



Video  
Frank Schneemann, MS EdTech

# EIGRP



## Routing Protocols and Concepts – Chapter 9

Cisco | Networking Academy®  
Mind Wide Open™

### **Enhanced Interior Gateway Routing Protocol (EIGRP) is a distance vector, classless routing protocol**

- Describe the background and history of Enhanced Interior Gateway Routing Protocol (EIGRP).
- Examine the basic EIGRP configuration commands and identify their purposes.
- Calculate the composite metric used by EIGRP.
- Describe the concepts and operation of DUAL.
- Describe the uses of additional configuration commands in EIGRP.

## 9.0.1 Introduction

### Interior Gateway Protocols

	Distance Vector Routing Protocols		Link State Routing Protocols	
Classful	<b>RIP</b>	IGRP		
Classless	RIPv2	EIGRP	<b>OSPFv2</b>	IS-IS
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6

In this chapter, you will learn to:

- Describe the background and history of EIGRP.
- Describe the features and operation of EIGRP.
- Examine the basic EIGRP configuration commands and identify their purposes.
- Calculate the composite metric used by EIGRP.
- Describe the concepts and operation of DUAL.
- Describe the uses of additional configuration commands in EIGRP.

## 9.1.1 - EIGRP An Enhanced Distance Vector Routing Protocol



### Summary of Operations

#### Traditional Distance Vector Routing Protocols

- Use the Bellman-Ford or Ford-Fulkerson algorithm.
- Age out routing entries and uses periodic updates.
- Keep track of only the best routes; the best path to a destination network.
- When a route becomes unavailable, the router must wait for a new routing update.
- Slower convergence due to holddown timers.

#### Enhanced Distance Vector Routing Protocol: EIGRP

- Uses the Diffusing Update Algorithm (DUAL).
- Does not age out routing entries nor uses periodic updates.
- Maintains a topology table separate from the routing table, which includes the best path and any loop-free backup paths.
- When a route becomes unavailable, DUAL will use a backup path if one exists in the topology table.
- Faster convergence due to the absence of holddown timers and a system of coordinated route calculations.

## 9.1.2 EIGRP Message Format

### Encapsulated EIGRP Message

Data Link Frame  
Header

IP Packet  
Header

EIGRP Packet  
Header

Type/Length/Values Types

#### Data Link Frame

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

MAC Source Address = Address of sending interface

#### 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

#### TLV Types

Some types include:

0x0001 EIGRP Parameters

0x0102 IP Internal Routes

0x0103 IP External Routes

## 9.1.2 EIGRP Message Format

### Encapsulated EIGRP Message



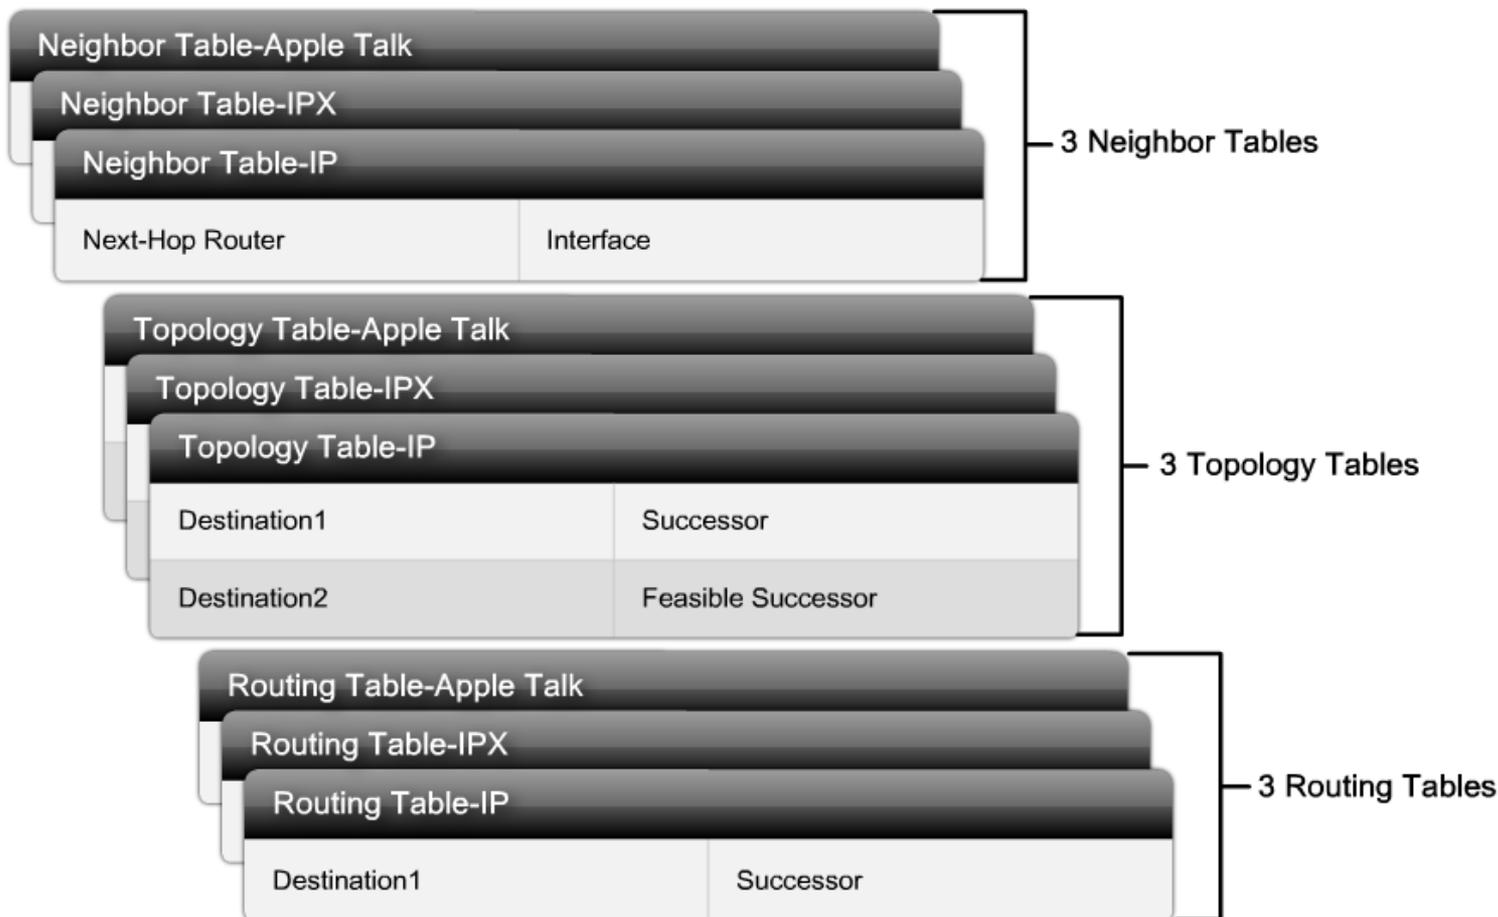
- Every **EIGRP message** includes the **header**. Important fields for our discussion include the Opcode field and the Autonomous System Number field. Opcode specifies the EIGRP packet type:
- The **EIGRP parameters** message includes the weights that EIGRP uses for its composite metric. By default, only bandwidth and delay are weighted
- The **IP Internal message** is used to advertise EIGRP routes within an autonomous system.
- The **IP External message** is used when external routes are imported into the EIGRP routing process.

The Autonomous System (AS) Number specifies the EIGRP routing process. Unlike RIP, Cisco routers can run multiple instances of EIGRP. The AS number is used to track multiple instances of EIGRP.

*In the curriculum you will be able to get more detailed information about the message format fields*

## 9.1.3 Protocol Dependent Modules (PDM)

EIGRP Protocol-Dependent Modules (PDM)

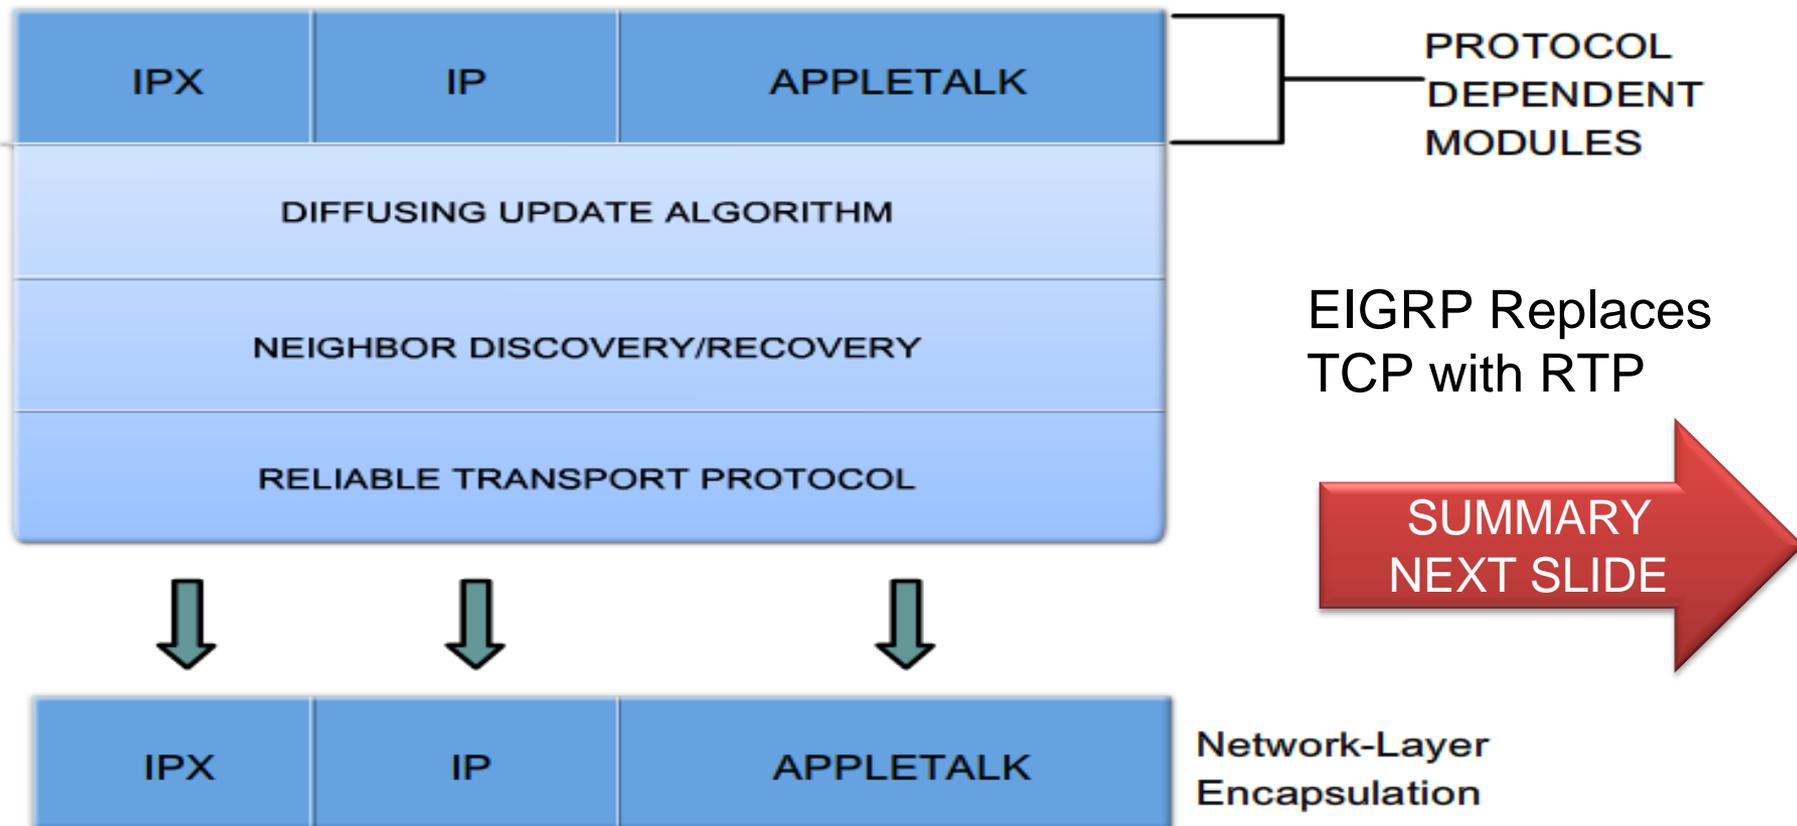


### Example

EIGRP uses different EIGRP packets and maintains separate neighbor, topology, and routing tables for each Network layer protocol.

EIGRP has the capability for routing several different protocols including IP, IPX, and AppleTalk using protocol-dependent modules (PDM). PDMs are responsible for the specific routing tasks for each Network layer protocol.

## 9.1.4 RTP and EIGRP Packet Types



**Reliable Transport Protocol (RTP)** is the protocol used by EIGRP for the delivery and reception of EIGRP packets. EIGRP was designed as a **Network layer independent routing protocol**; therefore, it cannot use the services of UDP or TCP because IPX and Appletalk do not use protocols from the TCP/IP protocol suite. The figure shows conceptually how RTP operates.

## 9.1.4 RTP and EIGRP Packet Types

### Reliable Transport Protocol (RTP)

#### ▪ Purpose of RTP

–Used by EIGRP to transmit and receive EIGRP packets

#### ▪ Characteristics of RTP

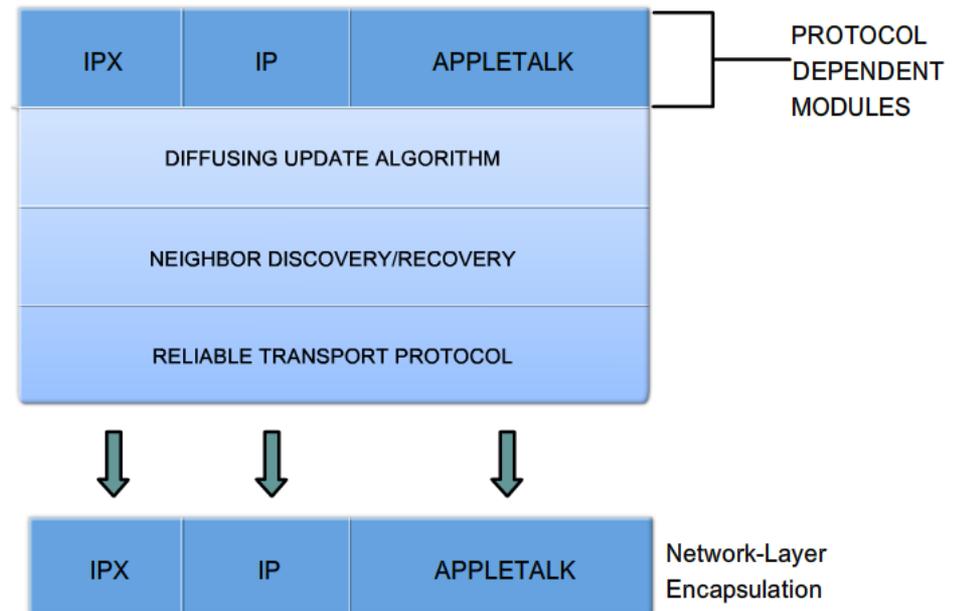
–Involves both reliable & unreliable delivery of EIGRP packet

- **Reliable** delivery requires acknowledgment from destination
- **Unreliable** delivery does not require an acknowledgement from destination

–Packets can be sent

- Unicast
- Multicast

Using address 224.0.0.10

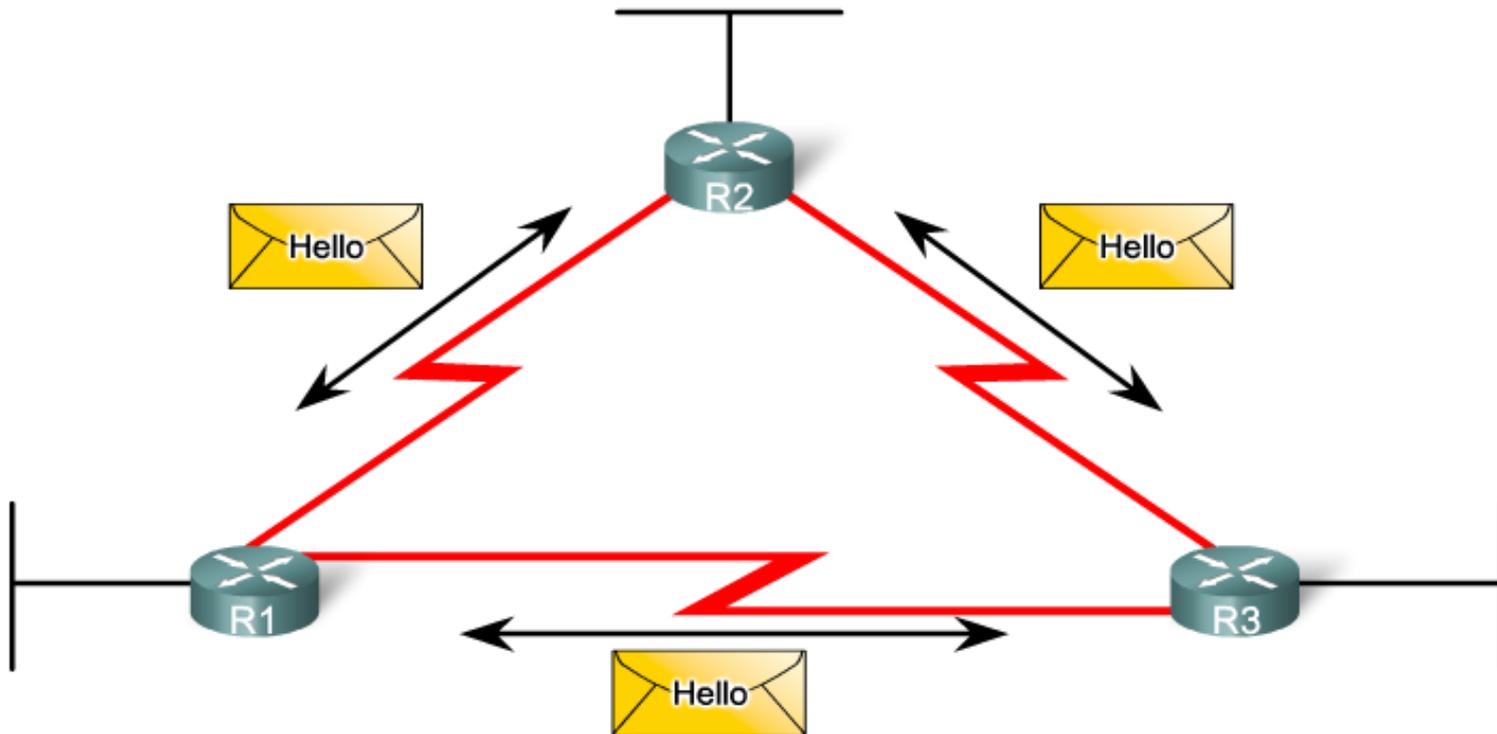


## 9.1.4 RTP and EIGRP Packet Types

### EIGRP's 5 Packet Types

- **Hello packets**

- Used to discover & form adjacencies with neighbors



#### Hello packet

- Use to discover neighbors & form adjacencies
- Unreliable so no response required from recipient

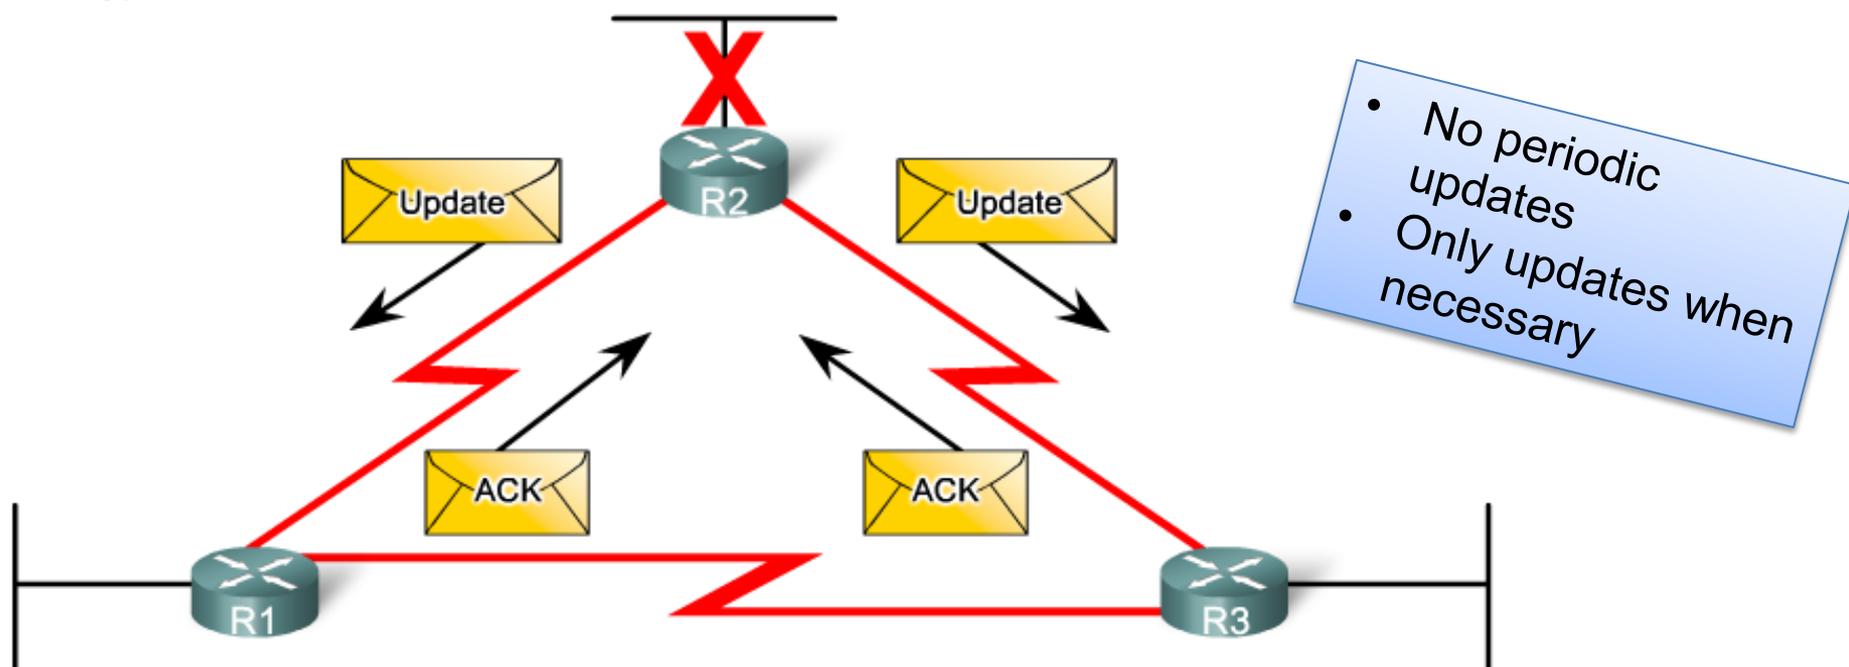
## 9.1.4 RTP and EIGRP Packet Types

- **Update packets**

- Used to propagate routing information

- **Acknowledgement packets**

- Used to acknowledge receipt of update, query & reply packets



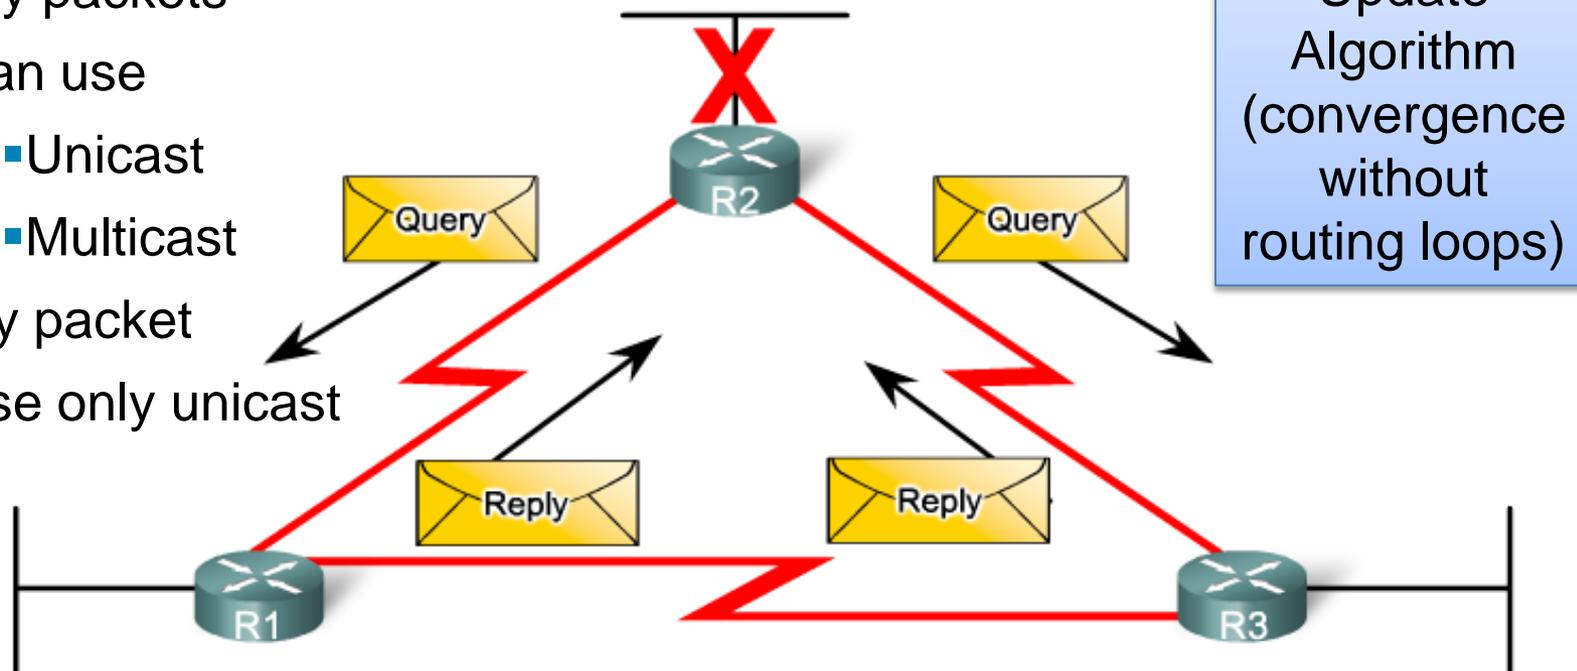
### Update packet

- Used to propagate routing information after a change
- Acknowledgement (ACK) packet
- Automatically sent back when reliable RTP is used

## 9.1.4 RTP and EIGRP Packet Types

### ■ Query & Reply packets

- Used by DUAL for searching for networks
- Query packets
  - Can use
    - Unicast
    - Multicast
- Reply packet
  - Use only unicast



DUAL  
Diffusing  
Update  
Algorithm  
(convergence  
without  
routing loops)

#### Query packet

- Used by DUAL when searching for networks or other tasks

#### Reply packet

- Automatically sent in response to Query packet

#### Acknowledgement (ACK) packet

- Automatically sent back when reliable RTP is used

## 9.1.5 Hello Protocol

*Before any EIGRP packets can be exchanged between routers, EIGRP must first discover its neighbors. EIGRP neighbors are other routers running EIGRP on shared, directly connected networks.*

### ■ Purpose of Hello Protocol

–To discover & establish adjacencies with neighbor routers

### ■ Characteristics of hello protocol

–Time interval for sending hello packet

- Most networks it is every 5 seconds
- Multipoint non broadcast multi-access networks

–Unicast every 60 seconds

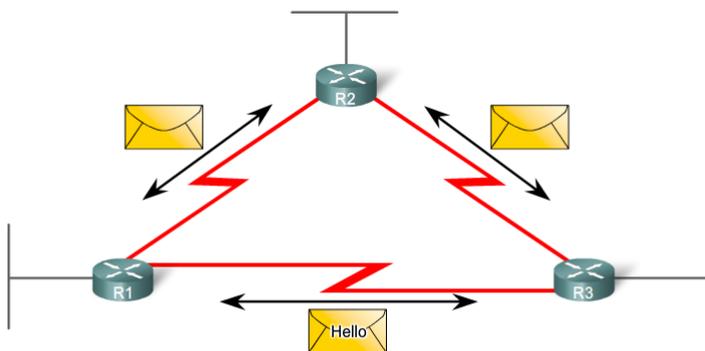
### · Holdtime

- This is the maximum time router should wait before declaring a neighbor down

- Default holdtime

–3 times hello interval

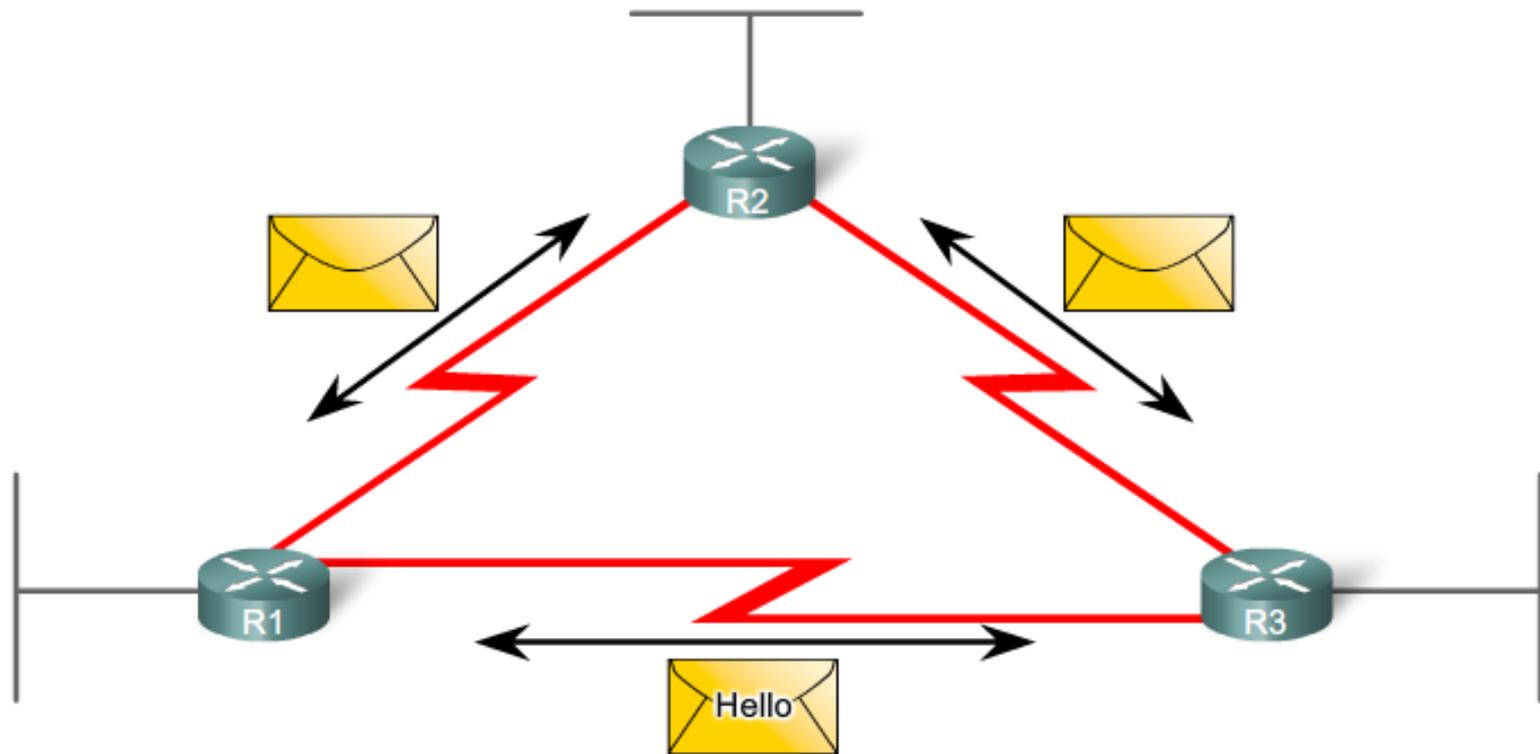
Default Hello Intervals and Hold Times for EIGRP



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

## 9.1.5 Hello Protocol

Default Hello Intervals and Hold Times for EIGRP



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

## 9.1.6 EIGRP Bounded Updates

### EIGRP Bounded Updates

- EIGRP only sends update when there is a change in route status
- **Partial update**
  - A partial update includes only the route information that has changed – the whole routing table is NOT sent
- **Bounded update**
  - When a route changes, only those devices that are impacted will be notified of the change
- EIGRP's use of partial bounded updates minimizes use of bandwidth

#### EIGRP Updates are partial and bounded:

*Partial* because the update only includes information about route changes.

*Bounded* because only those routers affected by the change will receive the update.

## 9.1.7 DUAL UPDATES

*Diffusing Update Algorithm (DUAL) is the convergence algorithm used by EIGRP instead of the Bellman-Ford or Ford Fulkerson algorithms used by other distance vector routing protocols, like RIP*

### Diffusing Update Algorithm (DUAL)

–Purpose

- EIGRP's primary method for preventing routing loops

–Advantage of using DUAL

- Provides for fast convergence time by keeping a list of loop-free backup routes

- Allows all routers involved in a topology change to synchronize at the same time.
- Routers that are not affected by the topology changes are not involved in the recompilation.
- This method provides EIGRP with faster convergence times than other distance vector routing protocols
- Uses a series of updates and acknowledgements

## 9.1.8 Default Administrative Distances

- Administrative Distance (AD)
  - Defined as the trustworthiness of the source route
- EIGRP default administrative distances
  - Summary routes = 5
  - Internal routes = 90
  - Imported routes = 170

*Administrative distance (AD) is the trustworthiness (or preference) of the route source*

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
RIP	120
External EIGRP	170
Internal BGP	200

When compared to other interior gateway protocols (IGPs), EIGRP is the most preferred by the Cisco IOS because it has the lowest administrative distance.

### Interior Gateway Protocols

Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.

### Border Gateway Routing (BGP)

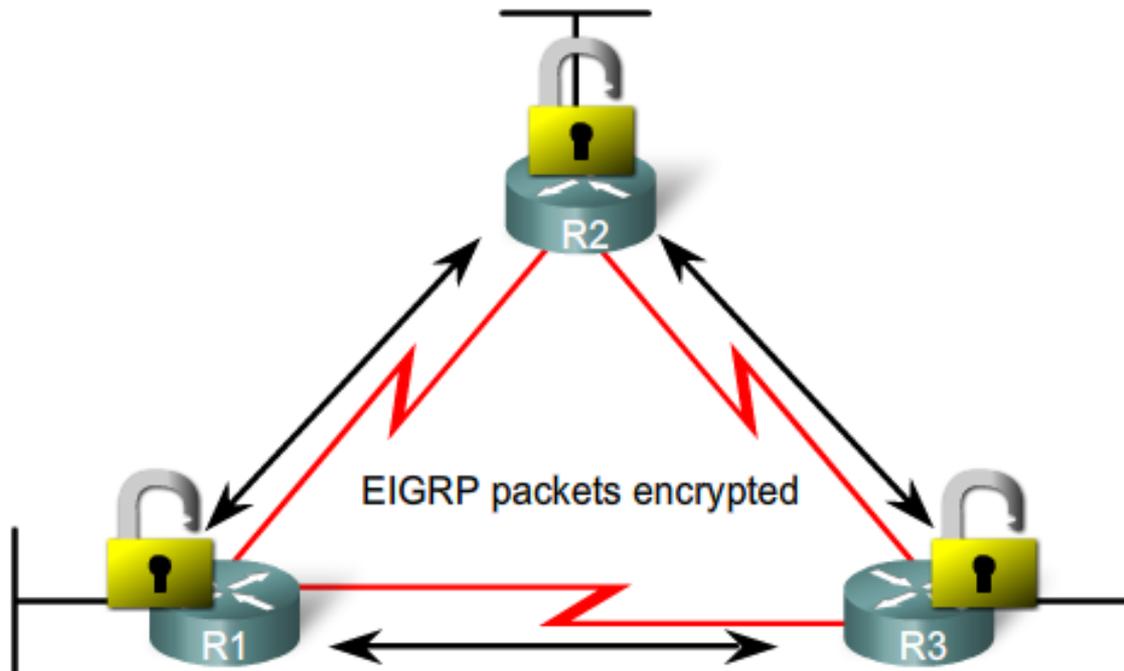
Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

## 9.1.9 Authentication

### Authentication

- EIGRP can
  - Encrypt routing information
  - Authenticate routing information

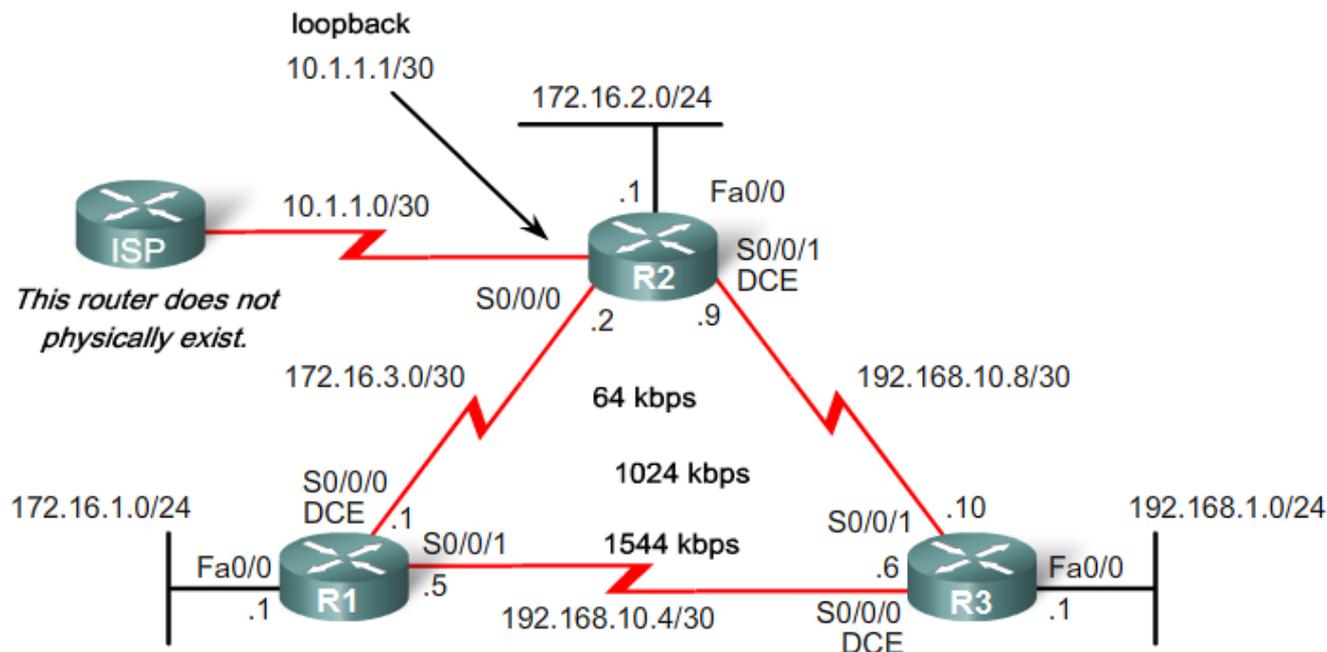
*Covered in a later course*



Ensures that routers will only accept routing information from other routers that have been configured with the same password or authentication information.

Note: Authentication does not encrypt the router's routing table

## 9.2.1 EIGRP Topology

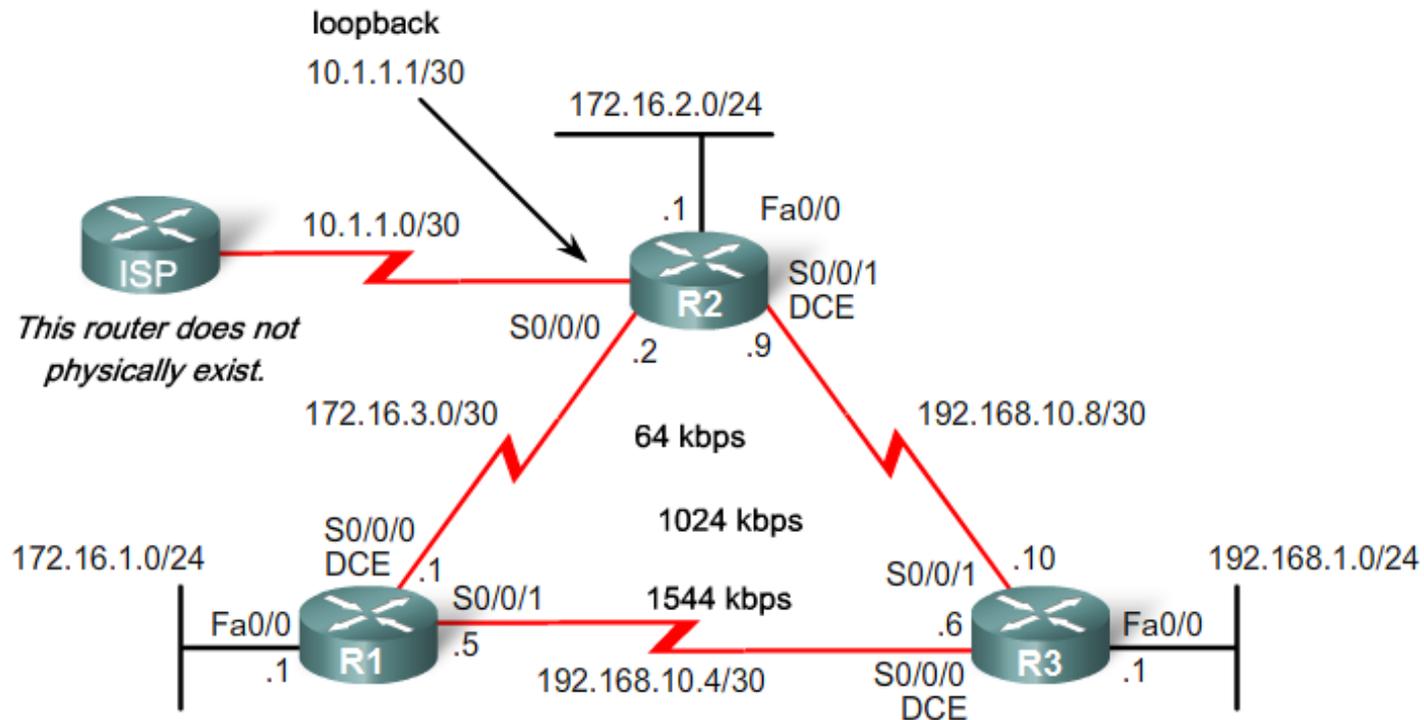


Device	Interface	IP Address	Subnet Mask
R1	Fa0/0	172.16.1.1	255.255.255.0
	S0/0/0	172.16.3.1	255.255.255.252
	S0/0/1	192.168.10.5	255.255.255.252
R2	Fa0/0	172.16.2.1	255.255.255.0
	S0/0/0	172.16.3.2	255.255.255.252
	S0/0/1	192.168.10.9	255.255.255.252
	Lo1	10.1.1.1	255.255.255.252
R3	Fa0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.10.6	255.255.255.252
	S0/0/1	192.168.10.10	255.255.255.252

The figure shows our topology from previous chapters, but now includes the addition of the ISP router. Notice that both the R1 and R2 routers have subnets that are part of the 172.16.0.0/16 classful network, a class B address. The fact that 172.16.0.0 is a class B address is only relevant because **EIGRP automatically summarizes at classful boundaries, similar to RIP.**

## 9.2.1 EIGRP Topology

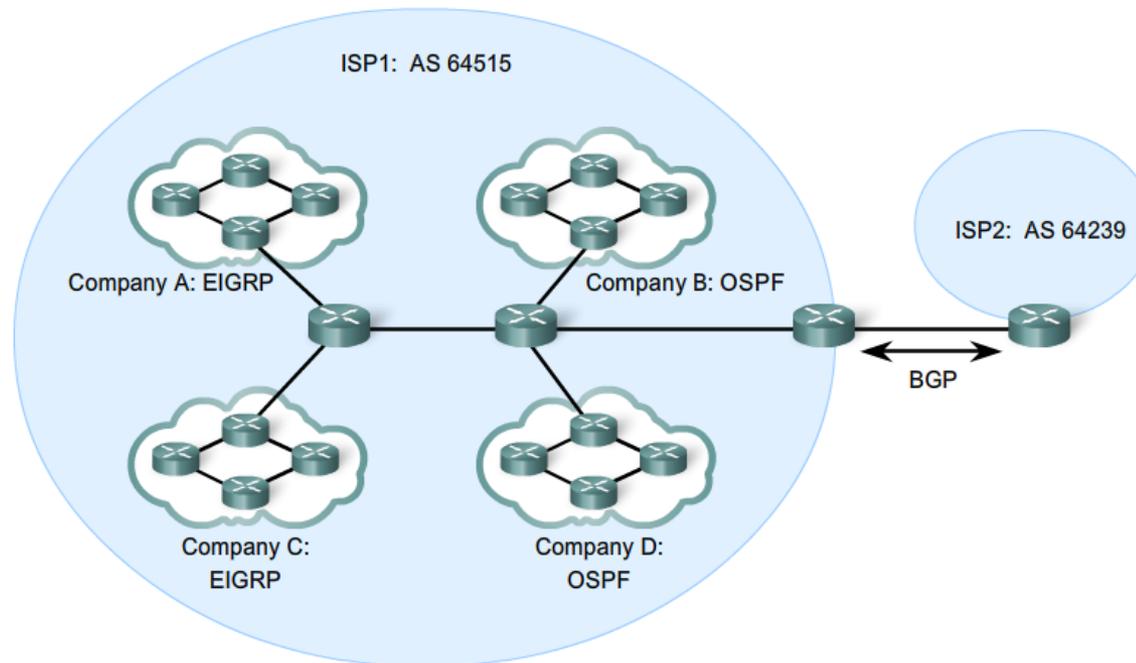
- The ISP router does not physically exist in our configurations. The connection between R2 and ISP is represented with a loopback interface on router R2.
- A loopback interface can be used to represent an interface on a router that does not have any actual connection to a **physical link** on the network.
- Loopback addresses can be verified with the ping command and included in routing updates.



## 9.2.2 Autonomous Systems and Process ID's

- **Autonomous System (AS) & Process IDs**

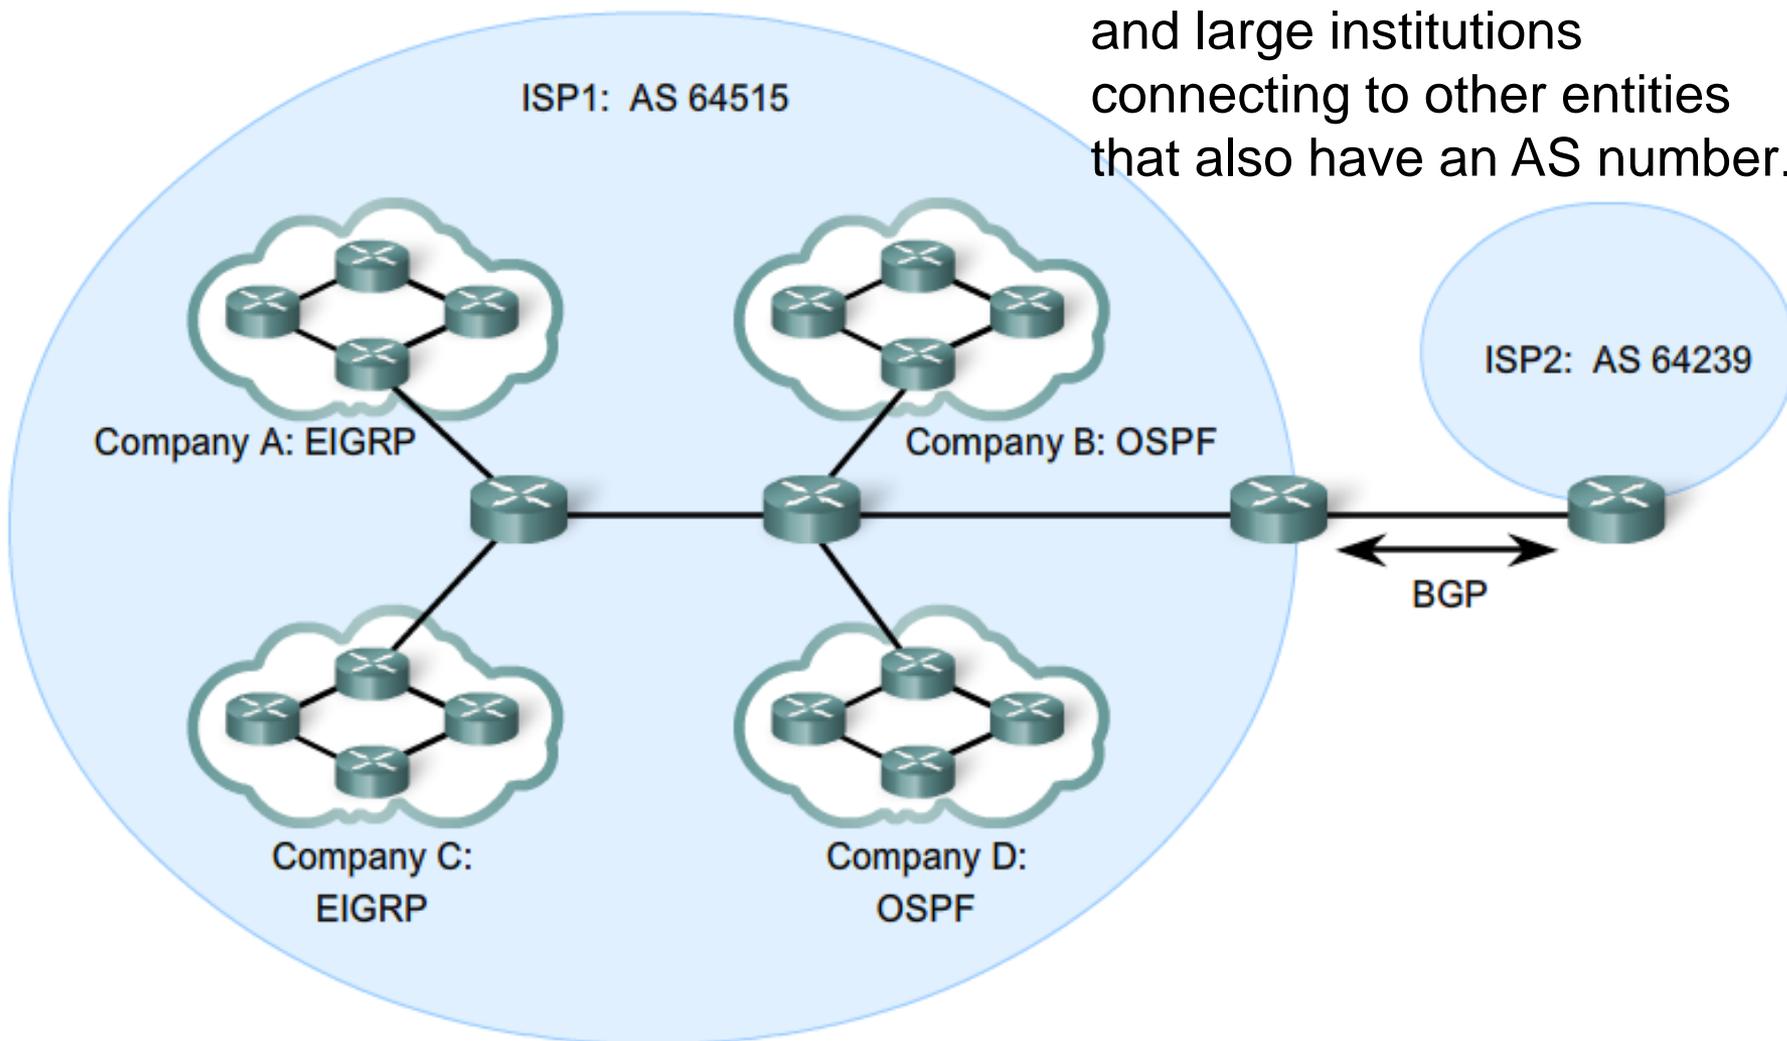
- This is a collection of networks under the control of a single authority
- AS Numbers are assigned by IANA
- Entities needing AS numbers
  - ISP
  - Internet Backbone providers
  - Institutions connecting to other institutions using AS numbers



## 9.2.2 Autonomous Systems and Process ID's

### Autonomous Systems

Usually used by ISPs, Internet backbone providers, and large institutions connecting to other entities that also have an AS number.



## 9.2.2 Autonomous Systems and Process ID's

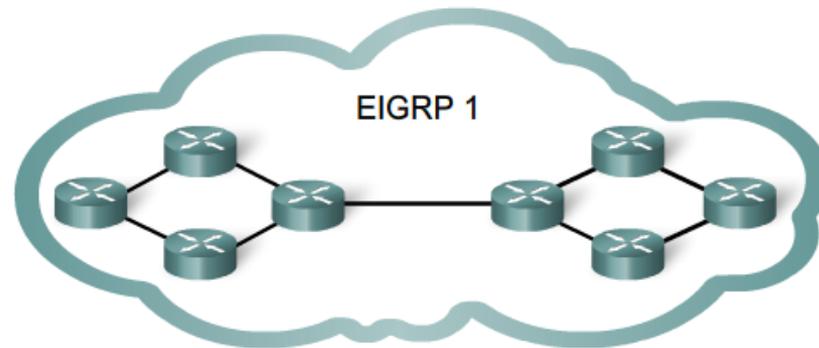
- EIGRP autonomous system number actually functions as a process ID
- Process ID represents an instance of the routing protocol running on a router
- Example

Router(config)#router

**eigrp** *autonomous-system*

EIGRP requires all routers in the same routing domain to be configured with the same process ID

Single Process IDs



```
R1(config)#router eigrp ?  
  <1-65535>  Autonomous system number  
R1(config)#router eigrp 1
```

Although the Cisco IOS refers to the router eigrp parameter as an "Autonomous system number", this parameter configures an EIGRP process—an instance of EIGRP running on the router—and has nothing to do with AS configurations in ISP routers.

## 9.2.3 The Router IGRP Command

The *router eigrp* command

- The global command that enables eigrp is

*router eigrp* *autonomous-system*

-All routers in the EIGRP routing domain **must use the same process ID number** (autonomous-system number)

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

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

```
R3(config)#router eigrp 1  
R3(config-router)#
```

## 9.2.4 The Network Command

### The Network Command

- Functions of the network command
  - Enables interfaces to transmit & receive EIGRP updates
  - Includes network or subnet in EIGRP updates
- Example
  - Router(config-router)#network *network-address*

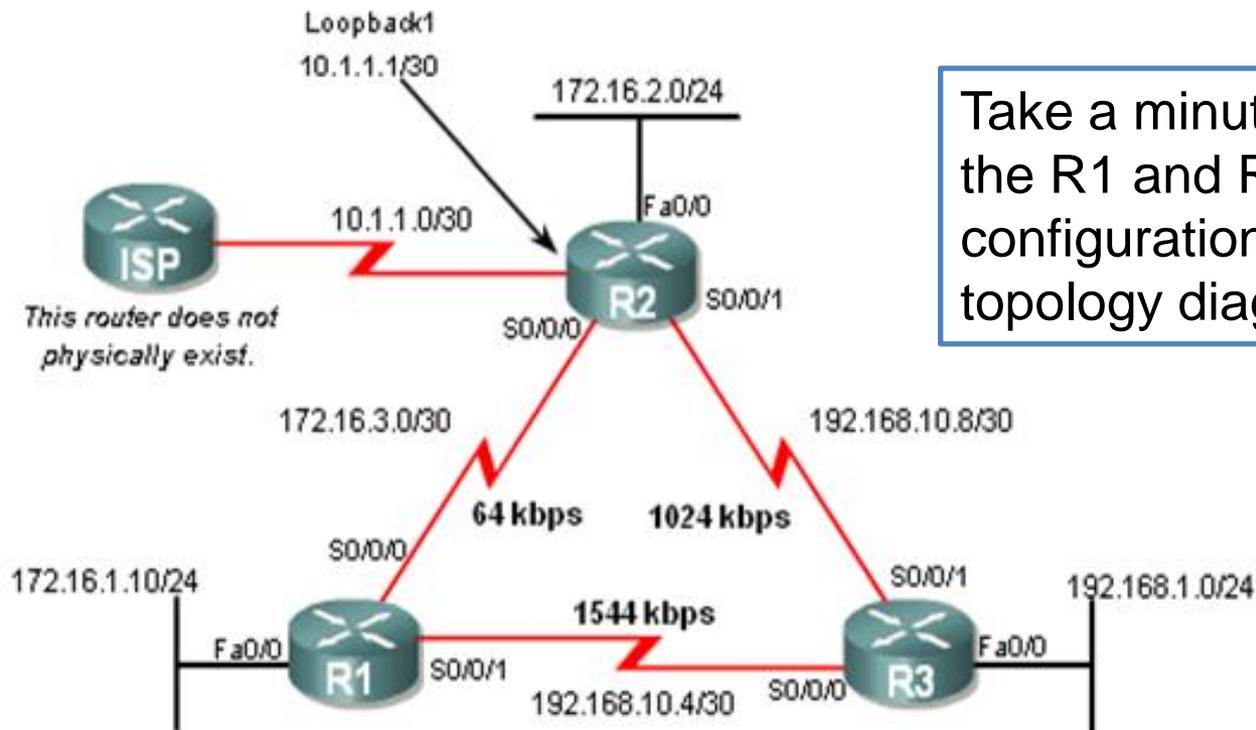
```
R1(config)#router eigrp 1
R1(config-router)#network 172.16.0.0
R1(config-router)#network 192.168.10.0
```

```
R2(config)#router eigrp 1
R2(config-router)#network 172.16.0.0
```

```
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 172.16.3.1 (Serial0/0/0) is up: new adjacency
```

## 9.2.4 The Network Command

### Enabling EIGRP Routing



Take a minute and relate the R1 and R2 configuration to the topology diagram

```
R1(config)#router eigrp 1
R1(config-router)#network 172.16.0.0
R1(config-router)#network 192.168.10.0
```

```
R2(config)#router eigrp 1
R2(config-router)#network 172.16.0.0
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 172.16.3.1 (Serial0/0/0) is up: new adjacency
```

## 9.2.4 The Network Command

- The network Command with a Wildcard Mask

- This option is used when you want to configure EIGRP to advertise specific subnets

- Example

Router(config-router)#network network-address [wildcard-mask]

```
R1(config)#router eigrp 1
R1(config-router)#network 172.16.0.0
R1(config-router)#network 192.168.10.0
```

```
R2(config)#router eigrp 1
R2(config-router)#network 172.16.0.0
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 172.16.3.1 (Serial0/0/0) is up: new adjacency
R2(config-router)#network 192.168.10.8 0.0.0.3
```

```
R3(config)#router eigrp 1
R3(config-router)#network 192.168.10.0
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 192.168.10.5 (Serial0/0/0) is up: new adjacency
R3(config-router)#
%DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 192.168.10.9 (Serial0/0/1) is up: new adjacency
R3(config-router)#network 192.168.1.0
```

## 9.2.4 The Network Command

By default, when using the network command and a classful network address such as 172.16.0.0, all interfaces on the router that belong to that classful network address will be enabled for EIGRP. However, there may be times when the network administrator does not want to include all interfaces within a network when enabling EIGRP. To configure EIGRP to advertise specific subnets only, use the wildcard-mask option with the network command:

```
Router(config-router)#network network-address [wildcard-mask]
```

Think of a wildcard mask as the inverse of a subnet mask. The inverse of subnet mask 255.255.255.252 is 0.0.0.3. To calculate the inverse of the subnet mask, subtract the subnet mask from 255.255.255.255:

255.255.255.255

- 255.255.255.252

Subtract the subnet mask

-----

0. 0. 0. 3

Wildcard mask

## 9.2.5 Verifying EIGRP

### Verifying EIGRP

- EIGRP routers must establish adjacencies with their neighbors before any updates can be sent or received
- Command used to view neighbor table and verify that EIGRP has established adjacencies with neighbors is

***show ip eigrp neighbors***

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address          Interface    Hold Uptime   SRTT  RTO  Q  Seq Type
   Address          Interface    (sec) (sec) (ms)  (ms) Cnt Num
1   192.168.10.10    Se0/0/1      10   00:01:41  20    200  0   7
0   172.16.3.1       Se0/0/0      10   00:09:49  25    200  0  28
```

Address of neighbors

Interface connected to neighbor

Amount of time left before neighbor is considered "down"

Amount of time since adjacency was established

## 9.2.5 Verifying EIGRP

- The ***show ip protocols*** command is also used to verify that EIGRP is enabled

```
R1#show ip protocols
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 metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  Automatic network summarization is in effect
  Automatic address summarization:
    192.168.10.0/24 for FastEthernet0/0, Serial0/0/0
      Summarizing with metric 2169856
    172.16.0.0/16 for Serial0/0/1
      Summarizing with metric 28160
  Maximum path: 4
  Routing for Networks:
    172.16.0.0
    192.168.10.0
  Routing Information Sources:
  Gateway         Distance      Last Update
  (this router)   90           00:03:29
  192.168.10.6    90           00:02:09
  Gateway         Distance      Last Update
  172.16.3.2      90           00:02:12
  Distance: internal 90 external 170
```

Notice that the output specifies the process ID used by EIGRP:

- Routing Protocol is "eigrp 1"

Remember, the process ID must be the same on all routers for EIGRP to establish neighbor adjacencies and share routing information.

EIGRP's internal and external administrative distances are also displayed:

- Distance: internal 90 external 170

## 9.2.6 Examining the Routing Table

### Examining the Routing Table

- The **show ip route** command is also used to verify EIGRP
- EIGRP routes are denoted in a routing table by the letter “D”
- By default , EIGRP automatically summarizes routes at major network boundary

```
R1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
<Output omitted>

Gateway of last resort is not set

    192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
D       192.168.10.0/24 is a summary, 00:03:50, Null0
C       192.168.10.4/30 is directly connected, Serial0/0/1
D       192.168.10.8/30 [90/2681856] via 192.168.10.6, 00:02:43, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
D       172.16.0.0/16 is a summary, 00:10:52, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
D       172.16.2.0/24 [90/2172416] via 172.16.3.2, 00:10:47, Serial0/0/0
C       172.16.3.0/30 is directly connected, Serial0/0/0
D       192.168.1.0/24 [90/2172416] via 192.168.10.6, 00:02:31, Serial0/0/1
```

## 9.2.6 Examining the Routing Table

### ■ Introducing the Null0 Summary Route

- Null0 is not a physical interface
- In the routing table summary routes are sourced from Null0
  - Reason: routes are used for advertisement purposes
- EIGRP will automatically include a null0 summary route as child route when 2 conditions are met
  - At least one subnet is learned via EIGRP
  - Automatic summarization is enabled

```
R2#show ip route

Gateway of last resort is not set

    192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
D    192.168.10.0/24 is a summary, 00:04:13, Null0
D    192.168.10.4/30 [90/2681856] via 192.168.10.10, 00:03:05, Serial0/0/1
C    192.168.10.8/30 is directly connected, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
D    172.16.0.0/16 is a summary, 00:04:07, Null0
D    172.16.1.0/24 [90/2172416] via 172.16.3.1, 00:11:11, Serial0/0/0
C    172.16.2.0/24 is directly connected, FastEthernet0/0
C    172.16.3.0/30 is directly connected, Serial0/0/0
    10.0.0.0/30 is subnetted, 1 subnets
C    10.1.1.0 is directly connected, Loopback1
D    192.168.1.0/24 [90/2172416] via 192.168.10.10, 00:02:54, Serial0/0/1
```

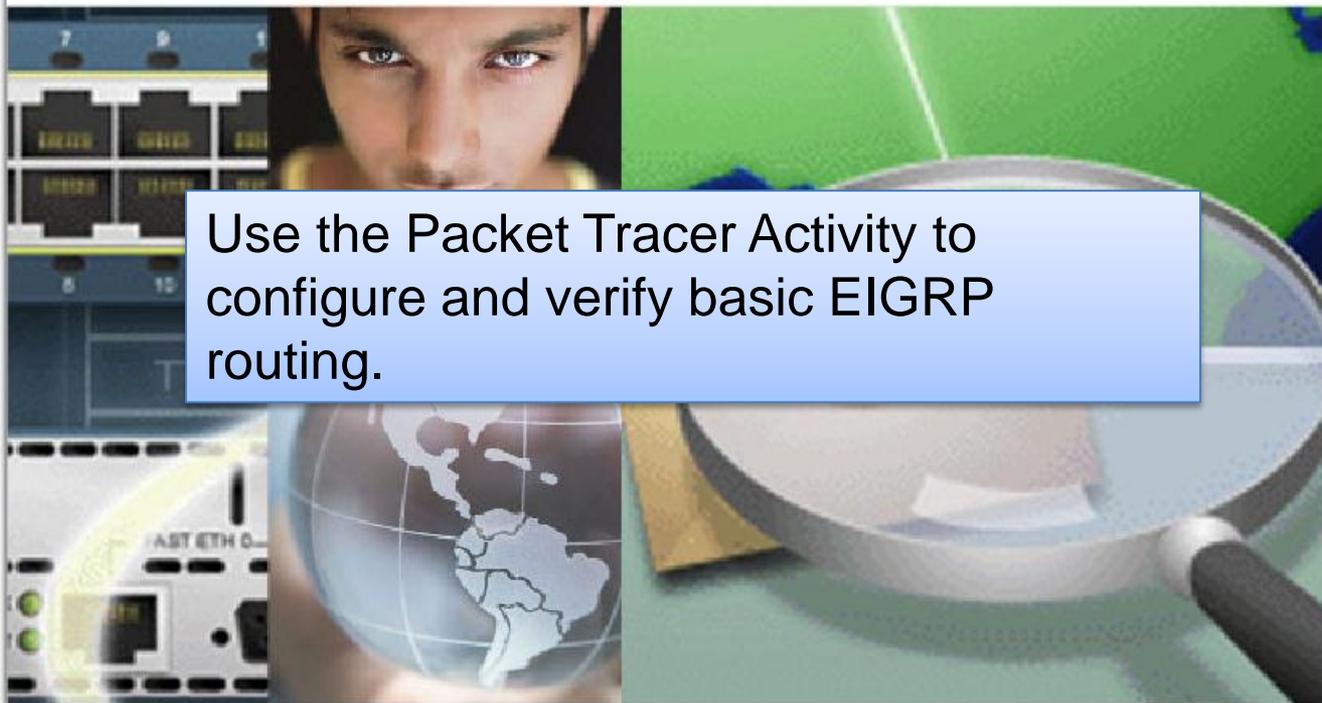
Summary Routes to Null0

## 9.2.6 Examining the Routing Table



### Packet Tracer Exploration: Configure and Verify EIGRP Routing

Use the Packet Tracer Activity to  
configure and verify basic EIGRP  
routing.



## 9.3.1 EIGRP Composite Metric and K Values

### EIGRP Composite Metric & the K Values

- EIGRP uses the following values in its composite metric
  - Bandwidth, delay, reliability, and load
- The composite metric used by EIGRP
  - formula used has values K1 →K5
    - K1 & K3 = 1
    - all other K value = 0

Default Composite Formula:

$\text{metric} = [K1 * \text{bandwidth} + K3 * \text{delay}]$

Complete Composite Formula:

$\text{metric} = [K1 * \text{bandwidth} + (K2 * \text{bandwidth}) / (256 - \text{load}) + K3 * \text{delay}] * [K5 / (\text{reliability} + K4)]$

(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
```

## 9.3.1 EIGRP Composite Metric and K Values

Use the **show ip protocols** command to verify the K values

### R1 K Values

```
R1#show ip protocols
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 metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  Automatic network summarization is in effect
  Automatic address summarization:
    192.168.10.0/24 for FastEthernet0/0, Serial10/0/0
```

*Changing these values to other than the default is not recommended unless the network administrator has a very good reason to do so.*

## 9.3.2 EIGRP Metrics

### EIGRP Metrics

- Use the **show interfaces** command to view metrics
- EIGRP Metrics

**Bandwidth** – EIGRP uses a static bandwidth to calculate metric

Most serial interfaces use a default bandwidth value of 1.544Mbos (T1)

```
R1#show interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to R2
  Internet address is 172.16.3.1/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation HDLC, loopback not set
  Keepalive set (10 sec)
  Last input 00:00:00, output 00:00:01, output hang never
  Last clearing of "show interface" counters 3d22h
```

## 9.3.2 EIGRP Metrics

### EIGRP Metrics

- **Delay** is defined as the measure of time it takes for a packet to traverse a route
  - it is a static value based on link type to which interface is connected

Media	Delay
100M ATM	100 $\mu$ s
Fast Ethernet	100 $\mu$ s
FDDI	100 $\mu$ s
1HSSI	20,000 $\mu$ s
16M Token Ring	630 $\mu$ s
Ethernet	1,000 $\mu$ s
T1 (Serial Default)	20,000 $\mu$ s
512K	20,000 $\mu$ s
DSO	20,000 $\mu$ s
56K	20,000 $\mu$ s

Delay Values  
in  
Microseconds

## 9.3.2 EIGRP Metrics

- **Reliability** (not a default EIGRP metric)
  - A measure of the likelihood that a link will fail
  - Measure dynamically & expressed as a fraction of 255  
the higher the fraction the better the reliability
- **Load** (not a default EIGRP metric)
  - A number that reflects how much traffic is using a link
  - Number is determined dynamically and is expressed as a fraction of 255
    - The lower the fraction the less the load on the link

```
MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
```

```
reliability 255/255, txload 1/255, rxload 1/255
```

Reliability Value

Load Value

## 9.3.3 Using the Bandwidth Command

### Using the Bandwidth Command

- Modifying the interface bandwidth
  - Use the **bandwidth** command
  - Example  
Router(config-if)#**bandwidth** kilobits
- Verifying bandwidth
  - Use the **show interface** command
- Note – bandwidth command **does not change** the link's physical bandwidth

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

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

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

## 9.3.4 Calculating the EIGRP Metric

The EIGRP metric can be determined by examining the bandwidth delay

### Calculating the EIGRP Default Metric

Default metric =  $[K1 * \text{bandwidth} + K3 * \text{delay}] * 256$

Since K1 and K3 both equal 1, the formula simplifies to: **bandwidth + delay**

bandwidth = speed of slowest link in route to the destination

delay = sum of the delays of each link in route to the destination

Slowest bandwidth:  $(10,000,000 / \text{bandwidth kbps}) * 256$

Plus the sum of the delays:  $+ (\text{sum of delay} / 10) * 256$

= **EIGRP metric**

```
R2#show ip route
```

```
D    192.168.1.0/24 [90/3014400] via 192.168.10.10, 00:02:14, Serial0/0/1
```



## 9.3.4 Calculating the EIGRP Metric

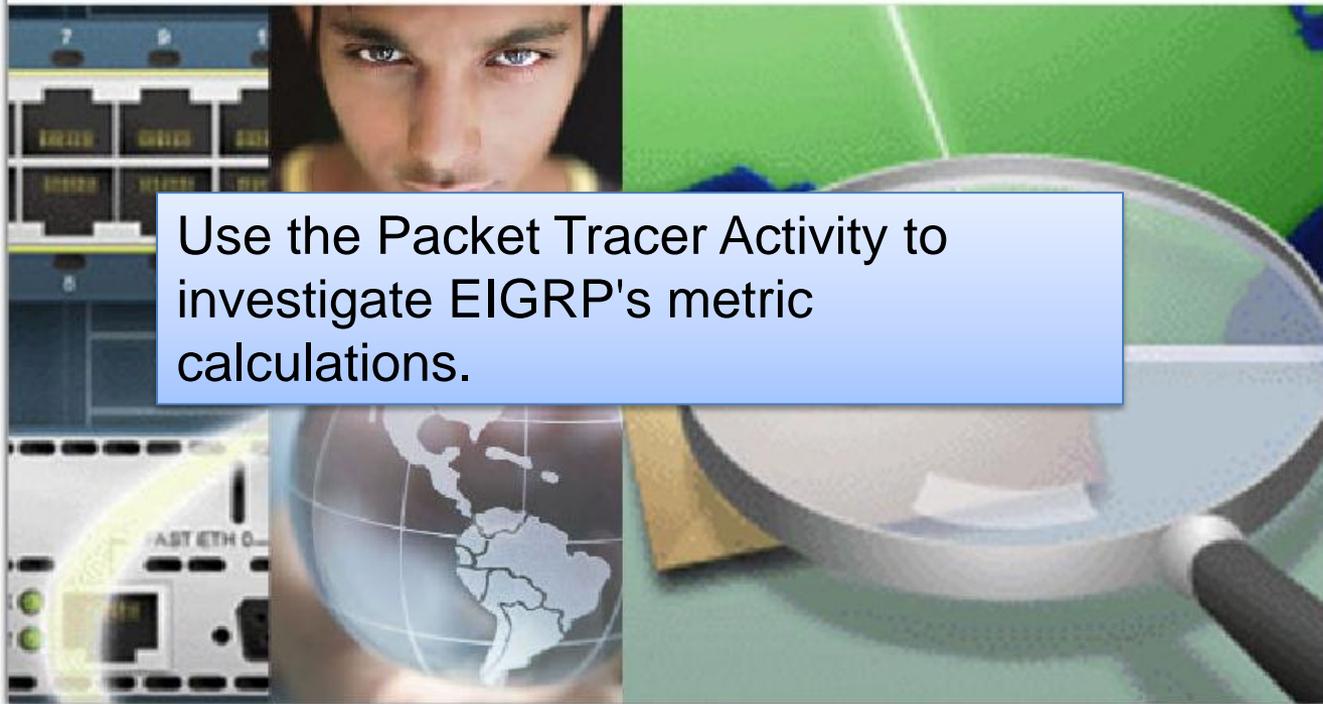
- EIGRP uses the lowest bandwidth (BW) in its metric calculation
  - Calculated BW = reference BW / lowest BW(kbps)
- Delay – EIGRP uses the cumulative sum of all outgoing interfaces
  - Calculated Delay = the sum of outgoing interface delays
- EIGRP Metric = calculated BW + calculated delay

## 9.3.4 Calculating the EIGRP Metric



Packet Tracer Exploration:  
Calculating the EIGRP metric

Use the Packet Tracer Activity to investigate EIGRP's metric calculations.



## 9.4.1 DUAL Concepts

The **D**iffusing **U**ppdate **A**lgorithm (DUAL) is used to prevent looping

### DUAL Concepts

#### DUAL provides:

- Loop-free paths
- Loop-free backup paths which can be used immediately
- Fast convergence
- Minimum bandwidth usage with bounded updates

## 9.4.2 Successor and Feasible Distance

- **Successor**

The **best least cost route** to a destination found in the routing table

- **Feasible distance**

The **lowest calculated metric** along a path to a destination network

```
R2#show ip route
***output omitted***

Gateway of last resort is not set

  192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
D    192.168.10.0/24 is a summary, 00:00:15, Null0
D    192.168.10.4/30 [90/21024000] via 192.168.10.10, 00:00:15, Serial0/0/1
C    192.168.10.8/30 is directly connected, Serial0/0/1
  172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
D    172.16.0.0/16 is a summary, 00:00:15, Null0
D    172.16.1.0/24 [90/40514560] via 172.16.3.1, 00:00:15, Serial0/0/0
C    172.16.2.0/24 is directly connected, FastEthernet0/0
C    172.16.3.0/30 is directly connected, Serial0/0/0
  10.0.0.0/30 is subnetted, 1 subnets
C    10.1.1.0 is directly connected, Loopback1
D    192.168.1.0/24 [90/3014400] via 192.168.10.10, 00:00:15, Serial0/0/1
```

Feasible Distance

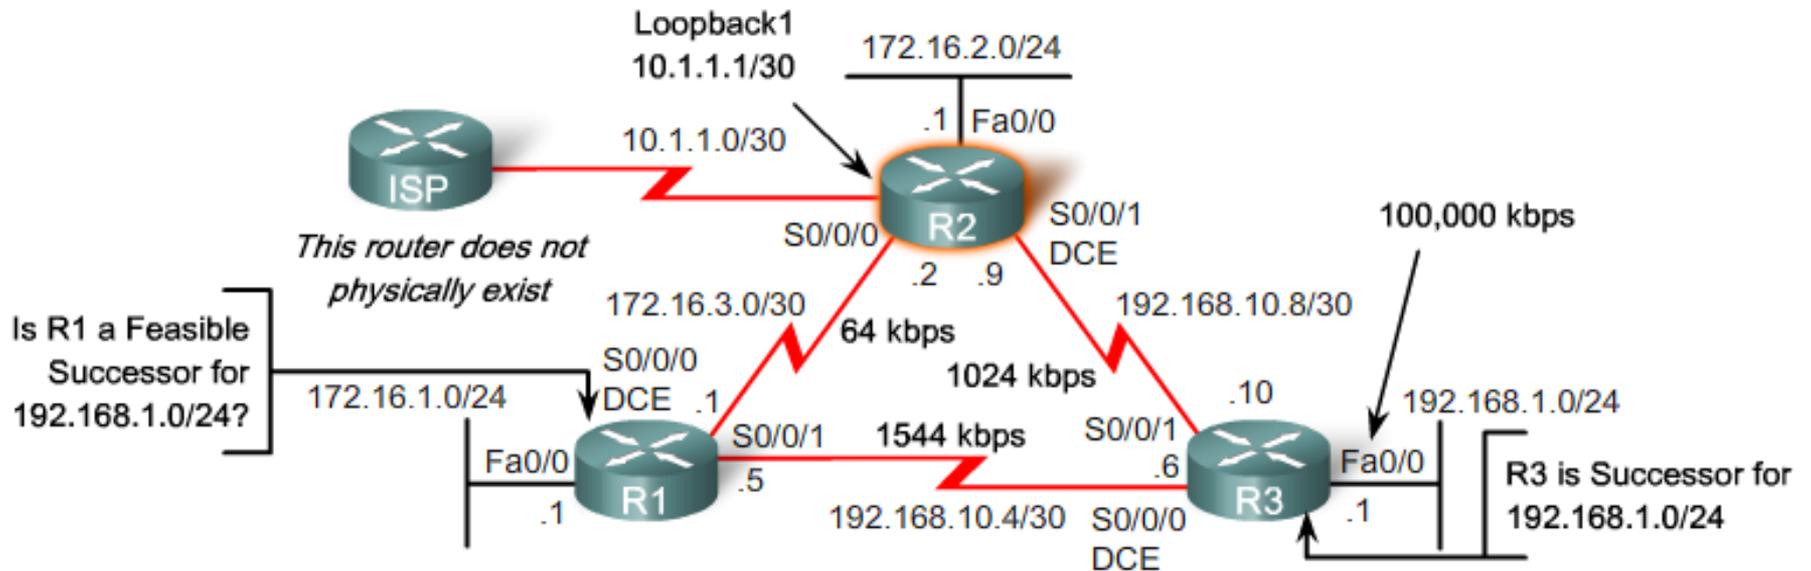
Successor

R3 at 192.168.10.10 is the successor for network 192.168.1.0/24. This route has a feasible distance of 3014400.



# 9.4.3 Feasible Successors, Feasibility Condition and Reported Distance

## Finding the Feasible Successor



**Feasible Successor**  
This is a loop free backup route to same destination as successor route

```
172.16.0.0/16 is variably subnetted, 4 subnets, 3 hosts
D   172.16.0.0/16 is a summary, 00:00:15, Null0
D   172.16.1.0/24 [90/40514560] via 172.16.3.1, 00:00:15, Serial0/0/0
C   172.16.2.0/24 is directly connected, FastEthernet0/0
C   172.16.3.0/30 is directly connected, Serial0/0/0
10.0.0.0/30 is subnetted, 1 subnets
C   10.1.1.0 is directly connected, Loopback1
D   192.168.1.0/24 [90/3014400] via 192.168.10.10, 00:00:15, Serial0/0/1
```

## 9.4.4 Topology Table: Successor and Feasible Successor

- The successor, feasible distance, and any feasible successors with their reported distances are kept by a router in its EIGRP topology table or topology database.
- As shown in the figure, the topology table can be viewed using the **show ip eigrp topology** command.
- The topology table lists all successors and feasible successors that DUAL has calculated to destination networks.

```
R2#show ip eigrp topology
IP-EIGRP Topology Table for AS(1)/ID(10.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

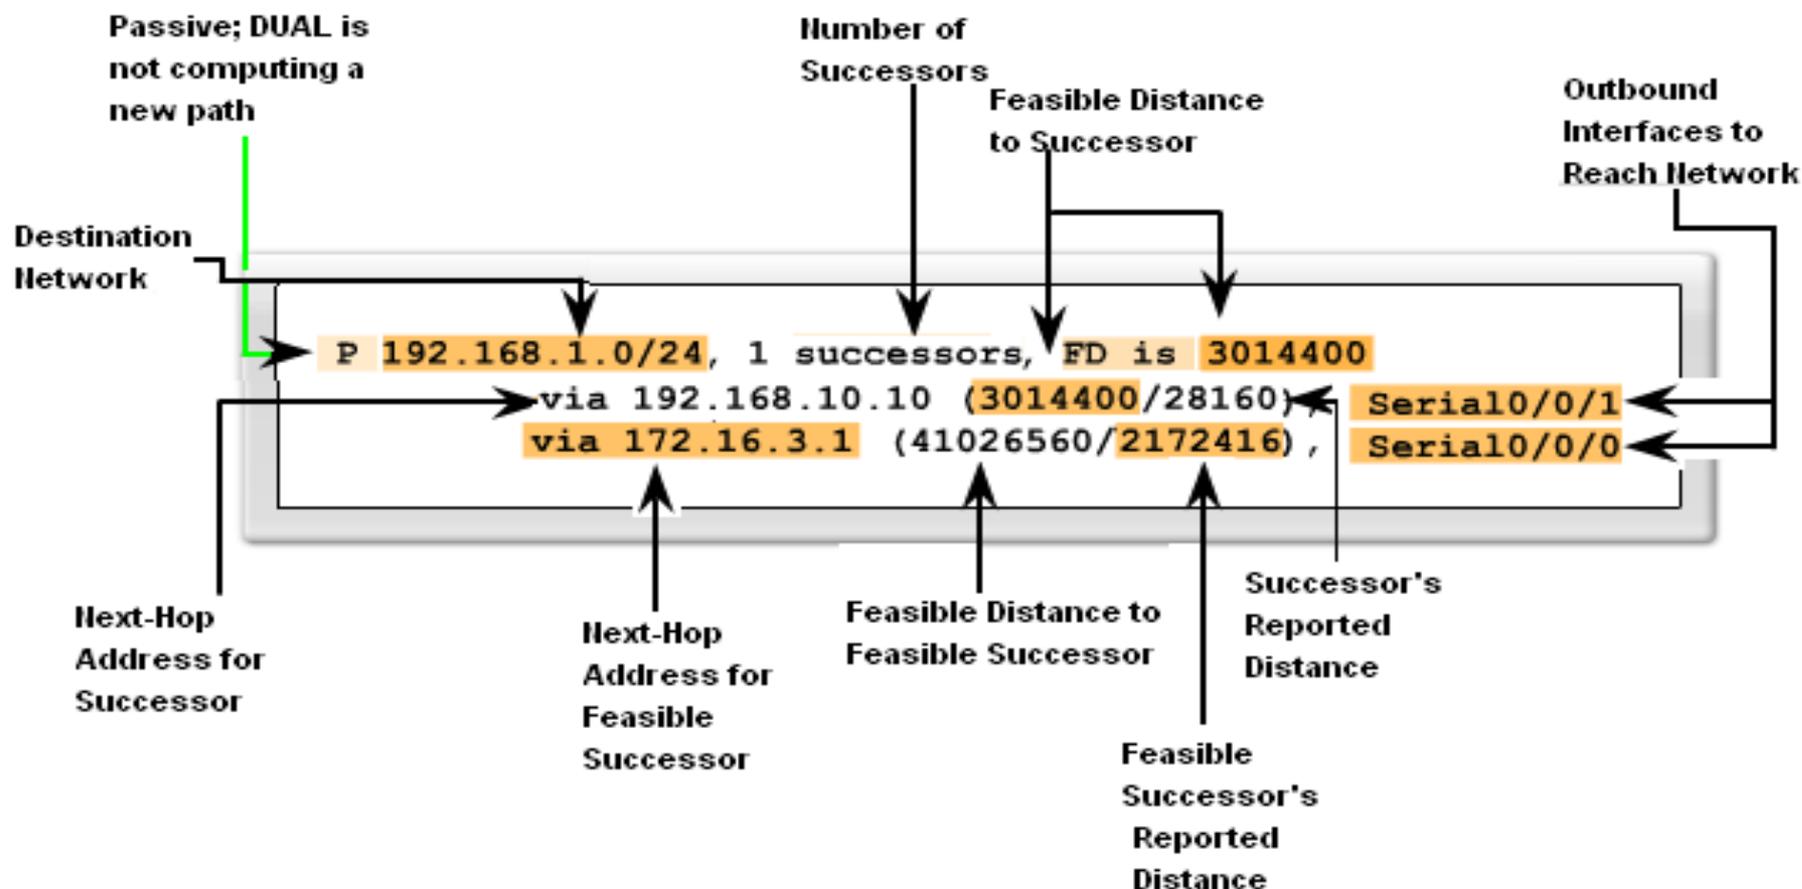
P 192.168.1.0/24, 1 successors, FD is 3014400
   via 192.168.10.10 (3014400/28160), Serial0/0/1
   via 172.16.3.1 (41026560/2172416), Serial0/0/0
P 192.168.10.8/30, 1 successors, FD is 3011840
   via Connected, Serial0/1
```

Routes for 192.168.1.0/24

## 9.4.4 Topology Table: Successor and Feasible Successor

### EIGRP Topology Table dissected

Table Entry for 192.168.1.0/24



## 9.4.4 Topology Table: Successor and Feasible Successor

To view detailed information about the metrics of a specific entry in the topology table, add the optional parameter [network] to the show ip eigrp topology command, as shown in the figure:

```
R2#show ip eigrp topology 192.168.1.0
```

Remember that EIGRP is a distance vector routing protocol. This command lists the full list of distance vector metrics available to EIGRP even though, by default, EIGRP only uses bandwidth and delay. It also displays other information included in the routing update, but not included in the composite metric: minimum MTU and hop count.

### Topology Table: **No Feasible Successor**

- A feasible successor may not be present because the feasibility condition may not be met
  - In other words, the reported distance of the neighbor is greater than or equal to the current feasible distance

## 9.4.4 Topology Table: Successor and Feasible Successor

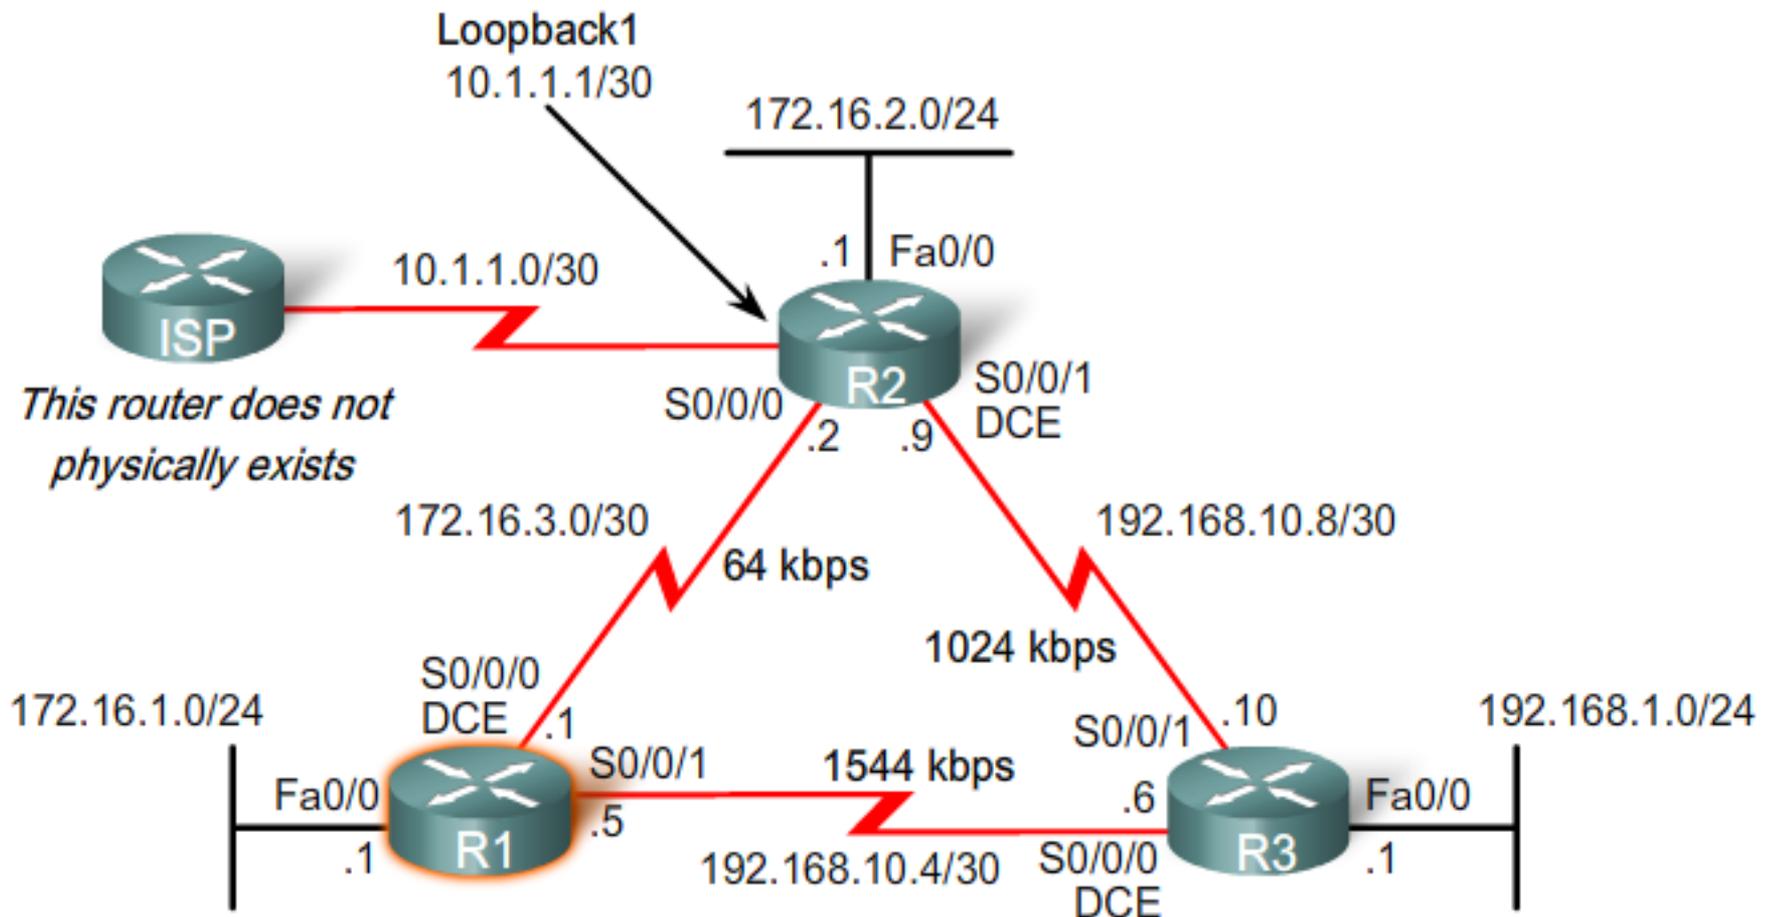
- **Successor**

The **best least cost route** to a destination found in the routing table

- **Feasible distance**

The **lowest calculated metric** along a path to a destination network

## 9.4.5 Topology Table: No Feasible Successor



```
D 192.168.1.0/24 [90/2172416] via 192.168.10.6, 00:44:55, Serial10/0/1
```

The route to 192.168.1.0/24 shows that the successor is R3 via 192.168.10.6 with a feasible distance of 2172416.

## 9.4.6 Finite State Machine

The centerpiece of EIGRP is DUAL and its EIGRP route-calculation engine. The actual name of this technology is DUAL Finite State Machine (FSM). This finite state machine contains all of the logic used to calculate and compare routes in an EIGRP network.

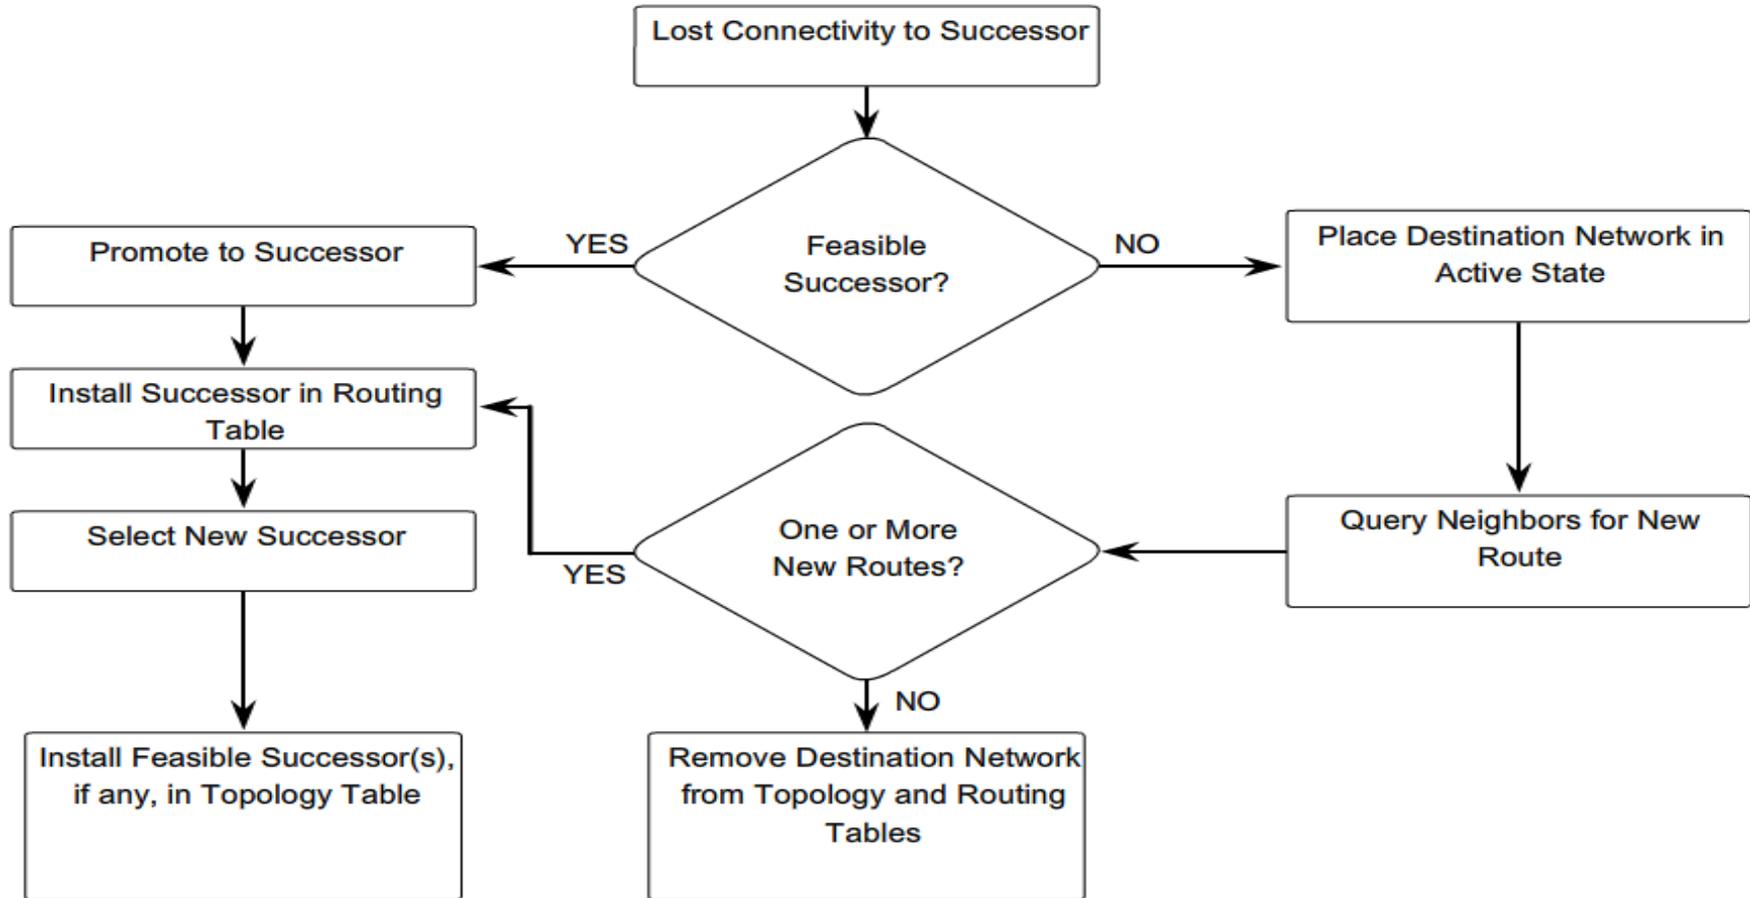
The figure on the next slide shows a simplified version of the DUAL FSM.

### Finite State Machine (FSM)

- **An abstract machine** that defines a set of possible states something can go through, what event causes those states and what events result from those states
- **FSMs are used to** describe how a device, computer program, or routing algorithm will react to a set of input events

## 9.4.6 Finite State Machine

### DUAL Finite State Machine



#### ■ DUAL FSM

- Selects a best loop-free path to a destination
- Selects alternate routes by using information in EIGRP tables

## 9.4.6 Finite State Machine

### Finite State Machines (FSM)

- To examine output from EIGRP's finite state machine us the ***debug eigrp fsm*** command

```
R2#debug eigrp fsm
EIGRP FSM Events/Actions debugging is on
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#int s0/0/1
R2(config-if)#shutdown
<some debug output omitted>

DUAL: Find FS for dest 192.168.1.0/24. FD is 3014400, RD is 3014400
DUAL: 192.168.10.10 metric 4294967295/4294967295
DUAL: 172.16.3.1 metric 41026560/2172416 found Dmin is 41026560
DUAL: Removing dest 192.168.1.0/24, nexthop 192.168.10.10
DUAL: RT installed 192.168.1.0/24 via 172.16.3.1

R2(config-if)#end
R2#undebug all
All possible debugging has been turned off

R2#show ip route
<some output omitted>

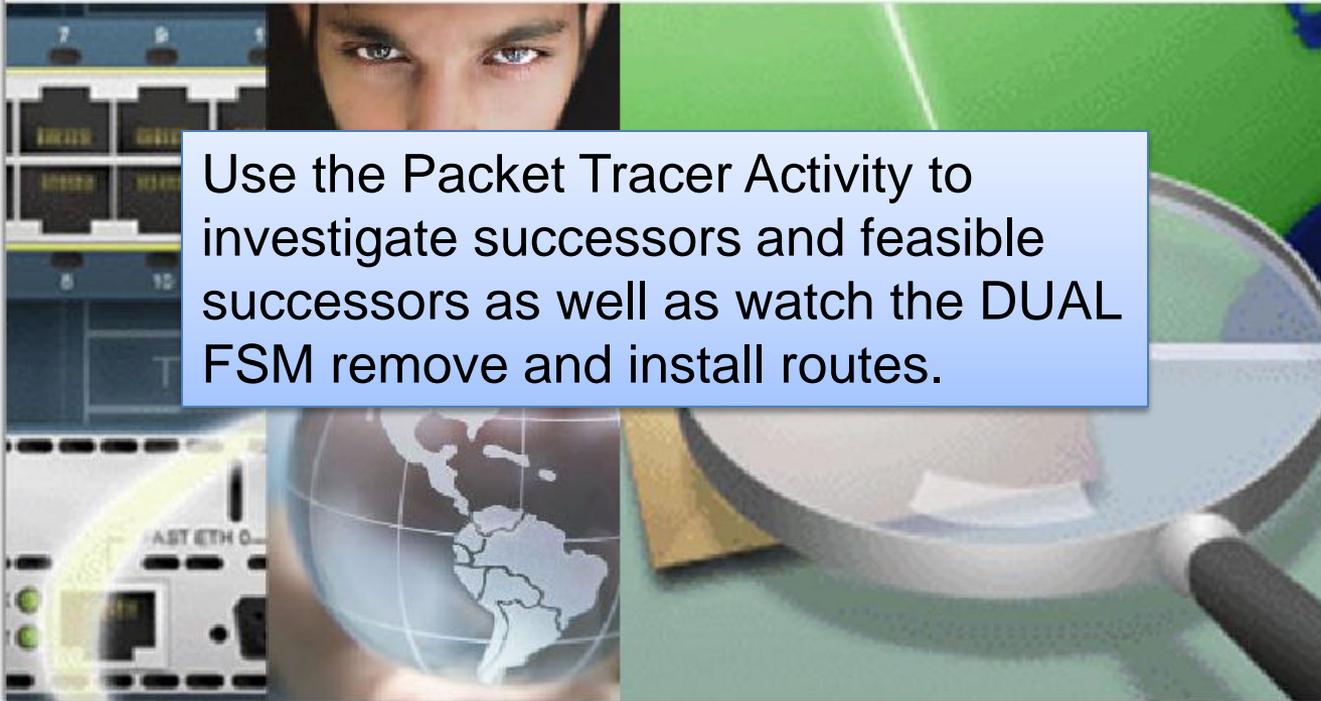
D    192.168.1.0/24 [ 90/41026560]  via 172.16.3.1, 00:08:58, Serial0/0
```

## 9.4.6 Finite State Machine



### Packet Tracer Exploration: Finite State Machine

Use the Packet Tracer Activity to investigate successors and feasible successors as well as watch the DUAL FSM remove and install routes.



## 9.5.1 The Null0 Summary Route

### The Null0 Summary Route

- **By default**, EIGRP uses the Null0 interface to discard any packets that match the parent route but do not match any of the child routes
- EIGRP automatically includes a null0 summary route as a child route whenever both of the following conditions exist
  - One or subnets exists that was learned via EIGRP
  - Automatic summarization is enabled. The Null0 interface is simply a route to nowhere, commonly known as "the bit bucket."

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

 192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
D    192.168.10.0/24 is a summary, 00:45:09, Null0
C    192.168.10.4/30 is directly connected, Serial0/0/1
D    192.168.10.8/30 [90/3523840] via 192.168.10.6, 00:44:56, Serial0/0/1
 172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
D    172.16.0.0/16 is a summary, 00:46:10, Null0
C    172.16.1.0/24 is directly connected, FastEthernet0/0
D    172.16.2.0/24 [90/40514560] via 172.16.3.2, 00:45:09, Serial0/0/0
C    172.16.3.0/30 is directly connected, Serial0/0/0
D    192.168.1.0/24 [90/2172416] via 192.168.10.6, 00:44:55, Serial0/0/1
```

**EIGRP installs a Null0 summary route for each parent route.  
Packets matching the Null0 summary route are discarded.**

## 9.5.2 Disabling Automatic Summarization

### Disabling Automatic Summarization

- The **auto-summary** command permits EIGRP to automatically summarize at major network boundaries
- The **no auto-summary** command is used to disable automatic summarization
  - This causes all EIGRP neighbors to send updates that will not be automatically summarized
    - this will cause changes to appear in both
      - routing tables
      - topology tables

```
R3#show ip route
```

```
    192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
```

```
D       192.168.10.0/24 is a summary, 01:08:35, Null0
```

```
C       192.168.10.4/30 is directly connected, Serial0/0/0
```

```
C       192.168.10.8/30 is directly connected, Serial0/0/1
```

```
D       172.16.0.0/16 [90/2172416] via 192.168.10.5, 01:08:30, Serial0/0/0
```

```
C       192.168.1.0/24 is directly connected, FastEthernet0/0
```

R3 uses the suboptimal route through R1 to reach 172.16.2.0.

## 9.5.3 Manual Summarization

EIGRP can be configured to summarize routes, whether or not automatic summarization (auto-summary) is enabled. Because EIGRP is a classless routing protocol and includes the subnet mask in the routing updates, manual summarization can include supernet routes. Remember, a supernet is an aggregation of multiple major classful network addresses.

### Manual Summarization

- Manual summarization can include supernets

Reason: EIGRP is a classless routing protocol & include subnet mask in update

- Command used to configure manual summarization

–Router(config-if)#ip summary-address eigrp as-number  
network-address subnet-mask

## 9.5.3 Manual Summarization

```
R3(config)#inter loopback 2
R3(config-if)#ip address 192.168.2.1 255.255.255.0
R3(config-if)#interface loopback 3
R3(config-if)#ip address 192.168.3.1 255.255.255.0
R3(config-if)#router eigrp 1
R3(config-router)#network 192.168.2.0
R3(config-router)#network 192.168.3.0
```

Two simulated LANs are added to R3 using loopback interfaces.  
The networks are then configured as part of the EIGRP process.

Suppose we added two more networks to router R3 using loopback interfaces: 192.168.2.0/24 and 192.168.3.0/24. We also configure networks in R3's EIGRP routing process with network commands so that R3 will propagate these networks to other routers.

## 9.5.3 Manual Summarization

```
192.168.1.0: 11000000 . 10101000 . 00000001 . 00000000
192.168.2.0: 11000000 . 10101000 . 00000010 . 00000000
192.168.3.0: 11000000 . 10101000 . 00000011 . 00000000
```

← Loopback1 →

```
R3(config)#interface serial 0/0/0
R3(config-if)#ip summary-address eigrp 1 192.168.0.0 255.255.252.0
R3(config-if)#interface serial 0/0/1
R3(config-if)#ip summary-address eigrp 1 192.168.0.0 255.255.252.0
```

Loopback 2

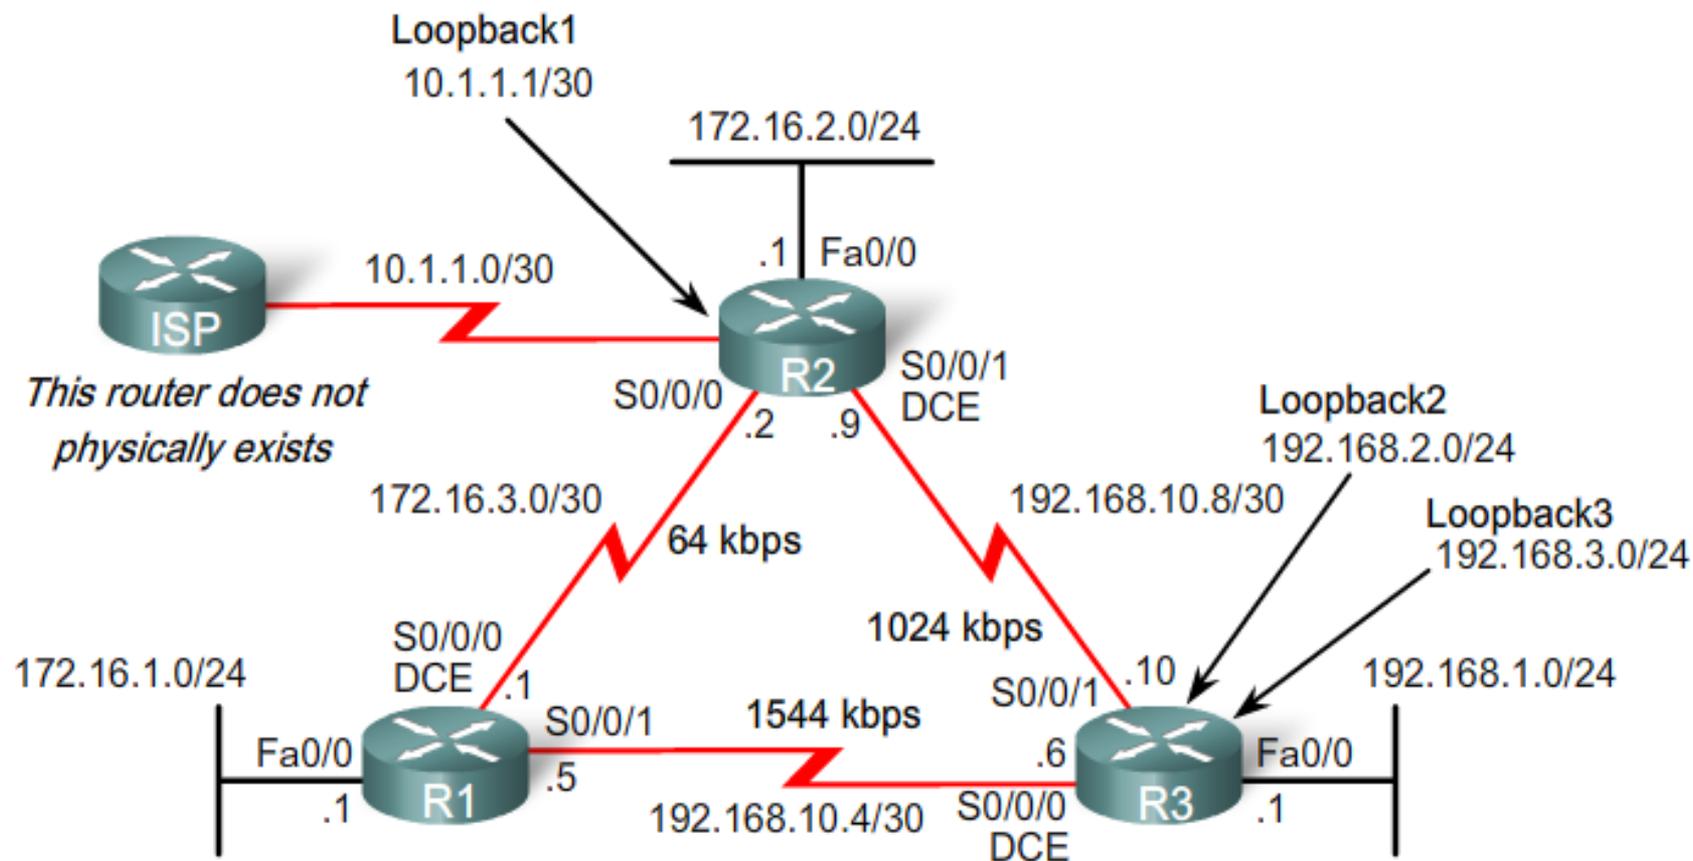
Configure the summary route on all interfaces that send EIGRP packets.

Find the Summary using the same method used to summarize static routes:

1. Write out the networks that you want to summarize in binary.
2. To find the subnet mask for summarization, start with the left-most bit.
3. Work your way to the right, finding all the bits that match consecutively.
4. The column of bits that do not match, are at the summary boundary.
5. Now, count the number of left-most matching bits, which in our example is 22.
6. To find the network address for summarization, copy the matching 22 bits and add all 0 bits to the end to make 32 bits.

**The result is the summary network address and mask for 192.168.0.0/22.**

## 9.5.4 EIGRP Default Routes



Using a static route to 0.0.0.0/0 as a default route is not routing protocol dependent.

```
R2(config)#ip route 0.0.0.0 0.0.0.0 loopback 1
R2(config)#router eigrp 1
R2(config-router)#redistribute static
```

## 9.5.4 EIGRP Default Routes

```
R2(config)#ip route 0.0.0.0 0.0.0.0 loopback 1
R2(config)#router eigrp 1
R2(config-router)#redistribute static
```

Using a static route to 0.0.0.0/0 as a default route is not routing protocol dependent. The "quad zero" static default route can be used with any currently supported routing protocols. The static default route is usually configured on the router that has a connection to a network outside the EIGRP routing domain, for example, to an ISP.

EIGRP requires the use of the redistribute static command to include this static default route with its EIGRP routing updates. The redistribute static command tells EIGRP to include this static route in its EIGRP updates to other routers. The figure shows the configuration of the static default route and the redistribute static command on router R2.

Note: The static default route is using the exit interface of Loopback1. This is because the ISP router in our topology does not physically exist. By using a loopback interface we can simulate a connection to another router.

### EIGRP Default Routes

- “quad zero” static default route
  - Can be used with any currently supported routing protocol
  - Is usually configured on a router that is connected a network outside the EIGRP domain
- EIGRP & the “Quad zero” static default route
  - Requires the use of the **redistribute static** command to disseminate default route in EIGRP updates

### EIGRP Bandwidth Utilization

```
R1 (config) #interface serial 0/0/0
R1 (config-if) #bandwidth 64
R1 (config-if) #ip bandwidth-percent eigrp 1 50
```

```
R2 (config) #interface serial 0/0/0
R2 (config-if) #bandwidth 64
R2 (config-if) #ip bandwidth-percent eigrp 1 50
```

By default, EIGRP will use only up to 50 percent of the bandwidth of an interface for EIGRP information. This prevents the EIGRP process from over-utilizing a link and not allowing enough bandwidth for the routing of normal traffic. The `ip bandwidth-percent eigrp` command can be used to configure the percentage of bandwidth that may be used by EIGRP on an interface.

## 9.5.5 Fine Tuning EIGRP

### Fine-Tuning EIGRP

- EIGRP bandwidth utilization

- By default, EIGRP uses only up to 50% of interface bandwidth for EIGRP information

- The command to change the percentage of bandwidth used by EIGRP is

Router(config-if)#*ip bandwidth-percent eigrp as-number percent*

```
R2 (config)#interface serial 0/0/0
R2 (config-if)#bandwidth 64
R2 (config-if)#ip bandwidth-percent eigrp 1 50
```

### Configuring Hello Intervals and Hold Times

- Hello intervals and hold times are configurable on a per-interface basis
- The command to configure hello interval is

Router(config-if)#*ip hello-interval eigrp as-number seconds*

- Changing the hello interval also requires changing the hold time to a value greater than or equal to the hello interval
- The command to configure hold time value is

Router(config-if)#*ip hold-time eigrp as-number seconds*

```
R2 (config) #int s0/0/0
R2 (config-if) #ip hello-interval eigrp 1 60
R2 (config-if) #ip hold-time eigrp 1 180
R2 (config-if) #end
```

## 9.6.1 Basic EIGRP Configuration Lab



### Hands-on Lab: Basic EIGRP Configuration Lab

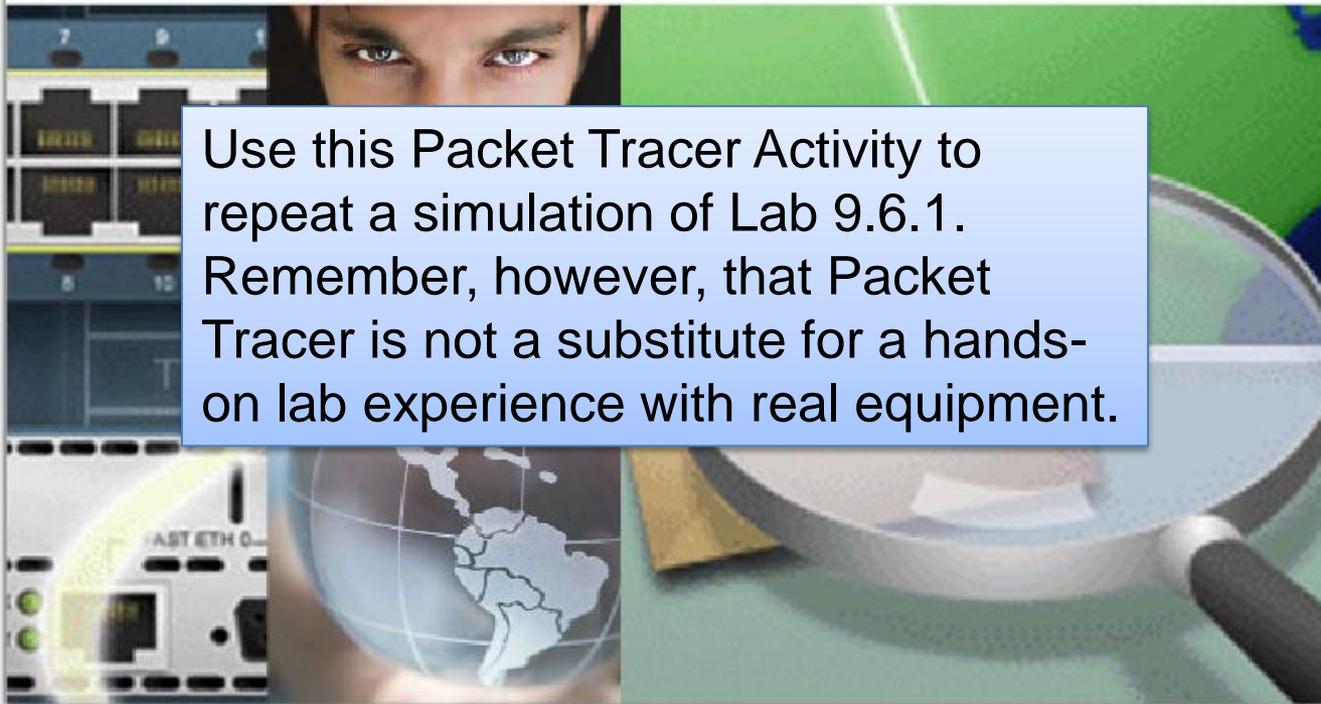
In this lab, you will learn how to configure the routing protocol EIGRP. A loopback address will be used on the R2 router to simulate a connection to an ISP, where all traffic that is not destined for the local network will be sent. Some segments of the network have been subnetted using VLSM. EIGRP is a classless routing protocol that can be used to provide subnet mask information in the routing updates. This will allow VLSM subnet information to be propagated throughout the network.

## 9.6.1 Basic EIGRP Configuration Lab



### Packet Tracer Exploration: Basic EIGRP Configuration Lab

Use this Packet Tracer Activity to repeat a simulation of Lab 9.6.1. Remember, however, that Packet Tracer is not a substitute for a hands-on lab experience with real equipment.



## 9.6.2 Challenge EIGRP Configuration Lab



### Hands-on Lab:

### Challenge EIGRP Configuration Lab

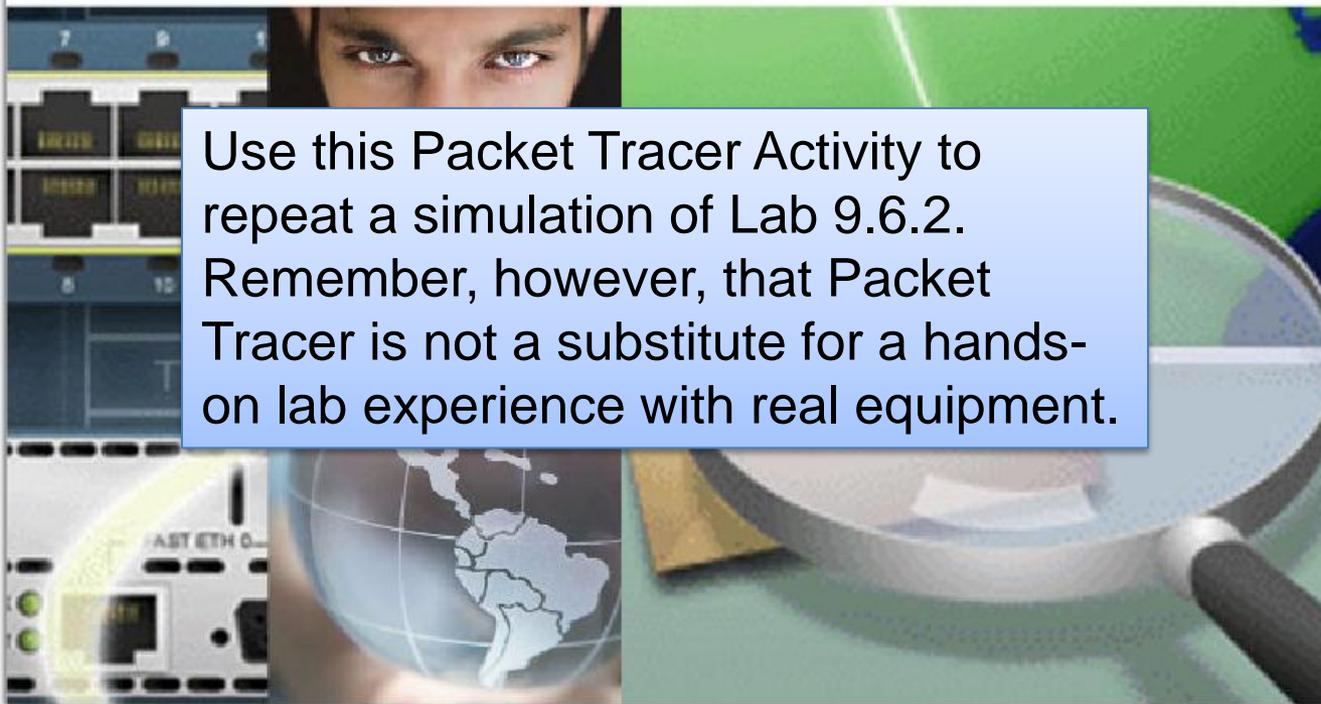
In this lab activity, you will be given a network address that must be subnetted using VLSM to complete the addressing of the network. A combination of EIGRP routing and static routing will be required so that hosts on networks that are not directly connected will be able to communicate with each other. EIGRP must be configured so that all IP traffic takes the shortest path to the destination address.

## 9.6.2 Challenge EIGRP Configuration Lab



### Packet Tracer Exploration: Challenge EIGRP Configuration Lab

Use this Packet Tracer Activity to repeat a simulation of Lab 9.6.2. Remember, however, that Packet Tracer is not a substitute for a hands-on lab experience with real equipment.



## 9.6.3 EIGRP Troubleshooting Lab



### Hands-on Lab: EIGRP Troubleshooting Lab

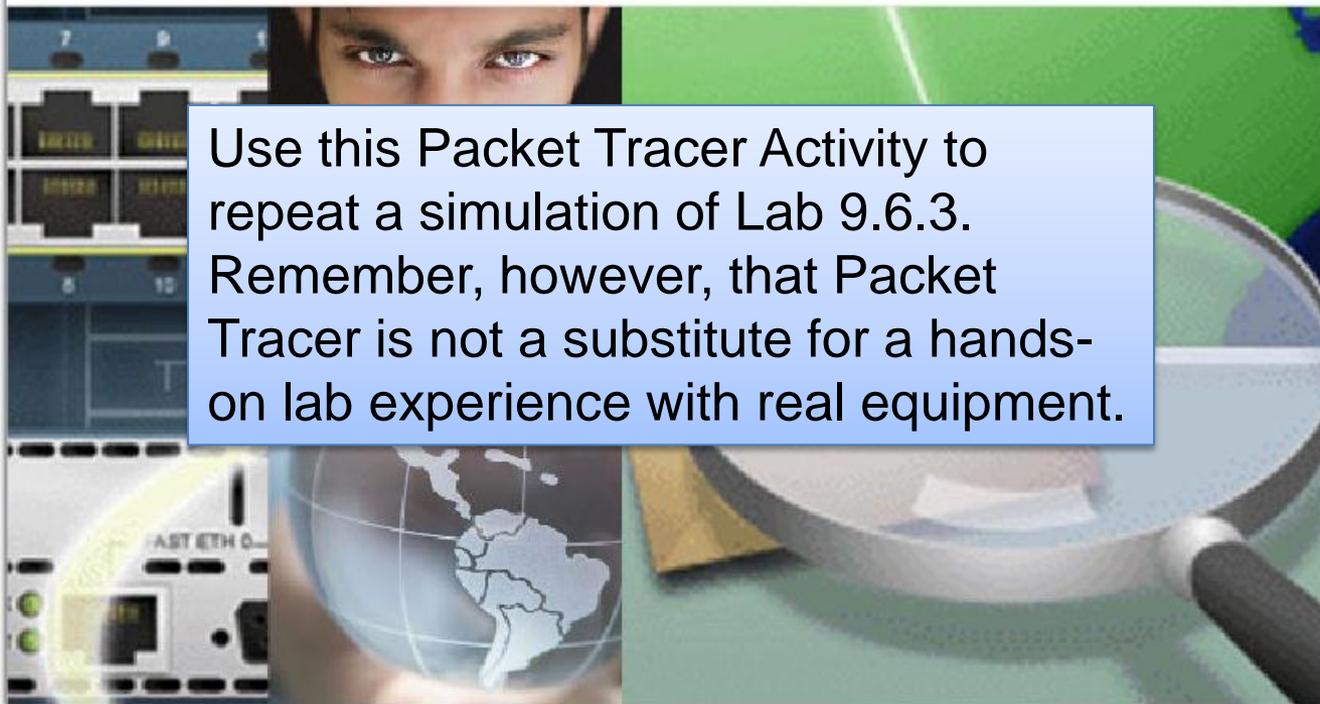
In this lab, you will begin by loading corrupted configuration scripts on each of the routers. These scripts contain errors that will prevent end-to-end communication across the network. You will need to troubleshoot each router to determine the configuration errors, and then use the appropriate commands to correct the configurations. When you have corrected all of the configuration errors, all of the hosts on the network should be able to communicate with each other.

## 9.6.3 EIGRP Troubleshooting Lab



Packet Tracer Exploration:  
EIGRP Troubleshooting Lab

Use this Packet Tracer Activity to repeat a simulation of Lab 9.6.3. Remember, however, that Packet Tracer is not a substitute for a hands-on lab experience with real equipment.



## 9.7.1 Summary and Review

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

In this chapter, you have learned to:

- Describe the background and history of EIGRP.
- Describe the features and operation of EIGRP.
- Examine the basic EIGRP configuration commands and identify their purposes.
- Calculate the composite metric used by EIGRP.
- Describe the concepts and operation of DUAL.
- Describe the uses of additional configuration commands in EIGRP.

- **Background & History**

- EIGRP is a derivative of IGRP
  - EIGRP is a Cisco proprietary distance vector routing protocol released in 1994

- **EIGRP terms and characteristics**

- EIGRP uses RTP to transmit & receive EIGRP packets
- EIGRP has 5 packet type:
  - Hello packets
  - Update packets
  - Acknowledgement packets
  - Query packets
  - Reply packets
- Supports VLSM & CIDR

- **EIGRP terms and characteristics**
  - EIGRP uses a hello protocol
    - Purpose of hello protocol is to discover & establish adjacencies
  - EIGRP routing updates
    - Aperiodic
    - Partial and bounded
    - Fast convergence

- **EIGRP commands**

- The following commands are used for EIGRP configuration

- RtrA(config)#router eigrp [autonomous-system #]

- RtrA(config-router)#network *network-number*

- The following commands can be used to verify EIGRP

- Show ip protocols

- Show ip eigrp neighbors

- Show ip route

- **EIGRP metrics include**
  - Bandwidth (default)
  - Delay (default)
  - Reliability
  - Load

### ■ DUAL

- Purpose of DUAL
  - To prevent routing loops
- Successor
  - Primary route to a destination
- Feasible successor
  - Backup route to a destination
- Feasible distance
  - Lowest calculated metric to a destination
- Reported distance
  - The distance towards a destination as advertised by an upstream neighbor

- **Choosing the best route**

- After router has received all updates from directly connected neighbors, it can calculate its DUAL
  - 1<sup>st</sup> metric is calculated for each route
  - 2<sup>nd</sup> route with lowest metric is designated successor & is placed in routing table
  - 3<sup>rd</sup> feasible successor is found
    - Criteria for feasible successor: it must have lower reported distance to the destination than the installed route's feasible distance
    - Feasible routes are maintained in topology table

- **Automatic summarization**
  - On by default
  - Summarizes routes on classful boundary
  - Summarization can be disabled using the following command
    - `RtrA(config-if)#no auto-summary`

