Chapter – 5

Distance Vector Routing Protocols

Distance Vector Routing Protocols

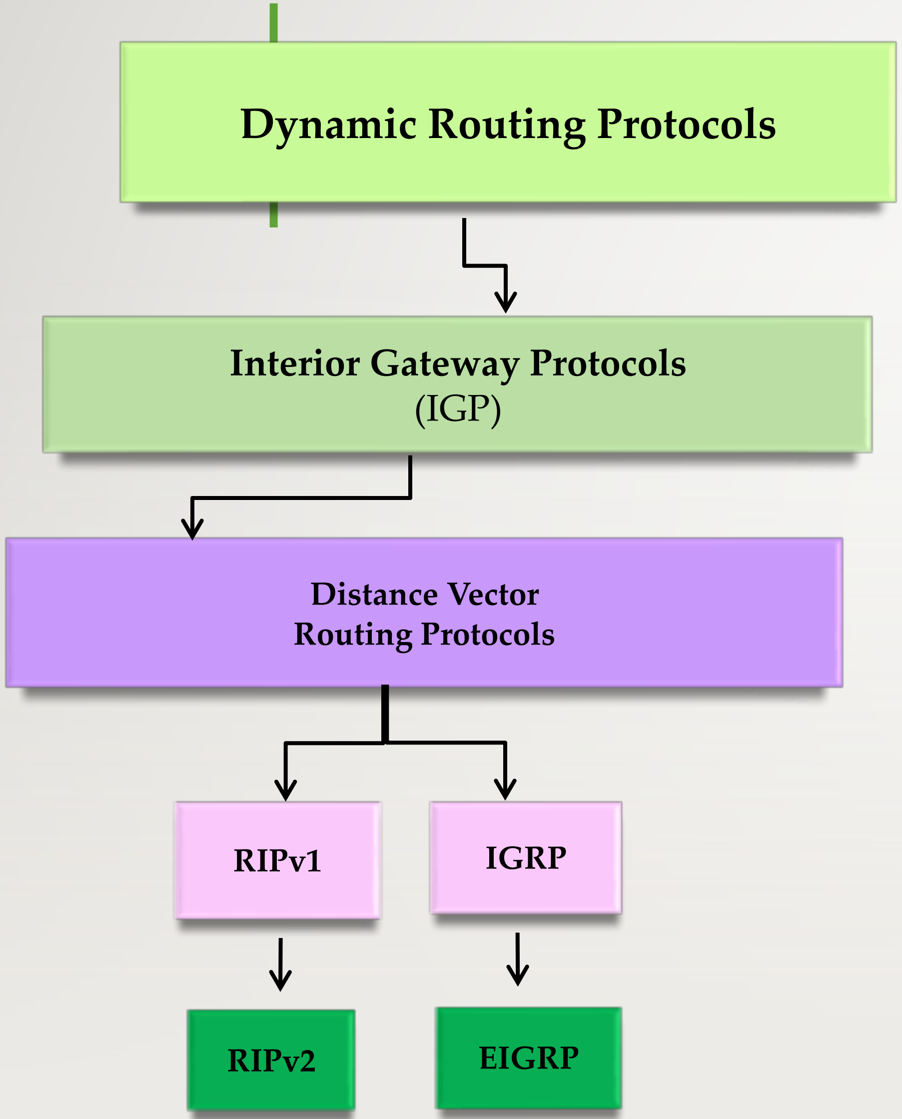

- **Distance vector routing protocols:**
  - Share updates between neighbors
  - Not aware of the network topology
- RIPv1 sends periodically broadcasts updates to IP **255.255.255.255** even if topology has not changed
  - Updates consume bandwidth and network device CPU resources
- EIGRP will only send an update when topology has changed
- RIPv2 and EIGRP use multicast addresses

Routing Information Protocol ver. 1 (RIPv1)

Data Link Frame Header	IP Packet Header	UDP Segment Header	RIP Message (512 Bytes: Up to 25 Routes)

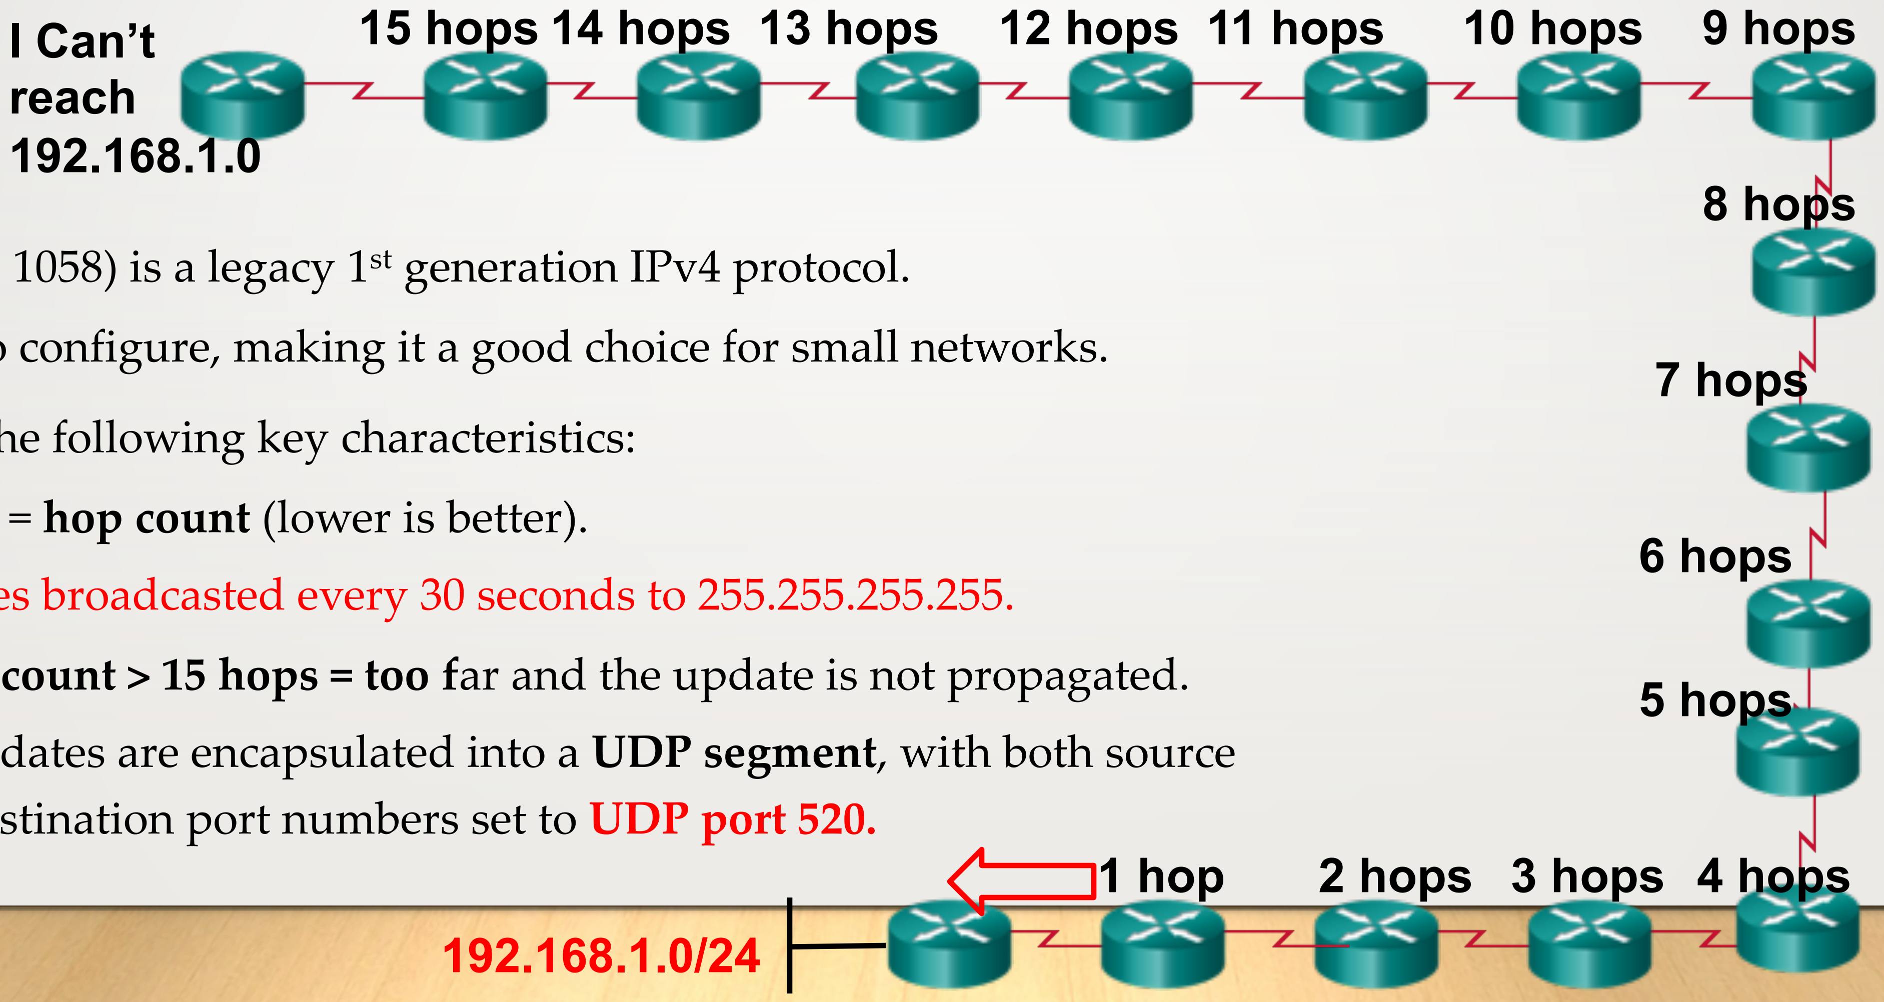

- RIPv1 (RFC 1058) is a legacy 1st generation IPv4 protocol.
  - Easy to configure, making it a good choice for small networks.
- RIPv1 has the following key characteristics:
  - **Metric** = **hop count** (lower is better).
  - Updates broadcasted every 30 seconds to 255.255.255.255.
  - **If hop count > 15 hops = too** far and the update is not propagated.
  - RIP updates are encapsulated into a **UDP segment**, with both source and destination port numbers set to **UDP port 520.**

Routing Information Protocol ver. 2 (RIPv2)

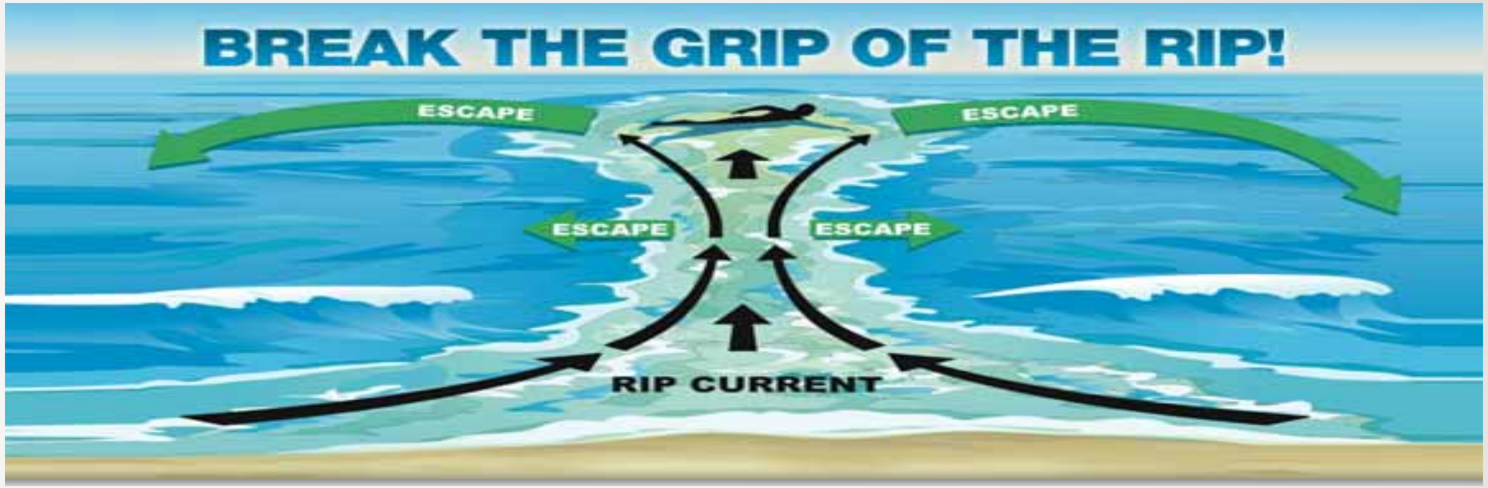

- RIPv2 (RFC 1058) replaced RIPv1and included the following improvements:
  - **Classless routing protocol:** Supports VLSM and CIDR, because it includes the subnet mask in the routing updates.
  - **Increased efficiency:** Forwards updates to **multicast address 224.0.0.9**, instead of the broadcast address 255.255.255.255.
  - **Reduced routing entries:** Supports manual route summarization.
  - **Secure:** Supports an authentication mechanism to secure routing table updates between neighbors.

Advertising the R1 Networks

- To enable RIP and advertise a network, use the routing configuration command **network** *network-address*
  - Enter the network address for ***each directly connected network***.
- Entering the command automatically :
  - Enables RIP on all interfaces that belong to a specific network.
    - Interfaces now both send and receive RIP updates.
  - Advertises the specified network in RIP updates every 30 seconds.

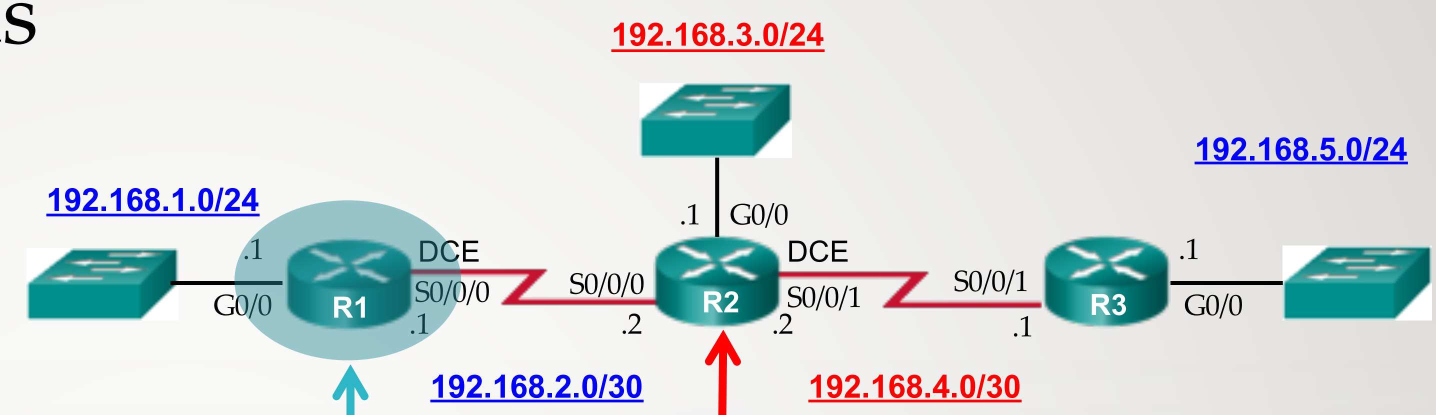


```
R1(config)# router rip
R1(config-router)# network 192.168.1.0
R1(config-router)# network 192.168.2.0
R2(config)# router rip
R2(config-router)# network 192.168.2.0
R2(config-router)# network 192.168.3.0
R2(config-router)# network 192.168.4.0
R2(config-router)#
```

Discover Directly Attached Networks

- R1 adds the 10.1.0.0 network available through interface FastEthernet 0/0 and 10.2.0.0 is available through interface Serial 0/0/0.
- R2 adds the 10.2.0.0 network available through interface Serial 0/0/0 and 10.3.0.0 is available through interface Serial 0/0/1.
- R3 adds the 10.3.0.0 network available through interface Serial 0/0/1 and 10.4.0.0 is available through interface FastEthernet 0/0.

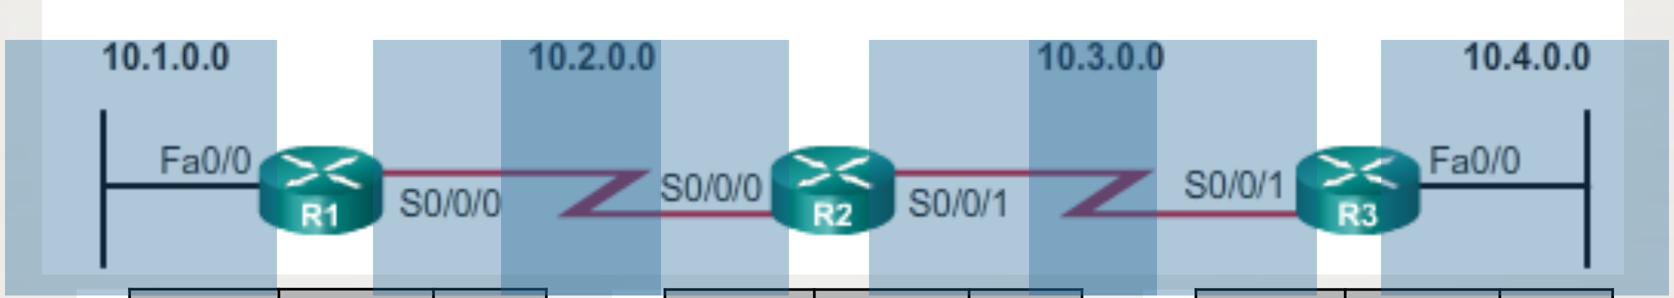

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0

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

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

Network Discovery – Initial Exchange

- R1 sends an update about network 10.1.0.0 out the Serial0/0/0 interface and sends an update about network 10.2.0.0 out of Fa0/0.
- R2 sends an update about network 10.3.0.0 out of Serial 0/0/0 and sends an update about network 10.2.0.0 out of Serial 0/0/1.
- R3 sends an update about network 10.4.0.0 out of Serial 0/0/1 and sends an update about network 10.3.0.0 out of FastEthernet0/0.

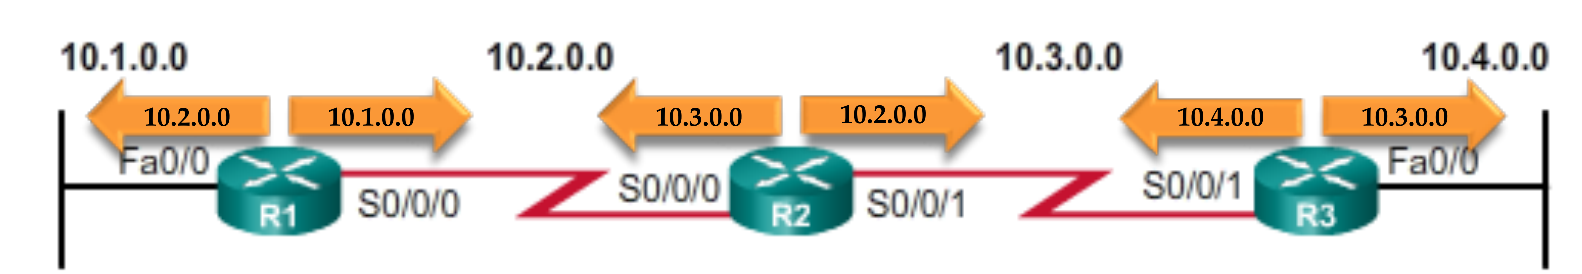

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0

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

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

This example is only for ***distance vector*** routing protocols (not link state routing protocols). More later!

Update Routing Table

- R1 receives the update from R2 about network 10.3.0.0, increments the hop count by 1, and stores the network in the routing table (metric of 1).
- R2 receives the update from R1 about network 10.1.0.0 and R3 about network 10.4.0.0, increments and stores both networks in the routing table (metric of 1).
- R3 receives the update from R2 about network 10.2.0.0, increments the hop count by 1, and stores the network in the routing table (metric of 1).

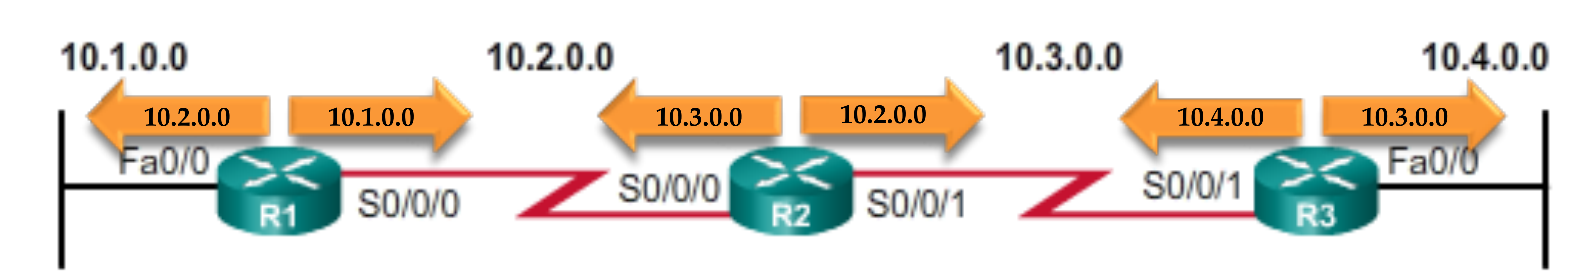

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0

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

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

This example is only for ***distance vector*** routing protocols (not link state routing protocols). More later!

Network Discovery – Next Exchange

- R1 sends an update about 10.1.0.0 out Serial0/0/0 and networks 10.2.0.0 metric 1) and 10.3.0.0 out of Fa0/0.
- R2 sends an update about 10.3.0.0 and 10.4.0.0 (metric 1) out Serial 0/0/0 and network 10.2.0.0 and 10.1.0.0 (metric 1) out of Serial 0/0/1.
- R3 sends an update about 10.4.0.0 out Serial 0/0/1 and networks 10.2.0.0 (metric 1) and 10.3.0.0 out FastEthernet0/0.

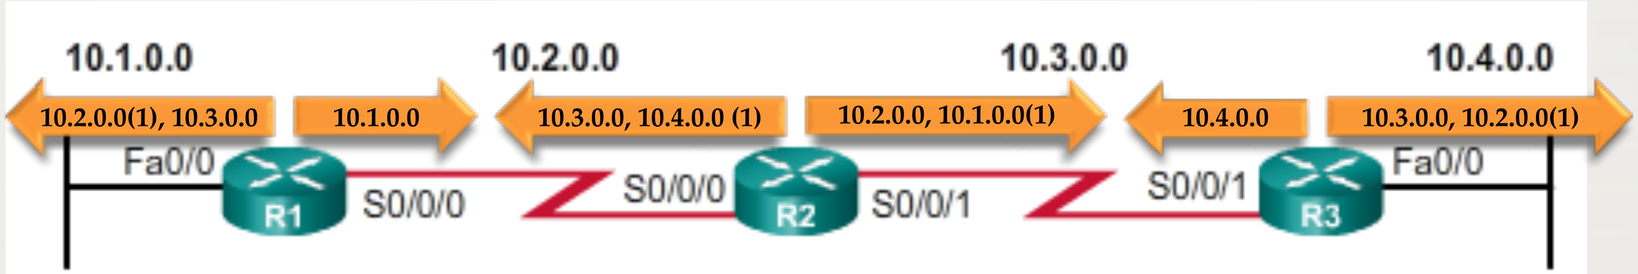

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0
R	10.3.0.0	S0/0/0	1

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

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

This example is only for ***distance vector*** routing protocols (not link state routing protocols). More later!

Update Routing Table

- R1 receives the update from R2 about network 10.3.0.0 and 10.4.0.0 (1).
  - Refreshes the information for 10.3.0.0.
  - It increments the 10.4.0.0 hop count by 1 and stores it in the routing table (metric 2).
  - *10.4.0.0 is 1 hop for R2, so 2 hops for R1*

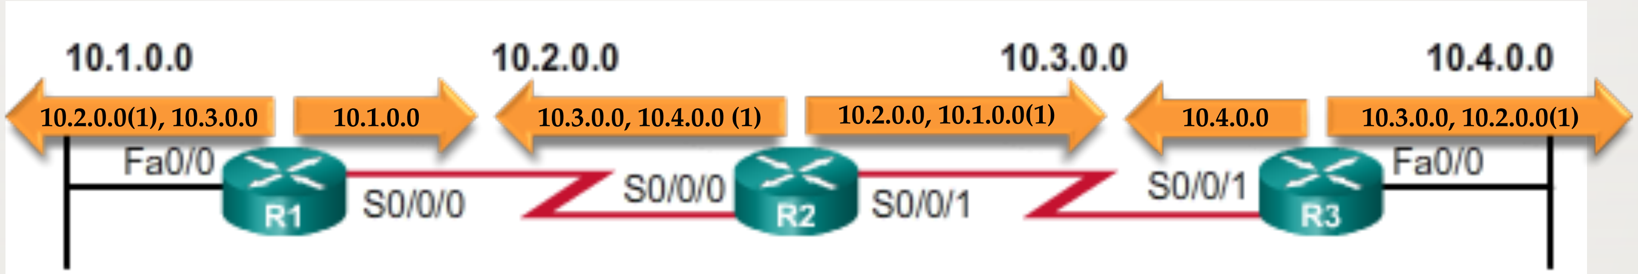

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0
R	10.3.0.0	S0/0/0	1

	Network	Interface	Hop
C	10.2.0.0	S0/0/0	0
C	10.3.0.0	S0/0/1	0
R	10.1.0.0	S0/0/0	1

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

This example is only for ***distance vector*** routing protocols (not link state routing protocols). More later!

Update Routing Table

- R2 receives the update from R1 about network 10.1.0.0 and from R3 about network 10.4.0.0 and refreshes the routing table.
- ***No changes to routing table because it didn't learn any new routes or of a better route to any networks.***

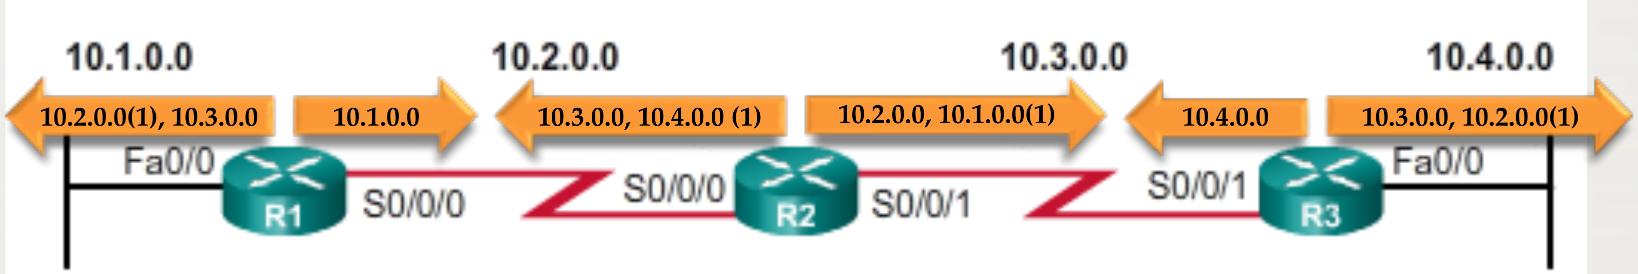

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0
R	10.3.0.0	S0/0/0	1
R	10.4.0.0	S0/0/0	2

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

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

This example is only for ***distance vector*** routing protocols (not link state routing protocols). More later!

Update Routing Table

- R3 receives the update from R2 about network 10.1.0.0 (1) and 10.2.0.0.
  - Refreshes the information for 10.2.0.0.
  - It increments the 10.1.0.0 hop count by 1 and stores it in the routing table (metric 2).
  - ***10.1.0.0 is 1 hop for R2 so 2 hops or R3***

Network Has Converged

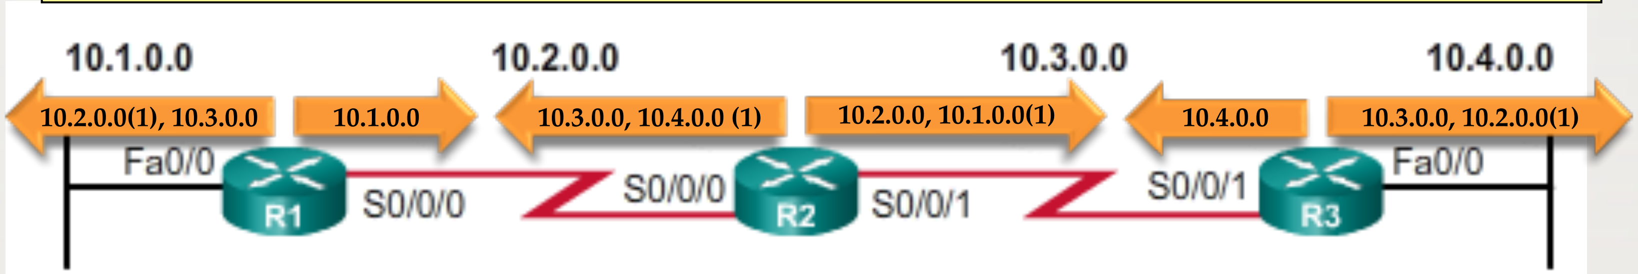

	Network	Interface	Hop
C	10.1.0.0	Fa0/0	0
C	10.2.0.0	S0/0/0	0
R	10.3.0.0	S0/0/0	1
R	10.4.0.0	S0/0/0	2

	Network	Interface	Hop
C	10.2.0.0	S0/0/0	0
C	10.3.0.0	S0/0/1	0
R	10.4.0.0	S0/0/1	1

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

Auto-Summary and RIPv2

What do you expect to see?

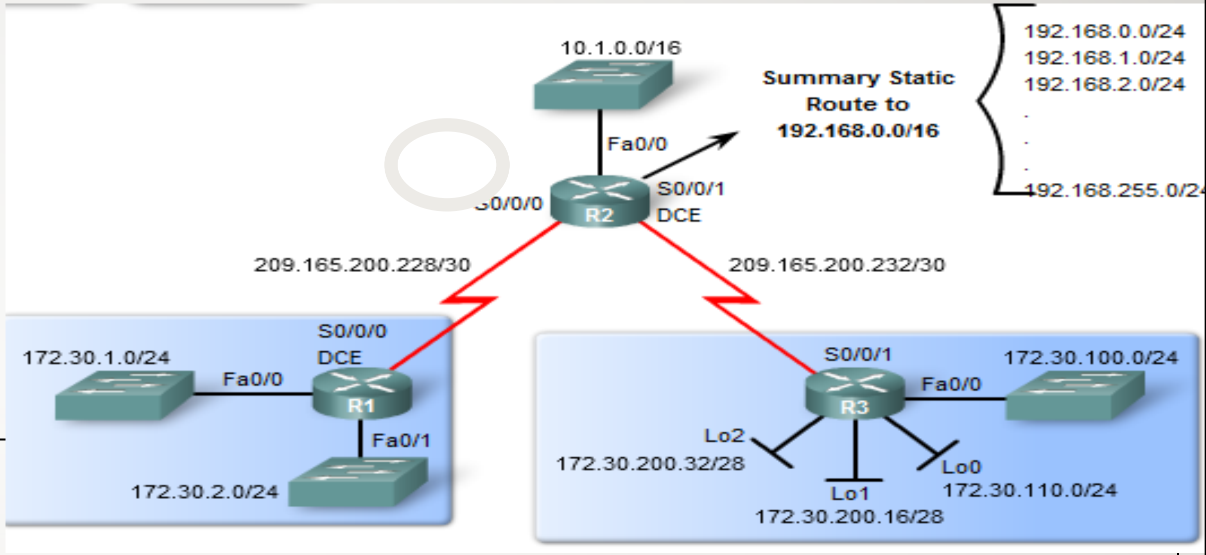


```
R2# show ip route

R    172.30.0.0/16 [120/1] via 209.165.200.230, 00:00:28, Serial0/0/0
                   [120/1] via 209.165.200.234, 00:00:18, Serial0/0/1
     209.165.200.0/30 is subnetted, 2 subnets
C       209.165.200.232 is directly connected, Serial0/0/1
C       209.165.200.228 is directly connected, Serial0/0/0
     10.0.0.0/16 is subnetted, 1 subnets
C       10.1.0.0 is directly connected, FastEthernet0/0
S    192.168.0.0/16 is directly connected, Null0
```

- You still see the **summarized 172.30.0.0/16 route** with the same **two equal-cost paths**.

Configuring the RIP Protocol

Disable Auto Summarization

```
R2(config)# router rip
R2(config-router)# no auto-summary

R3(config)# router rip
R3(config-router)# no auto-summary

R1(config)# router rip
R1(config-router)# no auto-summary
```

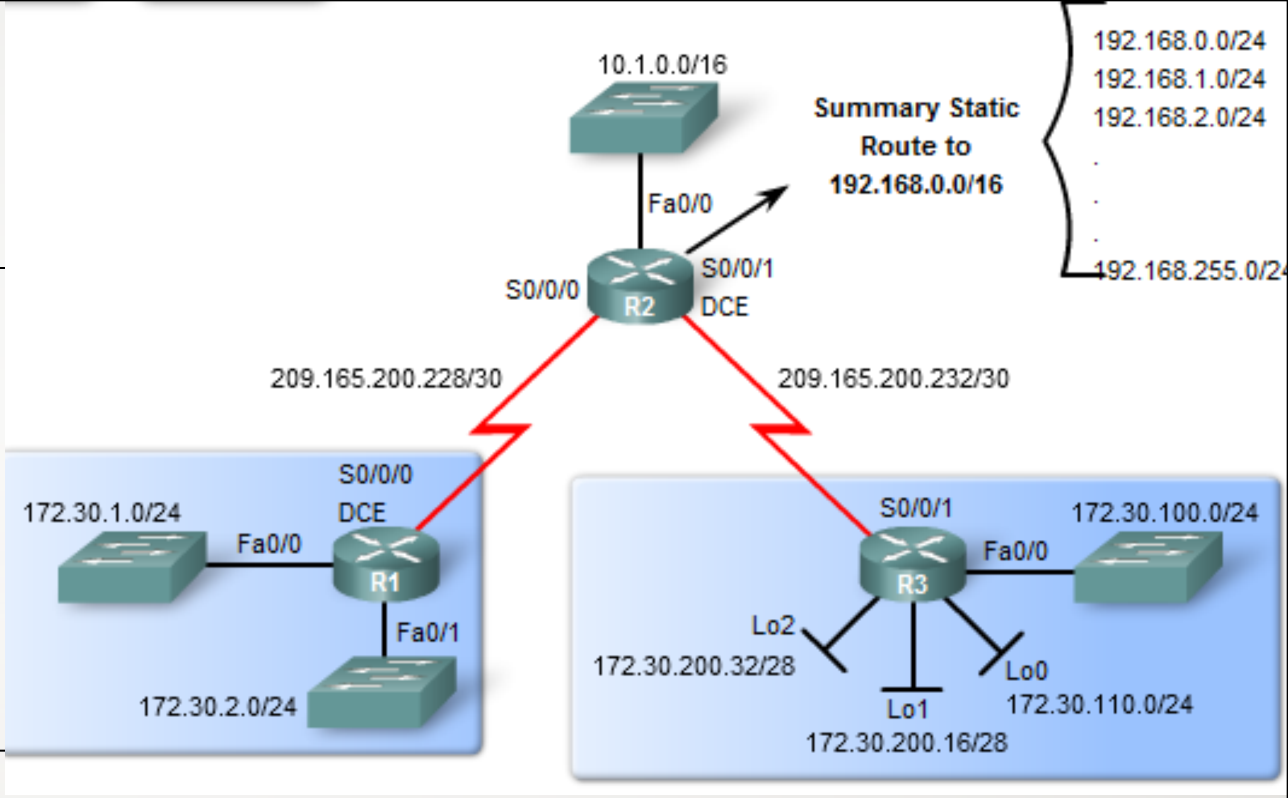

- RIPv2 automatically summarizes networks at major network boundaries.
- Use the **no auto-summary** router configuration mode command to disable auto summarization.
- Use the **show ip protocols** command to verify that auto summarization is off.

Configure Passive Interfaces

- RIP updates:
  - Are forwarded out all RIP-enabled interfaces by default.
  - Only need to be sent out interfaces that are connected to other RIP-enabled routers.
- Sending RIP update to LAN affects the network in three ways:
  - Inefficient Use of Bandwidth
  - Inefficient Use of Resources
  - Increased Security Risk
- Use the passive-interface router configuration command to stop routing updates out the interface. Still allows that network to be advertised to other routers.

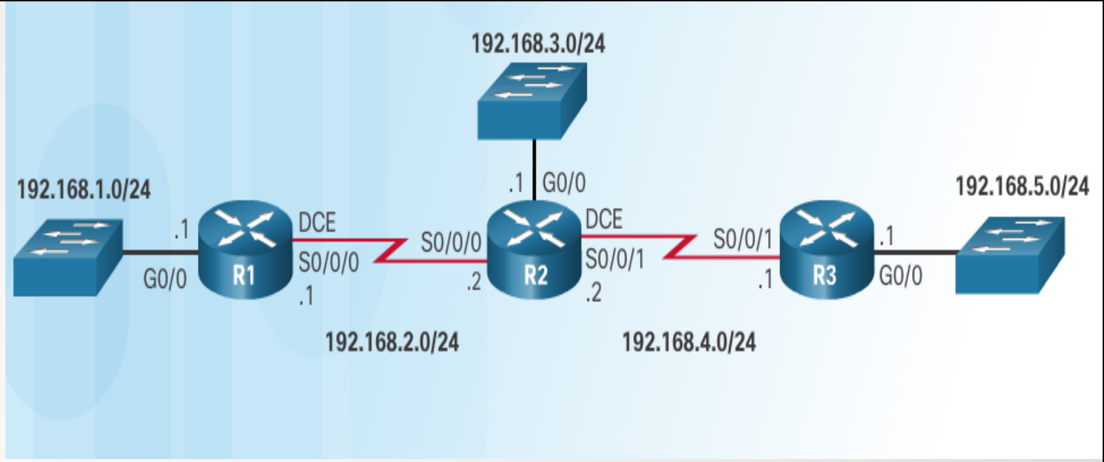


```
R1(config)# router rip
R1(config-router)# passive-interface g0/0
R1(config-router)# end
R1#
R1# show ip protocols | begin Default
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-
chain
    Serial0/0/0           2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    192.168.1.0
    192.168.2.0
  Passive Interface(s):
    GigabitEthernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.2.2          120      00:00:06
  Distance: (default is 120)

R1#
```

Enhanced Interior Gateway Routing Protocol

EIGRP

Overview of EIGRP

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

- **Enhanced IGRP** is a Cisco-proprietary routing protocol released in 1992.
  - EIGRP was created as a classless version of IGRP.
- EIGRP acts like a link-state routing protocol, but it's still a distance vector routing protocol.
- In 2013, Cisco released a basic functionality of EIGRP as an open standard to the IETF as an *informational RFC*.
  - Cisco will continue to maintain control of EIGRP.

Main Components of Routing Protocols

Data Structures

- Routing protocols create and maintain tables (databases) in RAM for its operations.

Routing Protocol Messages

- Routing protocols use messages to learn and maintain accurate information about the network.
- Specifically, messages are used to discover neighboring routers, exchange routing information, and other tasks.

Algorithm

- Routing protocols use algorithms to determine the best path to various destinations.

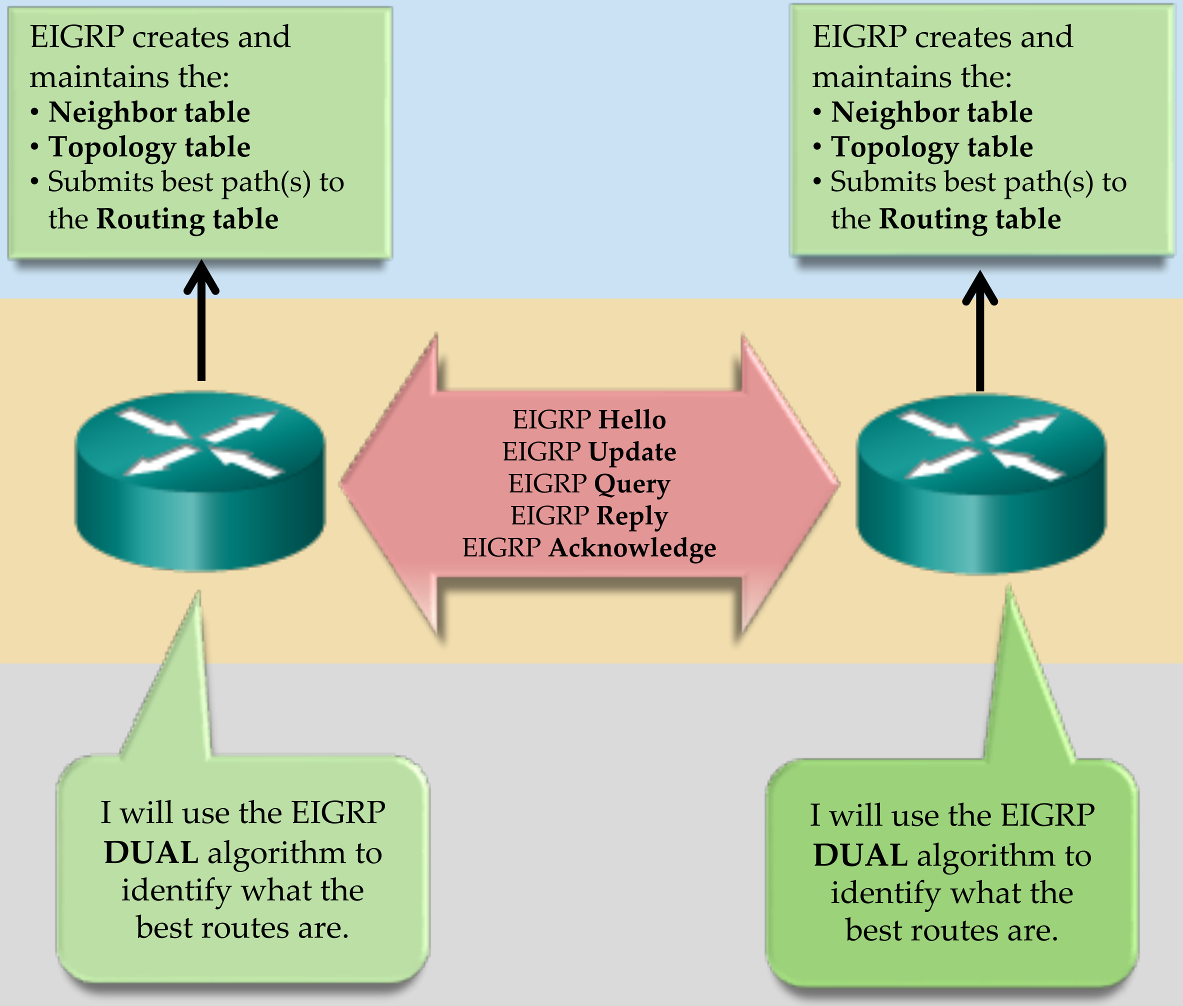

Features of EIGRP

- Administrative distance is 90
- Uses Diffusing Update Algorithm (DUAL)
- Does not age out routable entries
- Does not use periodic updates
- Maintains a topology table, separate from the routing table
  - Includes best path
  - Loop-free backup paths, if any
- Faster convergence
- Equal and unequal cost load balancing
- Partial and bounded updates:
  - **Partial**: Update only includes information about the route changes
  - **Bounded**: Updates are sent only to those affected routers.
- Used Protocol Dependent Modules (PDMs) to support different Layer 3 protocols

EIGRP Tables and Protocol-Dependent Modules

- EIGRP maintains a individual tables for each routed protocol.
- **Neighbor table:** Contains list of directly connected routers (Show ip eigrp Neighbor)
- **Topology Table :** list of all best routes learned from each neighbor (show ip eigrp topology)
- **Routing Table:** The best route to destination (show ip route)

EIGRP uses protocol-dependent modules (PDMs) to provide support for IPv4, IPv6 and legacy protocols IPX and AppleTalk.

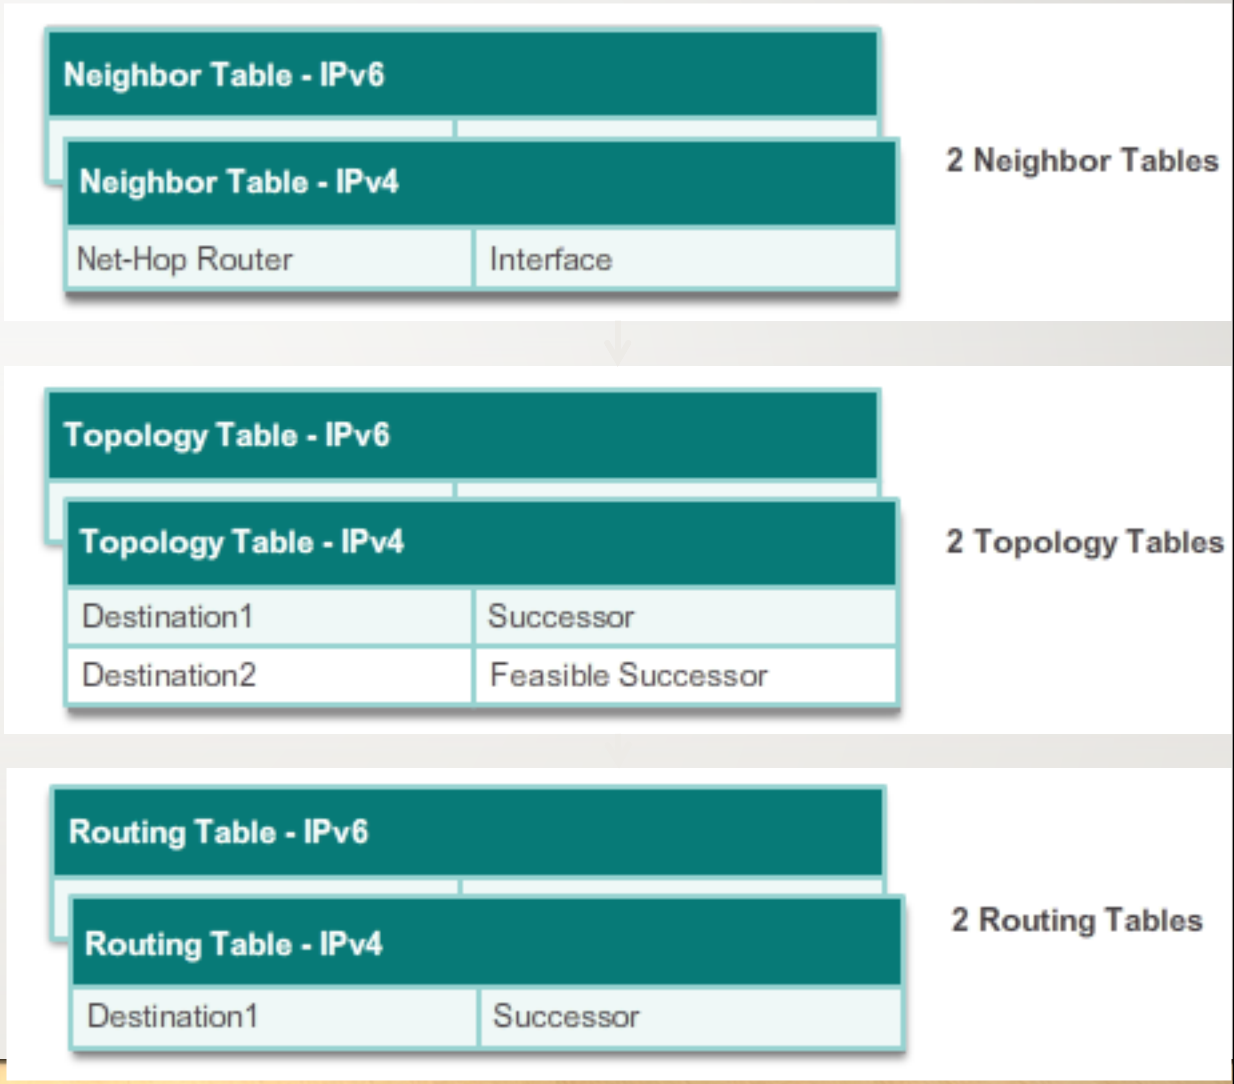

EIGRP RTP

- **Reliable Transport Protocol (RTP)**
- RTP is the EIGRP Transport layer protocol used for the delivery and reception of EIGRP packets.
  - However, not all RTP packets are sent reliably.
  - Why RTP?
  - Legacy protocols didn't use TCP/IP.

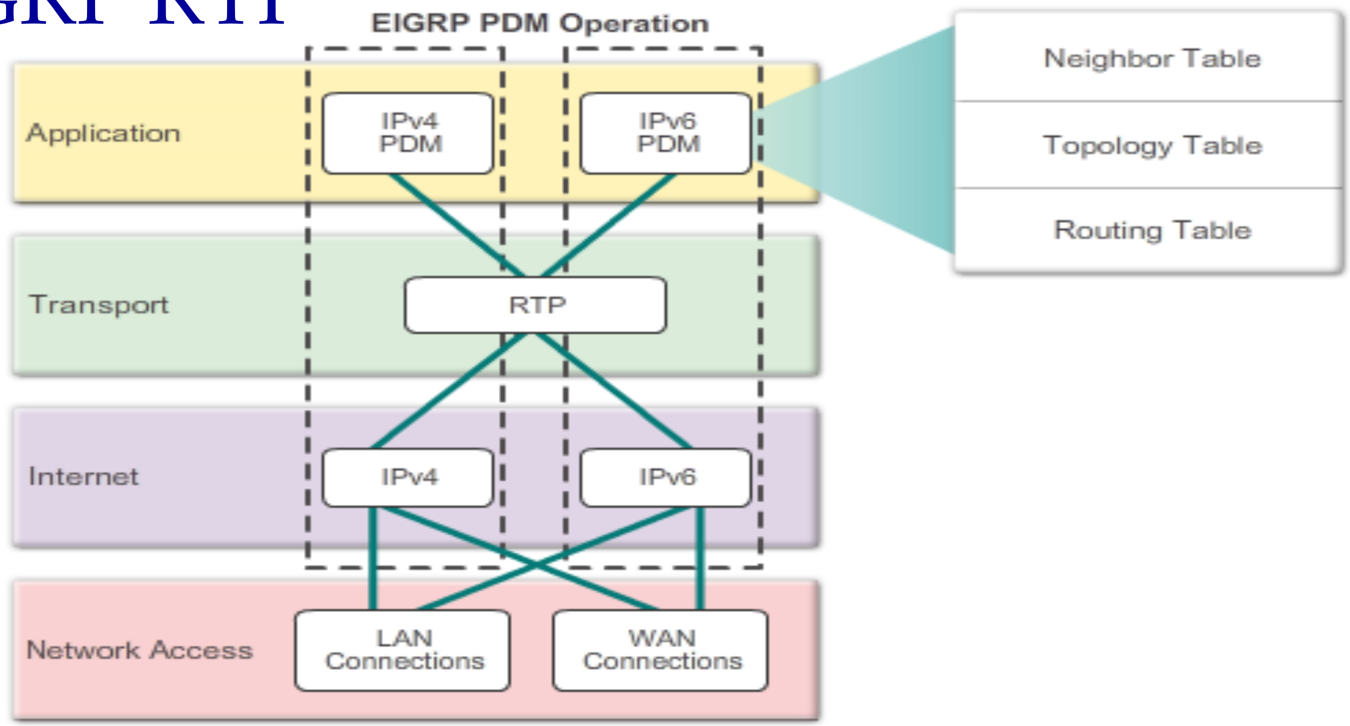

Reliable packets require explicit acknowledgement from destination

Update, Query, Reply

Unreliable packets do not require acknowledgement from destination

Hello, ACK

RTP can send EIGRP packets as unicast or multicast.

IPv4 EIGRP multicast address **224.0.0.10**.

EIGRP Packets

EIGRP Packets

- IP EIGRP relies on 5 types of packets to maintain its various tables and establish complex relationships with neighbor routers.

EIGRP Hello Packets

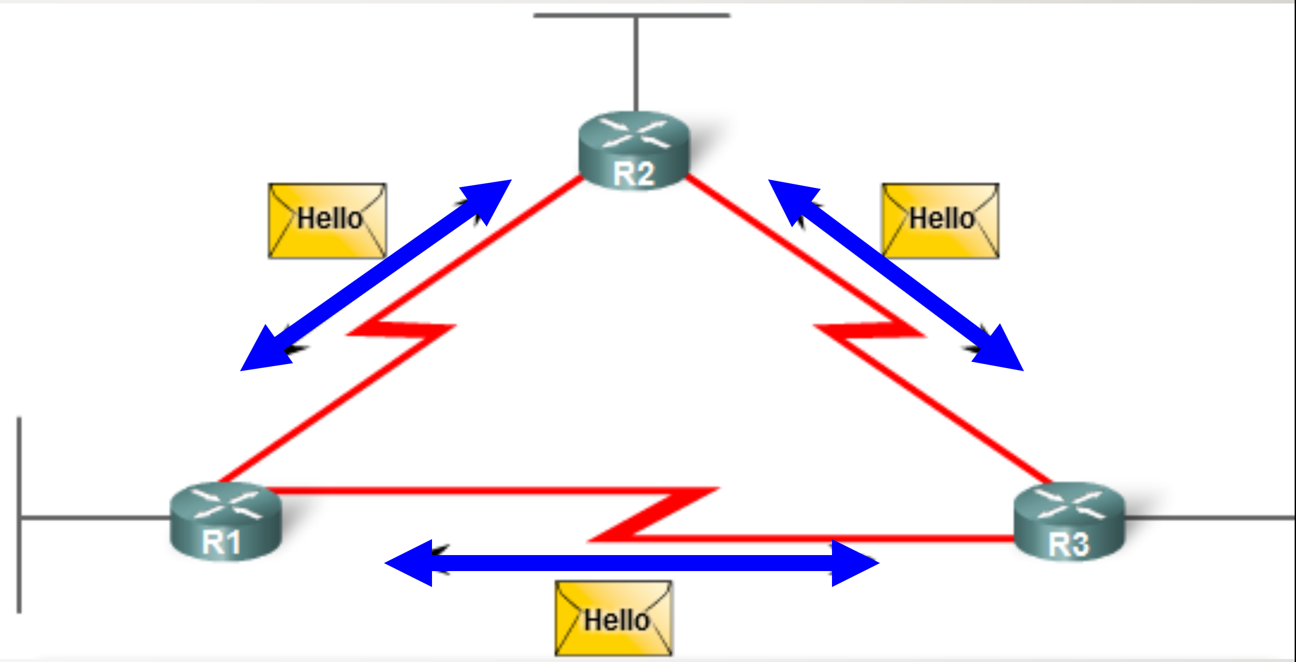

- **EIGRP routers** discover neighbors and establish adjacencies with neighbor routers using the **hello packet**.
- Multicasted to:
  - IPv4: **224.0.0.10**
  - IPv6: **FF02::A**
- Hello packets are always sent unreliably.
  - Therefore Hello packets do not require acknowledgment.

Hello Packets

Bandwidth	Example Link	Default Hello Interval	Default Hold Time
1.544 Mb/s	Multipoint Frame Relay	60 seconds	180 seconds
Greater than 1.544 Mb/s	T1, Ethernet	5 seconds	15 seconds

- Hello packets are sent on a regular interval.
  - **Router assumes that as long as it is receiving Hello packets from a neighbor, the neighbor and its routes remain viable.**
- The interval depends on the interface's bandwidth.
  - **Low Bandwidth = 60 seconds**
  - **High bandwidth = 5 seconds**
- **Hold time** - maximum time the router should wait to receive the next hello before declaring that neighbor as unreachable.
- Default hold time - **3 times the hello interval**
- If the hold time **expires**:
  - EIGRP declares the route as down
  - DUAL searches for a new path in the topology table ***or*** by sending out queries.

EIGRP Packet Types – Update and Acknowledgement Packet

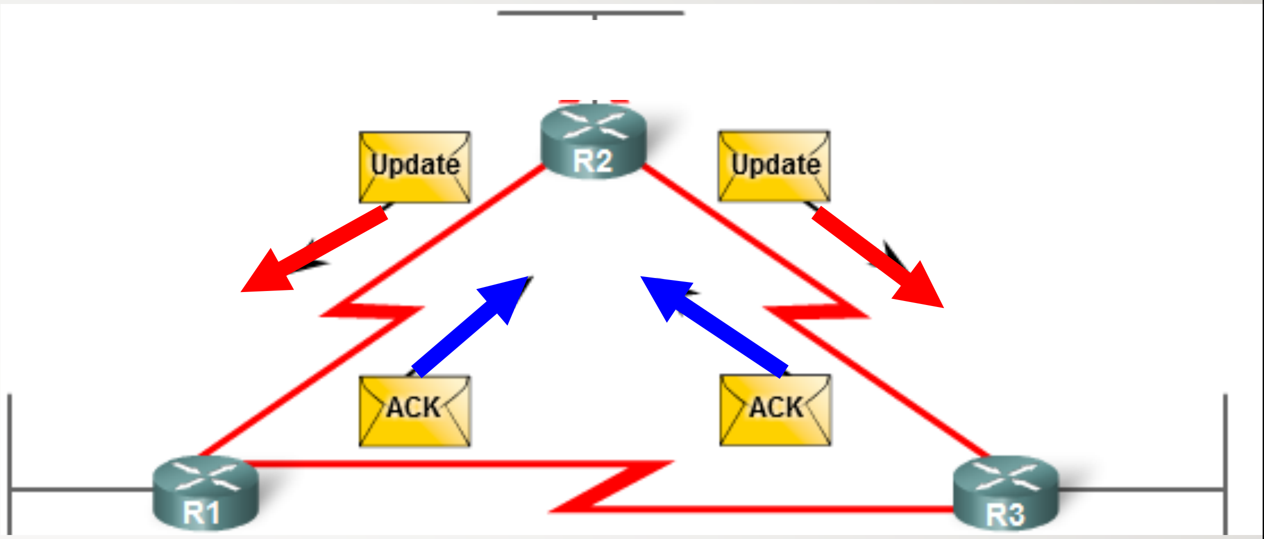

EIGRP uses triggered updates

- **Update Packets**
  - Contains only the routing information needed (a change occurs)
  - Sent only to those routers that require it.
  - Uses reliable delivery.
- **Acknowledgment (ACK) Packets**
  - Sent when reliable delivery is used (update, query, and reply packets).

EIGRP Packet Types – Query and Reply Packets

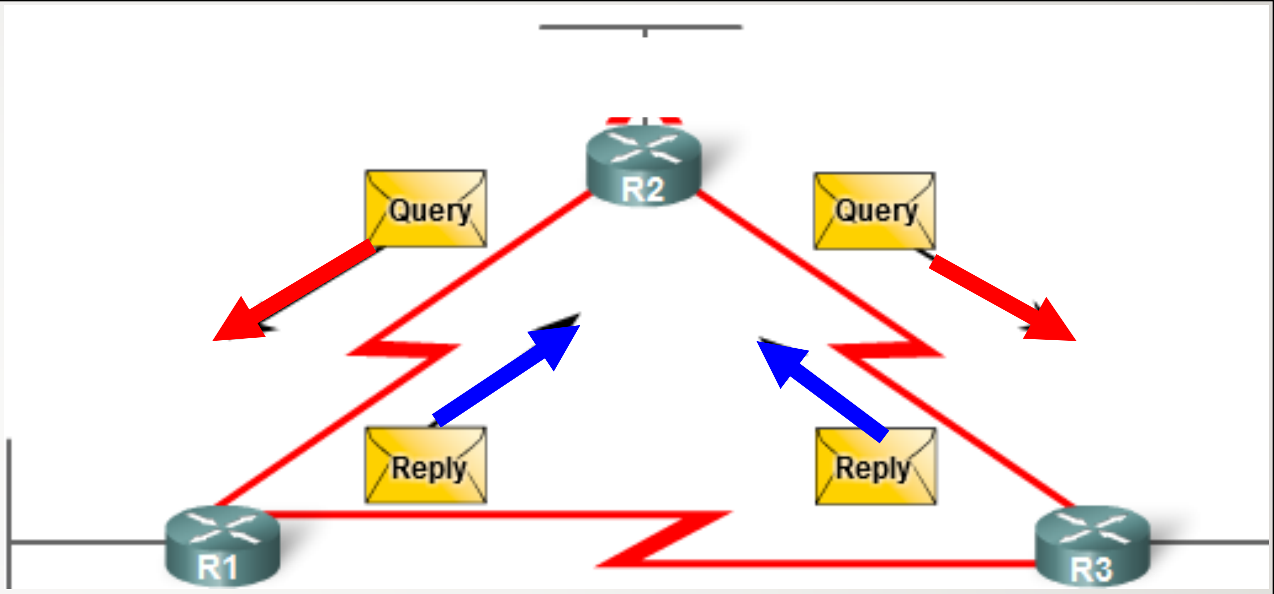

Why Query? Another router could be attached to the same LAN.

- Used by **DUAL** when searching for networks and other tasks.
- **Queries** and **replies** use reliable delivery.
- Queries can use multicast or unicast, whereas Replies are always sent as unicast.

EIGRP Packet Header Format

Data Link Frame Header	IP Packet Header	EIGRP Packet Header	Type/Length/Values Types

Data Link Frame

MAC Source Address = Address of sending interface

MAC Destination Address = Multicast: 01-00-5E-00-00-0A

IP Packet

IP Source Address = Address of sending interface

IP Destination Address = Multicast: 224.0.0.10

Protocol field = 88 for EIGRP

EIGRP Packet Header

Opcode for EIGRP packet type

AS Number

Operation of EIGRP

Neighbor Discovery

- The goal of any dynamic routing protocol is to:
  - Learn about remote networks
  - Reach convergence in the routing domain.
- To establish and maintain EIGRP neighbor adjacencies, EIGRP routers:
  - Use the Hello packet
    - Metric parameters must match (later)
    - AS number must match

Initial Route Discovery

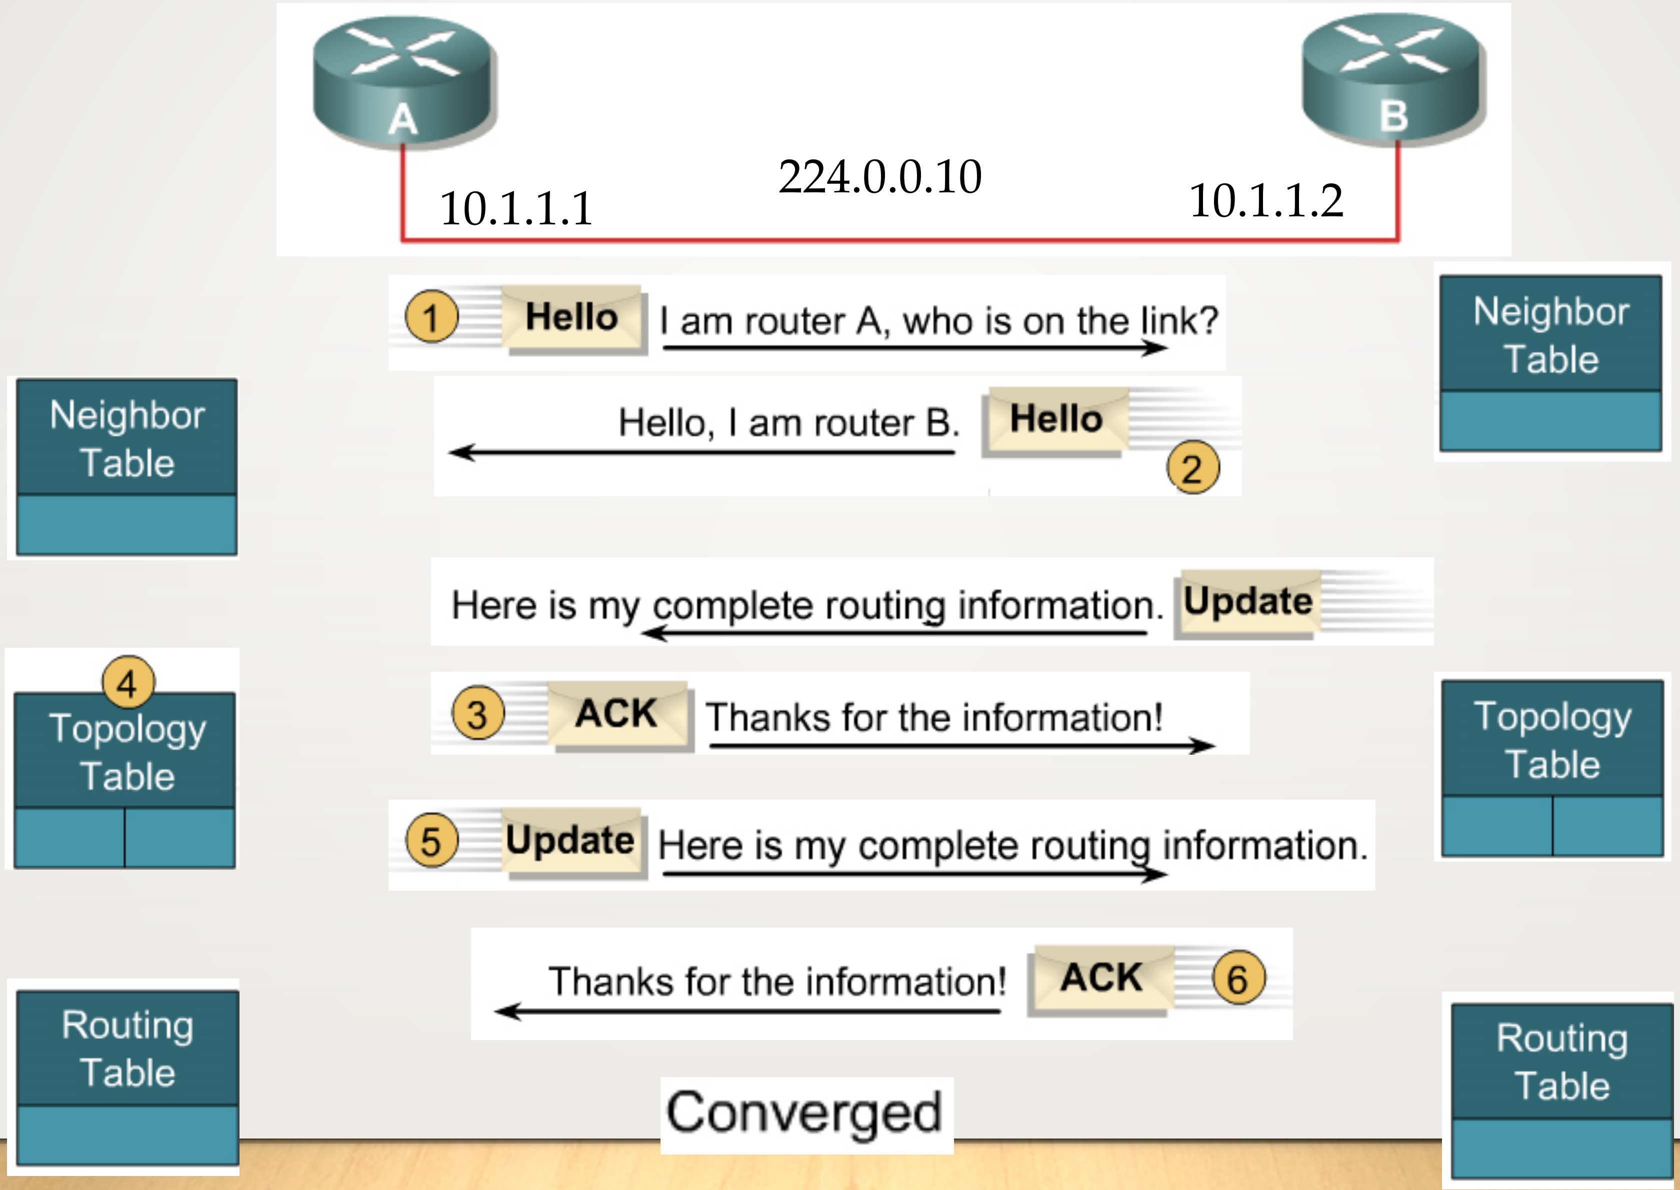

EIGRP Metrics

EIGRP Composite Metric and the K Values

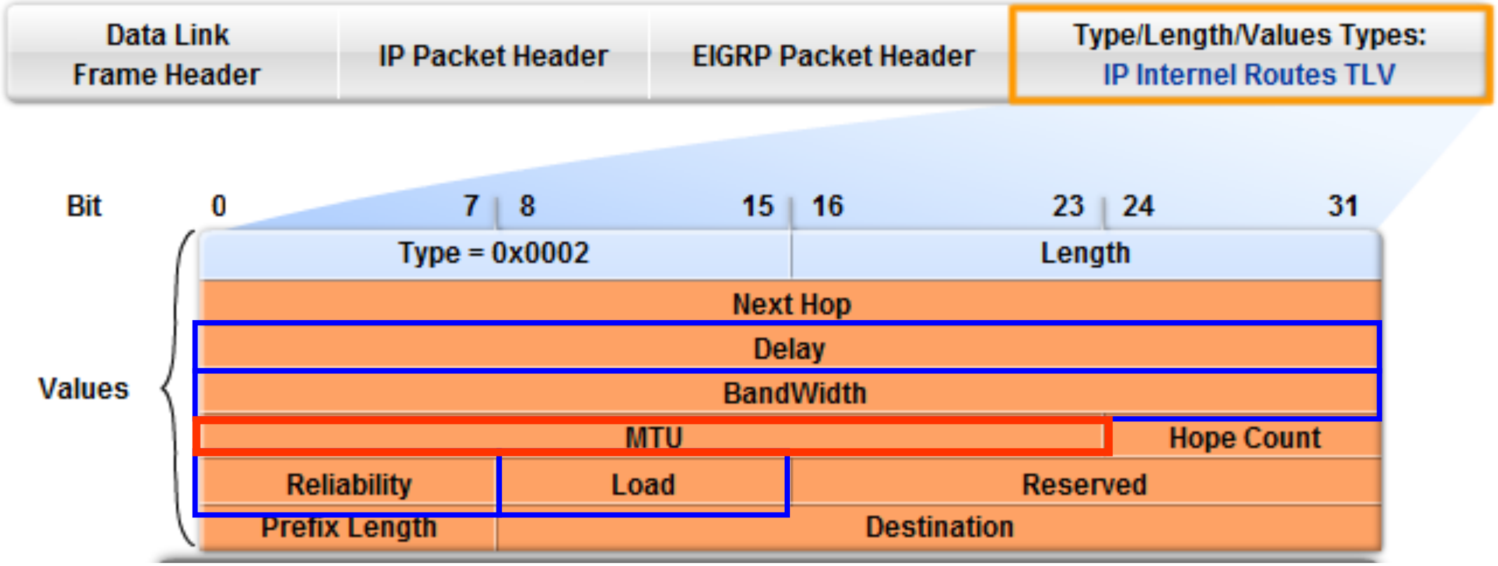

- EIGRP uses the following values in its composite metric to calculate the preferred path to a network:
  - **Bandwidth:** The lowest bandwidth between source and destination.
  - **Delay:** The cumulative interface delay along the path
  - **Reliability:** Worst reliability between source and destination, based on keepalives.
  - **Load**: Worst load on a link between source and destination, based on the packet rate and the configured bandwidth of the interface.
- **Note**: Although **MTU** is included in the routing table updates, it is not a routing metric used by EIGRP or IGRP.

The 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)

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
```

- By default:
  - K1 and K3 are set to 1,
  - K2, K4, and K5 are set to 0.
- The result is that only the bandwidth and delay values are used in the computation of the default composite metric.

EIGRP Operation

DUAL and the Topology Table

- EIGRP uses the Diffusing Update Algorithm (DUAL) to provide the best and backup loop-free paths.
- DUAL uses several terms, which are discussed in more detail throughout this section:

Term	Description
Successor	• Is a neighboring router that is used for packet forwarding and is the least-cost route to the destination network. • The IP address of a successor is shown in a routing table entry right after the word "via".
Feasible Successors (FS)	• These are the "Backup paths" that are a loop-free. • Must comply to a feasibility condition.
Reported Distance (RD)	• Also called "advertised distance", this is the reported metric from the neighbor advertising the route. • If the RD metric is less than the FD, then the next-hop router is downstream and there is no loop.
Feasible Distance (FD)	• This is the actual metric of a route from the current router. • Is the lowest calculated metric to reach the destination network. • FD is the metric listed in the routing table entry as the second number inside the brackets.

EIGRP Operation

DUAL and the Topology Table

- Routing loops, even temporary ones, can be detrimental to network performance and EIGRP prevents routing loops with the DUAL algorithm.
  - The DUAL algorithm is used to obtain loop-freedom at every instance throughout a route computation.
- DUAL uses EIGRP metrics to select efficient, loop-free paths, and to identify the routes with the least-cost path to be inserted into the routing table.

DUAL and the Topology Table

- A successor is a neighboring router with the least-cost route to the destination network.
- FD is the lowest calculated metric to reach the destination network.
  - Also known as the "*metric*" for the route.
- Notice that EIGRP's best path for the 192.168.1.0/24 network is through router R3, and that the feasible distance is 3,012,096.

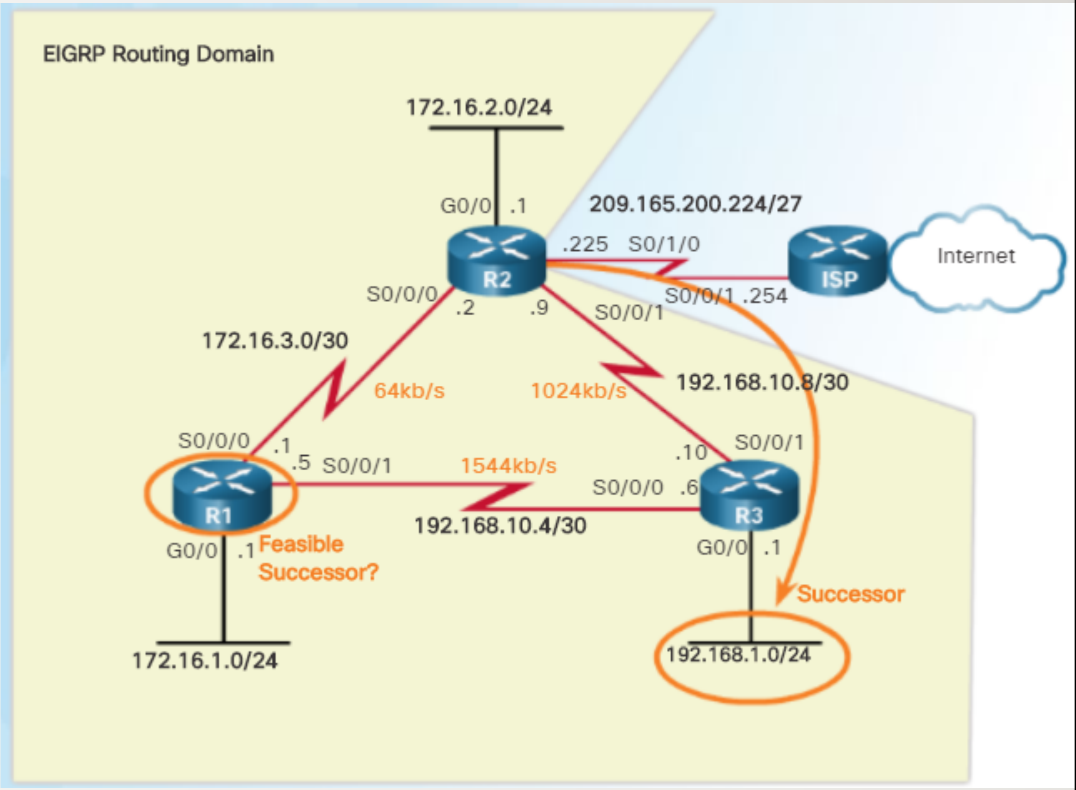

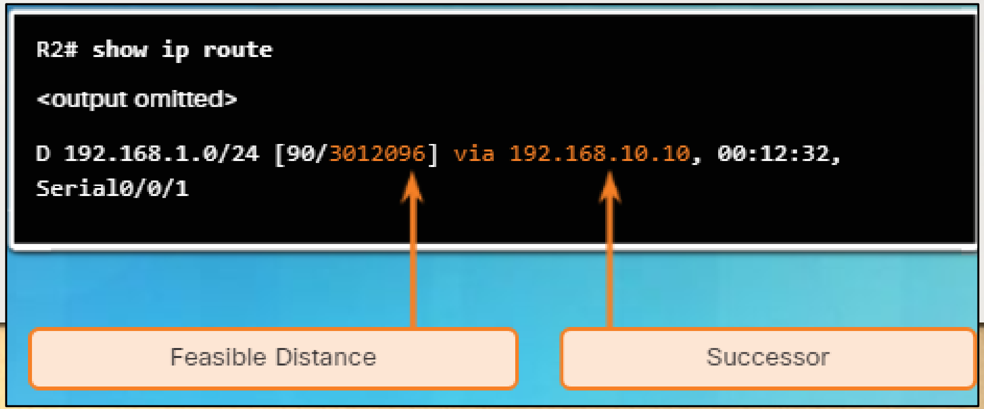

EIGRP Operation
DUAL and the Topology Table

- DUAL converges quickly because it can use backup paths known as Feasible Successors (FSs).

- A FS is a neighbor with a loop-free backup path to the same network as the successor.
    - A FS must satisfy the Feasibility Condition (FC).
    - **The FS is met when a neighbor's Reported Distance (RD) is less than the local router's feasible distance.**
    - If the reported distance is less, it represents a loop-free path.

- E.g., the RD of R1 (2,170,112) is less than R2's own FD (3,012,096) and therefore, R1 meets the FC and becomes the FS for R2 to the 192.168.1.0/24 network.

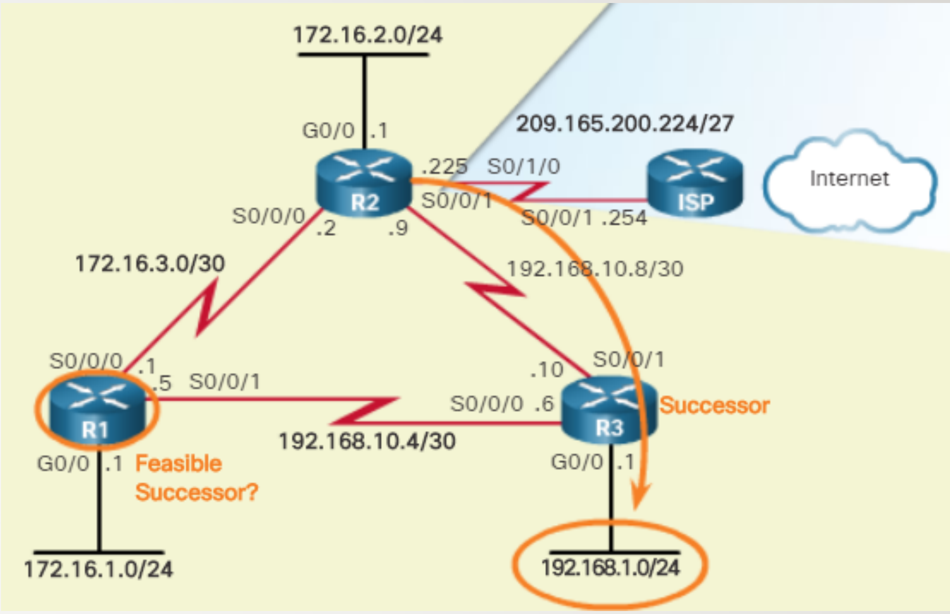

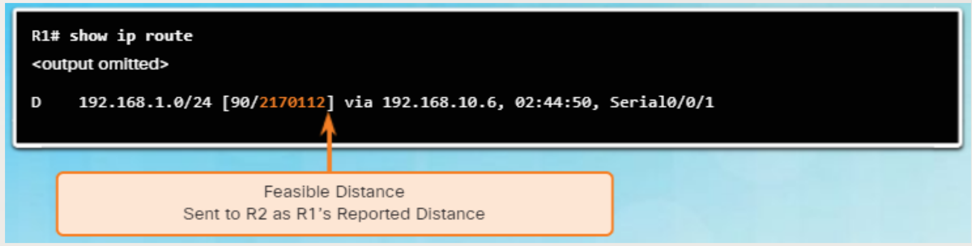

DUAL and the Topology Table

- Topology table stores the following information required by DUAL to calculate distances and vectors to destinations.
  - The **reported distance (RD)** that each neighbor advertises for each destination
  - The **feasible distance (FD)** that this router would use to reach the destination via that neighbor.

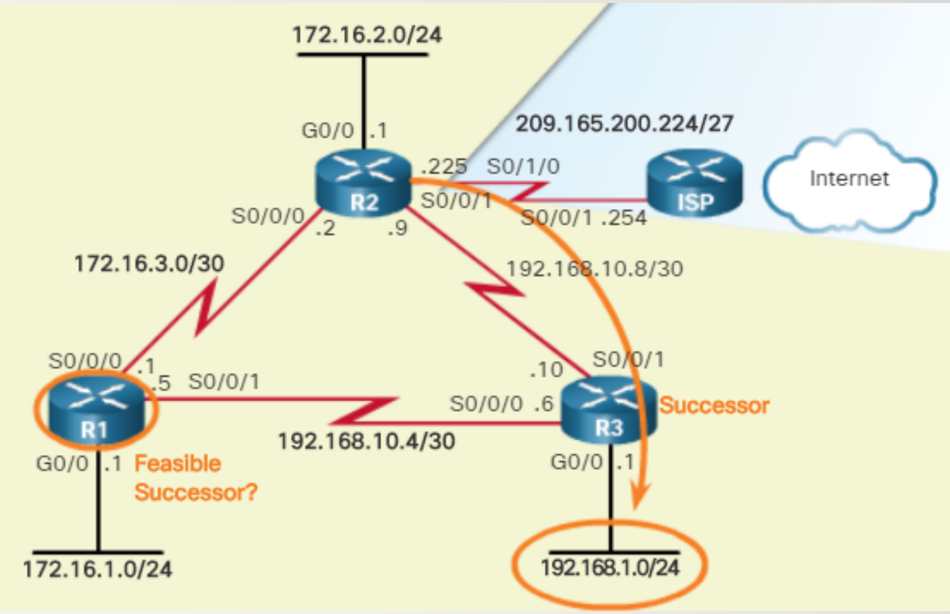


```
R2# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(2.2.2.2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.2.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/0
P 192.168.10.4/30, 1 successors, FD is 3523840
        via 192.168.10.10 (3523840/2169856), Serial0/0/1
        via 172.16.3.1 (41024000/2169856), Serial0/0/0
P 192.168.1.0/24, 1 successors, FD is 3012096
        via 192.168.10.10 (3012096/2816), Serial0/0/1
        via 172.16.3.1 (41024256/2170112), Serial0/0/0
P 172.16.3.0/30, 1 successors, FD is 40512000
        via Connected, Serial0/0/0
P 172.16.1.0/24, 1 successors, FD is 3524096
        via 192.168.10.10 (3524096/2170112), Serial0/0/1
        via 172.16.3.1 (40512256/2816), Serial0/0/0
P 192.168.10.8/30, 1 successors, FD is 3011840
        via Connected, Serial0/0/1

R2#
```

Process ID

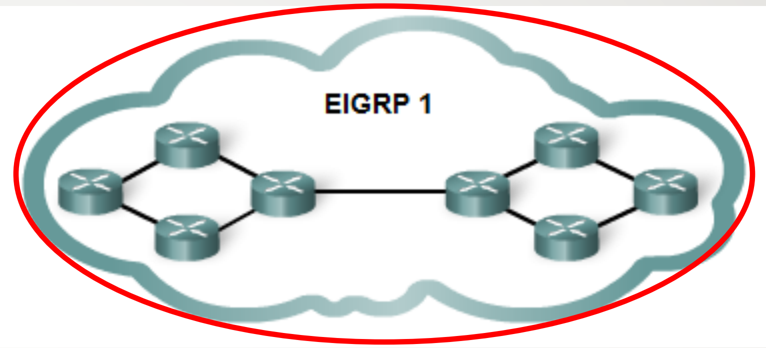


```
Router(config)# router eigrp autonomous-system
```

```
Router(config)# router eigrp 1
```

Must be same on all routers in EIGRP routing domain

- Both **EIGRP** and **OSPF** use a ***process ID*** to represent an instance of their respective routing protocol running on the router.
- **EIGRP** refers to "***autonomous-system***" number
  - Actually functions as a ***process ID***.
  - 1 and 65,535

End of Chapter 6