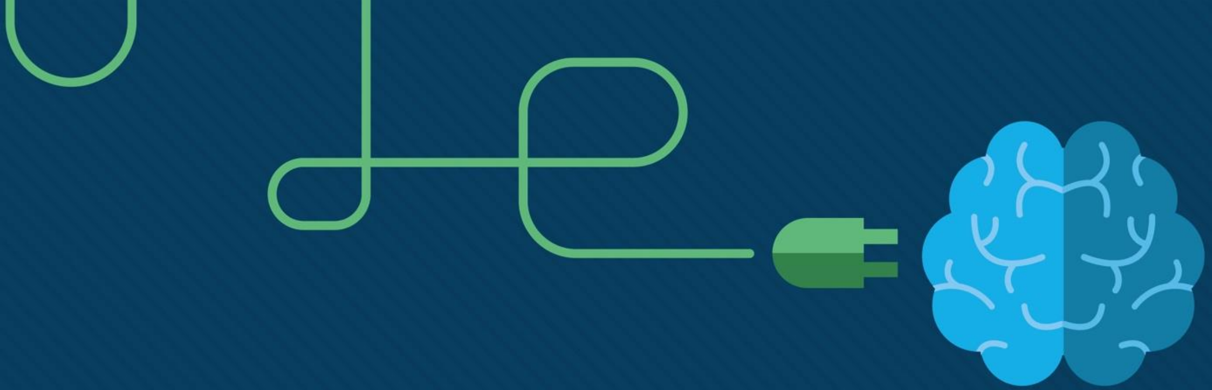

Chapter 7: EIGRP

Instructor Materials

CCNP Enterprise: Core Networking

Chapter 7 Content

This chapter covers the following content:

- **EIGRP Fundamentals -** This section explains how EIGRP establishes a neighbor adjacency with other routers and how routes are exchanged with other routers.
- **Path Metric Calculation -** This section explains how EIGRP calculates the path metric to identify the best and alternate loop-free paths.
- **Failure Detection and Timers -** This section explains how EIGRP detects the absence of a neighbor and the convergence process.
- **Route Summarization -** This section explains the logic and configuration related to summarizing routes on a router.

EIGRP Fundamentals

- EIGRP overcomes the deficiencies of other distance vector routing protocols with unequal-cost load balancing, support for networks 255 hops away, and rapid convergence features.
- EIGRP uses a diffusing update algorithm (DUAL) to identify network paths and enable fast convergence using precalculated loop-free backup paths.
- EIGRP adds to the route selection algorithm logic that uses factors outside hop count.

EIGRP Fundamentals
Autonomous Systems

A router can run multiple EIGRP processes. Each process operates under the context of an autonomous system, which represents a common routing domain. Routers within the same domain use the same metric calculation formula and exchange routes only with members of the same autonomous system, as shown in the figure.

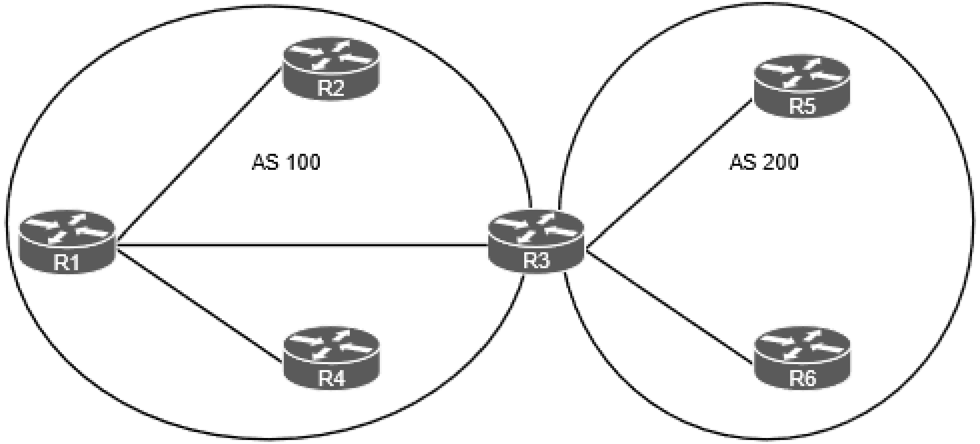

Figure 7-1 *EIGRP Autonomous Systems*

EIGRP Fundamentals

EIGRP Terminology

Term	Definition
Successor route	The route with the lowest path metric to reach a destination. The successor route for R1 to reach 10.4.4.0/24 on R4 is R1→R3→R4.
Successor	The first next-hop router for the successor route. The successor for 10.4.4.0/24 is R3.

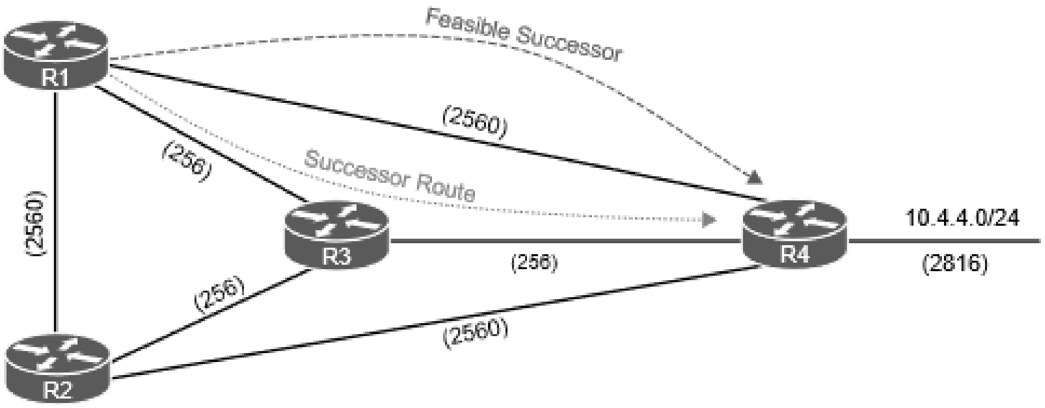

Figure 7-2 *EIGRP Reference Topology*

EIGRP Fundamentals
EIGRP Terminology (Cont.)

Term	Definition
Feasible distance (FD)	The metric value for the lowest-metric path to reach a destination. The FD calculated by R1 for the 10.4.4.0/24 network is 3328 (that is, 256+256+2816).
Reported distance (RD)	The distance reported by a router to reach a prefix. The reported distance value is the feasible distance for the advertising router. R3 advertises the 10.4.4.0/24 prefix with an RD of 3072. R4 advertises the 10.4.4.0/24 to R1 and R2 with an RD of 2816.

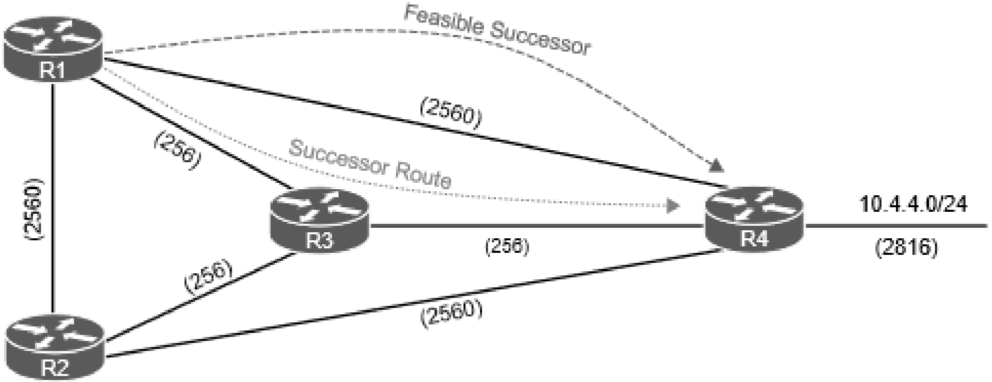

Figure 7-2 *EIGRP Reference Topology*

EIGRP Fundamentals

EIGRP Terminology (Cont.)

Term	Definition
Feasibility condition	A condition under which, for a route to be considered a backup route, the reported distance received for that route must be less than the feasible distance calculated locally. This logic guarantees a loop-free path.
Feasibility successor	A route that satisfies the feasibility condition and is maintained as a backup route. The route R1→R4 is the feasible successor because the RD 2816 is lower than the FD 3328 for the R1→R3→R4 path.

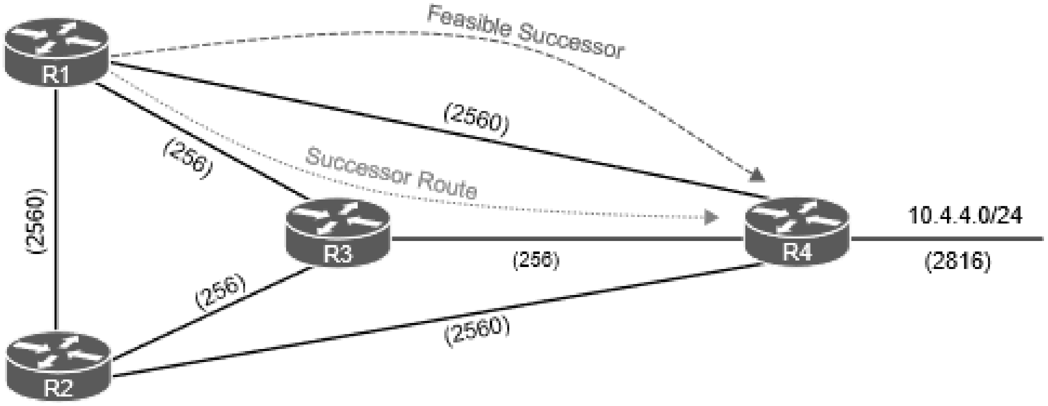

Figure 7-2 *EIGRP Reference Topology*

EIGRP Fundamentals

Topology Table

EIGRP contains a topology table that makes it different from a “true” distance vector routing protocol. EIGRP’s topology table is part of DUAL and helps to identify loop-free backup routes. It contains all the network prefixes advertised within an EIGRP autonomous system.

Each entry in the table contains the following:

- Network prefix
- EIGRP neighbors that have advertised that prefix
- Metrics from each neighbor (for example, reported distance, hop count)
- Values used for calculating the metric (for example, load, reliability, total delay, minimum bandwidth)

EIGRP Fundamentals
Topology Table (Cont.)

The show topology command will display both successor and feasible successors.

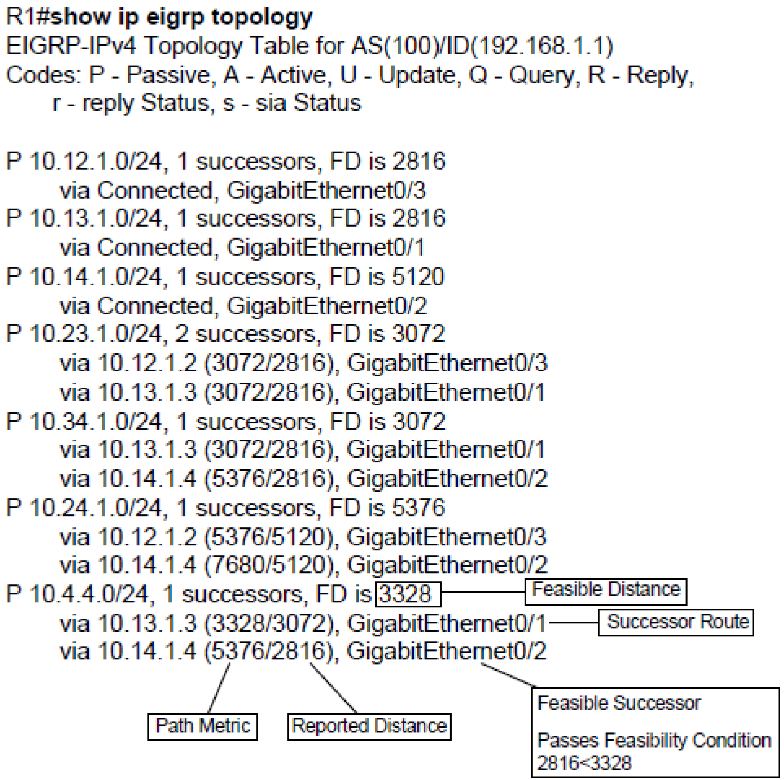

```
R1#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.12.1.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/3
P 10.13.1.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet0/1
P 10.14.1.0/24, 1 successors, FD is 5120
        via Connected, GigabitEthernet0/2
P 10.23.1.0/24, 2 successors, FD is 3072
        via 10.12.1.2 (3072/2816), GigabitEthernet0/3
        via 10.13.1.3 (3072/2816), GigabitEthernet0/1
P 10.34.1.0/24, 1 successors, FD is 3072
        via 10.13.1.3 (3072/2816), GigabitEthernet0/1
        via 10.14.1.4 (5376/2816), GigabitEthernet0/2
P 10.24.1.0/24, 1 successors, FD is 5376
        via 10.12.1.2 (5376/5120), GigabitEthernet0/3
        via 10.14.1.4 (7680/5120), GigabitEthernet0/2
P 10.4.4.0/24, 1 successors, FD is 3328
        via 10.13.1.3 (3328/3072), GigabitEthernet0/1
        via 10.14.1.4 (5376/2816), GigabitEthernet0/2
```

Figure 7-3 *EIGRP Topology Output*

EIGRP Fundamentals

EIGRP Neighbors

EIGRP neighbors exchange the entire routing table when forming an adjacency, and they advertise only incremental updates as topology changes occur within a network. EIGRP uses five different packet types to communicate with other routers.

Type	Packet Name	Function
1	Hello	Used for discovery of EIGRP neighbors and for detecting when a neighbor is no longer available
2	Request	Used to get specific information from one or more neighbors
3	Update	Used to transmit routing and reachability information with other EIGRP neighbors
4	Query	Sent out to search for another path during convergence
5	Reply	Sent in response to a query packet

Table 7-3 EIGRP Packet Types

Path Metric Calculation

- Metric calculation is a critical component for any routing protocol. EIGRP uses multiple factors to calculate the metric for a path.

Path Metric Calculation

EIGRP Classic Metric Formula

Metric calculation uses bandwidth and delay by default, but it can include interface load and reliability, too. The formula shown in Figure 7-4 illustrates the EIGRP classic metric formula.

$$\text{Metric} = \left[\left(K_1 * \text{BW} + \frac{K_2 * \text{BW}}{256 - \text{Load}} + K_3 * \text{Delay}\right) * \frac{K_5}{K_4 + \text{Reliability}}\right]$$

Figure 7-4 *EIGRP Classic Metric Formula*

Path Metric Calculation
EIGRP Classic Metric Formula with Definitions

- EIGRP uses K values to define which factors the formula uses and the associated impact of a factor when calculating the metric.
- BW represents the slowest link in the path scaled to a 10 Gbps link. Link speed is collected from the configured interface bandwidth on an interface.
- Delay is the total measure of delay in the path, measured in tens of microseconds (μs).
- The EIGRP formula is based on the IGRP metric formula, except the output is multiplied by 256 to change the metric from 24 bits to 32 bits.

$$\text{Metric} = 256 * \left[\left(K_1 * \frac{10^7}{\text{Min. Bandwidth}} + \frac{K_2 * \frac{10^7}{\text{Min. Bandwidth}}}{256 - \text{Load}} + \frac{K_3 * \text{Total Delay}}{10} \right) * \frac{K_5}{K_4 + \text{Reliability}} \right]$$

Figure 7-5 *EIGRP Classic Metric Formula with Definitions*

Path Metric Calculation

EIGRP Classic Metric Formula with Default K Values

By default, K_1 and K_3 have the value 1, and K_2, K_4, and K_5 are set to 0.

Figure 7-6 places default K values into the formula and then shows a streamlined version of the formula.

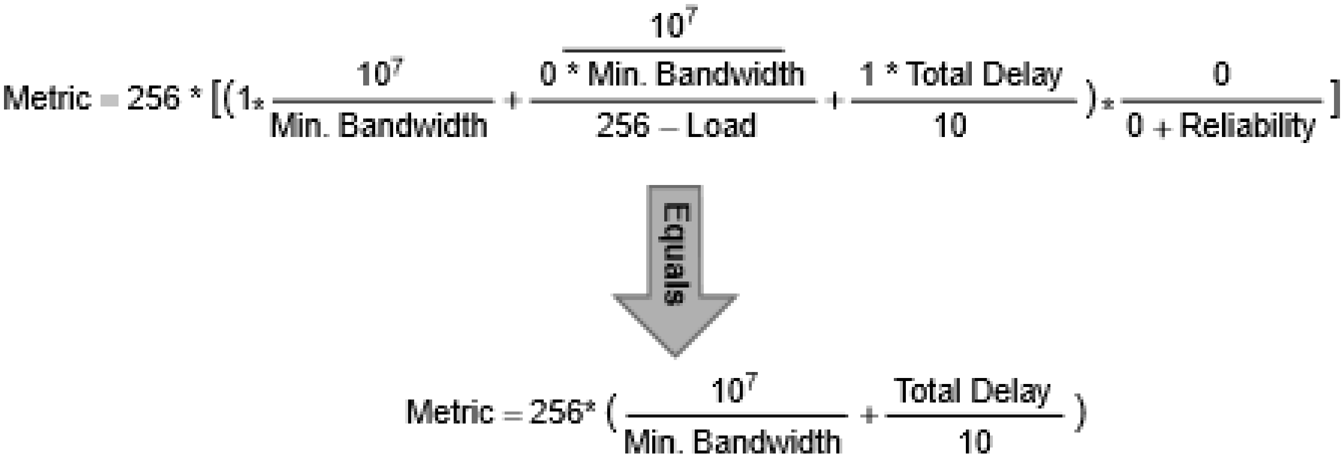

$$\text{Metric} = 256 * \left[\left(1 * \frac{10^7}{\text{Min. Bandwidth}} + \frac{\frac{10^7}{0 * \text{Min. Bandwidth}}}{256 - \text{Load}} + \frac{1 * \text{Total Delay}}{10}\right) * \frac{0}{0 + \text{Reliability}}\right]$$

$$\text{Metric} = 256 * \left(\frac{10^7}{\text{Min. Bandwidth}} + \frac{\text{Total Delay}}{10}\right)$$

Figure 7-6 *EIGRP Classic Metric Formula with Default K Values*

Path Metric Calculation

EIGRP Attribute Propagation

- The EIGRP update packet includes path attributes associated with each prefix.
- The EIGRP path attributes can include hop count, cumulative delay, minimum bandwidth link speed, and RD.
- The attributes are updated each hop along the way, allowing each router to independently identify the shortest path.

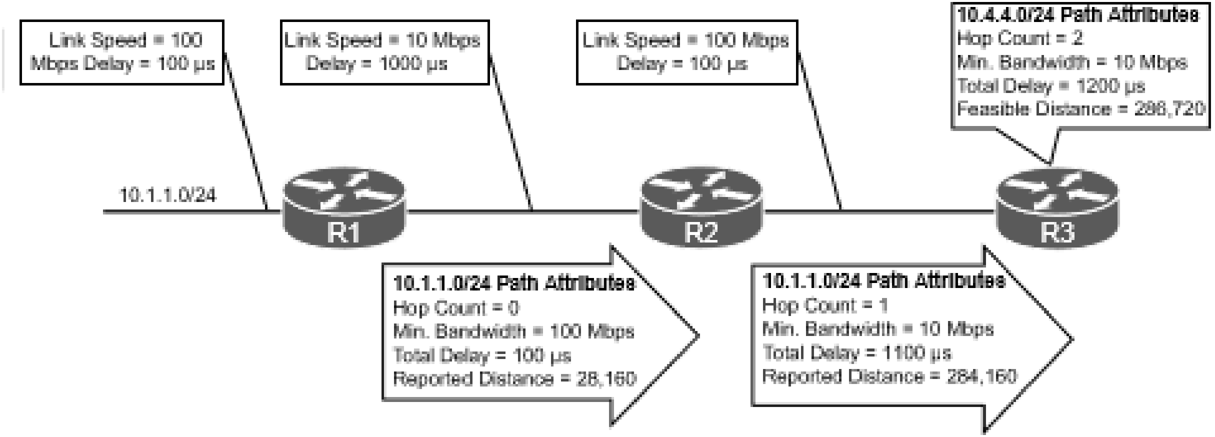

Figure 7-7 *EIGRP Attribute Propagation*

Path Metric Calculation
Default EIGRP Interface Metrics for Classic Metrics

Interface Type	Link Speed (kbps)	Delay	Metric
Serial	64	20,000 μs	40,512,000
T1	1544	20,000 μs	2,170,031
Ethernet	10,000	1000 μs	281,600
FastEthernet	100,000	100 μs	28,160
GigabitEthernet	1,000,000	10 μs	2816
10 GigabitEthernet	10,000,000	10 μs	512

Table 7.4 Default EIGRP Interface Metrics for Classic Metrics

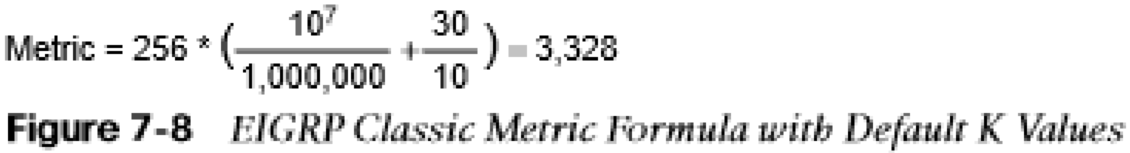

$$\text{Metric} = 256 * \left(\frac{10^7}{1{,}000{,}000} + \frac{30}{10}\right) = 3{,}328$$

Figure 7-8 *EIGRP Classic Metric Formula with Default K Values*

Path Metric Calculation

Wide Metrics

The original EIGRP specifications measured delay in 10 µs units and bandwidth in kilobytes per second, which did not scale well with higher-speed interfaces.

Example 7-1 provides some metric calculations for common LAN interface speeds. Notice that there is not a differentiation between an 11 Gbps interface and a 20 Gbps interface. The composite metric stays at 256, despite having different bandwidth rates.

Calculating Metrics for Common LAN Interface Speeds

GigabitEthernet: Scaled Bandwidth = 10,000,000 / 1000000
Scaled Delay = 10 / 10   Composite Metric = 10 + 1 x 256 = 2816

10 GigabitEthernet: Scaled Bandwidth = 10,000,000 / 10000000
Scaled Delay = 10 / 10   Composite Metric = 1 + 1 x 256 = 512

11 GigabitEthernet: Scaled Bandwidth = 10,000,000 / 11000000
Scaled Delay = 10 / 10   Composite Metric = 1 + 1 x 256 = 256

20 GigabitEthernet: Scaled Bandwidth = 10,000,000 / 20000000
Scaled Delay = 10 / 10   Composite Metric = 1 + 1 x 256 = 256

Example 7.1 Calculating Metrics for Common LAN Interface Speeds

Path Metric Calculation

EIGRP Wide Metrics Formula

EIGRP includes support for a second set of metrics, known as wide metrics, that addresses the issue of scalability with higher-capacity interfaces. Figure 7-9 shows the explicit EIGRP wide metrics formula.

Notice that an additional K value (K6) is included that adds an extended attribute to measure jitter, energy, or other future attributes.

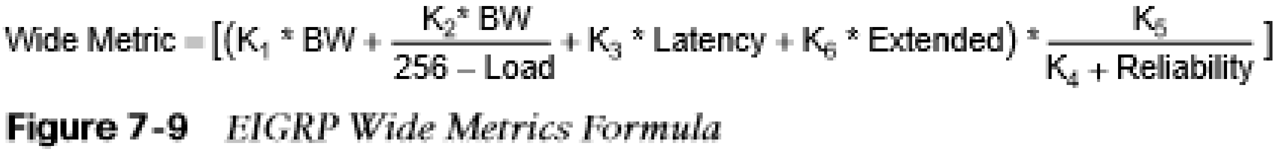

$$\text{Wide Metric} = \left[\left(K_1 * BW + \frac{K_2 * BW}{256 - \text{Load}} + K_3 * \text{Latency} + K_6 * \text{Extended}\right) * \frac{K_5}{K_4 + \text{Reliability}}\right]$$

Figure 7-9 *EIGRP Wide Metrics Formula*

Path Metric Calculation

EIGRP Wide Metrics Formula with Definitions

EIGRP wide metrics scale by 65,535 to accommodate higher-speed links. This provides support for interface speeds up to 655 Tbps (65,535 $\times 10^7$) without any scalability issues.

Figure 7-10 displays the updated formula that takes into account the conversions in latency and scalability.

$$\text{Wide Metric} = 65{,}535 * \left[\left(\frac{K_1 * 10^7}{\text{Min. Bandwidth}} + \frac{\frac{K_2 * 10^7}{\text{Min. Bandwidth}}}{256 - \text{Load}} + \frac{K_3 * \text{Latency}}{10^{-6}} + K_6 * \text{Extended}\right) * \frac{K_5}{K_4 + \text{Reliability}}\right]$$

Figure 7-10 *EIGRP Wide Metrics Formula with Definitions*

Path Metric Calculation

Metric Backward Compatibility

With EIGRP wide metrics, K_1 and K_3 are set to a value of 1, and K_2, K_4, K_5, and K_6 are set to 0, which allows backward compatibility because the K value metrics match with classic metrics.

EIGRP is able to detect when peering with a router is using classic metrics, and it unscales a metric from the formula in Figure 7-11.

$$\text{Unscaled Bandwidth} = \left(\frac{\text{EIGRP Bandwidth} * \text{EIGRP Classic Scale}}{\text{Scaled Bandwidth}} \right)$$

Figure 7-11 *Formula for Calculating Unscaled EIGRP Metrics*

Path Metric Calculation
Load Balancing

Installing multiple paths into the RIB for the same prefix is called equal-cost multipathing (ECMP).

EIGRP supports unequal-cost load balancing by changing EIGRP's variance multiplier. The EIGRP variance value is the feasible distance (FD) for a route multiplied by the EIGRP variance multiplier. Any feasible successor's FD with a metric below the EIGRP variance value is installed into the RIB.

Dividing the feasible successor metric by the successor route metric provides the variance multiplier. The variance multiplier is a whole number, so any remainders should always round up.

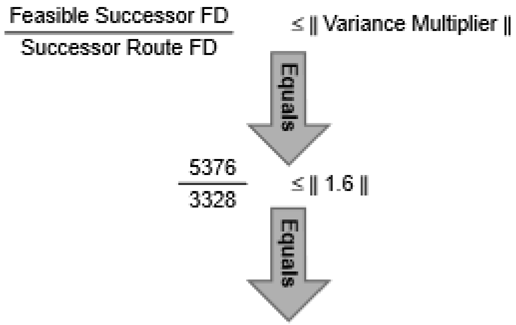

Figure 7-12 *EIGRP Variance Multiplier Formula*

Path Metric Calculation

Load Balancing (Cont.)

Verify unequal load balancing with the **show ip route** command.

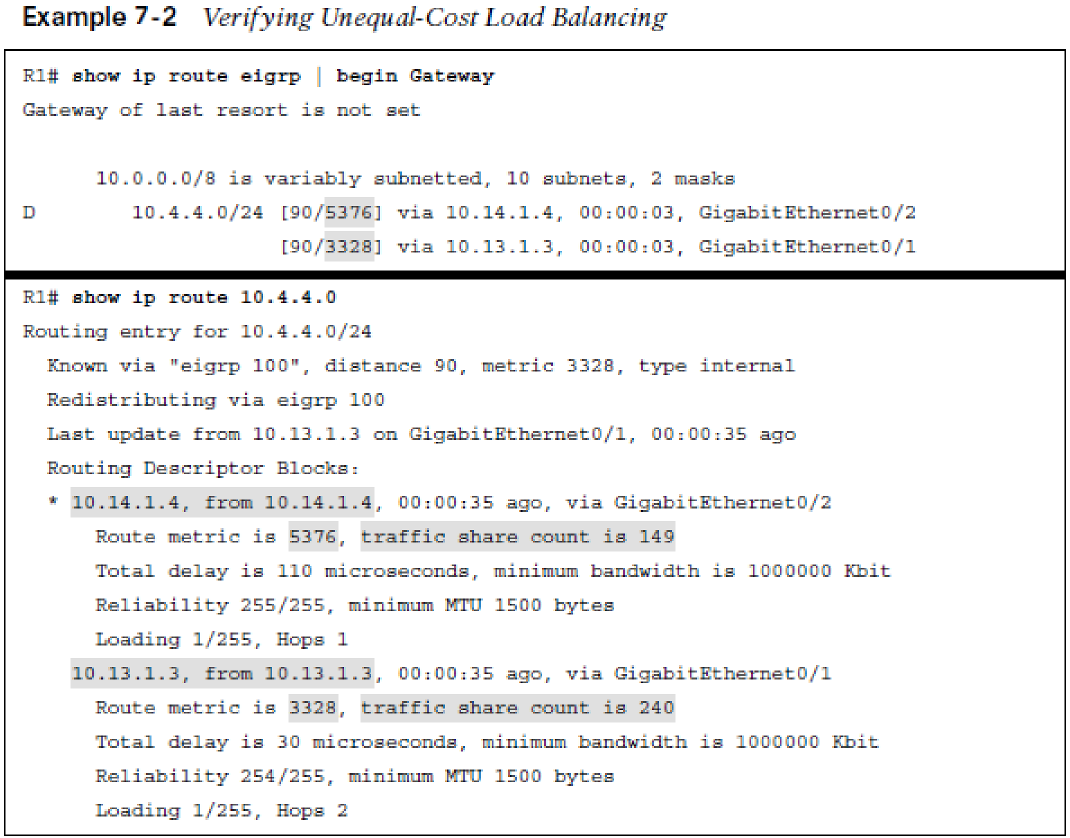

Example 7-2 *Verifying Unequal-Cost Load Balancing*

```
R1# show ip route eigrp | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D        10.4.4.0/24 [90/5376] via 10.14.1.4, 00:00:03, GigabitEthernet0/2
                     [90/3328] via 10.13.1.3, 00:00:03, GigabitEthernet0/1
```

```
R1# show ip route 10.4.4.0
Routing entry for 10.4.4.0/24
  Known via "eigrp 100", distance 90, metric 3328, type internal
  Redistributing via eigrp 100
  Last update from 10.13.1.3 on GigabitEthernet0/1, 00:00:35 ago
  Routing Descriptor Blocks:
  * 10.14.1.4, from 10.14.1.4, 00:00:35 ago, via GigabitEthernet0/2
      Route metric is 5376, traffic share count is 149
      Total delay is 110 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    10.13.1.3, from 10.13.1.3, 00:00:35 ago, via GigabitEthernet0/1
      Route metric is 3328, traffic share count is 240
      Total delay is 30 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 254/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

Failure Detection and Timers

- A secondary function for the EIGRP hello packets is to ensure that EIGRP neighbors are still healthy and available. EIGRP hello packets are sent out in intervals determined by the hello timer.
- EIGRP uses a second timer for the hold time, which is the amount of time EIGRP deems the router reachable and functioning. If the hold time reaches 0, EIGRP declares the neighbor unreachable and notifies DUAL of a topology change.
- The default EIGRP hello timer is 5 seconds, but it is 60 seconds on slow-speed interfaces (T1 or lower). The hold time value defaults to 3 times the hello interval.

Failure Detection and Timers

Convergence

When an EIGRP neighbor moves to a down state, paths are recomputed for any prefix where that EIGRP neighbor was a successor (upstream router).

When EIGRP detects that it has lost its successor for a path, the feasible successor instantly becomes the successor route. The router sends out an update packet for that path because of the new EIGRP path metrics. Downstream routers run their own DUAL for any impacted prefixes to account for the new EIGRP metrics. Figure 7-13 demonstrates such a scenario when the link between R1 and R3 fails.

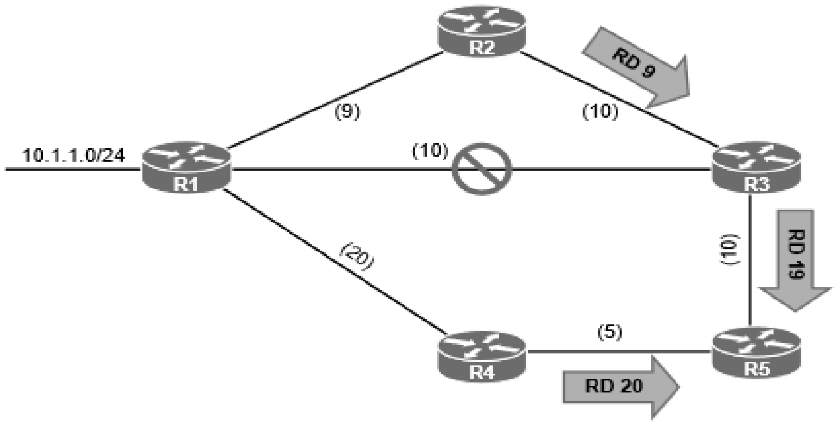

Figure 7-13 *EIGRP Topology with Link Failure*

Failure Detection and Timers

Convergence (Cont.)

If a feasible successor is not available for a prefix, DUAL must perform a new route calculation. The route state changes from passive (P) to active (A) in the EIGRP topology table.

The router detecting the topology change sends out query packets to EIGRP neighbors for the route. The query packet includes the network prefix with the delay set to infinity so that other routers are aware that it has gone active.

Upon receipt of a query packet, an EIGRP router does one of the following:

- It might reply to the query that the router does not have a route to the prefix.
- If the query did not come from the successor for that route, it detects the delay set for infinity but ignores it because it did not come from the successor. The receiving router replies with the EIGRP attributes for that route.
- If the query came from the successor for the route, the receiving router detects the delay set for infinity, sets the prefix as active in the EIGRP topology, and sends out a query packet to all downstream EIGRP neighbors for that route.

Failure Detection and Timers

Convergence (Cont.)

The query process continues from router to router until a router establishes the query boundary. A query boundary is established when a router does not mark the prefix as active, meaning that it responds to a query as follows:

- It says it does not have a route to the prefix.
- It replies with EIGRP attributes because the query did not come from the successor.

Failure Detection and Timers

Convergence (Cont.)

When a router receives a reply for every downstream query that was sent out, it completes the DUAL, changes the route to passive, and sends a reply packet to any upstream routers that sent a query packet to it. Figure 7-14 shows a topology where the link between R1 and R2 has failed.

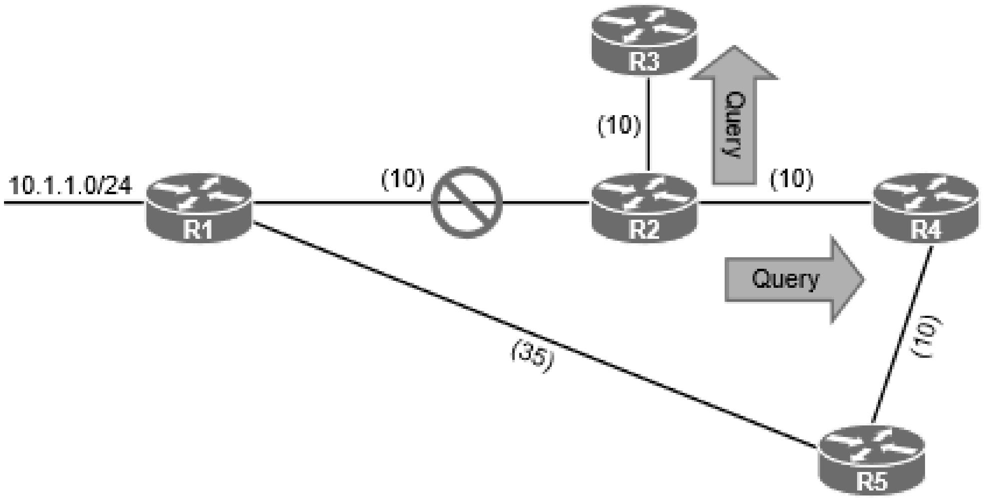

Figure 7-14 *EIGRP Convergence Topology*

Route Summarization

- Scalability of an EIGRP autonomous system depends on summarization. As the size of an EIGRP autonomous system increases, convergence may take longer. Scaling an EIGRP topology requires summarizing routes in a hierarchical fashion.

Route Summarization

EIGRP Summarization

EIGRP summarizes network prefixes on an interface basis. A summary aggregate is configured for the EIGRP interface. Prefixes within the summary aggregate are suppressed, and the summary aggregate prefix is advertised in lieu of the original prefixes. The summary aggregate prefix is not advertised until a prefix matches it. Interface-specific summarization can be performed in any portion of the network topology.

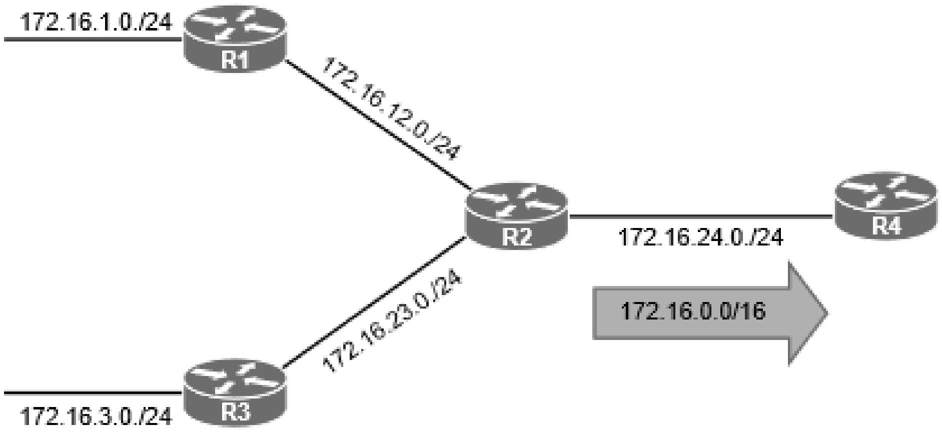

Figure 7-15 *EIGRP Summarization*

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 7

Description
EIGRP Terminology
Topology Table
EIGRP Packet Types
EIGRP Attribute Propagation
EIGRP Wide Metrics Formula
EIGRP unequal-cost load balancing
Convergence
Active route state

Prepare for the Exam
Key Terms for Chapter 7

Key Term	
Autonomous system	Reported distance
Feasible distance	Successor
Feasibility condition	Successor route
Feasibility successor	Summarization
Hello packets	Topology table
Hello timer	Variance value
K values	Wide metric