Chapter 6: EIGRP

CCNA Routing and Switching

Scaling Networks v6.0

Chapter 6 - Sections & Objectives

- 6.1 EIGRP Characteristics
  - Explain the features and characteristics of EIGRP.
    - Describe the basic features of EIGRP.
    - Describe the types of packets used to establish and maintain an EIGRP neighbor adjacency.
    - Describe the encapsulation of an EIGRP messages.

- 6.2 Implement EIGRP for IPv4
  - Implement EIGRP for IPv4 in a small to medium-sized business network.
    - Configure EIGRP for IPv4 in a small routed network.
    - Verify EIGRP for IPv4 operation in a small routed network.

Chapter 6 - Sections & Objectives (Cont.)

- 6.3 EIGRP Operation
  - Explain how EIGRP operates in a small to medium-sized business network.
    - Explain how EIGRP forms neighbor relationships.
    - Explain the metrics used by EIGRP.
    - Explain how DUAL operates and uses the topology table.
    - Describe events that trigger EIGRP updates.

- 6.4 Implement EIGRP for IPv6
  - Implement EIGRP for IPv6 in a small to medium-sized business network.
    - Compare characteristics and operation of EIGRP for IPv4 to EIGRP for IPv6.
    - Configure EIGRP for IPv6 in a small routed network.
    - Verify EIGRP for IPv6 implementation in a small routed network.

6.1 EIGRP Characteristics

EIGRP Basic Features

- Enhanced IGRP is a Cisco-proprietary distance-vector routing protocol released in 1992.
  - EIGRP was created as a classless version of IGRP.
  - Ideal choice for large, multiprotocol networks built primarily on Cisco routers.

EIGRP Feature	Description
Diffusing Update Algorithm (DUAL)	• EIGRP uses DUAL as its routing algorithm. • DUAL guarantees loop-free and backup paths throughout the routing domain.
Establishing Neighbor Adjacencies	• EIGRP establishes relationships with directly connected EIGRP routers. • Adjacencies are used to track the status of these neighbors.
Reliable Transport Protocol	• EIGRP RTP provides delivery of EIGRP packets to neighbors. • RTP and neighbor adjacencies are used by DUAL.
Partial and Bounded updates	• Instead of periodic updates, EIGRP sends partial triggered updates when a path or metric changes. • Only those routers that require the information are updated minimizing bandwidth use.
Equal and Unequal Cost Load Balancing	• EIGRP supports equal cost load balancing and unequal cost load balancing, which allows administrators to better distribute traffic flow in their networks.

EIGRP Basic Features

- EIGRP uses protocol-dependent modules (PDMs) to support different protocols such as IPv4, IPv6, and legacy protocols IPX and AppleTalk.
- PDMs are responsible for:
  - Maintaining EIGRP neighbor and topology tables
  - Computing the metric using DUAL
  - Interfacing DUAL and routing table
  - Implementing filtering and access lists
  - Performing redistribution with other routing protocols

EIGRP maintains individual tables for each routed protocol.

EIGRP Basic Features

- RTP is the EIGRP Transport layer protocol used for the delivery and reception of EIGRP packets.

- Not all RTP packets are sent reliably.
  - Reliable packets require explicit acknowledgement from destination
  - Update, Query, Reply
  - Unreliable packets do not require acknowledgement from destination
  - Hello, ACK

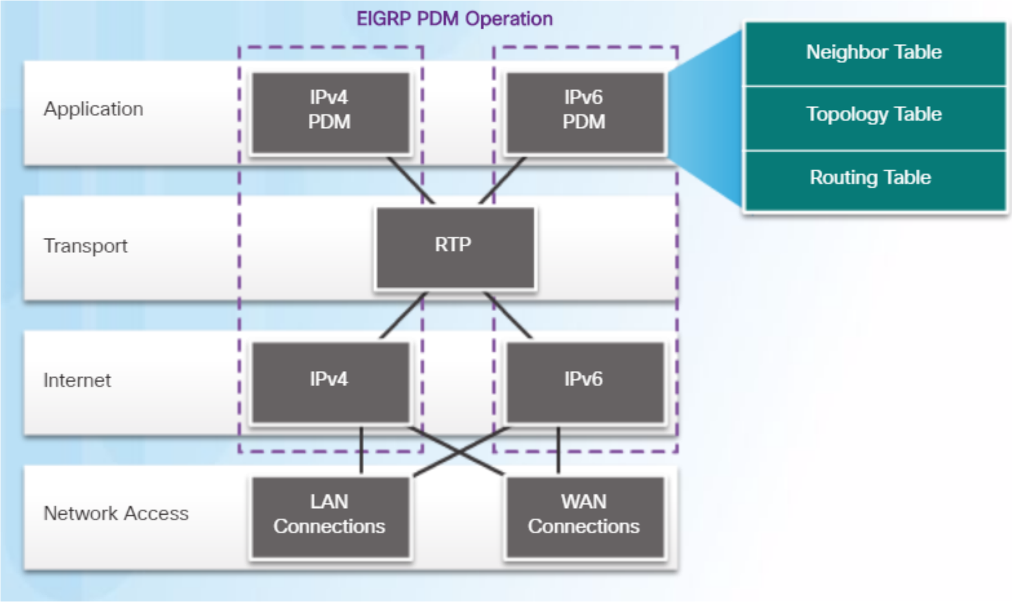

EIGRP Basic Features

- EIGRP supports authentication and is recommended.
  - EIGRP authentication ensures that routers only accept routing information from other routers that have been configured with the same password or authentication information.

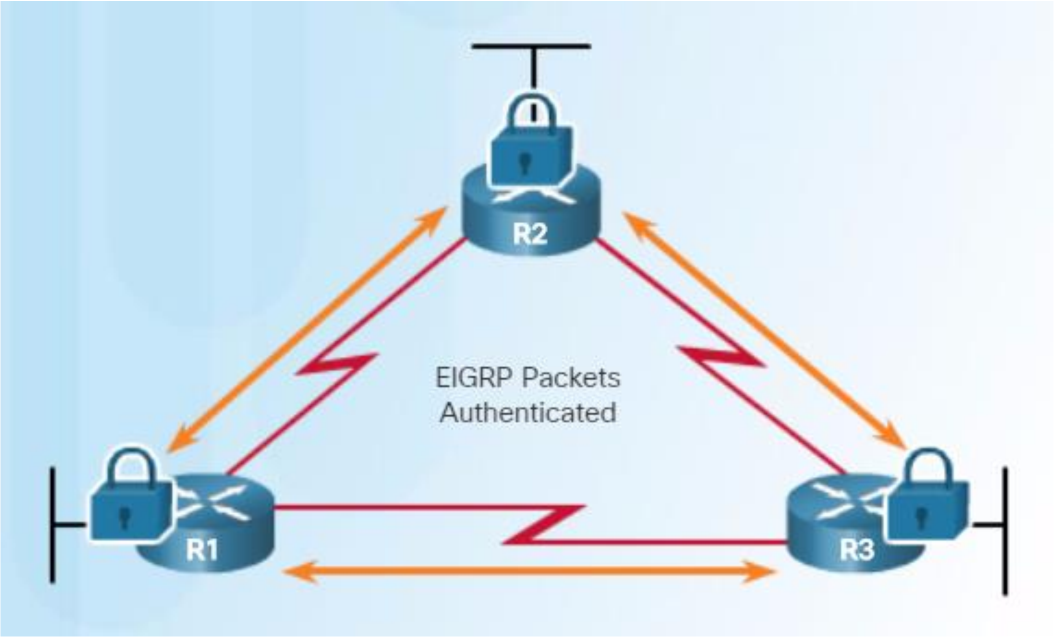

- **Note**:
  - Authentication does not encrypt the EIGRP routing updates.

EIGRP Packet Types

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

Packet Type	Description
Hello	• Used to discover other EIGRP routers in the network. • Sent unreliably to multicast address 224.0.0.5 (or 224.0.0.6).
Acknowledgement	• Used to acknowledge the receipt of any EIGRP packet. • Sent unreliably as unicasts.
Update	• Convey routing information to known destinations. • Sent reliably as unicasts or multicasts.
Query	• Used to get specific information from a neighbor router. • Sent reliably as unicasts or multicasts.
Reply	• Used to respond to a query. • Sent reliably as unicasts.

EIGRP Packet Types

- Hello packets are used to discover & form adjacencies with neighbors.
  - On hearing Hellos, a router creates a neighbor table and the continued receipt of Hellos maintains the table.
- Hello packets are always sent unreliably.
  - Therefore Hello packets do not require acknowledgment.

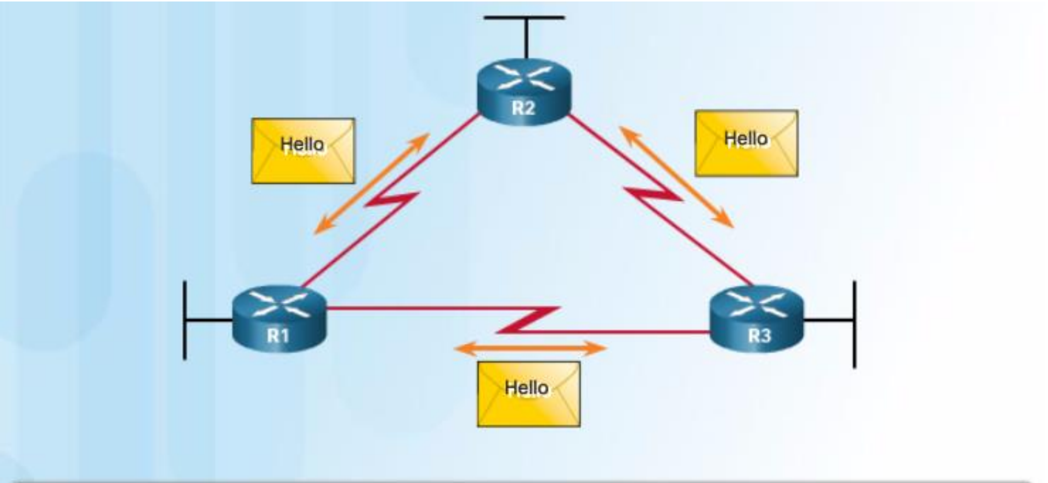

Bandwidth	Example Link	Default Hello Interval	Default Hold Time
1.544 Mb/s	Multipoint Frame Relay	60 seconds	180 seconds
Greater than 1.544 Mb/s	T1, Ethernet	5 seconds	15 seconds

EIGRP uses multicast and unicast rather than broadcast.
- As a result, end stations are unaffected by routing updates or queries.
- The EIGRP multicast IPv4 address is **224.0.0.10**
- The EIGRP multicast IPv6 address is **FF02::A**.

EIGRP Packet Types

- EIGRP Update packets are used to propagate routing information.
  - Sent to initially exchange topology information or topology change.
  - EIGRP updates only contain needed routing information and are unicast to routers that require it.
  - Update packets are sent reliably and therefore requires acknowledgements.

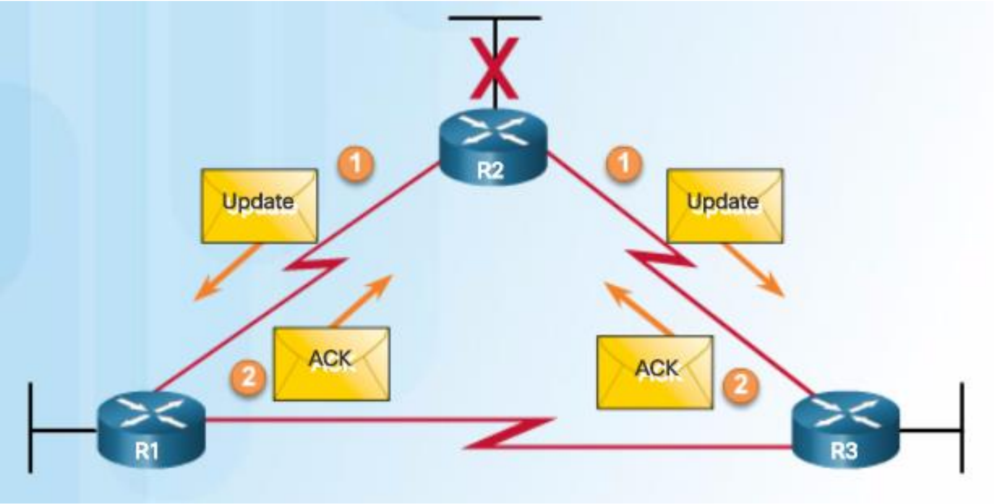

- Acknowledgements packets are “dataless” Hello packets used to indicate receipt of any EIGRP packet during a "reliable" (i.e., RTP) exchange.
  - Used to acknowledge the receipt of Update packets, Query packets, and Reply packets.

EIGRP Packet Types

- Query and reply packets are used by DUAL when searching for networks.
- They both use reliable delivery and therefore require acknowledgement.
- Queries can use multicast or unicast, whereas Replies are always sent as unicast.

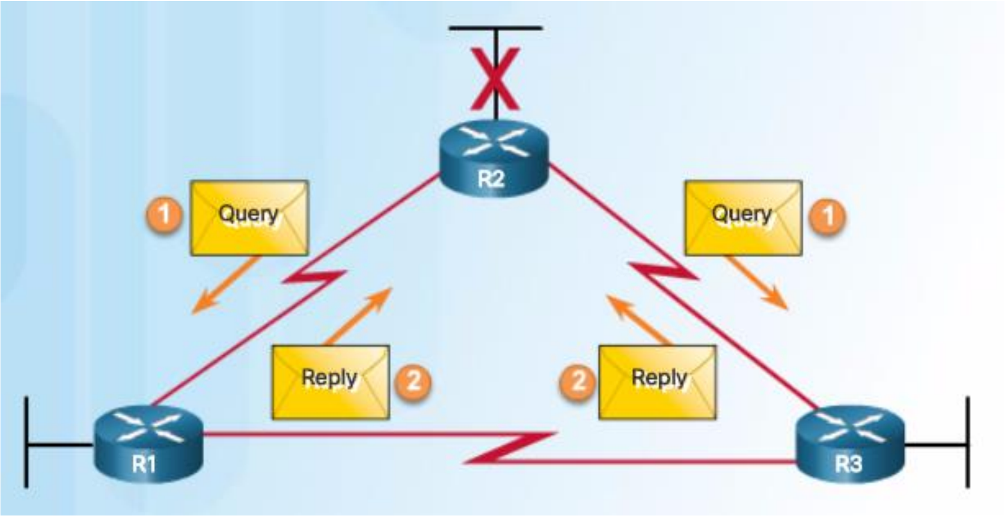

EIGRP Packet Types

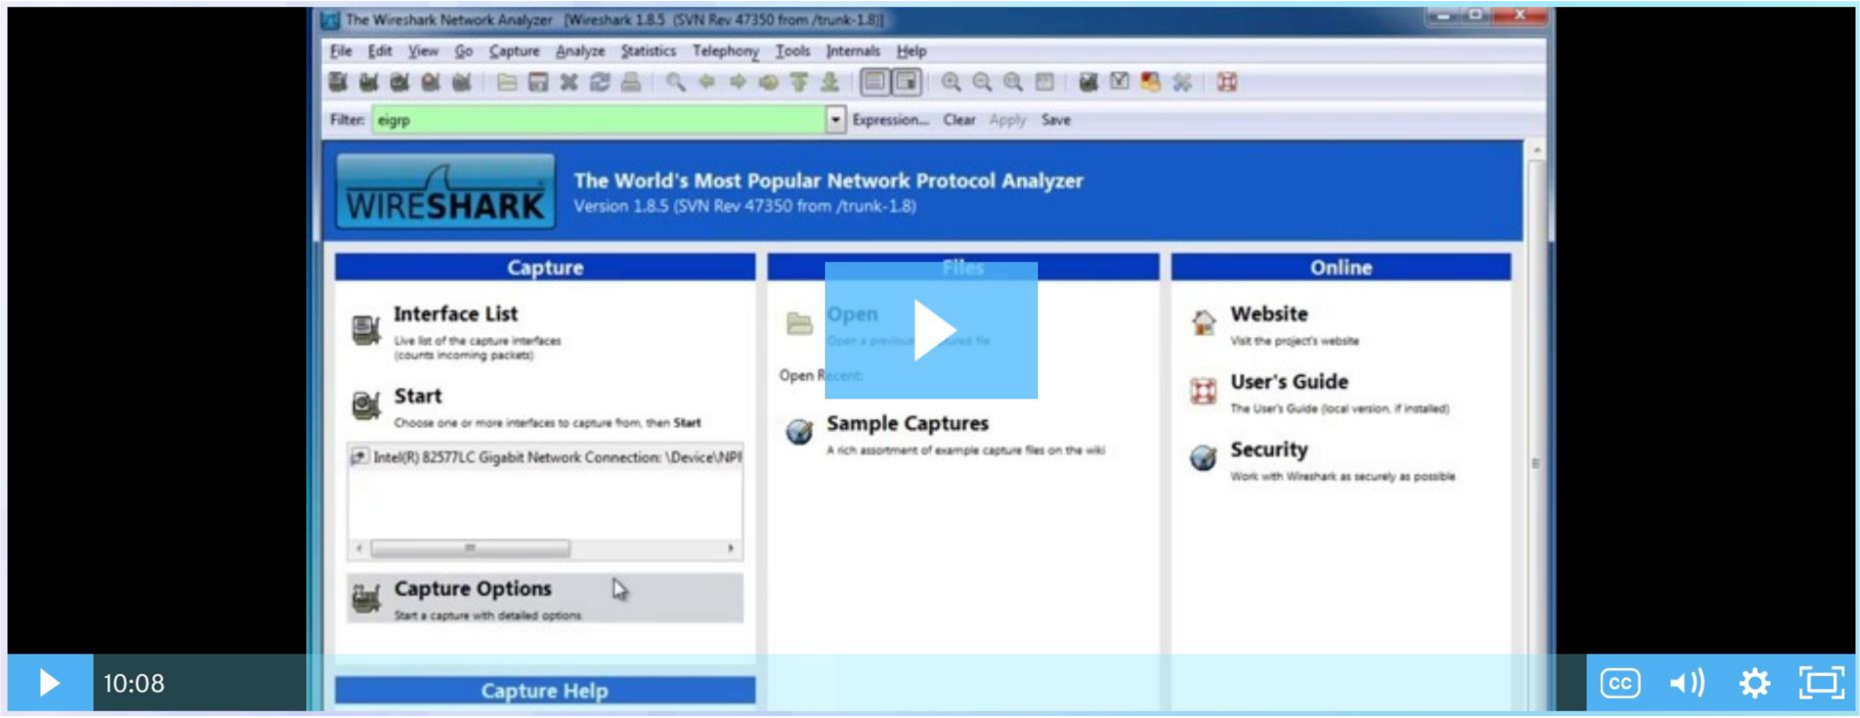

EIGRP Messages

- EIGRP frame contains destination multicast address 01-00-5E-00-00-0A.
- The IP packet header contains destination IP address 224.0.0.10 and identifies this packet as an EIGRP packet (protocol 88).
- The data portion of the EIGRP message includes:
  - **Packet header** - The EIGRP packet header identifies the type of EIGRP message.
  - **Type/Length/Value (TLV)** - The TLV field contains EIGRP parameters, IP internal and external routes.
- EIGRP for IPv6 is encapsulated using an IPv6 header with multicast address FF02::A and the next header field set to protocol 88.

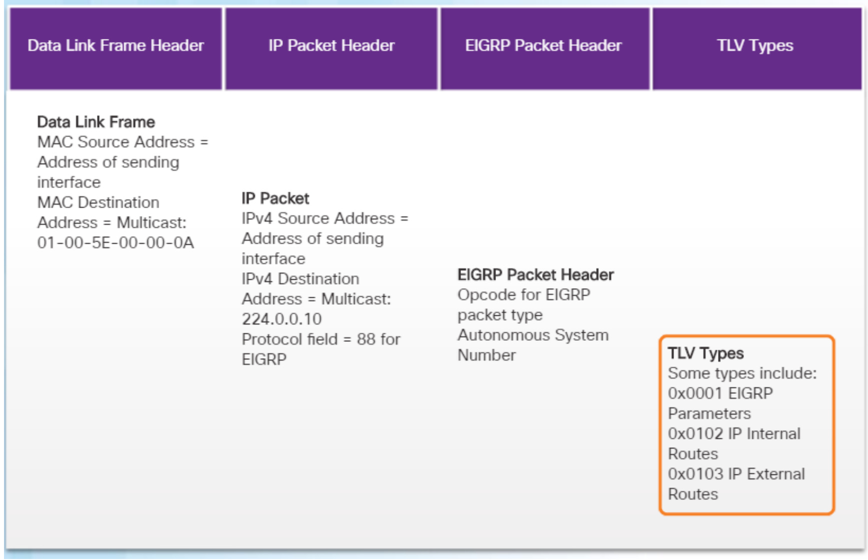

EIGRP Messages

- EIGRP messages include the header with an Opcode field that specifies the type of EIGRP packet (Hello, Ack, Update, Query, and Reply) and the AS number field.

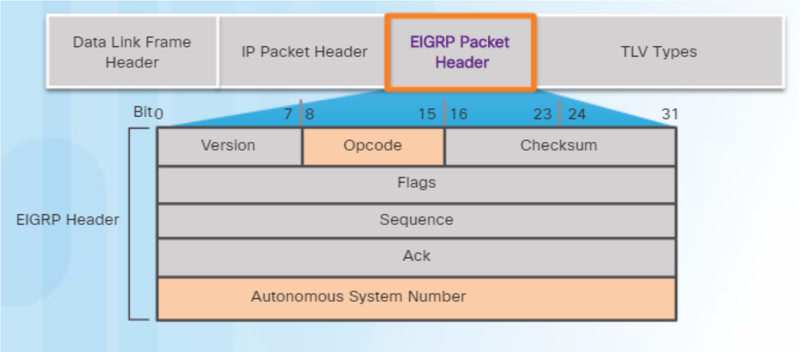

EIGRP TLV: EIGRP Parameters

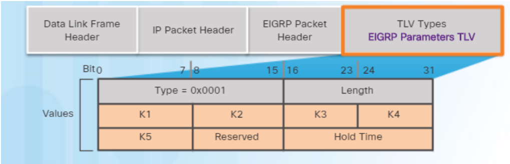

EIGRP TLV: Internal Routes

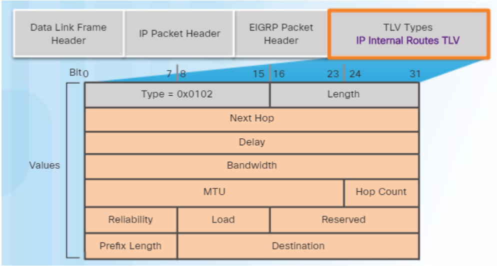

EIGRP TLV: External Routes

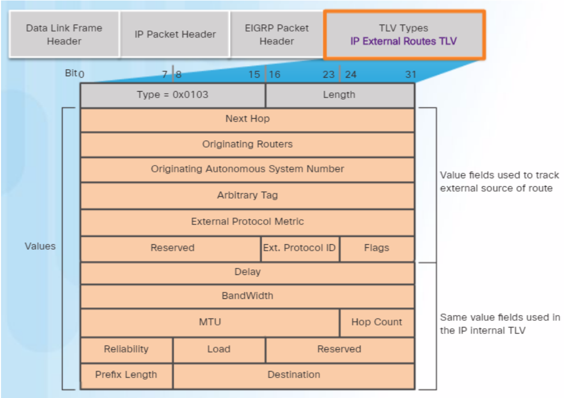

6.2 Implement EIGRP for IPv4

Configure EIGRP with IPv4

- The routers in the topology have a starting configuration that includes addresses on the interfaces. There is currently no static routing or dynamic routing configured on any of the routers.

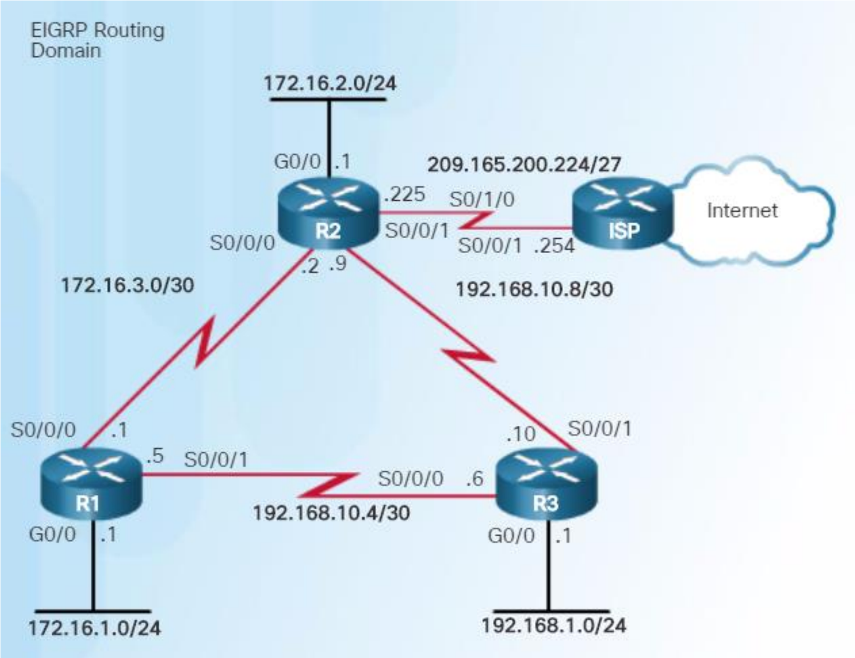

Configure EIGRP with IPv4

- An Autonomous System (AS) is a collection of networks under the control of a single authority (reference RFC 1930).
  - AS numbers are needed to exchange routes between AS.
  - AS numbers are managed by IANA and assigned by RIRs to ISPs, Internet Backbone providers, and institutions connecting to other institutions using AS numbers.

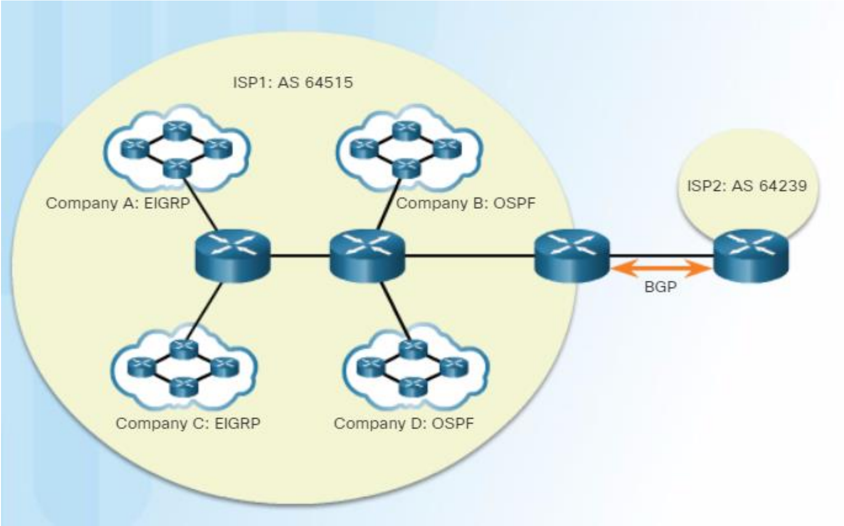

- AS numbers are usually 16-bit numbers, ranging from 0 to 65535.
  - Since 2007, AS numbers can now be 32 bits, therefore increasing the number of AS numbers to over 4 billion.

Configure EIGRP with IPv4

- To configure EIGRP, use the **router eigrp** *AS-#* command.
  - The *AS-#* functions as a process ID.
  - The AS number used for EIGRP configuration is only significant to the EIGRP routing domain.
  - All routers in the EIGRP routing domain must use the same AS number (process ID number).

- **Note:**
  - Do NOT configure multiple instances of EIGRP on the same router.

Configure EIGRP with IPv4

- The EIGRP router ID is used to uniquely identify each router in the EIGRP routing domain.
- Routers use the following three criteria to determine its router ID:

1. Use the address configured with the **eigrp router-id** *ipv4-address* router config command.
2. If the router ID is not configured, choose the highest IPv4 address of any of its loopback interfaces.
3. If no loopback interfaces are configured, choose the highest configured, choose the highest active IPv4 address of any of its physical interfaces.

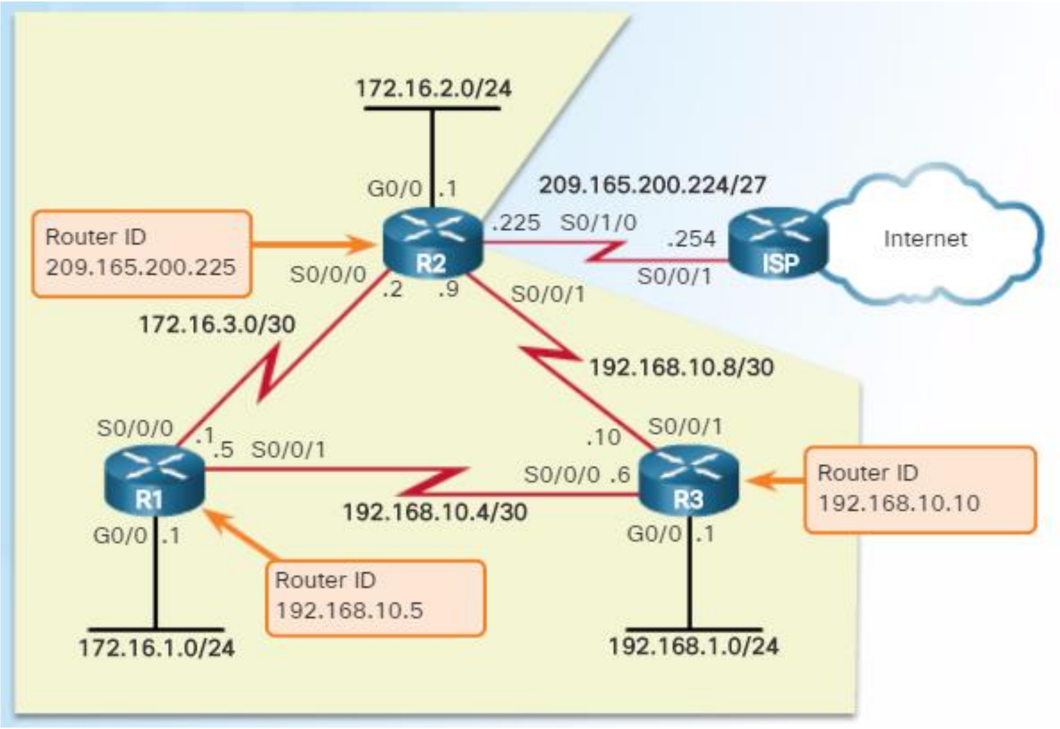

Configure EIGRP with IPv4

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

```
R2(config)# router eigrp 1
R2(config-router)# eigrp router-id 2.2.2.2
R2(config-router)#
```

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

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not
  set
  Incoming update filter list for all interfaces is not
  set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 1.1.1.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

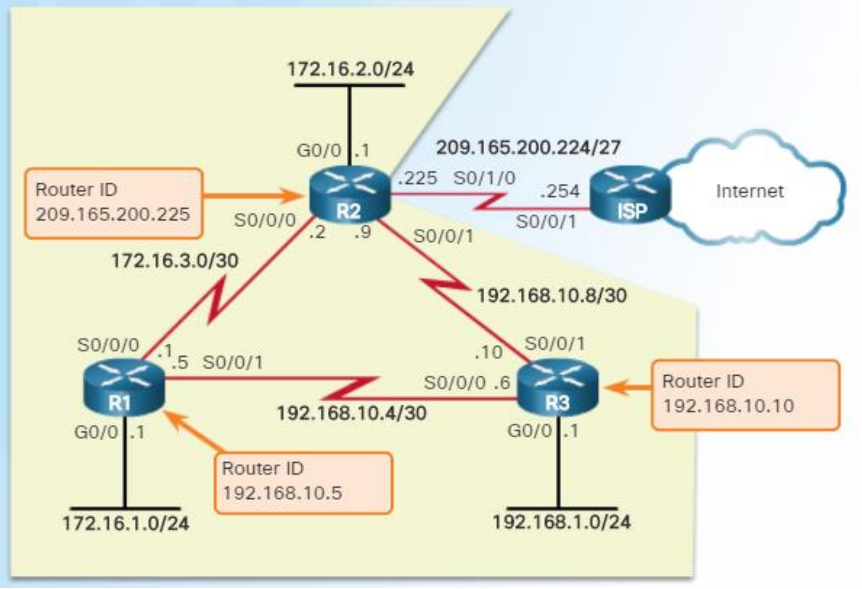

Configure EIGRP with IPv4

- Use the **network** *network-number* [*wildcard-mask*] router config command to enable and advertise a network in EIGRP.
  - It enables the interfaces configured for that network address to begin transmitting & receiving EIGRP updates
  - Includes network or subnet in EIGRP updates

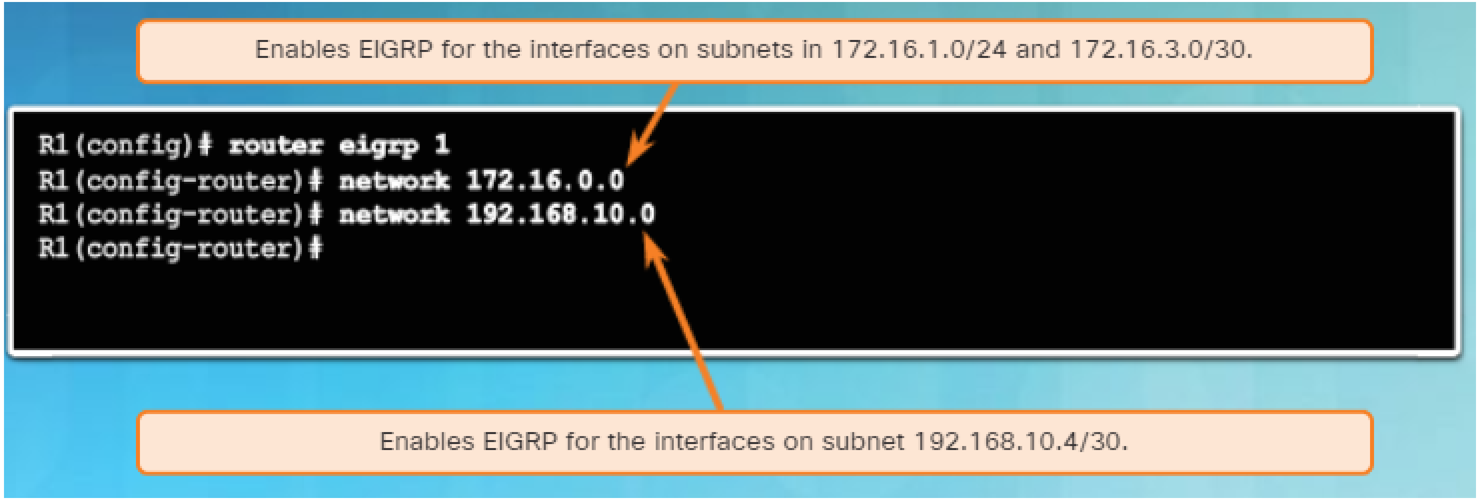

Configure EIGRP with IPv4

- A wildcard mask is similar to a subnet mask but is calculated by subtracting a SNM from 255.255.255.255.

- For example, if the SNM is 255.255.255.252:
  - 255.255.255.255
  - - 255.255.255.252
  - 0. 0. 0. 3 Wildcard mask

```
R2(config)# router eigrp 1
R2(config-router)# network 192.168.10.8 0.0.0.3
R2(config-router)
```

- EIGRP also automatically converts a subnet mask to its wildcard mask equivalent.
  - E.g., entering 192.168.10.8 **255.255.255.252** automatically converts to 192.168.10.8 **0.0.0.3**

```
R2(config)# router eigrp 1
R2(config-router)# network 192.168.10.8 255.255.255.252
R2(config-router)# end
R2# show running-config | section eigrp 1
router eigrp 1
 network 172.16.0.0
 network 192.168.10.8 0.0.0.3
 eigrp router-id 2.2.2.2
R2#
```

Configure EIGRP with IPv4

- Passive interfaces prevent EIGRP updates out a specified router interface.

Router(config-router)#

```
passive-interface type number [default]
```

```
R3(config)# router eigrp 1
R3(config-router)# passive-interface gigabitethernet 0/0
```

- Set a particular interface or all router interfaces to passive.
  - The **default** option sets all router interfaces to passive.
  - Prevents neighbor relationships from being established.
  - Routing updates from a neighbor are ignored.

```
R3# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 1"
<output omitted>
Routing for Networks:
    192.168.1.0
    192.168.10.4/30
    192.168.10.8/30
 Passive Interface(s):
   GigabitEthernet0/0
 Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.10.5          90      01:37:57
    192.168.10.9          90      01:37:57
  Distance: internal 90 external 170
R3#
```

Verify EIGRP with IPv4

- Use the **show ip eigrp neighbors** command to view the neighbor table and verify that EIGRP has established an adjacency with its neighbors.
  - The output displays a list of each adjacent neighbor.

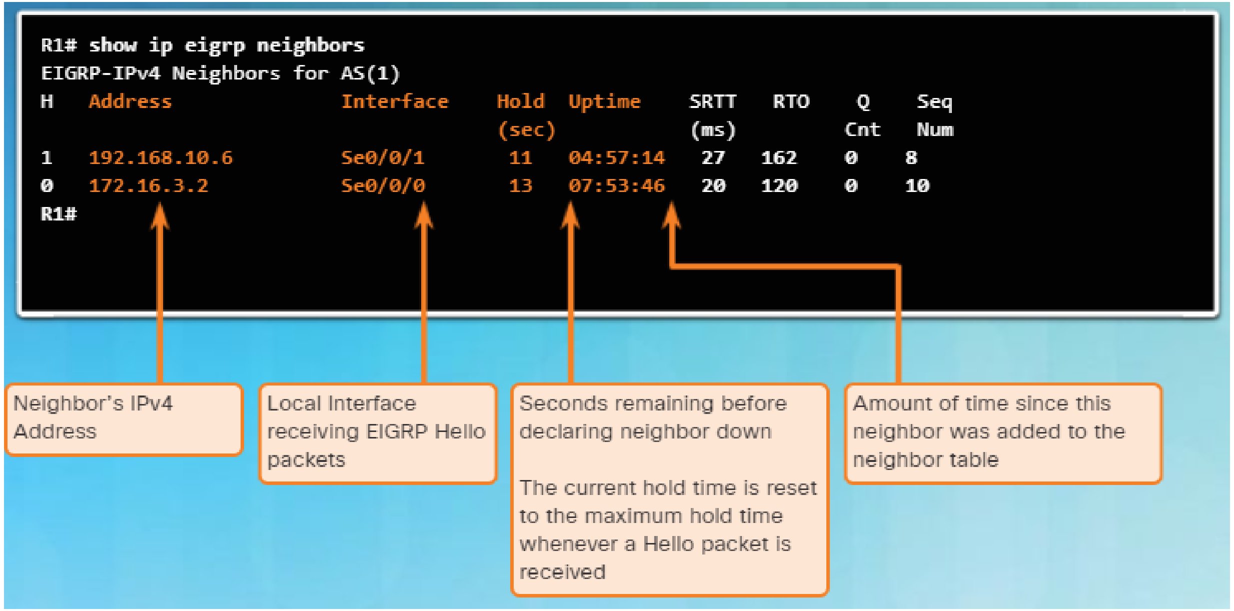

Verify EIGRP with IPv4

- The **show ip protocols** command is useful to identify the parameters and other information about the current state of any active IPv4 routing protocol processes configured on the router.
- For example, in the command output in the figure:

1. EIGRP is an active dynamic routing protocol on R1 configured with the autonomous system number 1.
2. The EIGRP router ID of R1 is 1.1.1.1.
3. The EIGRP administrative distances on R1 are internal AD of 90 and external of 170 (default values).
4. By default, EIGRP does not automatically summarize networks. Subnets are included in the routing updates.
5. The EIGRP neighbor adjacencies R1 has with other routers used to receive EIGRP routing updates.

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

Routing Protocol is "eigrp 1"  (1) Routing protocol and Process ID (AS
                                   Number)
 Outgoing update filter clist for all interfaces is not set
 Incoming update filter list for all interfaces is not set
 Default networks flagged in outgoing updates
 Default networks accepted from incoming updates
 EIGRP-IPv4 Protocol for AS(1)
   Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
   NSF-aware route hold timer is 240
   Router-ID: 1.1.1.1  (2)     EIGRP Router ID

   Topology : 0 (base)
     Active Timer: 3 min
Distance: internal 90 external 170   (3)   EIGRP Administrative
                                           Distances

     Maximum path: 4
     Maximum hopcount 100
     Maximum metric variance 1

  Automatic Summarization: disabled  (4)   EIGRP Automatic Summarization
                                           is disabled.

  Maximum path: 4
  Routing for Networks:
    172.16.0.0
    192.168.10.0
  Routing Information Sources:          (5)   EIGRP Routing
                                              Information Sources
    Gateway         Distance    Last Update   lists all the EIGRP
    192.168.10.6          90    00:40:20      routing sources the
    172.16.3.2            90    00:40:20      IOS uses to build its
                                              IPv4 routing table.

  Distance: internal 90 external 170

 R1#
```

Verify EIGRP with IPv4

```
R2# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
<output omitted>
Gateway of last resort is not set

 172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
D   172.16.1.0/24 [90/2170112] via 172.16.3.1, 00:11:05, Serial0/0/0
C   172.16.2.0/24 is directly connected, GigabitEthernet0/0
L   172.16.2.1/32 is directly connected, GigabitEthernet0/0
C   172.16.3.0/30 is directly connected, Serial0/0/0
L   172.16.3.2/32 is directly connected, Serial0/0/0
D  192.168.1.0/24 [90/2170112] via 192.168.10.10, 00:15:16,Serial0/0/1
  192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
D   192.168.10.4/30 [90/2681856] via 192.168.10.10, 00:52:00,Serial0/0/1
                    [90/2681856] via 172.16.3.1, 00:52:00,Serial0/0/0
C   192.168.10.8/30 is directly connected, Serial0/0/1
L   192.168.10.9/32 is directly connected, Serial0/0/1
  209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C   209.165.200.224/27 is directly connected, Loopback209
L   209.165.200.225/32 is directly connected, Loopback209
R2#
```

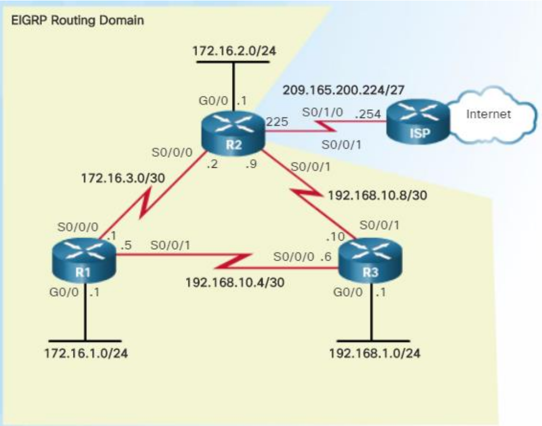


```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
<output omitted>
Gateway of last resort is not set

     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
D     172.16.2.0/24 [90/2170112] via 172.16.3.2,00:14:35, Serial0/0/0
C     172.16.3.0/30 is directly connected, Serial0/0/0
L     172.16.3.1/32 is directly connected, Serial0/0/0
D 192.168.1.0/24 [90/2170112] via 192.168.10.6,00:13:57, Serial0/0/1
  192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
C     192.168.10.4/30 is directly connected, Serial0/0/1
L     192.168.10.5/32 is directly connected, Serial0/0/1
D     192.168.10.8/30 [90/2681856] via 192.168.10.6,00:50:42, Serial0/0/1
                      [90/2681856] via 172.16.3.2,00:50:42, Serial0/0/0
R1#
```

```
R3# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
<output omitted>
Gateway of last resort is not set

 172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D    172.16.1.0/24 [90/2170112] via 192.168.10.5, 00:12:00, Serial0/0/0
D    172.16.2.0/24 [90/2170112] via 192.168.10.9, 00:16:49, Serial0/0/1
D    172.16.3.0/30 [90/2681856] via 192.168.10.9, 00:52:55, Serial0/0/1
                   [90/2681856] via 192.168.10.5, 00:52:55, Serial0/0/0
  192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, GigabitEthernet0/0
L    192.168.1.1/32 is directly connected, GigabitEthernet0/0
  192.168.10.0/24 is variably subnetted, 4 subnets, 2 masks
C    192.168.10.4/30 is directly connected, Serial0/0/0
L    192.168.10.6/32 is directly connected, Serial0/0/0
C    192.168.10.8/30 is directly connected, Serial0/0/1
L    192.168.10.10/32 is directly connected, Serial0/0/1
R3#
```

6.3 EIGRP Operation

EIGRP Initial Route Discovery

1. Router R1 starts has joined the EIGRP routing domain and sends an EIGRP Hello packet out all EIGRP enabled interfaces.
2. Router R2 receives the Hello packet and adds R1 to its neighbor table.
   - R2 sends an Update packet that contains all the routes it knows.
   - R2 also sends an EIGRP Hello packet to R1.
3. R1 updates its neighbor table with R2.

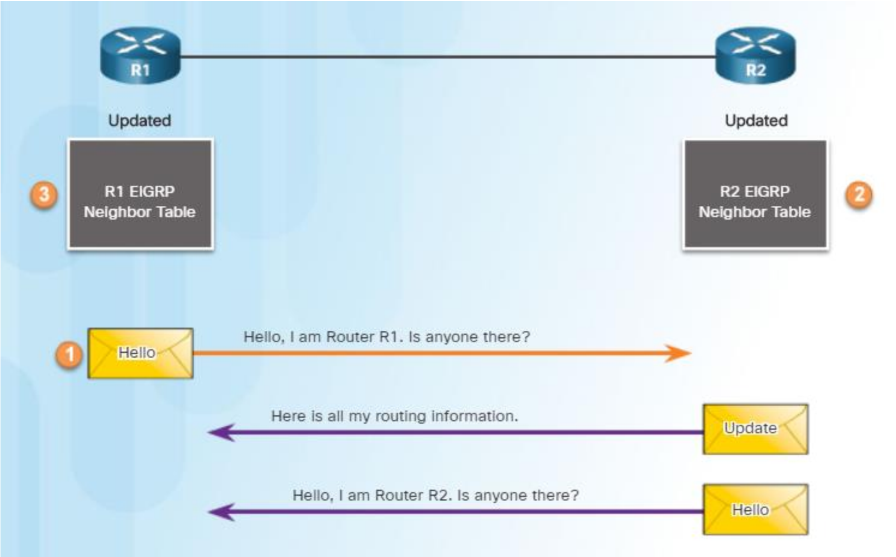

- After both routers have exchanged Hellos, the neighbor adjacency is established.

EIGRP Initial Route Discovery

1. R1 adds all update entries from R1 to its topology table.
   - The topology table includes all destinations advertised by neighboring (adjacent) routers and the cost (metric) to reach each network.

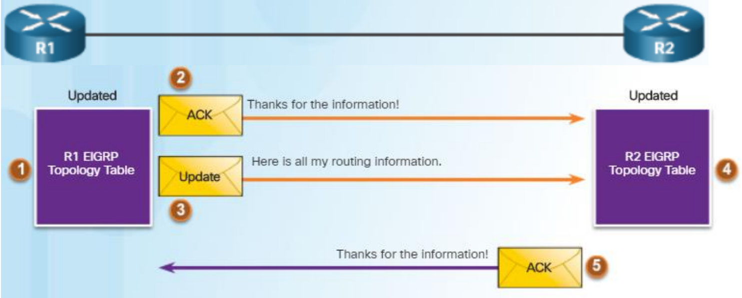

2. EIGRP update packets use reliable delivery; therefore, R1 replies with an EIGRP acknowledgment packet informing R2 that it has received the update.
3. R1 sends an EIGRP update to R2 advertising the routes that it is aware of, except those learned from R2 (split horizon).
4. R2 receives the EIGRP update from R1 and adds this information to its own topology table.
5. R2 responds to R1’s EIGRP update packet with an EIGRP acknowledgment.

EIGRP Initial Route Discovery

1. R1 uses DUAL to calculate the best routes to each destination, including the metric and the next-hop router and updates its routing table with the best routes.
2. Similarly, R2 uses DUAL and updates its routing table with the best newly discovered routes.

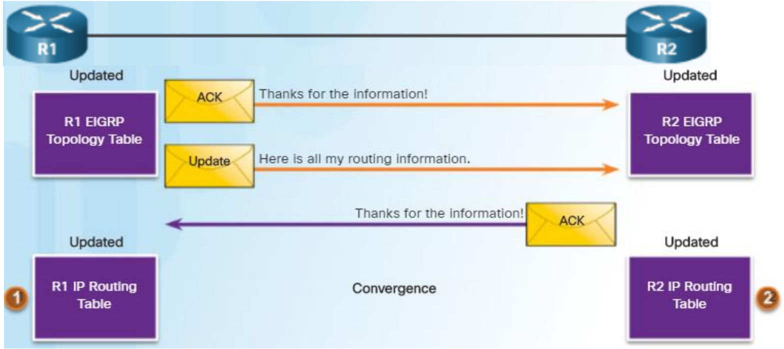

- At this point, EIGRP on both routers is considered to be in the converged state.

EIGRP Metrics

- EIGRP uses a composite metric which can be based on the following metrics:
  - **Bandwidth**: The lowest bandwidth between source and destination.
  - **Delay**: The cumulative interface delay along the path
  - **Reliability**: (Optional) Worst reliability between source and destination.
  - **Load**: (Optional) Worst load on a link between source and destination.

Note.
- It is often incorrectly stated that EIGRP can also use the smallest MTU in the path.

- The EIGRP composite metric formula consists metric weights with values K1 to K5.
  - K1 represents bandwidth, K3 delay, K4 load, and K5 reliability.

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

Default Composite Formula:

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

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] * 256

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

EIGRP Metrics

- Use the **show interfaces** command to examine the values used for bandwidth, delay, reliability, and load.

- **BW** - Bandwidth of the interface (in kb/s).
- **DLY** - Delay of the interface (in microseconds).
- **Reliability** - Reliability of the interface as a fraction of 255 (255/255 is 100% reliability).
- **Txload, Rxload** - Transmit and receive load on the interface as a fraction of 255 (255/255 is completely saturated), calculated as an exponential average over five minutes.

```
R1# show interfaces serial 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 1544 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation HDLC, loopback not set
<output omitted>
R1#

R1# show interfaces gigabitethernet 0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is CN Gigabit Ethernet, address is fc99.4775.c3e0
(bia fc99.4775.c3e0)
  Internet address is 172.16.1.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
<output omitted>
R1#
```

EIGRP Metrics

- Use the following interface configuration mode command to modify the bandwidth metric:
  - Router(config-if)# **bandwidth** *kilobits-bandwidth-value*

```
R2(config)# interface s 0/0/0
R2(config-if)# bandwidth 64
R2(config-if)# exit
R2(config)# interface s 0/0/1
R2(config-if)# bandwidth 1024
```

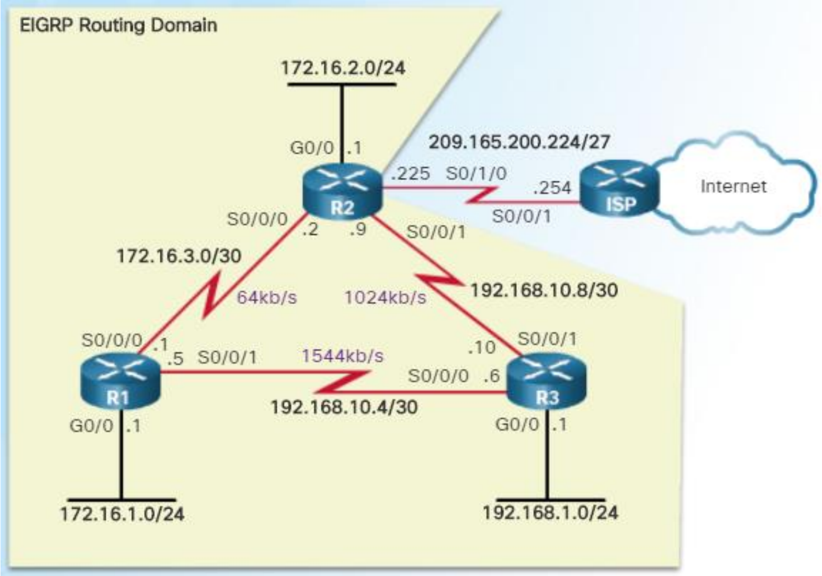

- Use the **show interfaces** command to verify the new bandwidth parameters.

```
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>
R1#
```

```
R2# show interface s 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Internet address is 172.16.3.2/30
  MTU 1500 bytes, BW 64 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
R2#
```

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

```
R3(config)# interface s 0/0/1
R3(config-if)# bandwidth 1024
```

EIGRP Metrics

- Delay is a measure of the time it takes for a packet to traverse a route.
- The delay (DLY) metric is not measured dynamically.
  - It is a static value measured in microseconds (μs or usec) based on the type of link to which the interface is connected.
- The delay value is calculated using the cumulative (sum) of all interface delays along the path, divided by 10.

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

EIGRP Metrics

- We can determine the EIGRP metric as follows:
  1. Determine the link with the slowest bandwidth and use that value to calculate bandwidth (10,000,000/bandwidth).
  2. Determine the delay value for each outgoing interface on the way to the destination and add the delay values and divide by 10 (sum of delay/10).
  3. This composite metric produces a 24-bit value which EIGRP multiplies with 256.

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

Because K1 and K3 both equal 1, the formula becomes:

(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
```

EIGRP Metrics

- How does EIGRP determine the following metric?

```
R2# show ip route
<output omitted>
D 192.168.1.0/24 [90/3012096] via 192.168.10.10, 00:12:32, Serial0/0/1
```

- **EIGRP Composite Metric** = (Bandwidth + Delay) x 256

- **Bandwidth** = 10,000,000 / slowest bandwidth

```
R2# show interface s 0/0/1
Serial0/0/1 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Internet address is 192.168.10.9/30
  MTU 1500 bytes, BW 1024 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
R2#
```

```
R3# show interface g 0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is CN Gigabit Ethernet, address is fc99.4771.7a20 (bia fc99.4771.7a20)
  Internet address is 192.168.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
R3#
```

- **Bandwidth** = 10,000,000 / 1024 = **9765**

- **Delay** = (Sum of all delays) / 10

```
R2# show interface s 0/0/1
Serial0/0/1 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Internet address is 192.168.10.9/30
  MTU 1500 bytes, BW 1024 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
R2#
```

```
R3# show interface g 0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is CN Gigabit Ethernet, address is fc99.4771.7a20 (bia fc99.4771.7a20)
  Internet address is 192.168.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
  R3#
```

- **Delay** = (20,000 + 10) / 1024 = **2001**

- **EIGRP Composite Metric** = (**9765** + **2001**) x 256 = **3,012,096**

cisco

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.

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.
- The decision process for all route computations is done by the DUAL Finite State Machine (FSM). An FSM is a workflow model, similar to a flow chart, which is composed of the following:
  - A finite number of stages (states)
  - Transitions between those stages
  - Operations
- The DUAL FSM tracks all routes and 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.
  - The successor IP address is shown right after “via”.
- FD is the lowest calculated metric to reach the destination network.
  - FD is the second number inside the brackets.
  - 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.

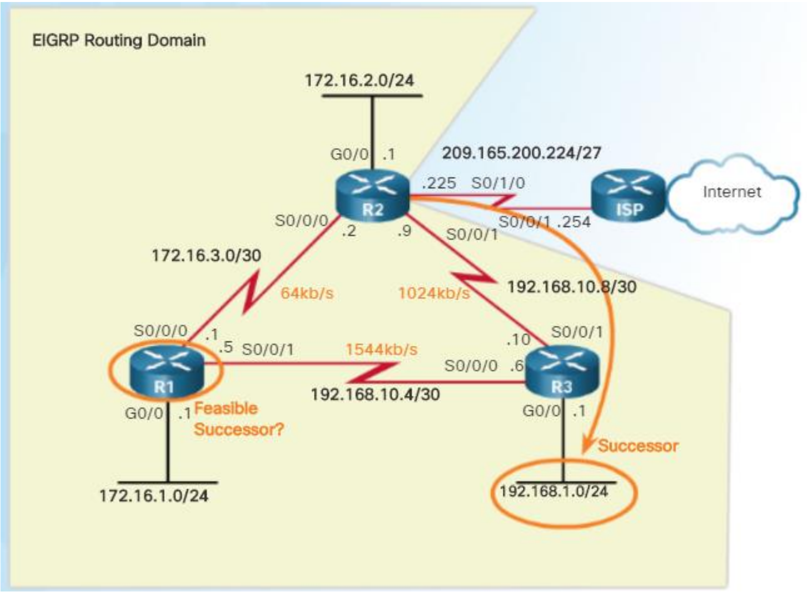

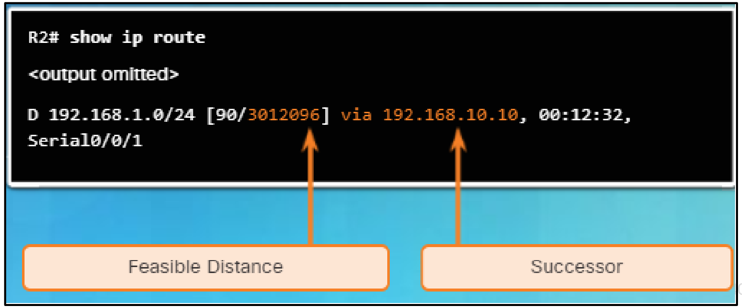

DUAL and the Topology Table

- 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 FC 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.

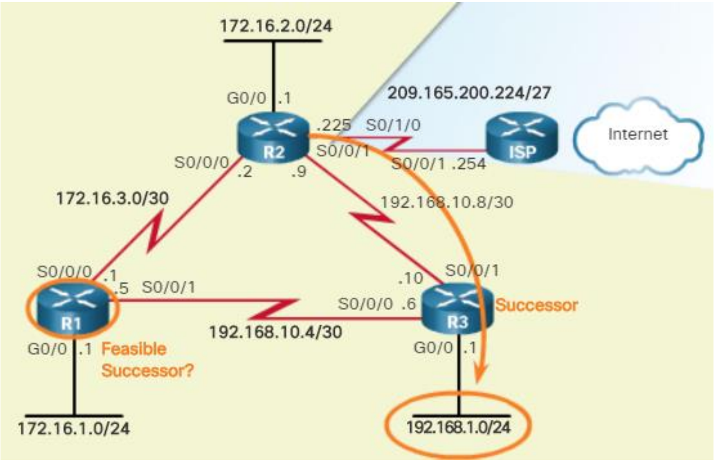

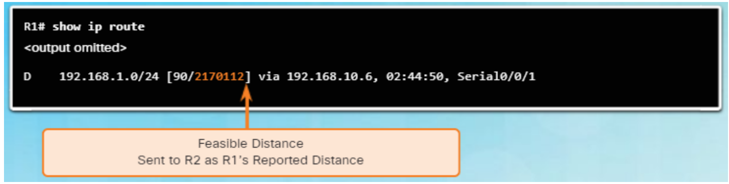

DUAL and the Topology Table

- 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.

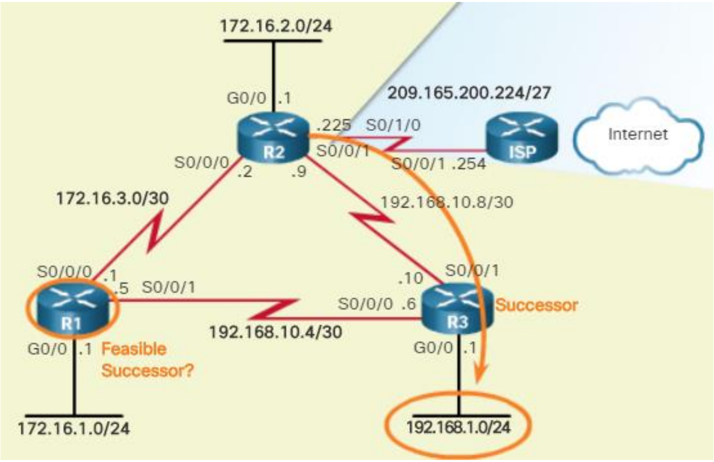

- Use the **show ip eigrp topology** command to list all successors and FSs to destination networks.
  - Only the successor is installed into the IP routing table.
  - **Passive State** - Route is in stable state and available for use.
  - **Active State** - Route is being recomputed by DUAL.

```
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#
```

DUAL and the Topology Table

- The first line in the topology table displays:
  - **P** - Route in the passive state (the route is in a stable mode). If DUAL recalculates or searches for a new path, the route is in an active state and displays an A.
  - **192.168.1.0/24** - Destination network is also found in the routing table.
  - **1 successors** - Displays the number of successors for this network. If there are multiple equal cost paths to this network, there are multiple successors.
  - **FD is 3012096** - FD, the EIGRP metric to reach the destination network. This is the metric displayed in the IP routing table.

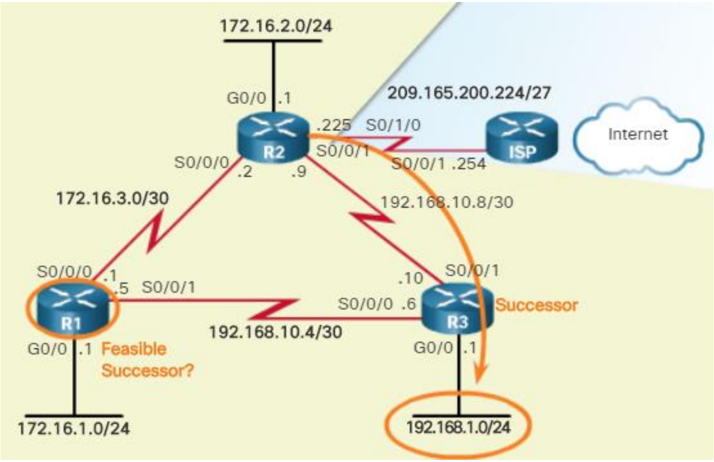


```
R2# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(2.2.2.2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

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#
```

DUAL and the Topology Table

- The partial output of the **show ip route** command displays the 192.168.1.0/24 route with the successor is R3 via 192.168.10.6 with an FD of 2,170,112.

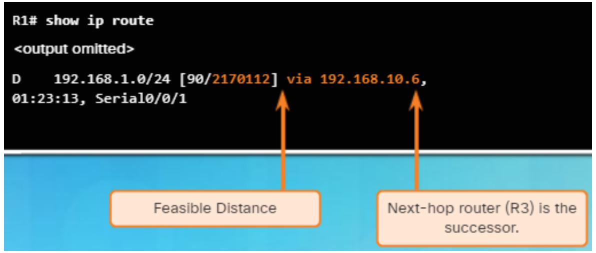

- The **show ip eigrp topology** command only shows the successor 192.168.10.6, which is R3.
  - Notice there are no FSs.

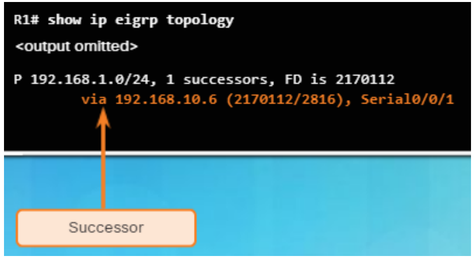

- The **show ip eigrp topology all-links** command shows all possible paths to a network, including successors, FSs, and even those routes that are not FSs.

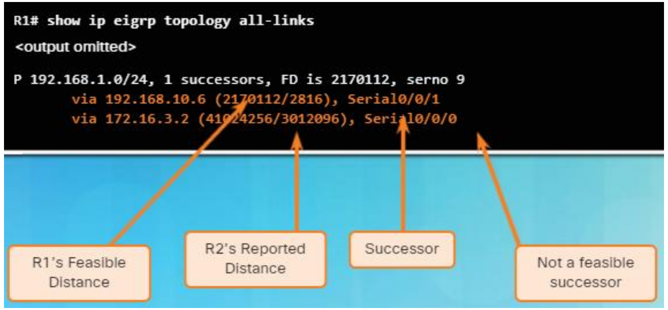

DUAL and Convergence

- The DUAL Finite State Machine (FSM) contains all of the logic used to calculate and compare routes in an EIGRP network.
- An FSM is an abstract machine, that defines a set of possible states that something can go through, what events cause those states, and what events result from those states.
  - Designers use FSMs to describe how a device, computer program, or routing algorithm reacts to a set of input events.

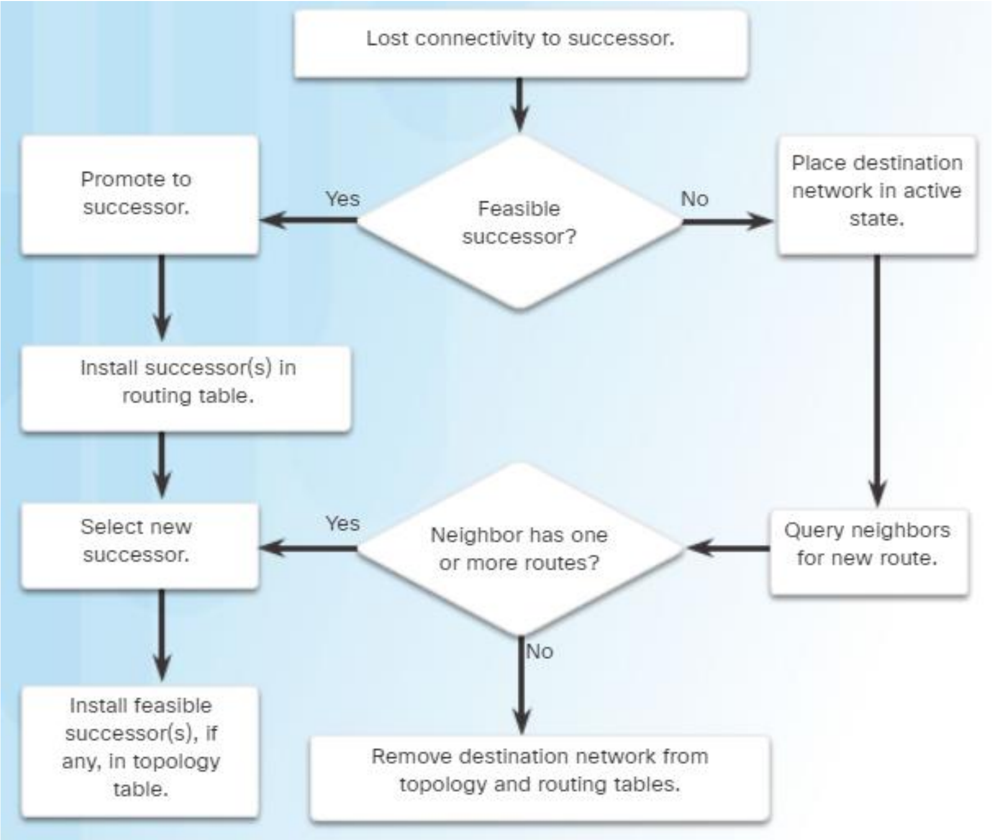

DUAL and Convergence

- IP EIGRP

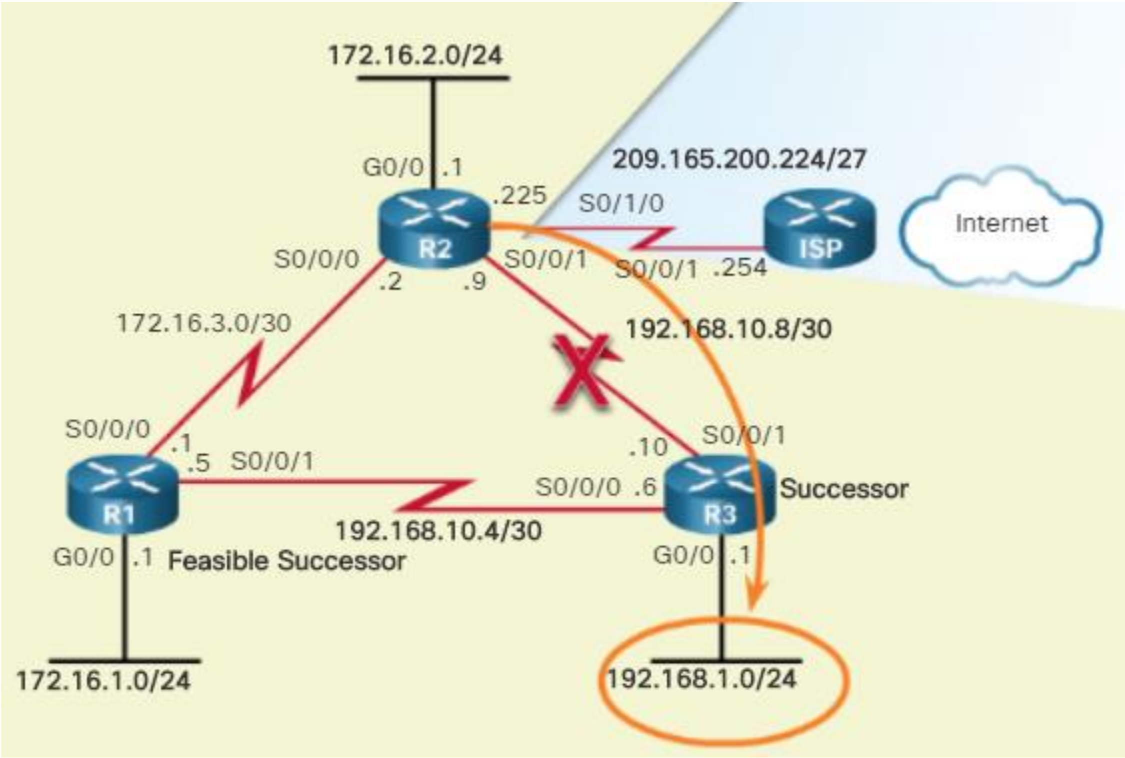

DUAL and Convergence

- If the path to the successor fails and there are no FSs, DUAL puts the network into the active state and actively queries its neighbors for a new successor.
  - DUAL sends EIGRP queries asking other routers for a path to the network.
  - Other routers return EIGRP replies, letting the sender of the EIGRP query know that they have a path to the requested network. If there is no reply, the sender of the query does not have a route to this network.
  - If the sender receives EIGRP replies with a path to the requested network, the preferred path is added as the new successor and also added to the routing table.

6.4 Implement EIGRP for IPv6

EIGRP for IPv6

- EIGRP for IPv6 is a distance-vector routing protocol.
  - The configuration and operation is similar to EIGRP for IPv4.
- The following remained the same as EIGRP for IPv4:
  - Uses the same protocol number (88)
  - Maintains a topology table and queries if no feasible successors are available.
  - Uses DUAL to calculate the successor routes

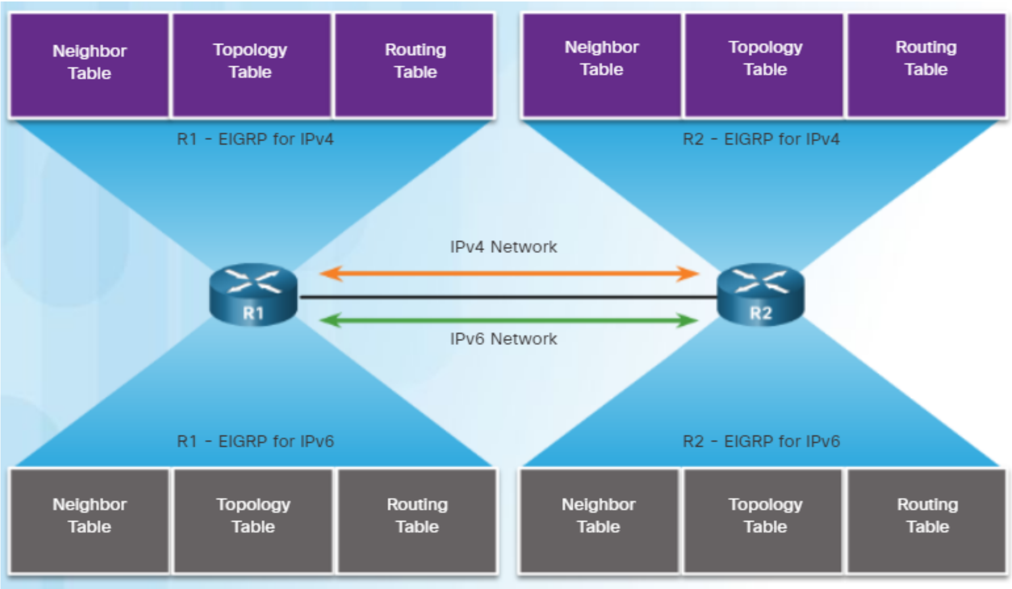

EIGRP for IPv6

- The following compares EIGRP for IPv4 and IPv6

	EIGRP for IPv4	EIGRP for IPv6
Advertised Routes	IPv4 networks	IPv6 prefixes
Distance Vector	Yes	Yes
Convergence Technology	DUAL	DUAL
Metric	Bandwidth and delay by default, reliability and load are optional	Bandwidth and delay by default, reliability and load are optional
Transport Protocol	RTP	RTP
Update Messages	Incremental, partial, and bounded updates	Incremental, partial, and bounded updates
Neighbor Discovery	Hello packets	Hello packets
Source and Destination Addresses	IPv4 source address and 224.0.0.10 IPv4 multicast destination address	IPv6 link-local source address and FF02::A IPv6 multicast destination address
Authentication	MD5, SHA256	MD5, SHA256
Router ID	32-bit router ID	32-bit router ID

EIGRP for IPv6

- EIGRP for IPv6 messages are sent using:
  - **Source IPv6 address** - This is the IPv6 link-local address of the exit interface.
  - **Destination IPv6 address** - When the packet needs to be sent to a multicast address, it is sent to the IPv6 multicast address FF02::A, the all-EIGRP-routers with link-local scope. If the packet can be sent as a unicast address, it is sent to the link-local address of the neighboring router.

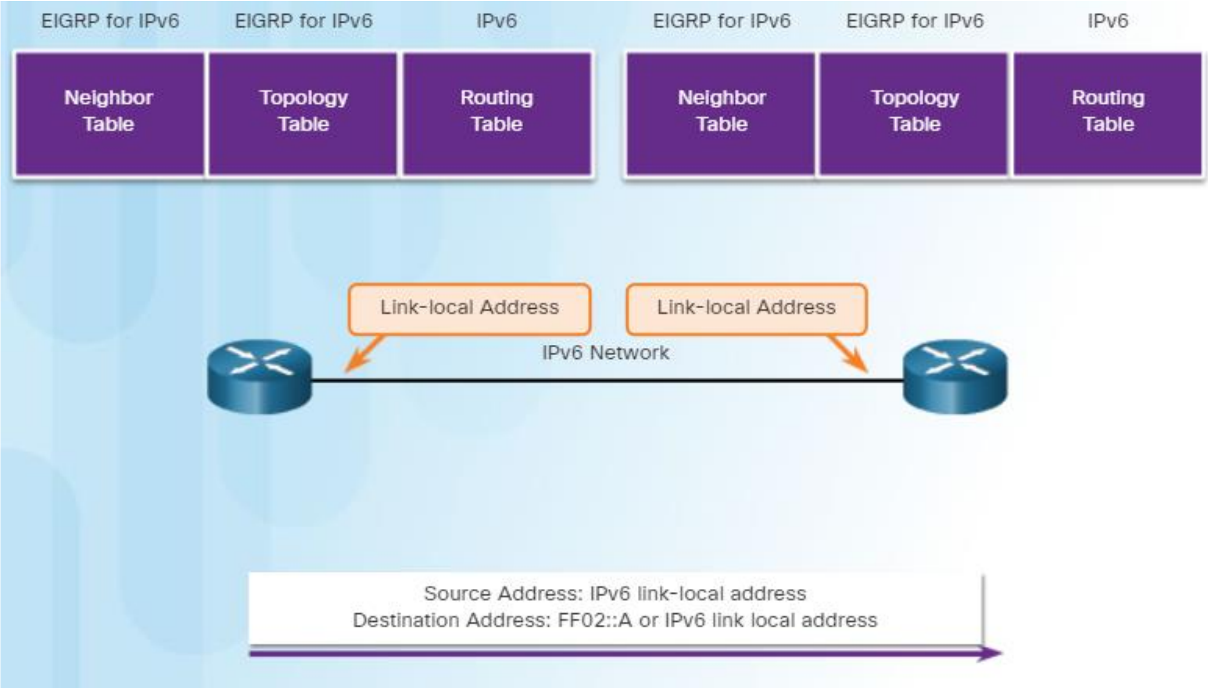

Implement EIGRP for IPv6

Configure EIGRP for IPv6

```
R2# show running-config
<output omitted>
!
interface GigabitEthernet0/0
 ipv6 address 2001:DB8:CAFE:2::1/64
!
interface Serial0/0/0
 ipv6 address 2001:DB8:CAFE:A001::2/64
!
interface Serial0/0/1
 ipv6 address 2001:DB8:CAFE:A002::1/64
 clock rate 64000
!
interface Serial0/1/0
 ipv6 address 2001:DB8:FEED:1::1/64
```

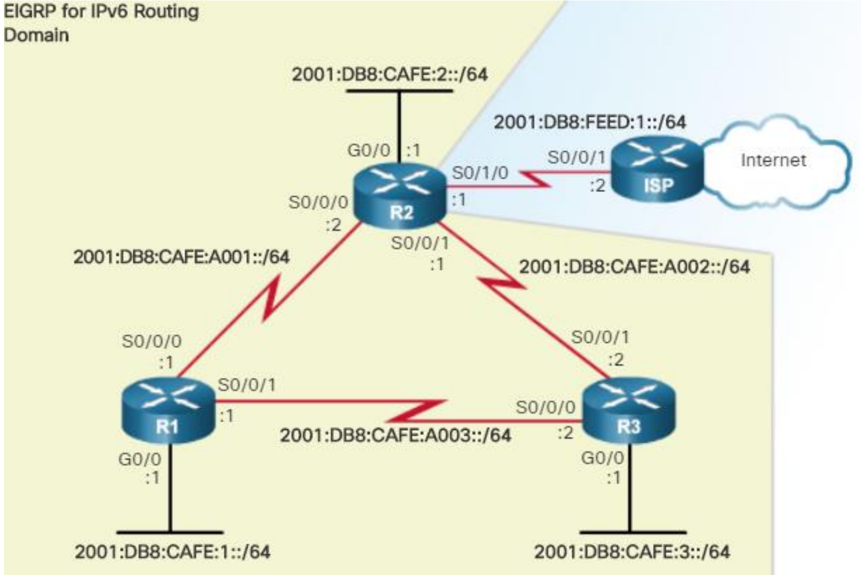


```
R1# show running-config
<output omitted>
!
interface GigabitEthernet0/0
 ipv6 address 2001:DB8:CAFE:1::1/64
!
interface Serial0/0/0
 ipv6 address 2001:DB8:CAFE:A001::1/64
 clock rate 64000
!
interface Serial0/0/1
 ipv6 address 2001:DB8:CAFE:A003::1/64
```

```
R3# show running-config
<output omitted>
!
interface GigabitEthernet0/0
 ipv6 address 2001:DB8:CAFE:3::1/64
!
interface Serial0/0/0
 ipv6 address 2001:DB8:CAFE:A003::2/64
 clock rate 64000
!
interface Serial0/0/1
 ipv6 address 2001:DB8:CAFE:A002::2/64
```

Configure EIGRP for IPv6

```
R2(config)# interface s 0/0/0
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# exit
R2(config)# interface s 0/0/1
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# exit
R2(config)# interface s 0/1/0
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)# exit
R2(config)# interface g 0/0
R2(config-if)# ipv6 address fe80::2 link-local
R2(config-if)#
```

```
R1(config)# interface s 0/0/0
R1(config-if)# ipv6 address fe80::1 ?
  link-local  Use link-local address

R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
R1(config)# interface s 0/0/1
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
R1(config)# interface g 0/0
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)#
```

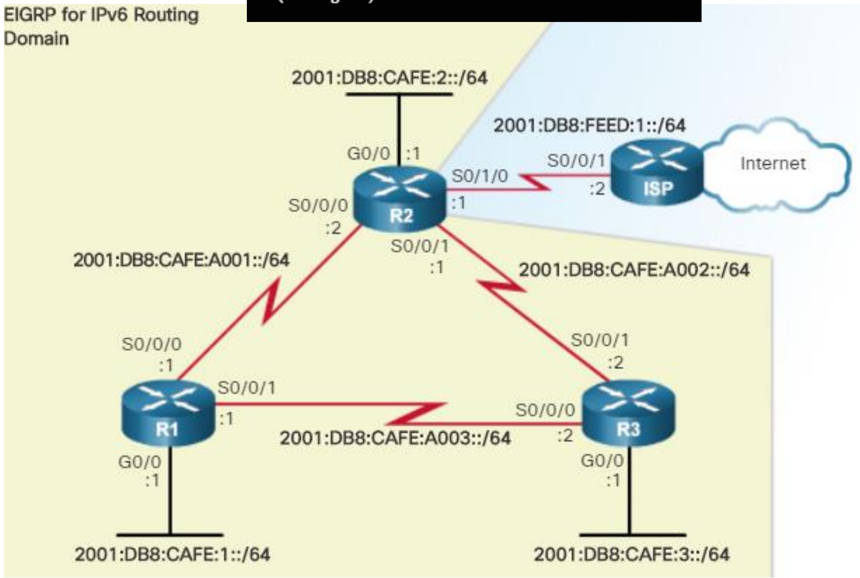


```
R1# show ipv6 interface brief
GigabitEthernet0/0      [up/up]
    FE80::1
    2001:DB8:CAFE:1::1
Serial0/0/0             [up/up]
    FE80::1
    2001:DB8:CAFE:A001::1
Serial0/0/1             [up/up]
    FE80::1
    2001:DB8:CAFE:A003::1
R1#
```

Same IPv6 link-local address is configured on all interfaces.

Configure EIGRP for IPv6

- The **ipv6 unicast-routing** global config mode command enables IPv6 routing on the router.
- Use the **ipv6 router eigrp** *autonomous-system* to enter EIGRP for IPv6 router configuration mode.
- Use the **eigrp router-id** *router-id* command is used to configure the router ID.
- By default, the EIGRP for IPv6 process is in a shutdown state and the **no shutdown** command is required to activate the EIGRP for IPv6 process.

```
R2(config)# ipv6 unicast-routing
R2(config)# ipv6 router eigrp 2
R2(config-rtr)# eigrp router-id 2.0.0.0
R2(config-rtr)# no shutdown
R2(config-rtr)#
```

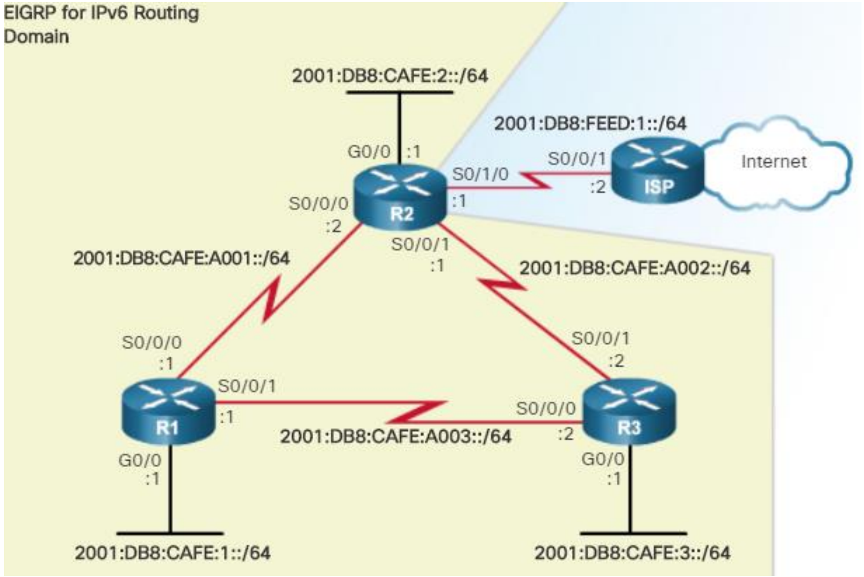

Implement EIGRP for IPv6
Configure EIGRP for IPv6

- Unlike EIGRP for IPv4 which uses the **network** command, EIGRP for IPv6 is configured directly on the interface using the **ipv6 eigrp** *autonomous-system* interface configuration command.

```
R2(config)# interface g 0/0
R2(config-if)# ipv6 eigrp 2
R2(config-if)# exit
R2(config)# interface s 0/0/0
R2(config-if)# ipv6 eigrp 2
R2(config-if)# exit
%DUAL-5-NBRCHANGE: EIGRP-IPv6 2: Neighbor FE80::1
 (Serial0/0/0) is up: new adjacency
R2(config)# interface s 0/0/1
R2(config-if)# ipv6 eigrp 2
R2(config-if)#
```

```
R1(config)# interface g0/0
R1(config-if)# ipv6 eigrp 2
R1(config-if)# exit
R1(config)# interface s 0/0/0
R1(config-if)# ipv6 eigrp 2
R1(config-if)# exit
R1(config)# interface s 0/0/1
R1(config-if)# ipv6 eigrp 2
R1(config-if)#
```

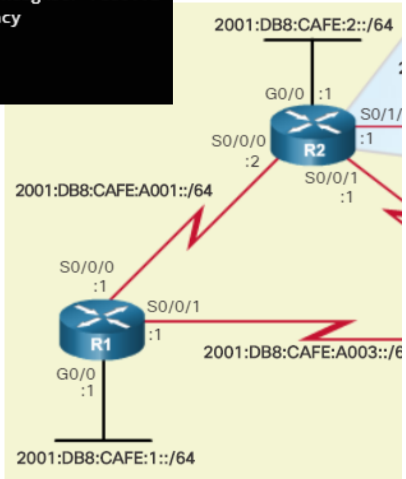

The same **passive-interface** command used for IPv4 is used with EIGRP for IPv6.

```
R1(config)# ipv6 router eigrp 2
R1(config-rtr)# passive-interface gigabitethernet 0/0
R1(config-rtr)# end

R1# show ipv6 protocols

IPv6 Routing Protocol is "eigrp 2"
EIGRP-IPv6 Protocol for AS(2)
<output omitted>

  Interfaces:
    Serial0/0/0
    Serial0/0/1
    GigabitEthernet0/0 (passive)
 Redistribution:
    None
R1#
```

Verifying EIGRP for IPv6

- Use the **show ipv6 eigrp neighbors** command to view the neighbor table and verify that EIGRP for IPv6 has established an adjacency with its neighbors.

  - **H** - Lists the neighbors in order they were learned.
  - **Address** - IPv6 link-local address of the neighbor.
  - **Interface** - Local interface that received the Hello.
  - **Hold** - Current hold time.
  - **Uptime** - Time since this neighbor was added.
  - **SRTT** and **RTO** - Used by RTP.
  - **Queue Count** - Should always be zero.
  - **Sequence Number** - Used to track updates, queries, and reply packets.

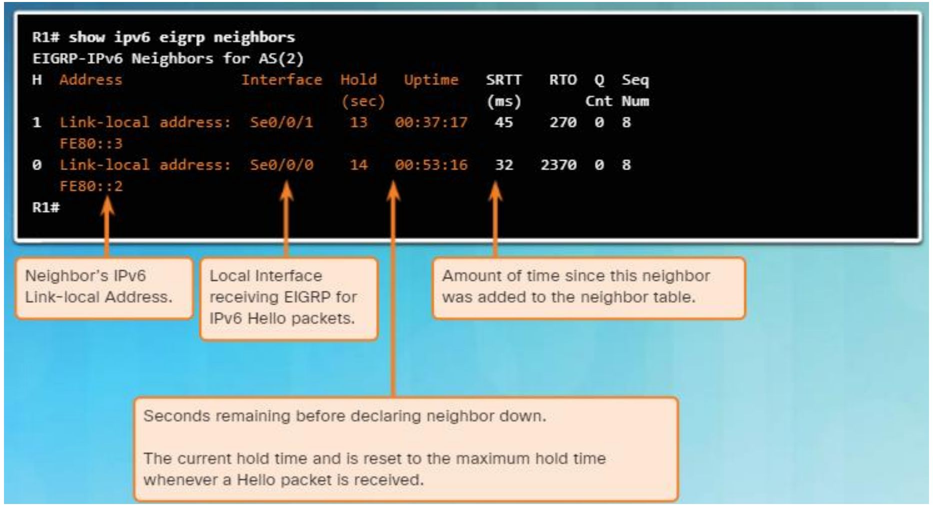

Verifying EIGRP for IPv6

- The **show ipv6 protocols** command displays the parameters and other information about the state of any active IPv6 routing protocol processes currently configured on the router.

1. EIGRP for IPv6 is an active dynamic routing protocol on R1.
2. These are the *k* values used to calculate the EIGRP composite metric.
3. The EIGRP for IPv6 router ID of R1 is 1.0.0.0.
4. Same as EIGRP for IPv4, EIGRP for IPv6 administrative distances have internal AD of 90 and external of 170 (default values).
5. The interfaces enabled for EIGRP for IPv6.

```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 2"
 EIGRP-IPv6 Protocol for AS(2)

  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0

  NSF-aware route hold timer is 240
  Router-ID: 1.0.0.0
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1

  Interfaces:
    GigabitEthernet0/0
    Serial0/0/0
    Serial0/0/1
  Redistribution:
    None
R1#
```

1 Routing protocol and Process ID (AS Number)
2 K values used in composite metric
3 EIGRP Router ID
4 EIGRP Administrative Distances
5 Interfaces enabled for EIGRP for IPv6

Verifying EIGRP for IPv6

- Use the **show ipv6 route** command to examine the IPv6 routing table.
  - EIGRP for IPv6 routes are denoted with a **D**.
- The figure shows that R1 has installed three EIGRP routes to remote IPv6 networks in its IPv6 routing table:
  - 2001:DB8:CAFE:2::/64 via R3 (FE80::3) using its Serial 0/0/1 interface
  - 2001:DB8:CAFE:3::/64 via R3 (FE80::3) using its Serial 0/0/1 interface
  - 2001:DB8:CAFE:A002::/64 via R3 (FE80::3) using its Serial 0/0/1 interface

```
R1# show ipv6 route
<output omitted>

C    2001:DB8:CAFE:1::/64 [0/0]
      via GigabitEthernet0/0, directly connected
L    2001:DB8:CAFE:1::1/128 [0/0]
      via GigabitEthernet0/0, receive
D    2001:DB8:CAFE:2::/64 [90/3524096]
      via FE80::3, Serial0/0/1
D    2001:DB8:CAFE:3::/64 [90/2170112]
      via FE80::3, Serial0/0/1
C    2001:DB8:CAFE:A001::/64 [0/0]
      via Serial0/0/0, directly connected
L    2001:DB8:CAFE:A001::1/128 [0/0]
      via Serial0/0/0, receive
D    2001:DB8:CAFE:A002::/64 [90/3523840]
      via FE80::3, Serial0/0/1
C    2001:DB8:CAFE:A003::/64 [0/0]
      via Serial0/0/1, directly connected
L    2001:DB8:CAFE:A003::1/128 [0/0]
      via Serial0/0/1, receive
L    FF00::/8 [0/0]
      via Null0, receive
R1#
```

6.5 Chapter Summary

Chapter 6: EIGRP

- EIGRP (Enhanced Interior Gateway Routing Protocol) is a classless, distance vector routing protocol.
- EIGRP uses the source code of "D" for DUAL in the routing table. EIGRP has a default administrative distance of 90 for internal routes and 170 for routes imported from an external source, such as default routes. These features include: Diffusing Update Algorithm (DUAL), establishing neighbor adjacencies, Reliable Transport Protocol (RTP), partial and bounded updates, and equal and unequal cost load balancing.
- EIGRP uses PDMs (Protocol Dependent Modules) giving it the capability to support different Layer 3 protocols including IPv4 and IPv6. EIGRP uses reliable delivery for EIGRP updates, queries and replies; and uses unreliable delivery for EIGRP Hellos and acknowledgments. Reliable RTP means an EIGRP acknowledgment must be returned.
- Before any EIGRP updates are sent, a router must first discover its neighbors using EIGRP Hello packets. The Hello and hold-down values do not need to match for two routers to become neighbors. The **show ip eigrp neighbors** command is used to view the neighbor table and verify that EIGRP has established an adjacency with its neighbors.

Chapter 6: EIGRP (Cont.)

- EIGRP sends partial or bounded updates, which include only route changes. Updates are sent only to those routers that are affected by the change. EIGRP composite metric uses bandwidth, delay, reliability, and load to determine the best path. By default only bandwidth and delay are used.
- At the center of EIGRP is DUAL (Diffusing Update Algorithm). The DUAL Finite State Machine is used to determine best path and potential backup paths to every destination network. The successor is a neighboring router that is used to forward the packet using the least-cost route to the destination network. Feasible distance (FD) is the lowest calculated metric to reach the destination network through the successor. A feasible successor (FS) is a neighbor who has a loop-free backup path to the same network as the successor, and also meets the feasibility condition. The 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. The reported distance is simply an EIGRP neighbor's feasible distance to the destination network.

Chapter 6: EIGRP (Cont.)

- EIGRP is configured with the **router eigrp** *autonomous-system* command. The autonomous-system value is actually a process-id and must be the same on all routers in the EIGRP routing domain. The **network** command is similar to that used with RIP. The network is the classful network address of the directly connected interfaces on the router. A wildcard mask is an optional parameter that can be used to include only specific interfaces.
- EIGRP for IPv6 shares many similarities with EIGRP for IPv4. However, unlike the IPv4 network command, IPv6 is enabled on the interface using **the ipv6 eigrp** *autonomous-system* interface configuration command.

cisco