Architecture des Réseaux

ESIEE AMIENS

L'innovation pour de nouveaux horizons

Part II: Routing Concepts and Algorithms

École d'ingénieurs

RÉFÉRENCES DU COURS

- Cours Obligatoire
- 8 heures de CM
- 2 heures de TD
- 4 heures de TP

I- Introduction
Characteristics of a Network

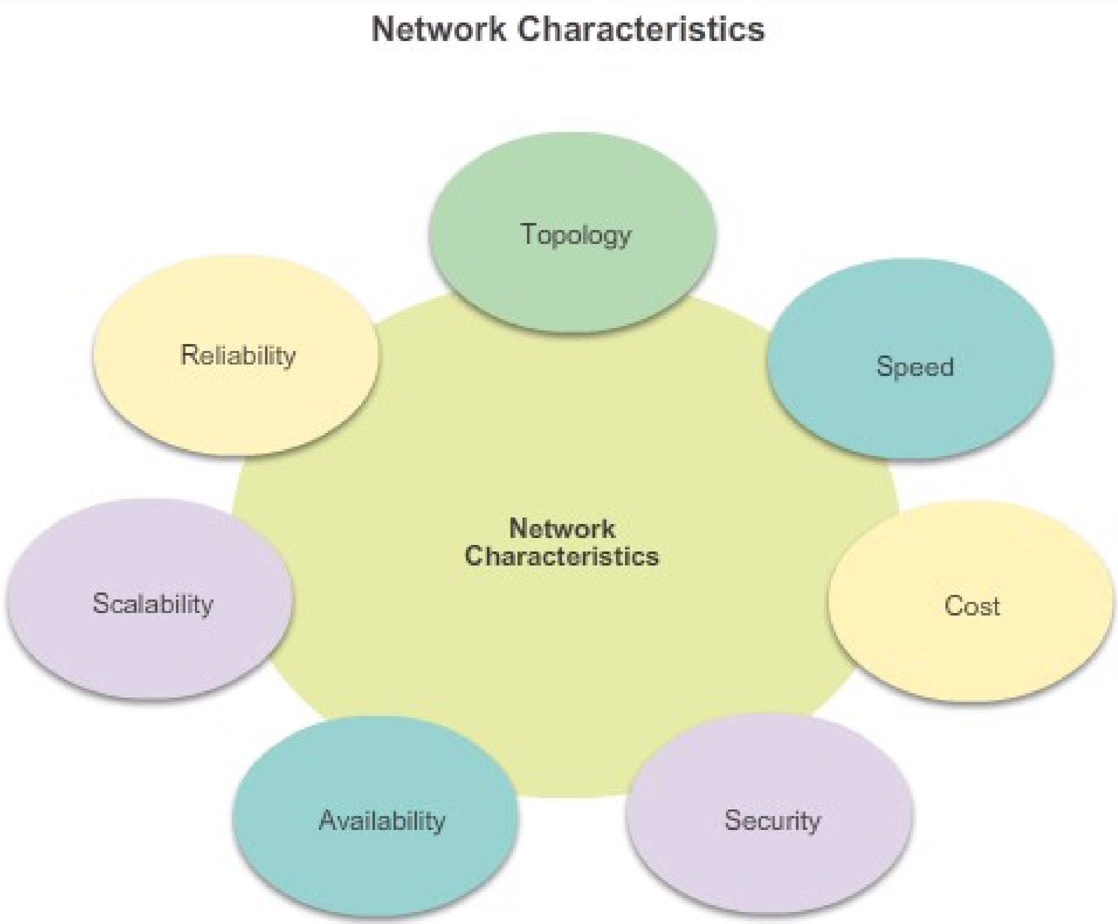

I- Introduction
Routers are Computers

- Routers are specialized computers containing the following required components to operate:
  - Central processing unit (CPU)
  - Operating system (OS) - Routers use Cisco IOS
  - Memory and storage (RAM, ROM, NVRAM, Flash, hard drive)

Memory	Volatile / Non-Volatile	Stores
RAM	Volatile	• Running IOS • Running configuration file • IP routing and ARP tables • Packet buffer
ROM	Non-Volatile	• Bootup instructions • Basic diagnostic software • Limited IOS
NVRAM	Non-Volatile	• Startup configuration file
Flash	Non-Volatile	• IOS • Other system files

I- Introduction
Routers Choose Best Paths

- Routers use static routes and dynamic routing protocols to learn about remote networks and build their routing tables.
- Routers use routing tables to determine the best path to send packets.
- Routers encapsulate the packet and forward it to the interface indicated in routing table.

I- Introduction
Packet Forwarding Methods

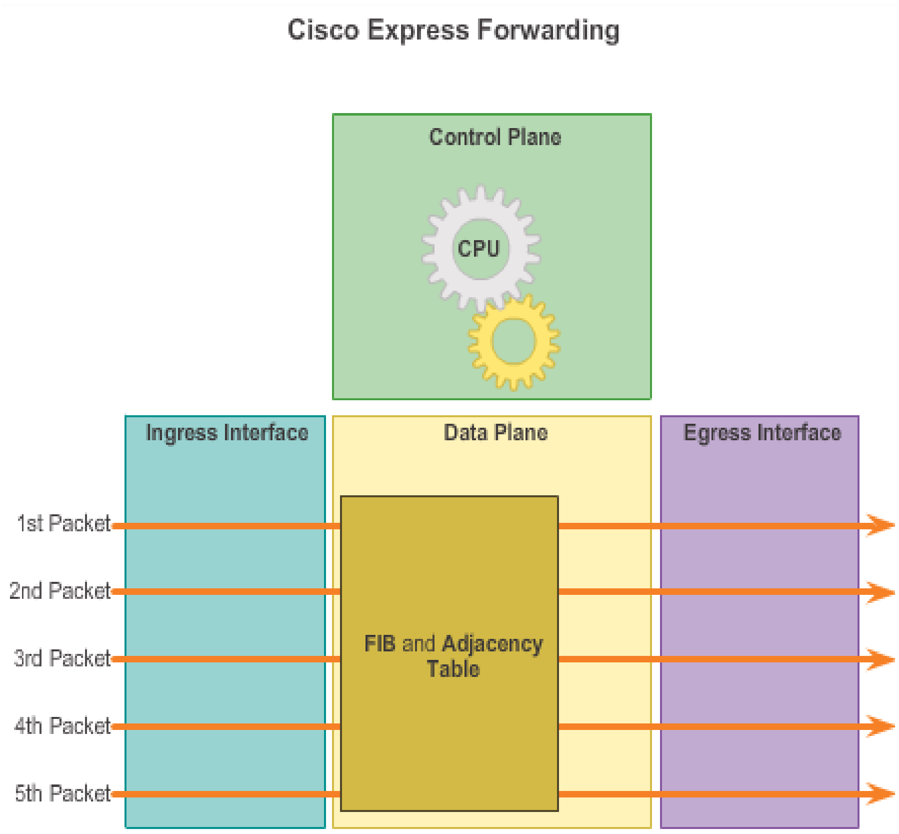

- **Process switching** – An older packet forwarding mechanism still available for Cisco routers.
- **Fast switching** – A common packet forwarding mechanism which uses a fast-switching cache to store next hop information.
- **Cisco Express Forwarding (CEF)** – The most recent, fastest, and preferred Cisco IOS packet-forwarding mechanism. Table entries are not packet-triggered like fast switching but change-triggered.

I- Introduction
Document Network Addressing

- Network Documentation should include at least the following in a topology diagram and addressing table:
  - Device names
  - Interfaces
  - IP addresses and subnet mask
  - Default gateways

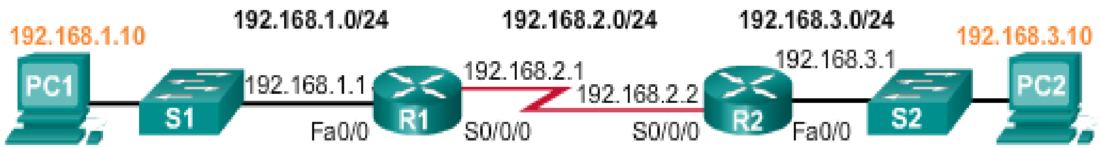

Device	Interface	IP Address	Subnet Mask	Default Gateway
R1	Fa0/0	192.168.1.1	255.255.255.0	N/A
	S0/0/0	192.168.2.1	255.255.255.0	N/A
R2	Fa0/0	192.168.3.1	255.255.255.0	N/A
	S0/0/0	192.168.2.2	255.255.255.0	N/A
PC1	N/A	192.168.1.10	255.255.255.0	192.168.1.1
PC2	N/A	192.168.3.10	255.255.255.0	192.168.3.1

I – Introduction
Router Switching Functions

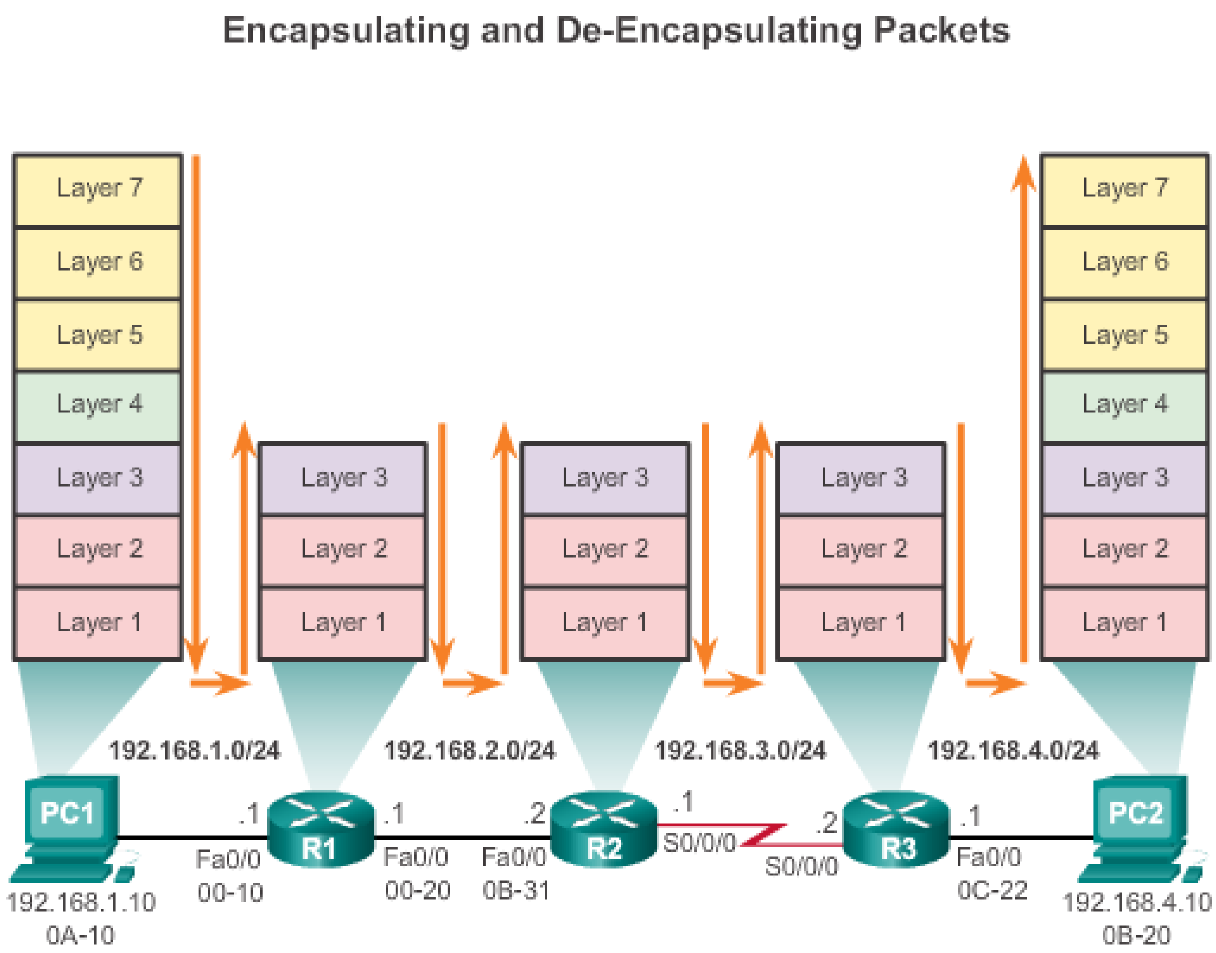

I – Introduction
Router Switching Functions

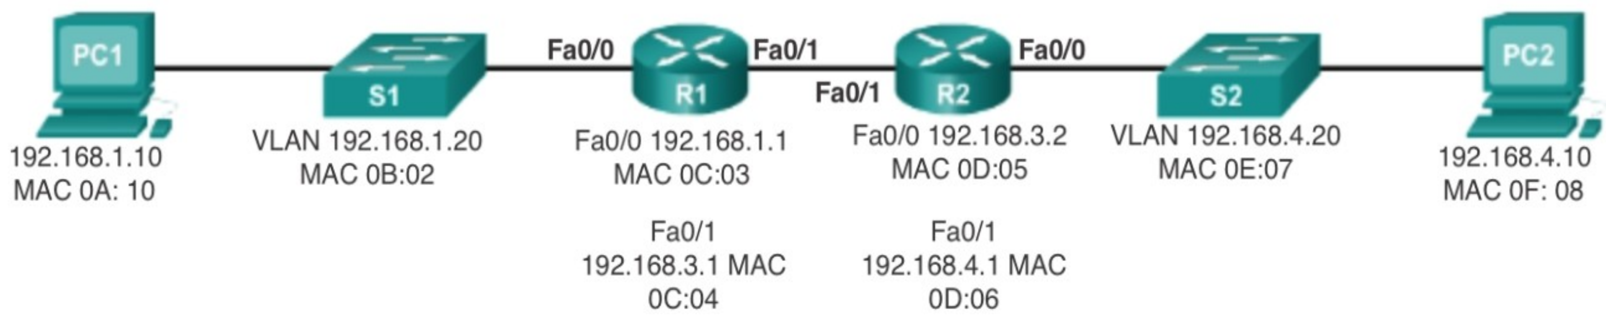

PC1 is sending data to PC2.
Use the blank frame to build the start data frame.
Noted that all divices are ARP complete.

Layer 2 Data Link Frame			Layer 3 Data Packet				
Destination MAC	Source MAC	Type 0x800	Source IP	Destination IP	IP Fields	Data	Trailer

I – Introduction
Routing Decisions

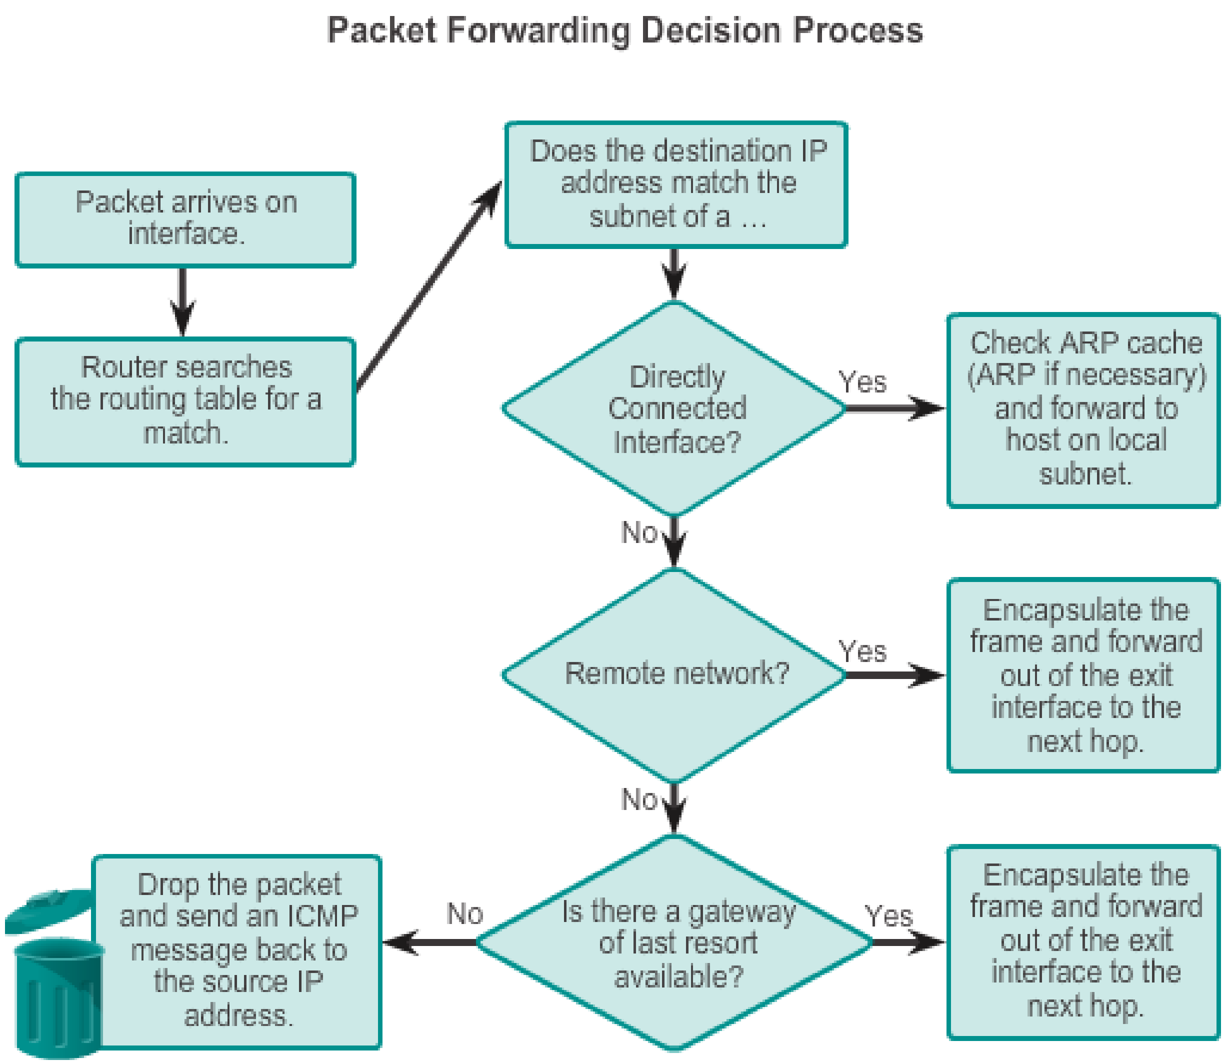

I – Introduction
Router Switching Functions

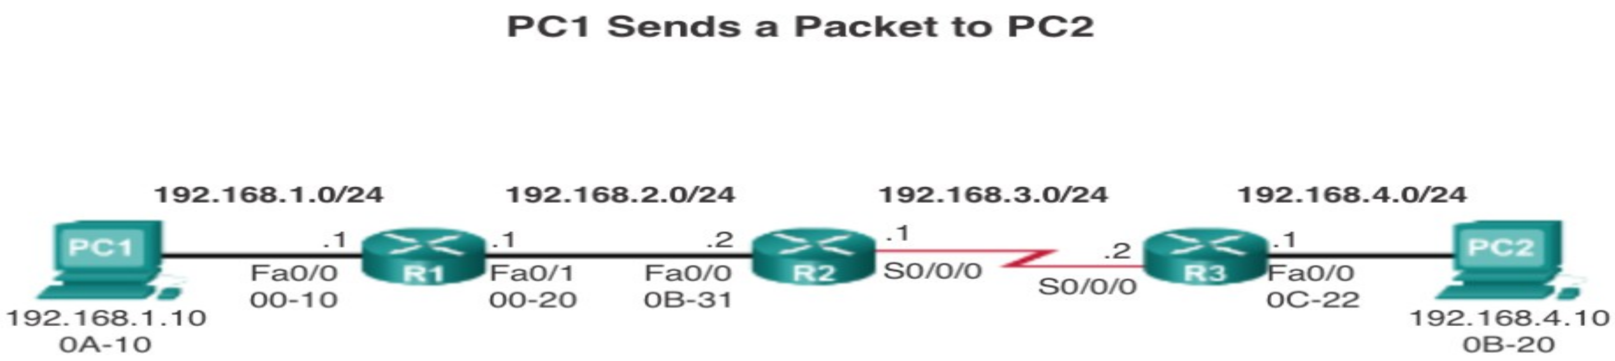

Build the frame between R1 and R2 ; R2 and R3

I – Introduction
Router Switching Functions

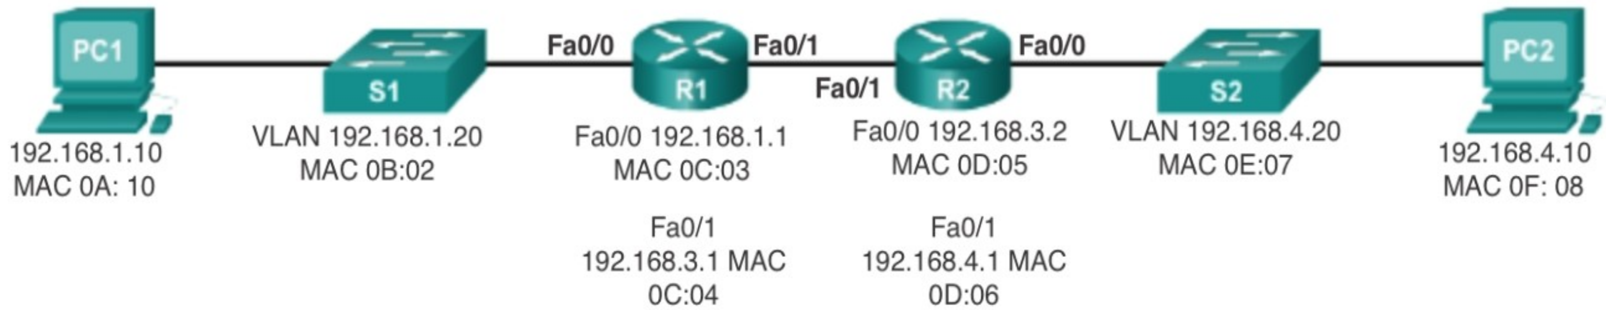

PC1 is sending data to PC2.
Build the frame to transfert data from R2 to PC2.

Layer 2 Data Link Frame			Layer 3 Data Packet				
Destination MAC	Source MAC	Type 0x800	Source IP	Destination IP	IP Fields	Data	Trailer

I – Introduction
Best Path

- Best path is selected by a routing protocol based on the value or metric it uses to determine the distance to reach a network:
  - A metric is the value used to measure the distance to a given network.
  - Best path to a network is the path with the lowest metric.

- Dynamic routing protocols use their own rules and metrics to build and update routing tables:
  - **Routing Information Protocol** (RIP) - Hop count
  - **Open Shortest Path First** (OSPF) - Cost based on cumulative bandwidth from source to destination
  - **Enhanced Interior Gateway Routing Protocol** (EIGRP) - Bandwidth, delay, load, reliability

Load Balancing

- When a router has two or more paths to a destination with equal cost metrics, then the router forwards the packets using both paths equally:
  - Equal cost load balancing can improve network performance.
  - Equal cost load balancing can be configured to use both dynamic routing protocols and static routes.
- RIP, OSPF and EIGRP support equal cost load balancing.

I – Introduction
Administrative Distance

- If multiple paths to a destination are configured on a router, the path installed in the routing table is the one with the lowest Administrative Distance (AD):
  - A static route with an AD of 1 is more reliable than an EIGRP-discovered route with an AD of 90.
  - A directly connected route with an AD of 0 is more reliable than a static route with an AD of 1.

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

I – Introduction
The Routing Table

- A routing table is a file stored in RAM that contains information about:
  1. Directly connected routes
  2. Remote routes
  3. Network or next hop associations

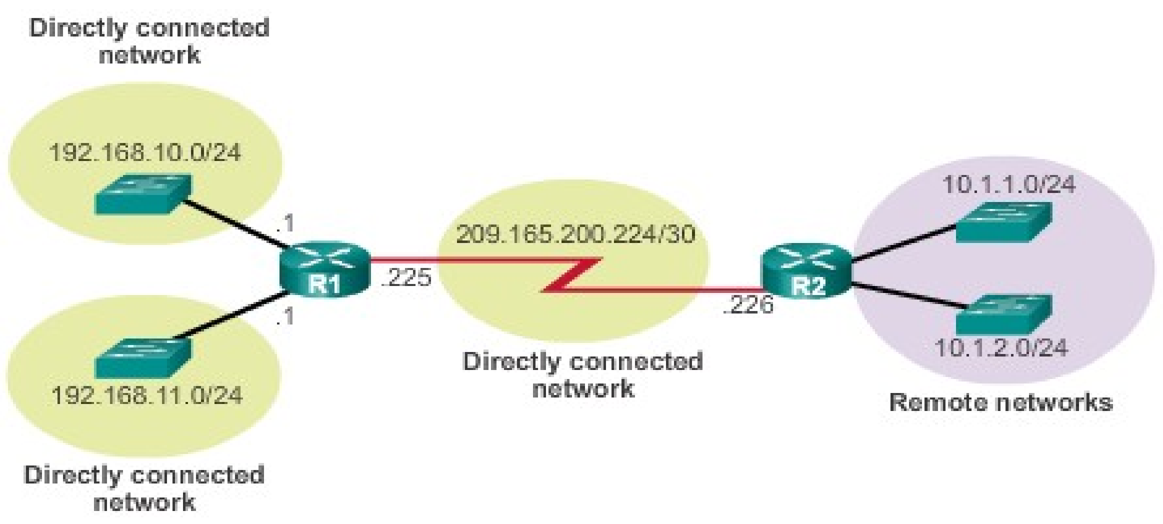

Static Routes

- Static routes:
  - define an explicit path between two networking devices.
  - must be manually updated if the topology changes.
  - improved security and control of resources.

- A default static route is used when the routing table does not contain a path for a destination network.

I – Introduction
Dynamic Routing

Dynamic routing is used by routers to share information about the reachability and status of remote networks. It performs network discovery and maintains routing tables.

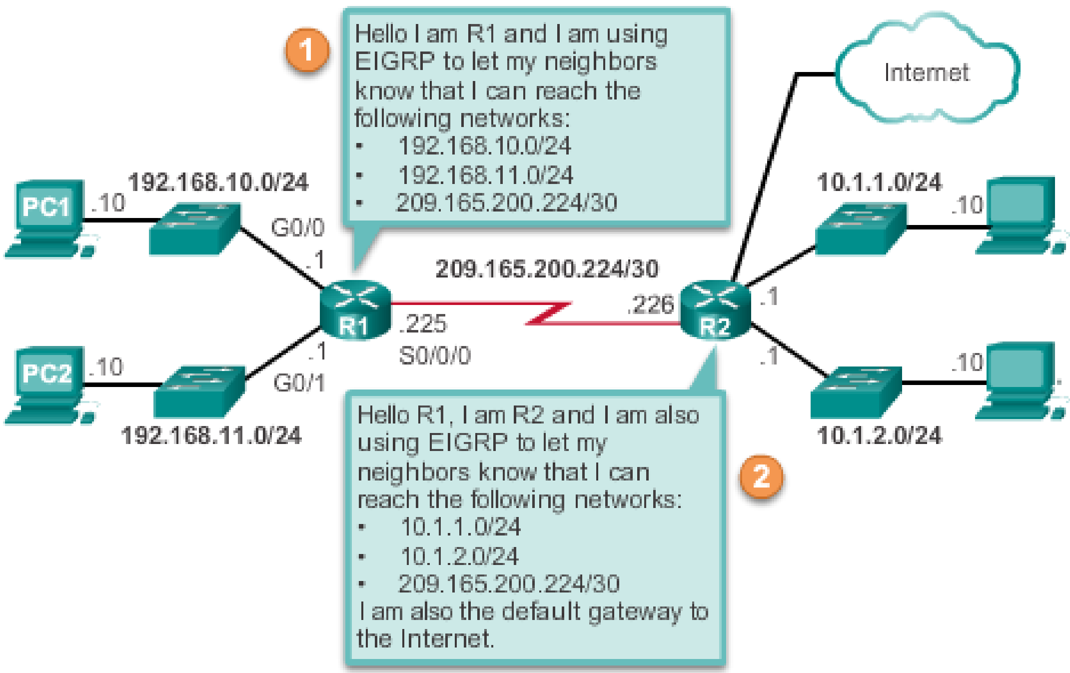

I – Introduction
Routing Protocols

- Recent routers can support a variety of dynamic routing protocols

- IPv4 routing Protocols
  - **RIP –** Routing Information Protocol
  - **EIGRP –** Enhanced Interior Gateway Routing Protocol
  - **OSPF –** Open Shortest Path First
  - **IS-IS –** Intermediate System-to-Intermediate System

- IPv6 routing Protocols :
  - **RIPng** - RIP next generation
  - **EIGRP for IPv6**
  - **OSPFv3**
  - **MP-BGP4** - Multicast Protocol-Border Gateway Protocol

II – Static Routing

II – Static Routing
Why Use Static Routing?

- Pros
  - Better security (no advertisement over the network)
  - Less bandwidth and CPU used
  - The path that a static route uses to send data is known.
- Cons
  - Time-consuming (initial configuration and maintenance).
  - Error-prone configuration
  - No auto-adaptation (Administrator intervention required to change route information)
  - Does not scale well with growing networks; maintenance becomes cumbersome.
  - Requires complete knowledge of the whole network for proper implementation.

II – Static Routing
Static Route Applications

- Static routing has three primary uses:
  - **Providing ease of routing table maintenance** in smaller networks that are not expected to grow significantly.
  - **Routing to and from stub networks**. A stub network is a network accessed by a single route, and the router has no other neighbors.
  - **Using a single default** route to represent a path to any network that does not have a more specific match with another route in the routing table. Default routes are used to send traffic to any destination beyond the next upstream router.

- Static Routes are often used to:
  - Connect to a specific network.
  - Provide a Gateway of Last Resort for a stub network.
  - Reduce the number of routes advertised by summarizing several contiguous networks as one static route.
  - Create a backup route in case a primary route link fails.

II – Static Routing
Static Route Applications

- Default static route:
  - matches all packets.
  - identifies the gateway IP address to which the router sends all IP packets that it does not have a learned or static route.
  - is simply a static route with 0.0.0.0/0 as the destination IPv4 address.

II – Static Routing
Static Route Applications

- Summary Static Route

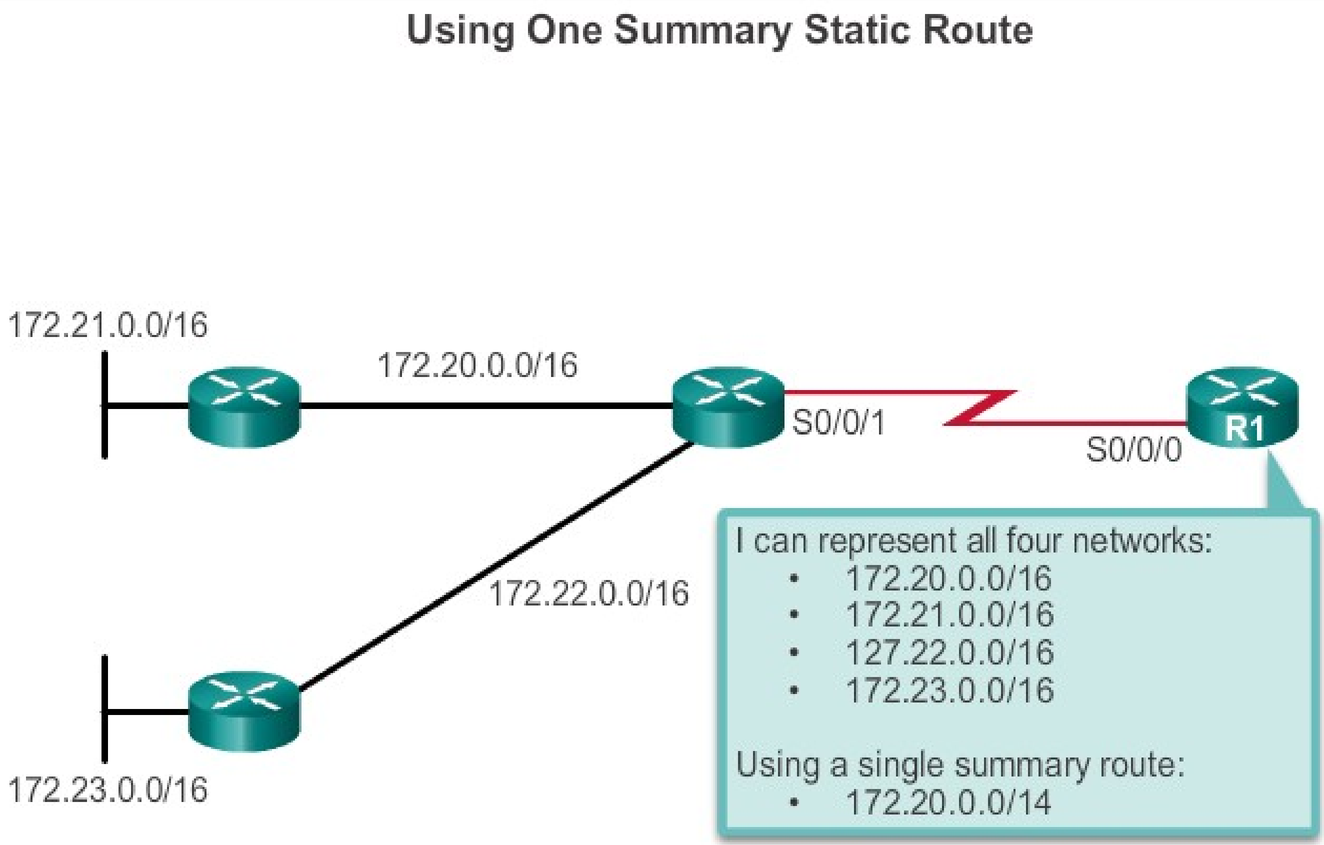

II – Static Routing
Static Route Applications

- Floating Static Route:
  - are static routes that are used to provide a backup path to a primary static or dynamic route, in the event of a link failure.
  - is only used when the primary route is not available.
  - is configured with a higher administrative distance than the primary route.

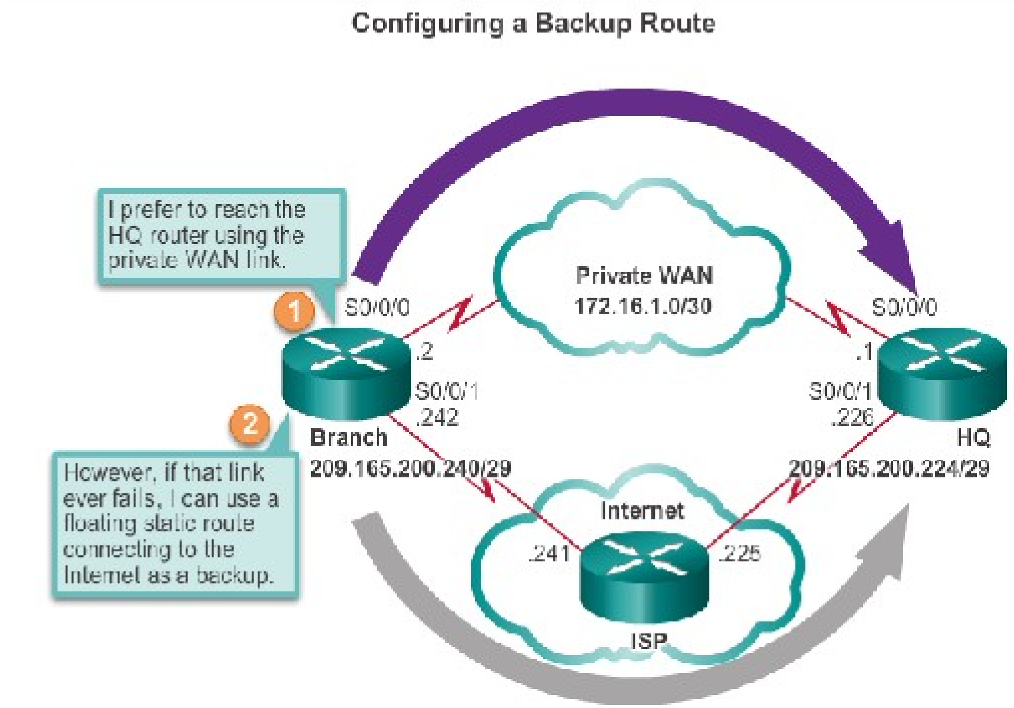

II – Static Routing
Next-Hop Options

- The next hop can be identified by an IP address, exit interface, or both.

- How the destination is specified creates one of the three following route types:
  - **Next-hop route** - Only the next-hop IP address is specified.
  - **Directly connected static route** - Only the router exit interface is specified.
  - **Fully specified static route** - The next-hop IP address and exit interface are specified.

II – Static Routing
Next-Hop Static Route

Verify the Routing Table of R1

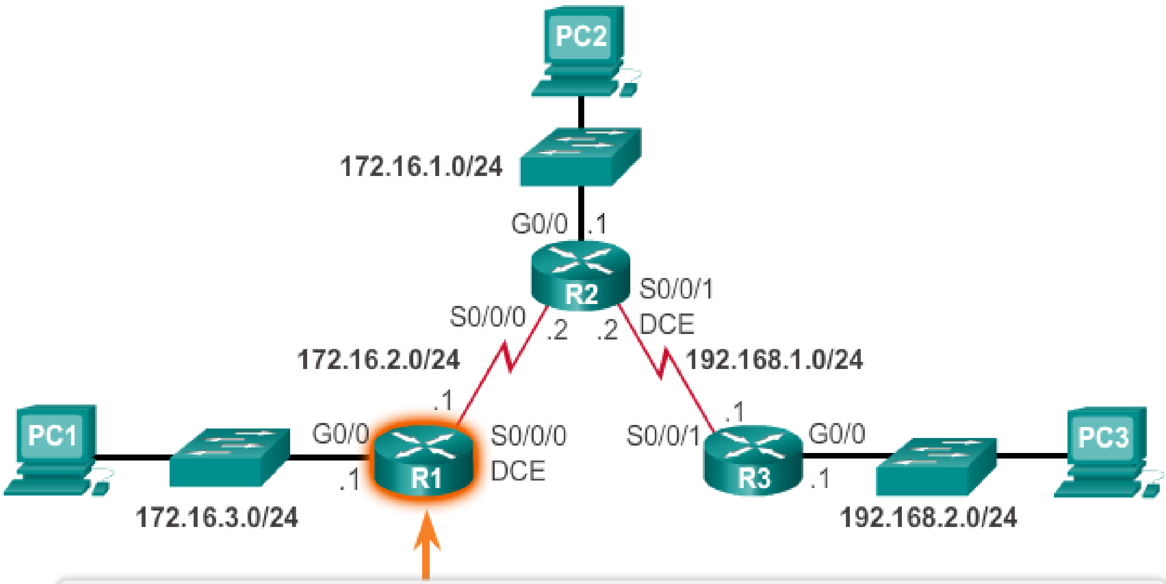


```
S    172.16.1.0/24 [1/0] via 172.16.2.2
C    172.16.2.0/24 is directly connected, Serial0/0/0
L    172.16.2.1/32 is directly connected, Serial0/0/0
C    172.16.3.0/24 is directly connected, GigabitEthernet0/0
L    172.16.3.1/32 is directly connected, GigabitEthernet0/0
S    192.168.1.0/24 [1/0] via 172.16.2.2
S    192.168.2.0/24 [1/0] via 172.16.2.2
R1#
```

II – Static Routing
Configure Directly Connected Static Route

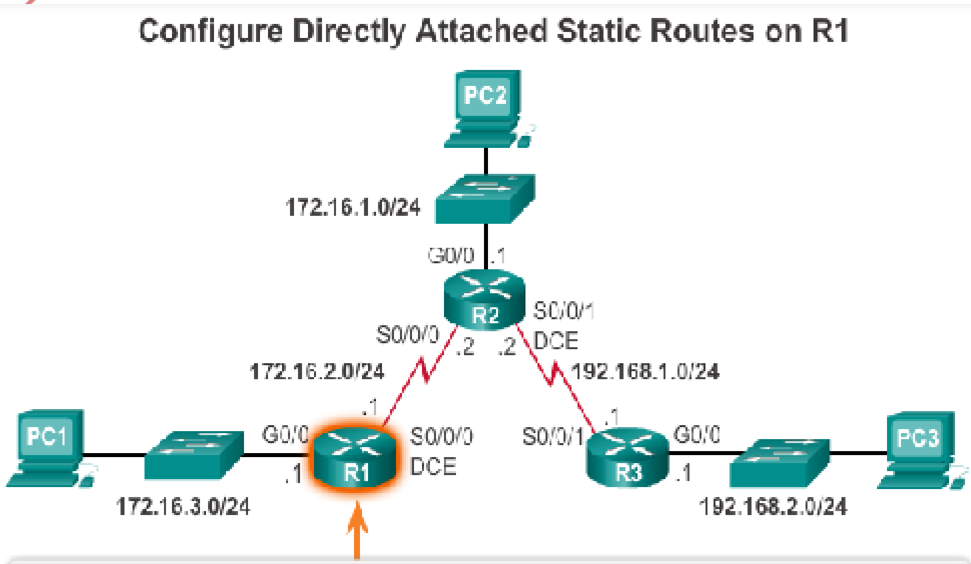


```
R1(config)#ip route 172.16.1.0 255.255.255.0 s0/0/0
R1(config)#ip route 192.168.1.0 255.255.255.0 s0/0/0
R1(config)#ip route 192.168.2.0 255.255.255.0 s0/0/0
R1(config)#
```

```
S        172.16.1.0/24 is directly connected, Serial0/0/0
C        172.16.2.0/24 is directly connected, Serial0/0/0
L        172.16.2.1/32 is directly connected, Serial0/0/0
C        172.16.3.0/24 is directly connected, GigabitEthernet0/0
L        172.16.3.1/32 is directly connected, GigabitEthernet0/0
S     192.168.1.0/24 is directly connected, Serial0/0/0
S     192.168.2.0/24 is directly connected, Serial0/0/0
R1#
```

II – Static Routing
Fully Specified Static Route

- In a fully specified static route:
  - Both the output interface and the next-hop IP address are specified.
  - This is another type of static route that is used in older IOSs, prior to CEF.
  - This form of static route is used when the output interface is a multi-access interface and it is necessary to explicitly identify the next hop.
  - The next hop must be directly connected to the specified exit interface.

II – Static Routing
ClassFul Addressing

- Classful Network Addressing

Class	High Order Bits	Start	End
Class A	0xxxxxxx	0.0.0.0	127.255.255.255
Class B	10xxxxxx	128.0.0.0	191.255.255.255
Class C	110xxxxx	192.0.0.0	223.255.255.255
Multicast	1110xxxx	224.0.0.0	239.255.255.255
Reserved	1111xxxx	240.0.0.0	255.255.255.255

II – Static Routing
ClassFul Addressing

- Classful Routing Protocol Example

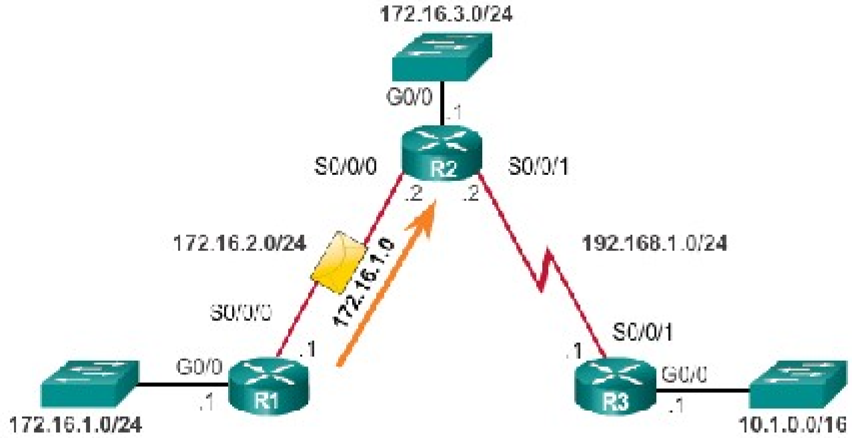

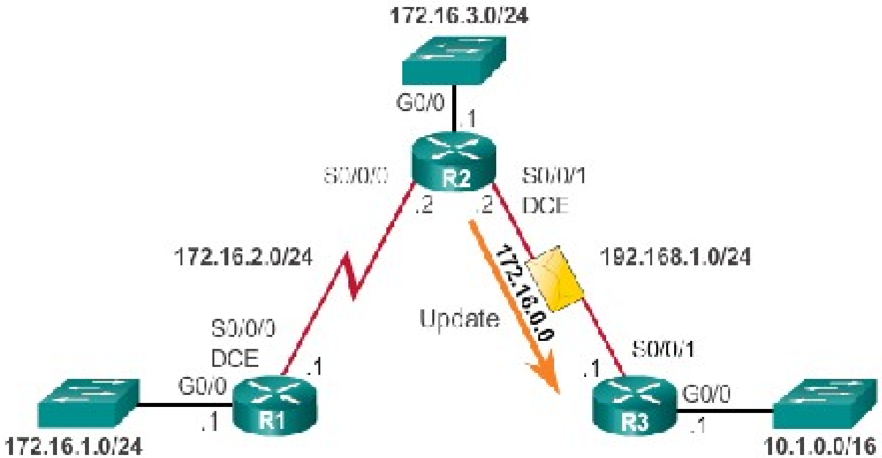

II – Static Routing
ClassFul Addressing

- Classful Network Addressing

Class A

	Network	Host	Host	Host
Subnet mask	255	.0	.0	.0

Class B

	Network	Network	Host	Host
Subnet mask	255	.255	.0	.0

Class C

	Network	Network	Network	Host
Subnet mask	255	.255	.255	.0

II – Static Routing
ClassFul Addressing

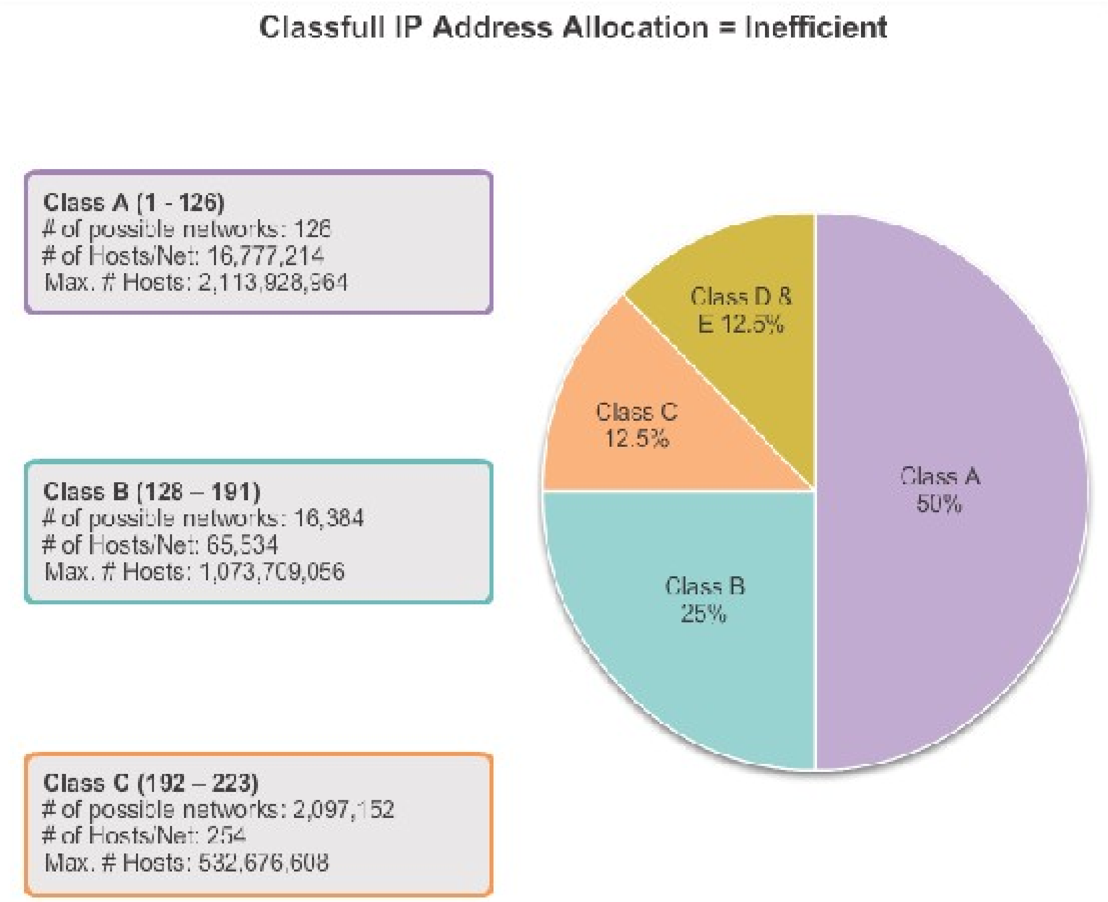

II – Static Routing
CIDR

- Classless Inter-Domain Routing (CIDR)

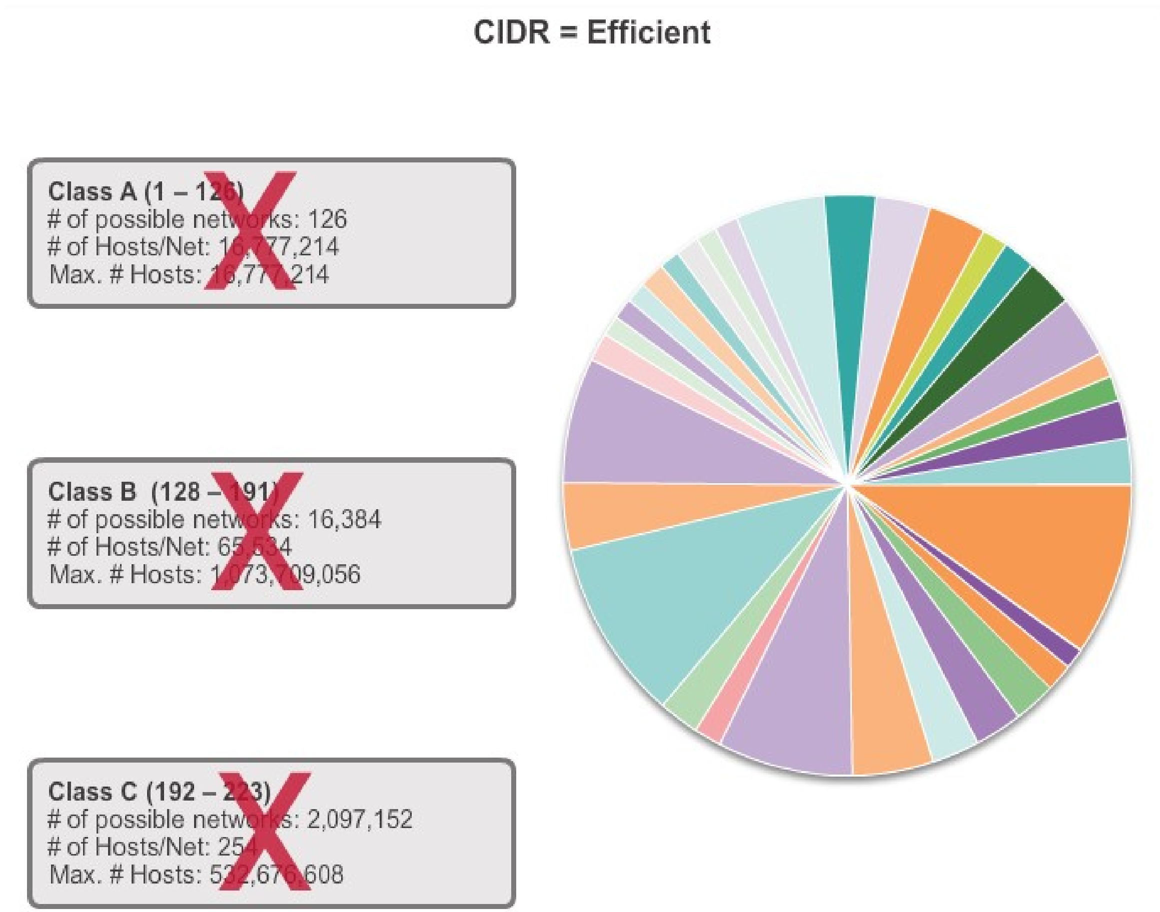

- CIDR and Route Summarization

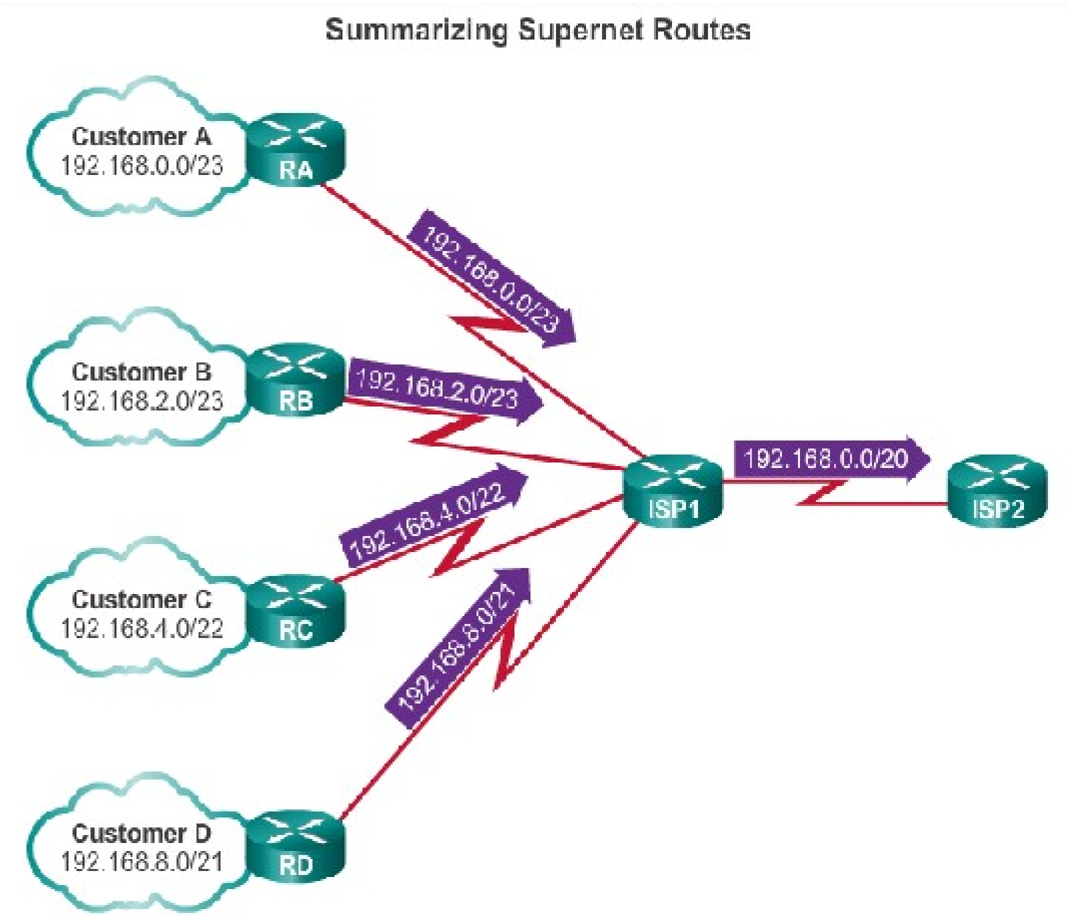

II – Static Routing
CIDR

One Summary Static Route

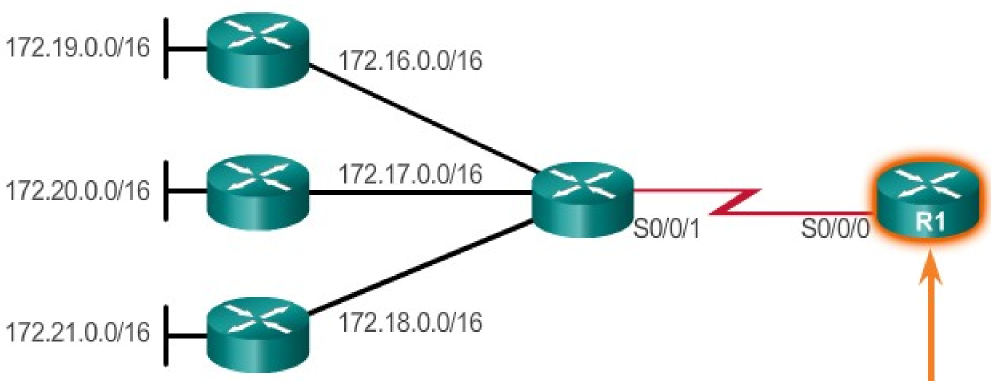


```
R1(config)#no ip route 172.16.0.0 255.255.0.0 s0/0/0
R1(config)#no ip route 172.17.0.0 255.255.0.0 s0/0/0
R1(config)#no ip route 172.18.0.0 255.255.0.0 s0/0/0
R1(config)#no ip route 172.19.0.0 255.255.0.0 s0/0/0
R1(config)#no ip route 172.20.0.0 255.255.0.0 s0/0/0
R1(config)#no ip route 172.21.0.0 255.255.0.0 s0/0/0
R1(config)#
R1(config)#ip route 172.16.0.0 255.248.0.0 s0/0/0
R1(config)#
```

II – Static Routing
CIDR

- Classless Routing Protocol Example

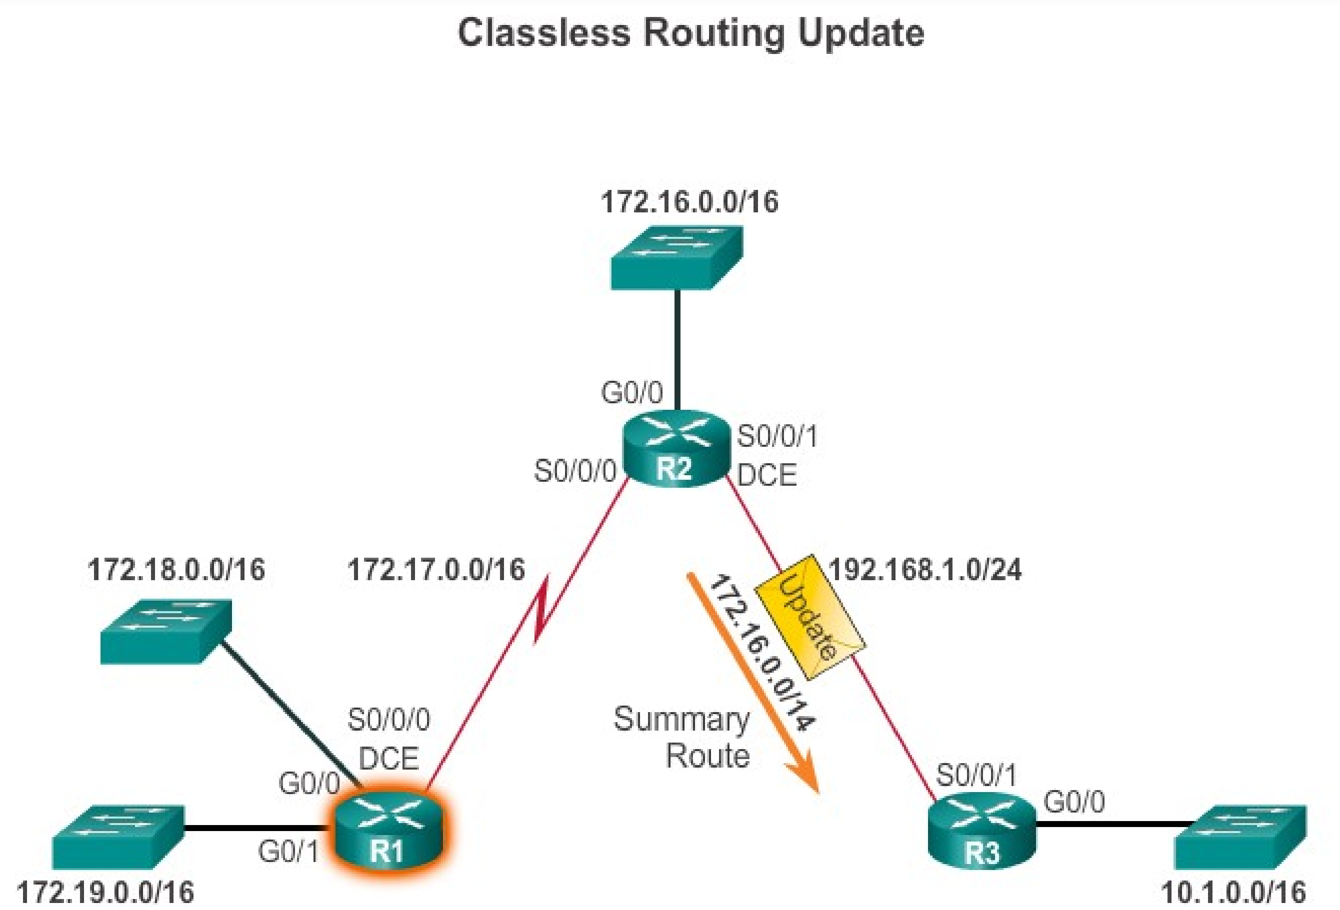

II – Static Routing
FLSM

- Fixed Length Subnet Masking

Basic Subnet Scheme

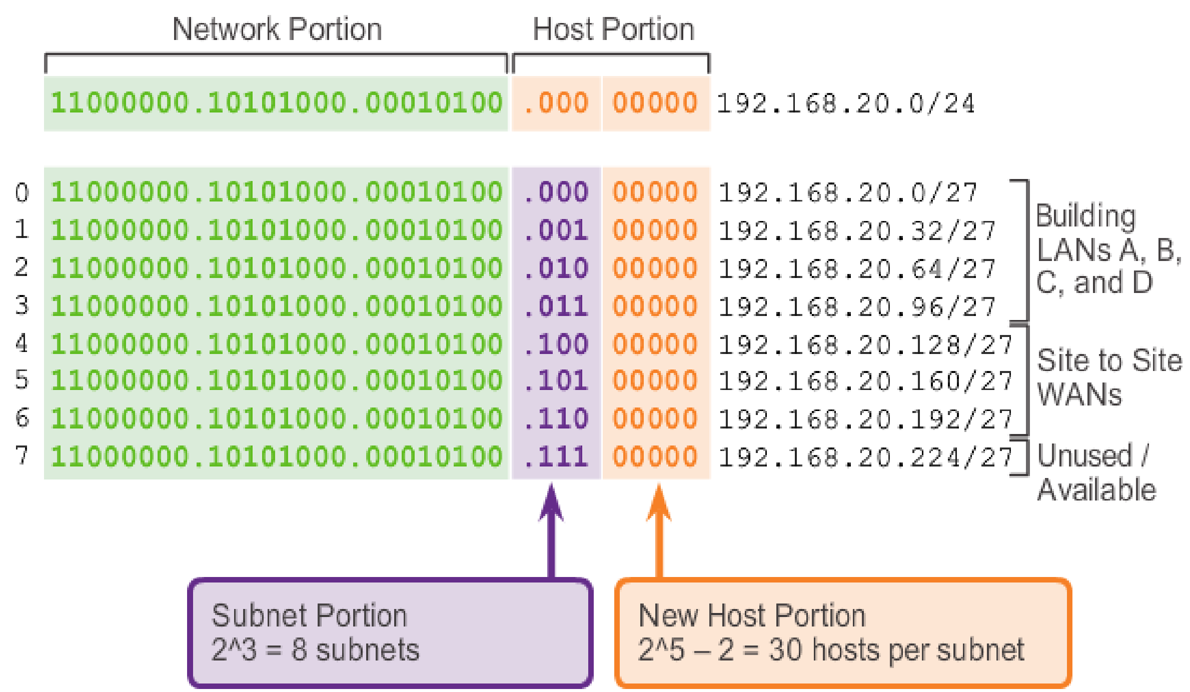

II – Static Routing
VLSM

- VLSM (Variable-Length Subnet Masking) allows the use of different masks for each subnet:
  - After a network address is subnetted, those subnets can be further subnetted.
  - VLSM is simply subnetting a subnet. VLSM can be thought of as sub-subnetting.
  - Individual host addresses are assigned from the addresses of "sub-subnets".

Subnetting the Subnet 10.2.0.0/16 to 10.2.0.0/24

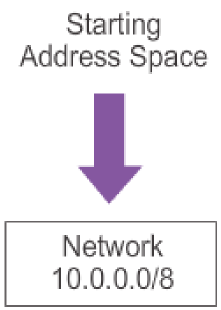

1st Round of Subnets
Subnets
10.0.0.0/16
10.1.0.0/16
10.2.0.0/16
10.3.0.0/16
10.4.0.0/16
10.5.0.0/16
.
.
.
10.255.0.0/16

256 Subnets

Subnets of the Subnet
Sub-Subnets
10.2.0.0/24
10.2.1.0/24
10.2.2.0/24
10.2.3.0/24
10.2.4.0/24
10.2.5.0/24
.
.
.
10.2.255.0/24

256 Subnets

II – Static Routing

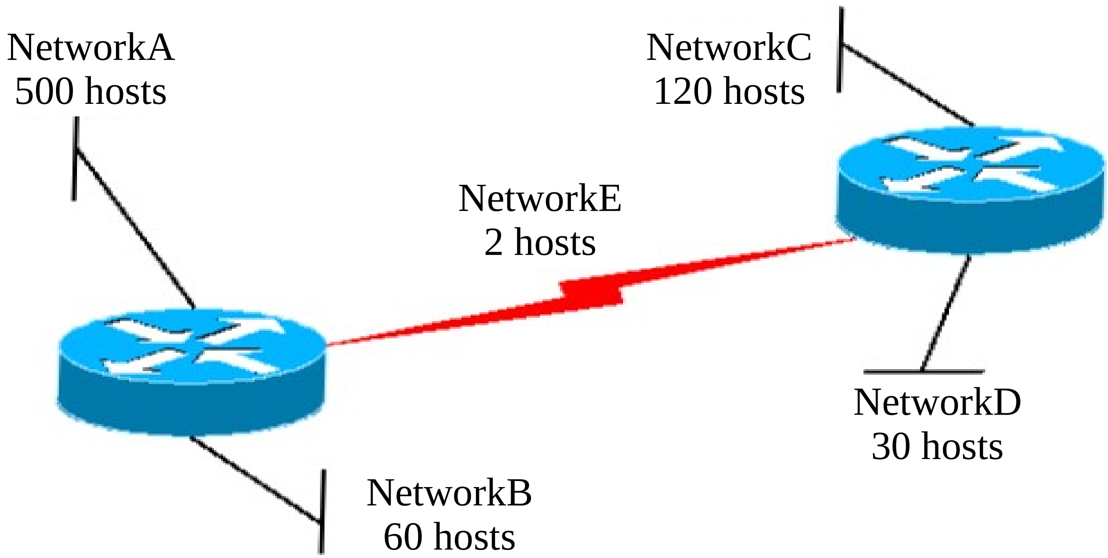

- The network address is 192.168.20.0/22

  Determine the addresses and masks of the different subnet using VLSM.

II – Static Routing

II – Static Routing
Route Summarization

- Route summarization, also known as route aggregation, is the process of advertising a contiguous set of addresses as a single address with a less-specific, shorter subnet mask:
  - CIDR is a form of route summarization and is synonymous with the term supernetting.
  - CIDR ignores the limitation of classful boundaries, and allows summarization with masks that are smaller than that of the default classful mask.
  - This type of summarization helps reduce the number of entries in routing updates and lowers the number of entries in local routing tables.

II – Static Routing
Route Summarization

Calculating a Route Summary

Step 1: List networks in binary format.

172.20.0.0	10101100 . 00010100 . 00000000 . 00000000
172.21.0.0	10101100 . 00010101 . 00000000 . 00000000
172.22.0.0	10101100 . 00010110 . 00000000 . 00000000
172.23.0.0	10101100 . 00010111 . 00000000 . 00000000

Step 2: Count the number of far-left matching bits to determine the mask.

Answer: 14 matching bits = /14 or 255.252.0.0

Step 3: Copy the matching bits and add zero bits to determine the network address.

10101100 . 00010100 . 00000000 . 00000000

Copy — Add zero bits

Answer: 172.20.0.0

II – Static Routing
Route Summarization

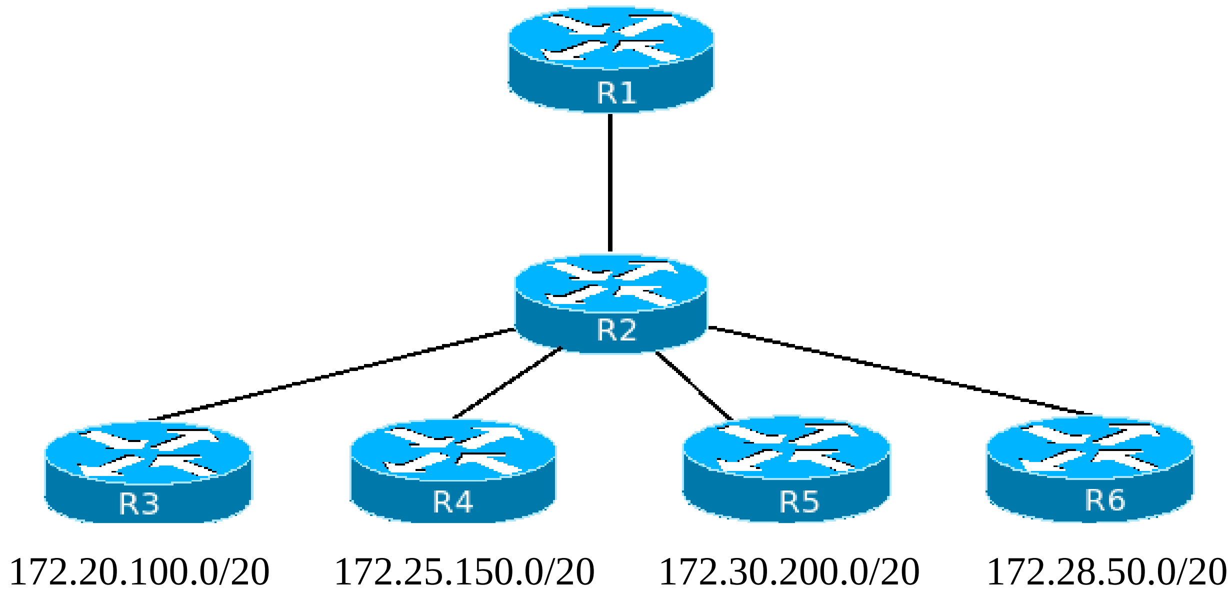

R1 needs to reach all networks shown. Determine the summary network address and prefix.

III – Dynamic Routing

III – Dynamic Routing
Evolution of Dynamic Routing Protocols

- Dynamic routing protocols used in networks since the late 1980s
- Newer versions support the communication based on IPv6

Routing Protocols Classification

	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector		Link-State		Path Vector
IPv4	RIPv2	EIGRP	OSPFv2	IS-IS	BGP-4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGP-MP

III – Dynamic Routing

Purpose of Dynamic Routing Protocols

- Routing Protocols are used to facilitate the exchange of routing information between routers.

- The purpose of dynamic routing protocols includes:
  - Discovery of remote networks
  - Maintaining up-to-date routing information
  - Choosing the best path to destination networks
  - Ability to find a new best path if the current path is no longer available

- Main components of dynamic routing protocols include:
  - **Data structures -** Routing protocols typically use tables or databases for its operations. This information is kept in RAM.
  - **Routing protocol messages -** Routing protocols use various types of messages to discover neighboring routers, exchange routing information, and other tasks to learn and maintain accurate information about the network.
  - **Algorithm -** Routing protocols use algorithms for facilitating routing information for best path determination.

Purpose of Dynamic Routing Protocols

Example of EIGRP: Enhanced Interior-Gateway Routing Protocol

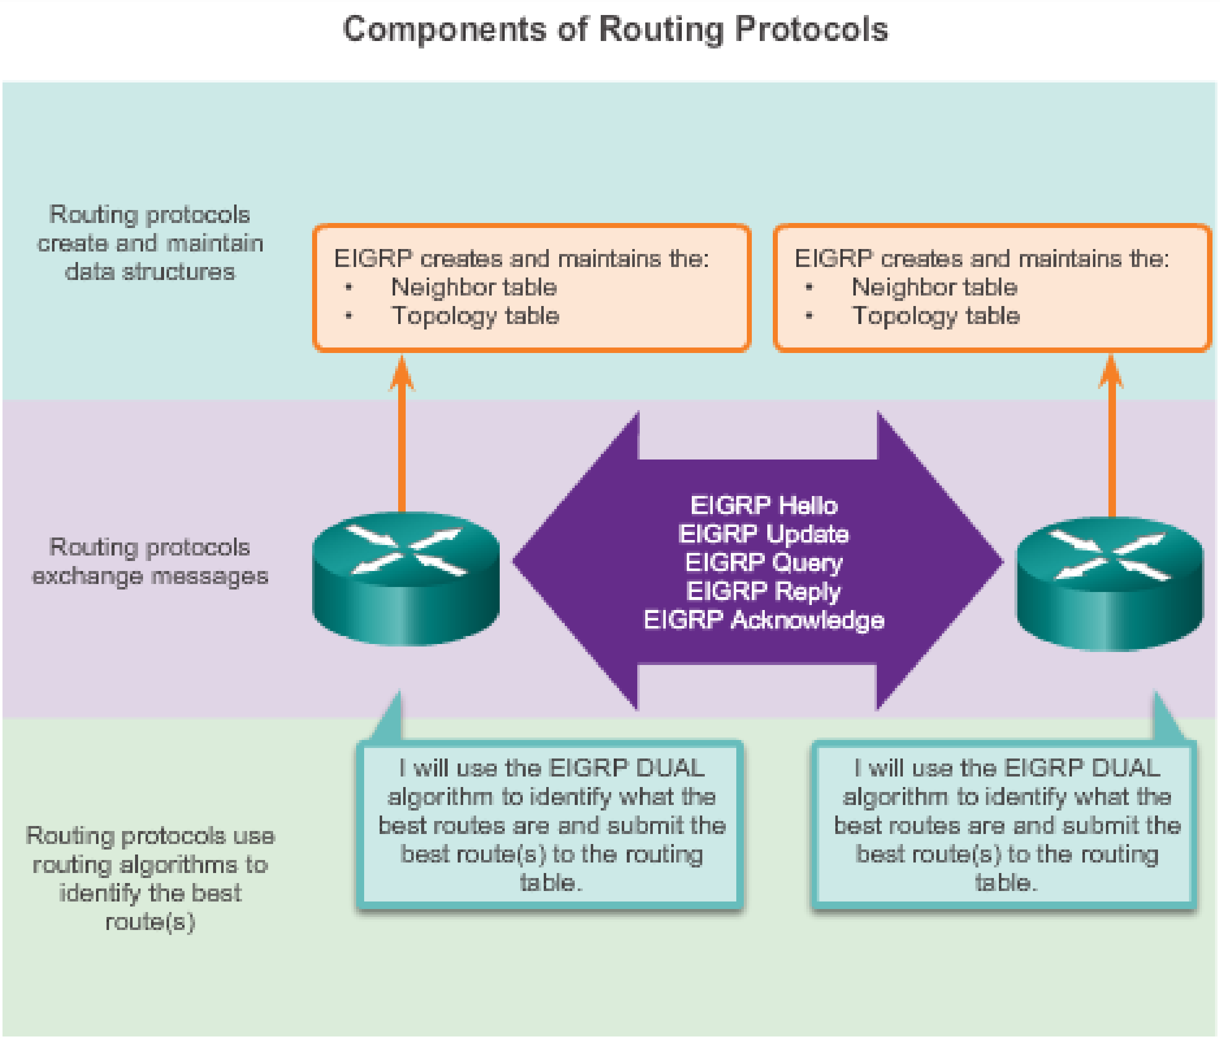

Role of Dynamic Routing Protocols

Pros

- Automatically share information about remote networks
- Determine the best path and add this information to their routing tables
- Less administrative overhead
- Less time-consuming process of configuring and maintaining static routes

Cons

- Resource consumption (CPU time and network bandwidth)
- Convergence Time

III – Dynamic Routing
Dynamic vs Static Routing

- Networks typically use a combination of both static and dynamic routing.

- Static routing has several primary uses:
  - Providing ease of routing table maintenance in smaller networks.
  - Routing to and from a stub network.
  - Accessing a single default router.

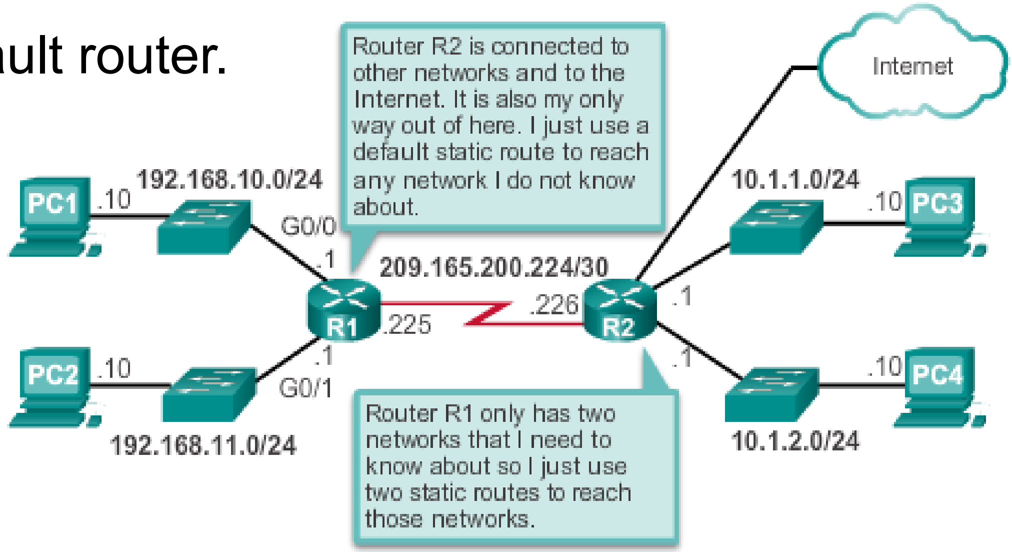

III – Dynamic Routing
Static Routing Scorecard

Static Routing Advantages and Disadvantages

Advantages	Disadvantages
Easy to implement in a small network.	Suitable only for simple topologies or for special purposes such as a default static route. Configuration complexity increases dramatically as network grows.
Very secure. No advertisements are sent as compared to dynamic routing protocols.	
Route to destination is always the same.	Manual intervention required to re-route traffic.
No routing algorithm or update mechanism required; therefore, extra resources (CPU or RAM) are not required.	

III – Dynamic Routing
Dynamic Routing Scorecard

Dynamic Routing Advantages and Disadvantages

Advantages	Disadvantages
Suitable in all topologies where multiple routers are required.	Can be more complex to implement.
Generally independent of the network size.	Less secure. Additional configuration settings are required to secure.
Automatically adapts topology to reroute traffic if possible.	Route depends on the current topology.
	Requires additional CPU, RAM, and link bandwidth.

Dynamic Routing Protocol Operation

In general, the operations of a dynamic routing protocol can be described as follows:

1. The router sends and receives routing messages on its interfaces.
2. The router shares routing messages and routing information with other routers that are using the same routing protocol.
3. Routers exchange routing information to learn about remote networks.
4. When a router detects a topology change the routing protocol can advertise this change to other routers.

III – Dynamic Routing
Cold Start

Directly Connected Networks Detected

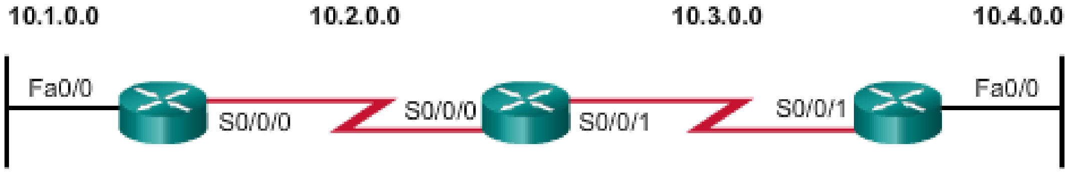

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0

Routers running RIPv2

1. Routers add the network available through local interfaces

III – Dynamic Routing
Network Discovery

1. Routers add the network available through local interfaces
2. Routers sends and receives update from and to neighbors

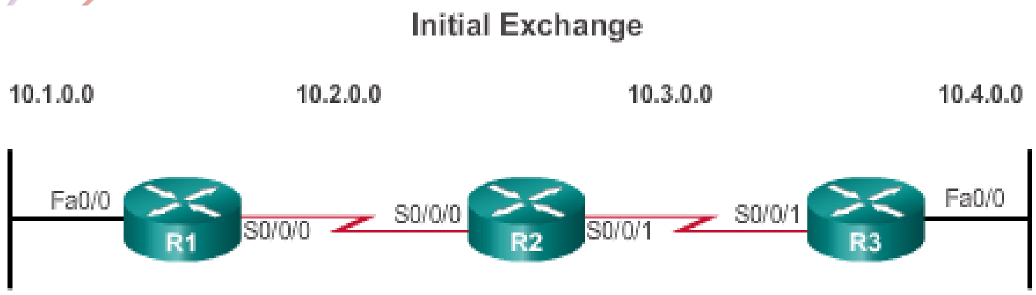

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/0	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1

Routers running RIPv2

III – Dynamic Routing
Exchanging the Routing Information

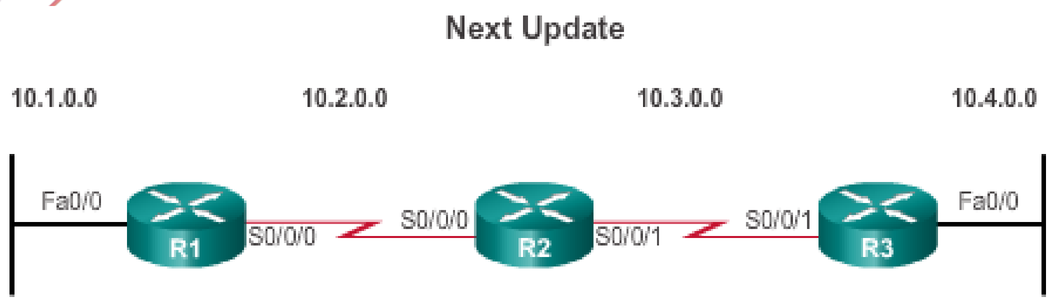

Network	Interface	Hop
10.1.0.0	Fa0/0	0
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/0	1
10.4.0.0	S0/0/0	2

Network	Interface	Hop
10.2.0.0	S0/0/0	0
10.3.0.0	S0/0/1	0
10.1.0.0	S0/0/0	1
10.4.0.0	S0/0/1	1

Network	Interface	Hop
10.3.0.0	S0/0/1	0
10.4.0.0	Fa0/0	0
10.2.0.0	S0/0/1	1
10.1.0.0	S0/0/1	2

Routers running RIPv2

1. Routers add the network available through local interfaces
2. Routers sends updates containing directed networks to neighbor Routers
3. Routers receives update messages from neighborhood and completes local routing table.

III – Dynamic Routing
Achieving Convergence

- The network is converged when all routers have complete and accurate information about the entire network:
  - Convergence time: the time it takes routers to share information, calculate best paths, and update their routing tables.
  - A network is not completely operable until the network has converged.
  - Convergence properties include:
    - speed of propagation of routing information
    - calculation of optimal paths.
  - Generally, older protocols, such as RIP, are slow to converge, whereas modern protocols, such as EIGRP and OSPF, converge more quickly.

III – Dynamic Routing
Classifying Routing Protocols

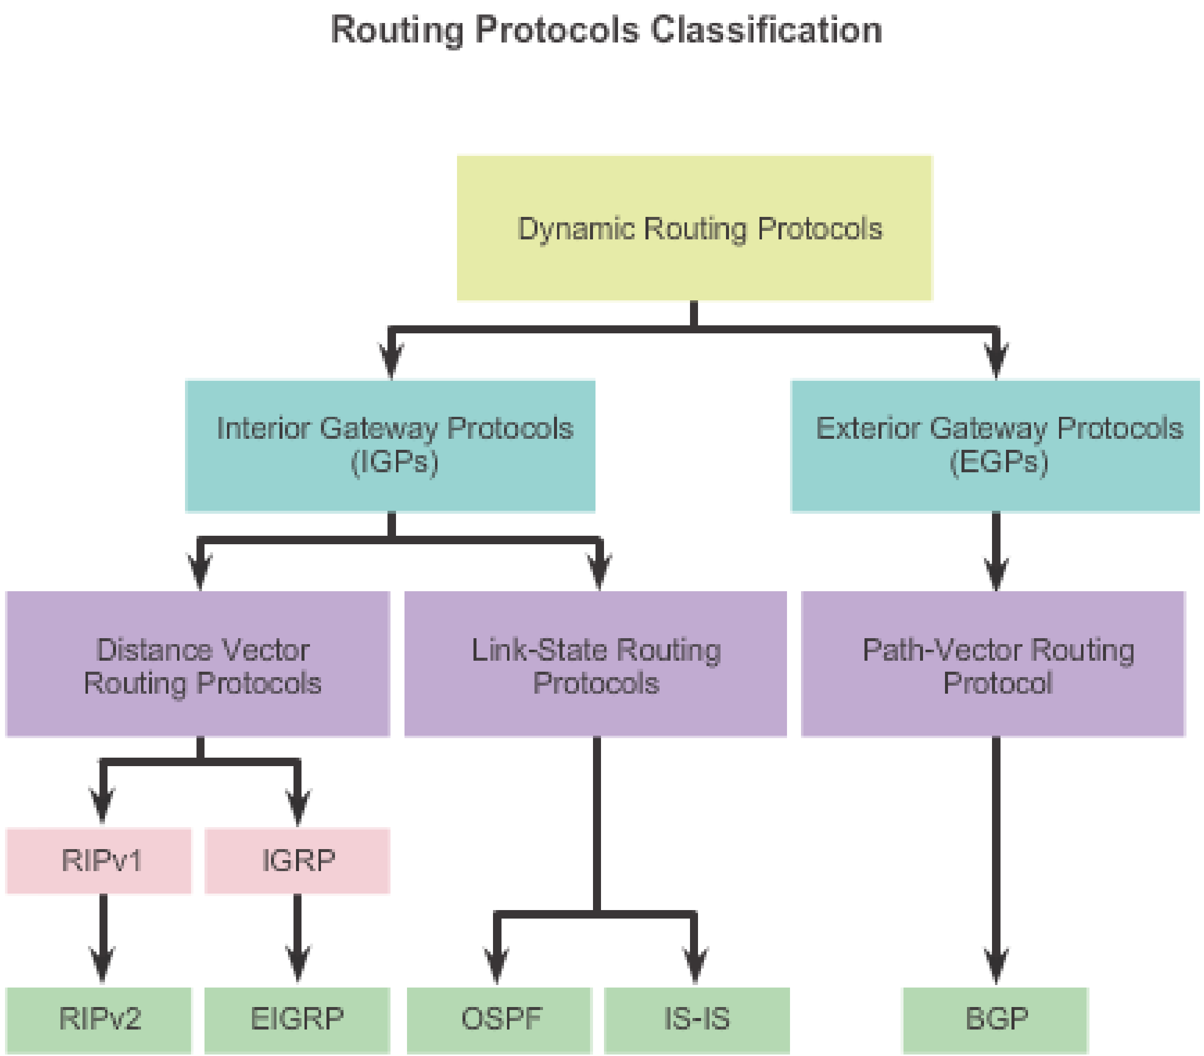

III – Dynamic Routing
IGP and EGP Routing Protocols

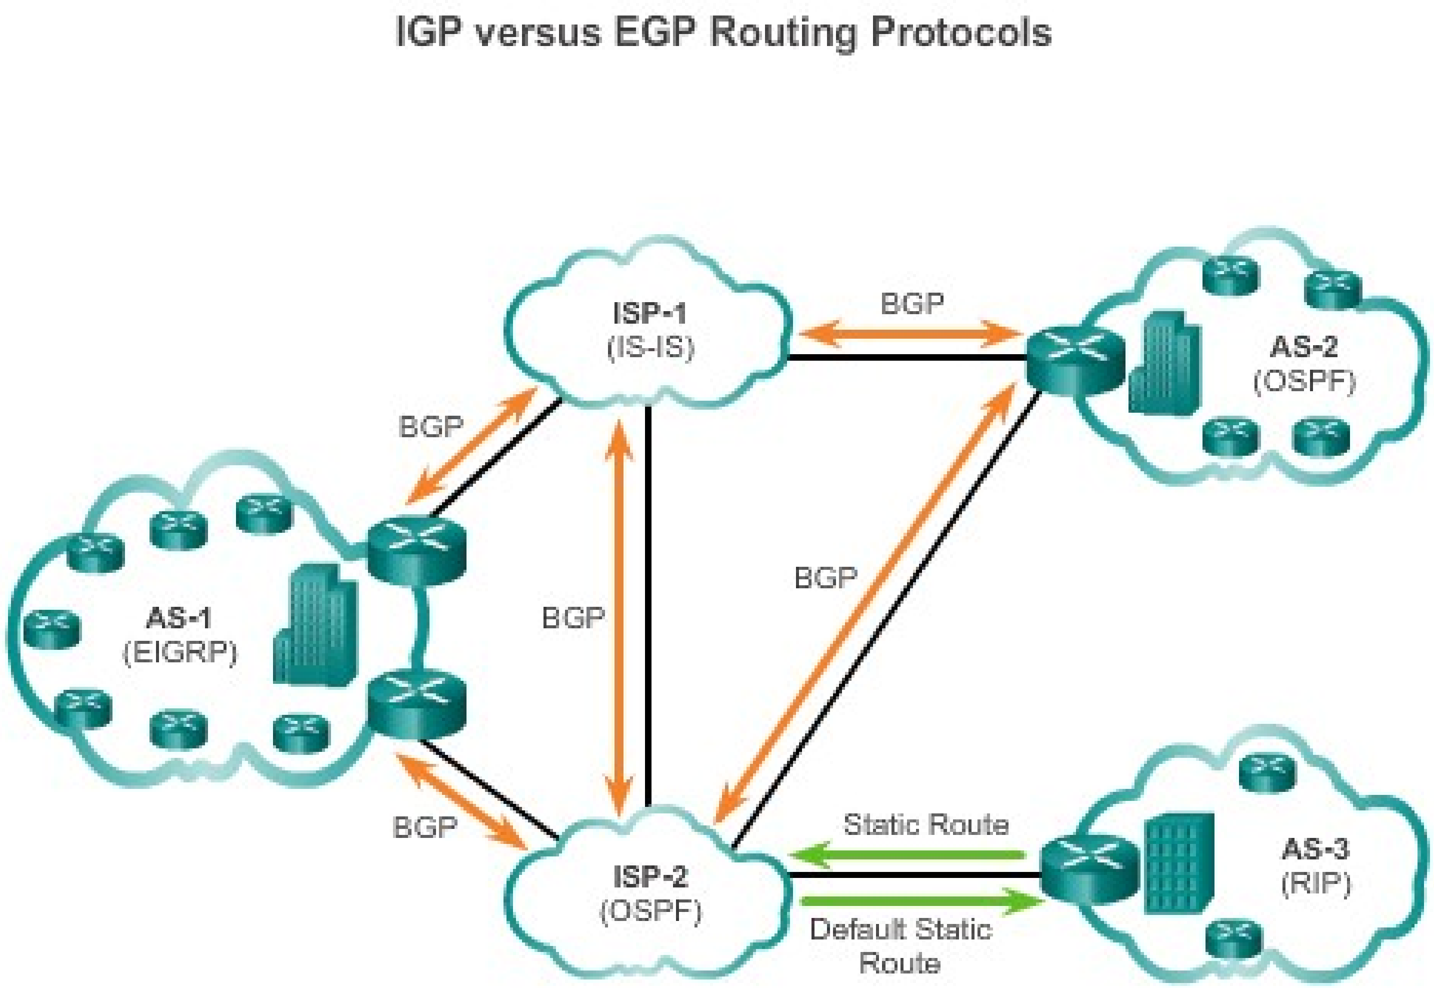

III – Dynamic Routing
Distance Vector or Link-State

Distance vector protocols use routers as sign posts along the path to the final destination.

A link-state routing protocol is like having a complete map of the network topology. The sign posts along the way from source to destination are not necessary, because all link-state routers are using an identical map of the network. A link-state router uses the link-state information to create a topology map and to select the best path to all destination networks in the topology.

III – Dynamic Routing
Distance Vector Routing Protocols

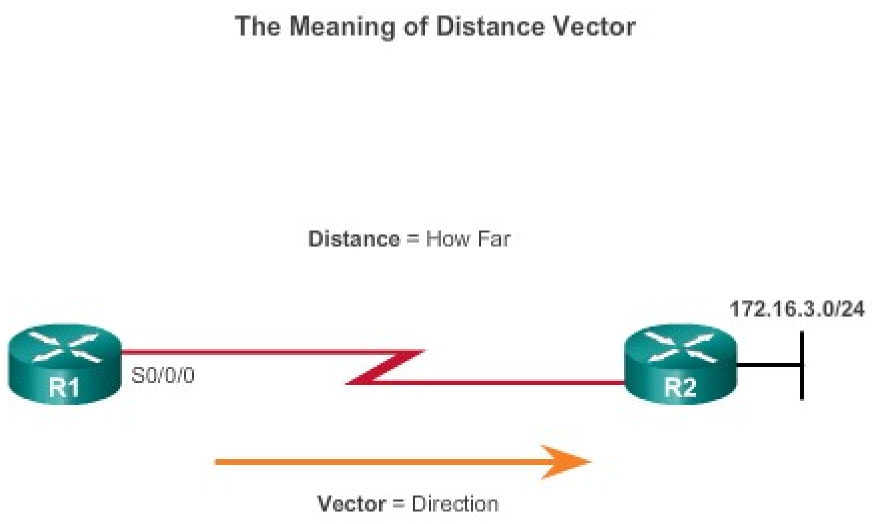

For R1, 172.16.3.0/24 is one hop away (distance). It can be reached through R2 (vector).

Distance vector IPv4 IGPs:

- **RIPv1** - First generation legacy protocol

- **RIPv2** - Simple distance vector routing protocol
- **IGRP** - First generation Cisco proprietary protocol (obsolete)
- **EIGRP** - Advanced version of distance vector routing

III – Dynamic Routing
Link- State Routing Protocols

Link-State Protocol Operation

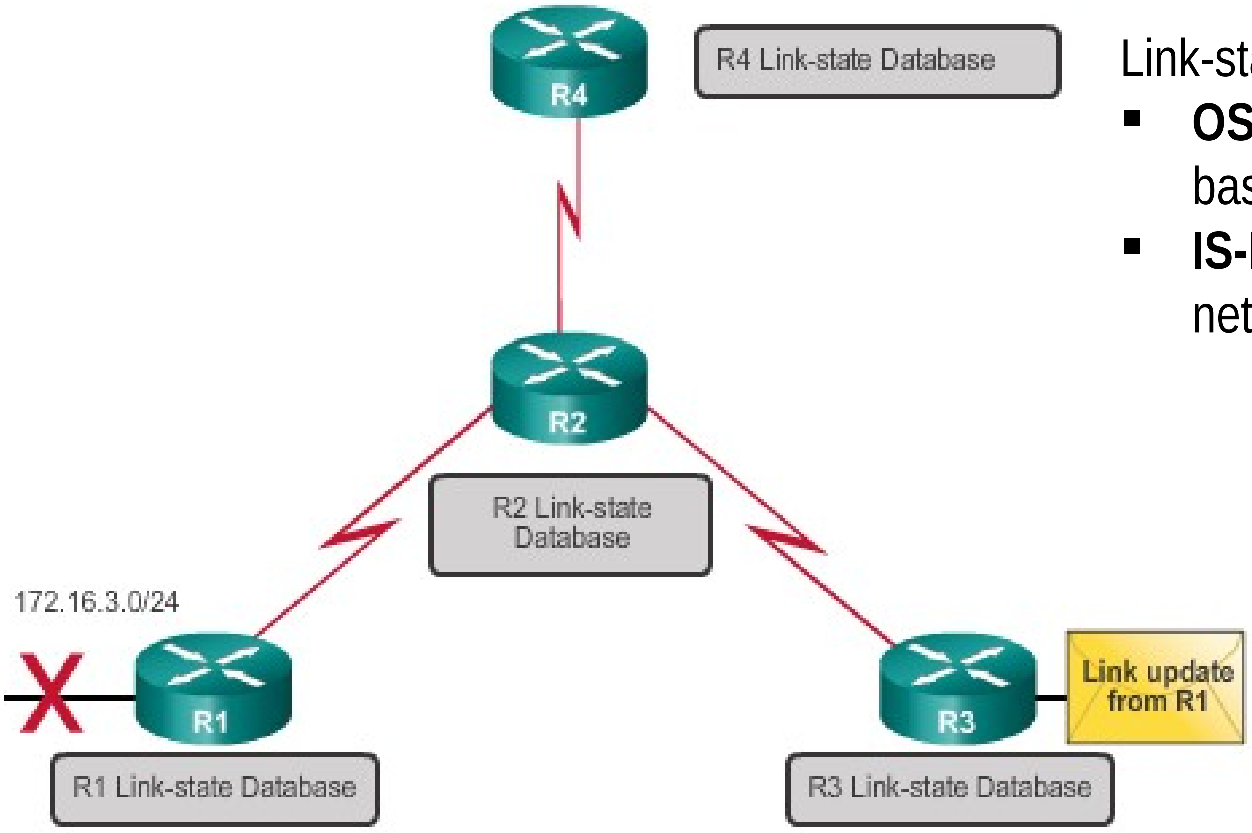

Link-state protocols forward updates when the state of a link changes.

Link-state IPv4 IGPs:

- **OSPF** - Popular standards based routing protocol
- **IS-IS** - Popular in provider networks.

III – Dynamic Routing
Routing Protocol Characteristics

	Distance Vector				Link State	
	RIPv1	RIPv2	IGRP	EIGRP	OSPF	IS-IS
Speed Convergence	Slow	Slow	Slow	Fast	Fast	Fast
Scalability - Size of Network	Small	Small	Small	Large	Large	Large
Use of VLSM	No	Yes	No	Yes	Yes	Yes
Resource Usage	Low	Low	Low	Medium	High	High
Implemenation and Maintenance	Simple	Simple	Simple	Complex	Complex	Complex

III – Dynamic Routing
Routing Protocol Characteristics

- A metric is a measurable value that is assigned by the routing protocol to different routes based on the usefulness of that route:
  - Used to determine the overall “cost” of a path from source to destination.
  - Routing protocols determine the best path based on the route with the lowest cost.

III – Dynamic Routing
Distance Vector Technologies

- Distance vector routing protocols:
  - Share updates between neighbors
  - Not aware of the network topology
  - Send periodic updates even if topology has not changed (except for EIGRP)
  - consume bandwidth and network device CPU resources for updates
  - use multicast addresses (RIPv2 and EIGRP)

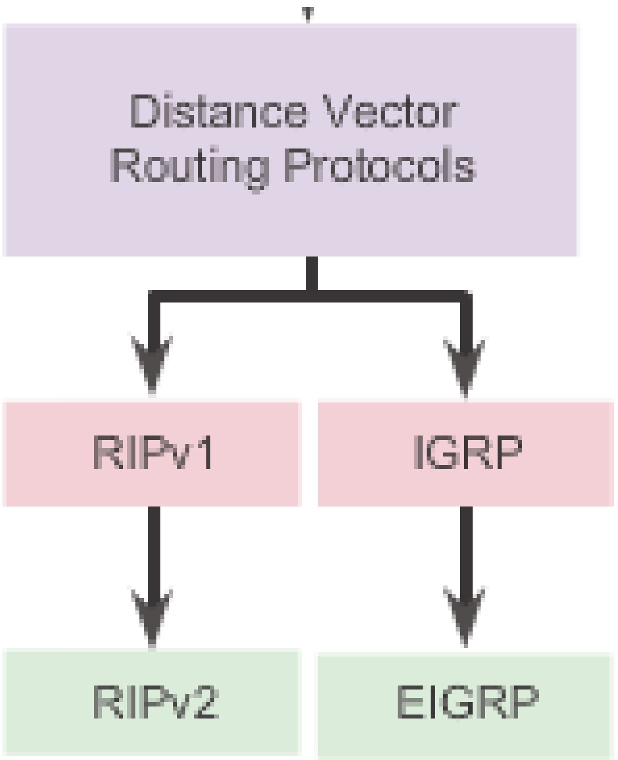

III – Dynamic Routing
Distance Vector Algorithm

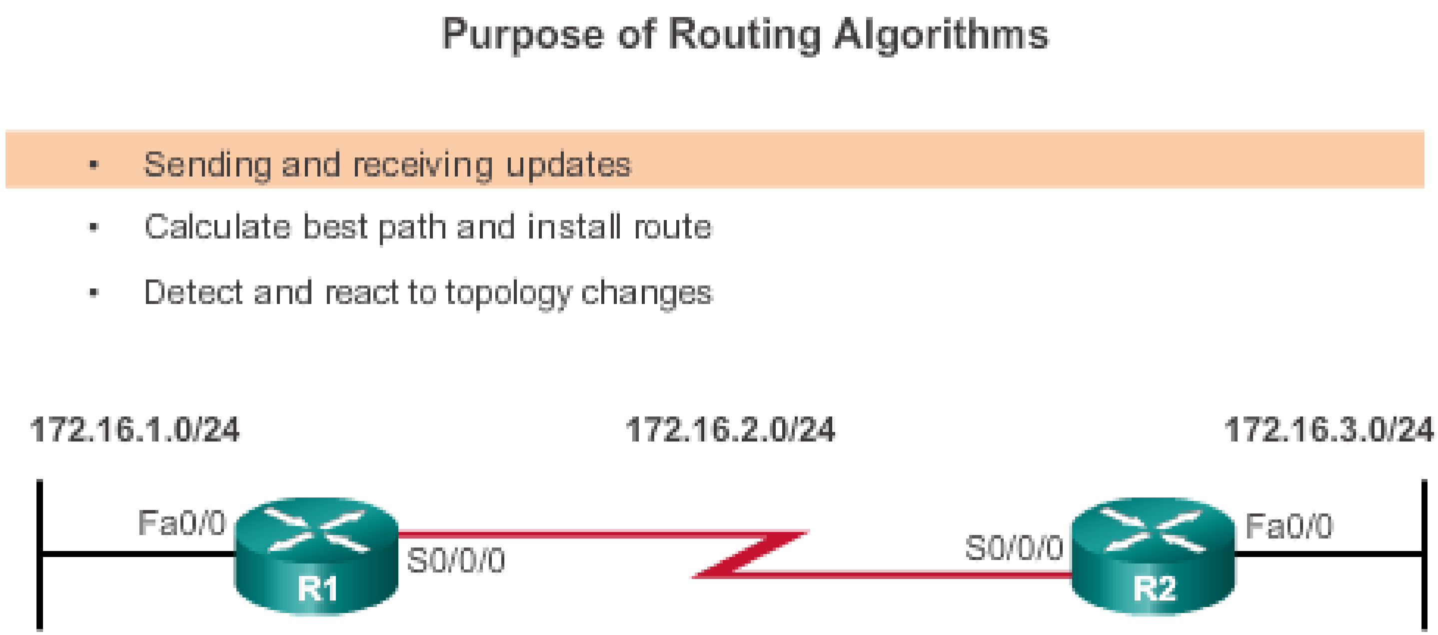

RIP uses the Bellman-Ford algorithm as its routing algorithm.

IGRP and EIGRP use the Diffusing Update Algorithm (DUAL) routing algorithm developed by Cisco.

III – Dynamic Routing
Routing Information Protocol

Routing updates broadcasted every 30 seconds

RIPv1 versus RIPv2

Characteristics and Features	RIPv1	RIPv2
Metric	Both use hop count as a simple metric. The maximum number of hops is 15.	
Updates Forwarded to Address	255.255.255.255	224.0.0.9
Supports VLSM	✗	✓
Supports CIDR	✗	✓
Supports Summarization	✗	✓
Supports Authentication	✗	✓

Updates use UDP port 520

RIPng is based on RIPv2 with a 15 hop limitation and the administrative distance of 120

III – Dynamic Routing
Enhanced Interior-Gateway Routing Protocol

IGRP versus EIGRP

Characteristics and Features	IGRP	EIGRP
Metric	Both use a composite metric consisting of bandwidth and delay. Reliability and load can also be included in the metric calculation.	
Updates Forwarded to Address	255.255.255.255	224.0.0.10
Supports VLSM	✗	✓
Supports CIDR	✗	✓
Supports Summarization	✗	✓
Supports Authentication	✗	✓

EIGRP:

- Is bounded triggered updates
- Uses a Hello keepalives mechanism
- Maintains a topology table
- Supports rapid convergence
- Is a multiple network layer protocol support

III – Dynamic Routing
Shortest Path First Protocols

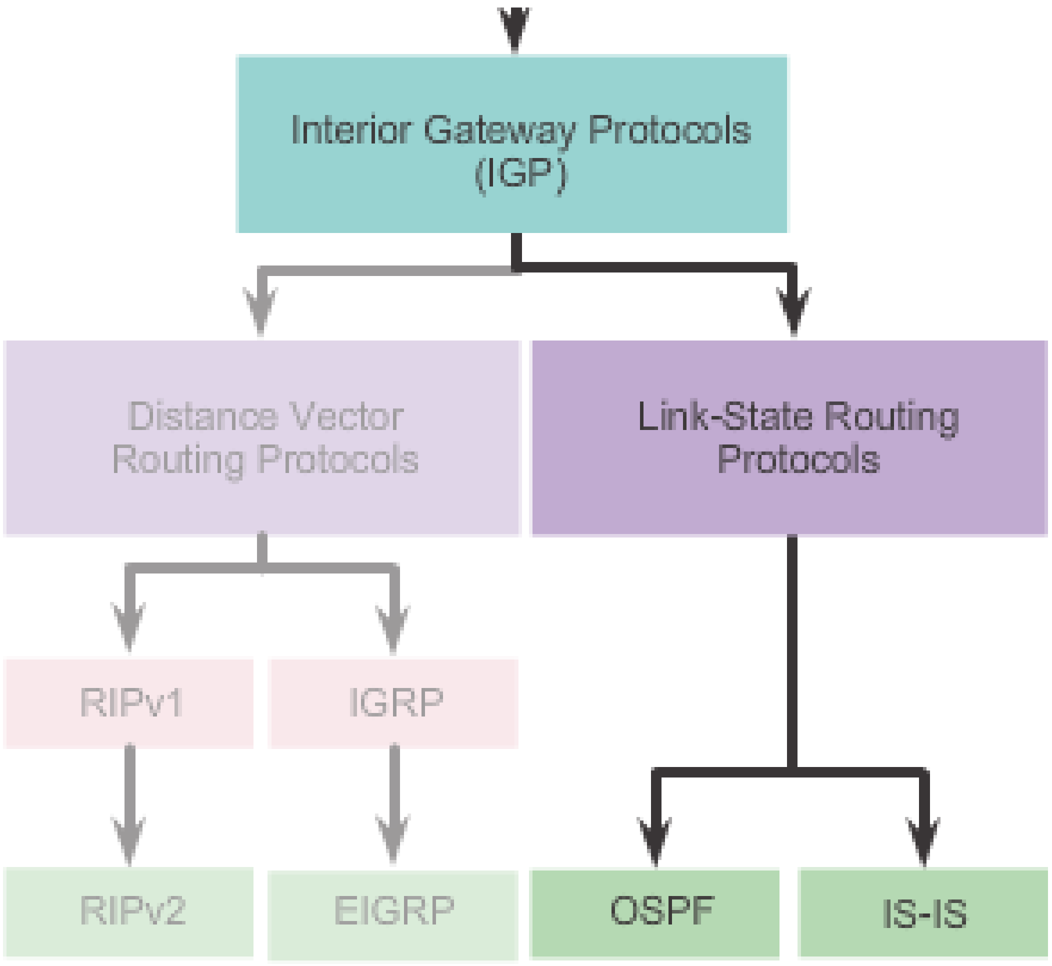

III – Dynamic Routing
Dijkstra's Algorithm

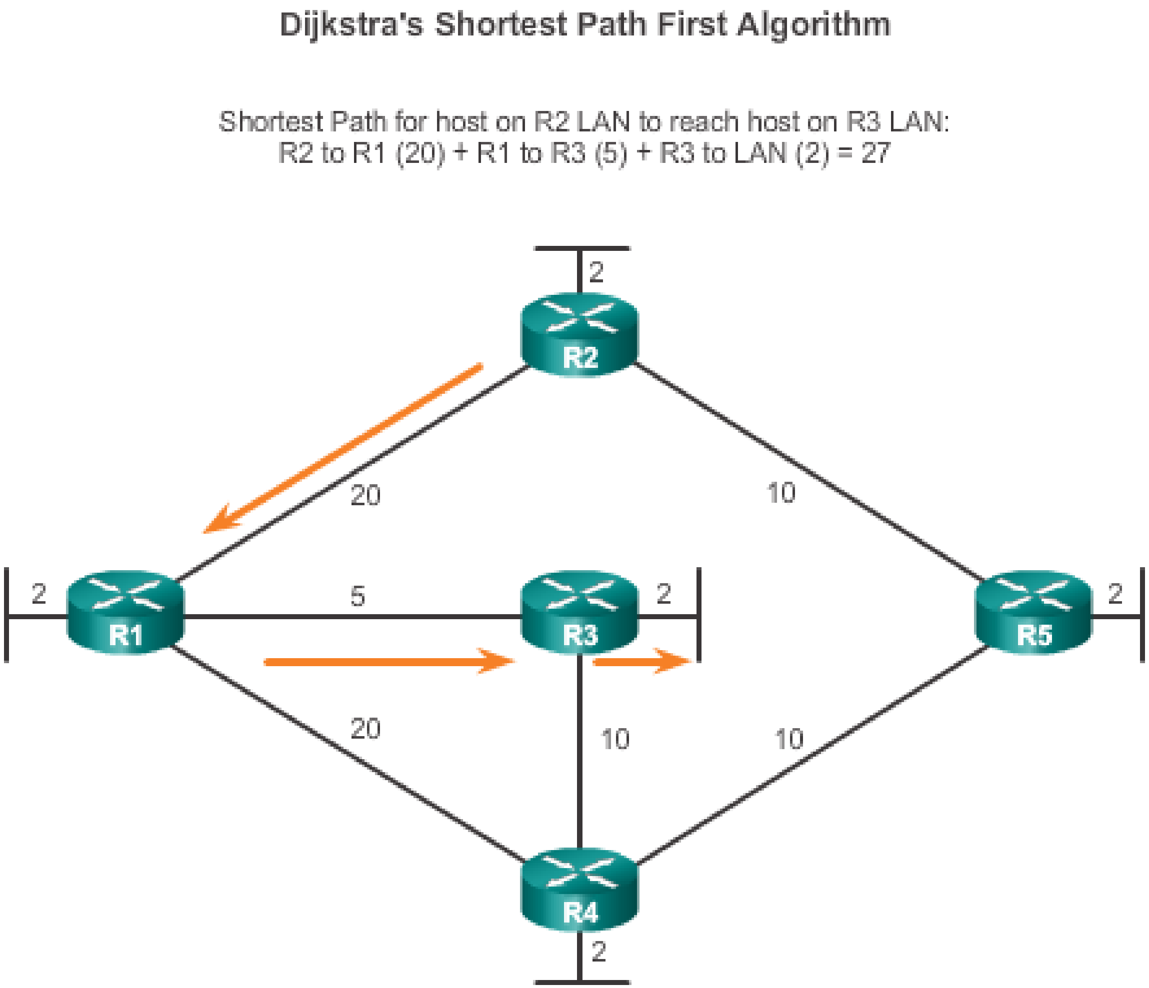

III – Dynamic Routing
Link-State Routing Process

Link-State Routing Process

- Each router learns about each of its own directly connected networks.
- Each router is responsible for "saying hello" to its neighbors on directly connected networks.
- Each router builds a Link State Packet (LSP) containing the state of each directly connected link.
- Each router floods the LSP to all neighbors who then store all LSP's received in a database.
- Each router uses the database to construct a complete map of the topology and computers the best path to each destination networks.

III – Dynamic Routing
Step 1 – Link and Link-State

The first step in the link-state routing process is that each router learns about its own links and its own directly connected networks.

Link-State of Interface Fa0/0

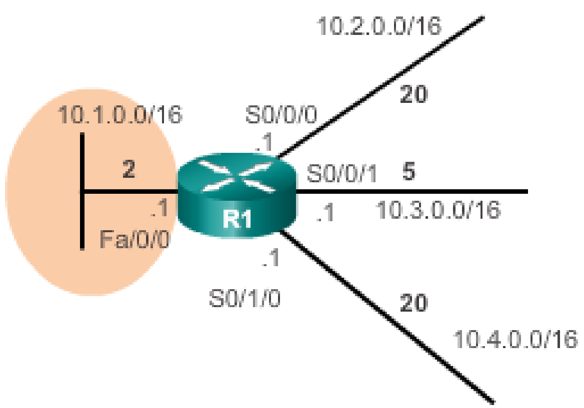

Link-State of Interface S0/0/0

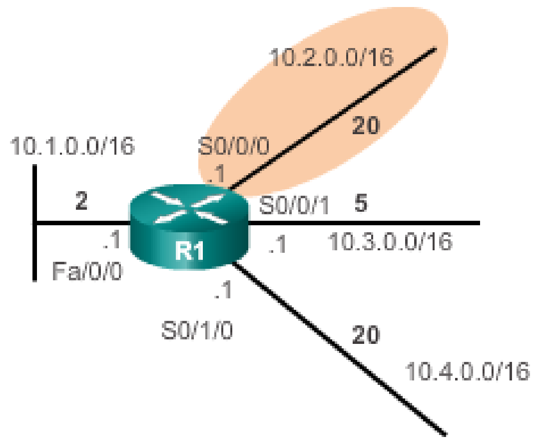

Link 1

- 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**

III – Dynamic Routing
Step 2 – Say Hello

The second step in the link-state routing process is that each router is responsible for meeting its neighbors on directly connected networks.

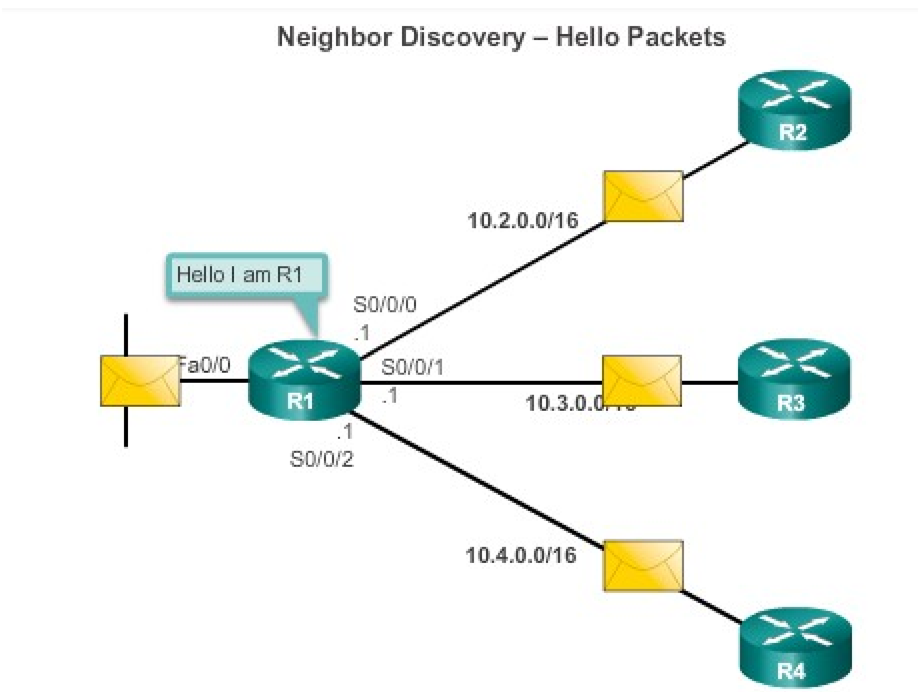

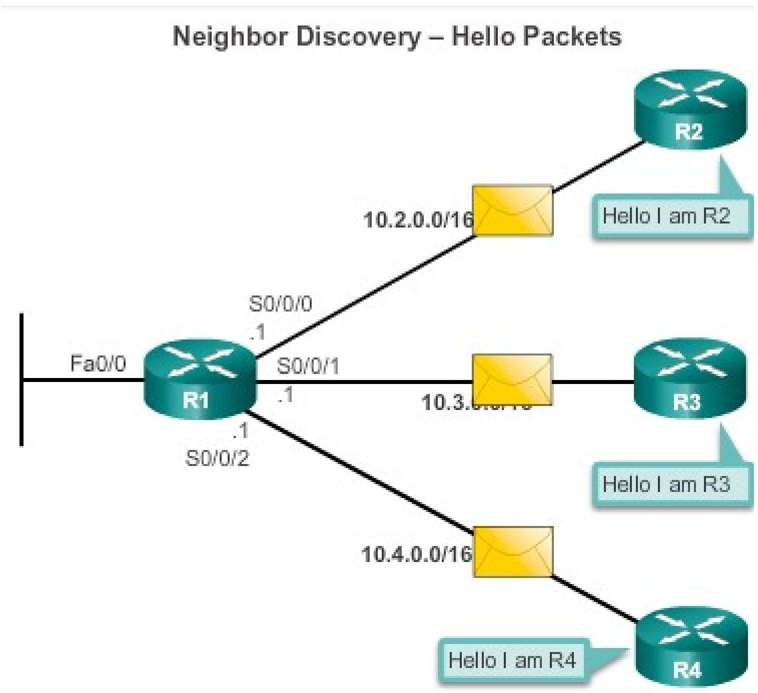

III – Dynamic Routing
Step 3 – Forging Link-State Packets

The third step in the link-state routing process is that each router builds a link-state packet (LSP) containing the state of each directly connected link.

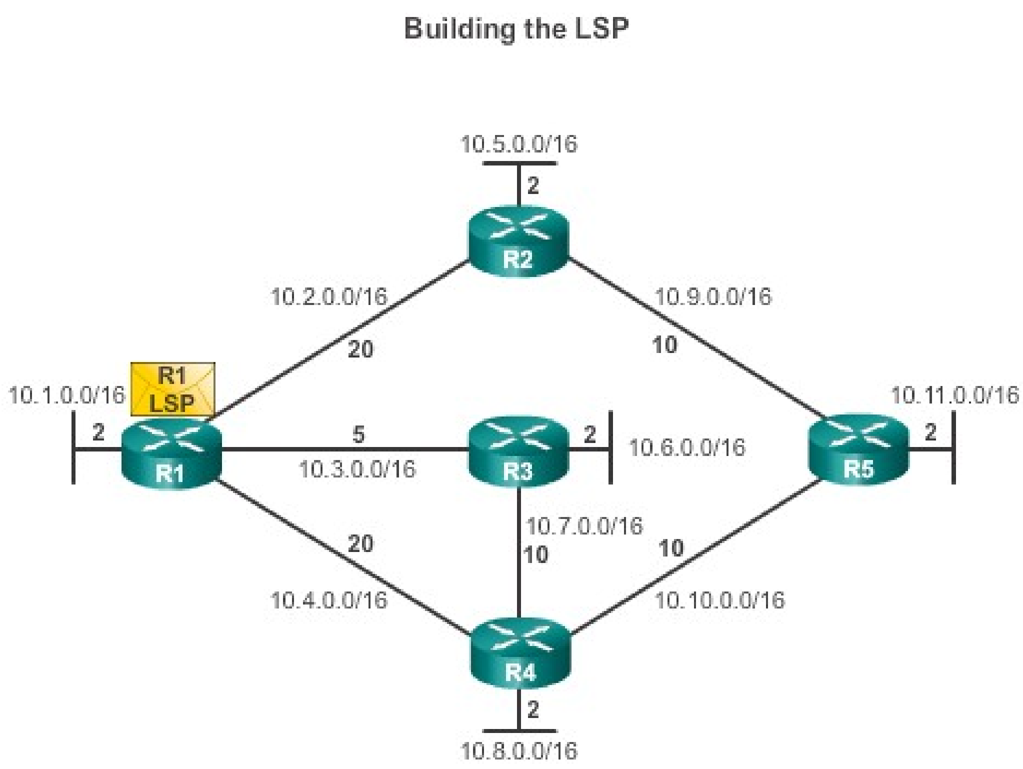

1. R1; Ethernet network 10.1.0.0/16; Cost 2
2. R1 -> R2; Serial point-to-point network; 10.2.0.0/16; Cost 20
3. R1 -> R3; Serial point-to-point network; 10.3.0.0/16; Cost 5
4. R1 -> R4; Serial point-to-point network; 10.4.0.0/16; Cost 20

III – Dynamic Routing
Step 4 – Flooding the LSP

The fourth step in the link-state routing process is that each router floods the LSP to all neighbors, who then store all LSPs received in a database.

Flooding the LSP

R1 Link State Contents

- R1; Ethernet network; 10.1.0.0/16; Cost 2
- R1 -> R2; Serial point-to-point network; 10.2.0.0/16; Cost 20
- R1 -> R3; Serial point-to-point network; 10.3.0.0/16; Cost 5
- R1 -> R4; Serial point-to-point network; 10.4.0.0/16; Cost 20

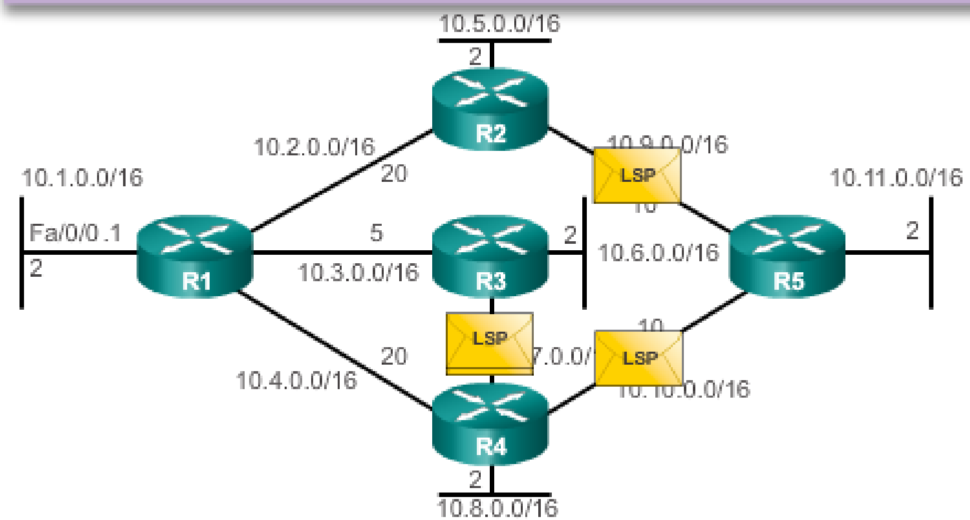

III – Dynamic Routing
Step 5 – Building the Link-State DB

The final step in the link-state routing process is that each router uses the database to construct a complete map of the topology and computes the best path to each destination network.

Contents of the Link-State Database

R1 Link-State Database

R1 Link-states:
- Connected to network 10.1.0.0/16, cost = 2
- Connected to R2 on network 10.2.0.0/16, cost = 20
- Connected to R3 on network 10.3.0.0/16, cost = 5
- Connected to R4 on network 10.4.0.0/16, cost = 20

R2 Link-states:
- Connected to network 10.5.0.0/16, cost = 2
- Connected to R1 on network 10.2.0.0/16, cost = 20
- Connected to R5 on network 10.9.0.0/16, cost = 10

R3 Link-states:
- Connected to network 10.6.0.0/16, cost = 2
- Connected to R1 on network 10.3.0.0/16, cost = 5
- Connected to R4 on network 10.7.0.0/16, cost = 10

R4 Link-states:
- Connected to network 10.8.0.0/16, cost = 2
- Connected to R1 on network 10.4.0.0/16, cost = 20
- Connected to R3 on network 10.7.0.0/16, cost = 10
- Connected to R5 on network 10.10.0.0/16, cost = 10

R5 Link-states:
- Connected to network 10.11.0.0/16, cost = 2
- Connected to R2 on network 10.9.0.0/16, cost = 10
- Connected to R4 on network 10.10.0.0/16, cost = 10

III – Dynamic Routing
Building the SPF Tree

ESIEE AMIENS

Identify the Directly Connected Networks

R1 Link-State Database	SPF Tree

R1 Link-states:
- Connected to network 10.1.0.0/16, cost = 2
- Connected to R2 on network 10.2.0.0/16, cost = 20
- Connected to R3 on network 10.3.0.0/16, cost = 5
- Connected to R4 on network 10.4.0.0/16, cost = 20

R2 Link-states:
- Connected to network 10.5.0.0/16, cost = 2
- Connected to R1 on network 10.2.0.0/16, cost = 20
- Connected to R5 on network 10.9.0.0/16, cost = 10

R3 Link-states:
- Connected to network 10.6.0.0/16, cost = 2
- Connected to R1 on network 10.3.0.0/16, cost = 5
- Connected to R4 on network 10.7.0.0/16, cost = 10

R4 Link-states:
- Connected to network 10.8.0.0/16, cost = 2
- Connected to R1 on network 10.4.0.0/16, cost = 20
- Connected to R3 on network 10.7.0.0/16, cost = 10
- Connected to R5 on network 10.10.0.0/16, cost = 10

R5 Link-states:
- Connected to network 10.11.0.0/16, cost = 2
- Connected to R2 on network 10.9.0.0/16, cost = 10
- Connected to R4 on network 10.10.0.0/16, cost = 10

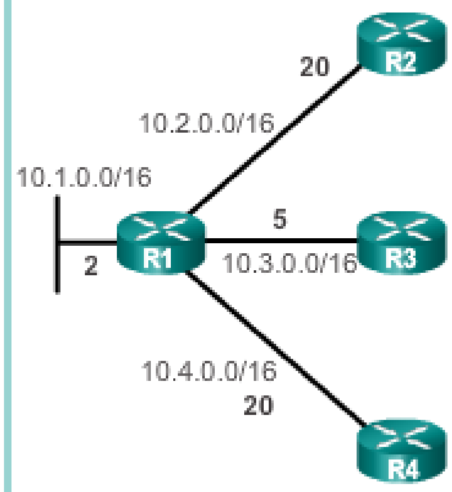

III – Dynamic Routing
Building the SPF Tree

Resulting SPF Tree of R1

Destination	Shortest Path	Cost
10.5.0.0/16	R1 → R2	22
10.6.0.0/16	R1 → R3	7
10.7.0.0/16	R1 → R3	15
10.8.0.0/16	R1 → R3 → R4	17
10.9.0.0/16	R1 → R2	30
10.10.0.0/16	R1 → R3 → R4	25
10.11.0.0/16	R1 → R3 → R4 → R5	27

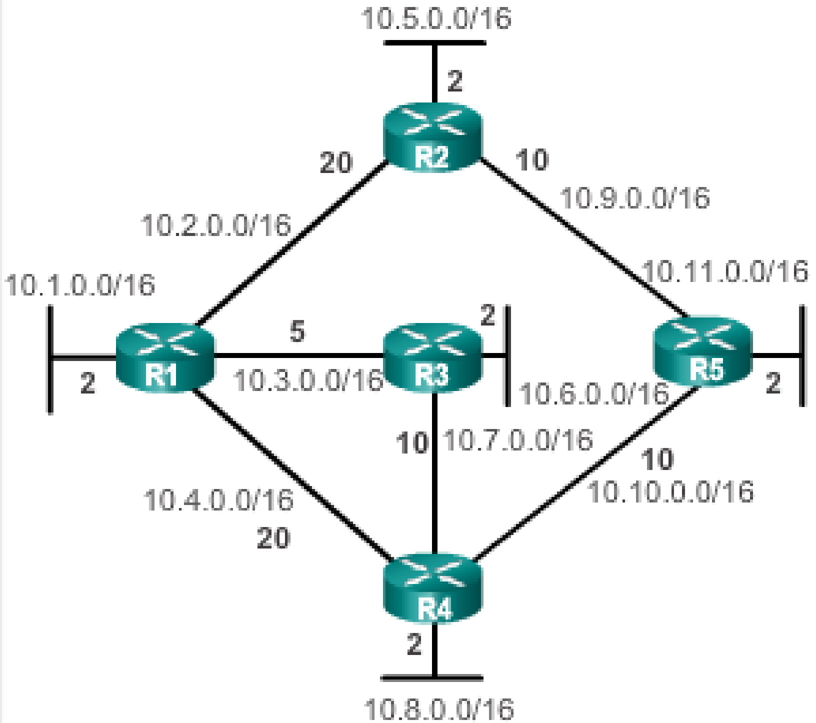

III – Dynamic Routing
Adding SPF Routes to the Routing Table

Populate the Routing Table

Destination	Shortest Path	Cost
10.5.0.0/16	R1 → R2	22
10.6.0.0/16	R1 → R3	7
10.7.0.0/16	R1 → R3	15
10.8.0.0/16	R1 → R3 → R4	17
10.9.0.0/16	R1 → R2	30
10.10.0.0/16	R1 → R3 → R4	25
10.11.0.0/16	R1 → R3 → R4 → R5	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

III – Dynamic Routing

Adding SPF Routes to the Routing Table

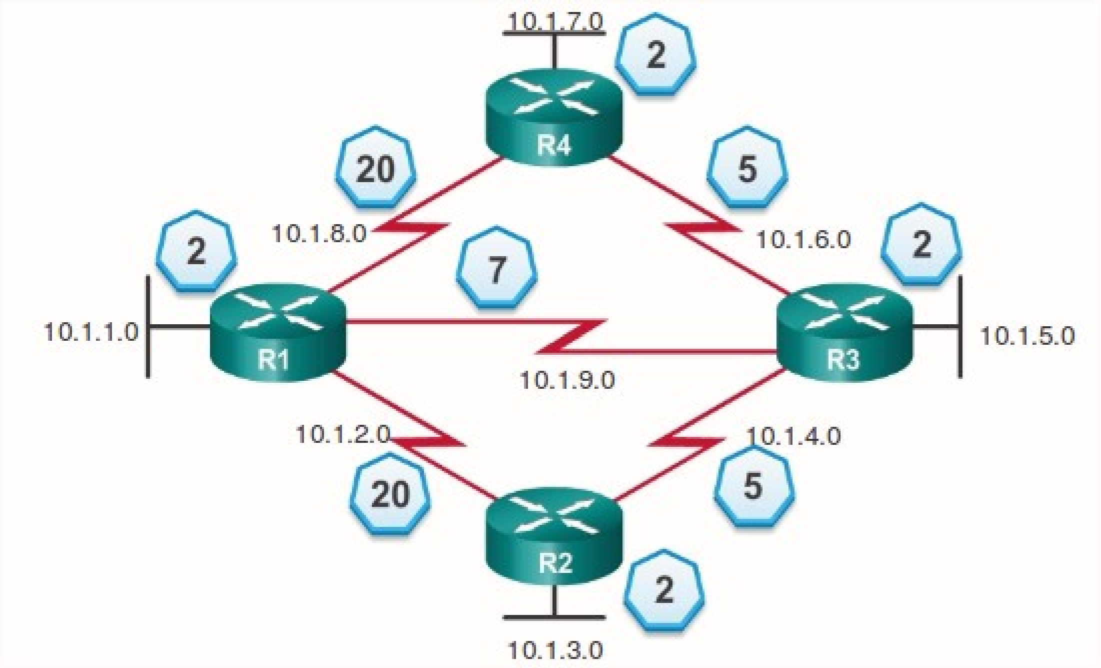

Using R1 SPF tree built from the topology, calculate the best cost for the destination route listed: 10.1.5.0, 10.1.6.0, 10.1.7.0, 10.1.8.0, 10.1.9.0,

And from R2 SPF tree built for: 10.1.1.0, 10.1.4.0, 10.1.5.0, 10.1.6.0, 10.1.7.0,

III – Dynamic Routing

Adding SPF Routes to the Routing Table

Why Use Link-State Protocols ?

Advantages of Link-State Routing Protocols

- Each router builds its own topological map of the network to determine the shortest path.
- Immediate flooding of LSPs achieves faster convergence.
- LSPs are sent only when there is a change in the topology and contain only the information regarding that change.
- Hierarchical design used when implementing multiple areas.

Disadvantages of Link-State Routing Protocols

- Maintaining a link-state database and SPF tree requires additional memory.
- Calculating the SPF algorithm also requires additional CPU processing.
- Bandwidth can be adversely affected by link-state packet flooding.

III – Dynamic Routing
Why Use Link-State Protocols ?

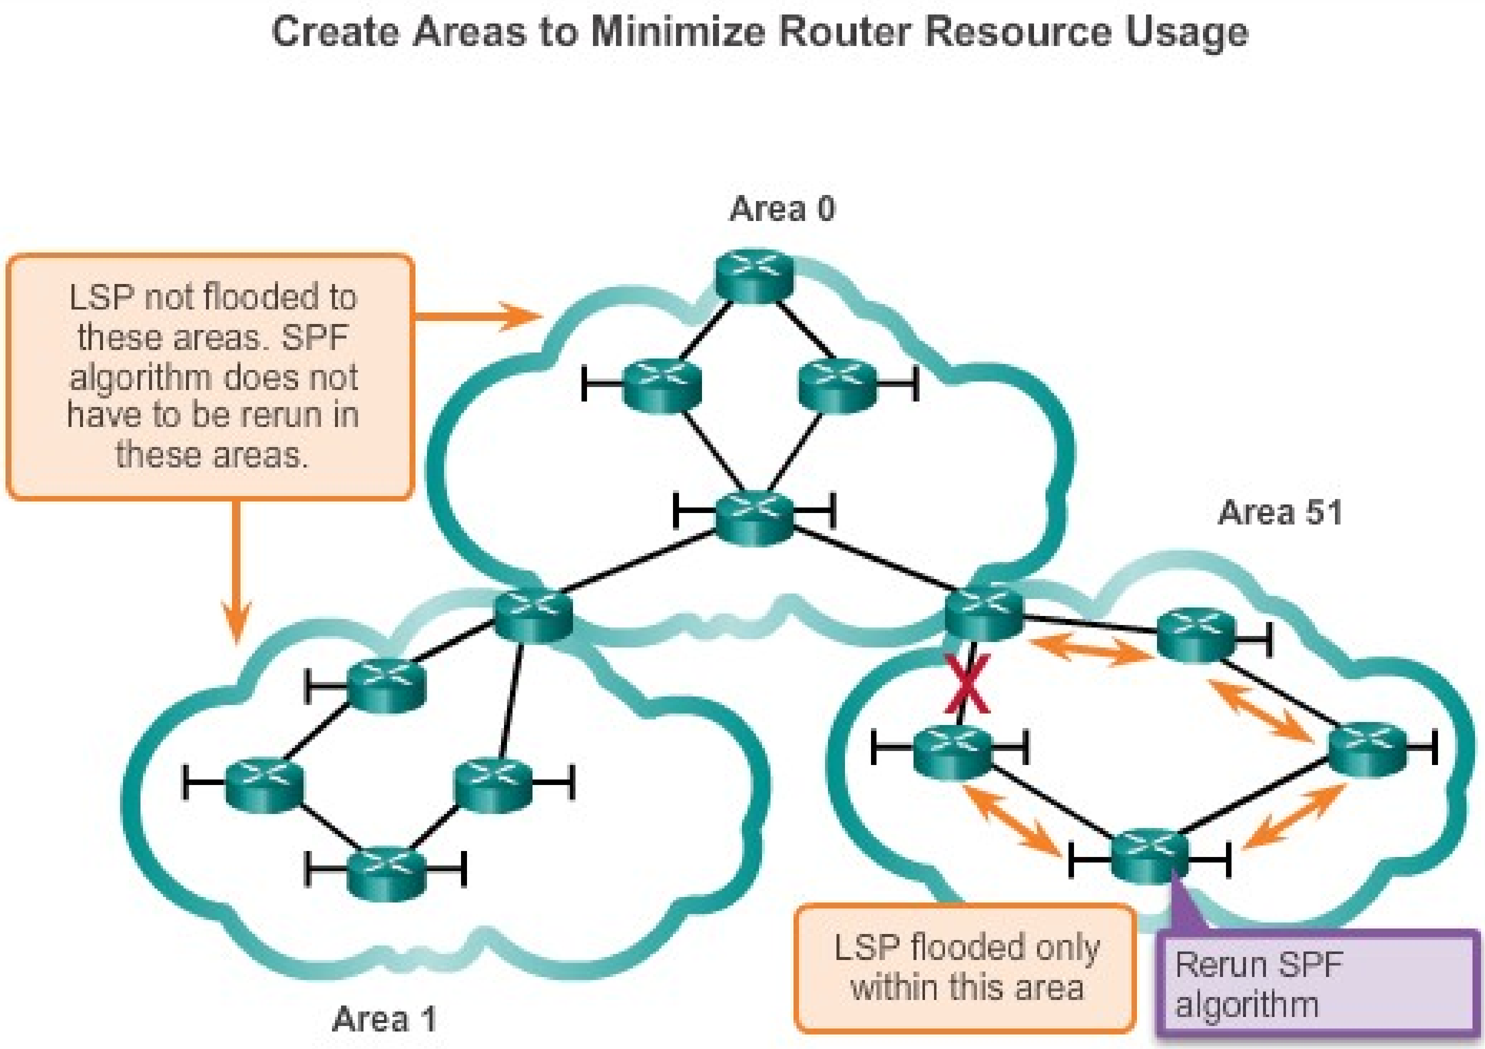

III – Dynamic Routing
Protocols that Use Link-State

There are only two link-state routing protocols:

- Open Shortest Path First (OSPF) most popular
  - began in 1987
  - two current versions
  - OSPFv2 - OSPF for IPv4 networks
  - OSPFv3 - OSPF for IPv6 networks
- IS-IS was designed by International Organization for Standardization (ISO)

III – Dynamic Routing
Routing table

Routing Table of R1

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                     is directly connected, Serial0/0/1
   172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C    172.16.1.0/24 is directly connected, GigabitEthernet0/0
L    172.16.1.1/32 is directly connected, GigabitEthernet0/0
R    172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12, Serial0/0/0
R    172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R    172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R  192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03, Serial0/0/0
   209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C    209.165.200.224/30 is directly connected, Serial0/0/0
L    209.165.200.225/32 is directly connected, Serial0/0/0
R    209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
                    Serial0/0/0
C    209.165.200.232/30 is directly connected, Serial0/0/1
L    209.165.200.233/30 is directly connected, Serial0/0/1
R1#
```

III – Dynamic Routing
Routing table

Directly Connected Interfaces of R1

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                is directly connected, Serial0/0/1
   172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C     172.16.1.0/24 is directly connected, GigabitEthernet0/0
L     172.16.1.1/32 is directly connected, GigabitEthernet0/0
R     172.16.2.0/24 [120/1] via 209.165.200.226,00:00:12, Serial0/0/0
R     172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R     172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03, Serial0/0/0
   209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C     209.165.200.224/30 is directly connected, Serial0/0/0
L     209.165.200.225/32 is directly connected, Serial0/0/0
R     209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12, Serial0/0/0
C     209.165.200.232/30 is directly connected, Serial0/0/1
L     209.165.200.233/32 is directly connected, Serial0/0/1
R1#
```

III – Dynamic Routing
Routing table – Remote Network Entries

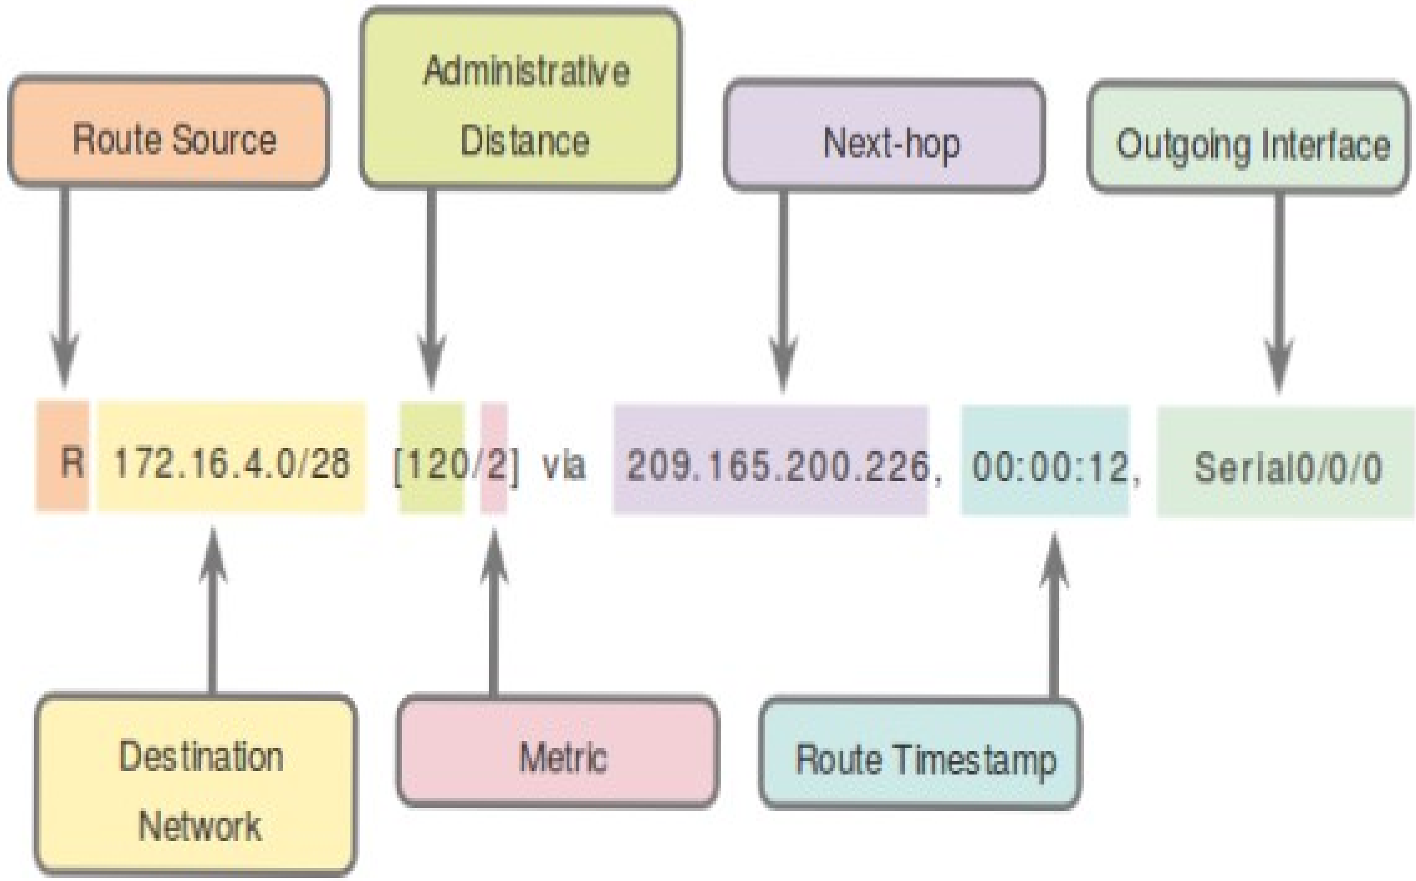

III – Dynamic Routing
Routing table – Terms

Routing Table of R1

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                is directly connected, Serial0/0/1
      172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C        172.16.1.0/24 is directly connected, GigabitEthernet0/0
L        172.16.1.1/32 is directly connected, GigabitEthernet0/0
R        172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12,
         Serial0/0/0
R        172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12,
         Serial0/0/0
R        172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12,
         Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
      Serial0/0/0
      209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C        209.165.200.224/30 is directly connected, Serial0/0/0
L        209.165.200.225/32 is directly connected, Serial0/0/0
R        209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
         Serial0/0/0
C        209.165.200.232/30 is directly connected, Serial0/0/1
L        209.165.200.233/32 is directly connected, Serial0/0/1
R1#
```

Routes are discussed in terms of:

- Ultimate route
- Level 1 route
- Level 1 parent route
- Level 2 child routes

III – Dynamic Routing
Routing table – Ultimate Route

Ultimate Routes of R1

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                is directly connected, Serial0/0/1
      172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C        172.16.1.0/24 is directly connected, GigabitEthernet0/0
L        172.16.1.1/32 is directly connected, GigabitEthernet0/0
R        172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12,
         Serial0/0/0
R        172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12,
          Serial0/0/0
R        172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12,
         Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
      Serial0/0/0
      209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C        209.165.200.224/30 is directly connected, Serial0/0/0
L        209.165.200.225/32 is directly connected, Serial0/0/0
R        209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
         Serial0/0/0
C        209.165.200.232/30 is directly connected, Serial0/0/1
L        209.165.200.233/32 is directly connected, Serial0/0/1
R1#
```

An ultimate route is a routing table entry that contains either a next-hop IP address or an exit interface. Directly connected, dynamically learned, and link local routes are ultimate routes.

III – Dynamic Routing
Routing table – Level 1 Route

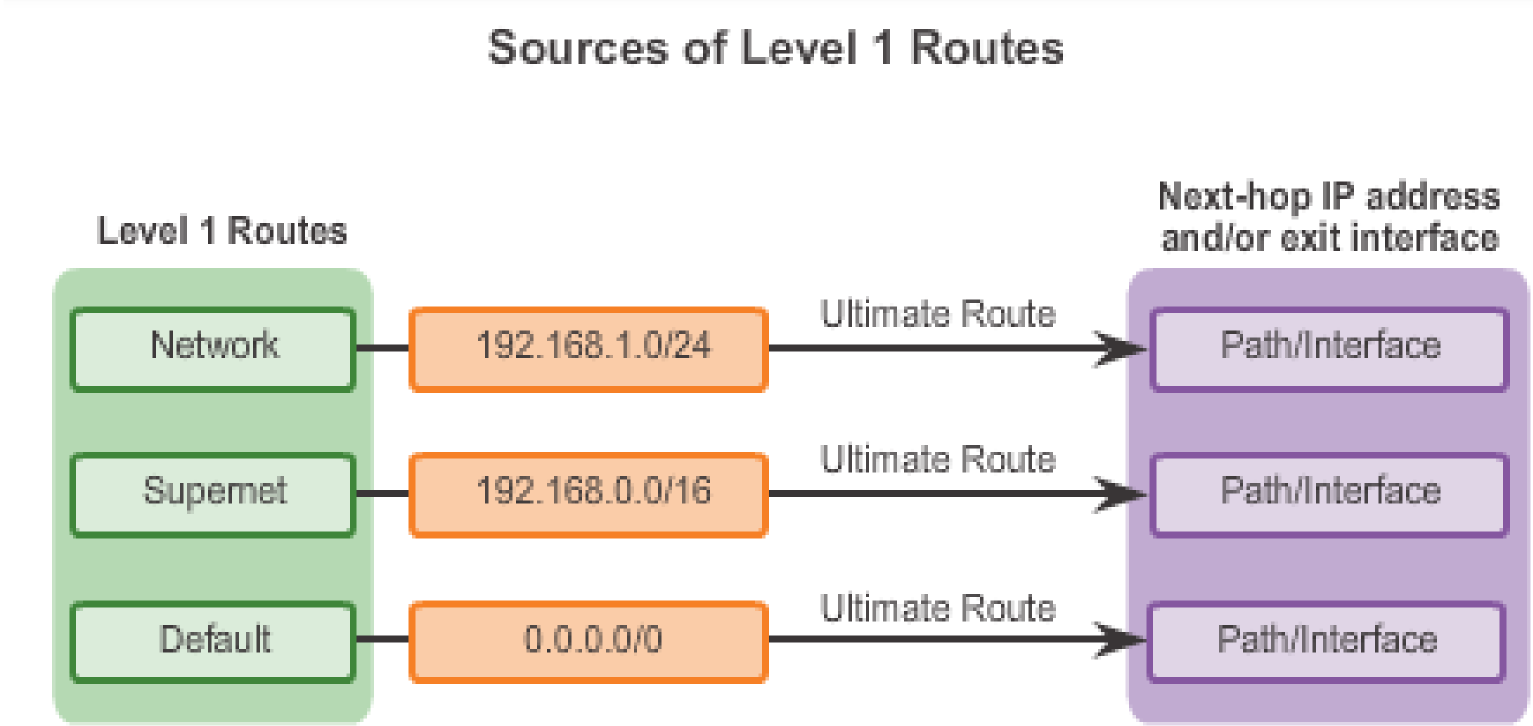

III – Dynamic Routing
Routing table – Level 1 Parent Route

Level 1 Parent Routes of R1

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network
0.0.0.0

S*    0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                is directly connected, Serial0/0/1
      172.16.0.0/16 is variably subnetted, 5 subnets, 3
masks
C        172.16.1.0/24 is directly connected,
GigabitEthernet0/0
L        172.16.1.1/32 is directly connected,
GigabitEthernet0/0
R        172.16.2.0/24 [120/1] via 209.165.200.226,
00:00:12, Serial0/0/0
R        172.16.3.0/24 [120/2] via 209.165.200.226,
00:00:12, Serial0/0/0
R        172.16.4.0/28 [120/2] via 209.165.200.226,
00:00:12, Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
Serial0/0/0
      209.165.200.0/24 is variably subnetted, 5 subnets, 2
masks
C        209.165.200.224/30 is directly connected,
Serial0/0/0
```

III – Dynamic Routing
Routing table – Level 2 Child Route

ESIEE AMIENS

Example of Level 2 Child Routes

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network
0.0.0.0

S*     0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
                is directly connected, Serial0/0/1
      172.16.0.0/16 is variably subnetted, 5 subnets, 3
masks
C        172.16.1.0/24 is directly connected,
GigabitEthernet0/0
L        172.16.1.1/32 is directly connected,
GigabitEthernet0/0
R        172.16.2.0/24 [120/1] via 209.165.200.226,
00:00:12, Serial0/0/0
R        172.16.3.0/24 [120/2] via 209.165.200.226,
00:00:12, Serial0/0/0
R        172.16.4.0/28 [120/2] via 209.165.200.226,
00:00:12, Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
Serial0/0/0
      209.165.200.0/24 is variably subnetted, 5 subnets, 2
masks
C        209.165.200.224/30 is directly connected,
Serial0/0/0
```

III – Dynamic Routing
Routing table – Route Lookup Process

1. If the best match is a level 1 ultimate route, then this route is used to forward the packet.

2. If the best match is a level 1 parent route, proceed to the next step.
   a) The router examines child routes (the subnet routes) of the parent route for a best match.
   b) If there is a match with a level 2 child route, that subnet is used to forward the packet.
   c) If there is not a match with any of the level 2 child routes, proceed to the next step.

3. The router continues searching level 1 supernet routes in the routing table for a match, including the default route, if there is one.

4. If there is now a lesser match with a level 1 supernet or default routes, the router uses that route to forward the packet.

5. If there is not a match with any route in the routing table, the router drops the packet.

III – Dynamic Routing
Routing table – Best Route

Matches for Packet Destined to 172.16.0.10

IP Packet Destination	172.16.0.10	10101100.00010000.00000000.00001010

Route 1	172.16.0.0/12	10101100.00010000.00000000.00000000
Route 2	172.16.0.0/18	10101100.00010000.00000000.00000000
Route 3	172.16.0.0/26	10101100.00010000.00000000.00000000

Longest Match to IP Packet Destination

IV - EIGRP

IV – EIGRP
Feature of EIGRP

- Released in 1992 as a Cisco proprietary protocol.
- 2013 basic functionality of EIGRP released as an open standard.
- Advanced Distance Vector routing protocol.
- Uses the Diffusing Update Algorithm (DUAL) to calculate paths and back-up paths.
- Establishes Neighbor Adjacencies.
- Uses the Reliable Transport Protocol to provide delivery of EIGRP packets to neighbors.
- Partial and Bounded Updates. Send updates only when there is a change and only to the routers that need the information.
- Supports Equal and Unequal Cost Load Balancing.

IV – EIGRP
Reliable Transport Protocol (RTP)

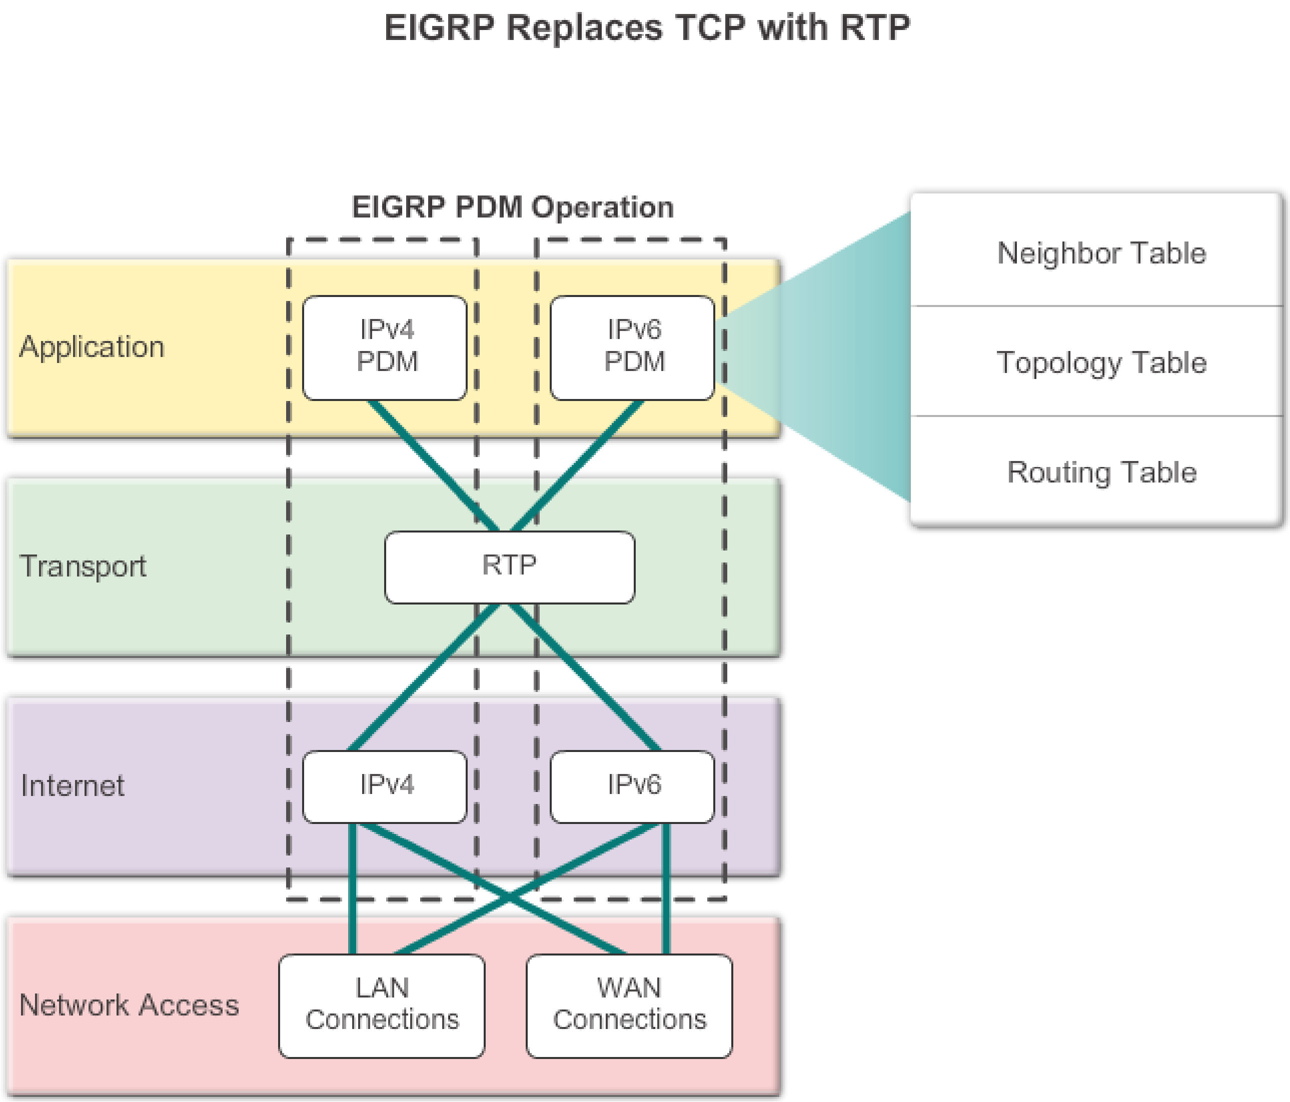

IV – EIGRP Authentication

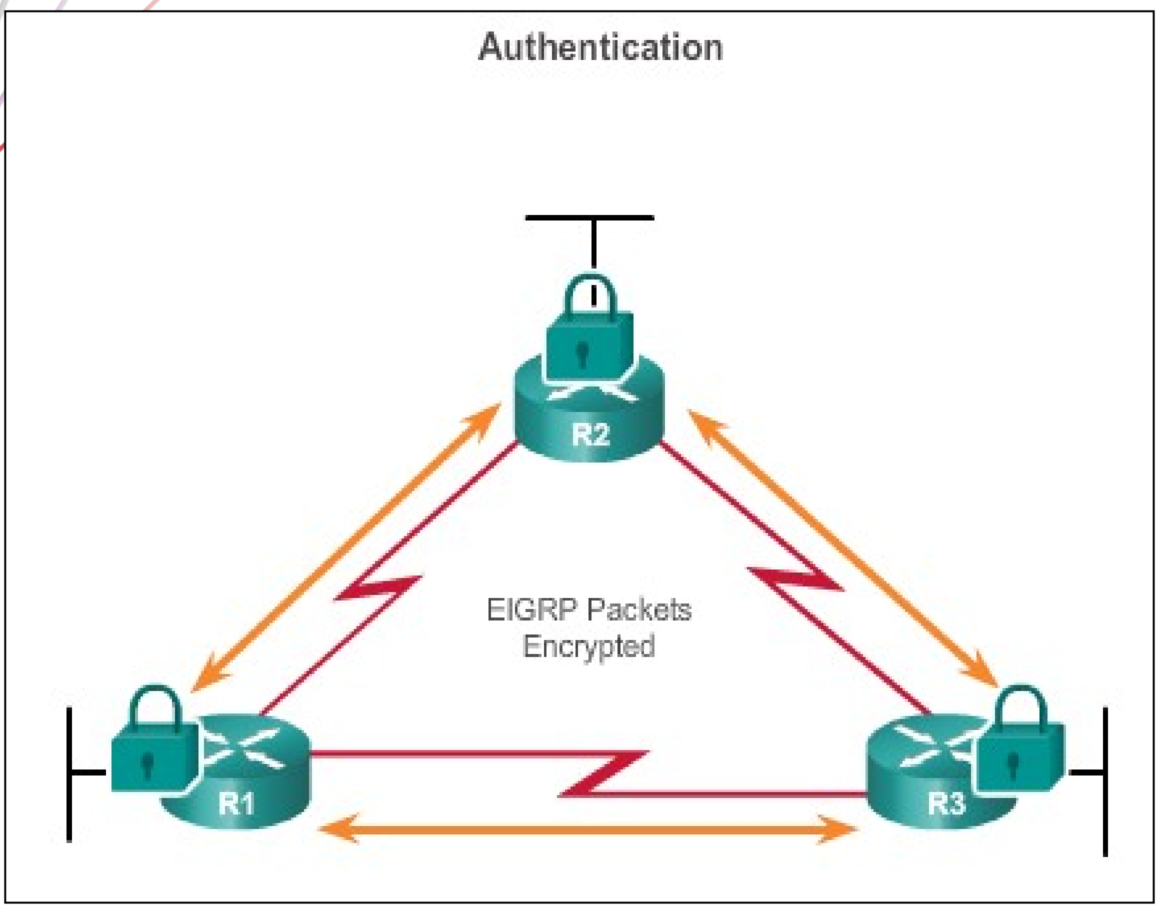

- EIGRP can be configured to authenticate routing information.
- Ensures routers only accept updates from routers that have been configured with the correct authentication information.

IV – EIGRP
Packet Types

Packet Type	Description
Hello	Used to discover other EIGRP routers in the network.
Acknowledgement	Used to acknowledge the receipt of any EIGRP packet.
Update	Convey routing information to known destinations.
Query	Used to request specific information from a neighbor router.
Reply	Used to respond to a query.

IV – EIGRP
Packet Types – Hello Packets

- EIGRP Hello Packets
  - Used to discover EIGRP neighbors.
  - Used to form and maintain EIGRP neighbor adjacencies.
  - Sent as IPv4 or IPv6 multicasts.
  - IPv4 multicast address 224.0.0.10.
  - IPv6 multicast address FF02::A.
  - Unreliable delivery.
  - Sent every 5 seconds (every 60 seconds on low-speed NBMA networks).
  - EIGRP uses a default Hold timer of three times the Hello interval before declaring neighbor unreachable.

IV – EIGRP
Packet Types – Query & Reply

- Used when searching for networks.
- Queries use reliable delivery, which can be multicast or unicast.
- Replies use reliable delivery.

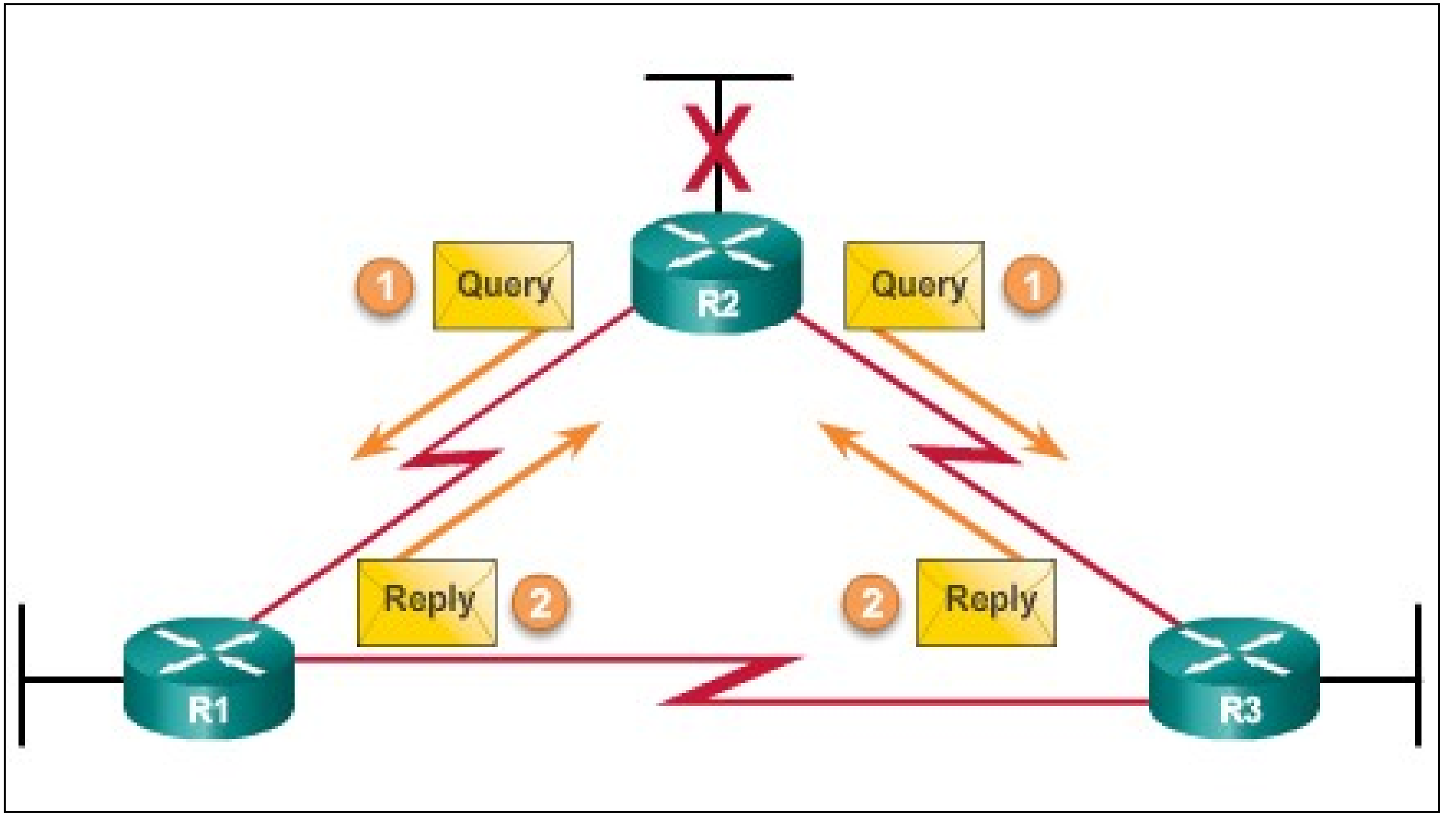

IV – EIGRP Encapsulation

Type/Length/Values Types

Data Link Frame Header	IP Packet Header	EIGRP Packet Header	TLV Types

Data Link Frame
MAC Source Address = Address of sending interface
MAC Destination Address = Multicast: 01-00-5E-00-00-0A

IP Packet
IPv4 Source Address = Address of sending interface
IPv4 Destination Address = Multicast: 224.0.0.10
Protocol field = 88 for EIGRP

EIGRP Packet Header
Opcode for EIGRP packet type
Autonomous System Number

TLV Types
Some types include:
0x0001 EIGRP Parameters
0x0102 IP Internal Routes
0x0103 IP External Routes

IV – EIGRP
Packet Header & TLV

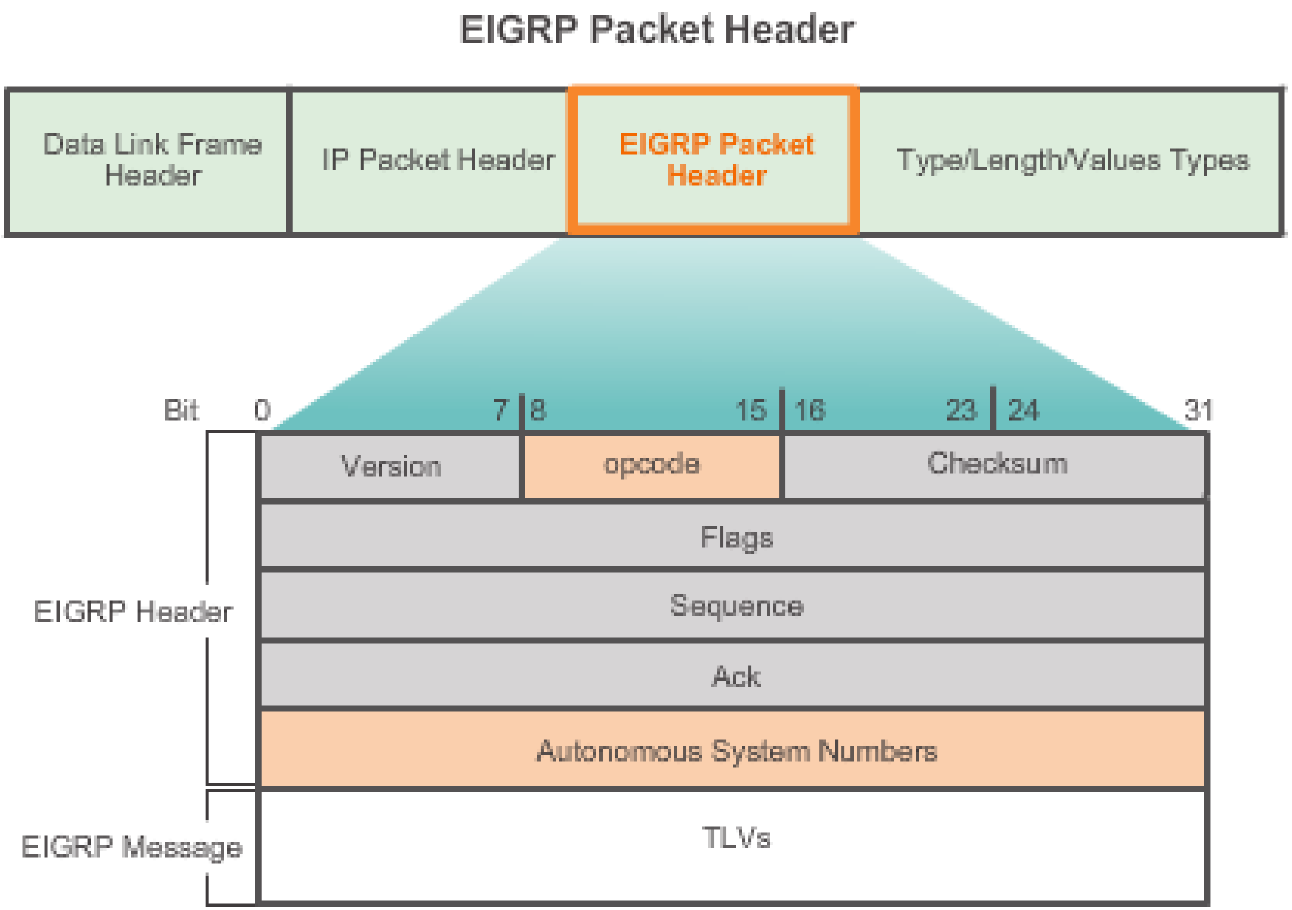

- **Opcode:** EIGRP Packet Type: Update (1), Query (3), Reply(4), Hello (5)
- **Autonomous System Number:** ID for this EIGRP routing process

IV – EIGRP
EIGRP Neighbor Table

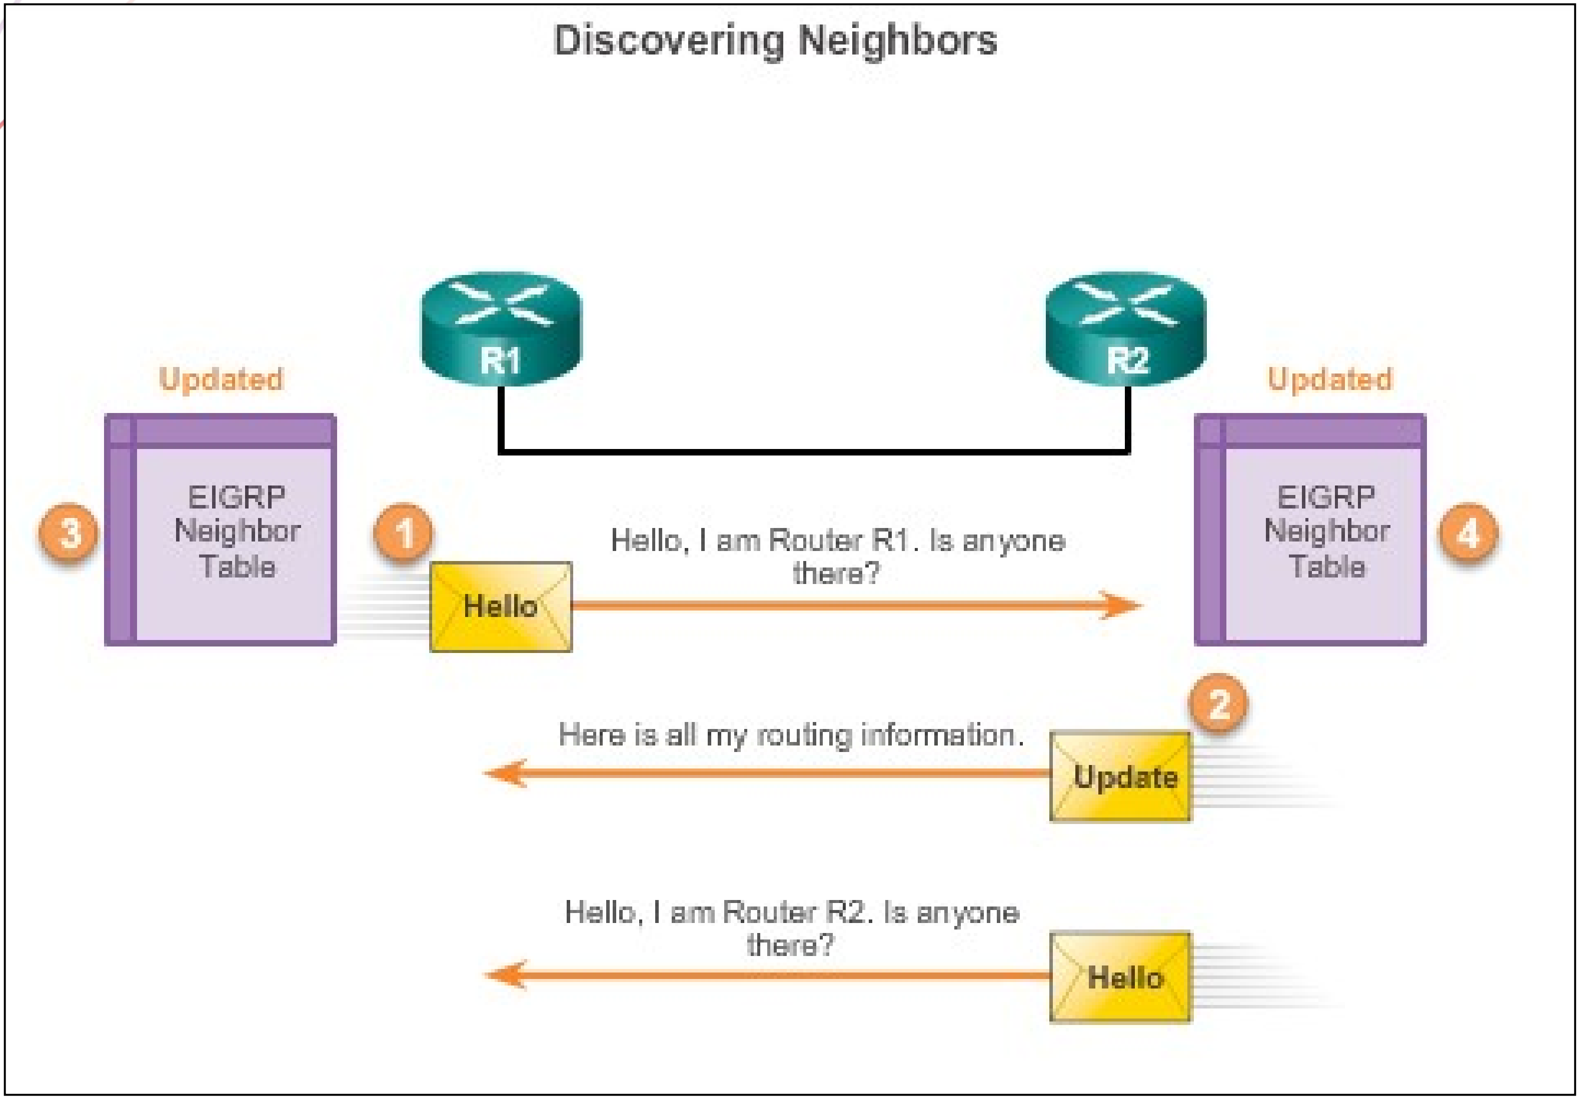

IV – EIGRP
EIGRP Topology Table

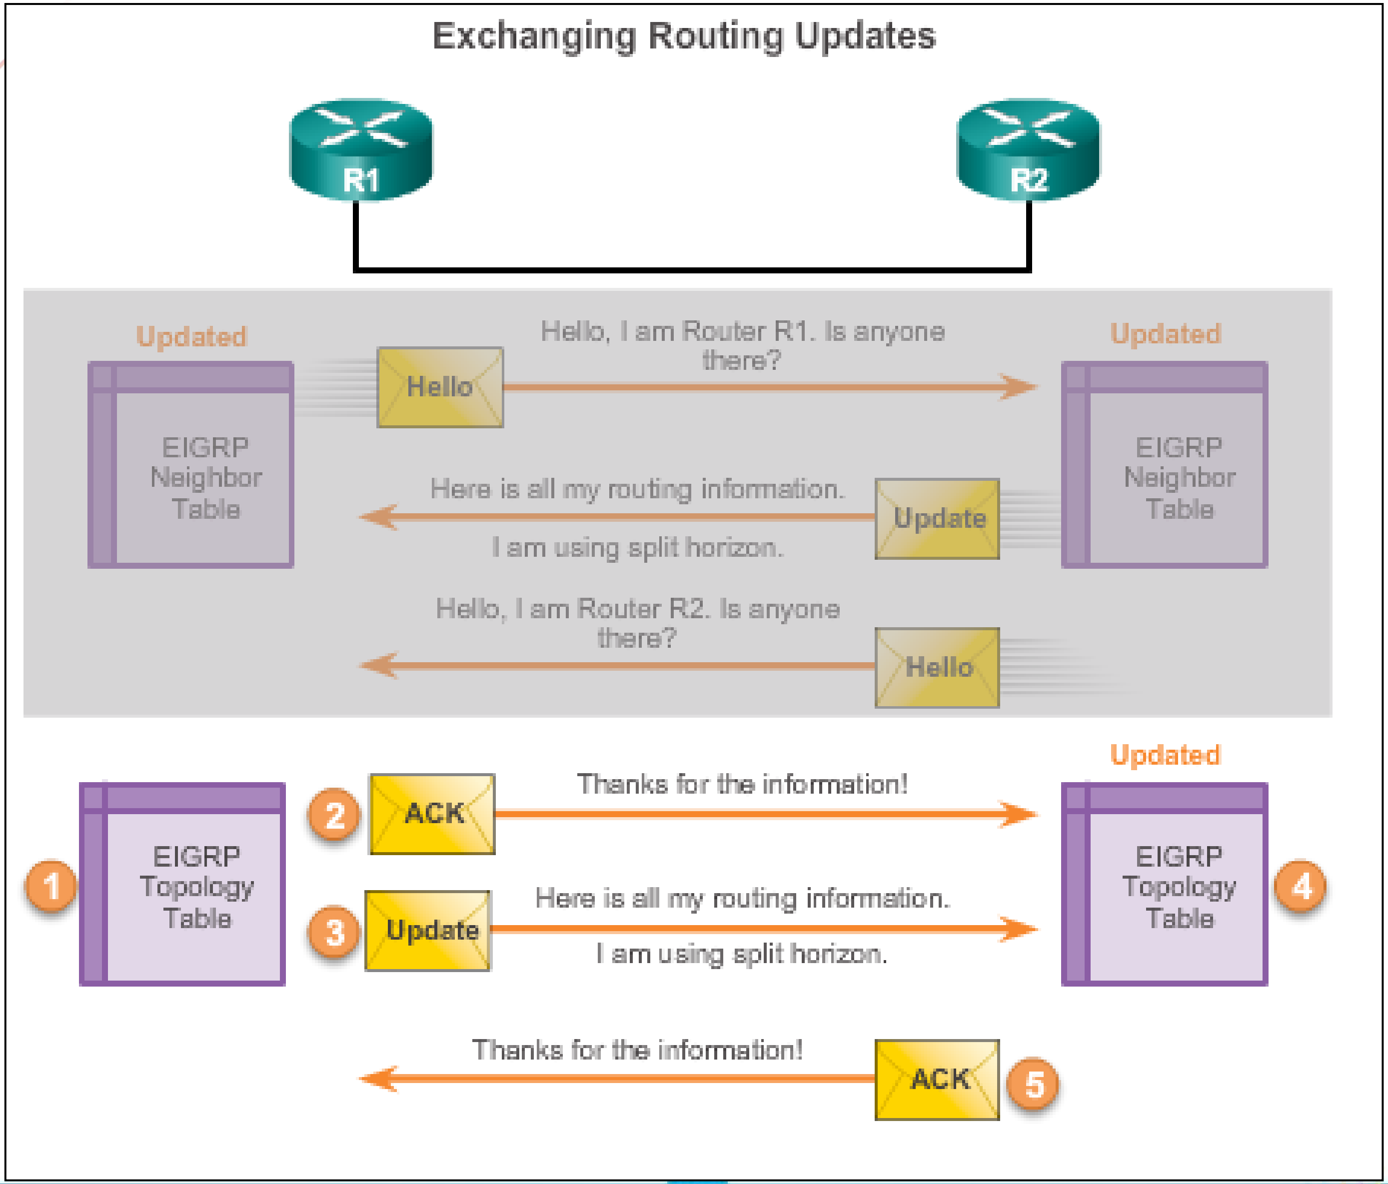

EIGRP Convergence – Routing Table update

Convergence – All routers have the correct, most up-to-date information about the network.

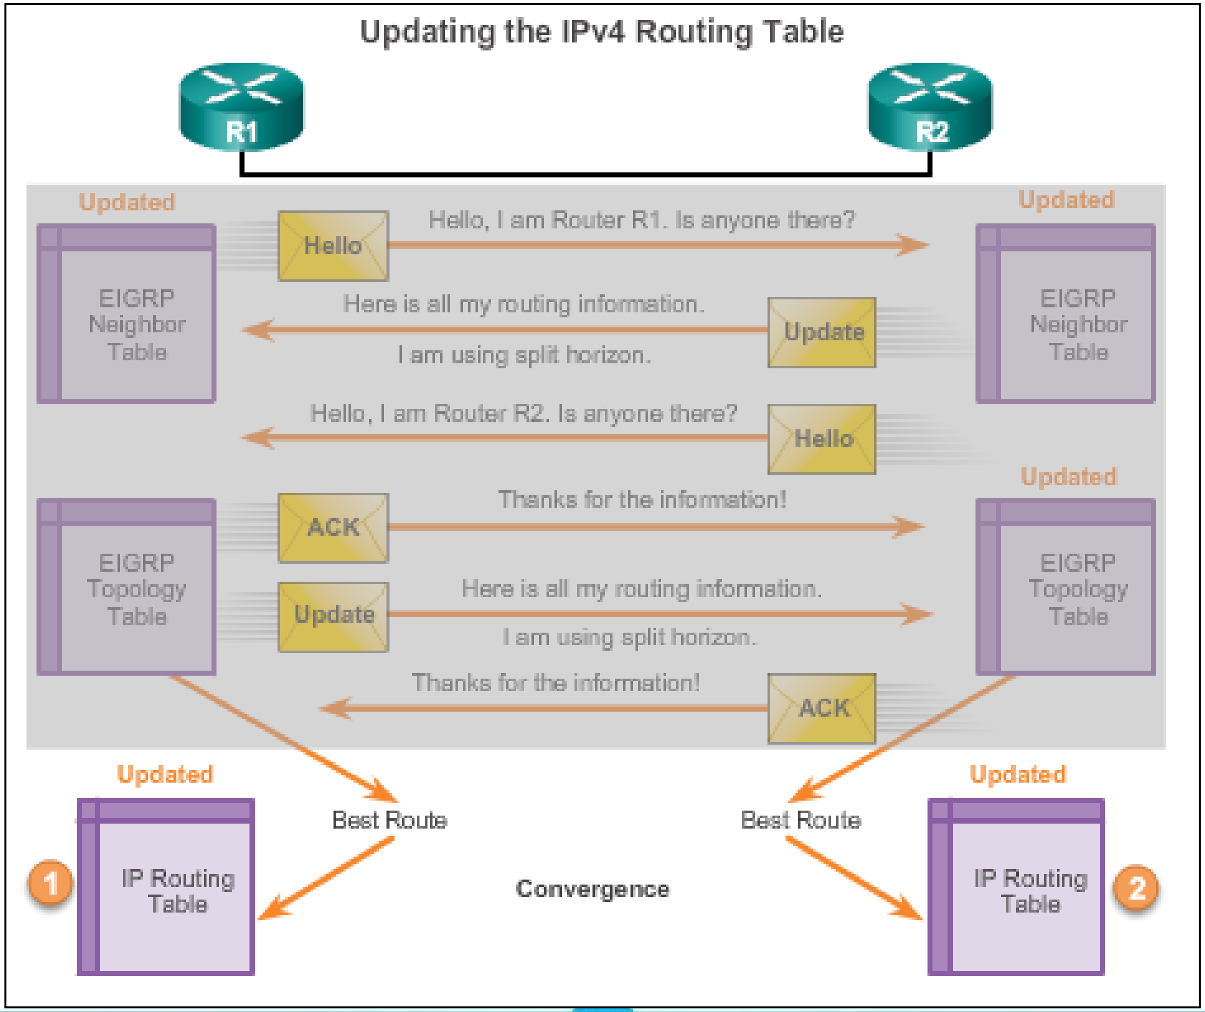

IV – EIGRP
Metric – EIGRP Composite Metric

ESIEE AMIENS

EIGRP Composite Metric

Default Composite Formula:
metric = **[K1*bandwidth + K3*delay]**

Complete Composite Formula:
metric = **[K1*bandwidth + (K2*bandwidth)/(256 - load)+K3*delay] * [K5/(reliability + K4)]**

(Not used if "K" values are 0)

Note: This is a conditional formula. If K5 = 0, the last term is replaced by 1 and the formula becomes: Metric = [K1 * bandwidth + (K2 * bandwidth)/(256–load) + K3 * delay]

Default values:
- K1 (bandwidth) = 1
- K2 (load) = 0
- K3 (delay) = 1
- K4 (reliability) = 0
- K5 (reliability) = 0

"K" values can be changed with the `metric weights` command

```
Router(config-router)# metric weights tos k1 k2 k3 k4 k5
```

IV – EIGRP
Metric – Bandwidth

- Use the **`show interfaces`** command to verify bandwidth.
- Most serial bandwidths are set to 1,544 kb/s (default).
- A correct value for bandwidth is very important in order to calculate the correct metric (both sides of link must have same bandwidth).

```
R1(config)# interface s 0/0/0
R1(config-if)# bandwidth 64
```

```
R1# show interface s 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Internet address is 172.16.3.1/30
  MTU 1500 bytes, BW 64 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<Output omitted>
```

Metric – Delay

Interface Delay Values

Media	Delay
Ethernet	1,000
Fast Ethernet	100
Gigabit Ethernet	10
16M Token Ring	630
FDDI	100
T1 (Serial Default)	20,000
DS0 (64 Kbps)	20,000
1024 Kbps	20,000
56 Kbps	20,000

IV – EIGRP
Metric – Calculating

Step 1. Determine the link with the slowest bandwidth. Use that value to calculate bandwidth (10,000,000/bandwidth).

Step 2. Determine the delay value for each outgoing interface on the way to the destination. Add the delay values and divide by 10 (sum of delay/10).

Step 3. Add the computed values for bandwidth and delay, and multiply the sum by 256 to obtain the EIGRP metric.

[K1 * bandwidth + K3 * delay] * 256 = Metric

Since K1 and K3 both equal 1, the formula simplifies to:

(Bandwidth + Delay) * 256 = Metric

((10,000,000 / bandwidth) + (sum of delay / 10)) * 256 = Metric

```
R2# show ip route

D 192.168.1.0/24 [90/3012096] via 192.168.10.10, 00:12:32, Serial0/0/1
```

IV – EIGRP
DUAL Concepts

- **Diffusing Update ALgorithm (DUAL)** provides the following:
  - Loop-free paths and loop-free backup paths
  - Fast convergence
  - Minimum bandwidth usage with bounded updates
- The decision process for all route computations is done by the **DUAL Finite State Machine (FSM)**
  - DUAL FSM tracks all routes.
  - Uses EIGRP metrics to select efficient, loop-free paths.
  - Identifies the routes with the least-cost path to be inserted into the routing table.
- EIGRP maintains a list of backup routes that DUAL has already determined that can be used immediately if the primary path fails.

IV – EIGRP
DUAL – Successor & Feasible Distance

- The **Successor** is the least-cost route to the destination network.
- The **Feasible Distance** (FD) is the lowest calculated metric to reach the destination network.

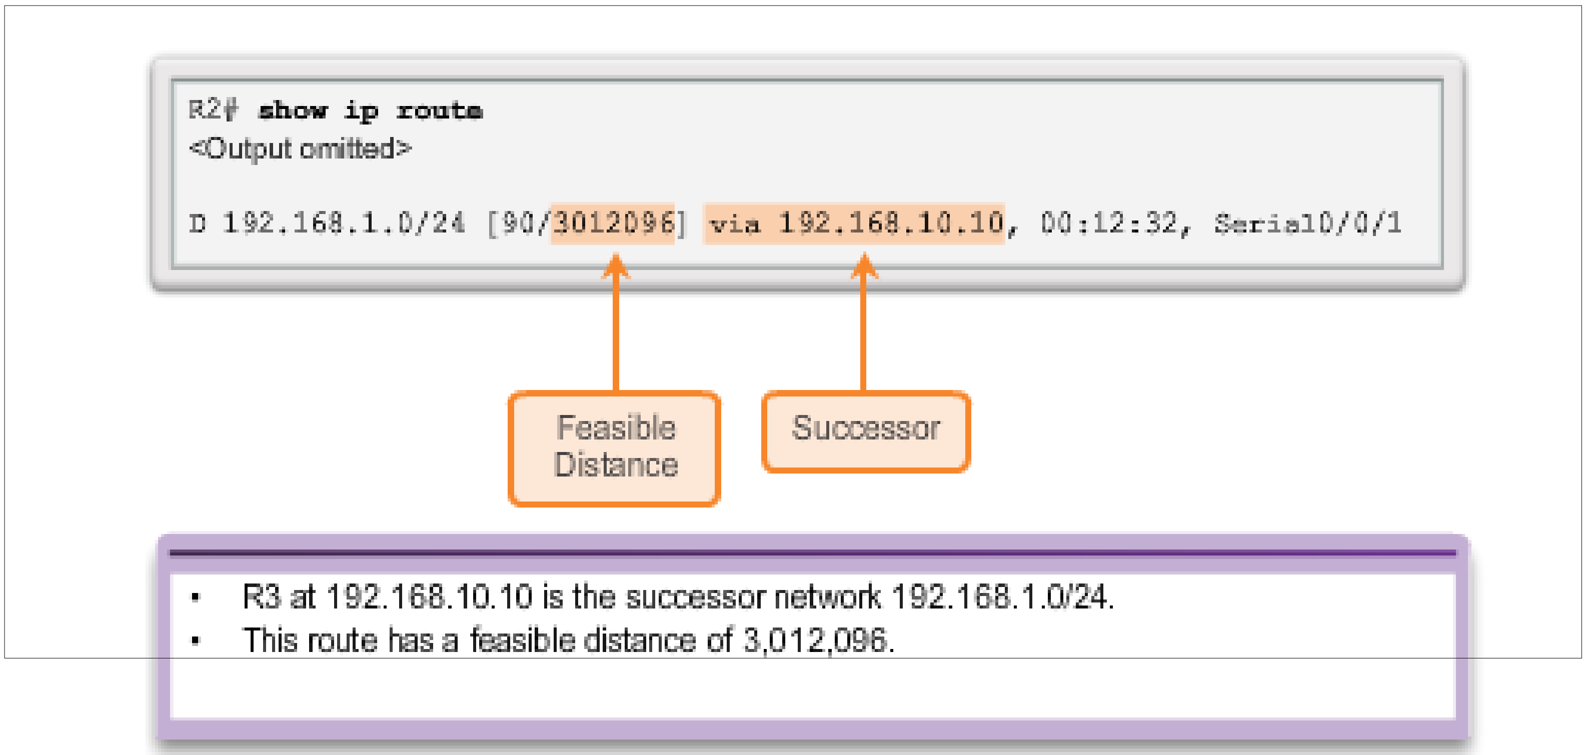

IV – EIGRP
DUAL – FS, FC & RD

- **Feasible Successor** (FS) is a neighbor that has a loop-free backup path to the same network as the successor, and it satisfies the Feasibility Condition (FC).
- **Feasibility Condition** (FC) is met when a neighbor's Reported Distance (RD) to a network is less than the local router's feasible distance to the same destination network.
- **Reported Distance** (RD) is an EIGRP neighbor's feasible distance to the same destination network.

IV – EIGRP
DUAL – Topology Table

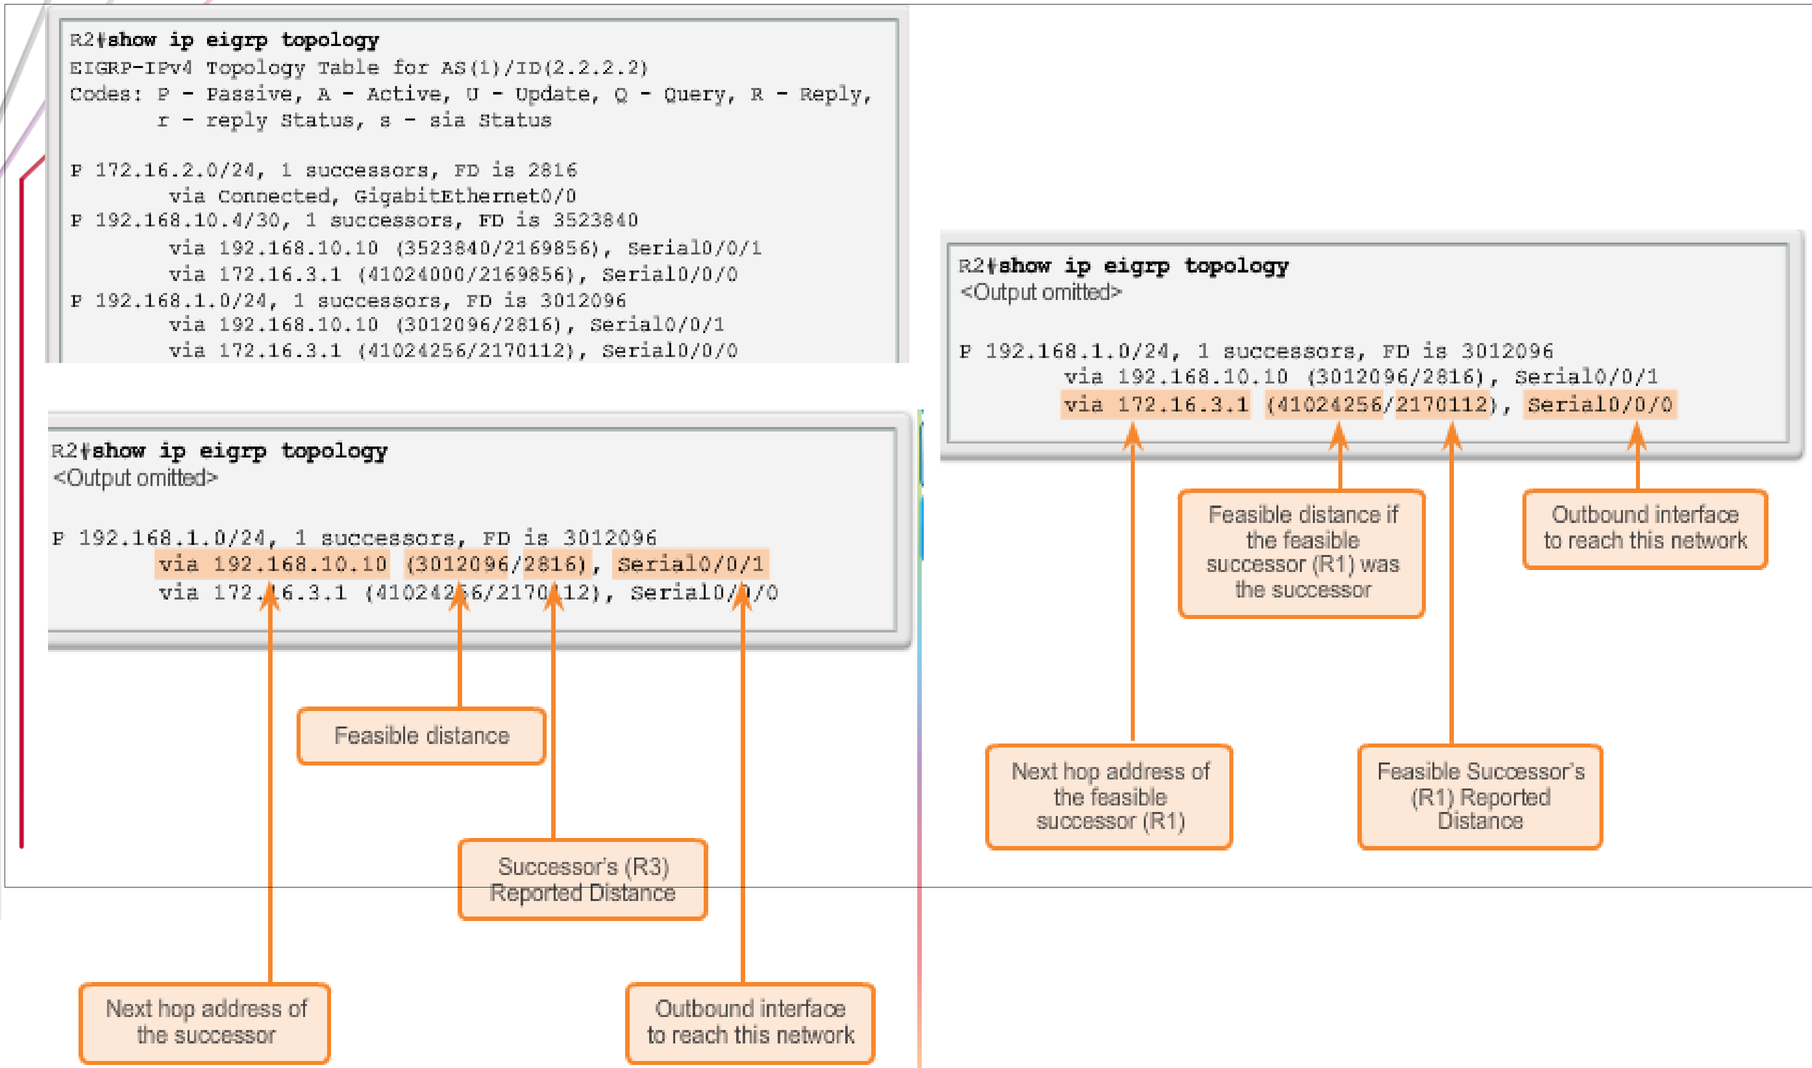

IV – EIGRP
DUAL – No Feasible Successor

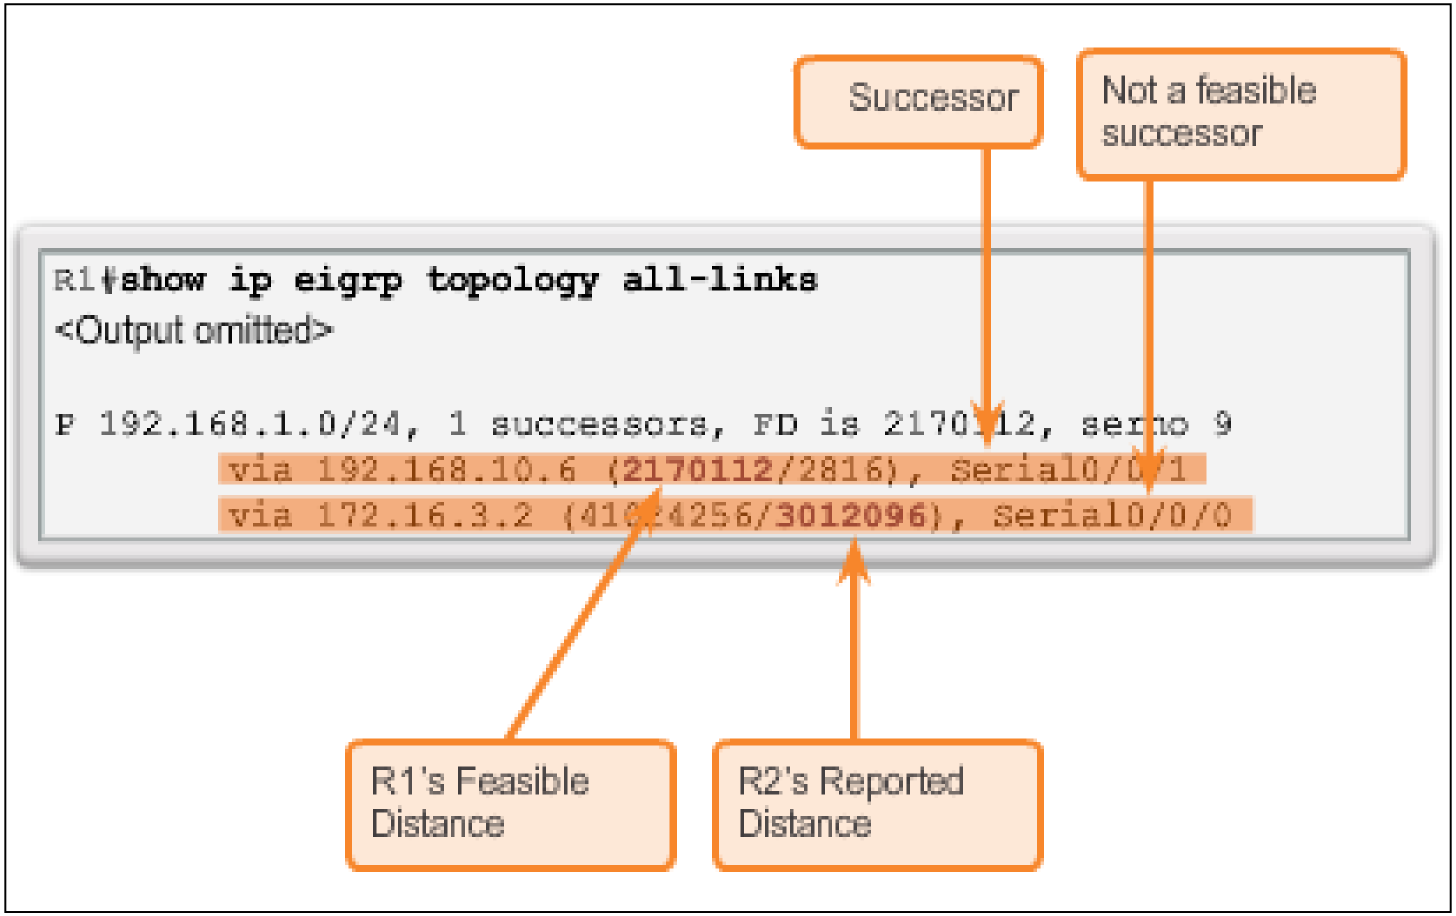

IV – EIGRP
DUAL – Finite State Machine (FSM)

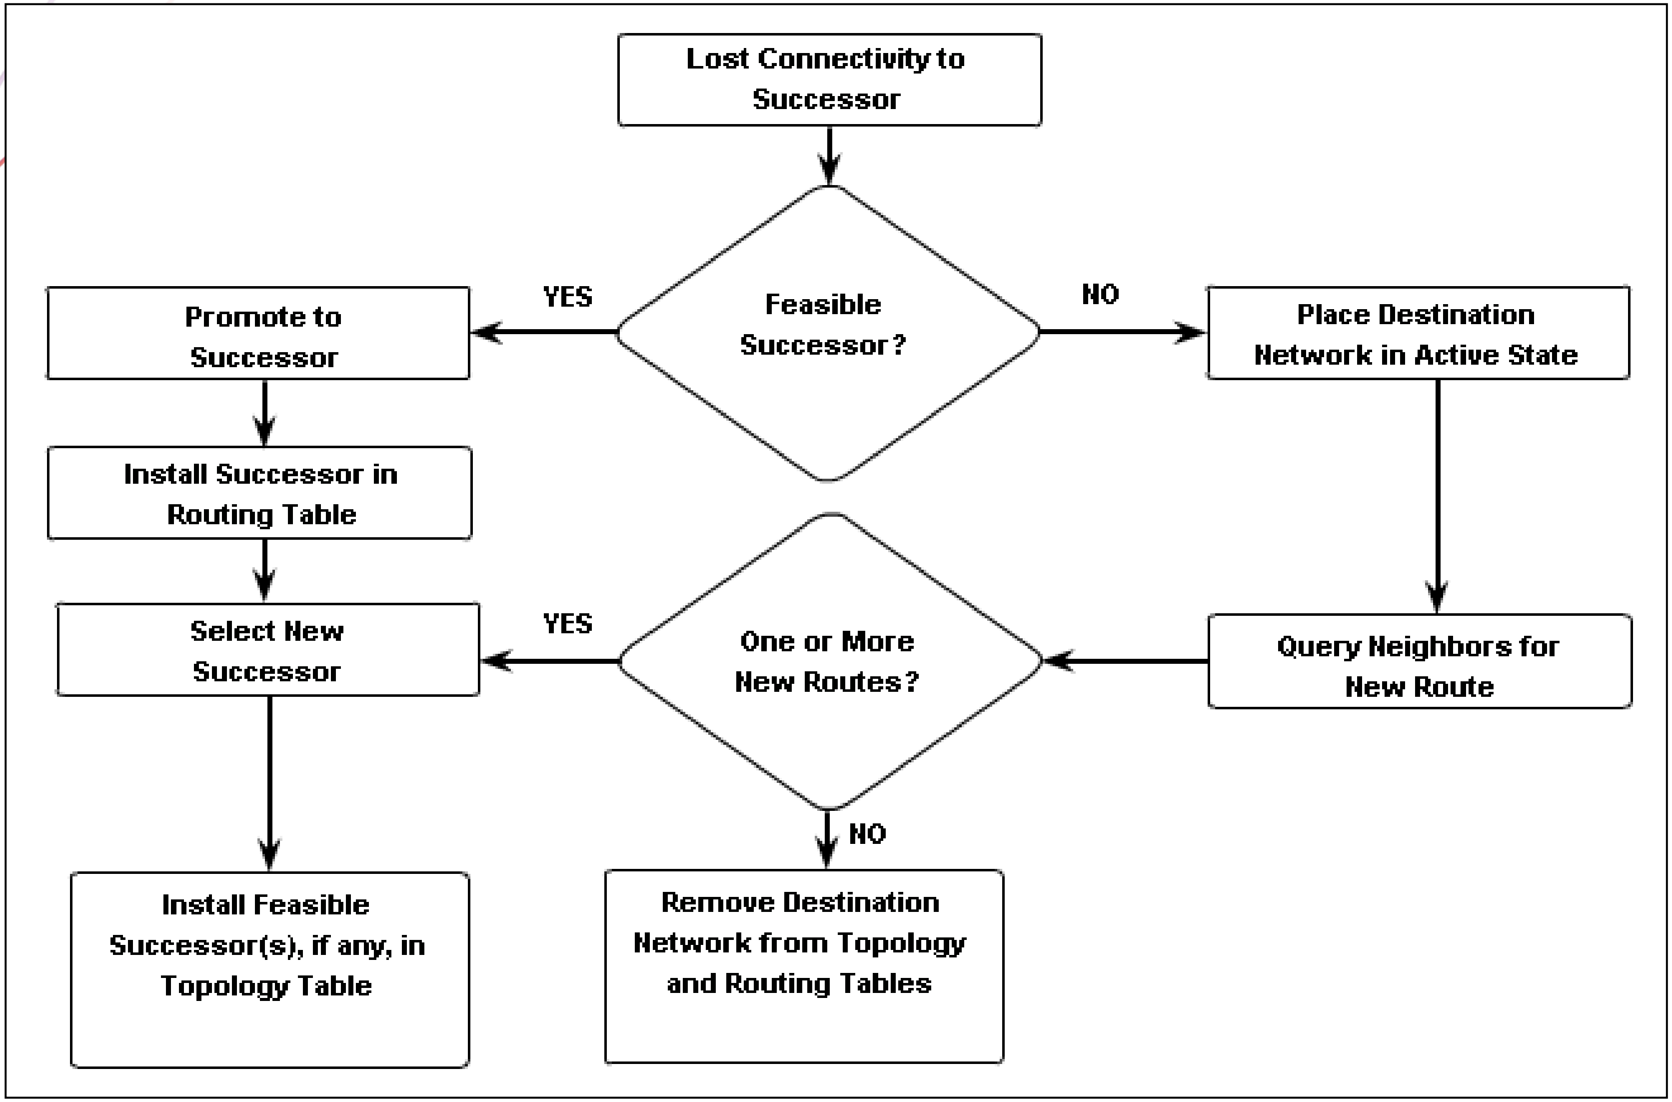

IV - OSPF

IV – OSPF
Features of OSPF

IV – OSPF
Components of OSPF

OSPF Data Structures

Database	Table	Description
Adjacency Database	Neighbor Table	• List of all neighbor routers to which a router has established bidirectional communication. • This table is unique for each router. • Can be viewed using the **`show ip ospf neighbor`** command.
Link-state Database (LSDB)	Topology Table	• Lists information about all other routers in the network. • The database shows the network topology. • All routers within an area have identical LSDB. • Can be viewed using the **`show ip ospf database`** command.
Forwarding Database	Routing Table	• List of routes generated when an algorithm is run on the link-state database. • Each router's routing table is unique and contains information on how and where to send packets to other routers. • Can be viewed using the **`show ip route`** command.

IV – OSPF
Components of OSPF

OSPF Routers Exchange Packets - These packets are used to discover neighboring routers and also to exchange routing information to maintain accurate information about the network.

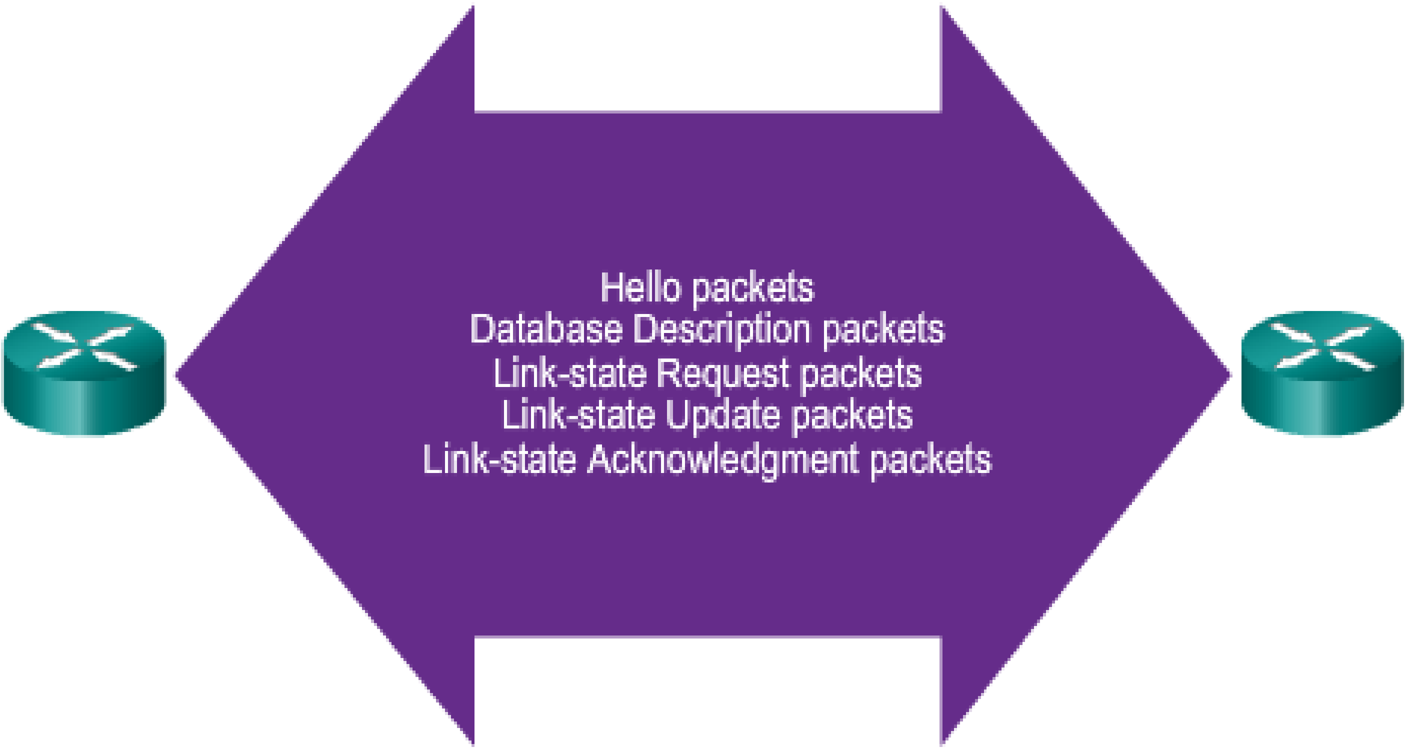

IV – OSPF
Link-State Operation – Hello Packets

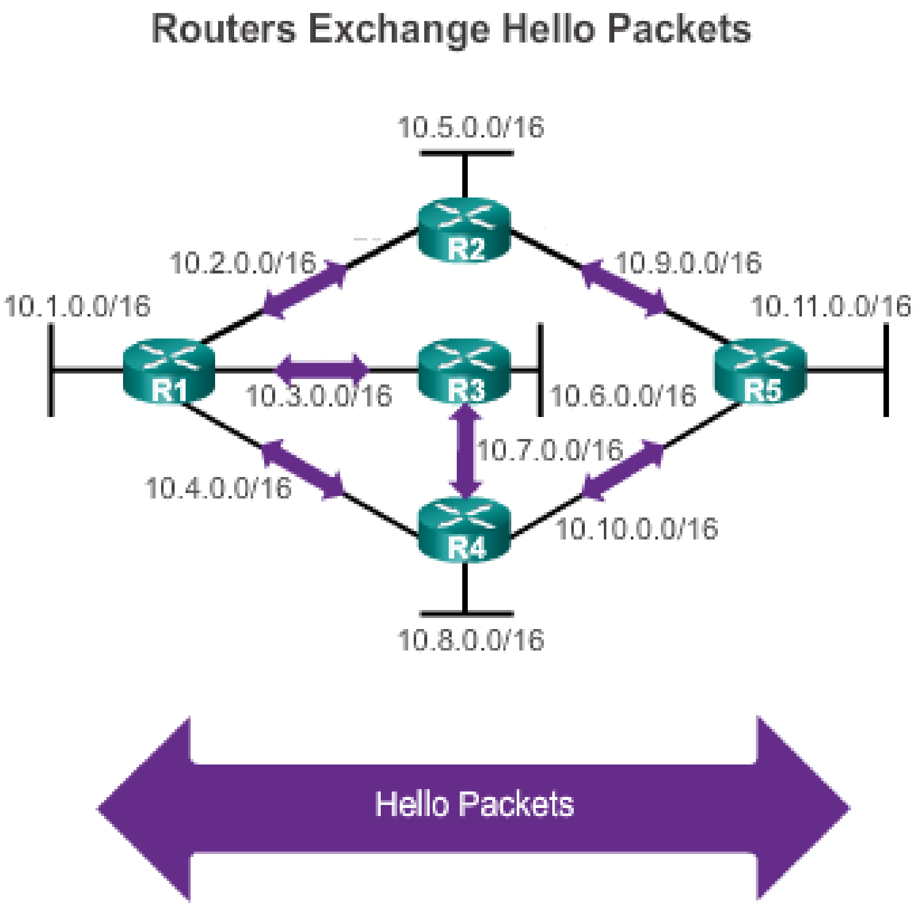

If a neighbor is present, the OSPF-enabled router attempts to establish a neighbor adjacency with that neighbor

IV – OSPF
Link-State Operation – LSA

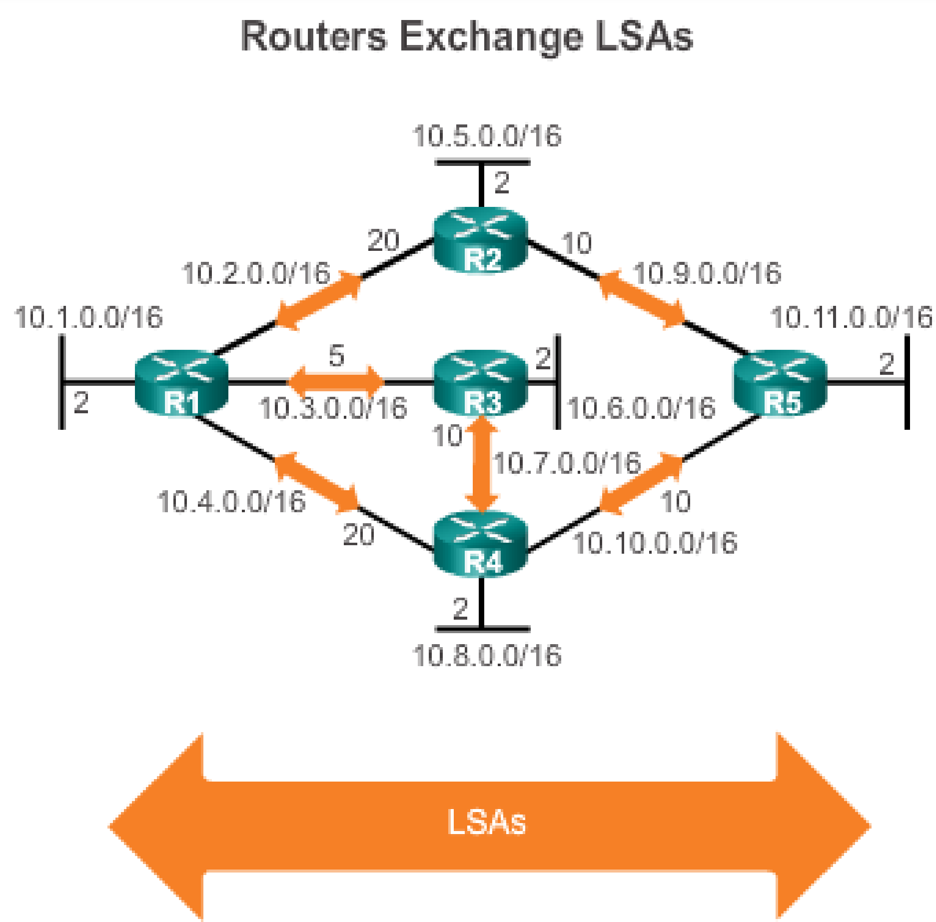

- LSAs contain the state and cost of each directly connected link.
- Routers flood their LSAs to adjacent neighbors.
- Adjacent neighbors receiving the LSA immediately flood the LSA to other directly connected neighbors, until all routers in the area have all LSAs.

IV – OSPF
Link-State Operation – SPF

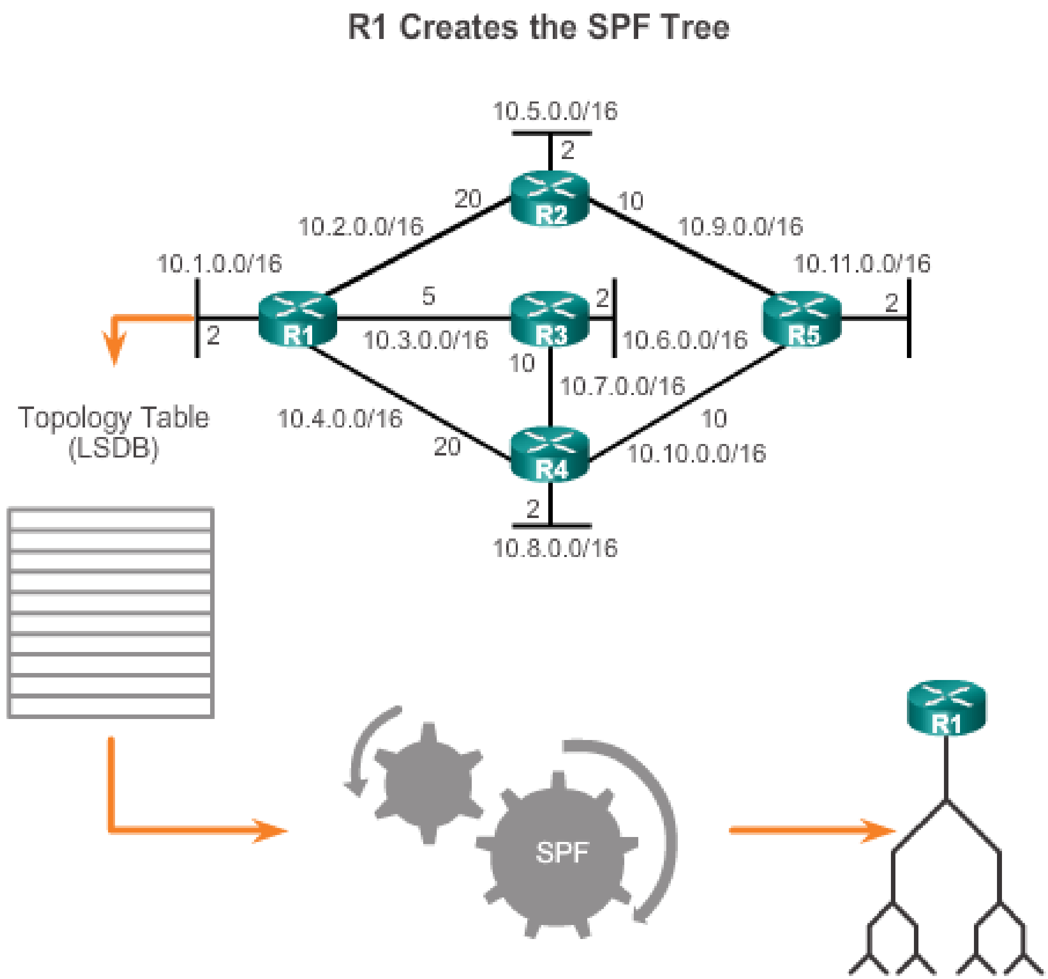

- Build the topology table based on the received LSAs.
- This database eventually holds all the information about the topology of the network.
- Execute the SPF Algorithm.

IV – OSPF
Link-State Operation – Best Paths

Content of the R1 SPF Tree

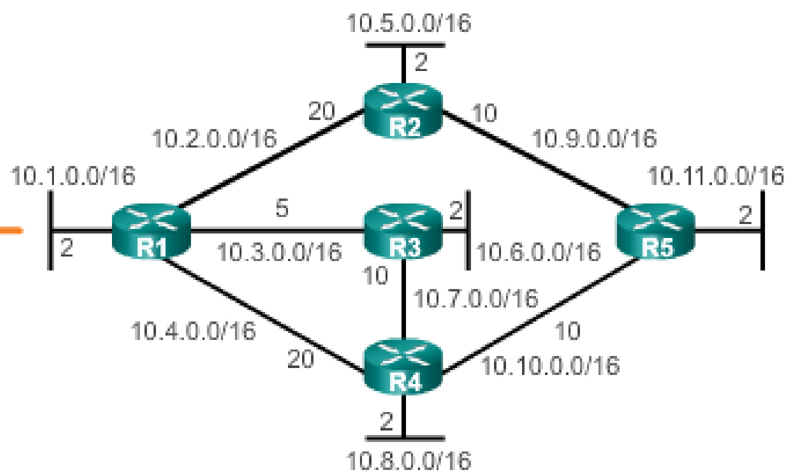

Destination	Shortest Path	Cost
10.5.0.0/16	R1 → R2	22
10.6.0.0/16	R1 → R3	7
10.7.0.0/16	R1 → R3	15
10.8.0.0/16	R1 → R3 → R4	17
10.9.0.0/16	R1 → R2	30
10.10.0.0/16	R1 → R3 → R4	25
10.11.0.0/16	R1 → R3 → R4 → R5	27

From the SPF tree, the best paths are inserted into the routing table.

IV – OSPF
Single-Area & Multiarea OSPF

Single-Area OSPF

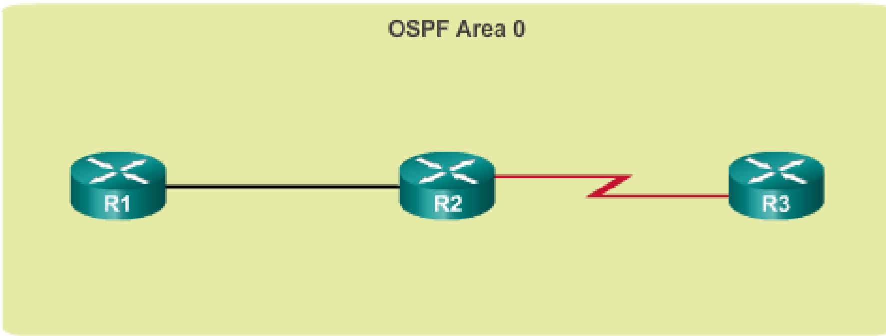

- Area 0 is also called the backbone area.
- Single-area OSPF is useful in smaller networks with few routers.

Multiarea OSPF

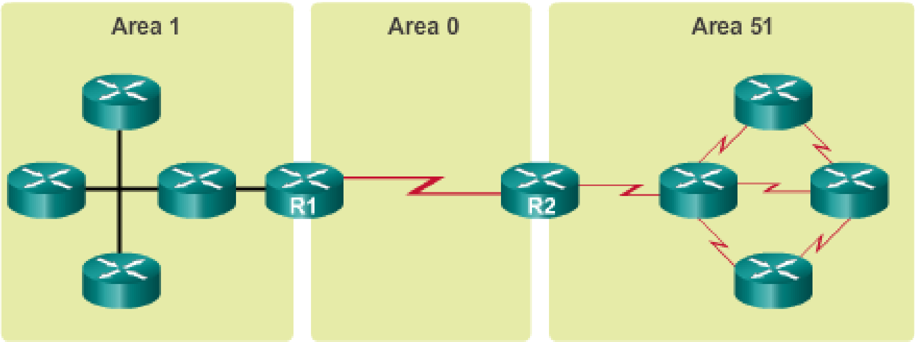

- Implemented using a two-layer area hierarchy as all areas must connect to the backbone area (area 0).
- Interconnecting routers are called Area Border Routers (ABR).
- Useful in larger network deployments to reduce processing and memory overhead.

IV – OSPF
Single-Area & Multiarea OSPF

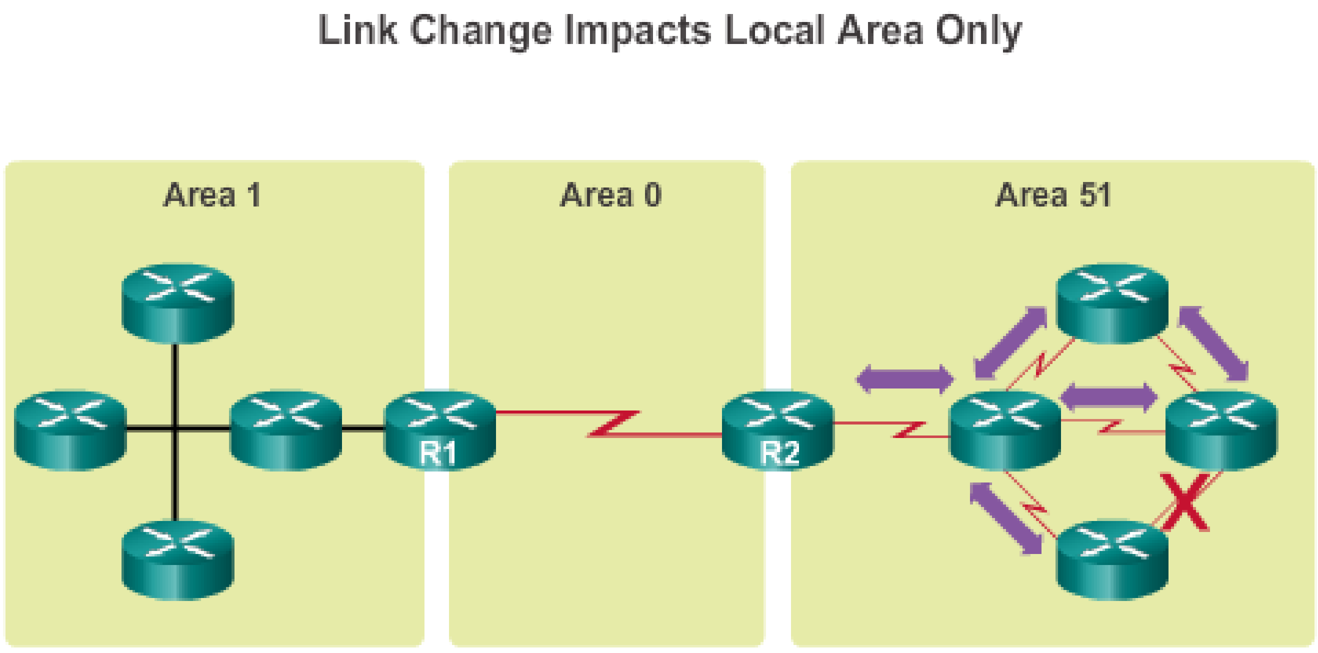

- Link failure affects the local area only (area 51).
- The ABR (R2) isolates the fault to area 51 only.
- Routers in areas 0 and 1 do not need the run the SPF algorithm.

Types of OSPF Packets

OSPF Packet Descriptions

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

OSPF Packets - Hello

OSPF Type 1 packet = Hello packet:

- Discover OSPF neighbors and establish neighbor adjacencies.
- Advertise parameters on which two routers must agree to become neighbors.
- Elect the Designated Router (DR) and Backup Designated Router (BDR) on multi access networks like Ethernet and Frame Relay.

IV – OSPF
OSPF Packets - Hello

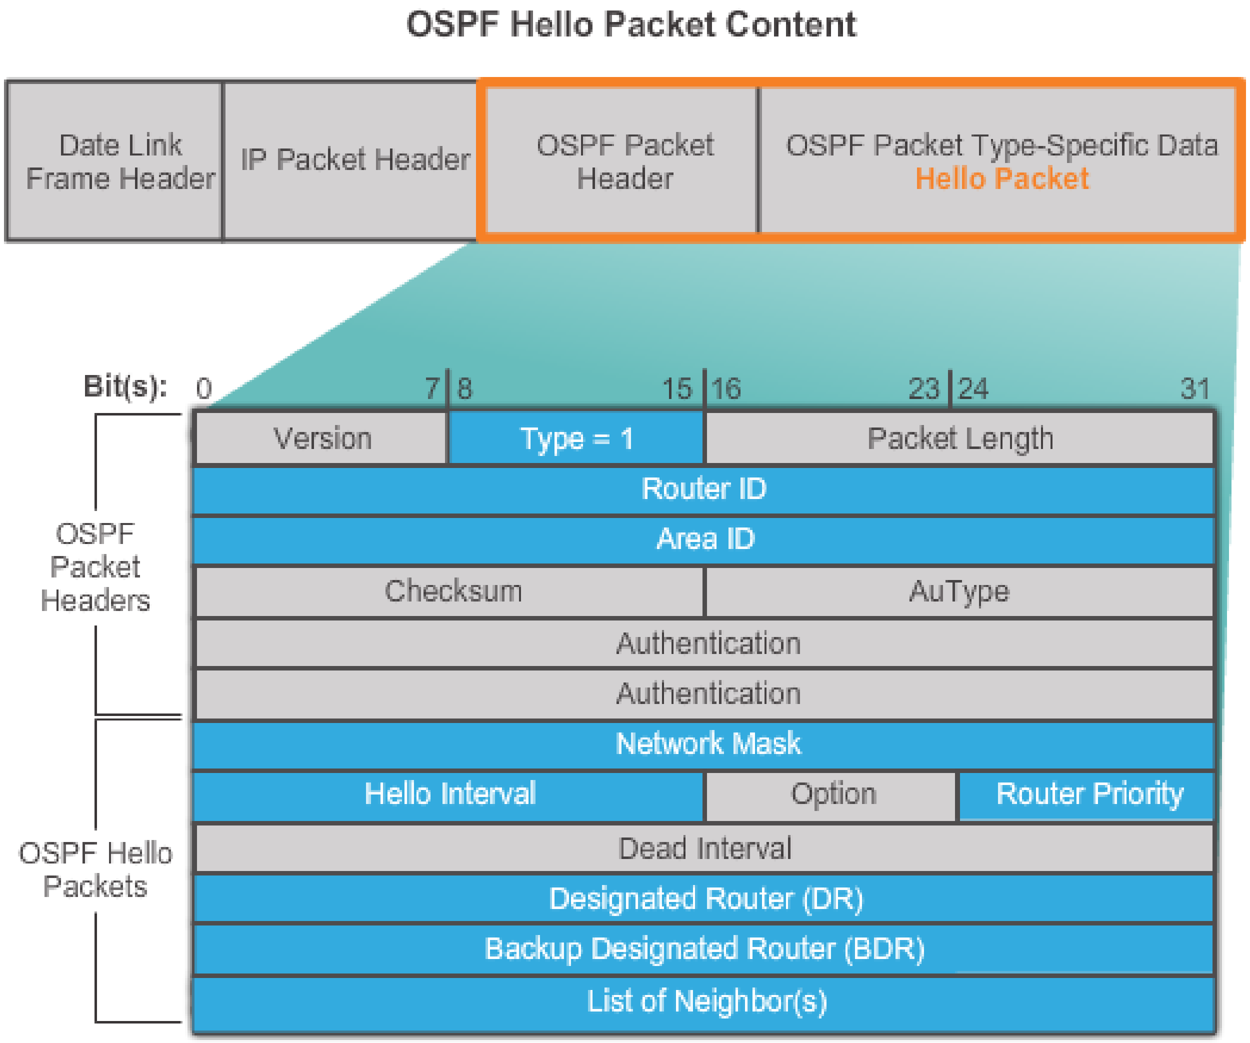

IV – OSPF
OSPF Packets – Link-State Updates

LSUs Contain 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

- An LSU contains one or more LSAs.
- LSAs contain route information for destination networks.

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

IV – OSPF
OSPF Operational States

When an OSPF router is initially connected to a network, it attempts to:

1. Create adjacencies with neighbors
2. Exchange routing information
3. Calculate the best routes
4. Reach convergence

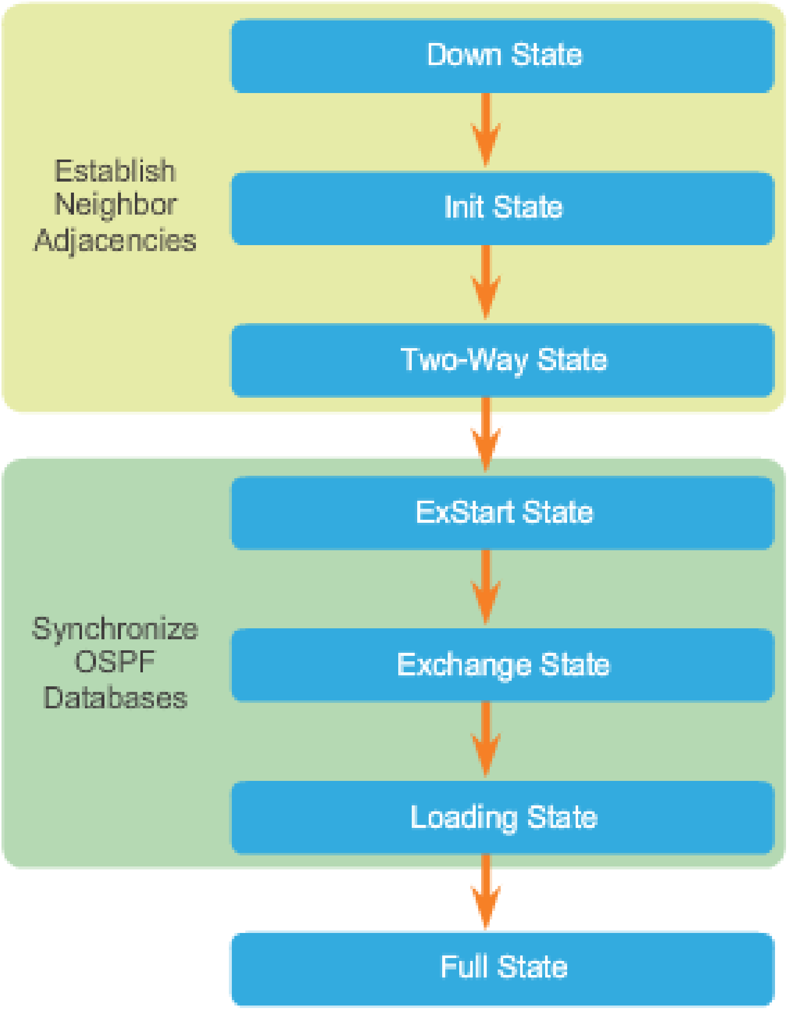

IV – OSPF
Establish Neighbor Adjacencies

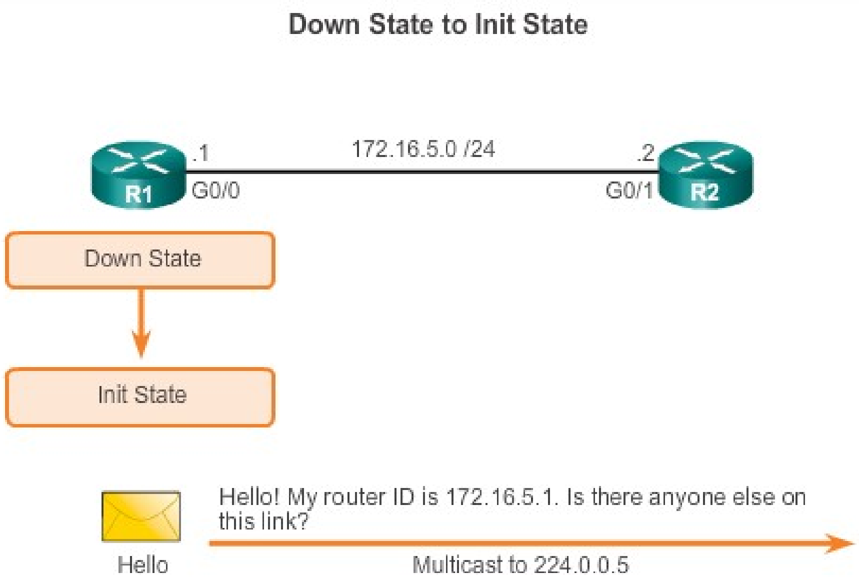

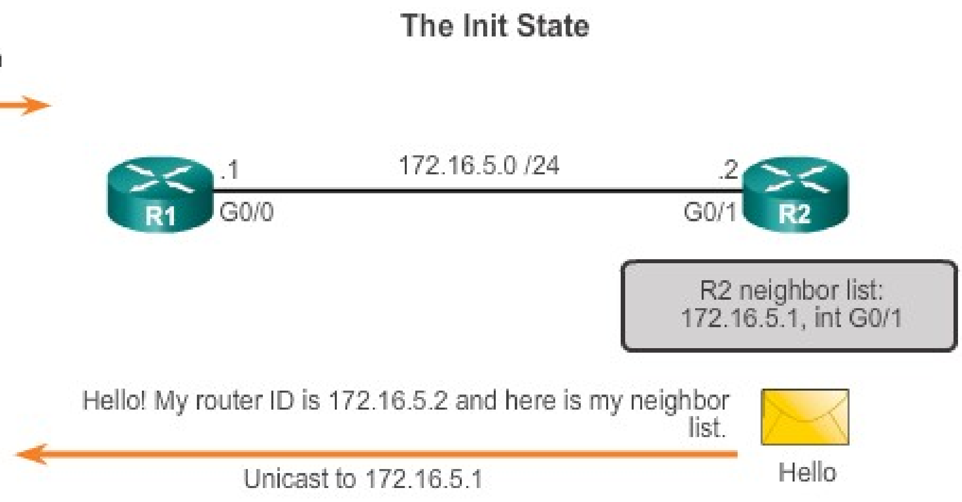

IV – OSPF
Establish Neighbor Adjacencies

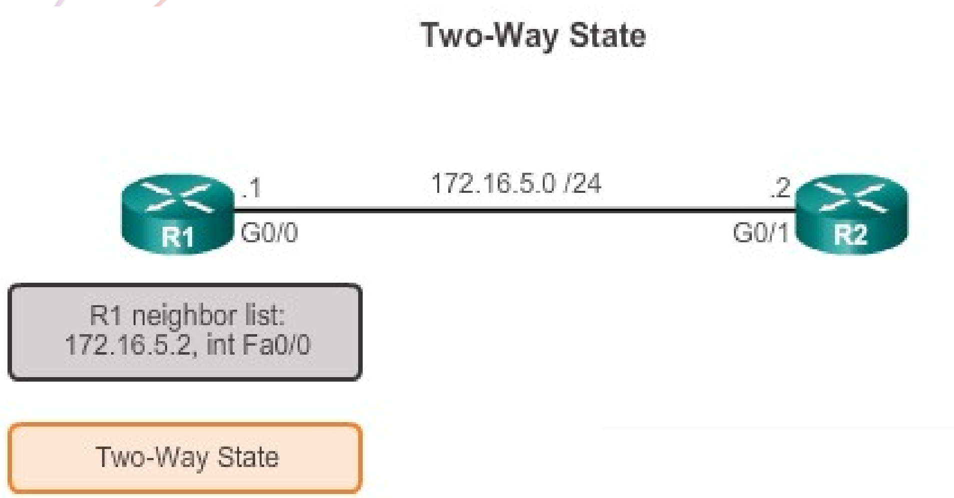

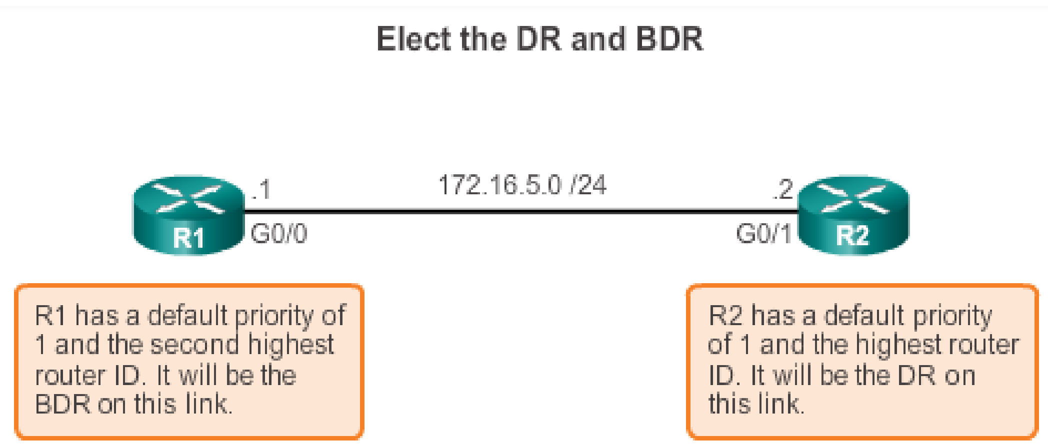

OSPF DR & BDR

Creating Adjacencies With Every Neighbor

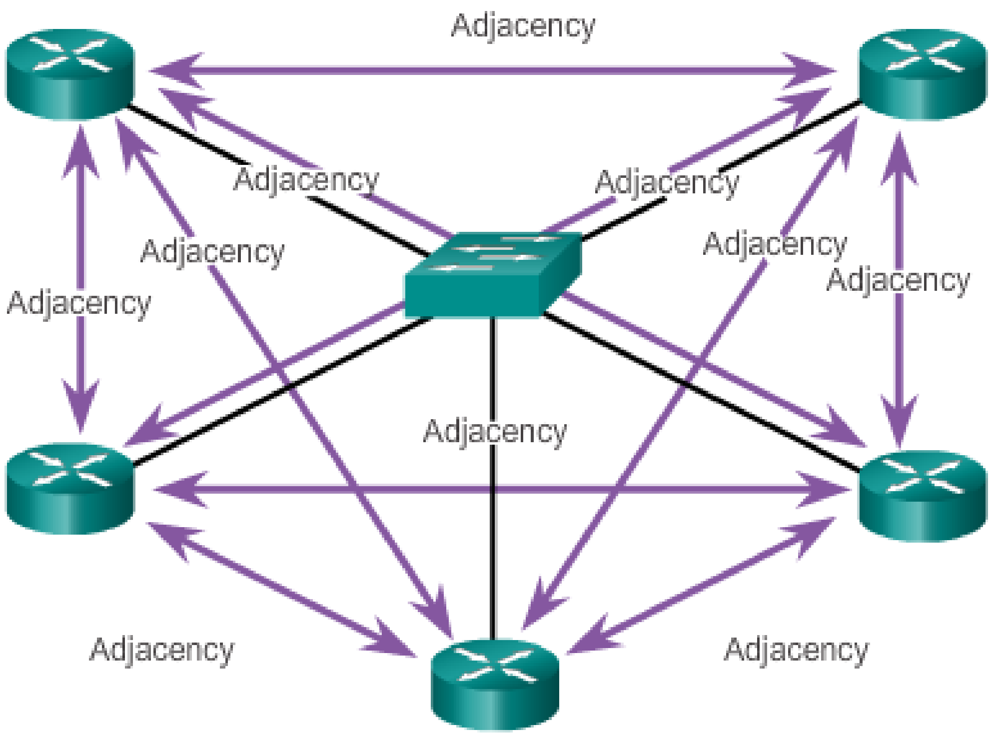

Number of Adjacencies=n(n-1)/2
n=number of routers
Example:5 routers (5-1)/2=10 adjacencies

IV – OSPF
Synchronizing OSPF Database

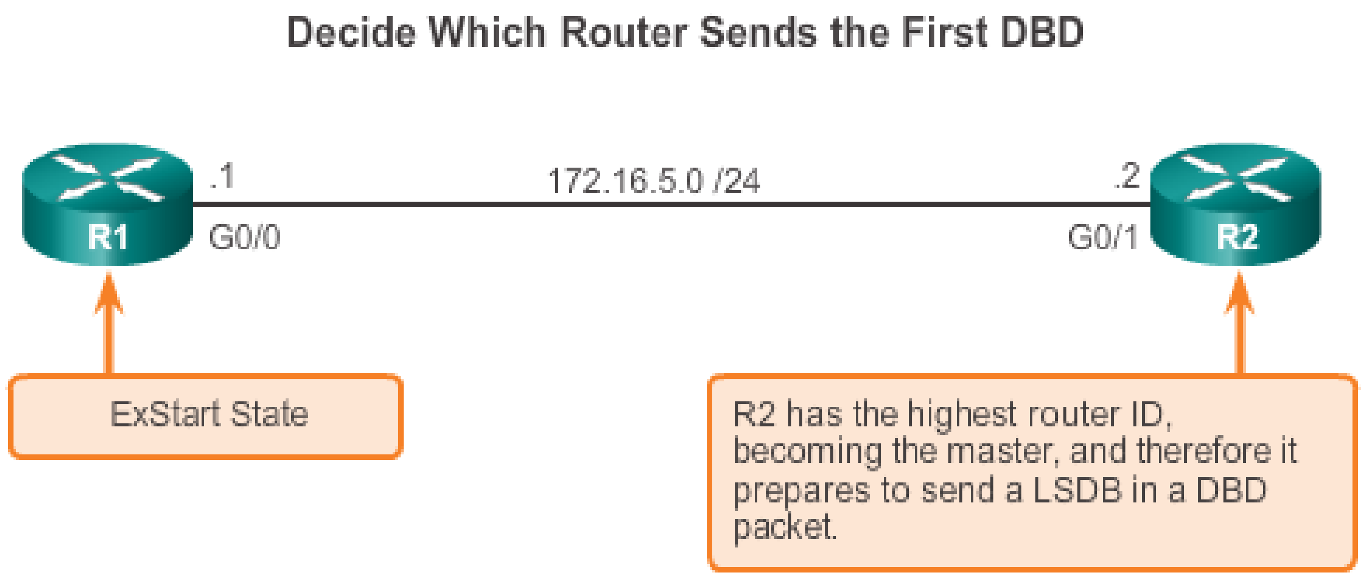

IV – OSPF
Synchronizing OSPF Database

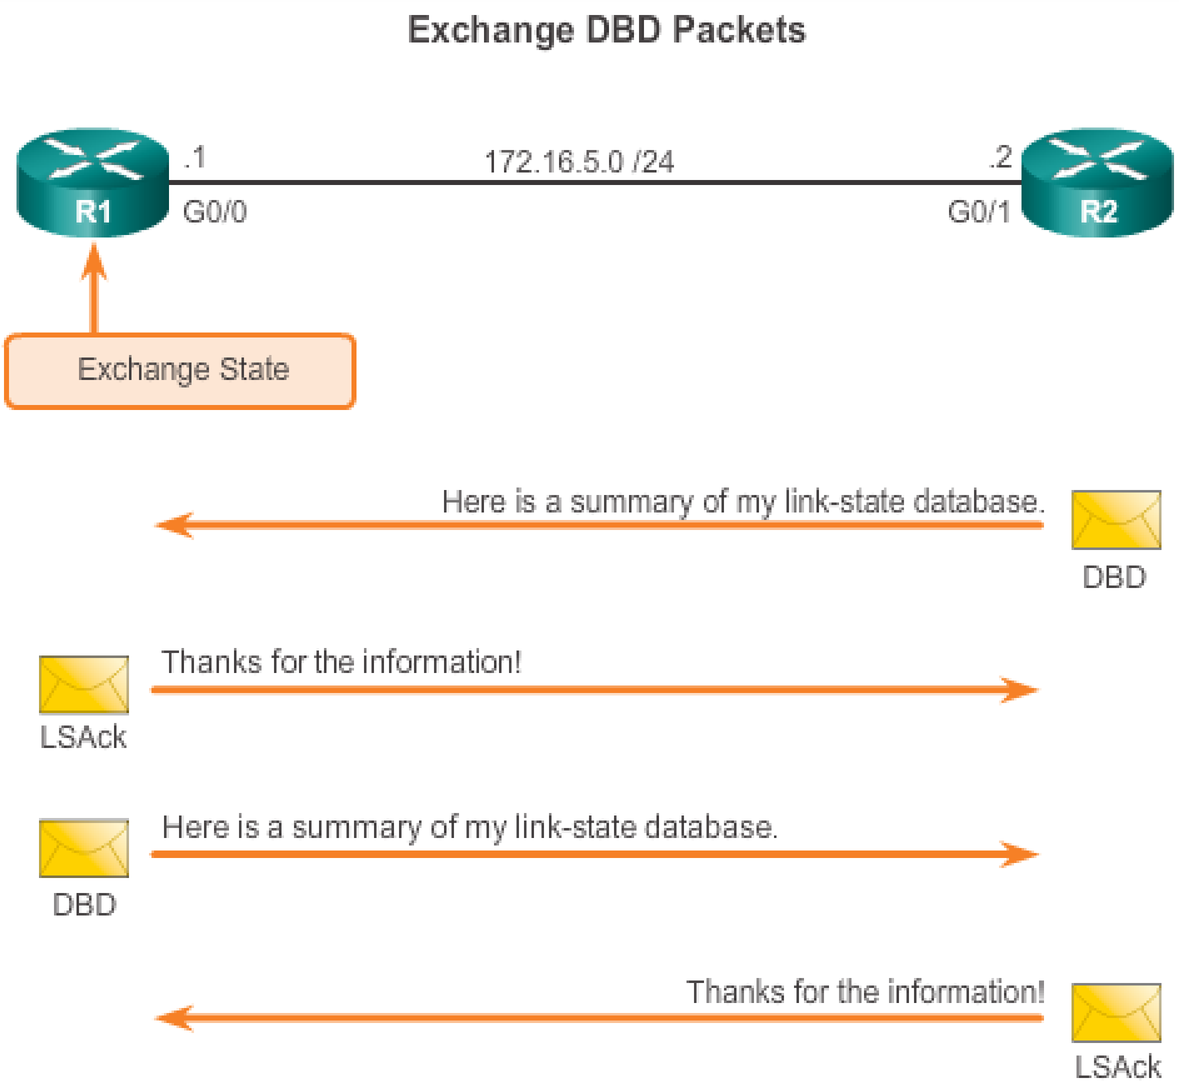

OSPF Network Topology

ESIEE AMIENS

Entering Router OSPF Configuration Mode on R1

```
R1(config)# router ospf 10
R1(config-router)# ?
Router configuration commands:
  auto-cost                 Calculate OSPF interface cost
                            according to bandwidth
  network                   Enable routing on an IP network
  no                        Negate a command or set its defaults
  passive-interface         Suppress routing updates on an
                            interface
  priority                  OSPF topology priority
  router-id                 router-id for this OSPF process
```

Note: Output has been altered to display only the commands that will be used in this chapter.

IV – OSPF
Router IDs

```
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#
*Mar 25 19:46:09.711: %SYS-5-CONFIG_I: Configured from
console by console
```

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

Router ID Order of Precedence

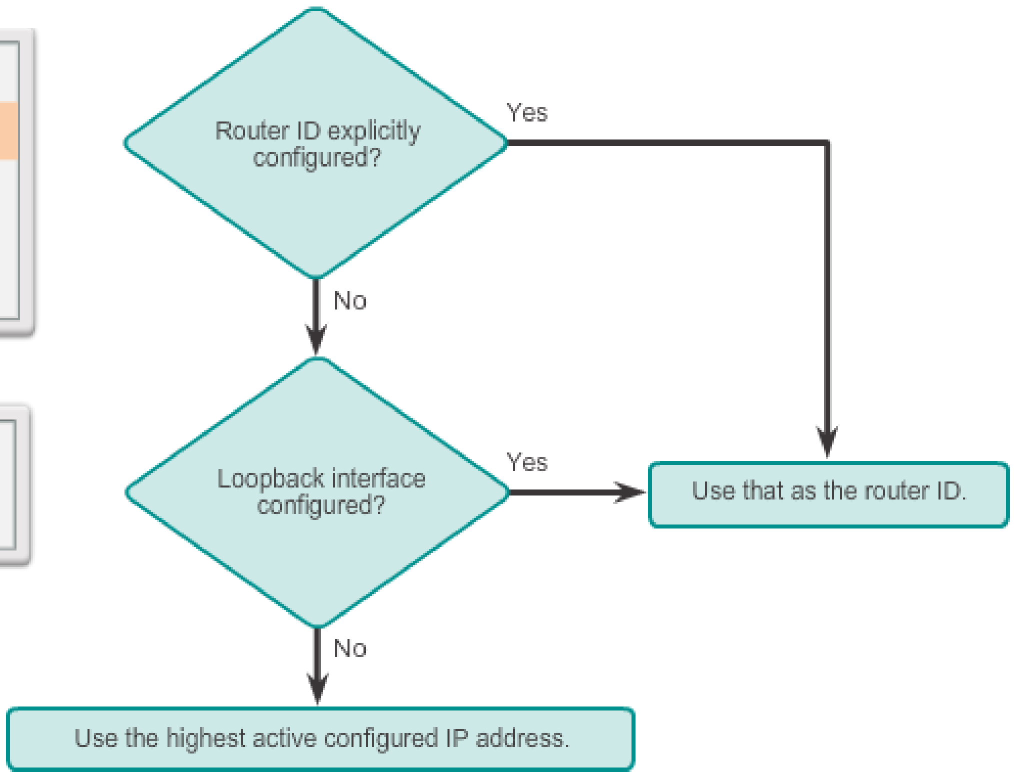

Clearing the OSPF Process

```
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
```

IV – OSPF
Passive Interface

- By default, OSPF messages are forwarded out all OSPF-enabled interfaces. However, these messages really only need to be sent out interfaces connecting to other OSPF-enabled routers.

- Sending out unneeded messages on a LAN affects the network in three ways:
  - Inefficient Use of Bandwidth
  - Inefficient Use of Resources
  - Increased Security Risk

- The Passive Interface feature helps limiting the scope of routing updates advertisements.

IV – OSPF
OSPF Metric = Cost

Cost = *reference bandwidth* / *interface bandwidth*
(default reference bandwidth is 10^8)
Cost = 100,000,000 bps / *interface bandwidth in bps*

Default Cisco OSPF Cost Values

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
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

Same Cos
due to
reference
bandwidtl

IV – OSPF
OSPF Accumulates Costs

Cost of an OSPF route is the accumulated value from one router to the destination network.

```
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#
```

IV – OSPF
SPF: Example

Determine the shortest path from A to the others.