# Cabrillo College

Cisco Systems
Networking Academy

## CCNP – Advanced Routing
## Ch. 6 EIGRP

*Rick Graziani, Instructor*

*Mar 24, 2002*

1

## EIGRP

- Cisco proprietary, released in 1994
- Based on IGRP
- EIGRP is an ***advanced distance-vector*** routing protocol that relies on features commonly associated ***with link-state protocols.*** (sometimes called a ***hybrid routing protocol***)
- Supports VLSM and CIDR
- Can route Novell IPX and Apple AppleTalk, including CIDR which gives you "Apple CIDR" - Not Really :-)
- Like IGRP, EIGRP supports unequal-cost load balancing (unlike OSPF).

2

## EIGRP

- Routing tables are not exchanged on a periodic basis.
- Uses Hello protocols to create and maintain neighbor relationships and to determine when a link is down. (Like link-state)
  - Every 5 seconds, with a hold timer of 15 seconds
  - Links less than T1 speed, Hellos are every 60 seconds, hold timer of 180 seconds
- When a change in topology occurs, EIGRP does not flood updates like link-state protocols, but immediately sends those changes to its neighbors.
- Has speed and efficiency of routing updates like a link-state protocol, along with a topology database.

3

## OSPF versus EIGRP

| OSPF | EIGRP |
|---|---|
| Supports CIDR and VLSM, rapid convergence, partial updates, neighbor discovery | Supports CIDR and VLSM, rapid convergence, partial updates, neighbor discovery |
| Administrator can define route summarization | Automatic route-summarization and user-defined route summaries |
| Open standard; multivendor support | Proprietary; Cisco routers only |
| Scalable; administratively defined "areas" provide manageable hierarchy | Scalable, but no hierarchical design |
| Difficult to implement | Easy to implement |
| Equal-cost load balancing | Unequal-cost load balancing |

4

## EIGRP vs IGRP

- IGRP and EIGRP are compatible with each other, although EIGRP offers multiprotocol support and IGRP does not.
  - EIGRP supports
    - TCP/IP
    - IPX/SPX
    - AppleTalk



## EIGRP and IGRP

- *Automatically* redistributes when using same AS number.
- We will see examples of this *later*.

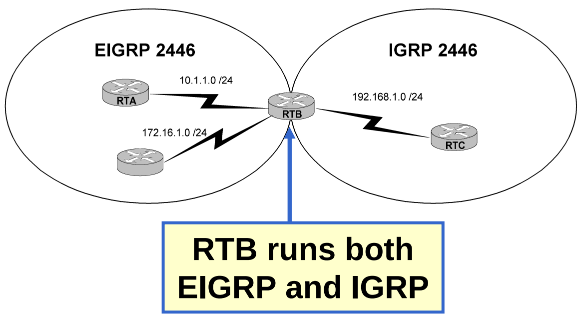

## Metric calculation: IGRP/EIGRP

*metric* = [K1 * *bandwidth* + ((K2 * *bandwidth*) / (256 * *load*)) + (K3 * *delay*)] * [K5/(*reliability* + K4)]

(with the following default constant values):

| Constant | Value |
|---|---|
| K1 | 1 |
| K2 | 0 |
| K3 | 1 |
| K4 | 0 |
| K5 | 0 |

**Notes**

- **k2** metric effects **LOAD**
- **k4** and **k5** effects **RELIABILITY**
- Red parenthesis added for clarity

7

## Metric Calculation

*metric* = [K1 * *bandwidth* + ~~((K2 * *bandwidth*) / (256 * *load*))~~ + (K3 * *delay*)] * [~~K5/(*reliability* + K4~~)]

When K2, K4 and K5 are 0, these portions of the equation is not factored in to the metric.

Thus, with the default constant values, K1=1 and K3=1, the metric equation boils down to this:

*metric* = [(1 * *bandwidth*) + (1 * *delay*)]

*metric* = ***bandwidth + delay***

**Actually:**

*metric* = **slowest *bandwidth* + sum of all *delays***

8

## Metric Calculation

- Information in the following slides is from CCNA Semester 3.
- Review the IGRP Metrics Presentation on my web site if you need a review or more information.

9

## Metric Calculation

The metrics used by EIGRP in making routing decisions are (lower the metric the better):

- bandwidth
- delay
- load
- reliability

By default, EIGRP uses only:

- Bandwidth
- Delay

***Analogies***:

Think of **bandwidth** as the *width of the pipe*

and

**delay** as the *length of the pipe*.

- **Bandwidth** is a the *carrying capacity*
- **Delay** is the *end-to-end travel time*.

10

## Metric Calculation

If these are the default:

- bandwidth (default)
- delay (default)

When are these used?

- load
- reliability

Only when configured by the network administrator to do so!

EIGRP also tracks (but does **not** use in its metric calculation):

- MTU (Maximum Transmission Unit)
- Hop Count

Use **show interface** command to view the metrics used on a specific interface that is routing EIGRP.

- These are the **raw values.**

11

## Metric Calculation – show interfaces

```
Router> show interfaces s1/0
Serial1/0 is up, line protocol is up
  Hardware is QUICC Serial
  Description: Out to VERIO
  Internet address is 207.21.113.186/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     rely 255/255, load 246/255
  Encapsulation PPP, loopback not set
  Keepalive set (10 sec)
<output omitted>
```

bandwidth

delay

reliability

load

12

## Metric Calculation – Bandwidth

**Bandwidth**

- Expressed in **kilobits** (**show interface**)
- This is a **static number** and used for metric calculations only.
- Does not necessarily reflect the actual bandwidth of the link.
- It is an **information parameter only**.
- You cannot adjust the actual bandwidth on an interface with this command.
- Use the **show interface** command to display the *raw value*

**The default values:**

- Default bandwidth of a Cisco interface depends on the type of interface.
- Default bandwidth of a Cisco **serial interface is 1544 kilobits** or 1,544,000 bps (T1), whether that interface is attached to a T1 line (1.544 Mbps) or a 56K line.
- IGRP/EIGRP metric uses the **slowest bandwidth** of all of the outbound interfaces to the destination network.
- *More on defaults soon!*

13

## Metric Calculation – Bandwidth

**Changing the bandwidth informational parameter:**

The bandwidth can be changed using:

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

To restore the default value:

```
Router(config-if)# no bandwidth
```

14

## Metric Calculation – Delay

**Delay**

- Like bandwidth, delay it is a **static number**.
- Expressed in **microseconds**, *millionths of a second*
- (Uses the Greek letter mu with an S, µS, NOT "ms" which is millisecond or *thousandths of a second*)
- Use the **show interface** command to display the *raw value*
- It is an **information parameter only**.

**The default values:**

- The default delay value of a Cisco interface depends upon the **type of interface**.
- Default delay of a Cisco **serial interface** is **20,000 microseconds**, that of a T1 line.
- IGRP/EIGRP metric uses the **sum of all of the delays** of all of the outbound interfaces to the destination network.
- *More on this coming!*

15

## Metric Calculation – Delay

**Changing the delay informational parameter:**

The delay can be changed using:

```
Router(config-if)#  delay tens-of- µS
                           (microseconds)
```

Example of changing the delay on a serial interface to 30,000 microseconds:

```
Router(config-if)# delay 3000
```

To restore the 20,000 microsecond default value:

```
Router(config-if)# no delay
```

16

## IGRP vs EIGRP Metric Calculation

Let's take a look a the default metrics used for bandwidth and delay, *depending upon the type of interface*….

17

## IGRP vs EIGRP Metric Calculation

**IGRP**

- bandwidth = (10,000,000/*bandwidth*)
- delay = *delay*/10

**EIGRP**

- bandwidth = (10,000,000/*bandwidth*) * 256
- delay = (*delay*/10) * 256

**Note: The reference-bandwidth**

**For both IGRP and EIGRP: $10^7$, (10,000,000/bandwidth), whereas with OSPF it was $10^8$ (100,000,000/bandwidth)**

**The difference**:

- IGRP metric is 20 bits long
- EIGRP metric is 32 bits long
- EIGRP metric is 256 times greater for the same route
- EIGRP allows for finer comparison of potential routes

18

# EIGRP Metrics

Values displayed in show interface commands and sent in routing updates.

| Media | Bandwidth K= kilobits | $BW_{EIGRP}$ 10,000,000/Bandwidth *256 | Delay | $DLY_{EIGRP}$ Delay/10 *256 |
|---|---|---|---|---|
| 100M ATM | **100,000K** | 25,600 | **100 µS** | 2,560 |
| Fast Ethernet | **100,000K** | 25,600 | **100 µS** | 2,560 |
| FDDI | **100,000K** | 25,600 | **100 µS** | 2,560 |
| HSSI | **45,045K** | 56,832 | **20,000 µS** | 512,000 |
| 16M Token Ring | **16,000K** | 160,000 | **630 µS** | 16,128 |
| Ethernet | **10,000K** | 256,000 | **1,000 µS** | 25,600 |
| T1 (Serial Default) | **1,544K** | 1,657,856 | **20,000 µS** | 512,000 |
| 512K | **512K** | 4,999,936 | **20,000 µS** | 512,000 |
| DS0 | **64K** | 40,000,000 | **20,000 µS** | 512,000 |
| 56K | **56K** | 45,714,176 | **20,000 µS** | 512,000 |

**$BW_{EIGRP}$** and **$DLY_{EIGRP}$** display values as sent in EIGRP updates and used in calculating the EIGRP metric.

Calculated values (cumulative) displayed in routing table (show ip route).

19

# IGRP Metrics

Values displayed in show interface commands and sent in routing updates.

| Media | Bandwidth K= kilobits | $BW_{IGRP}$ 10,000,000/Bandwidth | Delay | $DLY_{IGRP}$ Delay/10 |
|---|---|---|---|---|
| 100M ATM | **100,000K** | 100 | **100 µS** | 10 |
| Fast Ethernet | **100,000K** | 100 | **100 µS** | 10 |
| FDDI | **100,000** | 100 | **100 µS** | 10 |
| HSSI | **45,045K** | 222 | **20,000 µS** | 2,000 |
| 16M Token Ring | **16,000K** | 625 | **630 µS** | 63 |
| Ethernet | **10,000K** | 1,000 | **1,000 µS** | 100 |
| T1 (Serial Default) | **1,544K** | 6,476 | **20,000 µS** | 2,000 |
| 512K | **512K** | 19,531 | **20,000 µS** | 2,000 |
| DS0 | **64K** | 156,250 | **20,000 µS** | 2,000 |
| 56K | **56K** | 178,571 | **20,000 µS** | 2,000 |

**$BW_{IGRP}$** and **$DLY_{IGRP}$** display values as sent in IGRP updates and used in calculating the IGRP metric.

Calculated values (cumulative) displayed in routing table (show ip route). EIGRP values are 256 times greater.

# Metric Calculation – show interfaces

```
Router> show interfaces s1/0
Serial1/0 is up, line protocol is up
  Hardware is QUICC Serial
  Description: Out to VERIO
  Internet address is 207.21.113.186/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     rely 255/255, load 246/255
  Encapsulation PPP, loopback not set
  Keepalive set (10 sec)
<output omitted>
```

bandwidth
delay
reliability
load

21

## IGRP

| Media | Bandwidth K= kilobits | $BW_{IGRP}$ 10,000,000/Bandwidth | Delay | $DLY_{IGRP}$ Delay/10 |
|---|---|---|---|---|
| 100M ATM | **100,000K** | 100 | **100 µS** | 10 |
| Fast Ethernet | **100,000K** | 100 | **100 µS** | 10 |
| FDDI | **100,000** | 100 | **100 µS** | 10 |
| HSSI | **45,045K** | 222 | **20,000 µS** | 2,000 |
| 16M Token Ring | **16,000K** | 625 | **630 µS** | 63 |
| Ethernet | **10,000K** | 1,000 | **1,000 µS** | 100 |
| T1 (Serial Default) | **1,544K** | 6,476 | **20,000 µS** | 2,000 |
| 512K | **512K** | 19,531 | **20,000 µS** | 2,000 |
| DS0 | **64K** | 156,250 | **20,000 µS** | 2,000 |
| 56K | **56K** | 178,571 | **20,000 µS** | 2,000 |

**$BW_{IGRP}$** and **$DLY_{IGRP}$** display values as sent in IGRP updates and used in calculating the IGRP metric.

**Viva la difference!**

## EIGRP

| Media | Bandwidth K= kilobits | $BW_{EIGRP}$ 10,000,000/Bandwidth *256 | Delay | $DLY_{EIGRP}$ Delay/10 *256 |
|---|---|---|---|---|
| 100M ATM | **100,000K** | 25,600 | **100 µS** | 2,560 |
| Fast Ethernet | **100,000K** | 25,600 | **100 µS** | 2,560 |
| FDDI | **100,000K** | 25,600 | **100 µS** | 2,560 |
| HSSI | **45,045K** | 56,832 | **20,000 µS** | 512,000 |
| 16M Token Ring | **16,000K** | 160,000 | **630 µS** | 16,128 |
| Ethernet | **10,000K** | 256,000 | **1,000 µS** | 25,600 |
| T1 (Serial Default) | **1,544K** | 1,657,856 | **20,000 µS** | 512,000 |
| 512K | **512K** | 4,999,936 | **20,000 µS** | 512,000 |
| DS0 | **64K** | 40,000,000 | **20,000 µS** | 512,000 |
| 56K | **56K** | 45,714,176 | **20,000 µS** | 512,000 |

Calculated values (cumulative) displayed in routing table (show ip route).

EIGRP values are 256 times greater.

**$BW_{EIGRP}$** and **$DLY_{EIGRP}$** display values as sent in EIGRP updates and used in calculating the EIGRP metric.

## What about Reliability and Load?

The metrics used by EIGRP in making routing decisions are (lower the metric the better):

- bandwidth
- delay
- load
- reliability

By default, EIGRP uses only:

- Bandwidth
- Delay

23

## What about Reliability and Load?

**Reliability**

- Reliability is measure **dynamically**
- Uses **error rate** for measurement
- Reflects the **total outgoing error rates** of the **interfaces along the route**
- Calculated on a five minute weighted average, so not to allow sudden peaks and valleys to make a significant impact

Expressed as an 8 bit number

- **255** is a 100% reliable link
- **1** is a minimally reliable link

**Higher** the better!

24

## Metric Calculation – show interfaces

```
Router> show interfaces s1/0
Serial1/0 is up, line protocol is up
  Hardware is QUICC Serial
  Description: Out to VERIO
  Internet address is 207.21.113.186/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     rely 255/255, load 246/255
  Encapsulation PPP, loopback not set
  Keepalive set (10 sec)
<output omitted>
```

bandwidth
delay
reliability
load

**shows reliability as a fraction of 255, for example:**

**rely 190/255 (or 74% reliability)**

**rely 234/255 (or 92% reliability)**

**rely 255/255 (or 100% reliability)**

25

## What about Reliability and Load?

**Load**

- Load is measure **dynamically**
- Uses **channel occupancy** for measurement
- Reflects the **total outgoing load** of the **interfaces along the route**
- Calculated on a five minute weighted average, so not to allow sudden peaks and valleys to make a significant impact

Expressed as an 8 bit number

- **255** is a 100% loaded link
- **1** is a minimally loaded link

**Lower** the better!

Note: Even though load and reliability are dynamically changing values, EIGRP will not recalculate the route metric when these parameters change.

26

# Metric Calculation – show interfaces

```
Router> show interfaces s1/0
Serial1/0 is up, line protocol is up
  Hardware is QUICC Serial
  Description: Out to VERIO
  Internet address is 207.21.113.186/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
     rely 255/255, load 246/255
  Encapsulation PPP, loopback not set
  Keepalive set (10 sec)
<output omitted>
```

bandwidth  delay  reliability  load

**shows load as a fraction of 255, for example:**

- **load 10/255 (or 3% loaded link)**
- **load 40/255 (or 16% loaded link)**
- **load 255/255 (or 100% loaded link)**

27

# What about Reliability and Load?

**IGRP/EIGRP metric** =

[k1* **BW**$_{IGRP(minimum)}$ +
(k2* **BW**$_{IGRP(minimum)}$)/(256-**LOAD**) +
k3* **DLY**$_{IGRP(sum)}$ ] *
[k5/**RELIABILITY** + k4)]

- **k2** metric effects **LOAD**
- **k4** and **k5** effects **RELIABILITY**
- *Multiply Reliability only if > 0*

**Default**:

k1=k3=1 and k2=k4=k5=0

- You may change the k values to change what you want to give more or less weight to.
  - k1 for bandwidth
  - k2 for load
  - k3 for delay
  - k4 and k5 for Reliability
- Higher the k value, the more that part of the metric is used to calculate the overall IGRP metric

28

## What about Reliability and Load?

**Turning the knobs:**

We can use the other metrics of **Reliability** and **Load** by adjusting their k values to something greater than "0"

The command to adjust the k values is:

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

**Notes**:

- tos is always set to 0; at one time it was Cisco's intent to use it, but it was never implemented
- EIGRP neighbors must agree on K values to establish an adjacency and to avoid routing loops.

**Caution!**

- Know what the impact will be before changing the defaults.
- It can give you unexpected results if you do not know what you are doing!
- If you modify the weights, you should configure all routers so they are all using the same weight values.

29

## Now, let's see how EIGRP works …

30

## EIGRP Technologies

Four key technologies set EIGRP apart from IGRP

- Neighbor discovery and recovery
- Reliable Transport Protocol (RTP)
- DUAL finite-state machine (FSM)
- PDMs

31

## Neighbor Discovery/Recovery

- EIGRP routers establish adjacencies with neighbor routers by using small hello packets.
- Hellos are sent every **5 seconds** by default.
- K values must be the same between neighbors.
- An EIGRP router assumes that, as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable.
- Hold time tells the router how long it should consider the neighbor alive if it has not received any EIGRP packets (Hello, EIGRP updates, etc.)
- Hold time is normally three times the configured Hello interval.
- Both the Hello and Hold time intervals are configurable on a per interface basis.
- EIGRP routers exchange routing information the same way as other distance vector routing protocols, but do not send periodic updates.
- EIGRP updates are only sent when a network is added or removed from the topology database, when the successor for a given network changes, or when the locally used metric is updated. (later)
- EIGRP, like any other distance-vector routing protocol uses split-horizon.

32

## RTP

- RTP (Reliable Transport Protocol)
  - transport-layer protocol
- EIGRP is protocol-independent; that is, it doesn't rely on TCP/IP to exchange routing information the way RIP, IGRP, and OSPF do.
- To stay independent of IP, EIGRP uses the transport-layer protocol to guarantee delivery of routing information: **RTP**.
- RTP supports reliable and unreliable delivery
- RTP supports unicasting and multicasting
- Initial delivery of EIGRP messages are done using multicast packets, that is data is sent ot all neighbors on a segment, and every neighbor is expected to acknowledge it with a unicast Hello packet.
- After adjacency has been formed and added to neighbor table, routers exchange routing information which is stored in the topology table. (later)
- RTP uses a mechanism of sequence numbers and acknowledgements
- RTP is used for EIGRP queries, updates and replies
- RTP is **not** used for EIGRP Hello's and Ack's

33

## DUAL FSM

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

**What is FSM?**

- An FSM is an abstract machine, not a mechanical device with moving parts.
- FSMs define a set of possible states something can go through, what events causes those states, and what events result from those states.
- Designers use FSMs to describe how a device, computer program, or routing algorithm will react to a set of input events.

34

## Example FSM

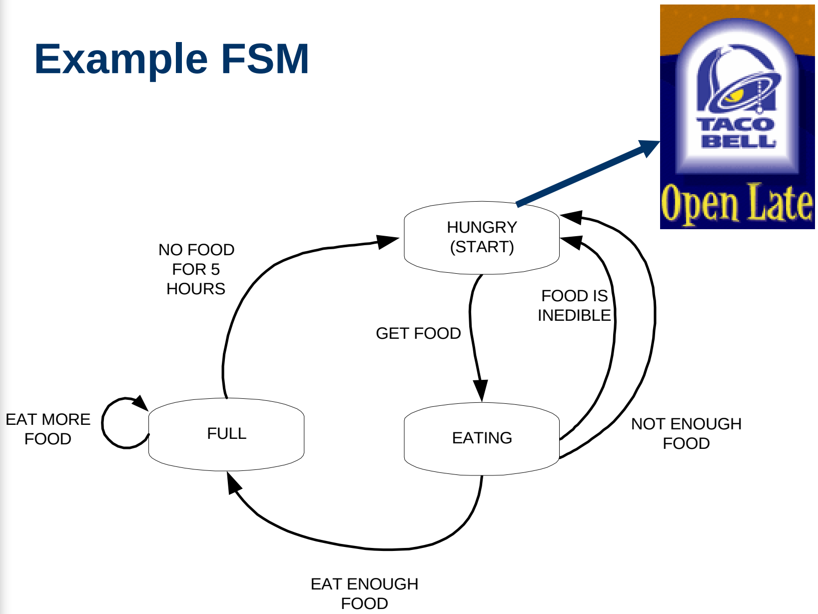




## DUAL FSM

- DUAL selects alternate routes quickly by using the information in the EIGRP tables.
- *If a link goes down, DUAL looks for a feasible successor in its neighbor and topology tables.*
- A **successor** is a neighboring router that is currently being used for packet forwarding, provides the ***least-cost route*** to the destination, and is not part of a routing loop.
- **Feasible successors** provide the ***next lowest-cost path*** without introducing routing loops.
  - Feasible successor routes can be used in case the existing route fails; packets to the destination network are immediately forwarded to the feasible successor, which at that point, is promoted to the status of successor.
- Selects a best loop-free path to a destination, the next hop being known as the **successor**.
- All other routers to the same destination, that also meet the **feasible condition**, meaning they are also loop-free (later), become **feasible successors**, or back-up routes.
- **debug eigrp fsm**

# PDMs

- PDM (Protocol-dependent module)
- EIGRP is modular (EIGRP is protocol-independent)
- Different PDMs can be added to EIGRP as new routed protocols are enhanced or developed:
  - **IPv4, IPv6, IPX, and AppleTalk**

Each PDM is responsible for all functions related to its specific routed protocol.

- The IP-EIGRP module is responsible for the following:
  - **Sending and receiving EIGRP packets that bear IP data**
  - **Notifying DUAL of new IP routing information that is received**
  - **Maintaining the results of DUAL's routing decisions in the IP routing table**
  - **Redistributing routing information that was learned by other IP-capable routing protocols**

37

# EIGRP Terminology

- EIGRP routers keep route and topology information readily available in RAM so that they can react quickly to changes.
- Like OSPF, EIGRP keeps this information in several tables, or databases.
  - Neighbor table
  - Topology table
  - Routing table
  - Successor
  - Feasible Successor
- We will first have an overview of all of the terminology and then see how it works and what it all means!

38

## Neighbor Table

- Neighbor table
  - Each EIGRP router maintains a neighbor table that lists adjacent routers.
  - This table is comparable to the adjacency database used by OSPF.
  - There is a neighbor table for each protocol that EIGRP supports
- Whenever a new neighbor is discovered, the address of that neighbor and the interface used to reach it are recorded in a new neighbor table entry.

```
RouterC#show ip eigrp neighbors
IP-EIGRP neighbors for process 44
H   Address        Interface    Hold Uptime    SRTT   RTO  Q   Seq
                                (sec)          (ms)       Cnt  Num
0   192.168.0.1    Se0            11 00:03:09 1138   5000  0   6
1   192.168.1.2    Et0            12 00:34:46    4    200  0   4
```

## Neighbor Table includes

```
RouterC#show ip eigrp neighbors
IP-EIGRP neighbors for process 44
H   Address        Interface    Hold Uptime    SRTT   RTO  Q   Seq
                                (sec)          (ms)       Cnt  Num
0   192.168.0.1    Se0            11 00:03:09 1138   5000  0   6
1   192.168.1.2    Et0            12 00:34:46    4    200  0   4
```

- ***Neighbor address*** The network-layer address of the neighbor router(s).
- ***Queue count*** The number of packets waiting in queue to be sent. If this value is constantly higher than zero, then there may be a congestion problem at the router. A zero means that there are no EIGRP packets in the queue.

40

# Neighbor Table includes

```
RouterC#show ip eigrp neighbors
IP-EIGRP neighbors for process 44
H    Address         Interface     Hold Uptime    SRTT   RTO  Q   Seq
                                   (sec)          (ms)        Cnt Num
0    192.168.0.1     Se0             11 00:03:09 1138   5000  0   6
1    192.168.1.2     Et0             12 00:34:46    4    200  0   4
```

***Smooth Round Trip Timer (SRTT)*** The average time it takes to send and receive packets from a neighbor.

This timer is used to determine the retransmit interval (RTO)

***Hold Time*** The interval to wait without receiving anything from a neighbor before considering the link unavailable.

Originally, the expected packet was a hello packet, but in current Cisco IOS software releases, **any EIGRP packets received after the first hello will reset the timer**.

41

# Neighbor Table notes

- Note that an EIGRP router can maintain multiple neighbor tables, one for each L3 protocol running (for example, IP, AppleTalk).
- A router must run a unique EIGRP process for each routed protocol.

```
RTX#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H    Address                 Interface     Hold Uptime    SRTT   RTO  Q   Seq
                                           (sec)          (ms)        Cnt Num
1    10.2.0.2                Se1             12 00:27:39  333   1998  0   10
0    10.1.0.1                Se0             14 01:17:14   40    240  0   27
```

```
RTX#show ipx eigrp neighbors
IPX EIGRP Neighbors for process 22
H    Address                 Interface     Hold Uptime    SRTT   RTO  Q   Seq
                                           (sec)          (ms)        Cnt Num
1    2000.0000.0c76.080c     Se1             14 00:04:21   28    200  0   22
0    1000.0000.0c38.6fa2     Se0             14 00:04:24   28    200  0   22
```

```
RTX#show appletalk eigrp neighbors
AT/EIGRP Neighbors for process 1, router id 2
H    Address                 Interface     Hold Uptime    SRTT   RTO  Q   Seq
                                           (sec)          (ms)        Cnt Num
0    1000.123                Se0             11 00:15:01    8    200  0   7
1    2000.28                 Se1             14 00:41:11   11    200  0   9
```

42

# EIGRP Topology Table

- ***Topology table***
  - Each EIGRP router maintains a topology table for each configured network protocol.
  - This table includes route entries for all destinations that the router has learned.
  - All learned routes to a destination are maintained in the topology table.
- EIGRP uses its **topology table** to store all the information it needs to calculate a set of distances and vectors to all reachable destinations.

**More about this table later!**

```
RouterB#show ip eigrp topology
IP-EIGRP Topology Table for process 44
Codes: P - Passive, A - Active, U - Update, Q - Query, R -
   Reply, r - Reply status
P 206.202.17.0/24, 1 successors, FD is 2195456
         via 206.202.16.1 (2195456/2169856), Ethernet0
P 206.202.18.0/24, 2 successors, FD is 2198016
         via 192.168.0.2 (2198016/284160), Serial0
         via 206.202.16.1 (2198016/2172416), Ethernet0
```

43

# Topology Table

- Not only does the topology table track information regarding route states, but it can also record special information for external routes, including the administrator tag.
- EIGRP classifies routes as either **internal** or **external**.
- EIGRP uses a process called route tagging to add special tags to each route.
- These tags identify a route as internal or external, and may include other information as well.

All external routes are included in the topology table, and are tagged with the following information:

- The identification number (**router ID**) of the EIGRP router that redistributed the route into the EIGRP network
- The **AS number** of the destination
- The **protocol** used in that external network
- The **cost or metric** received from that external protocol
- The configurable **administrator tag**

44

## Topology Table

```
RTX#sh ip eigrp top 204.100.50.0
IP-EIGRP topology entry for 204.100.50.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD
   is 2297856
  Routing Descriptor Blocks:
  10.1.0.1 (Serial0), from 10.1.0.1, Send flag is 0x0
      Composite metric is (2297856/128256), Route is External
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      External data:
        Originating router is 192.168.1.1
        AS number of route is 0
        External protocol is Connected, external metric is 0
        Administrator tag is 0 (0x00000000)
```

FD/RD

45

## Topology Table

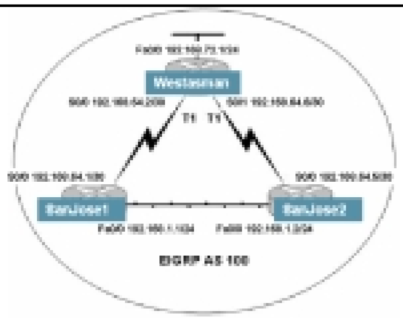


Much more on about this table and scenario later after we discuss a few more terms.

```
SanJose2#show ip eigrp topology all-links
P 192.168.72.0/24, 1 successors, FD is 2172416, serno 93
         via 192.168.64.6 (2172416/28160), Serial0
         via 192.168.1.1 (2174976/2172416), FastEthernet0
P 192.168.64.0/30, 1 successors, FD is 2172416, serno 91
         via 192.168.1.1 (2172416/2169856), FastEthernet0
         via 192.168.64.6 (2681856/2169856), Serial0
P 192.168.64.4/30, 1 successors, FD is 2169856, serno 72
         via Connected, Serial0
P 192.168.1.0/24, 1 successors, FD is 28160, serno 1
         via Connected, FastEthernet0
```

## Topology Table

Question: Since EIGRP has a topology table, does this make it a link-state routing protocol?

Answer:

- No, the information in the topology table is not in the form of PDUs describing the complete network topology.
- The EIGRP topology table contains information about paths through the router's adjacent neighbors.
- Also, EIGPR does not perform shortest-path calculation by caclulating the shortest-path tree, but instead uses the DUAL algorithm.

*Alex Zinin, Cisco IP Routing*

47

## IP Routing Table

- *Routing table*
  - EIGRP chooses the best (that is, successor) routes to a destination from the topology table and places these routes in the routing table.
  - Each EIGRP router maintains a topology table for each network protocol.

```
RouterB#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
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
   default U - per-user static route
Gateway of last resort is not set
C    10.1.1.0 is directly connected, Serial0
D    172.16.0.0 [90/2681856] via 10.1.1.0, Serial0
D EX 192.168.1.0 [170/2681856] via 10.1.1.1, 00:00:04, Serial0
```

- EIGRP displays both internal EIGRP routes and external EIGRP routes.
- External EIGRP routes are routes external to EIGRP (redistributed)

## Routing Tables

- The routing table contains the routes installed by DUAL as the best loop-free paths to a given destination.
- EIGRP will maintain up to four routes per destination.
- These routes can be of equal, or unequal cost (if using the **variance** command). (later)

49

## Routing Tables

```
RTX#show ip route
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
D     192.168.5.0/24 [90/3219456] via 10.2.0.2, 00:12:19, Serial1
      D     192.168.1.0/24 [90/2195456] via 10.1.0.1, 00:12:19, Serial0
C     192.168.2.0/24 is directly connected, Ethernet0
D     192.168.3.0/24 [90/2195456] via 10.2.0.2, 00:12:19, Serial1
```

```
RTX#show ipx route
11 Total IPX routes. Up to 1 parallel paths and 16 hops allowed.
No default route known.
C       1000 (HDLC),          Se0
E       3000 [2681856/0] via     2000.0000.0c76.080c, age 00:10:49, 1u, Se1
E       4000 [276864000/2] via   2000.0000.0c76.080c, age 00:10:41, 1u, Se1
```

```
RTX#show appletalk route
Codes: R - RTMP derived, E - EIGRP derived, C - connected, A - AURP
       S - static  P - proxy
6 routes in internet
The first zone listed for each entry is its default (primary) zone.
E Net 100-101 [1/G] via 1000.123, 1400 sec, Serial0, zone san fran
E Net 300-301 [1/G] via 2000.28, 3016 sec, Serial1, zone san jose
E Net 400-401 [2/G] via 2000.28, 3016 sec, Serial1, zone antioch
C Net 1000-1001 directly connected, Serial0, zone wan one
```

50

# EIGRP Terminology

- ***Successor***
  - A successor is a route selected as the primary route to use to reach a destination.
  - Successors are the entries kept in the routing table.

- ***Feasible Successor*** - A backup route
  - A feasible successor is a backup route.
  - These routes are selected at the same time the successors are identified, but they are kept in the topology table.
  - Multiple feasible successors for a destination can be retained in the topology table.

**Let's see how this works!**

51

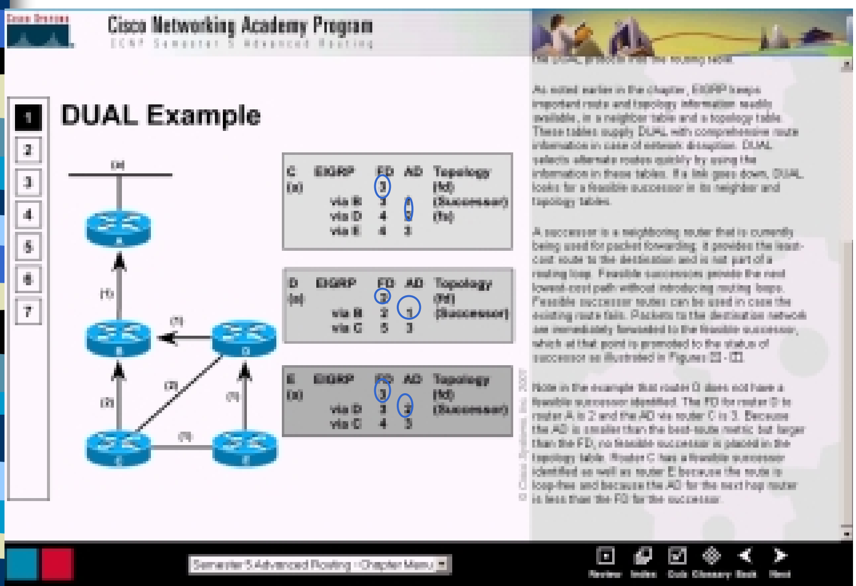


52

1 2 3 4 5 6 7

# DUAL Example

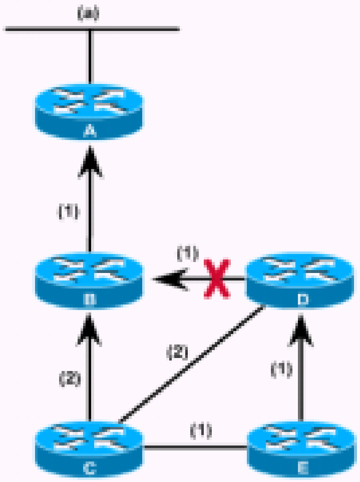


| C (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | | 3 | | (fd) |
| | via B | 3 | 1 | (Successor) |
| | via D | 4 | 2 | (fs) |
| | via E | 4 | 3 | |

| D | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | | 2 | | (fd) |
| | ~~via B~~ | ~~2~~ | ~~1~~ | ~~(Successor)~~ |
| | via C | 5 | 3 | |

| E | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | | 3 | | (fd) |
| | via D | 3 | 2 | (Successor) |
| | via C | 4 | 3 | |

3

1 2 3 4 5 6 7

# DUAL Example

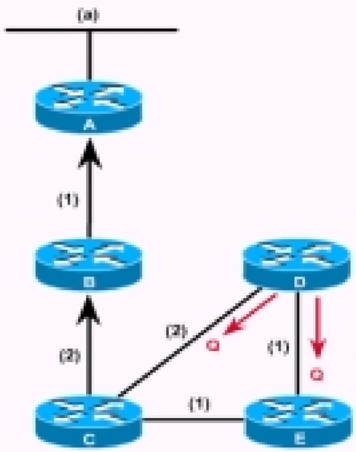


| C | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | | 3 | | (fd) |
| | via B | 3 | 1 | (Successor) |
| | via D | | | |
| | via E | 4 | 3 | |

| D | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | **ACTIVE** | -1 | | (fd) |
| | via E | | | (q) |
| | via C | 5 | 3 | (q) |

| E | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| (a) | | 3 | | (fd) |
| | ~~via D~~ | ~~3~~ | ~~2~~ | ~~(Successor)~~ |
| | via C | 4 | 3 | |

54

1 2 3 **4** 5 6 7

# DUAL Example

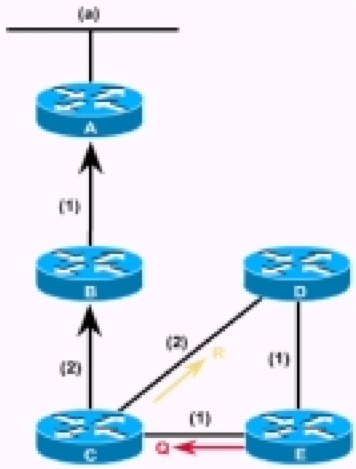


```
C    EIGRP      FD  AD  Topology
(a)             3       (fd)
       via B    3   1   (Successor)
       via D
       via E
```

```
D    EIGRP          FD  AD  Topology
(a) **ACTIVE**      -1      (fd)
       via E                (q)
       via C        5   3
```

```
E    EIGRP          FD  AD  Topology
(a) **ACTIVE**      -1      (fd)
       via D
       via C        4   3   (q)
```



1 2 3 4 **5** 6 7

# DUAL Example

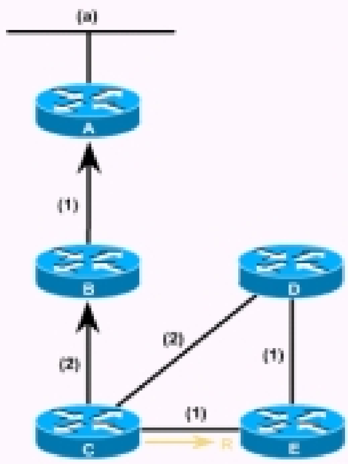


```
C    EIGRP      FD  AD  Topology
(a)             3       (fd)
       via B    3   1   (Successor)
       via D
       via E
```

```
D    EIGRP          FD  AD  Topology
(a) **ACTIVE**      -1      (fd)
       via E                (q)
       via C        5   3
```

```
E    EIGRP      FD  AD  Topology
(a)             4       (fd)
       via C    4   3   (Successor)
       via D
```

# DUAL Example

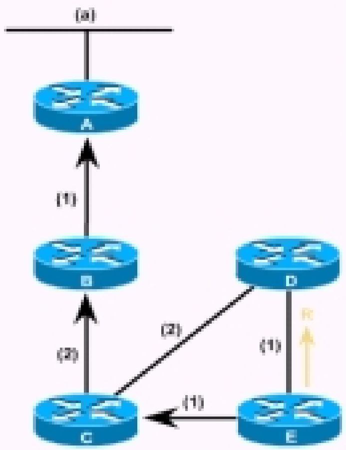


| C (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 3 | | (fd) |
| | via B | 3 | 1 | (Successor) |
| | via D | | | |
| | via E | | | |

| D (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 5 | | (fd) |
| | via C | 5 | 3 | (Successor) |
| | via E | 5 | 4 | |

| E (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 4 | | (fd) |
| | via C | 4 | 3 | (Successor) |
| | via D | | | |

57



# DUAL Example

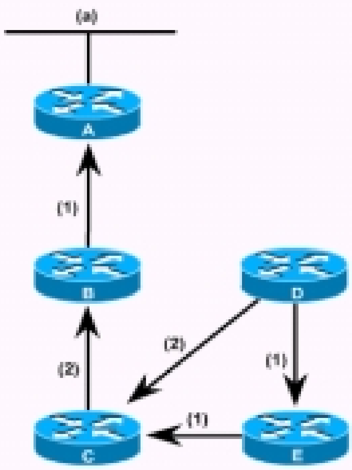


| C (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 3 | | (fd) |
| | via B | 3 | 1 | (Successor) |
| | via D | | | |
| | via E | | | |

| D (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 5 | | (fd) |
| | via C | 5 | 3 | (Successor) |
| | via E | 5 | 4 | (Successor) |

| E (a) | EIGRP | FD | AD | Topology |
|---|---|---|---|---|
| | | 4 | | (fd) |
| | via C | 4 | 3 | (Successor) |
| | via D | | | |

58

## Successors and Feasible Successors

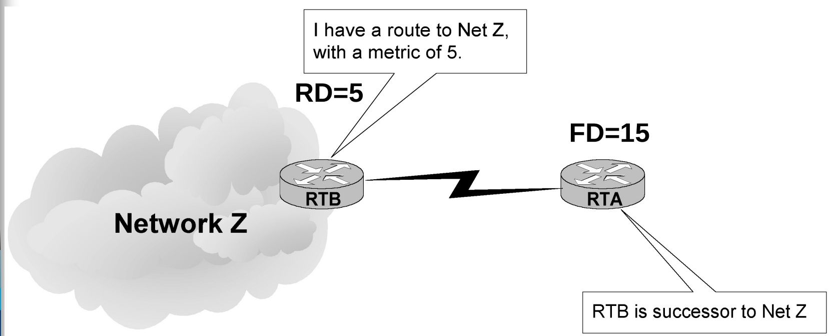


59

## Successors and Feasible Successors

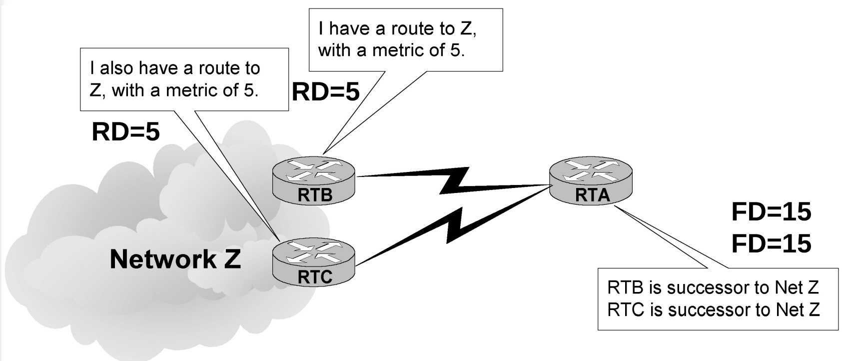


60

## Successors and Feasible Successors

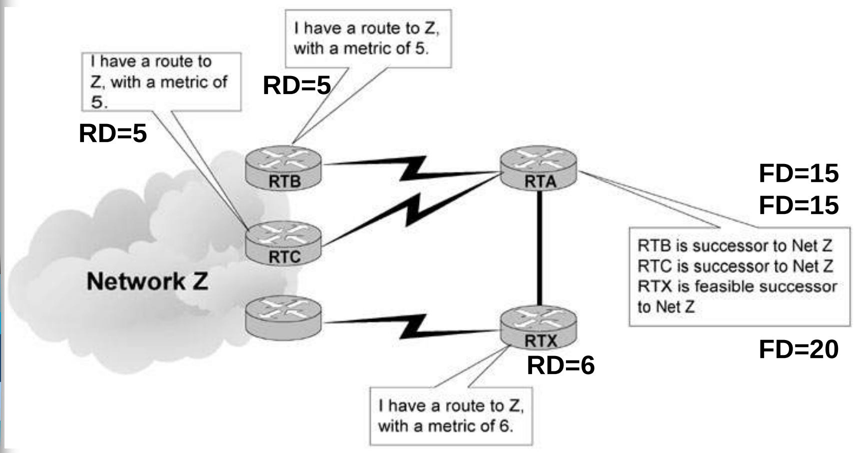


61

## Successors and Feasible Successors

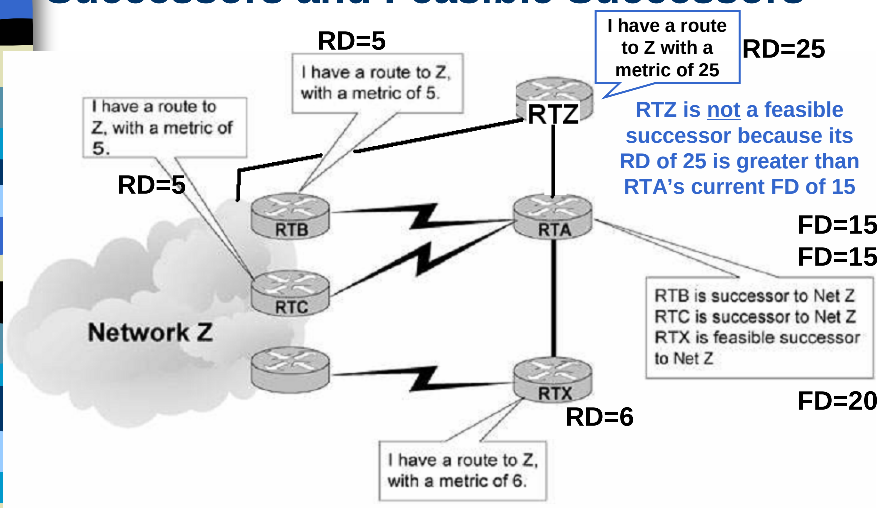


62

**Feasible distance (FD)** is the minimum distance (metric) along a path to a destination network.

**Reported distance (RD)** is the distance (metric) towards a destination as advertised by an upstream neighbor. Reported distance is the distance reported in the queries, the replies and the updates.

A neighbor meets the **feasible condition(FC)** if the reported distance by the neighbor is smaller than the current feasible distance (FD) of this router. "If a neighbors metric is less than mine, then I know the neighbor doesn't have a loop going through me."

A **feasible successor** is a neighbor whose reported distance (RD) is less than the current feasible distance (FD). Feasible successor is one who meets the feasible condition (FC).

Your route (metric) to the network (RD to me) must be LESS than my current route (my total metric) to that same network. If your route (metric) to the network (RD to me) is LESS than my current route (my total metric), I will include you as a **FEASIBLE SUCCESSOR**.

If your route (metric) to the network (RD to me) is MORE than my current route (my total metric), I will **NOT** include you as a **FEASIBLE SUCCESSOR**.

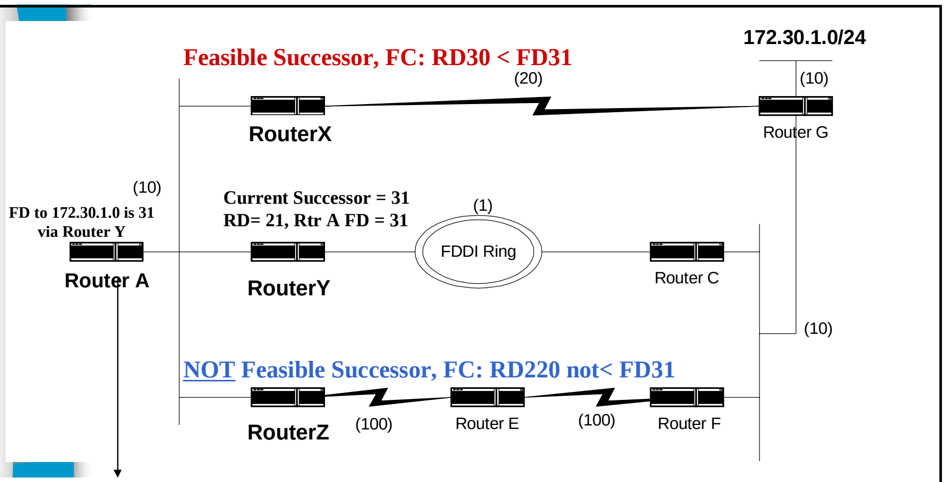


**Router A's Topology Table**

| Destination | Feasible Dist. | Advertised or Reported. Dist. | Neighbor | |
|---|---|---|---|---|
| 172.30.1.0 | 40 | 30 | X | In Topology Table |
| 172.30.1.0 | **31** | 21 | Y | <- **Routing Table** |
| 172.30.1.0 | 230 | 220 | Z | Not in Topology Table |

## What if Successor fails?

**Feasible Successor exists:**

- If current successor route fails, feasible successor becomes the current successor, i.e. the current route.
- Routing of packets continue with little delay.

**No Feasible Successor exists:**

- This may be because the Reported Distance (RD) is greater than the Feasible Distance (FD).
- Before this route can be installed, it must be placed in the ***active state*** and recomputed. (later)
- Routing of packets continue but with more of a delay.

65

## Convergence Using EIGRP

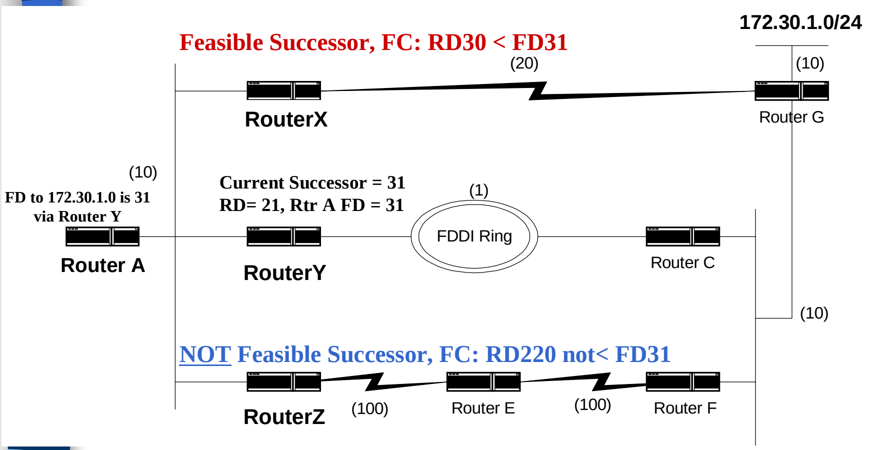


- RTY is successor with a computed cost of 31.
- “31” is the Feasible Distance (FD).
- RTX is a feasible successor because its RD is less than the FD.

66

# Convergence Using EIGRP

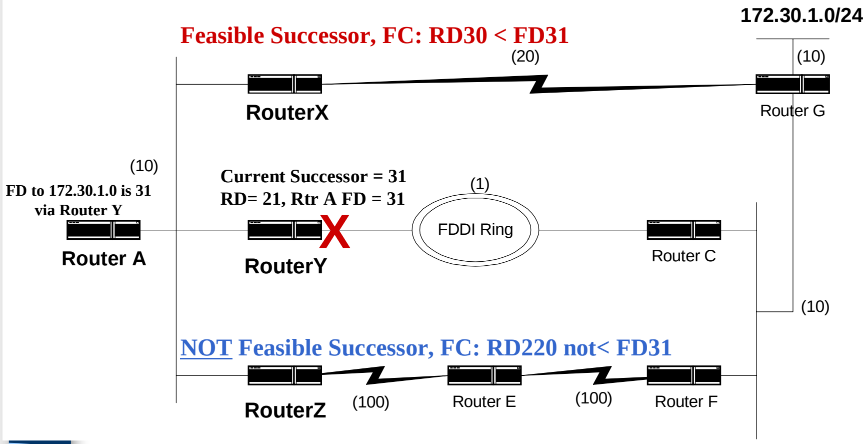


- Since RTX is a feasible successor, it is installed in the routing table immediately (no recomputation).
- It's RD (30) is less than the FD (31).

67

# Convergence Using EIGRP

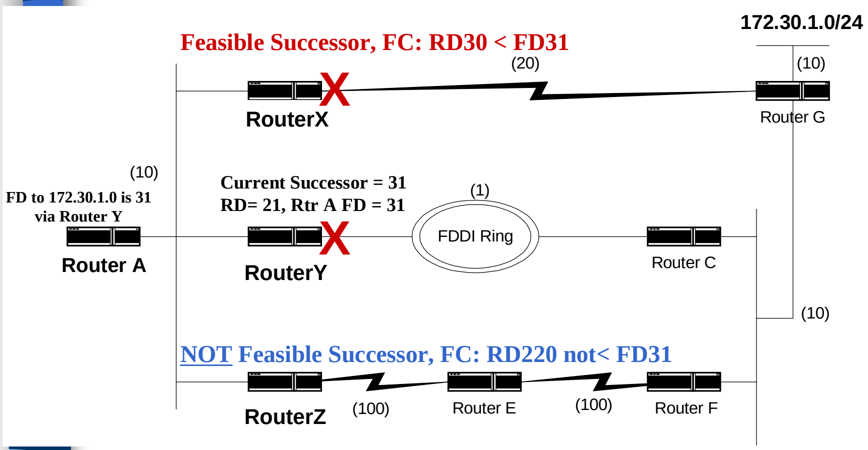


- RTZ is not a feasible successor.
- It's RD (220) is greater than the FD (31) for 172.30.1.0/24.
- Before this route can be installed, the route to net 24 must be placed in the ***active state*** and recomputed.
- Coming soon!

68

*Once again...*

***Topology table*** --- Each EIGRP router maintains a topology table for each configured network protocol. This table includes route entries for all destinations that the router has learned. All learned routes to a destination are maintained in the topology table.

- **show ip eigrp topology**
- (Feasible Distance/Reported Distance)
- 1 successor successors (routes) if FDs are different
  - smaller FD metric, that route is the the only successor
  - larger FD metric, those routes are possible feasible successor
- 2 or more successors (routes) if FDs are the same
  - Load balancing happens automatically

69

# Determining the costs

**What is the cost (metric) for 192.168.72.0/24 from SanJose2?**

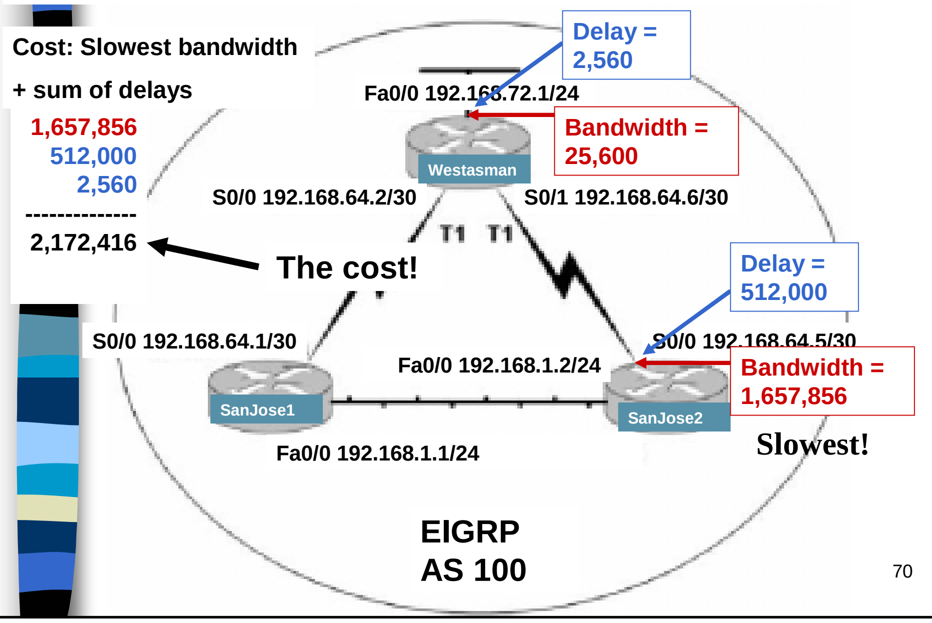


70

## Showing the cost in the Routing Table

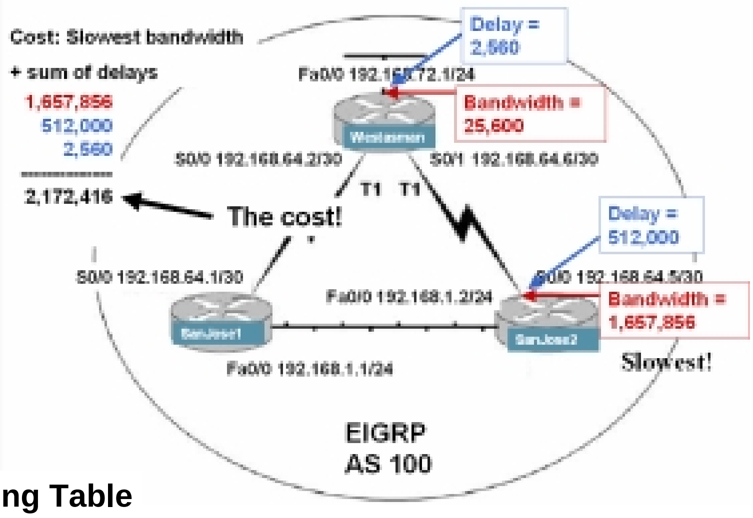


**Routing Table**

```
SanJose2#show ip route
D    192.168.72.0/24 [90/2172416]
               via 192.168.64.6, 00:28:26, Serial0
```

## Understanding the Topology Table

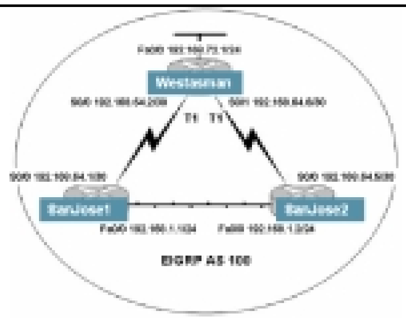


```
SanJose2#show ip eigrp top all
P 192.168.72.0/24, 1 successors, FD is 2172416, serno 93
         via 192.168.64.6 (2172416/28160), Serial0
         via 192.168.1.1 (2174976/2172416), FastEthernet0
P 192.168.64.0/30, 1 successors, FD is 2172416, serno 91
         via 192.168.1.1 (2172416/2169856), FastEthernet0
         via 192.168.64.6 (2681856/2169856), Serial0
P 192.168.64.4/30, 1 successors, FD is 2169856, serno 72
         via Connected, Serial0
P 192.168.1.0/24, 1 successors, FD is 28160, serno 1
         via Connected, FastEthernet0
```

## Understanding the Topology Table

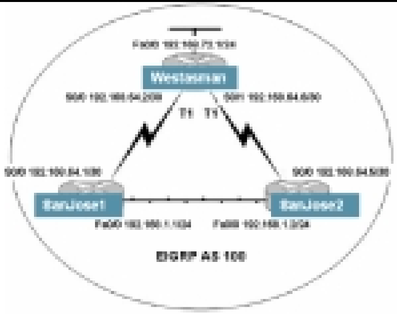


```
SanJose2#show ip eigrp top all-links
IP-EIGRP Topology Table for AS(100)/ID(192.168.64.5)

Codes: P - Passive, A - Active, U - Update, Q - Query, R
   - Reply,r - Reply status
```

Feasible distance

```
P 192.168.72.0/24, 1 successors, FD is 2172416  serno 22
      via 192.168.64.6 (2172416/28160), Serial0           successor
      via 192.168.1.1 (2174976/2172416), FastEthernet0    feasible successor
```

Feasible distance from this router if it was the successor.

```
SanJose2#show ip eigrp top all-links
IP-EIGRP Topology Table for AS(100)/ID(192.168.64.5)

Codes: P - Passive, A - Active, U - Update, Q - Query, R
   - Reply,r - Reply status
```

Feasible distance

```
P 192.168.72.0/24, 1 successors, FD is 2172416  serno 22
      via 192.168.64.6 (2172416/28160), Serial0           successor
      via 192.168.1.1 (2174976/2172416), FastEthernet0    feasible successor
```

Reported Distance

**Reported Distance**: This is the distance (cost) reported by the neighboring router, ie. Westasman and SanJose1.

- Westasman's Reported Distance is 28,160 = 25,600 (BW) + 2,560 (DLY)
- SanJose1's Reported Distance is 2,172,416 = 1,657,856 (BW) + 512,000 (DLY) + 2,560 (DLY)

**Feasible Successor**: Since the **Reported Distance** is **less than (or equal to)** the **Feasible Distance**, it is this neighbor is included as a **Feasible Successor**!

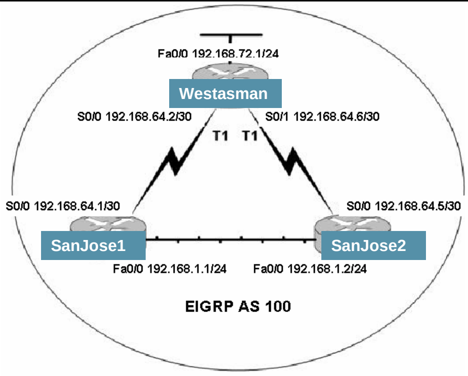


**Question**: What if there is only one entry in the Topology Table?

**Answer**: There is only a successor and no feasible successors.

**Why**?

The reported distance of that neighbor is greater than the current feasible distance. (You will see this in the lab.)

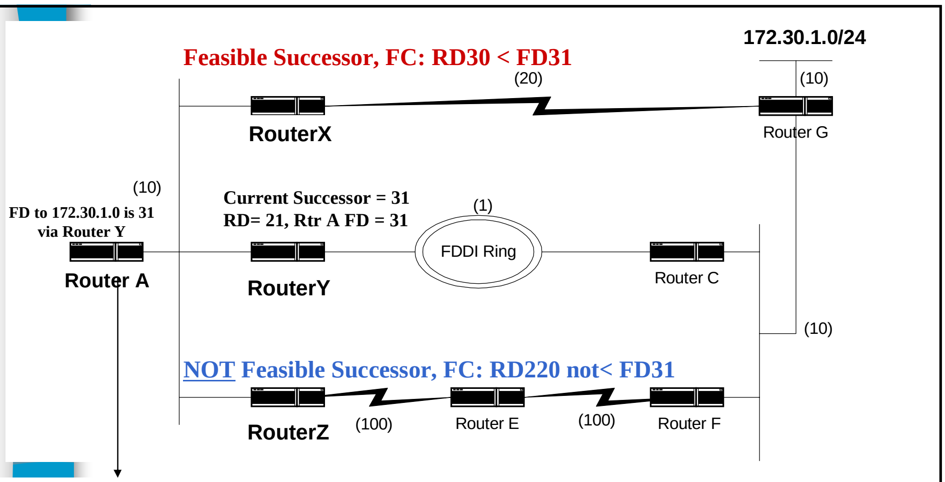


**Router A's Topology Table**

| Destination | Feasible Dist. | Advertised or Reported. Dist. | Neighbor |
|---|---|---|---|
| 172.30.1.0 | 40 | 30 | X **In Topology Table** |
| 172.30.1.0 | **31** | 21 | Y **<- Routing Table** |
| 172.30.1.0 | 230 | 220 | Z **Not in Topology Table** |

## Finding a feasible successor

No feasible successor in the topology table. EIGRP domain still finds another route (later).

```
SanJose2#debug eigrp fsm
EIGRP FSM Events/Actions debugging is on
SanJose2(config)#inter s 0
SanJose2(config-if)#shut
03:11:44: DUAL: Destination 192.168.72.0/24
03:11:44: DUAL: Find FS for dest 192.168.72.0/24. FD is
   2172416, RD is 2172416
03:11:44: DUAL:         192.168.64.6 metric
   4294967295/4294967295 not found Dmin is 4294967295
03:11:44: DUAL: Dest 192.168.72.0/24 entering active state.
```

Feasible successor is in the topology table. Backup route takes over right away.

```
Westasman#debug eigrp fsm
02:21:42: DUAL: Find FS for dest 192.168.64.4/30. FD is
   2169856, RD is 2169856
02:21:42: DUAL:         0.0.0.0 metric 2169856/0
02:21:42: DUAL:         192.168.64.1 metric
   4294967295/4294967295 found Dmin is 216985
```

## EIGRP Packet Types

The five EIGRP packet types are:

- Hello
- Acknowledgement
- Update (RTP)
- Query (RTP)
- Reply (RTP)

## EIGRP Hello Packets

- ***Hello packets to discover, verify, and rediscover neighbor routers.***
- EIGRP routers send hellos at a fixed (and configurable) interval, called the *hello interval*.
- The default hello interval depends on the bandwidth of the interface.
  - 5 seconds, hold time 15 seconds for T1 and faster
  - 60 seconds, hold time 180 seconds for slower than T1
- EIGRP hello packets are multicast.
- On IP networks, EIGRP routers send hellos to the multicast IP address **224.0.0.10**.
- If a neighbor is not heard from for the duration of the ***hold time (three times hello interval)***, EIGRP considers that neighbor down, and DUAL must step in to reevaluate the routing table.
  - By default, the hold time is three times the hello interval, but an administrator can configure both timers as desired.
- Unlike OSPF routers, EIGRP routers do not need to have the same hello intervals and hold down intervals.

79

## Acknowledgement Packets

- ***Acknowledgement packets, which are "data-less" hello packets, are used to ensure reliable communication.***
  - Unlike multicast hellos, acknowledgement packets are unicast.
  - Acknowledgements can be made by piggybacking on other kinds of EIGRP packets, such as reply packets.

80

# Update Packet

- ***Update packets are used when a router discovers a new neighbor.***
  - An EIGRP sends unicast update packets to that new neighbor so that it can add to its topology table.
  - More than one update packet may be needed to convey all of the topology information to the newly discovered neighbor.
- ***Update packets are also used when a router detects a topology change.***
- The EIGRP router sends a multicast update packet to all neighbors alerting them to the change.
- All update packets are sent reliably. (RTP)
- EIGRP routers exchange routing information the same way as other distance vector routing protocols, but do not send periodic updates.
- EIGRP updates are only sent when a network is added or removed from the topology database, when the successor for a given network changes, or when the locally used metric is updated.

81

# EIGRP Packet Format - FYI

| Bits | 8 | 8 | 16 | 32 | 32 | 32 | 32 | Variable |
|---|---|---|---|---|---|---|---|---|
| | Version | Opcode | Checksum | Flags | Sequence | ACK | AS Number | TLVs |

**General TLV**

| Bits | 16 | 16 | 8 | 8 | 8 | 8 | 8 | 8 | 16 |
|---|---|---|---|---|---|---|---|---|---|
| | Type 0x0001 | Length | K1 | K2 | K3 | K4 | K5 | Reserved | Hold Time |

**IP Internal Routes TLV**

| Bits | 16 | 16 | 32 | 32 | 32 | 24 | 8 | 8 | 8 | 16 | 8 | 24 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Type 0x0102 | Length | Next Hop | Delay | Bandwidth | MTU | Hop Count | Reliability | Load | Reserved | Prefix Length | Destination |

**IP External Routes TLV**

| Bits | 16 | 16 | 32 | 32 | 32 | 32 | 32 | 16 | 8 | 8 | 32 | 24 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Type 0x0103 | Length | Next Hop | Orig. router | Orig. AS | Tag | Ext. metric | Reserved | Ext. ID | Flags | Delay | Bandwidth |

| 24 | 8 | 8 | 8 | 16 | 8 | 24 |
|---|---|---|---|---|---|---|
| MTU | Hop Count | Reliability | Load | Reserved | Prefix Length | Destination |

**http://www.rware.demon.co.uk/eigrp.htm**

82

## EIGRP Packet Format - FYI

http://www.rware.demon.co.uk/eigrp.htm

- **Version** - there has only been one version
- **Opcode** - this is the EIGRP packet type:
  - **1** - Update
  - **3** - Query
  - **4** - Reply
  - **5** - Hello
  - **6** - IPX SAP
- **Checksum** - this is calculated for the whole EIGRP portion of the IP datagram
- **Flags** - The LSB (0x00000001) is the **Init** bit meaning that the route in this packet is the first in a new neighbor relationship. The next bit (0x00000002) is the **Conditional Receive** bit used in Cisco's Reliable Multicasting algorithm.
- **Sequence** - the 32-bit sequence number used by RTP.
- **ACK** - the 32-bit sequence last heard from the neighbor. A Hello packet with a non-zero value is an ACK.
- **AS Number** - the Autonomous System number of the EIGRP domain.
- **Type/Length/Value (TLV)** - There are a number of TLVs, all of them begin with a 16 bit Type field and a 16 bit Length field. There then follows a number of fields that vary depending on the type as given below.
- General TLVs
  - **0x0001** - General EIGRP parameters (applies to any EIGRP packet regardless of protocol)
  - **0x0003** - Sequence (used by Cisco's Reliable Multicast)
  - **0x0004** - EIGRP software version, the original version being **0** and the current version being **1** (used by Cisco's Reliable Multicast)
  - **0x0005** - Next Multicast Sequence (used by Cisco's Reliable Multicast)
- IP TLVs
  - **0x0102** - IP internal routes
  - **0x0103** - IP external routes
- AppleTalk TLVs
  - **0x0202** - AppleTalk internal routes
  - **0x0203** - AppleTalk external routes
  - **0x0204** - AppleTalk cable setup
- IPX TLVs
  - **0x0302** - IPX internal routes
  - **0x0303** - IPX external routes

## Query and Reply Packets

- ***EIGRP routers use query packets whenever it needs specific information from one, or all, of its neighbors.***
  - ***A reply packet is used to respond to a query.***
- If an EIGRP router loses its successor and cannot find a feasible successor for a route, DUAL places the route in the active state.
  - the router multicasts a query to all neighbors, searching for a successor to the destination network.
  - Neighbors must send replies that either provide information on successors, or indicate that no successor information is available.
- Queries can be multicast or unicast, while replies are always unicast.
- Both packet types are sent reliably. (RTP)

84

## Successors and Feasible Successors

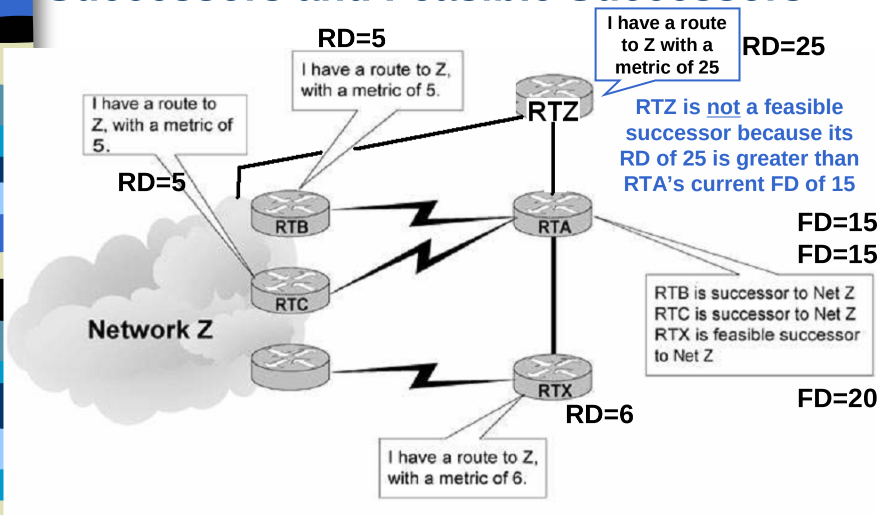




## Successors and Feasible Successors

- A router views its feasible successors as neighbors that are downstream, or closer, to the destination than it is.
- If something goes wrong with the successor, DUAL can quickly identify a feasible successor from the topology table, and install a new route to the destination.
- If no feasible successors to the destination exist, DUAL places the route in the active state.

- Entries in the topology table can be in one of two states: *active* or *passive*.
- A ***passive route*** is one that is stable and available for use.
- An ***active route*** is a route in the process of being recomputed by DUAL.

- **Recomputation** happens if a route becomes unavailable and DUAL can't find any feasible successors.
- Another route may exist, it is just that their Reported Distance was greater than your Feasible Distance.

## Query and Reply Packets

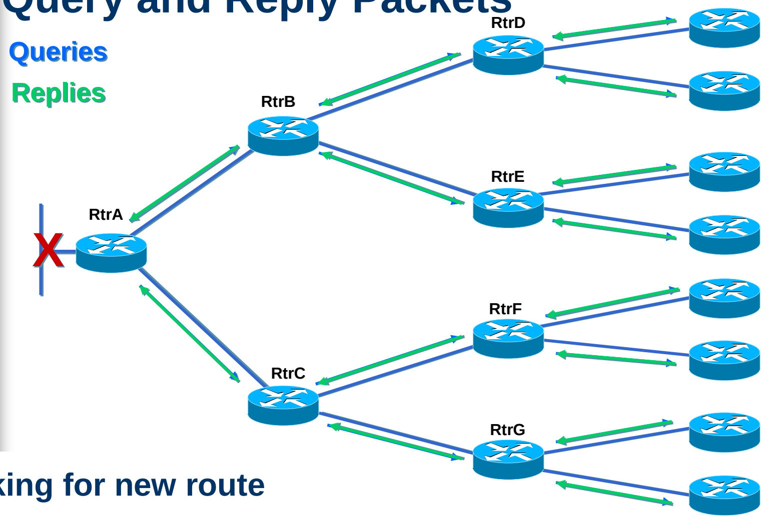


### Looking for new route

- If there were no Feasible Successors, the router must ask neighbors for help in hope of finding a new, loop-free path to the destination.
- Neighbor routers are compelled to reply to this query.
  - If a neighbor has a route, it will reply with information about the successor(s).
  - If not, the neighbor notifies the sender that it doesn't have a route to the destination either.

## Return Route or Forward Query

If a feasible successor does <u>not</u> exist:

- 1. The router flags the route as **active**.
- 2. The router looks for an alternate path by sending out a query packet to all neighbors to learn if they have a path to the given destination.
  - The query packets are multicast out every interface except the one which the dead link was learned, adhering to the split horizon rule.
- 3. If a neighbor does have a path that does not involve the querying router, or no path at all to the destination, it unicasts a reply with this information.
  - If a neighbor that receives the query is using the querying router as its feasible successor, then it multicasts its own query packet to its neighbor, which creates a "ripple effect" through the network until a new path is found or a major network boundary is met.
- 4. When the query router receives replies, it reacts based on the answer in the reply:
  - If the reply included a successor or feasible successor, the information is put into its topology table, and the querying router waits until all replies are received. It then recalculates the topology table, and adds the successr(s) to the routing table. The route returns to a passive state in the topolgy table and routing can continue.
  - If none of the replies includes a successor or feasible successor, the querying router removes the active route from its topology table and routing tables.
- If a neighbor router to which a query is sent does not reply within the active time of 180 seconds, EIGRP tears down the neighbor relationship with the offending router and puts routes learned from that router into an active state.

## Query and Reply Packets

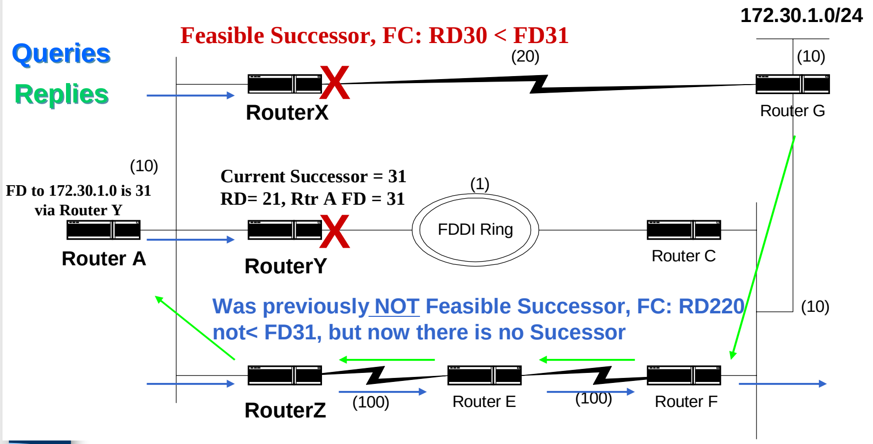


- RouterZ now becomes the successor for network 172.30.1.0/24

89

## Stuck in Active (SIA) Routes

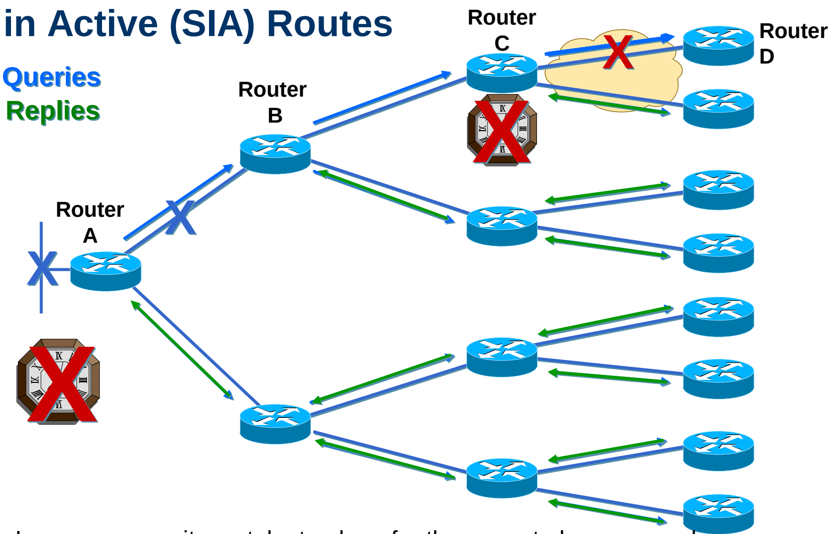


- In some cases, it can take too long for the query to be answered.
- When this happens, the router that issued the query gives up and resets its neighbor relationship with the router that didn't answer.
- The most basic situation where this occurs is when it simply takes too long for a query to reach the other end of the network and a reply to travel back.

90

## Stuck in Active (SIA) Routes

- Typically, SIAs results when a router cannot answer a query because:
  - the router is too busy to answer the query (generally high cpu utilization)
  - the router cannot allocate the memory to process the query or build the reply packet
  - the circuit between the two routers is not good (packet loss)
  - unidirectional links (a link on which traffic can only flow in one direction due to a failure)

91

## Stuck in Active (SIA) Routes

**Troubleshooting Steps:**

**Step 1**: find the routes which are consistently being reported as stuck in active.

- If you are logging console messages, a quick perusal of the log will indicate which routes are being marked as stuck in active most often.

**Step 2**: find out which routers are consistently failing to answer queries (not always easy).

- Use the **show ip eigrp topology active** command.
  - Any neighbors which have the **r** beside them are neighbors that the router is waiting on replies from
  - the active timer is how long the route has been active.
  - pay particular attention to routes that have replies outstanding and have been active for 2 to 3 minute

**Step 3**: find the reason why that router is not receiving or answering queries

- One you have found the router that is consistently not answering queries, look for problems on the link to this neighbor, memory or CPU utilization problems with this neighbor, etc.

92

## Stuck in Active (SIA) Routes

```
SanJose2#show ip eigrp top all-links
Codes: P - Passive, A - Active, U - Update, Q - Query,
  R - Reply, r - Reply status

A 10.2.4.0/24, 0 successors, FD is 512640000, Q
    1 replies, active 00:00:01, query-origin: Local
  origin
         via 10.1.2.2 (Infinity/Infinity), Serial1
    1 replies, active 00:00:01, query-origin: Local
  origin
         via 10.1.3.2 (Infinity/Infinity), r, Serial3
    Remaining replies:
         via 10.1.1.2, r, Serial0
```

- Any neighbors that show an **R** have yet to reply (the active timer shows how long the route has been active).
- Note that these neighbors may not show up in the Remaining replies section; they may appear among the other RDBs.
- Pay particular attention to routes that have outstanding replies and have been active for some time, generally two to three minutes.

## Fixing SIA

- Depending on the cause, you may have to restrict the "query range" of EIGRP routers.
  - You may have to redesign parts of the network
- Although not recommended, you can also increase the amount of time the router will wait after sending a query out before declaring the route stuck in active.
- This can be changed using the command:
  **timers active-time** ***time in minutes***

94

## Configuring EIGRP

- **Good link: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfeigrp.htm**

95

## EIGRP Configuration

```
RTA(config)#router eigrp AS
RTA(config-router)#network network
RTA(config-router)#network 1.0.0.0
```

AS must be the same on all routers in the AS

If AS's are identical on a router running both IGRP and EIGRP, routing tables are redistributed automatically.

IOS 12.0 added the stub command (separate line) for stub networks and allows for a wildcard mask in the network statement. Stub routers only announce only their directly connected networks or static routes to the hub routers.

```
router(config-router)# eigrp stub
```

96

## EIGRP Configuration

```
RTA(config)#router eigrp 123
RTA(config-router)#eigrp log-neighbor-changes
RTA(config-router)#network 1.0.0.0
```

Use the EIGRP **network** command as you would the IGRP or RIP network command.

Always issue the **eigrp log-neighbor-changes** command when first configuring EIGRP. Without this command, critical neighbor information will not be logged (console, buffers, syslog, etc). You will need this neighbor information in order to troubleshoot EIGRP.

97

## EIGRP Configuration

**Optional Interface Command:**

```
RTA(config-if)#ip bandwidth-percent eigrp 123 40
```

By default, EIGRP uses no more than 50% of a link's bandwidth for hellos, updates, queries, and acknowledgements.

EIGRP determines a link's bandwidth by using the bandwidth value assigned to the interface. On a low speed link or multipoint connection, you may need to tweak the default bandwidth percentage.

This example command configures EIGRP for AS 123 to use no more than 40% of the stated bandwidth.

98

## EIGRP Route Summarization

**Optional Interface Command:**

```
RTA(config-if)# no auto-summary
RTA(config-if)# ip summary-address eigrp
  autonomous-system-number ip-address mask
```

> You can configure a summary aggregate address for a specified interface, RTA(config-if), or for all interfaces, RTA(config-router). (Same command.)
>
> If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

99

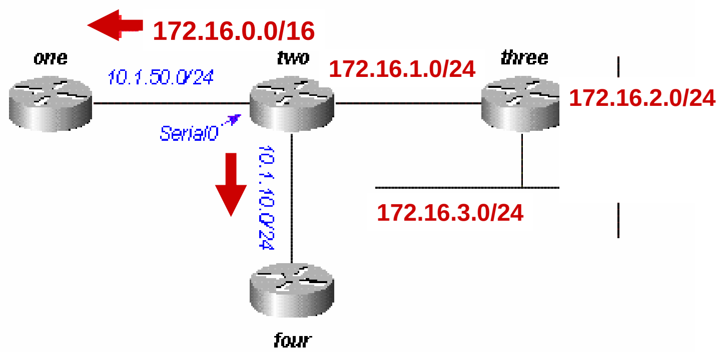


*Figure 14*

**Default**

- **By default, EIGRP summarizes at classful network boundaries, like RIP.**
- **Summarization can be disabled using: (RouterTwo)**

```
router eigrp 2000
    no auto-summary
```

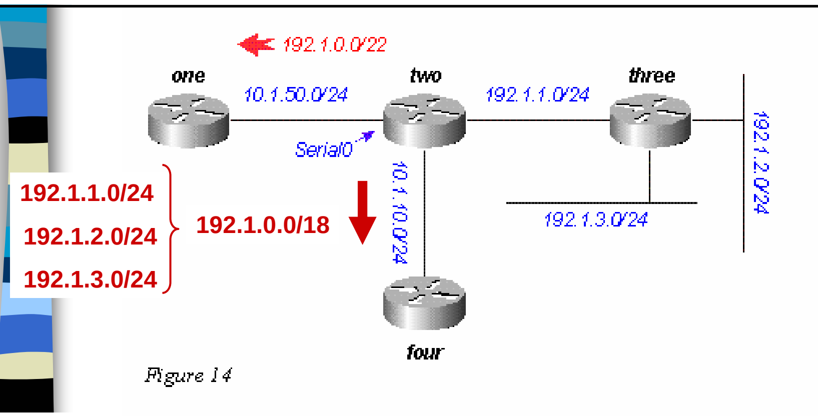

Figure 14

**ip summary-address (interface specific only) - choose your mask!**

```
interface Serial0
  ip address 10.1.50.1 255.255.255.0
  ip summary-address eigrp 2000 192.1.0.0 255.255.252.0
router eigrp 2000
  no auto-summary
  network ...
```

# EIGRP Variance Command

**Optional Interface Command:**

RTA(config-router)# **variance** *number*

IGRP and EIGRP also support unequal cost path load balancing, which is known as variance.

The **variance** command instructs the router to include routes with a metric less than or equal to ***n*** times the minimum metric route for that destination, where ***n*** is the number specified by the variance command.

**Note**: If a path isn't a feasible successor, then it isn't used in load balancing.

102

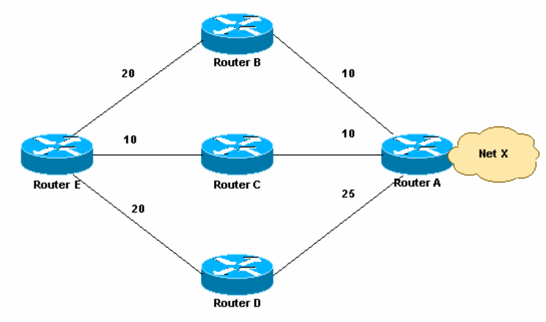


Let's look at an example, there are three ways to get to Network X, from Router E: (Note: metrics incorrectly not including outgoing interface to Net X.)

- E-B-A with a metric of 30
- E-C-A with a metric of 20 **<<<-- Minimum Metric**
- E-D-A with a metric of 45

103

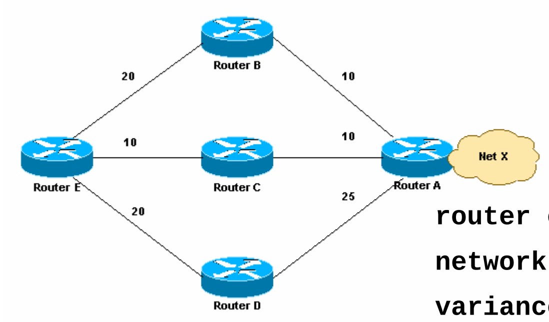


```
router eigrp 1
network x.x.x.x
variance 2
```

variance 2

- This increases the minimum metric to 40 (2 * 20 = 40).
- EIGRP includes all the routes that have a metric less than or equal to 40, **and** are feasible successors.
- In the above configuration, EIGRP now uses two paths to get to Network X, E-C-A and E-B-A, because both paths have a metric under 40.

104

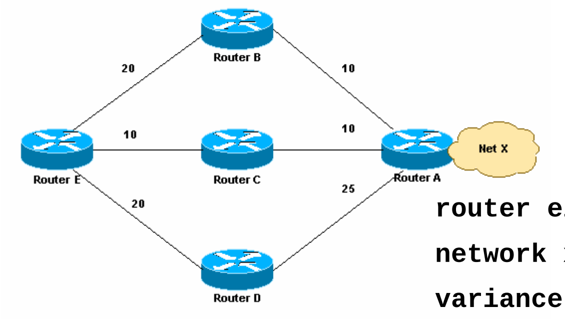


```
router eigrp 1
network x.x.x.x
variance 2
```

Two unequal-cost paths:

- E-B-A with a metric of 30  <<<-- Also below 2*20
- E-C-A with a metric of 20  **<<<-- Minimum Metric**
- E-D-A with a metric of 45

EIGRP doesn't use path E-D-A because it has a metric of 45, **and** it's not a feasible successor. Can have up to 6 unequal cost paths.

105

## Redistribution b/t EIGRP & IGRP

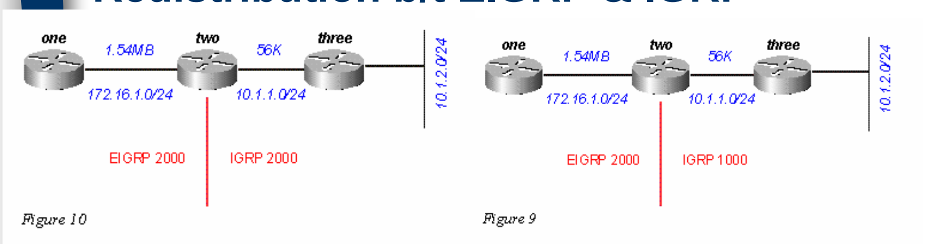


**Same AS numbers**

Router Two

```
router eigrp 2000
 network 172.16.1.0
!
router igrp 2000
 network 10.0.0.0
```

**(automatic redistribution)**

**Different AS numbers**

Router Two

```
router eigrp 2000
 redistribute igrp 1000
 network 172.16.1.0
!
router igrp 1000
 redistribute eigrp 2000
 network 10.0.0.0
```

**So what is the difference?  We'll see later when discussing Redistribution.**

## EIGRP and default routes

There are three ways to inject a default route into EIGRP:

- Redistribute a static route
- IP default-network
- Summarize to 0.0.0.0/0

107

## EIGRP and default routes

**Redistribute a static route**

- Use the first method when you want to draw all traffic to unknown destinations to a default route at the core of the network.
- This method is effective for advertising connections to the Internet. For example:

```
Gateway Router
ip route 0.0.0.0 0.0.0.0 x.x.x.x (next hop)
!
router eigrp 100
 redistribute static
 <text omitted>
```

108

## EIGRP and default routes

**Ip default-network**

- Propagates a default route to other routers, but needs to have a route or default route out once the packets arrive.

```
Gateway Router
router igrp 24
   <text omitted>
   network 207.21.20.0
ip route 0.0.0.0 0.0.0.0 207.21.20.1
ip default-network 207.21.20.0
```

109

## EIGRP and default routes

**Extra: Summarize to 0.0.0.0/0**

- Summarizing to a default route is effective only when you want to provide remote sites with a default route.
- Since summaries are configured per interface, you don't need to worry about using distribute-lists or other mechanisms to prevent the default route from being propagated toward the core of your network.

```
router eigrp 100
 network 10.0.0.0
!
interface serial 0
 ip address 10.1.1.1
 ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

110

# Cabrillo College

CISCO SYSTEMS
NETWORKING ACADEMY

CCNP – Advanced Routing

Ch. 6 EIGRP

*Rick Graziani, Instructor*

111