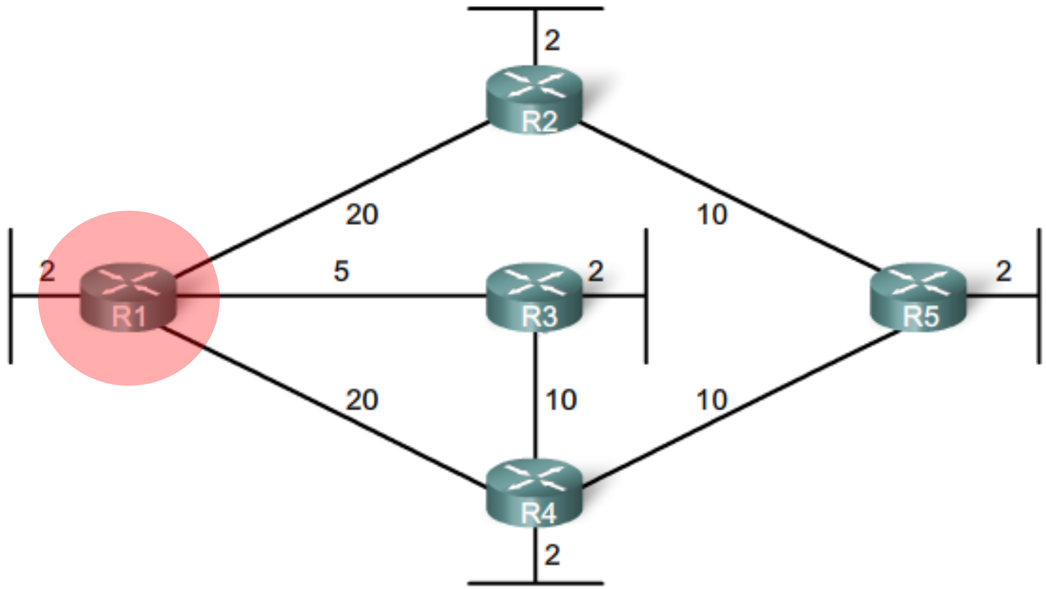

Destination	Shortest Path	Cost
R2 LAN	R1 to R2	22
R3 LAN	R1 to R3	7
R4 LAN	R1 to R3 to R4	17
R5 LAN	R1 to R3 to R4 to R5	27

Link 1
- Network: **10.1.0.0/16**
- IP address: **10.1.0.1**
- Type of network: **Ethernet**
- Cost of that link: **2**
- Neighbors: **None**

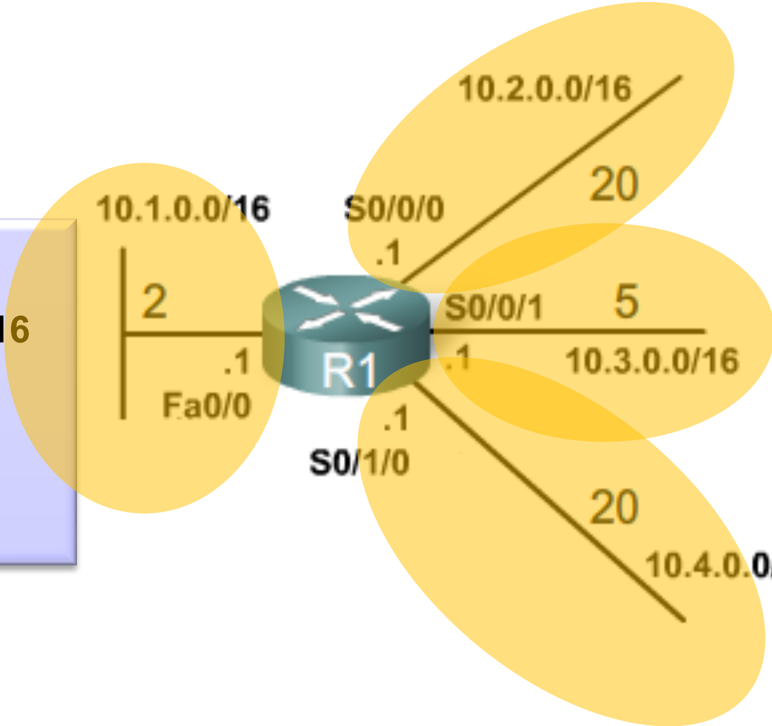

Link 2
- Network: **10.2.0.0/16**
- IP address: **10.2.0.1**
- Type of network: **Serial**
- Cost of that link: **20**
- Neighbors: **R2**

Link 3
- Network: **10.3.0.0/16**
- IP address: **10.3.0.1**
- Type of network: **Serial**
- Cost of that link: **5**
- Neighbors: **R3**

Link 4
- Network: **10.4.0.0/16**
- IP address: **10.4.0.1**
- Type of network: **Serial**
- Cost of that link: **20**
- Neighbors: **R4**

- ***Link states -*** Information about the state of a router' s links
- This information includes interface' s:
  - IP address/mask
  - Type of network
    - Ethernet (broadcast) or serial point-to-point link
  - Cost of that link
  - Any neighbor routers on that link

Building the Shortest Path First (SPF) Tree

Link State Database for R1

LSPs from R2	Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
	Has a network 10.5.0.0/16, cost of 2
LSPs from R3	Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
	Has a network 10.6.0.0/16, cost of 2
LSPs from R4	Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
	Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
	Has a network 10.8.0.0/16, cost of 2
LSPs from R5	Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
	Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
	Has a network 10.11.0.0/16, cost of 2
R1 link states	Connected to neighbor R2 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.4.0.0/16, cost of 20
	Has a network 10.1.0.0/16, cost of 2

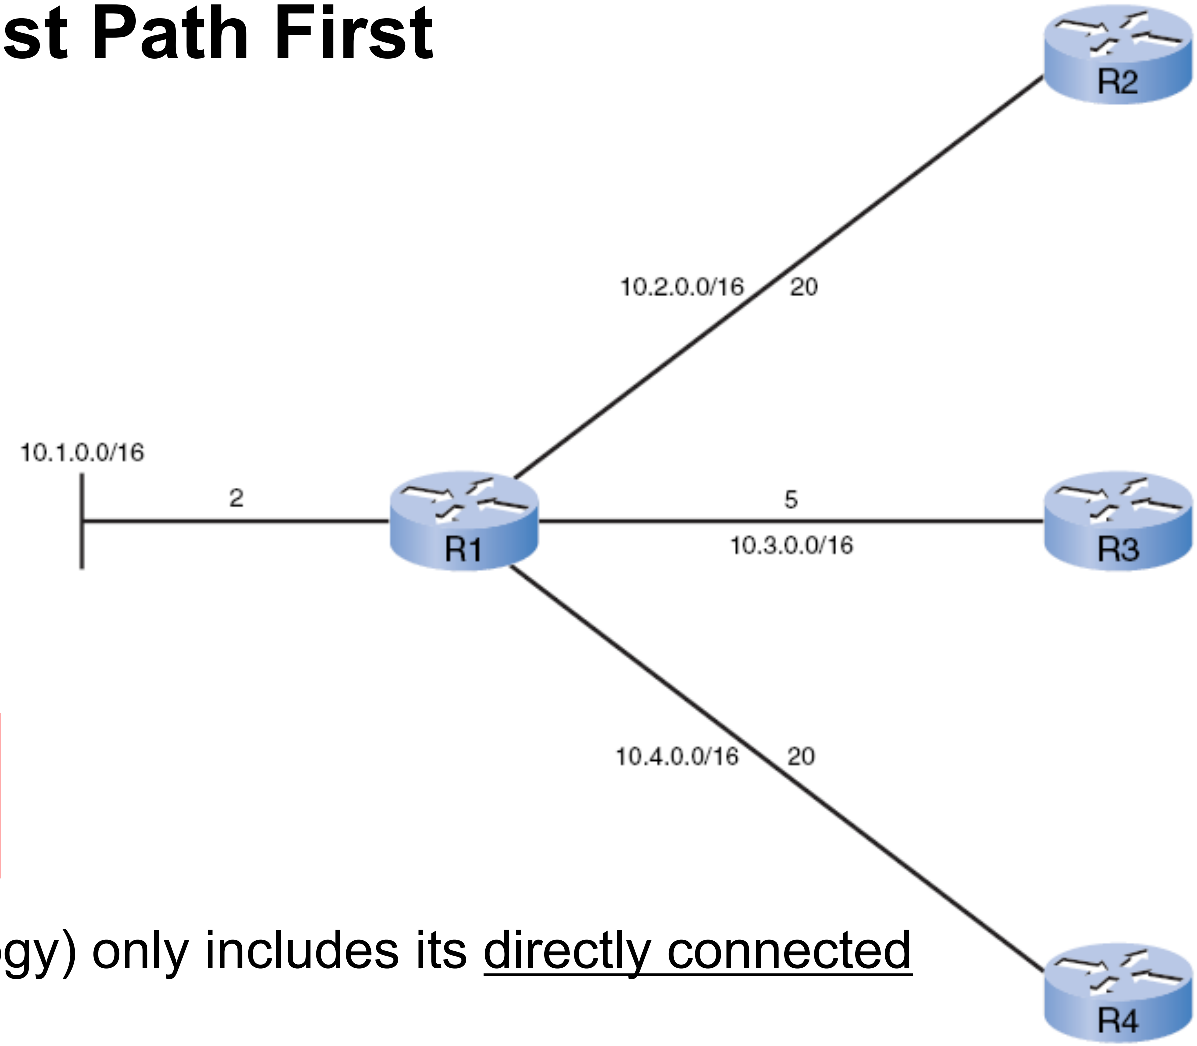

- <u>At first</u>, the tree (topology) only includes its <u>directly connected neighbors</u>.
- Using the <u>link-state information</u> from all other routers, R1 can now <u>begin to construct an SPF tree</u> of the network with <u>itself at the root</u> of the tree.

Link State Database for R1

LSPs from R2	Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
	Has a network 10.5.0.0/16, cost of 2
LSPs from R3	Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
	Has a network 10.6.0.0/16, cost of 2
LSPs from R4	Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
	Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
	Has a network 10.8.0.0/16, cost of 2
LSPs from R5	Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
	Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
	Has a network 10.11.0.0/16, cost of 2
R1 link states	Connected to neighbor R2 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.4.0.0/16, cost of 20
	Has a network 10.1.0.0/16, cost of 2

R1 Processes the LSPs from R2

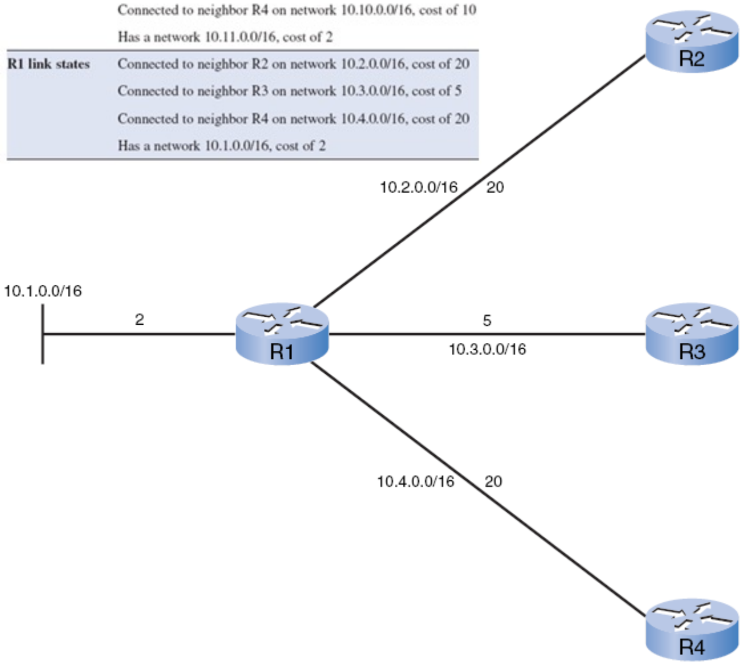

Red: New information for tree.

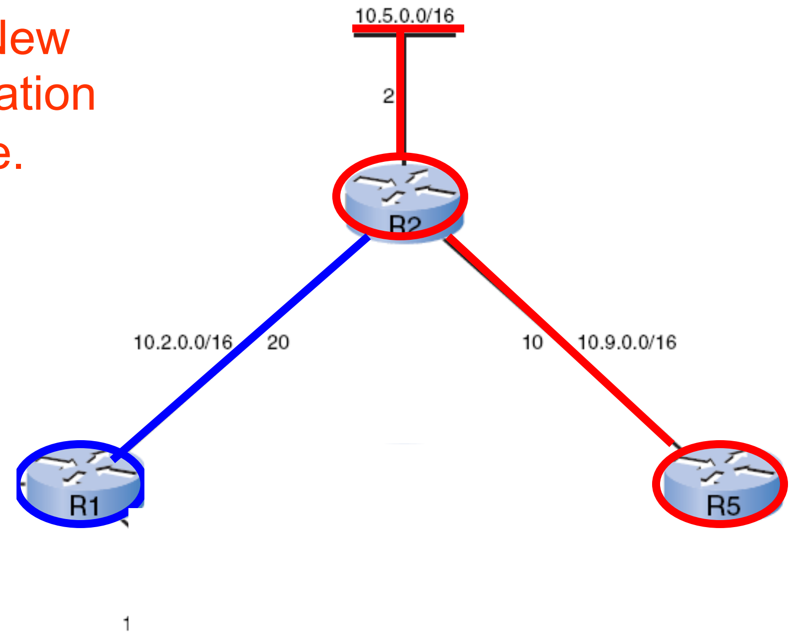

- The SPF algorithm begins by <u>processing the following LSP information from R2</u>:
  - Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
  - Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
  - Has a network 10.5.0.0/16, cost of 2

Link State Database for R1

LSPs from R2	Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
	Has a network 10.5.0.0/16, cost of 2
LSPs from R3	Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
	Has a network 10.6.0.0/16, cost of 2
LSPs from R4	Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
	Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
	Has a network 10.8.0.0/16, cost of 2
LSPs from R5	Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
	Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
	Has a network 10.11.0.0/16, cost of 2
R1 link states	Connected to neighbor R2 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.4.0.0/16, cost of 20
	Has a network 10.1.0.0/16, cost of 2

R1 Processes the LSPs from R3

Red: New information for tree.

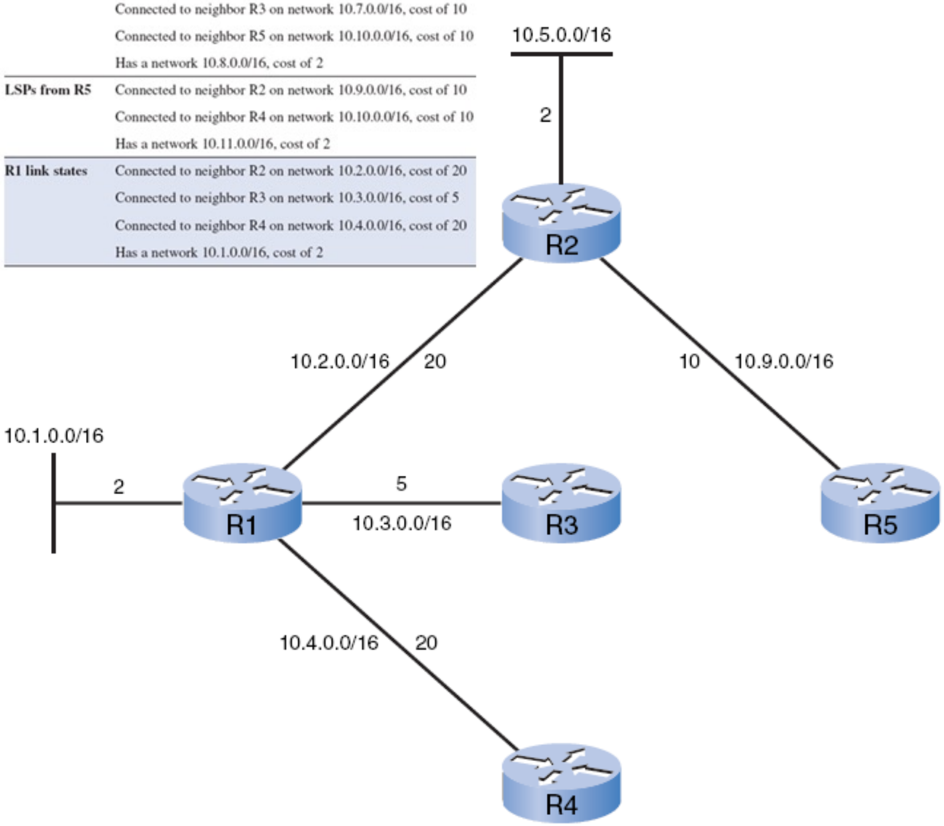

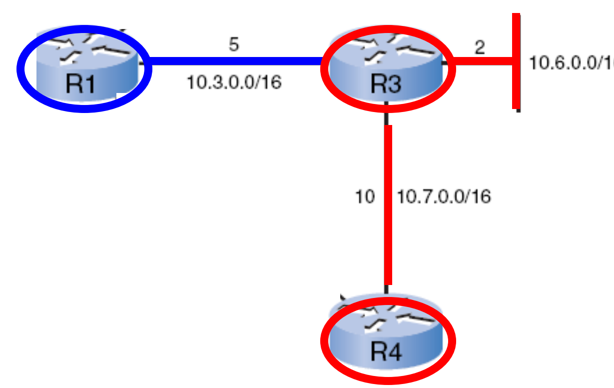

- The SPF algorithm begins by processing the following LSP information from R3:
  - Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
  - Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
  - Has a network 10.6.0.0/16, cost of 2

R1 Processes the LSPs from R4

Link State Database for R1

LSPs from R2	Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
	Has a network 10.5.0.0/16, cost of 2
LSPs from R3	Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
	Has a network 10.6.0.0/16, cost of 2
LSPs from R4	Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
	Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
	Has a network 10.8.0.0/16, cost of 2
LSPs from R5	Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
	Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
	Has a network 10.11.0.0/16, cost of 2
R1 link states	Connected to neighbor R2 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.4.0.0/16, cost of 20
	Has a network 10.1.0.0/16, cost of 2

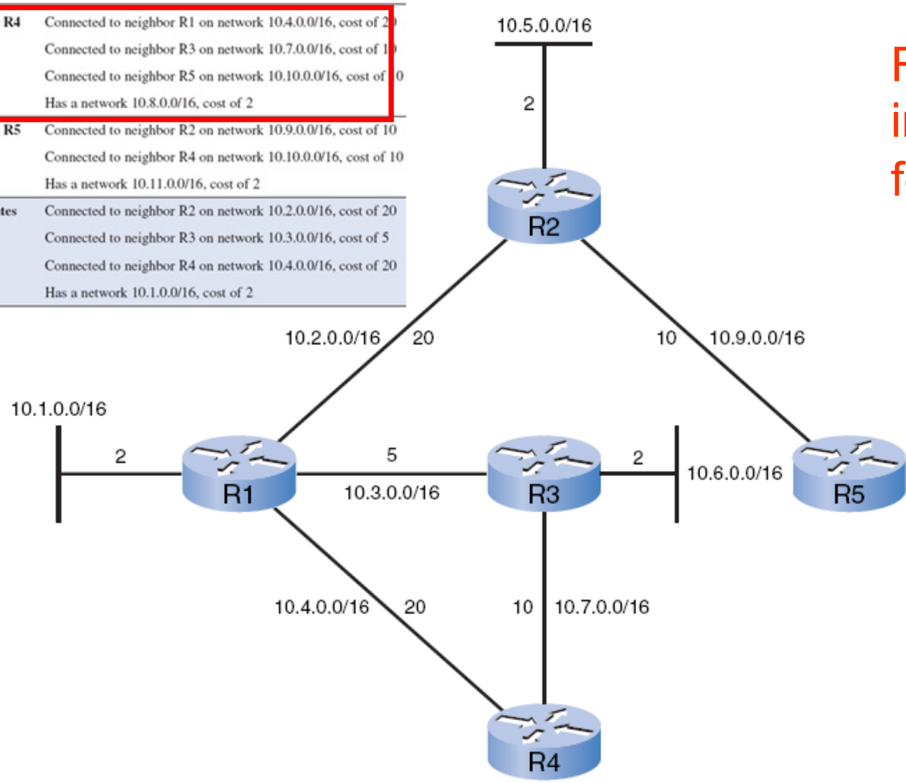

Red: New information for tree.

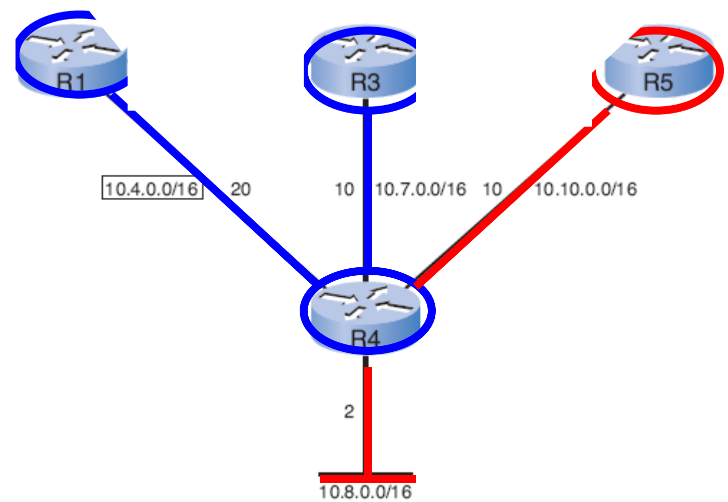

- The SPF algorithm begins by processing the following LSP information from R4:
  - Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
  - Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
  - Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
  - Has a network 10.8.0.0/16, cost of 2

Link State Database for R1

LSPs from R2	Connected to neighbor R1 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R5 on network 10.9.0.0/16, cost of 10
	Has a network 10.5.0.0/16, cost of 2
LSPs from R3	Connected to neighbor R1 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.7.0.0/16, cost of 10
	Has a network 10.6.0.0/16, cost of 2
LSPs from R4	Connected to neighbor R1 on network 10.4.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.7.0.0/16, cost of 10
	Connected to neighbor R5 on network 10.10.0.0/16, cost of 10
	Has a network 10.8.0.0/16, cost of 2
LSPs from R5	Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
	Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
	Has a network 10.11.0.0/16, cost of 2
R1 link states	Connected to neighbor R2 on network 10.2.0.0/16, cost of 20
	Connected to neighbor R3 on network 10.3.0.0/16, cost of 5
	Connected to neighbor R4 on network 10.4.0.0/16, cost of 20
	Has a network 10.1.0.0/16, cost of 2

R1 Processes the LSPs from R5

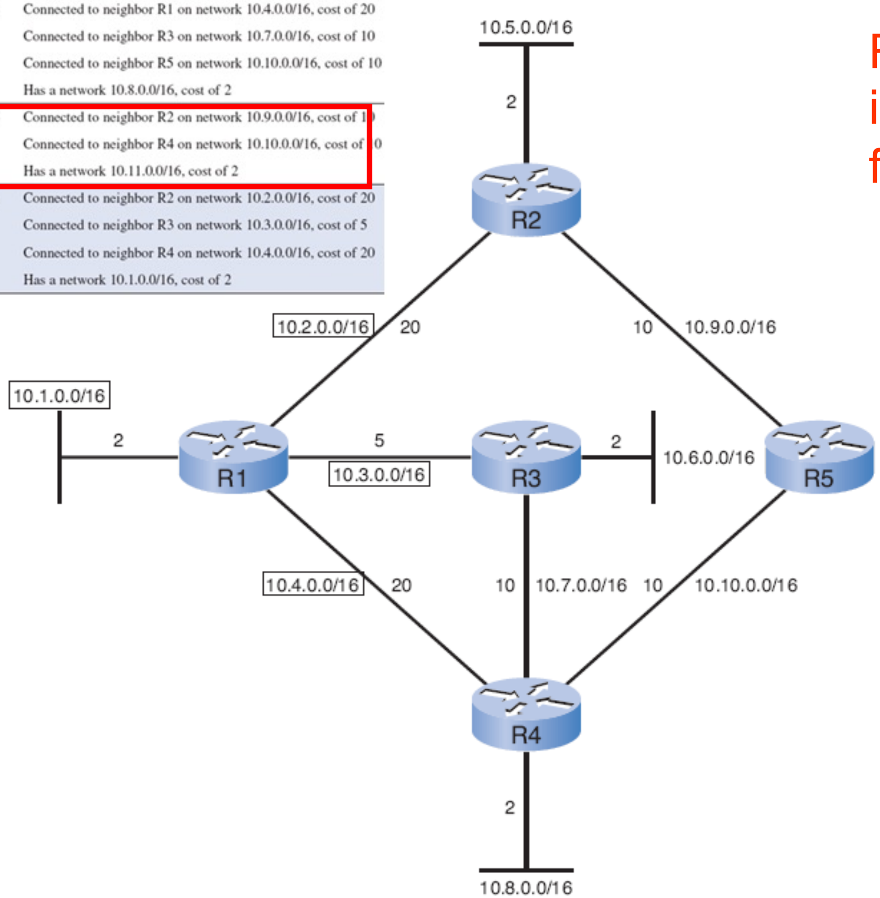

Red: New information for tree.

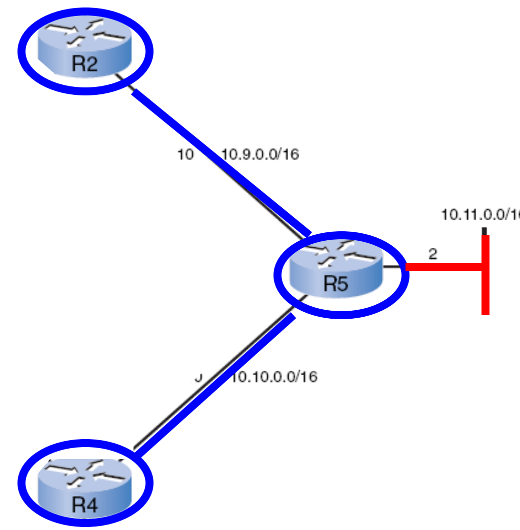

- The SPF algorithm begins by processing the following LSP information from R5:
  - Connected to neighbor R2 on network 10.9.0.0/16, cost of 10
  - Connected to neighbor R4 on network 10.10.0.0/16, cost of 10
  - Has a network 10.11.0.0/16, cost of 2

Graph of Network

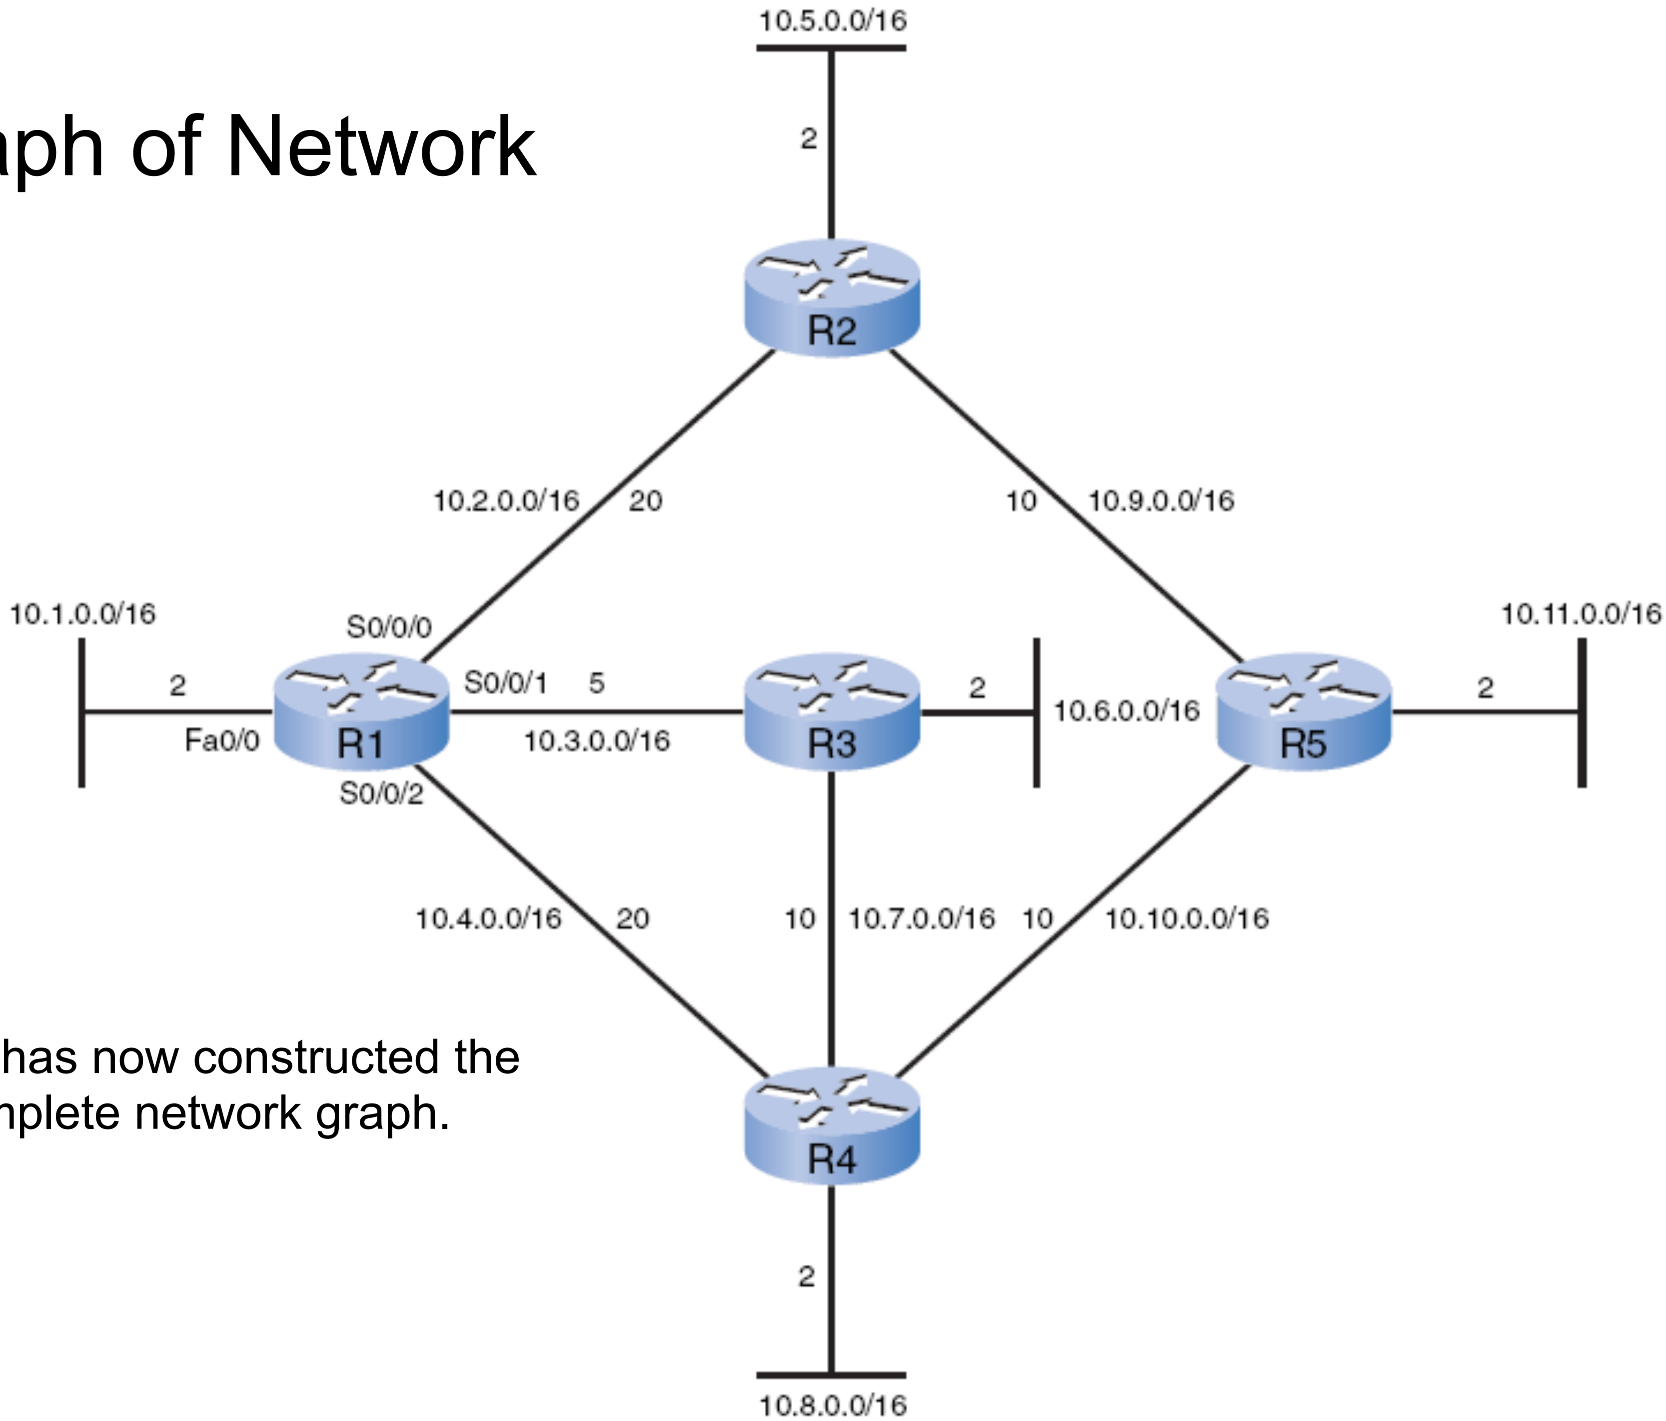

- R1 has now constructed the complete network graph.

Determining the Shortest Path

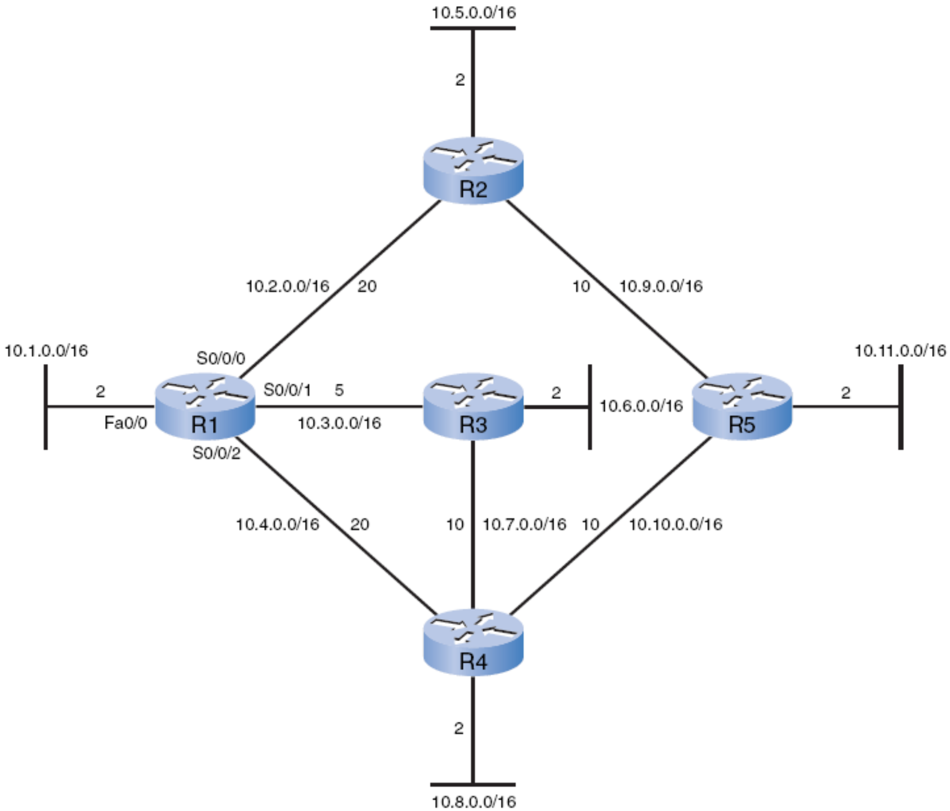

- Using the network graph, SPF algorithm results in the shortest path to each network.
  - Note: Only the LANs are shown in the table, but SPF can also be used to determine the shortest path to each WAN link network.

Determining the Shortest Path

Network 10.5.0.0/16
via R2 Serial 0/0/0
at a cost of 22

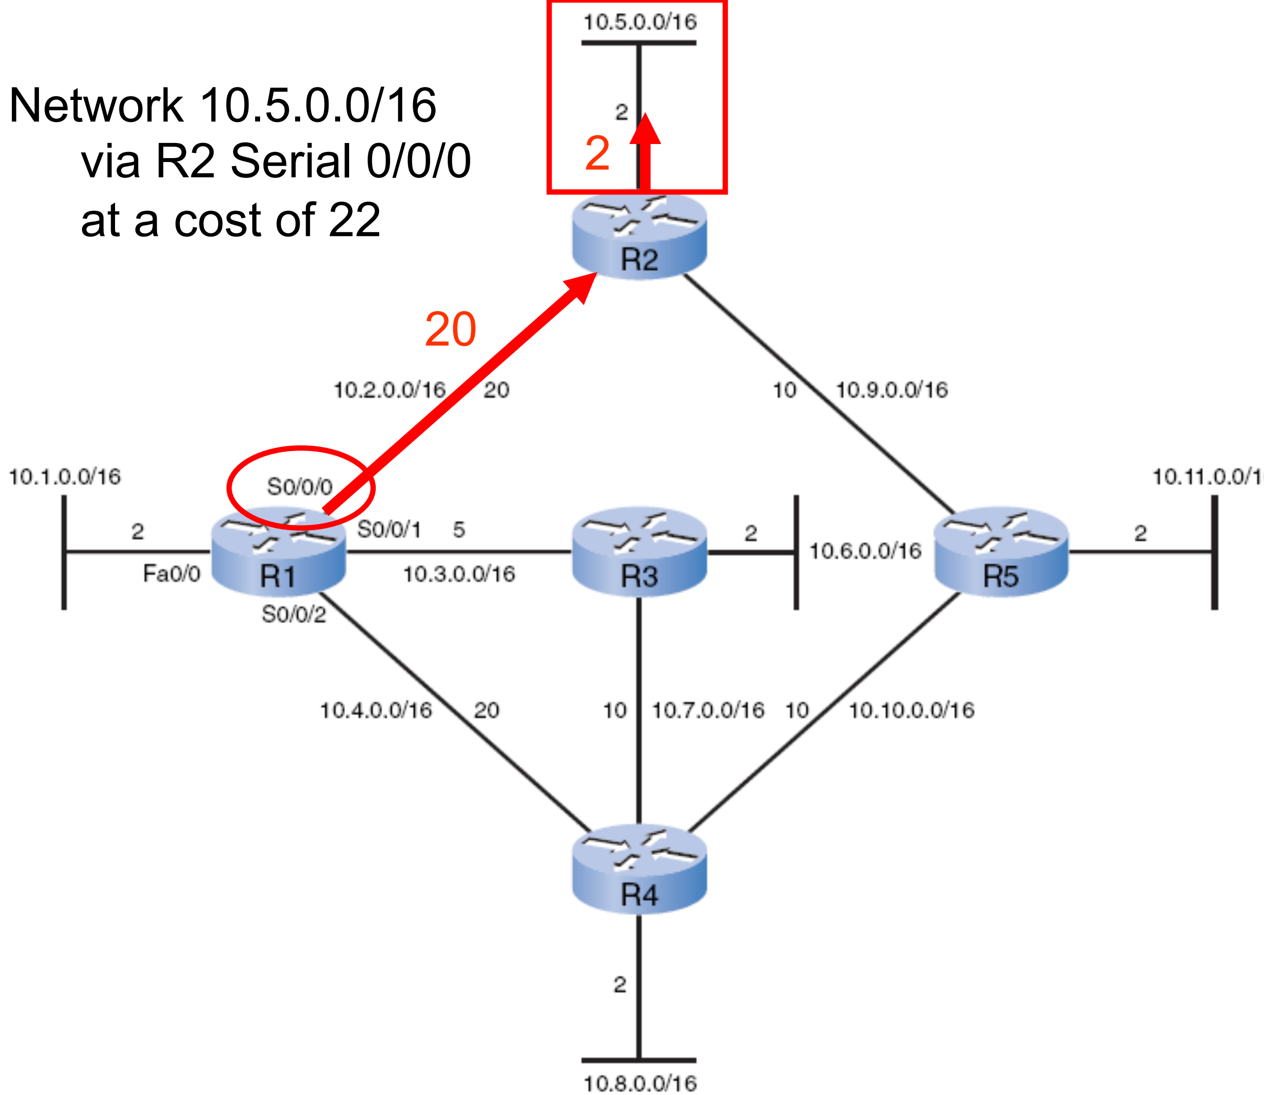

Determining the Shortest Path

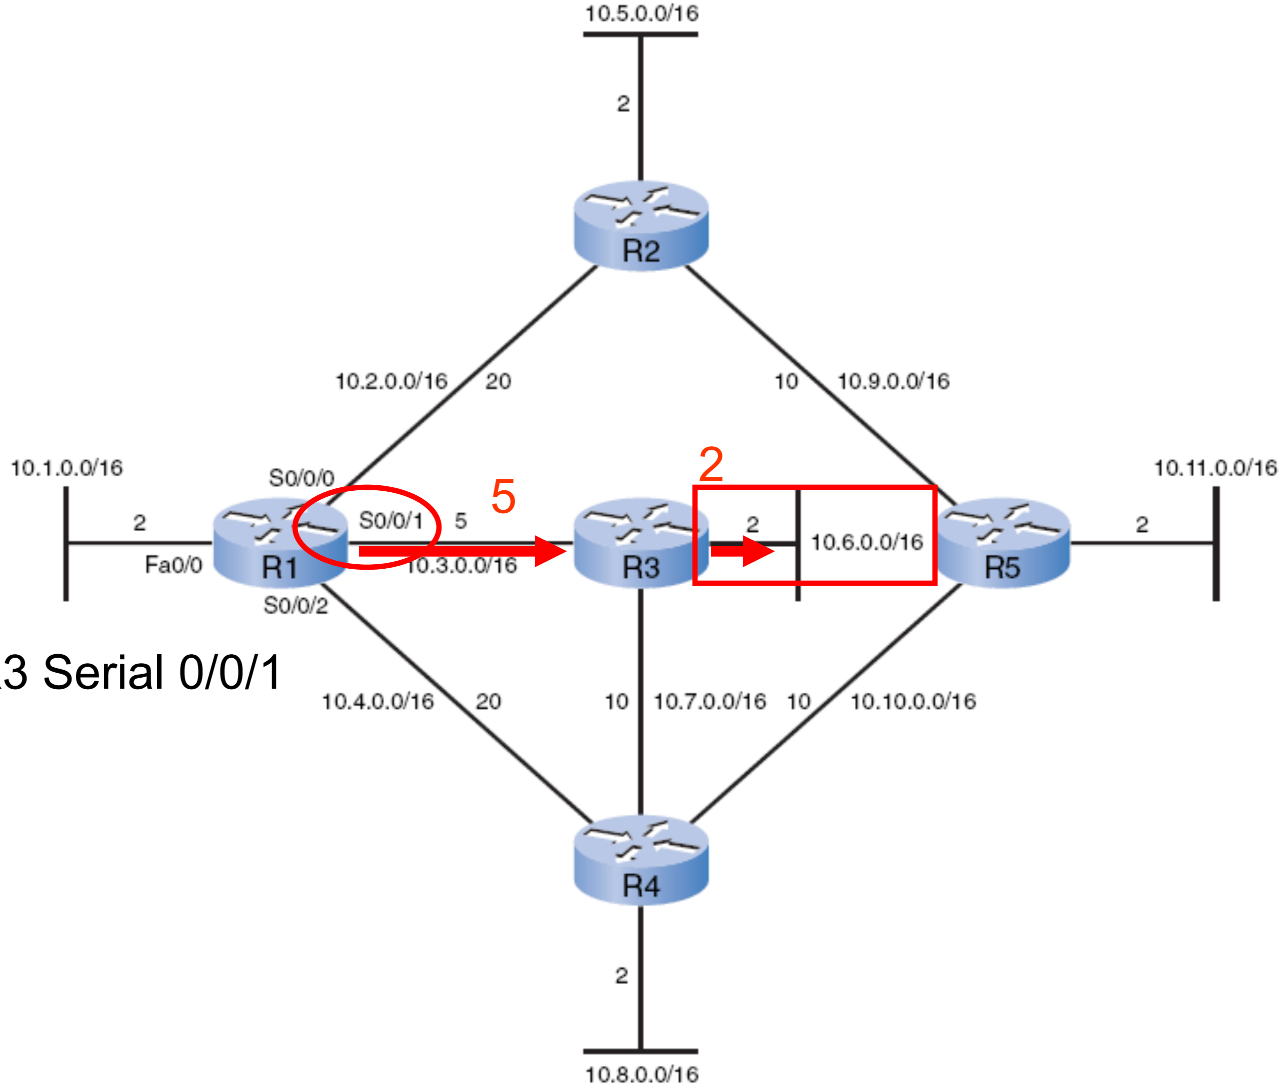

Network 10.6.0.0/16 via R3 Serial 0/0/1 at a cost of 7

Determining the Shortest Path

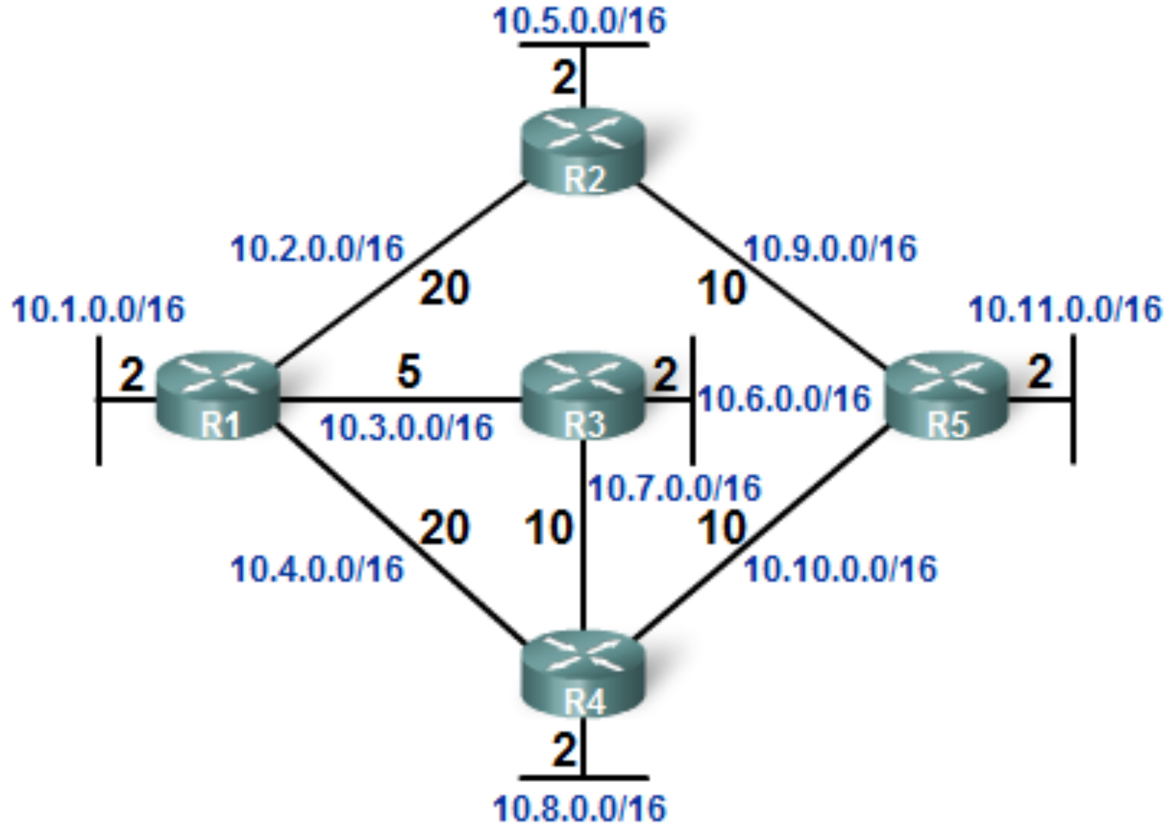

- <u>Each router</u> <u>constructs its own SPF tree independently</u> from all other routers.
- Link-state databases must be <u>identical on all routers</u>.

Generating a Routing Table from the SPF Tree

SPF Tree for R1

SPF Information

- Network 10.5.0.0/16 via R2 serial 0/0/0 at a cost of 22
- Network 10.6.0.0/16 via R3 serial 0/0/1 at a cost of 7
- Network 10.7.0.0/16 via R3 serial 0/0/1 at a cost of 15
- Network 10.8.0.0/16 via R3 serial 0/0/1 at a cost of 17
- Network 10.9.0.0/16 via R2 serial 0/0/0 at a cost of 30
- Network 10.10.0.0/16 via R3 serial 0/0/1 at a cost of 25
- Network 10.11.0.0/16 via R3 serial 0/0/1 at a cost of 27

R1 Routing Table

Directly Connected Networks

- 10.1.0.0/16 Directly Connected Network
- 10.2.0.0/16 Directly Connected Network
- 10.3.0.0/16 Directly Connected Network
- 10.4.0.0/16 Directly Connected Network

Remote Networks

- 10.5.0.0/16 via R2 serial 0/0/0, cost = 22
- 10.6.0.0/16 via R3 serial 0/0/1, cost = 7
- 10.7.0.0/16 via R3 serial 0/0/1, cost = 15
- 10.8.0.0/16 via R3 serial 0/0/1, cost = 17
- 10.9.0.0/16 via R2 serial 0/0/0, cost = 30
- 10.10.0.0/16 via R3 serial 0/0/1, cost = 25
- 10.11.0.0/16 via R3 serial 0/0/1, cost = 27

- These paths listed previously can now be added to the routing table.
- The routing table will also include
  - Directly connected networks
  - Routes from any other sources, such as static routes.
- Packets will now be forwarded according to these entries in the routing table.

Addressing Disadvantages

- Create areas to minimize the router memory requirements, processing requirements, and bandwidth requirenments.

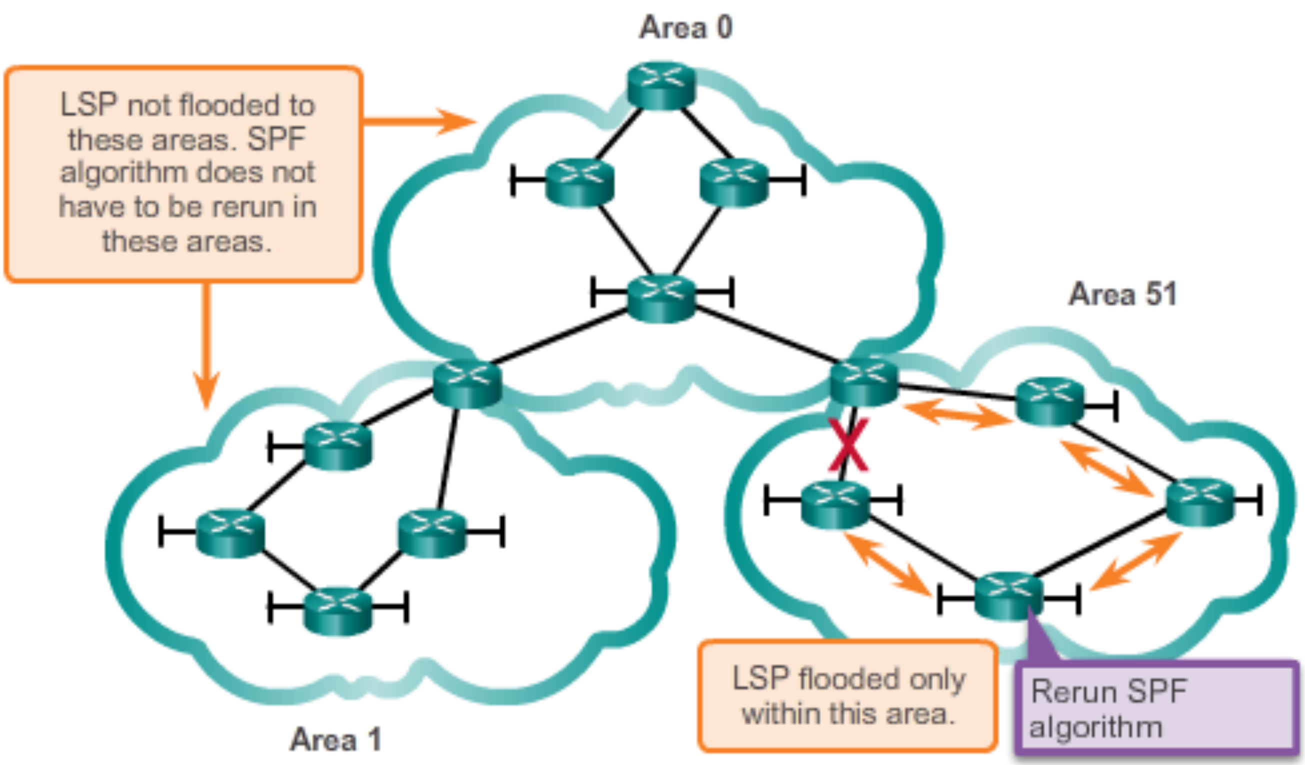

OSPF Messages

OSPF Message Encapsulation

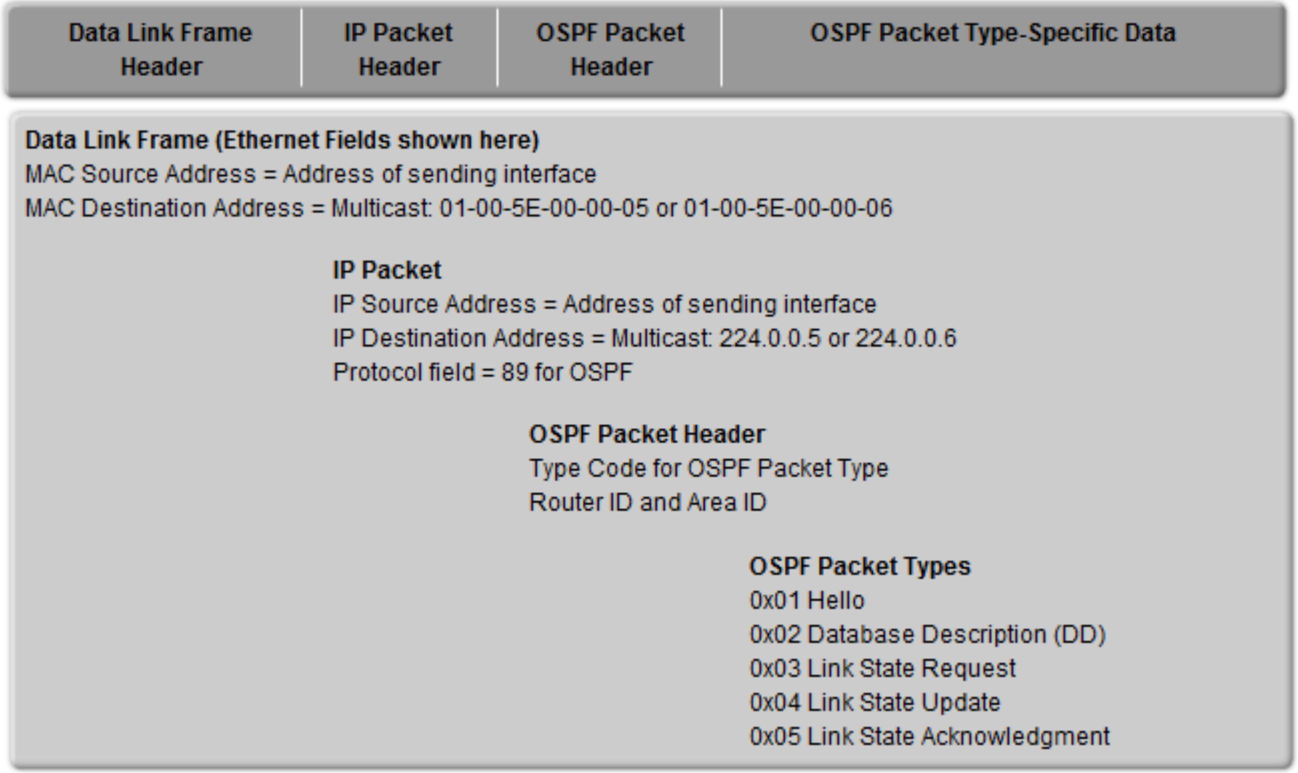

- In the IP packet header:
  - **Protocol field** is set to 89 (OSPF)
  - **Destination address** is typically set to one of two multicast addresses:
    - 224.0.0.5
    - 224.0.0.6
- Destination MAC address is also a multicast address:
  - 01-00-5E-00-00-05
  - 01-00-5E-00-00-06

OSPF Packet Types

Figure includes CCNP information

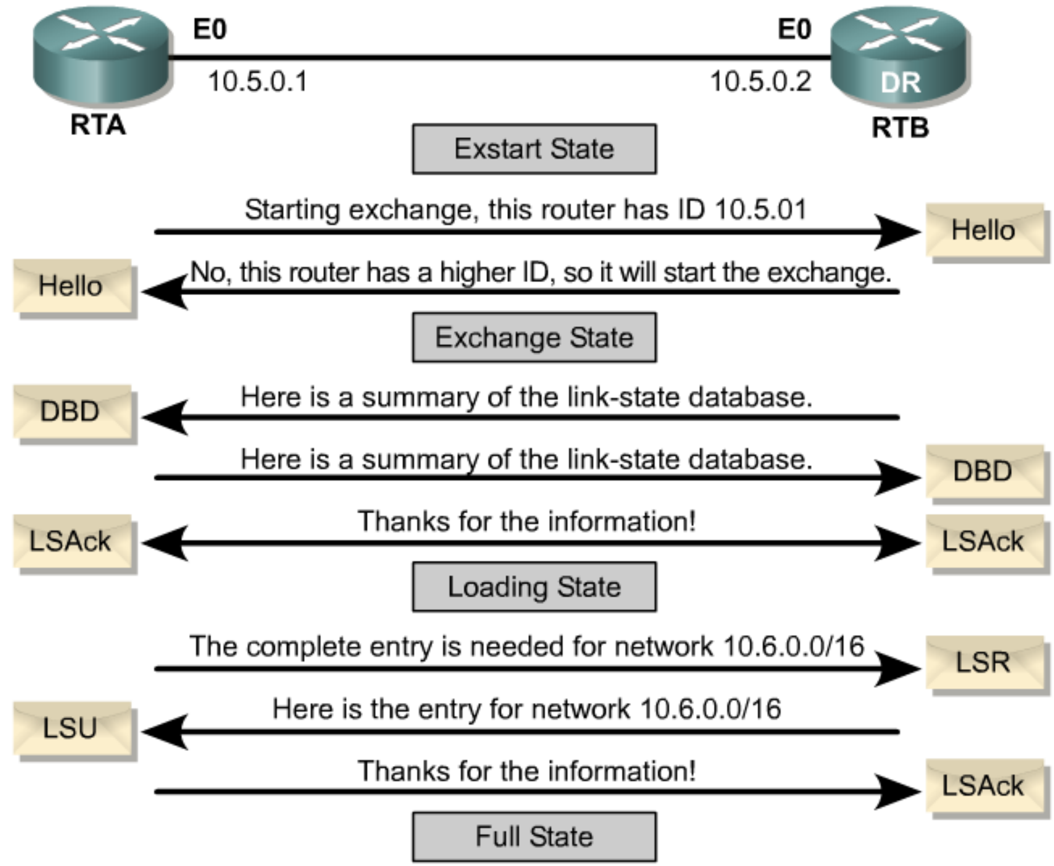

- Five types of OSPF LSPs (link-state packets).
  - **Hello**: Used to establish and maintain adjacency.
  - **DBD (Database Description):** Abbreviated list of the sending router's link-state database.
  - **LSR (Link-State Request)** : Used by routers to request more information about any entry in the DBD.
  - **LSU**: (**Link-State Update):** Link-state information.
  - **LSAck (LSA Acknowledgment)**: Router sends a link-state (LSAck) to confirm receipt of the LSU.

Hello Protocol

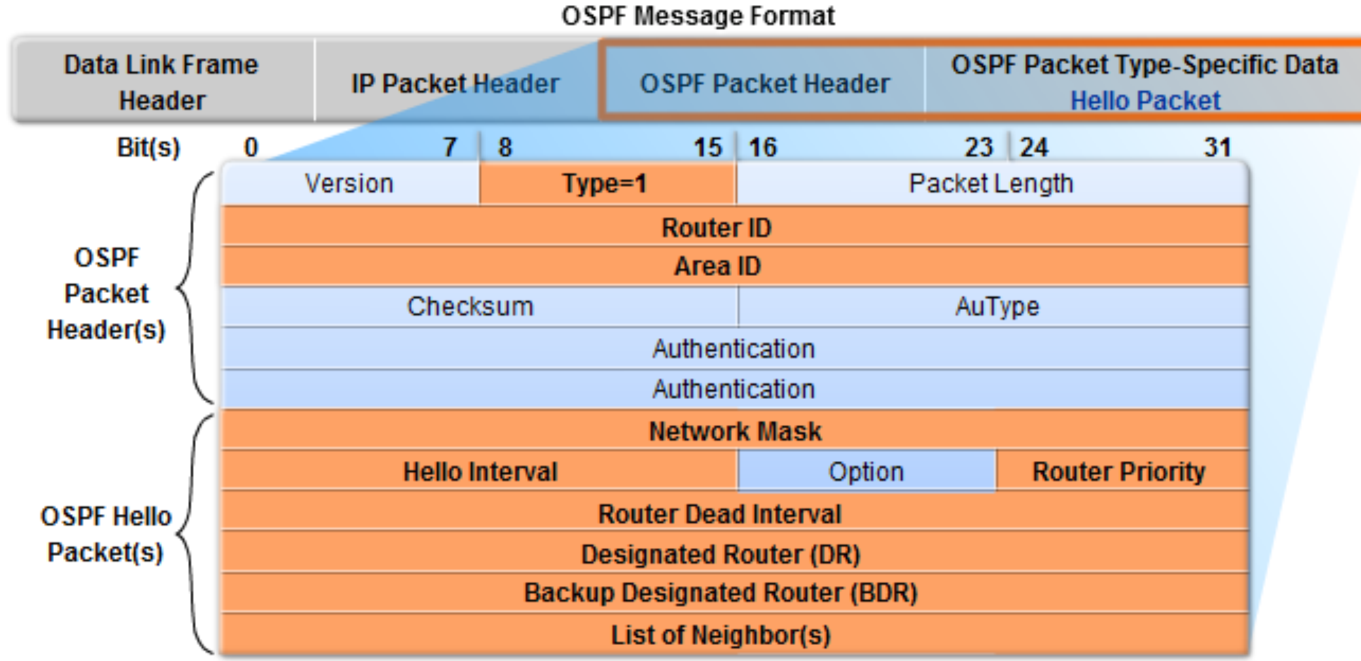

- **Type:** OSPF Packet Type: Hello (1), DD (2), LS Request (3), LS Update (4), LS ACK (5)
- **Router ID:** ID of the originating router
- **Area ID:** area from which the packet originated
- **Network Mask:**Subnet mask associated with the sending interface
- **Hello Interval:**number of seconds between the sending router's hellos
- **Router Priority:** Used in DR/BDR election (discussed later)
- **Designated Router (DR):**Router ID of the DR, if any
- **Backup Designated Router (BDR):**Router ID of the BDR, if any
- **List of Neighbors:** lists the OSPF Router ID of the neighboring router(s)

More in later

Hello packets :

- Discover neighbors (OSPF neighbors)
- Establish adjacencies
- Advertise parameters on which two routers must agree to become neighbors
  - Hello Interval, Dead Interval, Network Type
- Elect the Designated Router and Backup Designated Router on ***multiaccess networks*** such as Ethernet and Frame Relay

Hello Protocol

These will be discussed throughout this chapter.

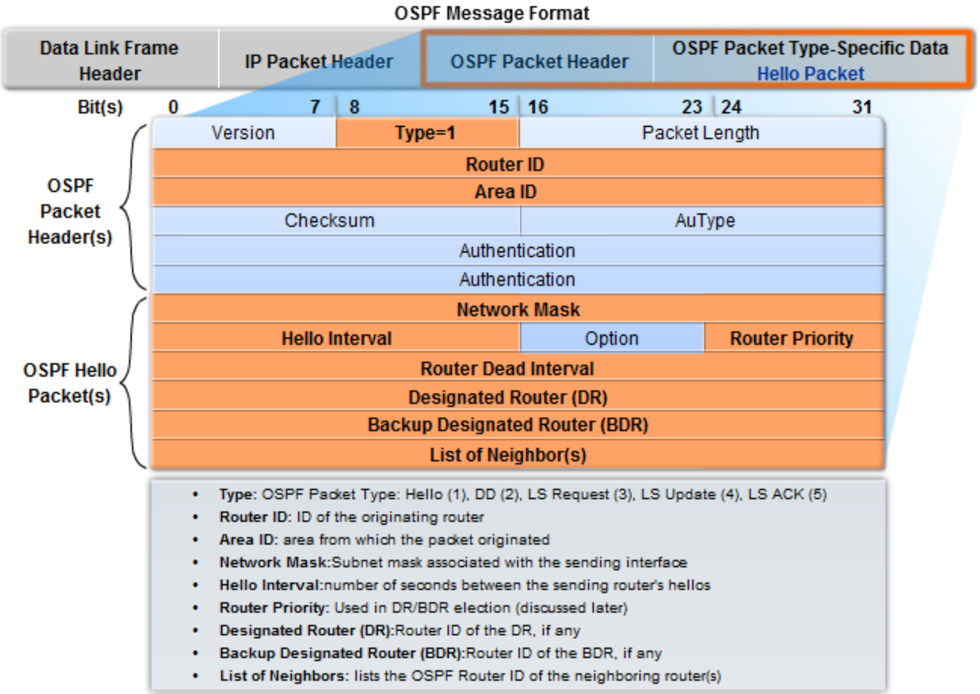

- **Type:** OSPF packet type: Hello (Type 1), DBD (Type 2), LS Request (Type 3), LS Update (Type 4), LS ACK (Type 5)
- **Router ID:** ID of the originating router
- **Area ID:** Area from which the packet originated
- **Network Mask:** Subnet mask associated with the sending interface
- **Hello Interval:** Number of seconds between the sending router's Hellos
- **Router Priority:** Used in DR/BDR election (discussed later)
- **Designated Router (DR):** Router ID of the DR, if any
- **Backup Designated Router (BDR):** Router ID of the BDR, if any
- **List of Neighbors:** Lists the OSPF Router ID of the neighboring router(s)

Neighbor Establishment, OSPF Hello and Dead Intervals

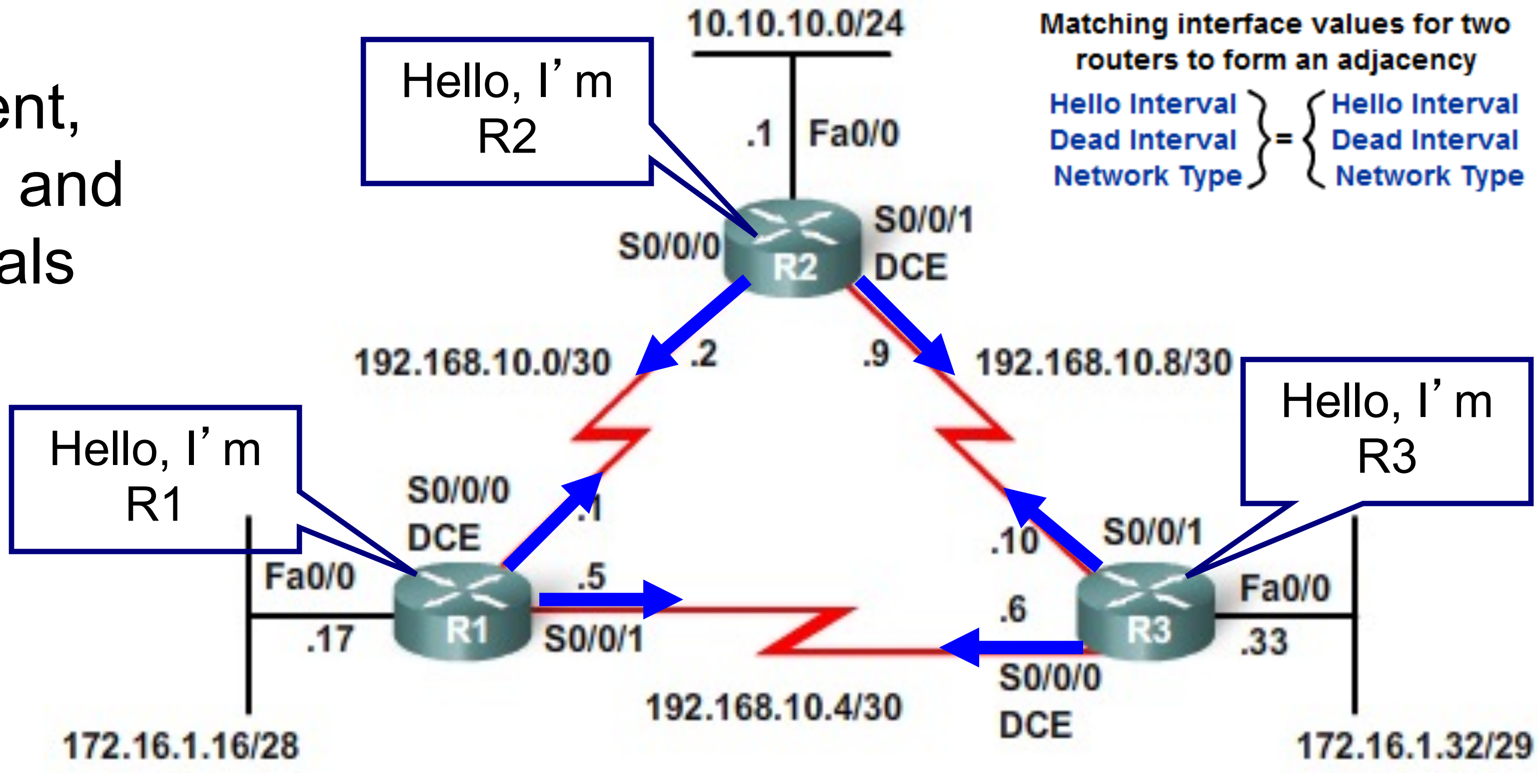

More later

- Before an OSPF router can flood its link states, must discover neighbors.
- **Before two routers can form an OSPF neighbor adjacency, they must agree on three values:**
  - Hello interval
  - Dead interval
  - Network type
- Both the interfaces must be part of the same network, including having the same subnet mask.

Hello Intervals

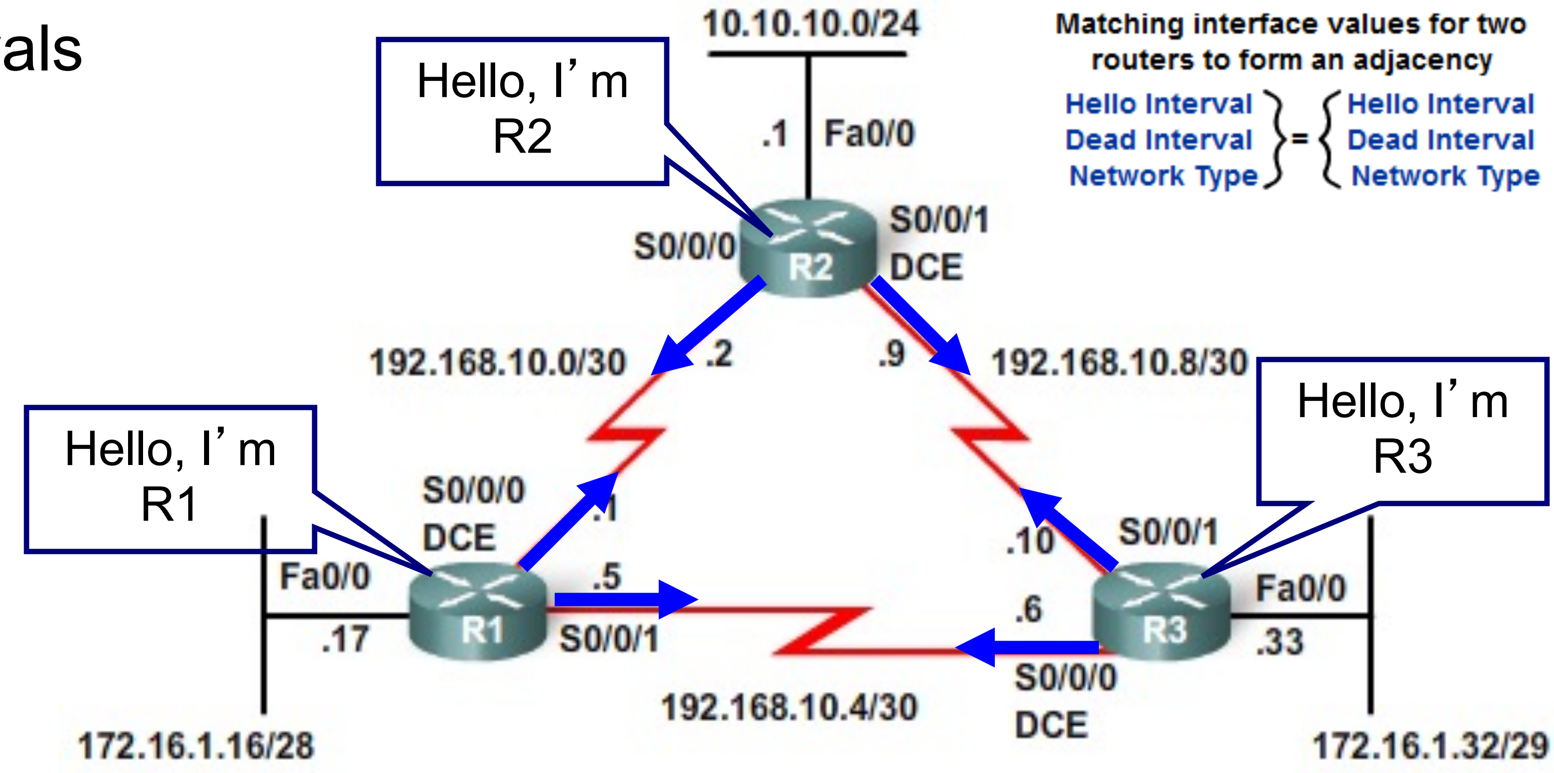

- By default, OSPF **Hello packets** are sent:
  - **10 seconds** on multiaccess and point-to-point segments
  - **30 seconds** on *nonbroadcast multiaccess (NBMA)* segments (Frame Relay, X.25, ATM).
- Sent to *ALLSPFRouters* at 224.0.0.5

Dead Intervals

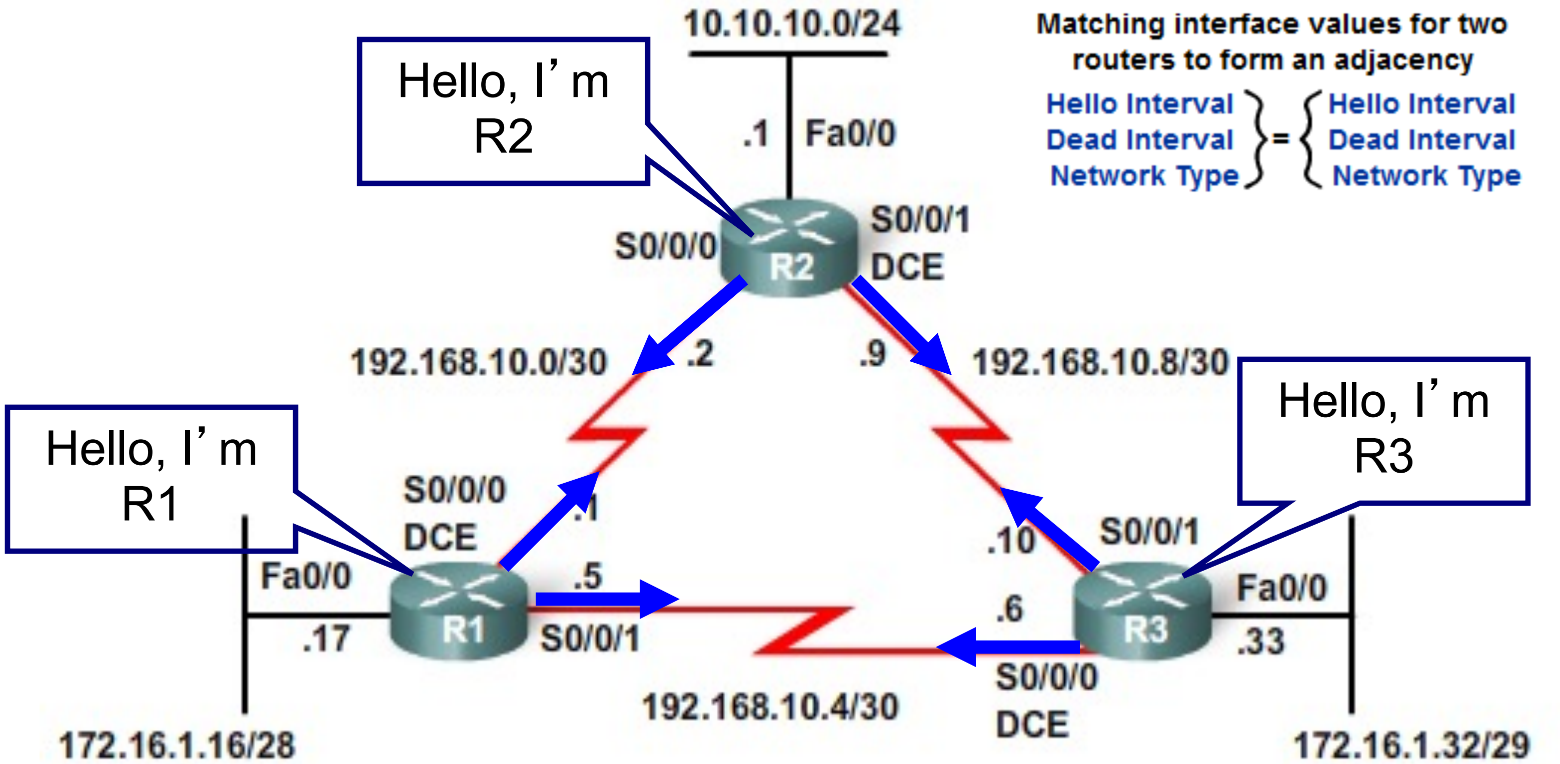

- **Dead interval** - Period, expressed in seconds, that the router will wait to receive a Hello packet before declaring the neighbor "down."
- **Cisco** uses a default of four times the Hello interval.
  - **40 seconds** - Multiaccess and point-to-point segments.
  - **120 seconds** - NBMA networks.
- Dead interval expires
  - OSPF removes that neighbor from its link-state database.
  - Floods the link-state information about the "down" neighbor out all OSPF-enabled interfaces.
- *Network types are discussed later in the chapter.*

Electing a DR and BDR

More later

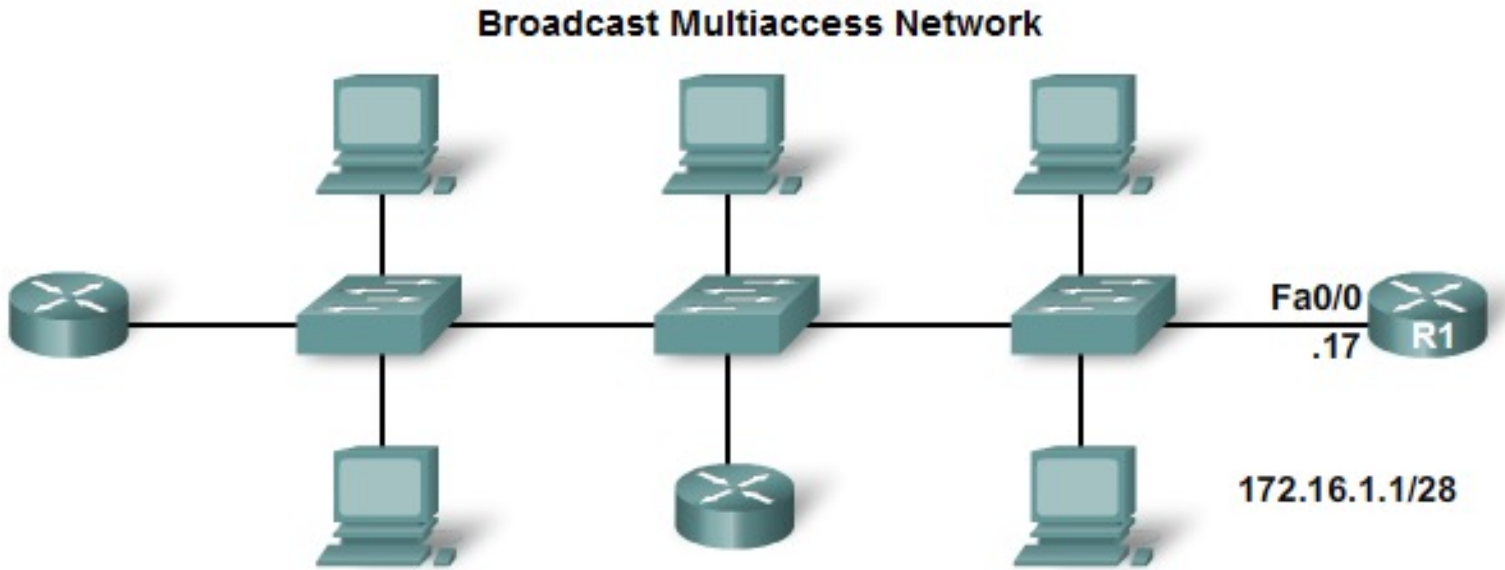

- Election of ***Designated Router (DR)*** and ***Backup Designated Router (BDR)***.
  - Used to reduce the amount of OSPF traffic on multiaccess networks
  - **DR** is responsible for updating all other OSPF routers.
  - **BDR** is the backup if the current DR fails.

OSPF LSUs

LSUs Contain Link-State Advertisements (LSAs)

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between router
3	LSR	Requests specific link-state records from router to router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

The acronyms LSA and LSU are often used interchangeably.

An LSU contains one or more LSAs.

LSAs contain route information for destination networks.

LSA specifics are discussed in CCNP.

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous System Extrenal LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gatway Protocol(BGP)
9,10,11	Opaque LSAs

- **Link-State Updates (LSU)** are the packets used for OSPF routing updates.
  - Can contain 11 different types of **LSAs (Link-State Advertisements)** (CCNP)
- At times, these terms are used interchangeably.

OSPF Algorithm

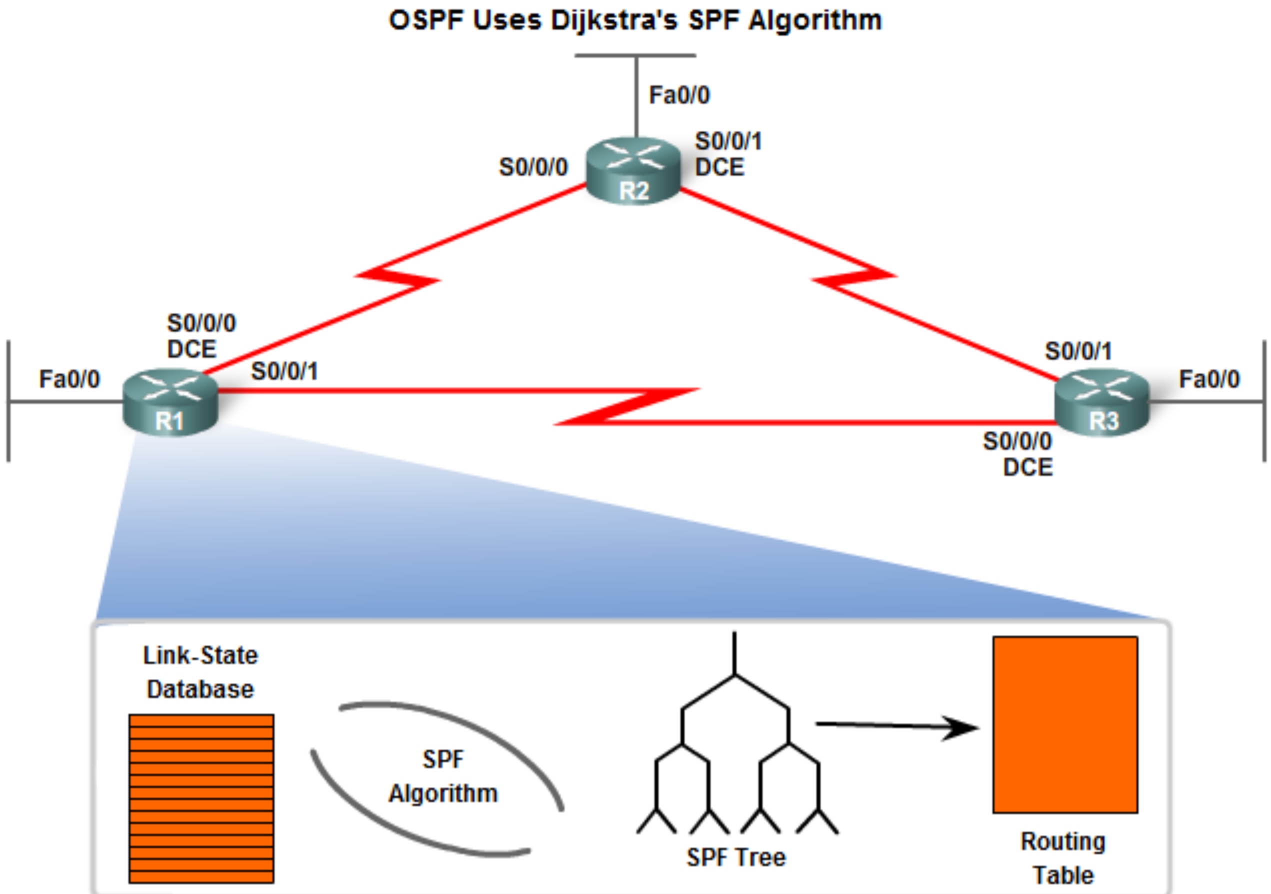

Administrative Distance

- Administrative distance (AD) is the trustworthiness (or preference) of the route source.
- **OSPF** has a default **AD of 110.**

Route Source	AD
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200

Authentication

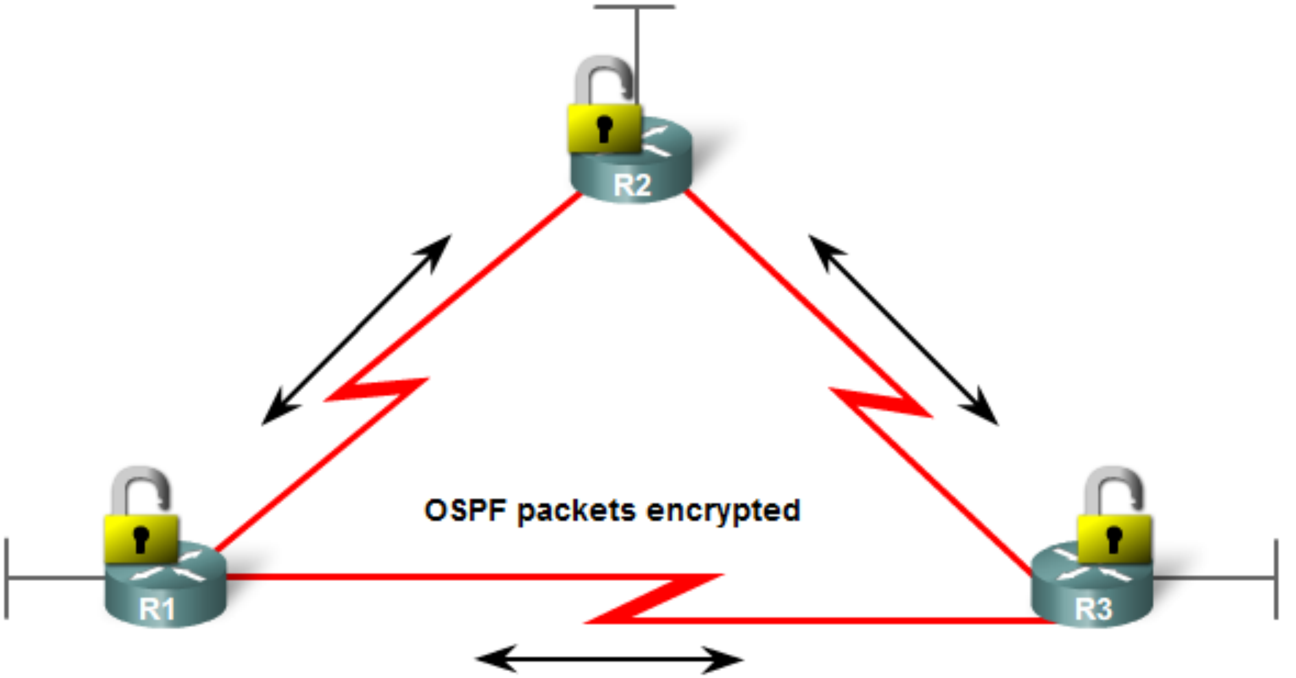

- OSPF can be configured to authenticate OSPF messages.

OSPF Operations

Steps to OSPF Operation with States

1. **Establishing router adjacencies (Routers are adjacent)**
   - Down State – No Hello received
   - Init State – Hello received, but not with this router's Router ID
     - "Hi, my name is Carlos."
     - "Hi, my name is Maria."
   - Two-way State – Hello received, and with this router's Router ID
     - "Hi, Maria, my name is Carlos."
     - "Hi, Carlos, my name is Maria."
2. **Electing DR and BDR – Multi-access (broadcast) segments only**
   - ExStart State with DR and BDR
   - Two-way State with all other routers
3. **Discovering Routes**
   - ExStart State
   - Exchange State
   - Loading State
   - Full State (Routers are "fully adjacent")

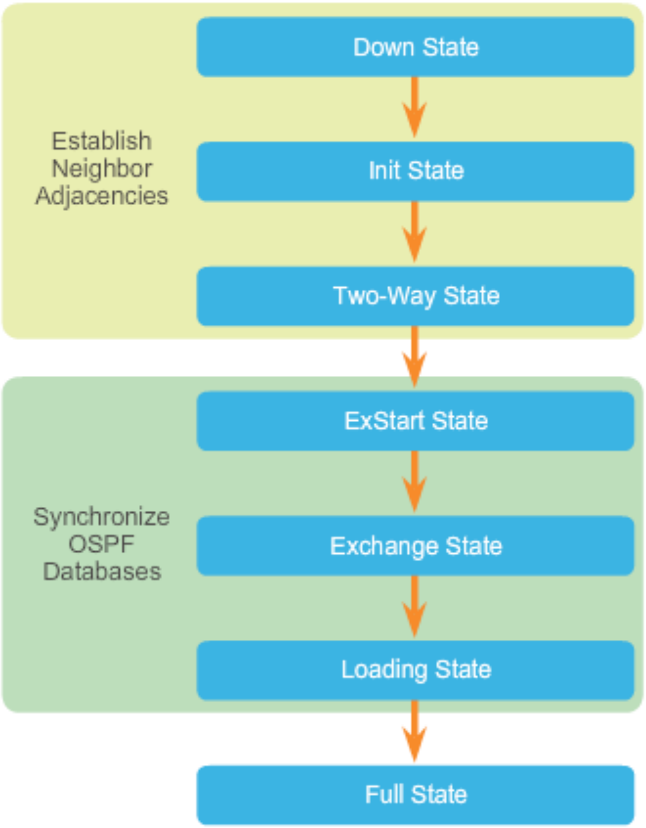

4. **Calculating the Routing Table**
5. **Maintaining the LSDB and Routing Table**

OSPF Packet Types

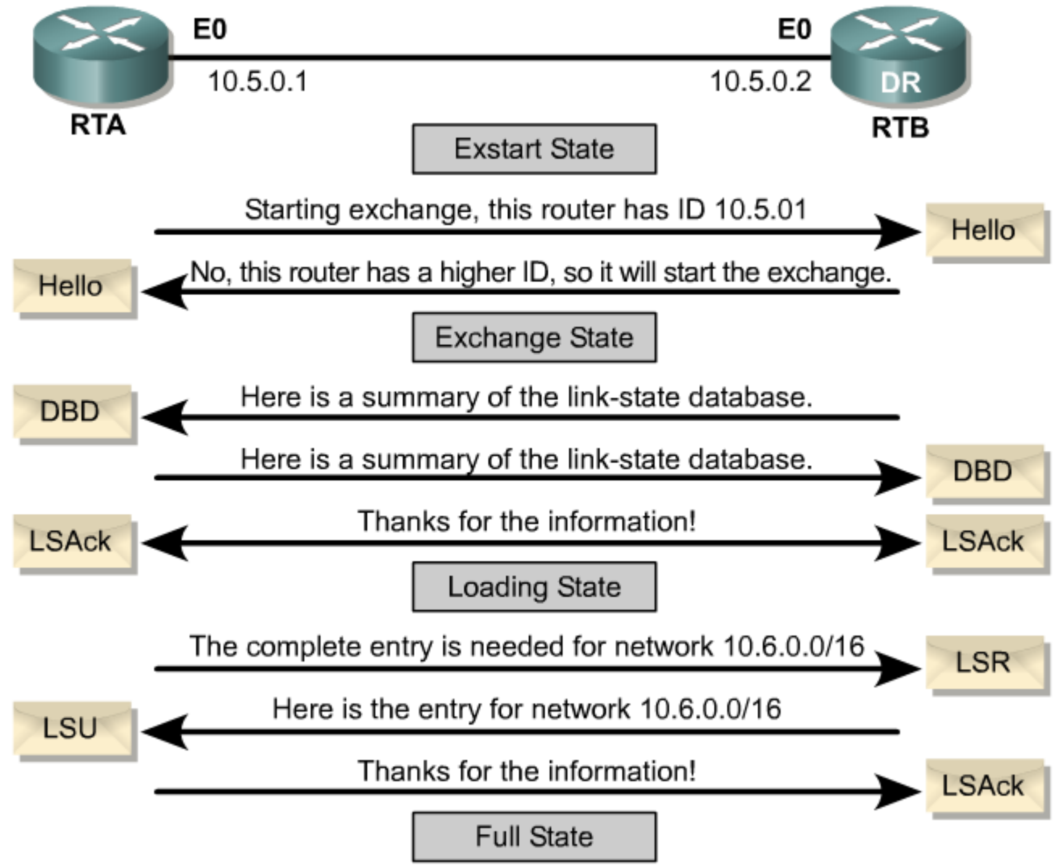

Figure includes CCNP information

- Five types of OSPF LSPs (link-state packets).
  - **Hello**: Used to establish and maintain adjacency.
  - **DBD (Database Description):** Abbreviated list of the sending router's link-state database.
  - **LSR (Link-State Request)** : Used by routers to request more information about any entry in the DBD.
  - **LSU**: (**Link-State Update):** Link-state information.
  - **LSAck (LSA Acknowledgment)**: Router sends a link-state (LSAck) to confirm receipt of the LSU.

Establishing Adjacencies (FYI)

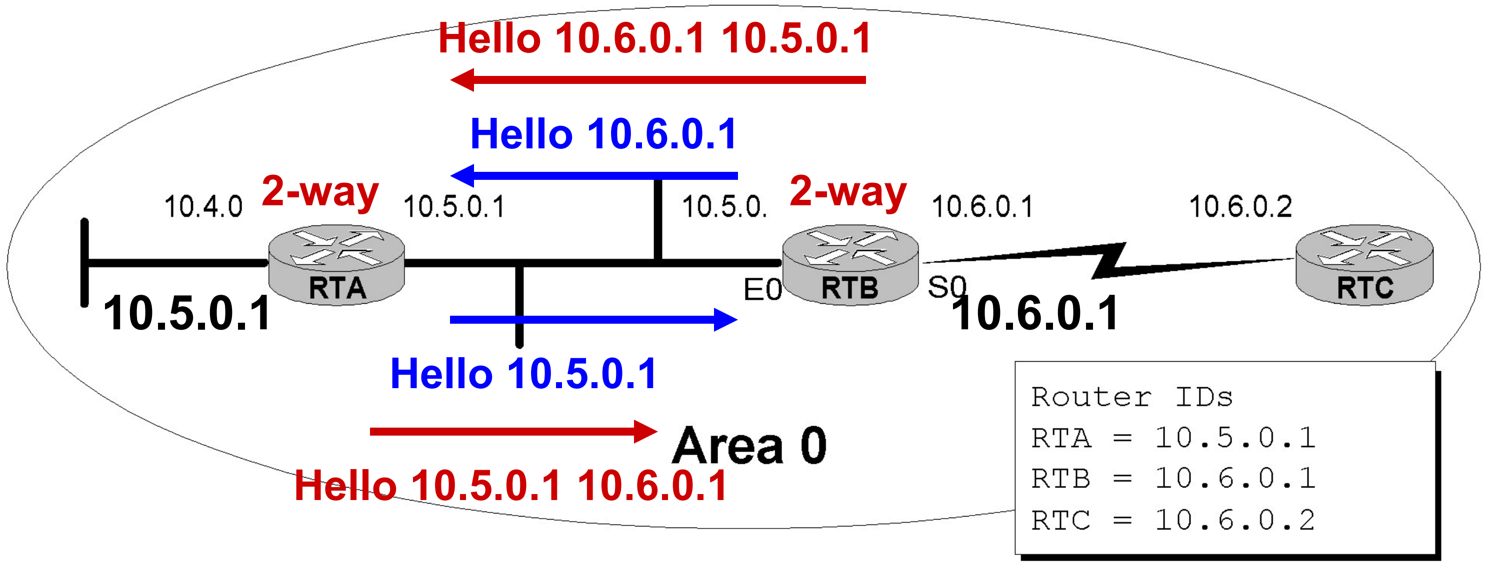

Down State - Init State – Two Way State

- **Down State** - OSPF routers send Hello packets at regular intervals (10 sec.) to establish neighbors.
- When a router (sends or) receives its first **Hello packet**, it enters the **init state**.
  - Hello packet contains a list of known neighbors.
- When the router sends a Hello packet (unicast reply) to the neighbor with its RouterID and the neighbor sends a Hello packet packet back with that Router ID, the router's interface will transition to the **two-way state**.
- Now, the router is ready to take the relationship to the next level.

Configuring Single Area OSPFv2

OSPF Reference Topology

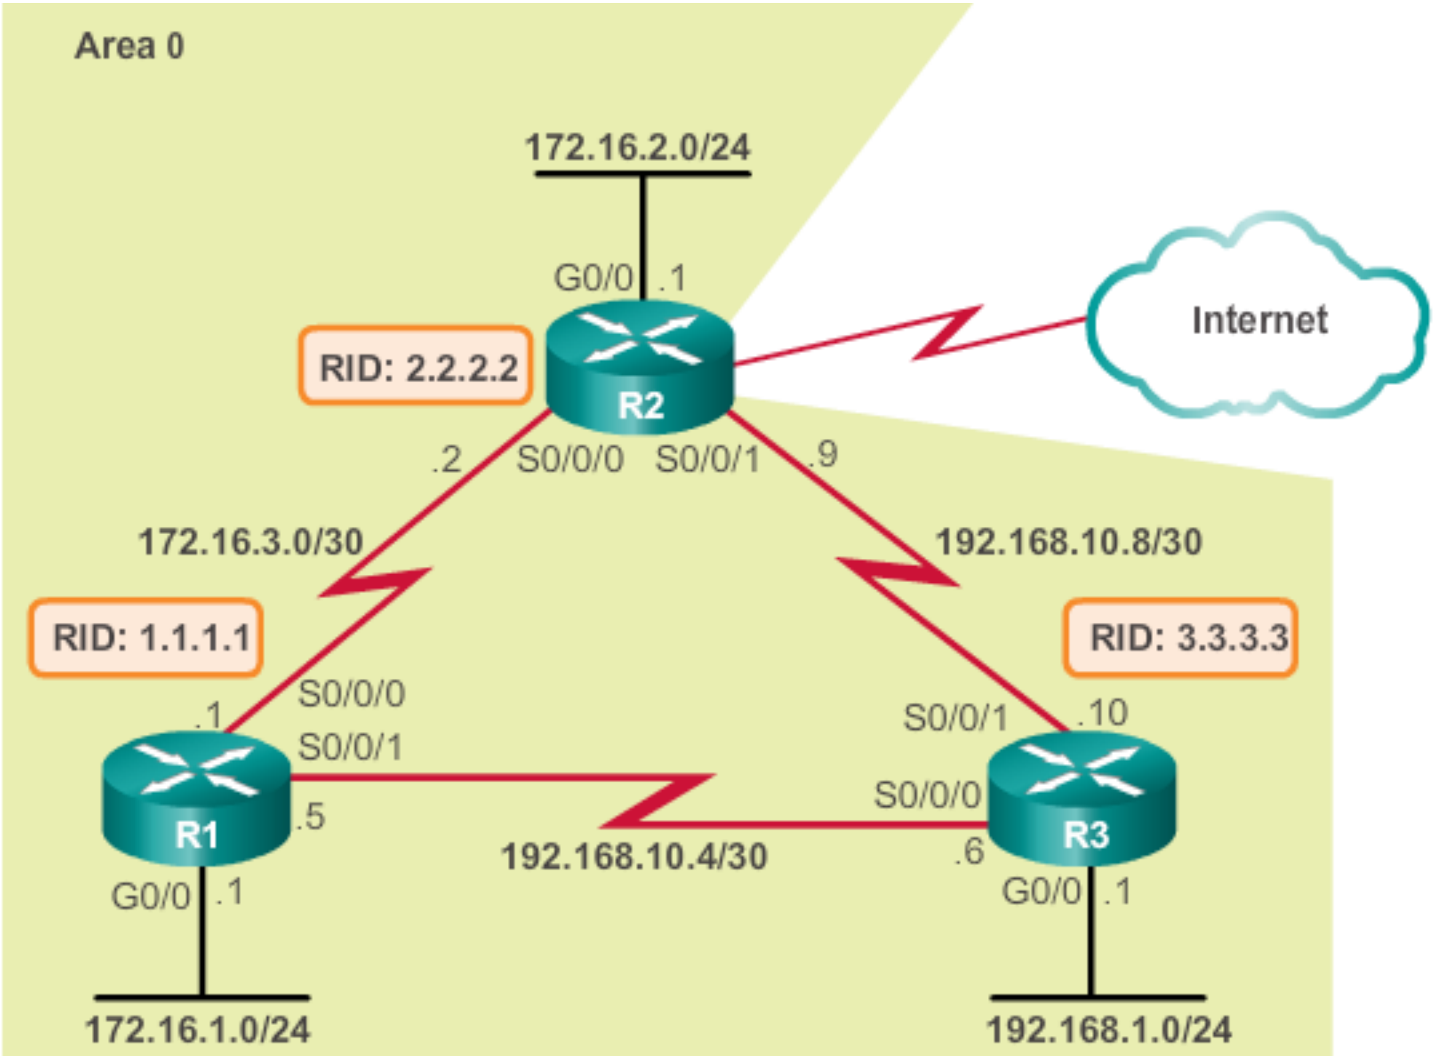

The router ospf Command

```
R1(config)# router ospf 10
R1(config-router)#
```

- The ***process-id***
  - Between 1 and 65,535
  - Chosen by the network administrator.
- Locally significant:
  - Does not have to match other OSPF routers.
  - This differs from EIGRP.
- We are using the same process ID simply for consistency.

OSPF Router ID

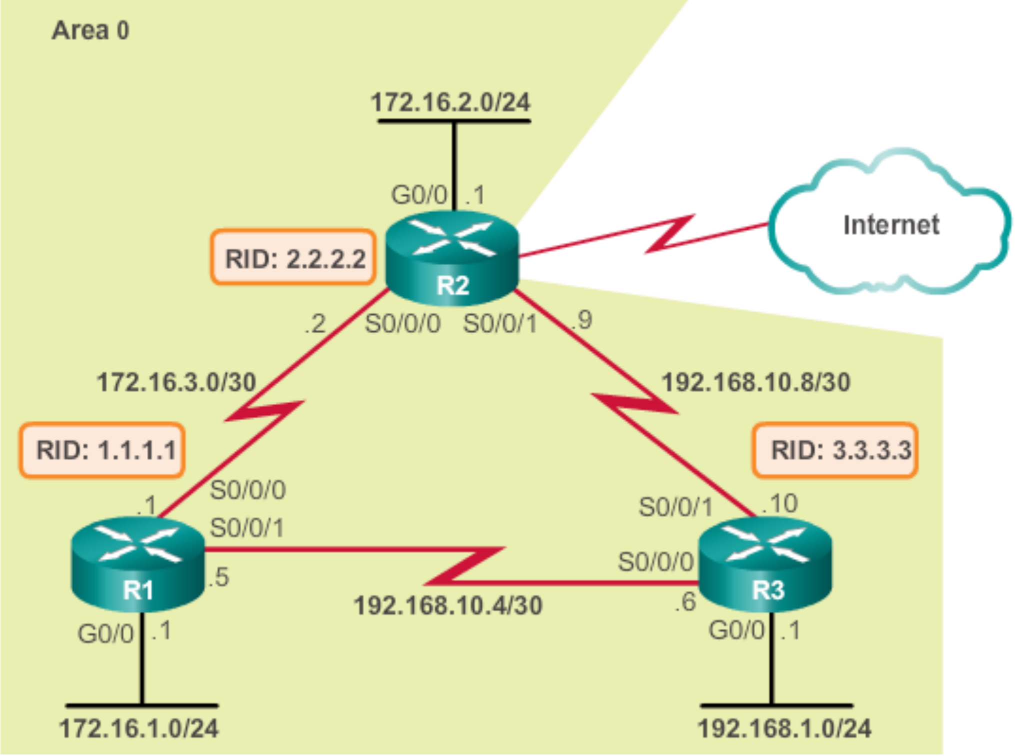

- A router is known to OSPF by the OSPF router ID number.
  - LSDBs use the OSPF router ID to differentiate one router from the next.
- By default, the router ID is the <u>highest IP address on an active interface</u> at the moment of OSPF process startup.
- However, for stability reason, it is recommended that the `router-id` command or a loopback interface be configured.

OSPF Router ID

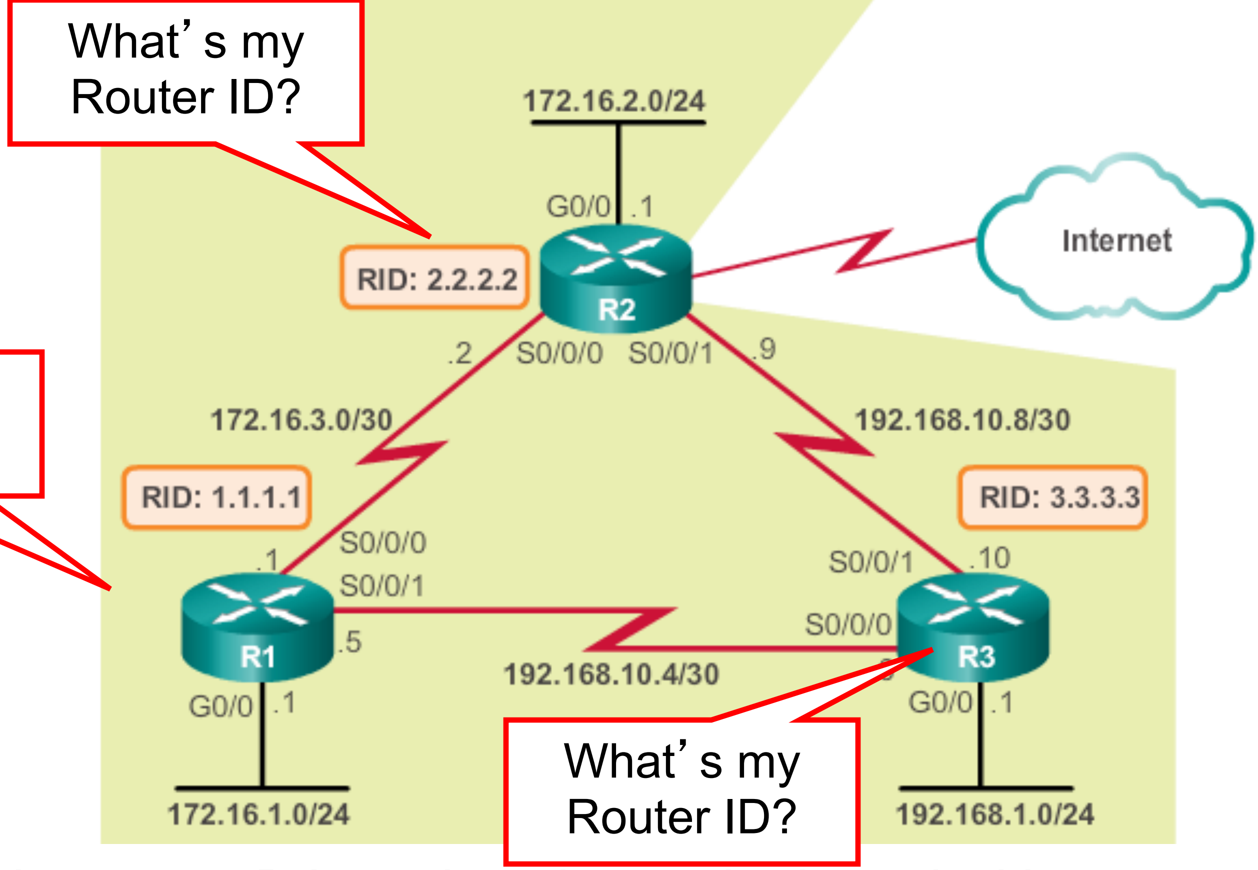

- Cisco routers derive the router ID based on three criteria and with the following precedence:
  1. IP address configured with the **OSPF `router-id` command**.
  2. Highest IP address of any of its **loopback interfaces.**
  3. Highest *active* IP address of any of its physical interfaces.
    - The interface does not need to be enabled for OSPF, i.e. it does not need to be included in one of the OSPF **`network`** commands.

Define the Router ID

- Assign a specific router ID to the router.

```
Router(config)# router ospf process-id
Router(config-router)# router-id ip-address
```

- Any unique arbitrary 32-bit value in an IP address format (dotted decimal) can be used.
- If this command is used on an OSPF process that is already active, then the new router ID takes effect:
  - After the next router reload.
  - After a manual restarting of the OSPF process using the **`clear ip ospf process`** privileged EXEC command.

```
Router# clear ip ospf process
```

Define the Router ID

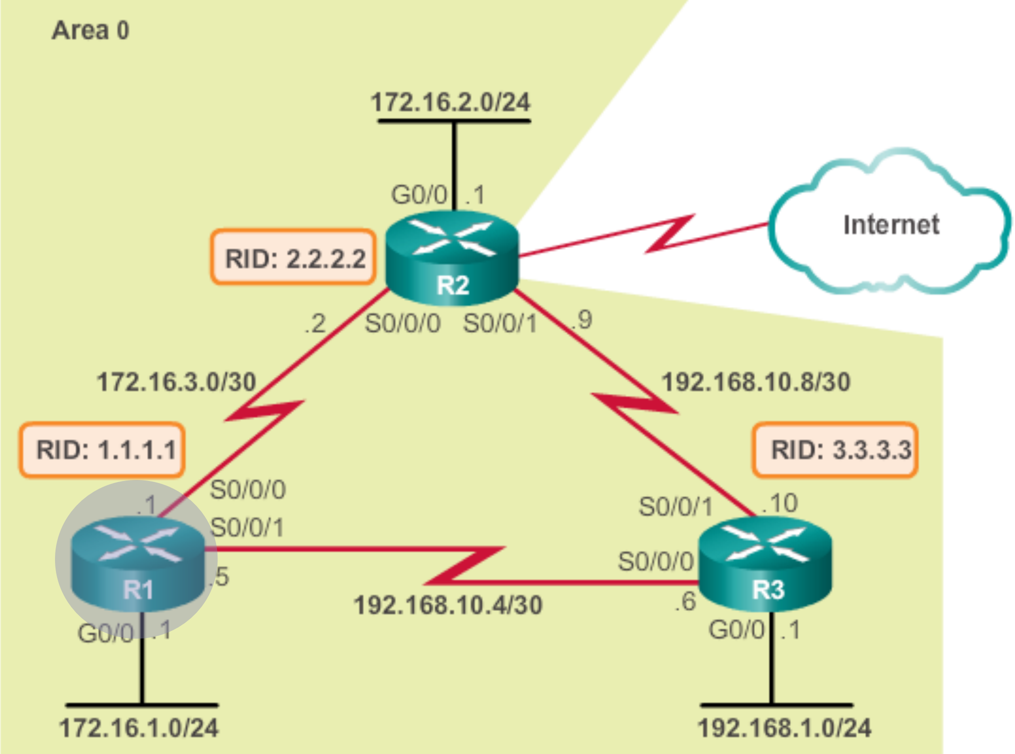


```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# end
R1#
*Mar 25 19:50:36.595: %SYS-5-CONFIG_I: Configured from
console by console
R1#
```

Define the Router ID

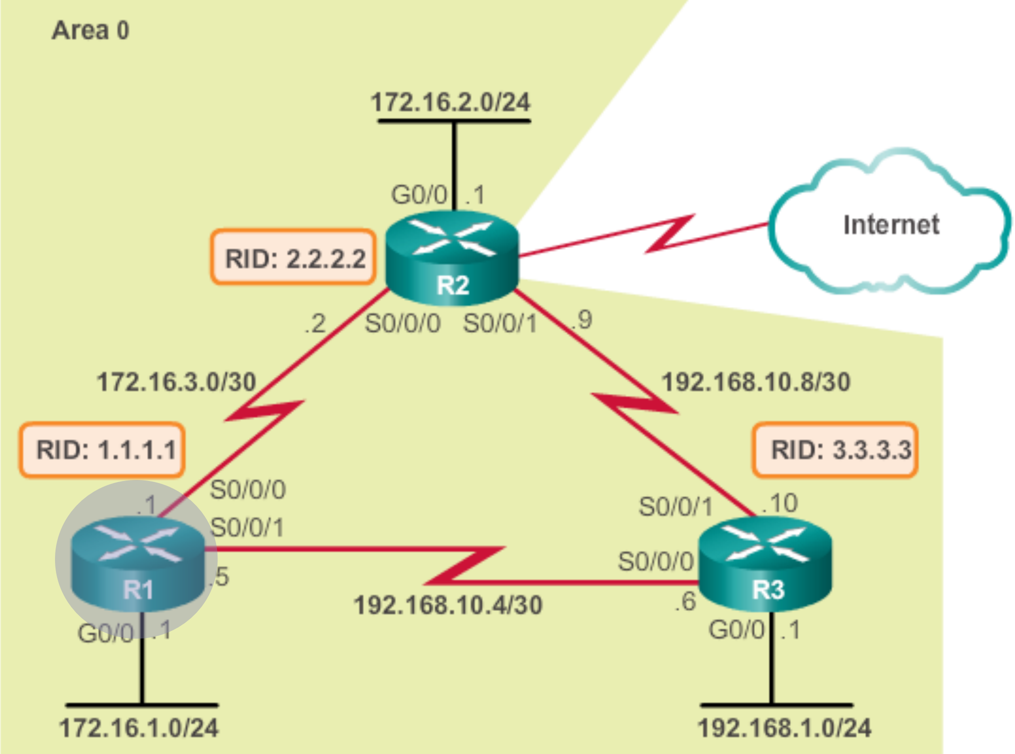


```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 0. 0 normal 0 stub 0 nssa
  Maximum path: 4
 <Output omitted>
```

Changing the OSPF Router-ID and Clearing the OSPF Process

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
% OSPF: Reload or use "clear ip ospf process" command, for this to
take effect
R1(config-router)# end
R1#
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R1#
*Mar 25 19:46:22.423: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
Serial0/0/1 from FULL to DOWN, Neighbor Down: Interface down or
detached
*Mar 25 19:46:22.423: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
Serial0/0/0 from FULL to DOWN, Neighbor Down: Interface down or
detached
*Mar 25 19:46:22.475: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3 on
Serial0/0/1 from LOADING to FULL, Loading Done
*Mar 25 19:46:22.475: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
Serial0/0/0 from LOADING to FULL, Loading Done
R1#
R1# show ip protocols | section Router ID
  Router ID 1.1.1.1
R1#
```

Alternative: Configure a Loopback

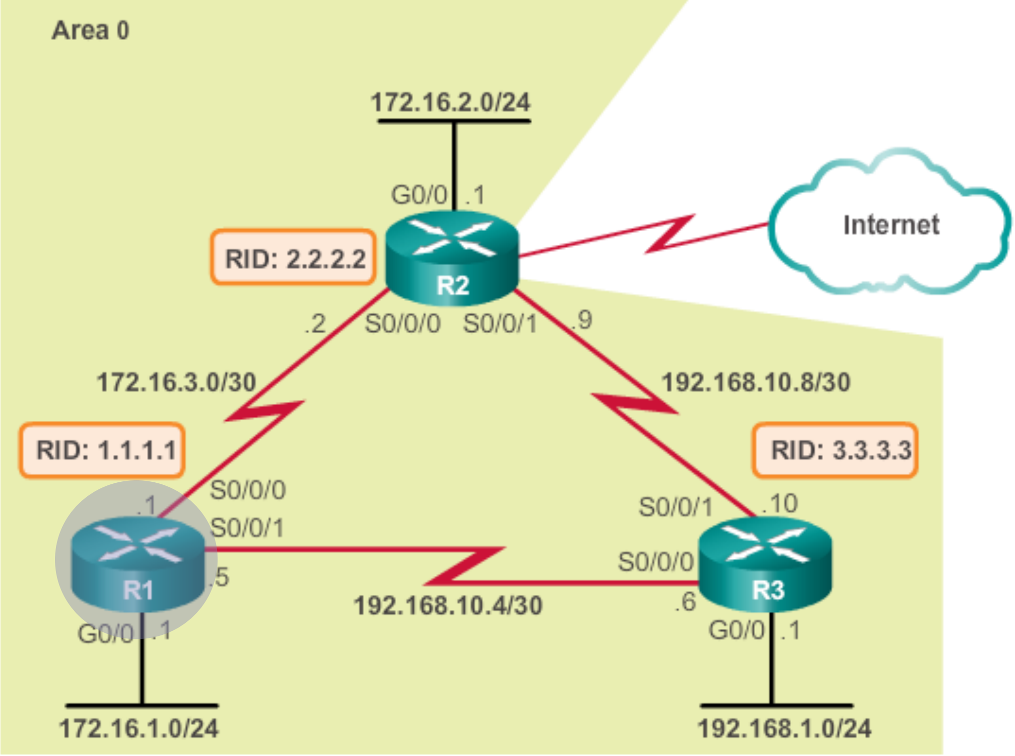


```
R1(config)# interface loopback 0
R1(config-if)# ip address 1.1.1.1 255.255.255.255
R1(config-if)# end
R1#
```

The network Command

```
Router(config-router)# network network-address wildcard-mask area area-id
```

- The **network** command (same function as when used with other IGP routing protocols)
  - Any interfaces on a router that match the network address in the **network** command will be enabled to send and receive OSPF packets.
  - This network (or subnet) will be included in OSPF routing updates.
- Requires the wildcard mask.
  - Used to specify the interface or range of interfaces enabled for OSPF.

The network Command

```
Router(config-router)# network network-address wildcard-mask area area-id
```

```
  255.255.255.255
- 255.255.255.240 Subtract the subnet mask
---------------
    0.   0.   0. 15 Wildcard mask
```

- The wildcard mask can be configured as the inverse of a subnet mask.
- Note:
  - Like EIGRP, some Cisco IOS software versions allow you to simply enter the subnet mask instead of the wildcard mask.
  - The Cisco IOS software then converts the subnet mask to the wildcard mask format.

The network Command

```
Router(config-router)# network network-address wildcard-mask area area-id
```

- The **area** *area-id* refers to the ***OSPF area***.
- A group of routers that <u>share link-state information</u>.
  - Identical <u>link-state databases</u>.
- In this chapter, we configure all the OSPF routers within a single area.
  - This is known as ***single-area OSPF***.
- The **network** commands must be configured with the <u>same area ID on all routers.</u>
  - <u>Good practice</u> to use an <u>area ID of 0</u> with single-area OSPF.

Advertising OSPF Networks

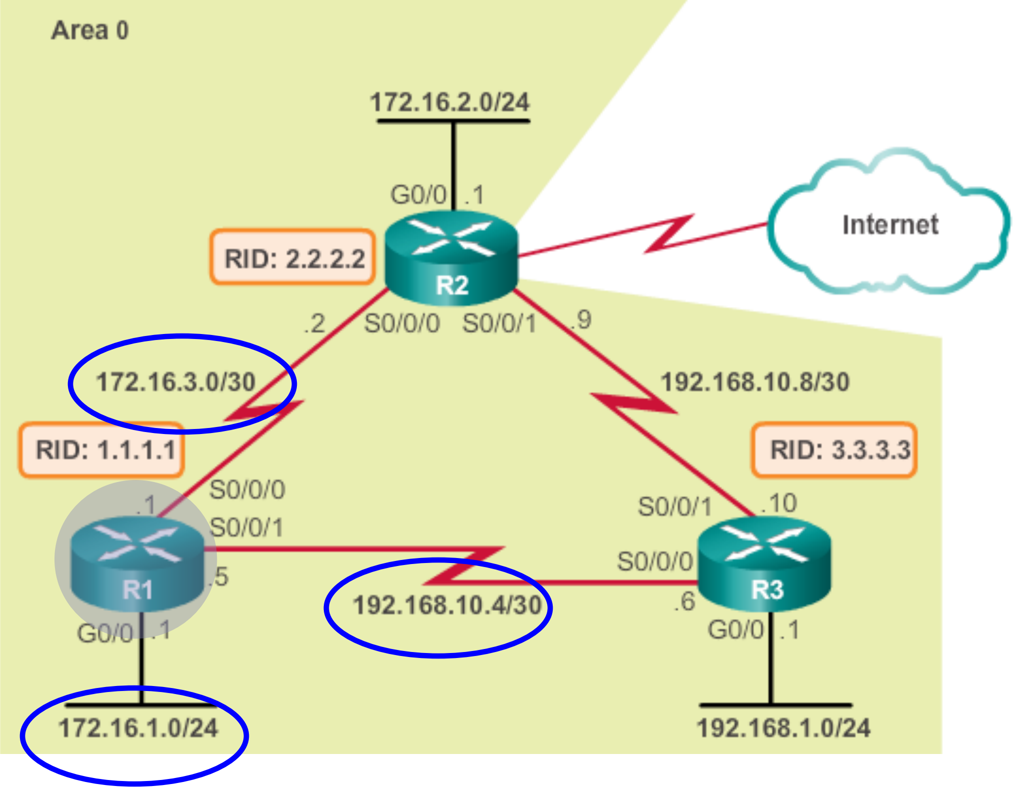


```
R1(config)# router ospf 10
R1(config-router)# route-id 1.1.1.1
R1(config-router)# network 172.16.1.0 0.0.0.255 area 0
R1(config-router)# network 172.16.3.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.4 0.0.0.3 area 0
R1(config-router)# end
R1#
```

Advertising OSPF Networks

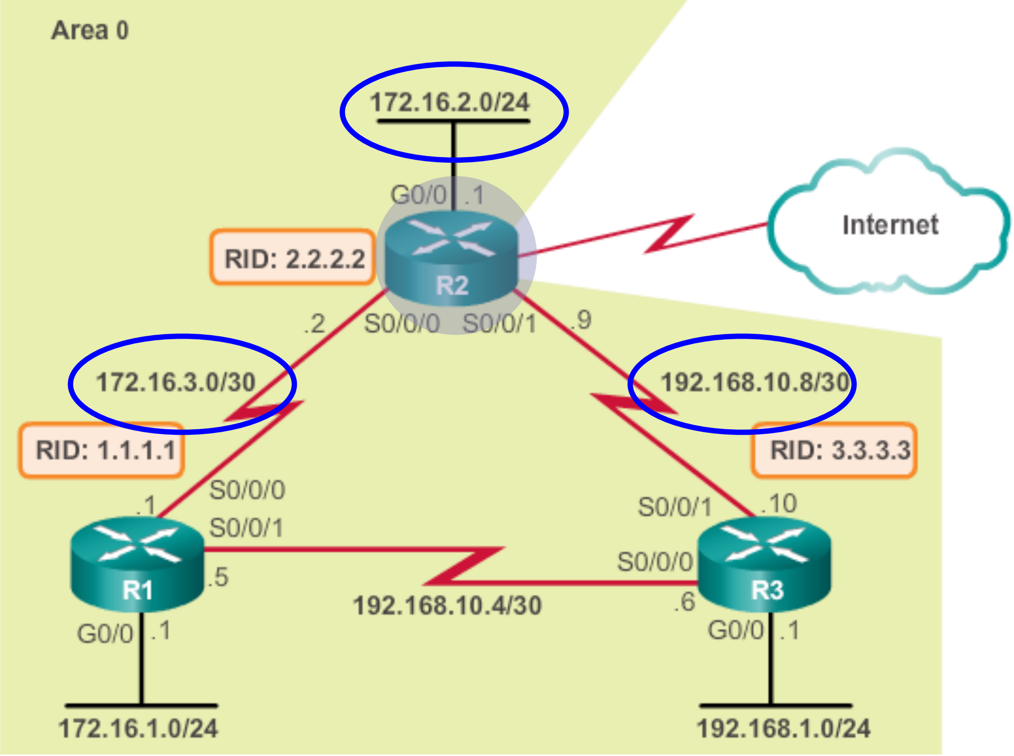


```
R2(config)# router ospf 10
R2(config-router)# route-id 2.2.2.2
R2(config-router)# network 172.16.2.0 0.0.0.255 area 0
R2(config-router)# network 172.16.3.0 0.0.0.3 area 0
R2(config-router)# network 192.168.10.8 0.0.0.3 area 0
R2(config-router)#
*Mar 25 21:19:21.938: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1
on Serial0/0/0 from LOADING to FULL, Loading Done
R2(config-router)# end
R2#
```

Advertising OSPF Networks

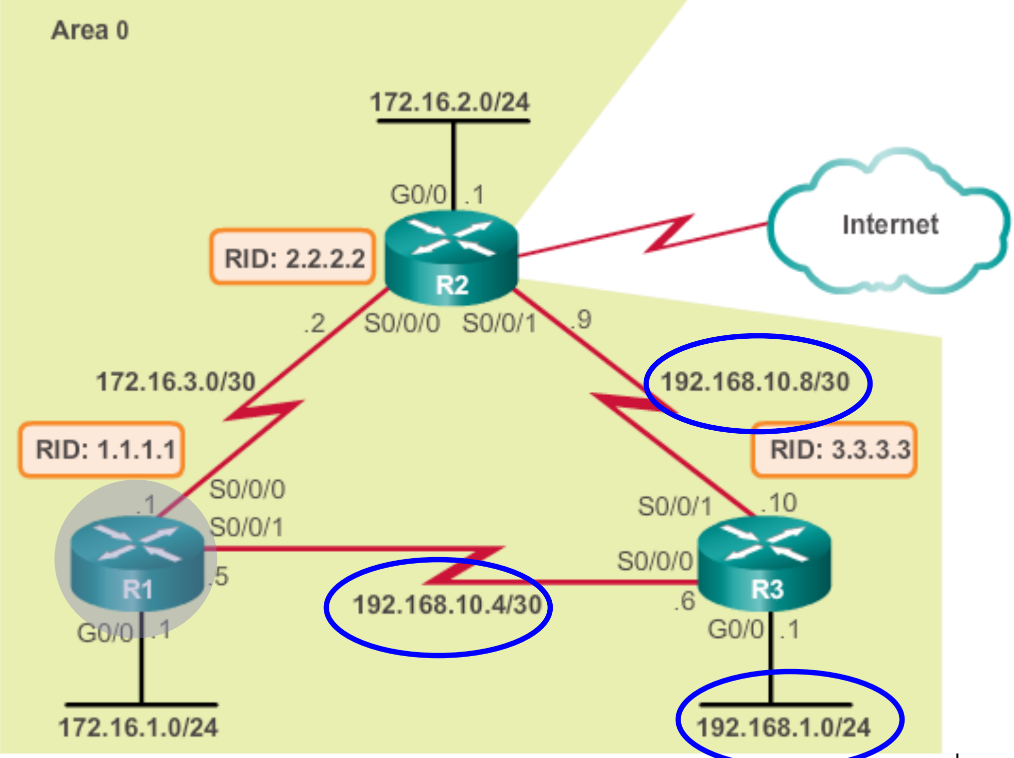


```
R3(config)# router ospf 10
R3(config-router)# router-id 3.3.3.3
R3(config-router)# network 192.168.1.1 0.0.0.0 area 0
R3(config-router)# network 192.168.10.6 0.0.0.0 area 0
R3(config-router)# network 192.168.10.10 0.0.0.0 area 0
R3(config-router)#
*Mar 26 14:00:55.183: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1
on Serial0/0/0 from LOADING to FULL, Loading Done
*Mar 26 14:00:55.243: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2
on Serial0/0/1 from LOADING to FULL, Loading Done
R3(config-router)# end
R3#
```

Optional Method: Identify OSPF Networks

- No network command is needed.
- Because this command is configured explicitly for the interface, it takes precedence over the **network area** command.

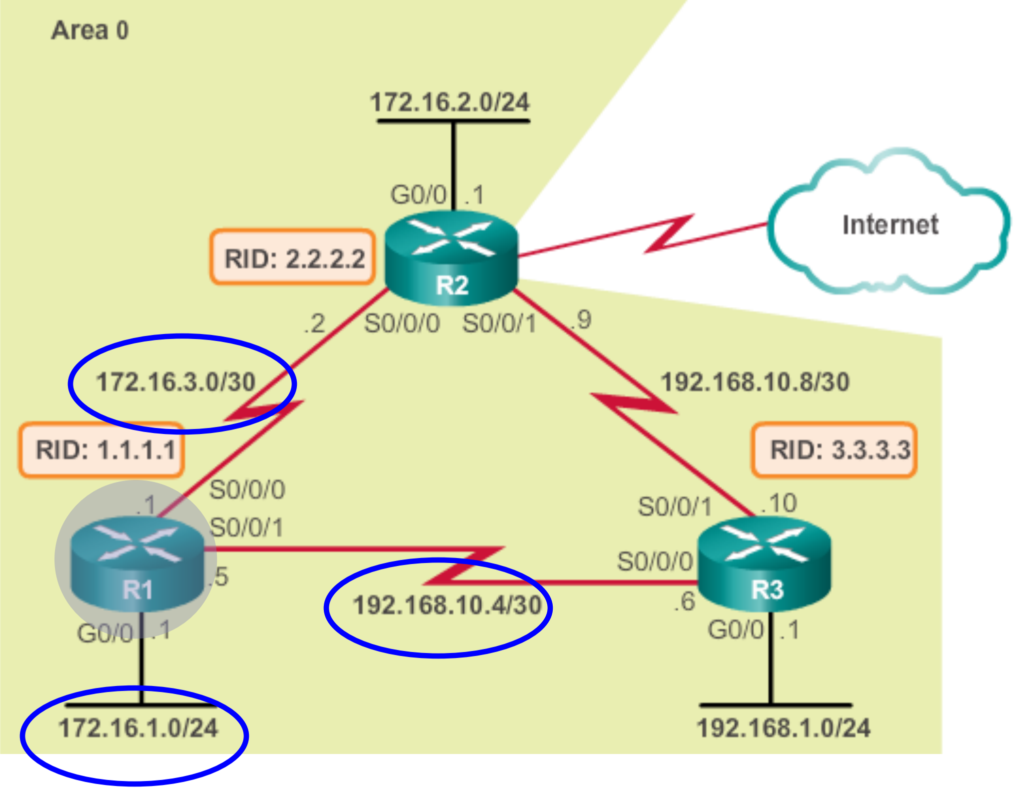


```
R1(config)# interface gig 0/0
R1(config-router)# ip ospf 10 area 0
R1(config-router)# exit
R1(config)# interface serial 0/0/0
R1(config-router)# ip ospf 10 area 0
R1(config-router)# exit
R1(config)# interface serial 0/0/1
R1(config-router)# ip ospf 10 area 0
R1(config-router)# end
R1#
```

Passive Interfaces

OSPF Metric (Cost)

Passive Interface

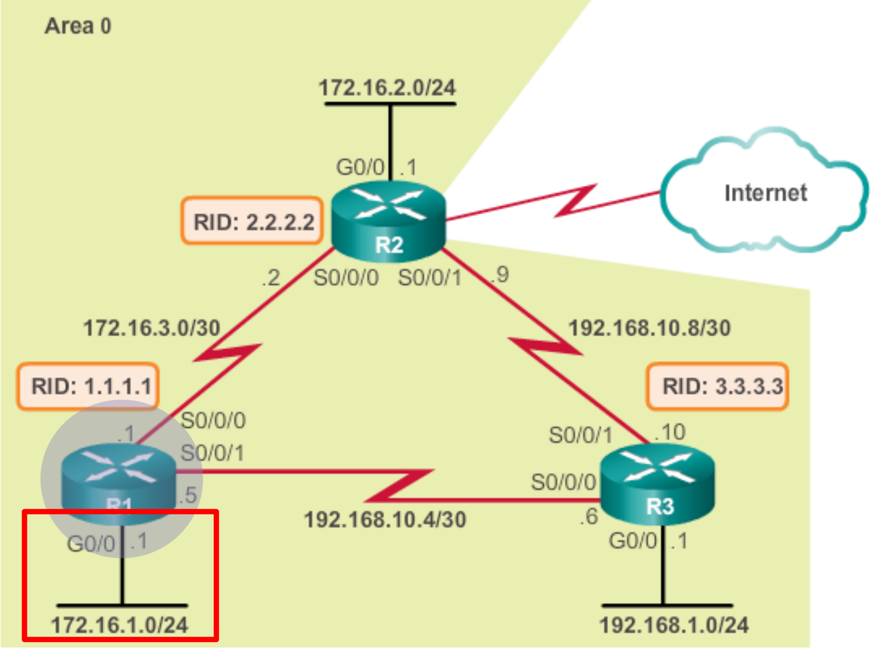

- By default, OSPF messages are forwarded out all OSPF-enabled interfaces.
- Sending out unneeded messages on a LAN affects the network in three ways:
  - Inefficient Use of Bandwidth
  - Inefficient Use of Resources
  - Increased Security Risk
- Interfaces which do not connect to another OSPF neighbor should be rendered as passive.

Configuring Passive Interfaces on R1 & R2

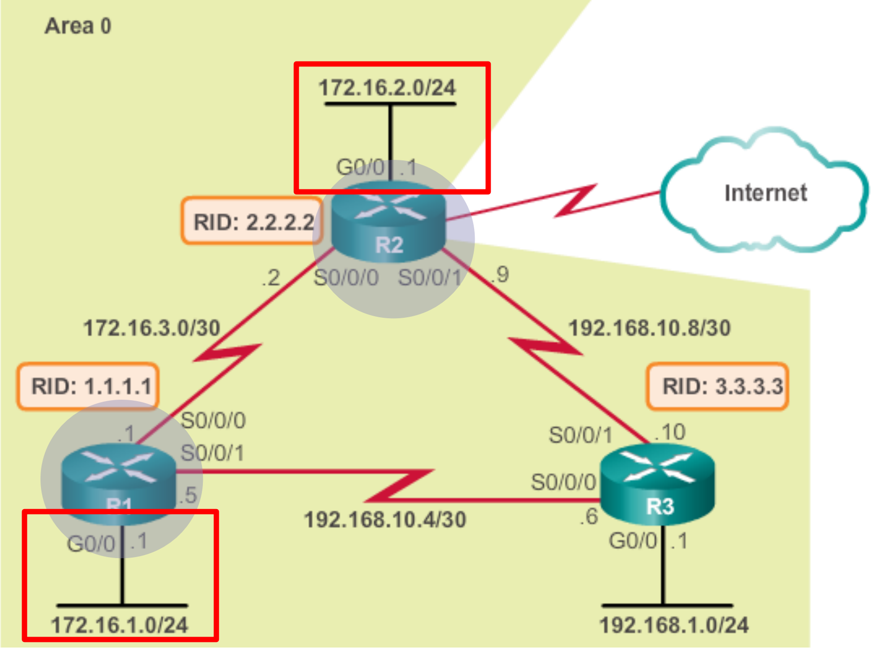


```
R1(config)# router ospf 10
R1(config-router)# passive-interface GigabitEthernet 0/0
R1(config-router)# end
R1#
```

```
R2(config)# router ospf 10
R2(config-router)# passive-interface GigabitEthernet 0/0
R2(config-router)# end
R2#
```

Verifying Passive Interfaces on R1 and R2

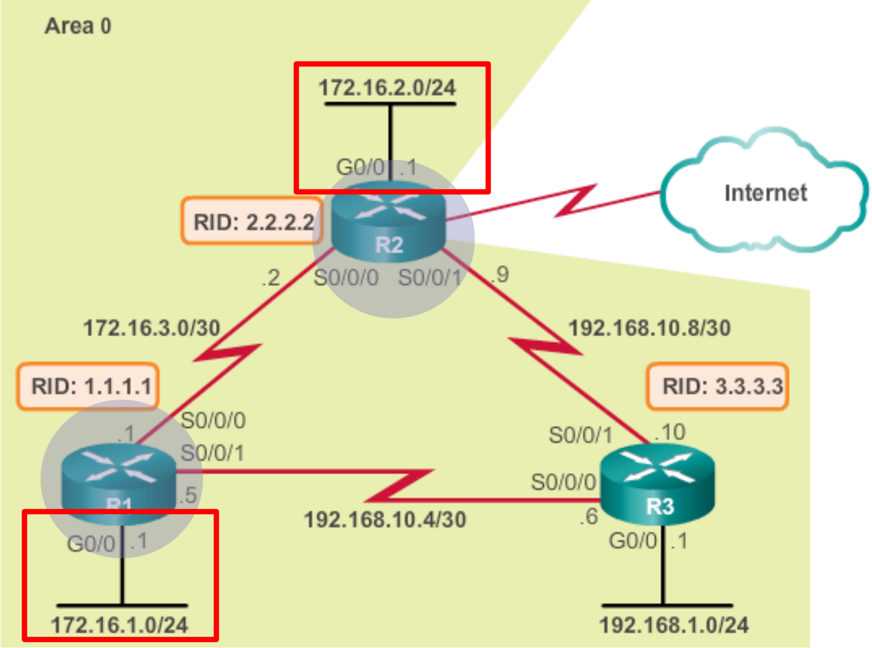


```
R1# show ip protocols | section Passive
Passive Interface(s):
    GigabitEthernet0/0

R1#
```

```
R2# show ip protocols | section Passive
Passive Interface(s):
    GigabitEthernet0/0

R2#
```

Configuring Passive Interfaces on R3

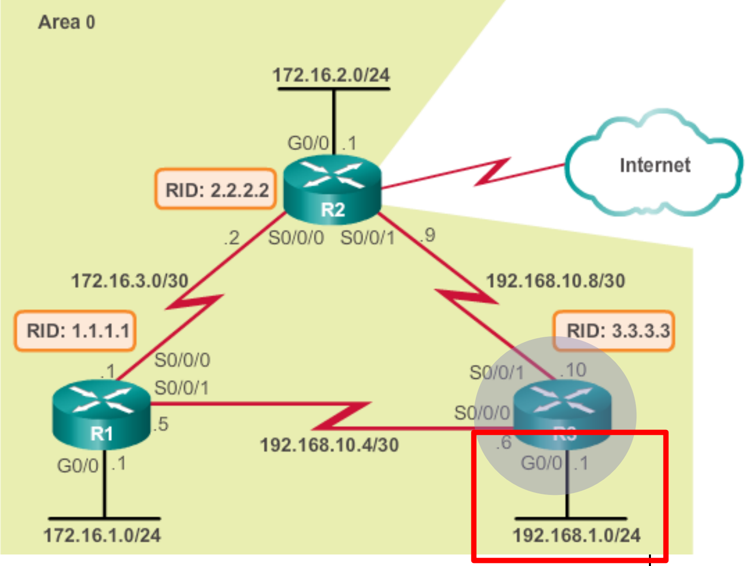


```
R3(config)# router ospf 10
R3(config-router)# passive-interface default
R3(config-router)#
*Apr  7 16:22:58.090: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
Serial0/0/0 from FULL to DOWN, Neighbor Down: Interface down or
detached
*Apr  7 16:22:58.090: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
Serial0/0/1 from FULL to DOWN, Neighbor Down: Interface down or
detached
R3(config-router)#
R3(config-router)# no passive-interface serial 0/0/0
*Apr  7 16:23:18.590: %OSPF-5-ADJCHG: Process 10, Nbr 1.1.1.1 on
Serial0/0/0 from LOADING to FULL, Loading Done
R3(config-router)#
R3(config-router)# no passive-interface serial 0/0/1
R3(config-router)#
*Apr  7 16:23:24.462: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2 on
Serial0/0/1 from LOADING to FULL, Loading Done
```

Network Working Group J. Moy
Request for Comments: 2328 Ascend Communications, Inc.
STD: 54 April 1998
Obsoletes: 2178
Category: Standards Track

OSPF Version 2

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Abstract

This memo documents version 2 of the OSPF protocol. OSPF is a link-state routing protocol. It is designed to be run internal to a single Autonomous System. Each OSPF router maintains an identical database describing the Autonomous System's topology. From this database, a routing table is calculated by constructing a shortest-path tree.

OSPF recalculates routes quickly in the face of topological changes, utilizing a minimum of routing protocol traffic. OSPF provides support for equal-cost multipath. An area routing capability is provided, enabling an additional level of routing protection and a reduction in routing protocol traffic. In addition, all OSPF routing protocol exchanges are authenticated.

OSPF Metric

- The OSPF metric is called **cost**. The following passage is from RFC 2328:
  - *A cost is associated with the output side of each router interface. This cost is configurable by the system administrator. The lower the cost, the more likely the interface is to be used to forward data traffic.*
- RFC 2328 does not specify which values should be used to determine the cost.

OSPF Metric

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	100,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	100,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	100,000,000	÷	100,000,000	1
Ethernet 10 Mbps	100,000,000	÷	10,000,000	10
Serial 1.544 Mbps	100,000,000	÷	1,544,000	64
Serial 128 kbps	100,000,000	÷	128,000	781
Serial 64 kbps	100,000,000	÷	64,000	1562

Cisco IOS Cost for OSPF = 10^8/bandwidth in bps

- Cisco IOS software uses the cumulative bandwidths of the outgoing interfaces from the router to the destination network as the cost value.
- **10^8** is known as the *reference bandwidth*
  - So that interfaces with the higher bandwidth values will have a lower calculated cost.

Reference Bandwidth

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	100,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	100,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	100,000,000	÷	100,000,000	1
Ethernet 10 Mbps	100,000,000	÷	10,000,000	10
Serial 1.544 Mbps	100,000,000	÷	1,544,000	64
Serial 128 kbps	100,000,000	÷	128,000	781
Serial 64 kbps	100,000,000	÷	64,000	1562

- The **reference bandwidth**
  - Defaults to 10^8, which is 100,000,000 bps or 100 Mbps.
- This results in interfaces with a bandwidth of 100 Mbps and higher having the **same OSPF cost of 1.**
- It can be modified using the OSPF command

  Router(config-router)# auto-cost *referencebandwidth Mb/s*

  - *When necessary, use on all routers so the OSPF routing metric remains consistent.*

Accommodating 10Gig Interfaces

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	10,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	10,000,000,000	÷	1,000,000,000	10
Fast Ethernet 100 Mbps	10,000,000,000	÷	100,000,000	100
Ethernet 10 Mbps	10,000,000,000	÷	10,000,000	1000
Serial 1.544 Mbps	10,000,000,000	÷	1,544,000	6477
Serial 128 kbps	10,000,000,000	÷	128,000	78125
Serial 64 kbps	10,000,000,000	÷	64,000	156250

In Mb/s

```
R1(config-router)# auto-cost reference-bandwidth 10000
```

Adjusting to Reference Bandwidth for Gig

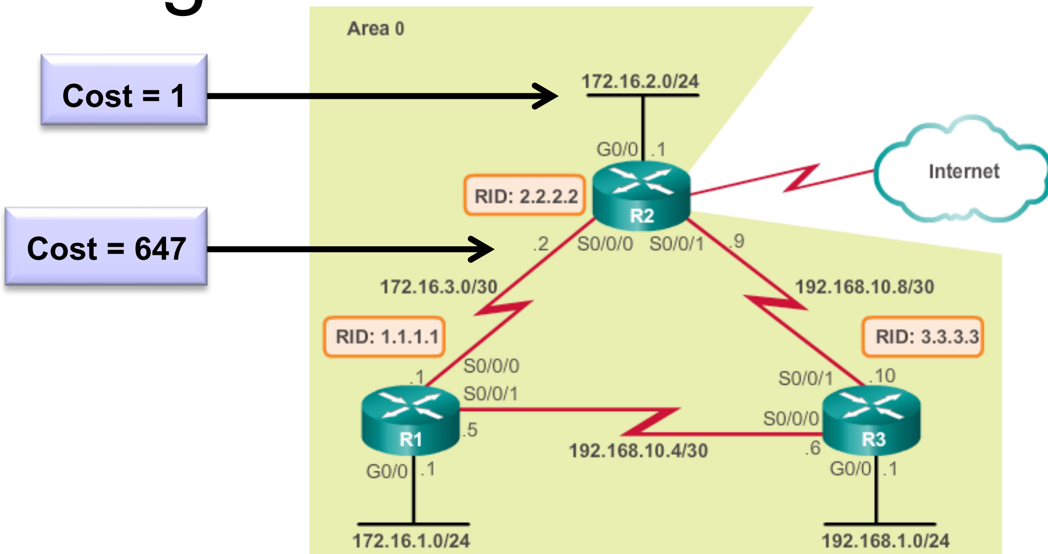

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	1,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	1,000,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	1,000,000,000	÷	100,000,000	10
Ethernet 10 Mbps	1,000,000,000	÷	10,000,000	100
Serial 1.544 Mbps	1,000,000,000	÷	1,544,000	647
Serial 128 kbps	1,000,000,000	÷	128,000	7812
Serial 64 kbps	1,000,000,000	÷	64,000	15625

```
R1(config)# router ospf 10
R1(config-router)# auto-cost reference-bandwidth 1000
R1(config-router)#
```

```
R2(config)# router ospf 10
R2(config-router)# auto-cost reference-bandwidth 1000
R2(config-router)#
```

```
R3(config)# router ospf 10
R3(config-router)# auto-cost reference-bandwidth 1000
R3(config-router)#
```

Adjusting to Reference Bandwidth for Gig

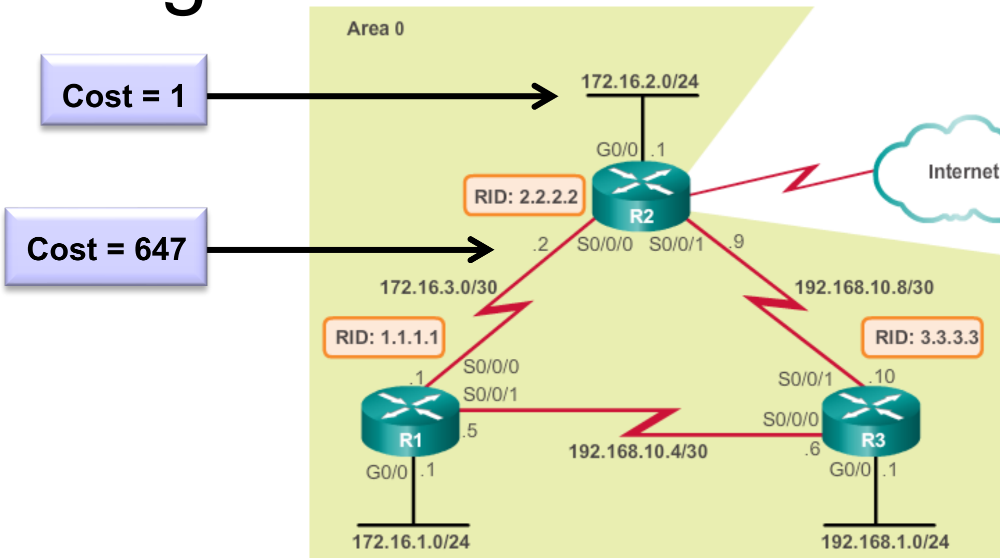

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	1,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	1,000,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	1,000,000,000	÷	100,000,000	10
Ethernet 10 Mbps	1,000,000,000	÷	10,000,000	100
Serial 1.544 Mbps	1,000,000,000	÷	1,544,000	647
Serial 128 kbps	1,000,000,000	÷	128,000	7812
Serial 64 kbps	1,000,000,000	÷	64,000	15625

```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 172.16.3.1/30, Area 0, Attached via Network Statement
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 647
  Topology-MTID     Cost     Disabled     Shutdown      Topology Name
        0            647        no           no            Base
  <Output omitted>
R1#
```

Verify the Adjusted Reference Bandwidth

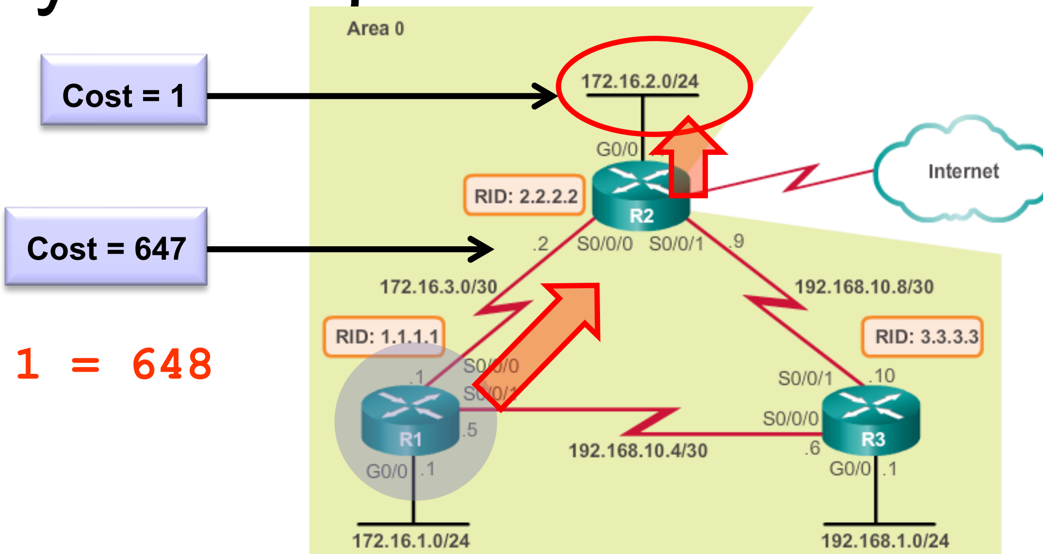

647 + 1 = 648

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	1,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	1,000,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	1,000,000,000	÷	100,000,000	10
Ethernet 10 Mbps	1,000,000,000	÷	10,000,000	100
Serial 1.544 Mbps	1,000,000,000	÷	1,544,000	647
Serial 128 kbps	1,000,000,000	÷	128,000	7812
Serial 64 kbps	1,000,000,000	÷	64,000	15625

```
R1# show ip route | include 172.16.2.0
O        172.16.2.0/24 [110/648] via 172.16.3.2, 00:06:03, Serial0/0/0
R1#
R1# show ip route 172.16.2.0
Routing entry for 172.16.2.0/24
  Known via "ospf 10", distance 110, metric 648, type intra area
  Last update from 172.16.3.2 on Serial0/0/0, 00:06:17 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 2.2.2.2, 00:06:17 ago, via Serial0/0/0
      Route metric is 648, traffic share count is 1
R1#
```

Default Bandwidth on Serial Interfaces

```
R1# show interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to R2
  Internet address is 192.168.10.1/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

- On Cisco routers, the **bandwidth value** on many serial interfaces defaults to T1 (1.544 Mbps).
  - Always check this with the **show interface** command.
  - Tip – Always use the bandwidth command on serial interfaces.
- Modified bandwidth values are shown in running-config.
- **Bandwidth value** does not actually affect the speed of the link

Default Bandwidth on Serial Interfaces

Serial interfaces bandwidth value defaults to T1 or 1544 Kbps.

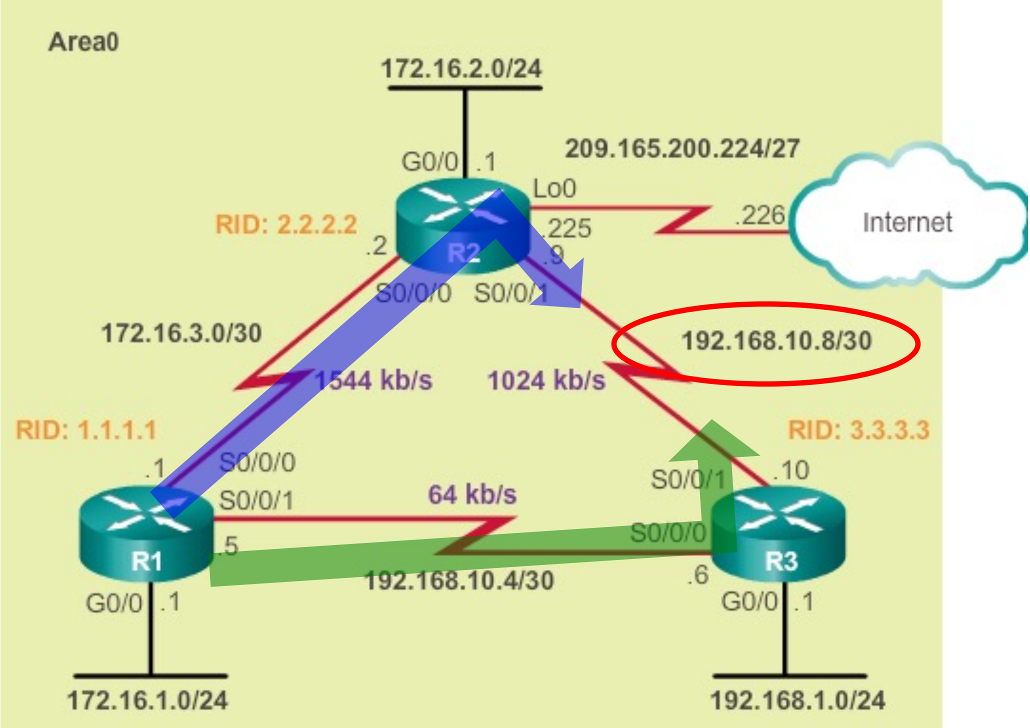


```
R1# show ip route
<route ouput omitted>
O    192.168.10.8 [110/1294] via 192.168.10.6, 14:27:57, Serial0/0/1
                  [110/1294] via 192.168.10.2, 14:27:57, Serial0/0/0
```

- **R1** believes that both of its serial interfaces are connected to T1 links.
- **R1's routing table** having **two equal-cost paths** to the 192.168.8.0/30 network.
  - Serial 0/0/0 is actually the better path.

Adjusting Interface Bandwidth

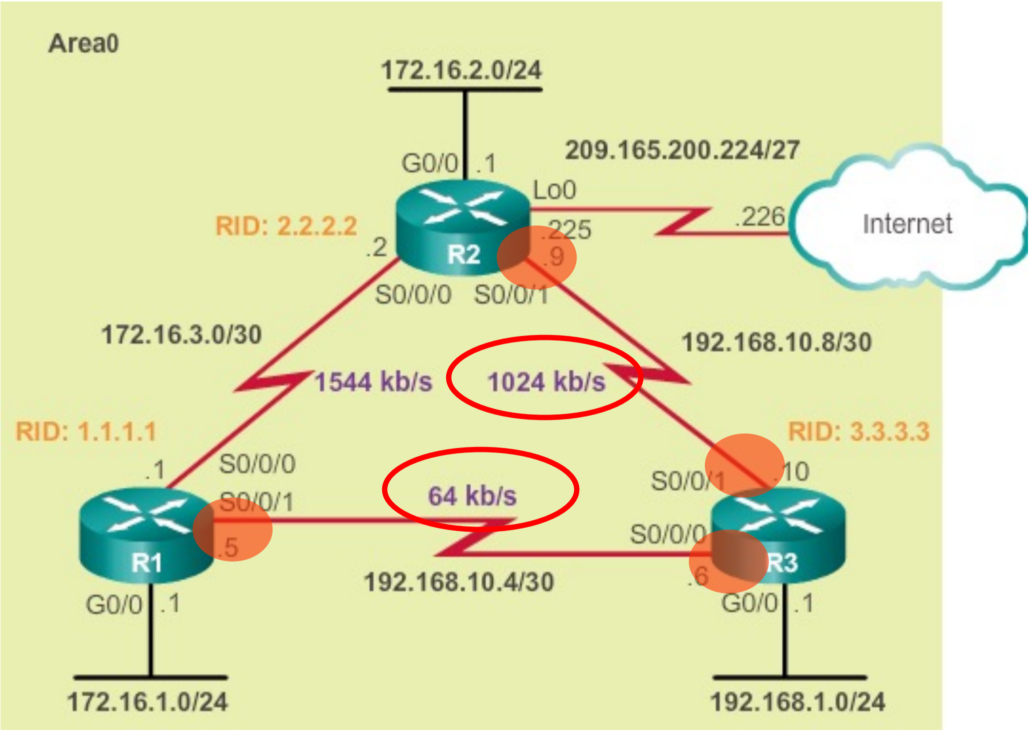

- To adjust the interface bandwidth use the **`bandwidth`** *`kilobits`* interface configuration command.

Adjusting Interface Bandwidth on R1

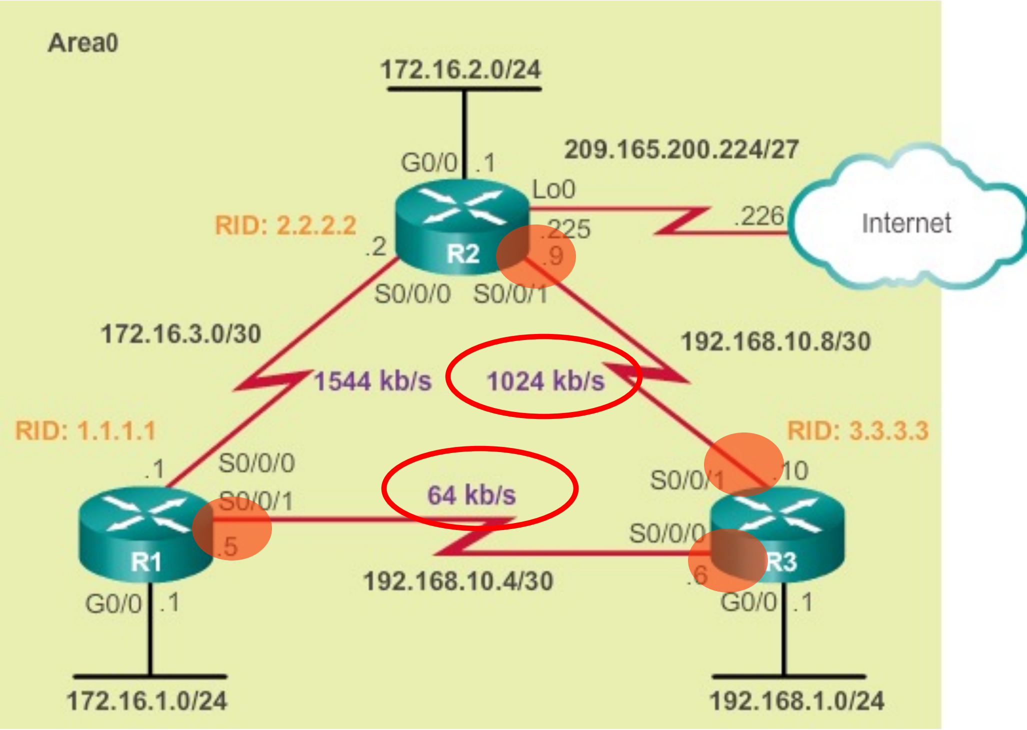

Bandwidth parameter is in Kb/s

```
R1(config)# int s0/0/1
R1(config-if)# bandwidth 64
R1(config-if)# end
```

Serial 64 kbps	1,000,000,000 ÷	64,000	15625

```
R1# show interfaces serial 0/0/1 | include BW
  MTU 1500 bytes, BW 64 Kbit/sec, DLY 20000 usec,
R1#
R1# show ip ospf interface serial 0/0/1 | include Cost:
 Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 15625
R1#
```

Modifying the Cost of the Link

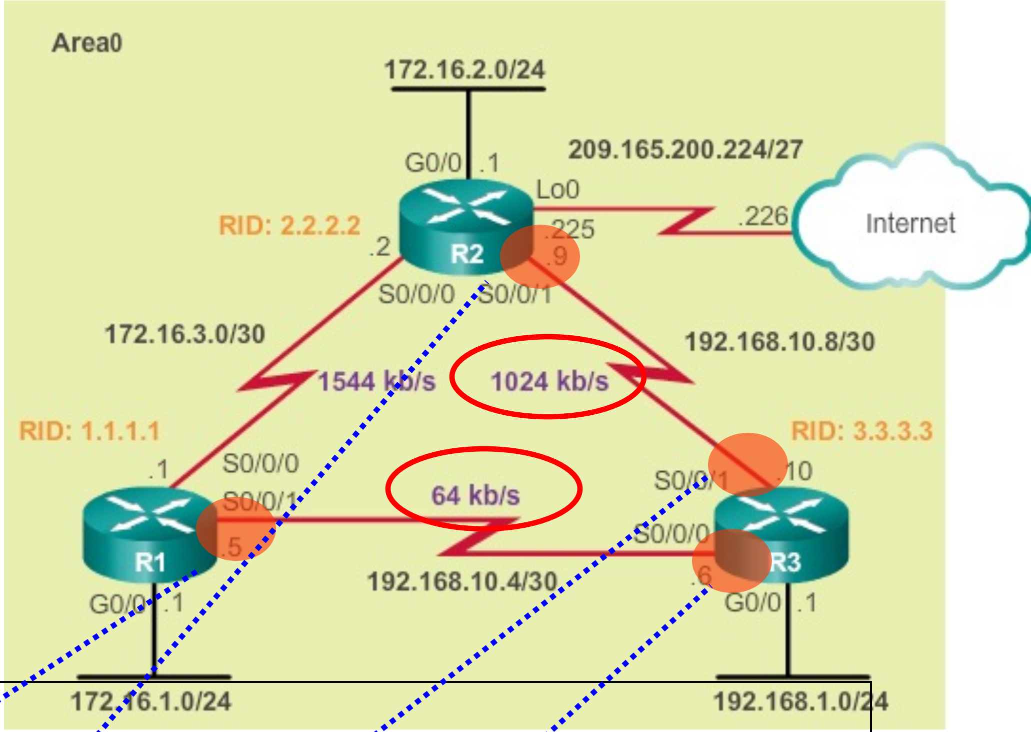

- Both sides of the link should be configured to have the same value.

```
R1(config)# inter serial 0/0/1
R2(config-if)# bandwidth 64
```

```
R2(config-if)# inter serial 0/0/1
R2(config-if)# bandwidth 1024
```

```
R3(config)# inter serial 0/0/1
R2(config-if)# bandwidth 1024
R3(config-if)# inter serial 0/0/0
R3(config-if)# bandwidth 64
```

The **ip ospf cost** Command

```
R1(config)# inter serial 0/0/1
R1(config-if)# bandwidth 64
R1(config-if)# end
R1# show ip ospf interface serial 0/0/0
Serial0/0 is up, line protocol is up
Internet Address 192.168.10.1/30, Area 0
Process ID 1, Router ID 10.1.1.1, Network Type POINT_TO_POINT,Cost: 15625
<output omitted>
```

1,000,000,000/64,000 = 15625

```
R1(config)# interface serial 0/0/1
R1(config-if)# ip ospf cost 15625
```

- An alternative method to using the **`bandwidth`** command is to use the **`ip ospf cost`** command, which allows you to directly specify the cost of an interface.
- This will not change the output of the **`show ip ospf interface`** command.

Changing Bandwidth Versus Cost

Adjusting the Interface Bandwidth	=	Manually Setting the OSPF Cost
R1(config)#**interface S0/0/1** R1(config-if)#**bandwidth 64**	=	R1(config)#**interface S0/0/1** R1(config-if)#**ip ospf cost 15625**
R2(config)#**interface S0/0/1** R2(config-if)#**bandwidth 1024**	=	R2(config)#**interface S0/0/1** R2(config-if)#**ip ospf cost 976**
R3(config)#**interface S0/0/0** R3(config-if)#**bandwidth 64**	=	R3(config)#**interface S0/0/0** R3(config-if)#**ip ospf cost 15625**
R3(config)#**interface S0/0/1** R3(config-if)#**bandwidth 1024**	=	R3(config)#**interface S0/0/1** R3(config-if)#**ip ospf cost 976**

Verifying OSPF

Verifying OSPF Settings

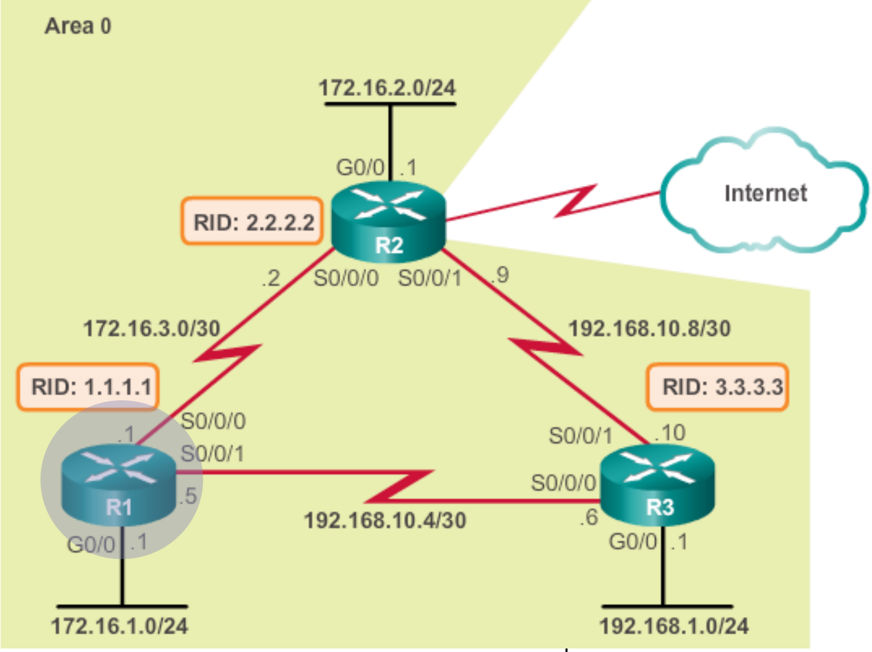


```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    172.16.3.0 0.0.0.3 area 0
    192.168.10.4 0.0.0.3 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    2.2.2.2              110      00:17:18
    3.3.3.3              110      00:14:49
  Distance: (default is 110)
R1#
```

Network command

Neighbors

Verifying OSPF Neighbors

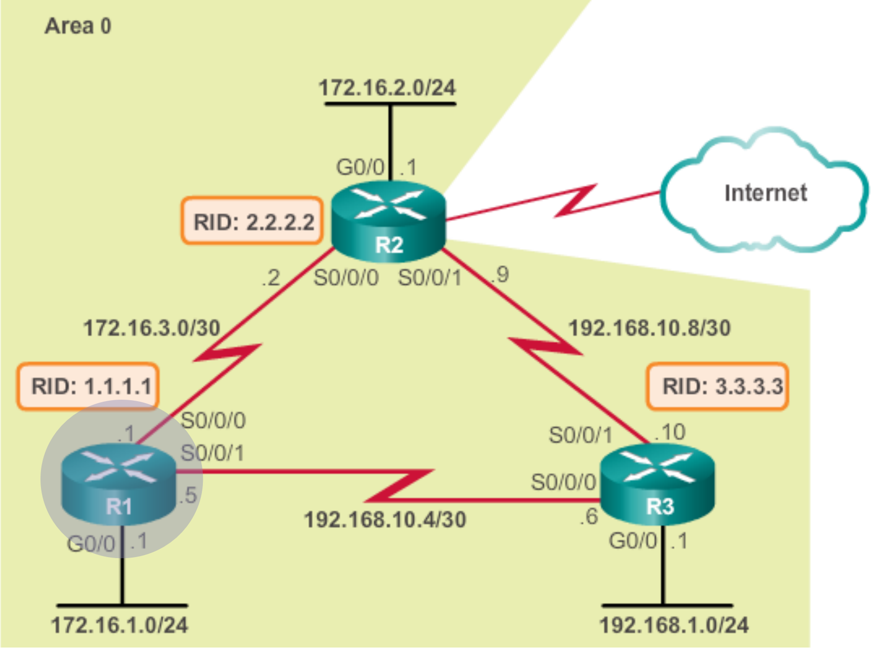

Lists of OSPF neighbors in the order they were learned.

The amount of time remaining before declaring the neighbor down.

The local interface to reach this neighbor.

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:37    192.168.10.6    Serial0/0/1
2.2.2.2           0   FULL/  -        00:00:30    172.16.3.2      Serial0/0/0
R1#
```

The OSPF priority of the interface.

- Used in DR/BDR election.

The state of the OSPF enabled interface.

- **FULL** state means that the router and its neighbor have identical OSPF LSDB.

The neighbor's IP address

Verifying OSPF Routes

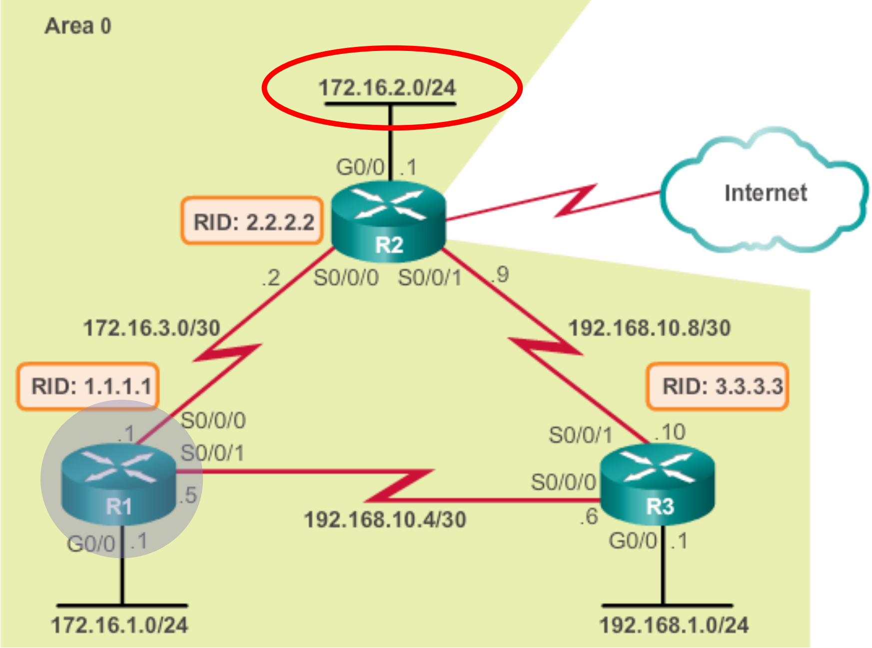


```
R1# show ip route | include 172.16.2.0
O        172.16.2.0/24 [110/65] via 172.16.3.2, 03:39:07,
Serial0/0/0
R1#
R1# show ip route 172.16.2.0
Routing entry for 172.16.2.0/24
  Known via "ospf 10", distance 110, metric 65, type intra area
  Last update from 172.16.3.2 on Serial0/0/0, 03:39:15 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 2.2.2.2, 03:39:15 ago, via Serial0/0/0
      Route metric is 65, traffic share count is 1
R1#
```

Clearing the OSPF Routing Table

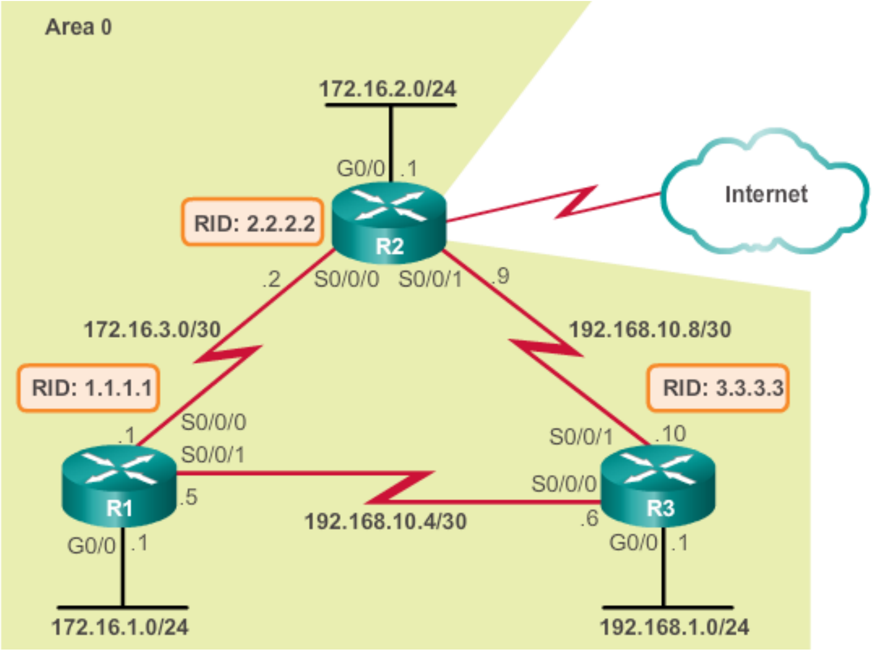

- To clear all routes from the IP routing table, use:
  - Router# **`clear ip route *`**
- To clear a specific route from the IP routing table, use:
  - Router# **`clear ip route A.B.C.D`**

Verifying OSPF Enabled Interfaces

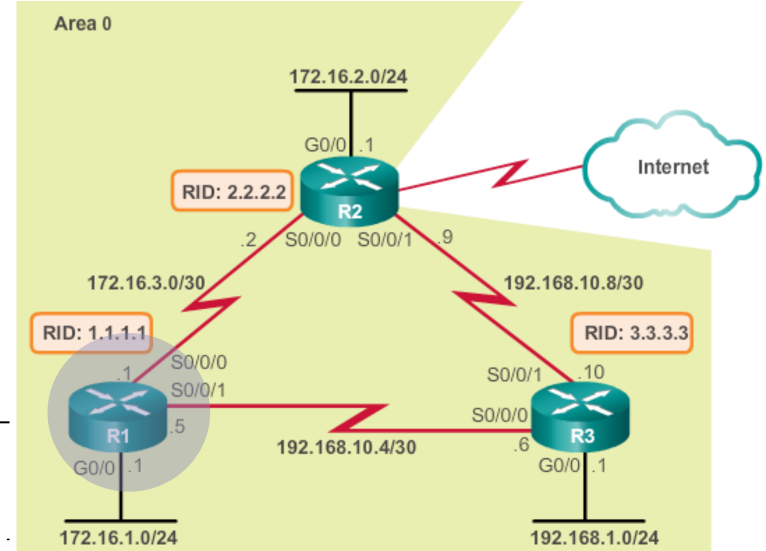


```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 172.16.3.1/30, Area 0, Attached
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 647
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           647       no          no             Base
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  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 2.2.2.2
Suppress hello for 0 neighbor(s)
R1#
```

Verifying OSPF Enabled Interfaces

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
10 Gigabit Ethernet 10 Gbps	1,000,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	1,000,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	1,000,000,000	÷	100,000,000	10
Ethernet 10 Mbps	1,000,000,000	÷	10,000,000	100
Serial 1.544 Mbps	1,000,000,000	÷	1,544,000	647
Serial 128 kbps	1,000,000,000	÷	128,000	7812
Serial 64 kbps	1,000,000,000	÷	64,000	15625

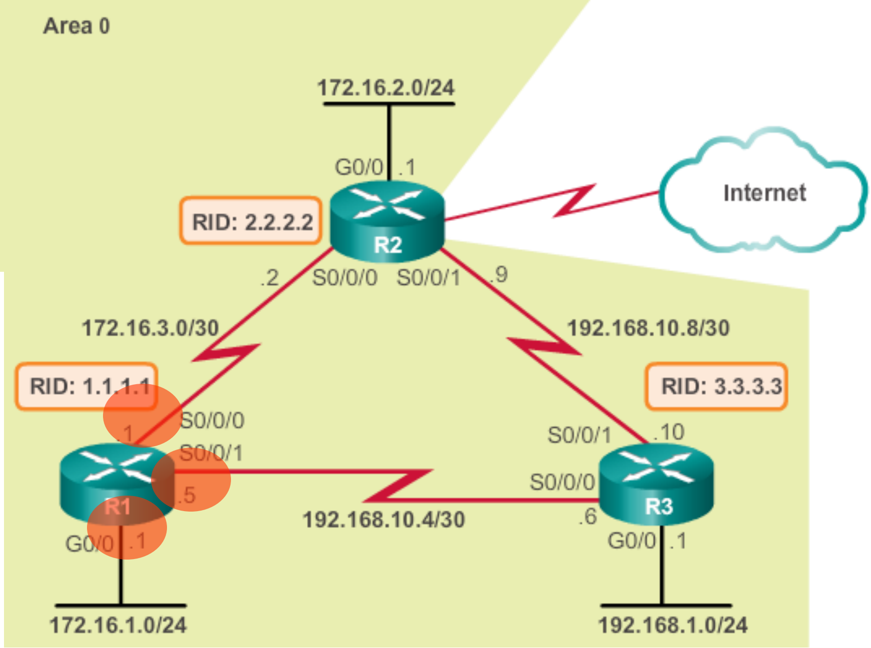


```
R1# show ip ospf interface brief
Interface      PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Se0/0/1        10    0               192.168.10.5/30    15625 P2P   1/1
Se0/0/0        10    0               172.16.3.1/30      647   P2P   1/1
Gi0/0          10    0               172.16.1.1/24      1     DR    0/0
R1#
```

Verifying the OSPF Process

```
R1# show ip ospf
 Routing Process "ospf 10" with ID 1.1.1.1
 Start time: 01:37:15.156, Time elapsed: 01:32:57.776

<Output omitted>

Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Number of areas transit capable is 0
 External flood list length 0
 IETF NSF helper support enabled
 Cisco NSF helper support enabled
 Reference bandwidth unit is 1000 mbps
    Area BACKBONE(0)
        Number of interfaces in this area is 3
       Area has no authentication
       SPF algorithm last executed 01:30:45.364 ago
       SPF algorithm executed 3 times
       Area ranges are
       Number of LSA 3. Checksum Sum 0x02033A

<Output omitted>
```

Verifying OSPF

Command	Description
`show ip protocols`	• Displays OSPF process ID, router ID, networks router is advertising & administrative distance
`show ip ospf neighbors`	• Displays OSPF neighbor relationships.
`show ip route`	• Displays the routing table.
`show ip ospf interface`	• Displays hello interval and dead interval
`show ip ospf`	• Displays OSPF process ID, router ID, OSPF area information & the last time SPF algorithm calculated

OSPFv3

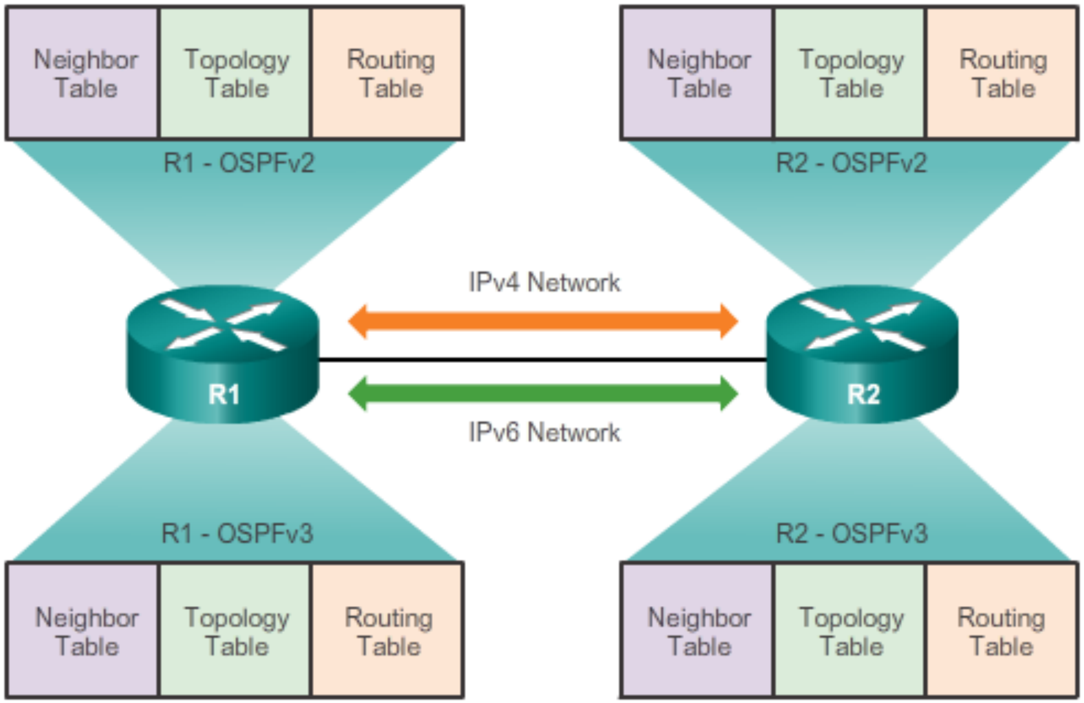

Neighbor Table
Topology Table
Routing Table
R1 - OSPFv2
Neighbor Table
Topology Table
Routing Table
R2 - OSPFv2
IPv4 Network
R1
R2
IPv6 Network
R1 - OSPFv3
Neighbor Table
Topology Table
Routing Table
R2 - OSPFv3
Neighbor Table
Topology Table
Routing Table

Note: OSPFv3 supports both IPv4 and IPv6 with the use of Address Families (beyond the scope of CCNA but in CCNP)

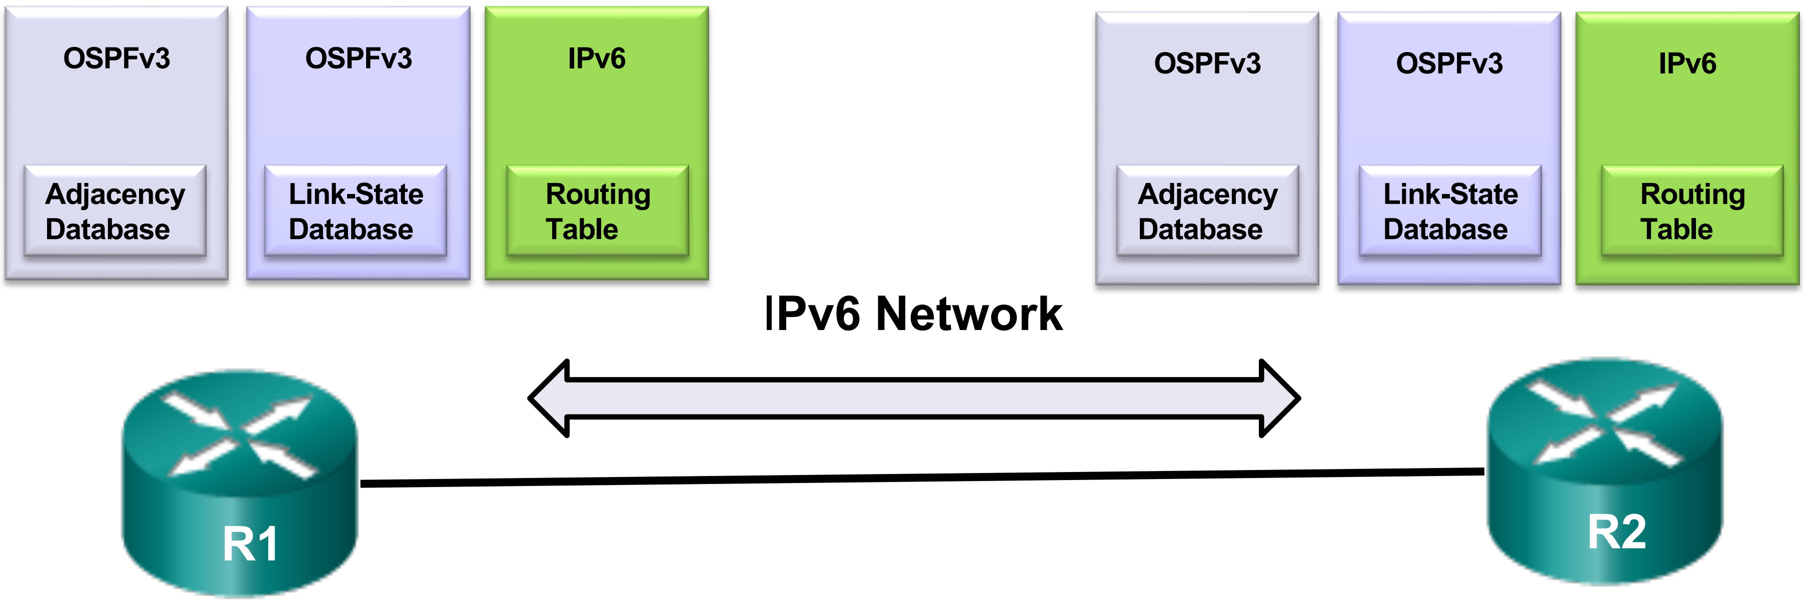

Note:

- IPv6 link-local addresses are in the FE80::/10 range.
- The /10 indicates that the first 10 bits are 1111 1110 10xx xxxx, which results in the first hextet having a range of 1111 1110 1000 0000 (FE80) to 1111 1110 1011 1111 (FEBF).

OSPFv2 and OSPFv3 Similarities

Link-State	
Routing Algorithm	
Metric	
Areas	
Packet types	
Neighbor discovery	
DR and BDR	
Router ID	

OSPFv2 vs. OSPFv3

	OSPFv2	OSPFv3
Advertises		
Source address		
Destination address		
IP unicast routing		
Authentication		

OSPFv3 Topology

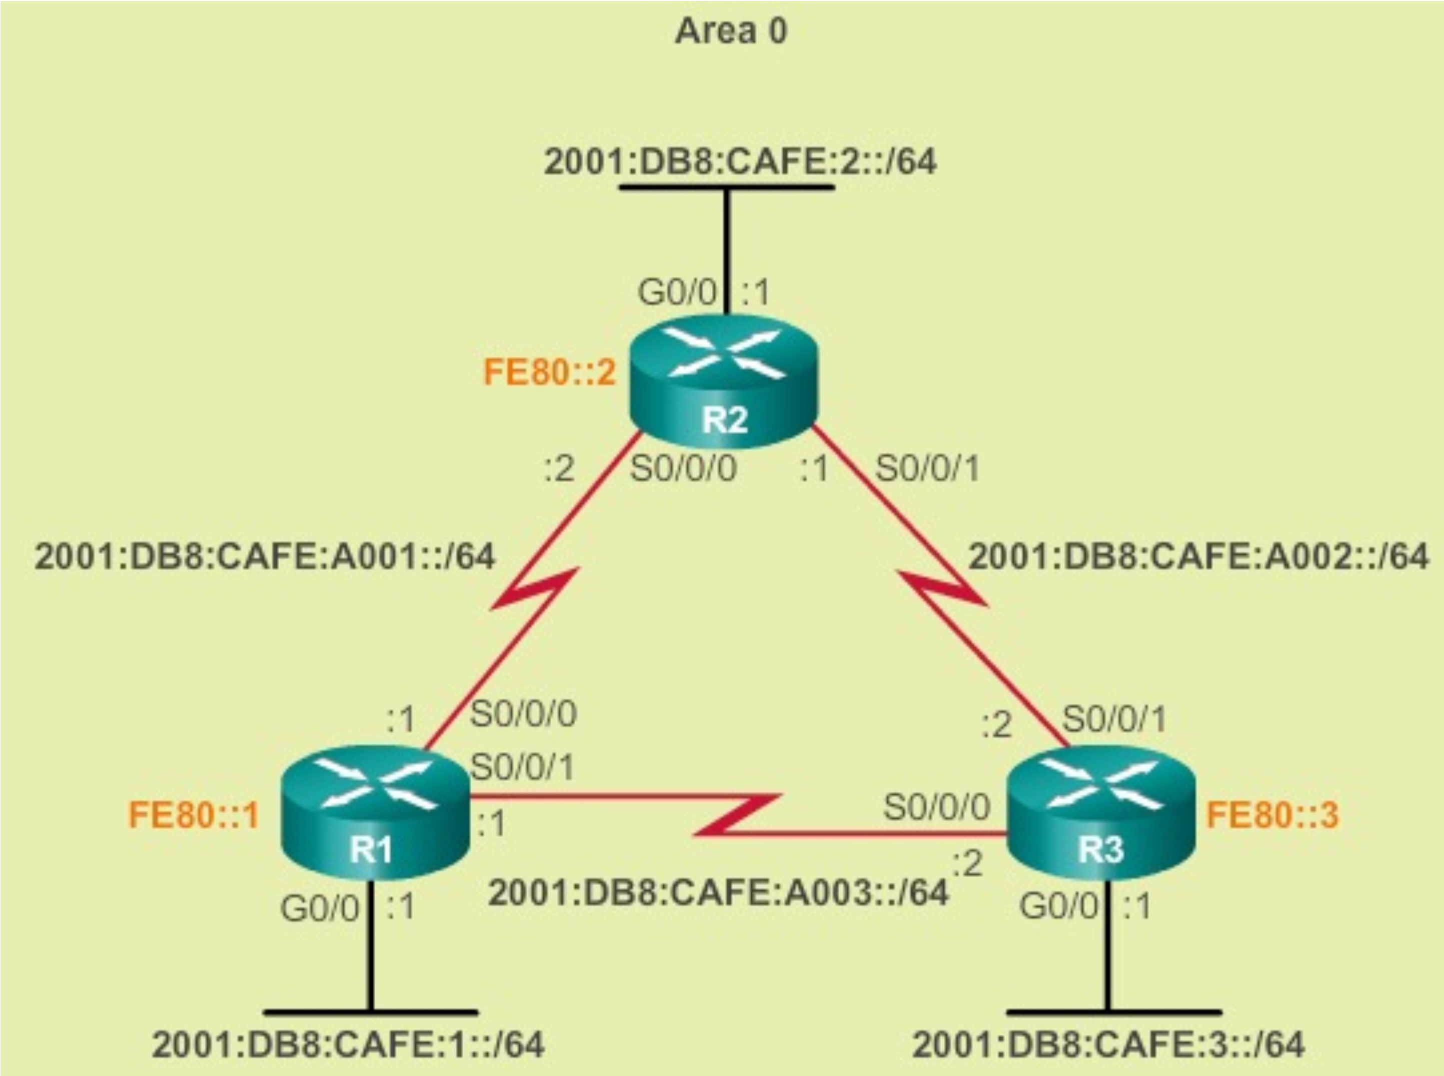

Verify OSPFv3 Interfaces on R1

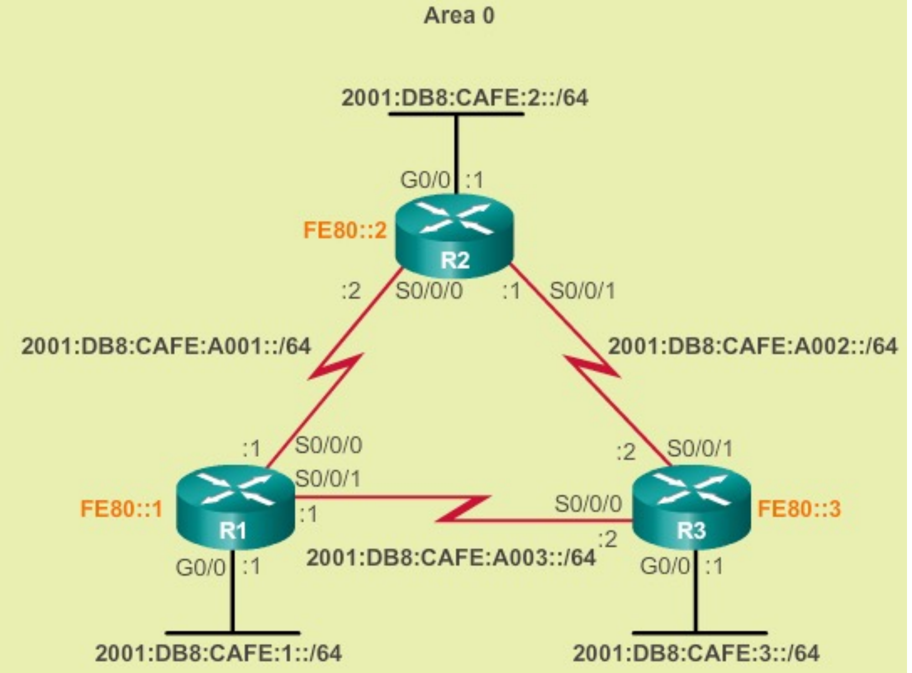


```
R1# show ipv6 interface brief

GigabitEthernet0/0      [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:1::1
GigabitEthernet0/1      [administratively down/down]
    unassigned
Serial0/0/0             [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:A001::1
Serial0/0/1             [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:A003::1
R1#
```

Configuring the OSPFv3 Routing Process

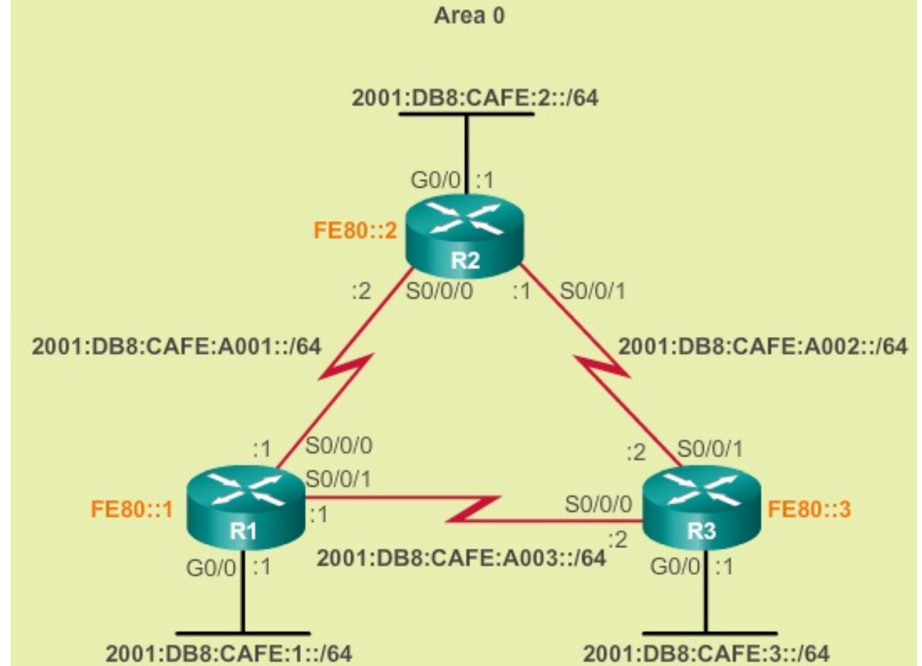


```
R1(config)#ipv6 unicast-routing
R1(config)#ipv6 router ospf 10
R1(config-rtr)#
*Mar 29 11:21:53.739: %OSPFv3-4-NORTRID: Process OSPFv3-1-IPv6
could not pick a router-id, please configure manually
R1(config-rtr)#
R1(config-rtr)#router-id 1.1.1.1
R1(config-rtr)#auto-cost reference-bandwidth 1000
% OSPFv3-1-IPv6:  Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across
        all routers.
R1(config-rtr)#end
R1#
```

Same process as OSPFv2

32-bit Router ID similar to OSPFv2

There is no "**no shutdown**". ☺

Verify Router ID

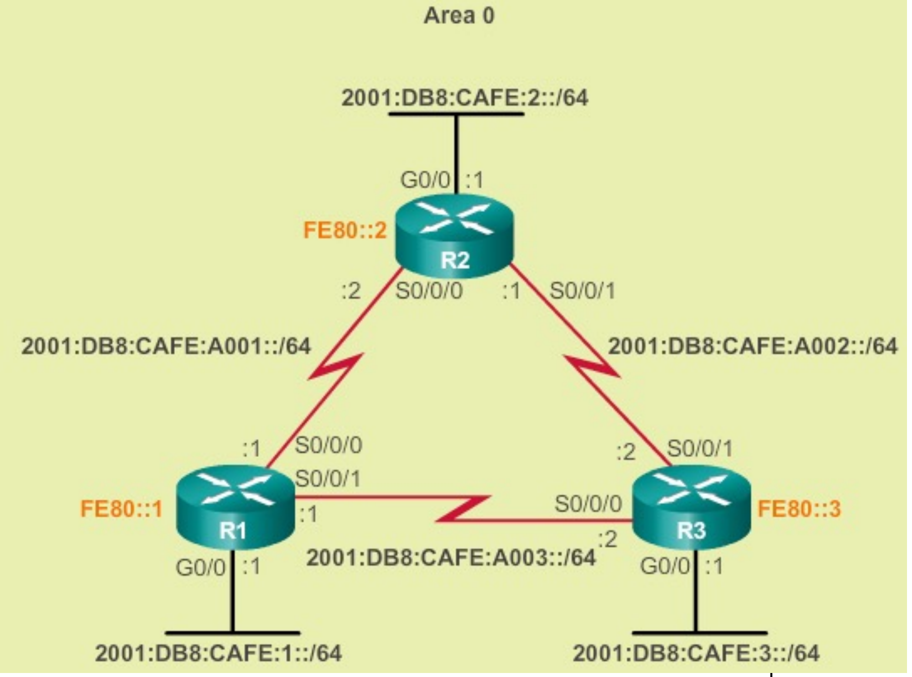


```
R1# clear ipv6 ospf process
Reset selected OSPFv3 processes? [no]: y
R1#
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 0 normal, 0 stub, 0 nssa
  Redistribution:
    None
R1#
```

Configure R1

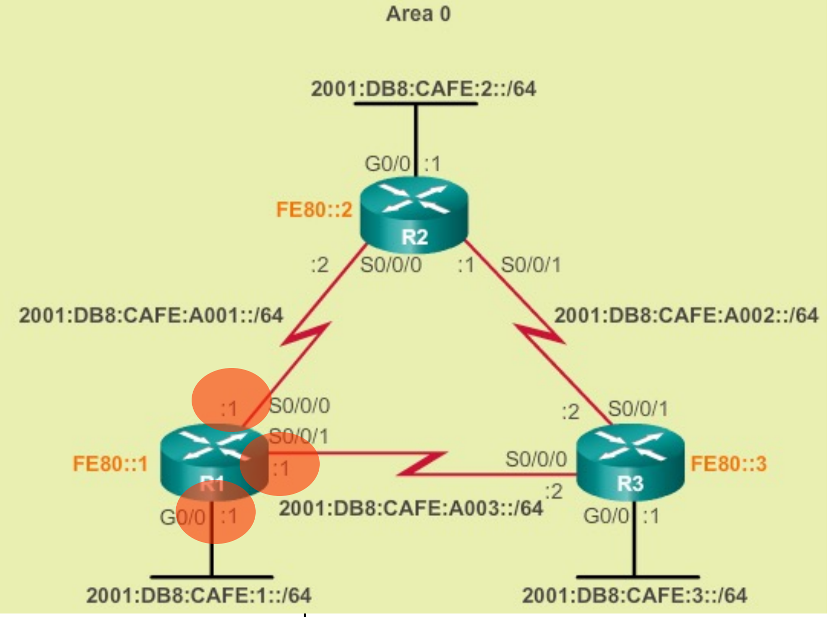


```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# exit
R1(config)# interface Serial0/0/0
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# exit
R1(config)# interface Serial0/0/1
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# bandwidth 64
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# end
R1#
```

Configure R2

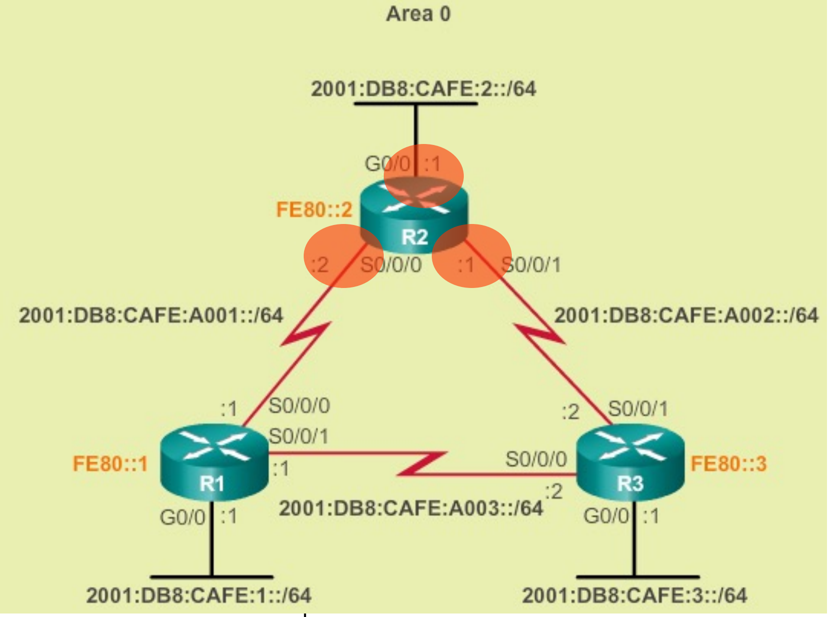


```
R2(config)# interface GigabitEthernet 0/0
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/0
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/1
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# bandwidth 1024
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# end
R2#
```

Configure R3

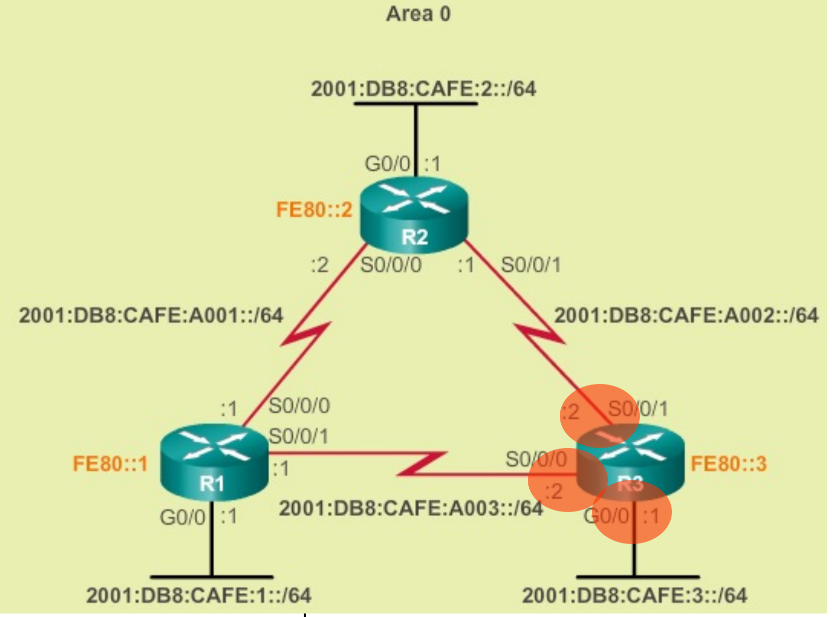


```
R3(config)# interface GigabitEthernet 0/0
R3(config-if)# ipv6 address fe80::3 link-local
R3(config-if)# ipv6 ospf 10 area 0
R3(config-if)# exit
R3(config)# interface Serial0/0/0
R3(config-if)# ipv6 address fe80::3 link-local
R3(config-if)# bandwidth 64
R3(config-if)# ipv6 ospf 10 area 0
R3(config-if)# exit
R3(config)# interface Serial0/0/1
R3(config-if)# ipv6 address fe80::3 link-local
R3(config-if)# bandwidth 1024
R3(config-if)# ipv6 ospf 10 area 0
R3(config-if)# end
```

Verify R1

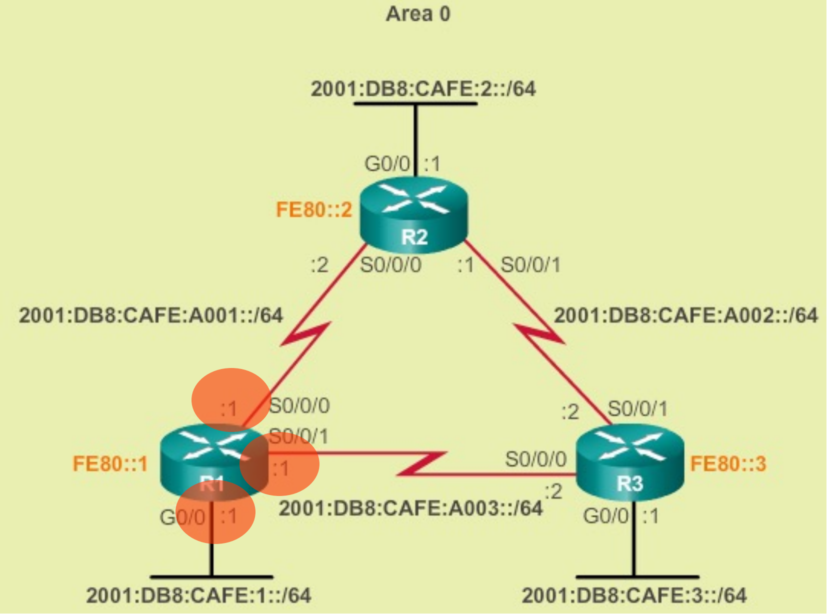


```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/1
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#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

Verify R1

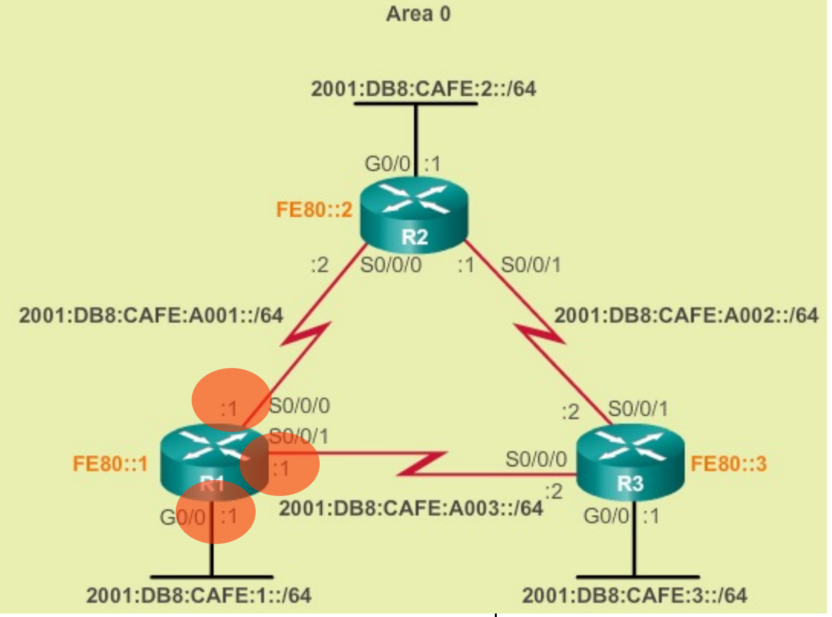


```
R1# show ipv6 interface brief
Em0/0                      [administratively down/down]
    unassigned
GigabitEthernet0/0         [up/up]
    FE80::1
    2001:DB8:CAFE:1::1
GigabitEthernet0/1         [administratively down/down]
    unassigned
Serial0/0/0                [up/up]
    FE80::1
    2001:DB8:CAFE:A001::1
Serial0/0/1                [up/up]
    FE80::1
    2001:DB8:CAFE:A003::1
R1#
```

Verify R1

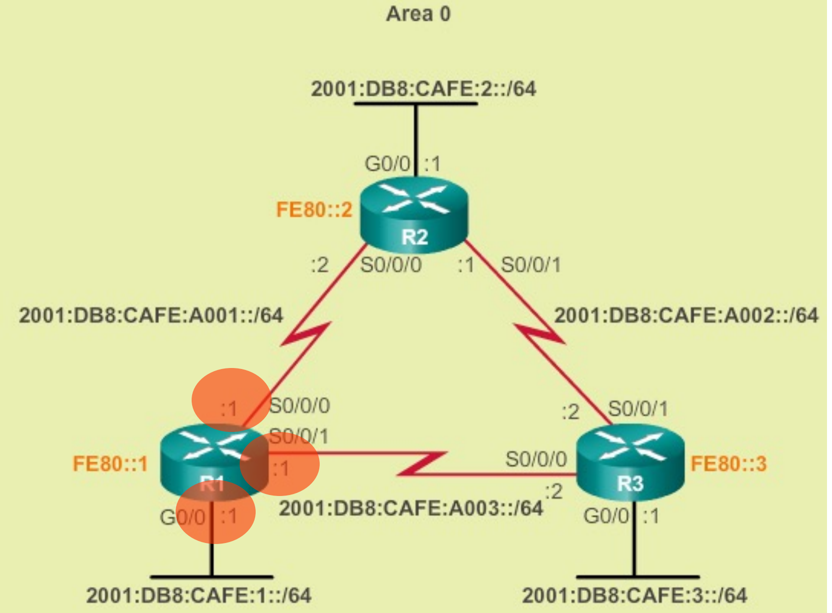


```
R1#show ipv6 ospf interface brief
Interface    PID   Area           Intf ID    Cost  State Nbrs F/C
Se0/0/1      10    0              7          15625 P2P   1/1
Se0/0/0      10    0              6          647   P2P   1/1
Gi0/0        10    0              3          1     DR    0/0
R1#
R1# show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

Neighbor ID     Pri   State        Dead Time   Interface ID
Interface
3.3.3.3           0   FULL/  -     00:00:39    6
Serial0/0/1
2.2.2.2           0   FULL/  -     00:00:36    6
Serial0/0/0
R1#
```

Verify R1

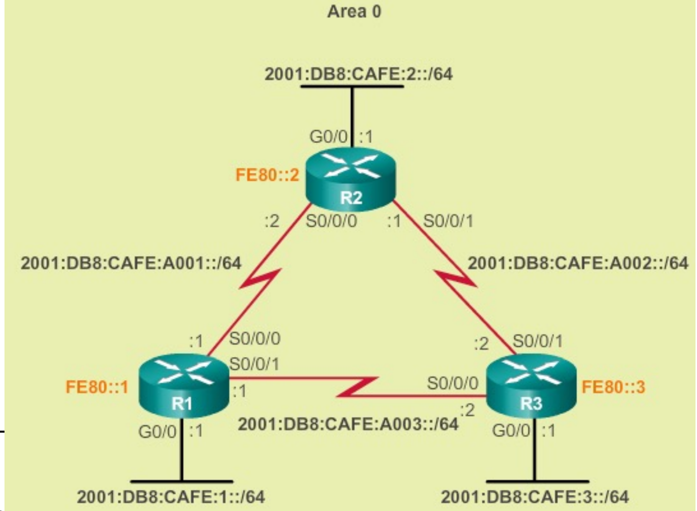

Interface bandwidth previously modified

```
Router(config-if)# bandwidth 64
```

```
R1# show interfaces serial 0/0/1 | include BW
  MTU 1500 bytes, BW 64 Kbit/sec, DLY 20000 usec,
R1#
R1# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix,
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
O   2001:DB8:CAFE:2::/64 [110/657]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:3::/64 [110/1304]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:A002::/64 [110/1294]
     via FE80::2, Serial0/0/0
R1#
```

Summary

Chapter 8: Summary

- OSPF for IPv4 is OSPFv2 and for IPv6 is OSPFv3
- Classless, link-state routing protocol with a default administrative distance of 110, and is denoted in the routing table with a route source code of O
- OSPFv2 is enabled with the router ospf process-id global configuration mode command.
  - The process-id value is locally significant, which means that it does not need to match other OSPF routers to establish adjacencies with those neighbors.
- Network command uses the wildcard-mask value which is the inverse of the subnet mask, and the area-id value

Chapter 8: Summary (cont.)

- By default, OSPF Hello packets are sent every 10 seconds on multiaccess and point-to-point segments and every 30 seconds on NBMA segments (Frame Relay, X.25, ATM), and are used by OSPF to establish neighbor adjacencies.
  - The Dead interval is four times the Hello interval, by default.
- For routers to become adjacent, their Hello interval, Dead interval, network types, and subnet masks must match.
  - Use the show ip ospf neighbors command to verify OSPF adjacencies.

Chapter 8: Summary (cont.)

- In multiaccess networks, the router with the highest router ID is the DR, and the router with the second highest router ID is the BDR. This can be superseded by the ip ospf priority command on that interface. The router with the highest priority value is the DR, and next-highest the BDR.

- The show ip protocols command is used to verify important OSPF configuration information, including the OSPF process ID, the router ID, and the networks the router is advertising.

Chapter 8: Summary (cont.)

- OSPFv3
  - Enabled on an interface and not under router configuration mode
  - Needs link-local addresses to be configured. IPv6
  - Unicast routing must be enabled for OSPFv3
  - 32-bit router-ID is required before an interface can be enabled for OSPFv3
  - show ipv6 protocols command is a quick way to verify configuration information (OSPF process ID, the router ID, and the interfaces enabled for OSPFv3)