

EIGRP Principles

Enhanced Interior Gateway Routing Protocol (EIGRP) is a distance vector protocol that uses the same sophisticated metric as IGRP and uses the Diffusing Update Algorithm (DUAL) convergence algorithm. EIGRP is able to converge quickly and use little bandwidth because it—like OSPF—has separate keepalives and reliable updates. EIGRP is sometimes referred to as a *hybrid routing protocol*, although *advanced distance vector routing protocol* is probably a more accurate description.

EIGRP is an efficient but proprietary solution for large networks. Its ability to scale is limited only by the design of the network.

This chapter covers EIGRP used with IPv4. You can find a short discussion of EIGRP support for IPv6 in Chapter 21, “IPv6 Routing Protocols, Configuration, and Transitioning from IPv4,” but the concepts of EIGRP are essentially the same.

This chapter provides a theoretical description of how EIGRP works. The operation of EIGRP, available options, and design considerations are explained in this chapter, particularly in reference to scaling EIGRP and its use over a nonbroadcast multiaccess (NBMA) WAN environment.

The topics in this chapter directly reflect questions on the BSCI exam. EIGRP is designed for use in large networks. As a proprietary routing protocol for Cisco, it is therefore an obligatory subject in a Cisco exam on IP routing protocols.

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide which parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions.

The 13-question quiz, derived from the major sections in the “Foundation Topics” portion of the chapter, helps you determine how to spend your limited study time.

Table 3-1 outlines the major topics discussed in this chapter and the corresponding quiz questions.

Table 3-1 “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundation Topics Section	Questions Covered in This Section	Score
EIGRP Features and Advantages	1–7	
Understanding EIGRP Tables	8–12	
EIGRP Network Design	13	
Total Score		

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark the question wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guessed skews your results and might provide you with a false sense of security.

1. Which of the following describes how a reply to an EIGRP query is sent?
 - a. Multicast
 - b. Best effort unicast
 - c. A reliable unicast
 - d. A reliable multicast

2. Which of the following are advantages to using EIGRP as a routing protocol instead of RIP?
 - a. Rapid convergence
 - b. Reduced bandwidth consumption
 - c. Link state routing
 - d. Increased vendor support

3. Cisco identifies four main components of EIGRP. Select the two correct components from the following list.
 - a. Neighbor discovery
 - b. SPF algorithm
 - c. Areas
 - d. RTP

4. Given the default constant values, which variables are included in the EIGRP metric?
 - a. Bandwidth
 - b. Load

- c. Reliability
 - d. Delay
5. For EIGRP to form a neighbor relationship, which of the following must be true?
- a. Authentication must be enabled.
 - b. The K-values of the metric must be the same on both routers.
 - c. The autonomous system number must be the same on both routers.
 - d. The holddown timer must be the same.
6. What do the letters SIA stand for?
- a. Stuck in Active
 - b. Shortest IP Address
 - c. Stuck in Area
 - d. Simple IP Address
7. What is the EIGRP reserved Multicast address?
- a. 224.0.0.1
 - b. 224.0.0.5
 - c. 224.0.0.9
 - d. 224.0.0.10
8. The neighbor table contains which of the following information?
- a. The administrative distance
 - b. The metric of neighbors
 - c. The feasible distance between neighbors
 - d. The status of the links between neighbors
9. What do the letters SRTT stand for?
- a. Shortest Remote-Trip Time
 - b. Smooth Round-Trip Time
 - c. Shortest Reliable-Trip Time
 - d. Single Remote Timer Test

10. Which of the following would trigger the topology table to recalculate?
 - a. LSP received
 - b. SRT packet received
 - c. A new router coming online
 - d. Link loss detected

11. Feasible successor information is stored in which table?
 - a. Topology
 - b. Routing
 - c. Neighbor
 - d. Autonomous System

12. What is EIGRP Active Mode?
 - a. When the router is actively forwarding
 - b. When the router is actively recalculating paths
 - c. When the router is searching for replacement paths
 - d. When the router is discovering neighbors

13. The factors that can affect the scaling of EIGRP include which of the following?
 - a. The amount of information sent between neighbors
 - b. The number of routers that receive updates
 - c. The distance between neighboring routers
 - d. The number of alternative paths to remote networks

You can find the answers to the “Do I Know This Already?” quiz in Appendix A, “Answers to Chapter ‘Do I Know This Already?’ Quizzes and Q&A Sections.” The suggested choices for your next step are as follows:

- **9 or less correct**—Read the entire chapter. This includes the “Foundation Topics,” “Foundation Summary,” and “Q&A” sections.
- **10 or 11**—Begin with the “Foundation Summary” section, and then go to the “Q&A” section at the end of the chapter. If you have trouble answering the Q&A questions, read the appropriate sections in “Foundation Topics.”
- **12 or 13 overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the “Q&A” section at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

EIGRP Features and Advantages

The goal of EIGRP is to solve the scaling limitations that IGRP faces, while keeping the advantages of distance vector routing protocols: simplicity, economy of memory, and economy of processor resources. EIGRP is scalable in terms of hardware resources and network capacity. EIGRP is also lightning fast.

Cisco identifies four principal components of EIGRP:

- **Protocol-dependent modules**—EIGRP supports several routed protocols independently. The two that are of interest today are IP and IPv6.
- **Reliable Transport Protocol**—EIGRP sends some packets reliably using a reliable transport protocol.
- **Neighbor discovery and recovery**—EIGRP uses hellos to identify its neighbors quickly and to recognize when those neighbors are down.
- **Diffusing Update Algorithm (DUAL)**—DUAL identifies the procedure used to sort the list of available paths and select best paths and feasible fail-over routes.

The following sections describe the key features and advantages of EIGRP, including neighborship and reliable incremental updates, neighbor discovery and recovery, the sophisticated metric used by EIGRP, DUAL, and queries.

Neighborship and Reliable Incremental Updates

EIGRP produces reliable updates by identifying its packets using IP protocol 88. *Reliable*, in a networking context, means that the receiver acknowledges that the transmission was received and understood. EIGRP only repeats itself if an advertisement is lost, so EIGRP is less “chatty” than other protocols.

EIGRP uses the following five types of packets to communicate. These packets are directly encapsulated by IP.

- **Hello**—Identifies neighbors. Hellos are sent as periodic multicasts and are not acknowledged directly.
- **Update**—Advertises routes. Updates are sent as multicasts only when there is a change.

- **Ack**—Acknowledges receipt of an update.
- **Query**—Used to ask about routes for which the previous best path has been lost. If an update indicates that a path is down, multicast queries are used to ask other neighbors if they still have a path. If the querier does not receive a reply from each of its neighbors, it repeats the query as a unicast to each unresponsive neighbor until it either gets a reply or gives up after sixteen tries.
- **Reply**—Used to answer a query. Each neighbor responds to the query with a unicast reply indicating an alternative path or the fact that it does not have a path.

Neighbor Discovery and Recovery

Using reliable updates produces two new problems:

- The router needs to know how many other routers exist, so it knows how many acknowledgements to expect.
- The router needs to know whether a missing advertisement should be interpreted as “no new information” or “neighbor disconnected.”

EIGRP uses the concept of neighborhood to address these problems. EIGRP produces hellos periodically. The first hellos are used to build a list of neighbors; thereafter, hellos indicate that the neighbor is still alive. If hellos are missed over a long period of time—the *hold time*—then the neighbor is removed from the EIGRP table and routing reconverges.

EIGRP starts by discovering its neighbors. Advertisements are multicast, and individual unicast acknowledgements come back. The neighbor table is used to make sure that each neighbor responds. Unresponsive neighbors receive a follow-up unicast copy, repeatedly, until they acknowledge. If a neighbor is still unresponsive after 16 attempts, the neighbor is removed from the neighbor table and EIGRP continues with its next task.

Presumably, the neighbor will at some point be able to communicate. When it is able to do so, it will send a hello and the process of routing with that neighbor will begin again.

Sophisticated Metric

EIGRP uses a sophisticated metric that considers bandwidth, load, reliability, and delay. That metric is

$$\text{metric} = 256 * (K_1 * \text{bandwidth} + \frac{K_2 * \text{bandwidth}}{256 - \text{load}} + K_3 * \text{delay}) * \frac{K_5}{\text{reliability} + K_4}$$

Although this equation looks intimidating, a little work will help you understand the math and the impact the metric has on route selection.

You first need to understand that EIGRP selects path based on the fastest path. To do that it uses K-values to balance bandwidth and delay. The K-values are constants that are used to adjust the relative contribution of the various parameters to the total metric. In other words, if you wanted *delay* to be much more relatively important than *bandwidth*, you might set K_3 to a much larger number.

You next need to understand the variables:

- **Bandwidth**—Bandwidth is defined as 107 kbps divided by the slowest link along the path. Because routing protocols select the lowest metric, inverting the bandwidth (using it as the divisor) makes faster paths have lower costs.
- **Load and reliability**—Load and reliability are 8-bit calculated values based on the performance of the link. Both are multiplied by a zero K-value, so neither is used.
- **Delay**—Delay is a constant value on every interface type, and is stored in terms of microseconds. For example, serial links have a delay of 20,000 microseconds and Ethernet lines have a delay of 1000 microseconds. EIGRP uses the sum of all delays along the path, in tens of microseconds.

By default, $K_1=K_3=1$ and $K_2=K_4=K_5=0$. Those who followed the math will note that when $K_5=0$ the metric is always zero. Because this is not useful, EIGRP simply ignores everything outside the parentheses. Therefore, given the default K-values the equation becomes

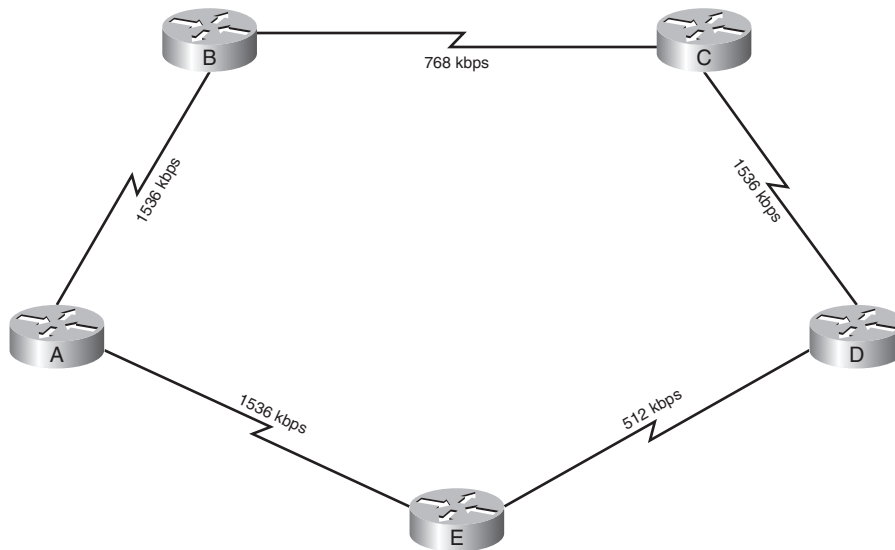
$$\begin{aligned} \text{metric} &= 256 * (1 * \text{bandwidth} + \frac{0 * \text{bandwidth}}{256 - \text{load}} + 1 * \text{delay}) * \frac{0}{\text{reliability} + 0} \\ &= 256 * (\text{bandwidth} + \text{delay}) \end{aligned}$$

Substituting the earlier description of variables, the equation becomes 10,000,000 divided by the chokepoint bandwidth plus the sum of the delays:

$$\text{metric} = 256 * \left(\frac{10^7}{\min(\text{bandwidth})} + \sum \frac{\text{delays}}{10} \right)$$

An example of the metric in context will make its application clear. Figure 3-1 shows a simple network topology, with routers labeled A, B, C, D, and E. Using the metric equation, which path would be used to pass traffic from Router A to Router D?

Figure 3-1 EIGRP Metric Topology



The top path (ABCD) metric would have a chokepoint bandwidth of 768 Kbps and would go across three serial lines:

$$metric = 256 * \left(\frac{10^7}{768} + \frac{20,000 + 20,000 + 20,000}{10} \right) = 256(13,020 + 6000) = 4,869,120$$

The bottom path (AED) metric would have a chokepoint bandwidth of 512 Kbps and would go across two serial lines:

$$metric = 256 * \left(\frac{10^7}{512} + \frac{20,000 + 20,000}{10} \right) = 256(19,531 + 4000) = 6,023,936$$

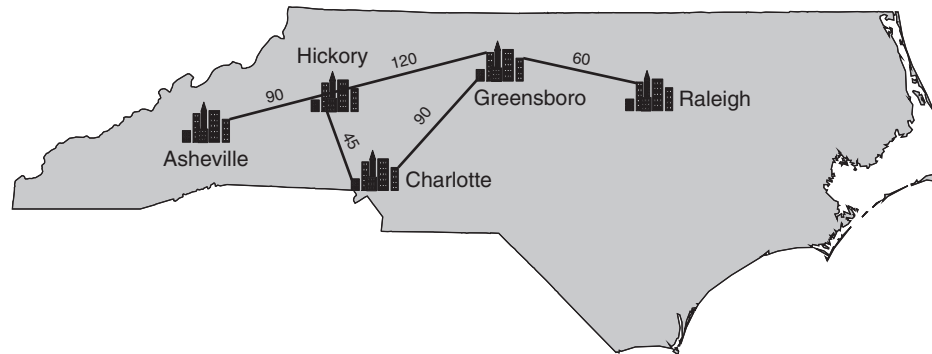
EIGRP chooses the top path based on bandwidth.

NOTE Routers will not become EIGRP neighbors unless they share K-values. There really is not a compelling reason to change the default K-values and Cisco does not recommend it.

DUAL

The Diffusing Update Algorithm (DUAL) is a modification of the way distance-vector routing typically works that allows the router to identify loop-free failover paths. This concept is easier to grasp if you imagine it geographically. Consider the map of North Carolina shown in Figure 3-2. The numbers show approximate travel time by car, in minutes. Pretend that you live in Hickory. From Hickory, you need to determine the best path to Raleigh.

Figure 3-2 North Carolina



Imagine that each of Hickory’s neighbors advertises a path to Raleigh. Each neighbor advertises its cost (travel time) to get to Raleigh and the cost Hickory would use. The cost from the neighbor to the destination is called the *advertised distance*. The cost from Hickory is called the *feasible distance*.

In this example, Greensboro reports that if Hickory routed to Raleigh through Greensboro, the total cost (feasible distance) is 180 minutes, and that the remaining cost once the traffic gets to Greensboro is only 60 minutes. Table 3-2 shows distances reported from Hickory to Raleigh going through each of Hickory’s neighbors.

Table 3-2 Feasible and Advertised Distance

City	Feasible Distance	Advertised Distance
Asheville	360	270
Charlotte	195	150
Greensboro	180	60

Hickory will select the route with the lowest feasible distance, which is the path through Greensboro.

If the Hickory-Greensboro link goes down, Hickory knows it may fail-over to Charlotte without creating a loop. Notice that the distance from Charlotte to Raleigh (150 minutes) is less than the distance from Hickory to Raleigh (180 minutes). Because Charlotte is closer to Raleigh, routing through Charlotte does not involve driving to Charlotte and then driving back to Hickory before going to Raleigh. Charlotte is a guaranteed loop-free path.

This idea that a path through a neighbor is loop free if the neighbor is closer is called the *feasibility requirement* and can be restated as “using a path where the neighbor’s advertised distance is less than our feasible distance will not result in a loop.”

The neighbor with the best path is referred to as the *successor*. Neighbors that meet the feasibility requirement are called *feasible successors*. In emergencies, EIGRP understands that using feasible successors will not cause a routing loop and instantly switches to the backup paths.

Notice that Asheville is not a feasible successor. Asheville's AD (270) is higher than Hickory's FD (180). For all we know, driving to Raleigh through Asheville involves driving from Hickory to Asheville, then turning around and driving back to Hickory before continuing on to Raleigh (in fact, it does). Asheville will still be queried if the best path is lost and no feasible successors are available because potentially there could be a path that way; however, paths that do not meet the feasibility requirement will not be inserted into the routing table without careful consideration.

Now consider how DUAL works in terms of routers and networks. Figure 3-3 shows the EIGRP topology while Table 3-3 shows each router's options for routing to 192.168.5.0/24 (the Ethernet network on the right).

Figure 3-3 EIGRP DUAL

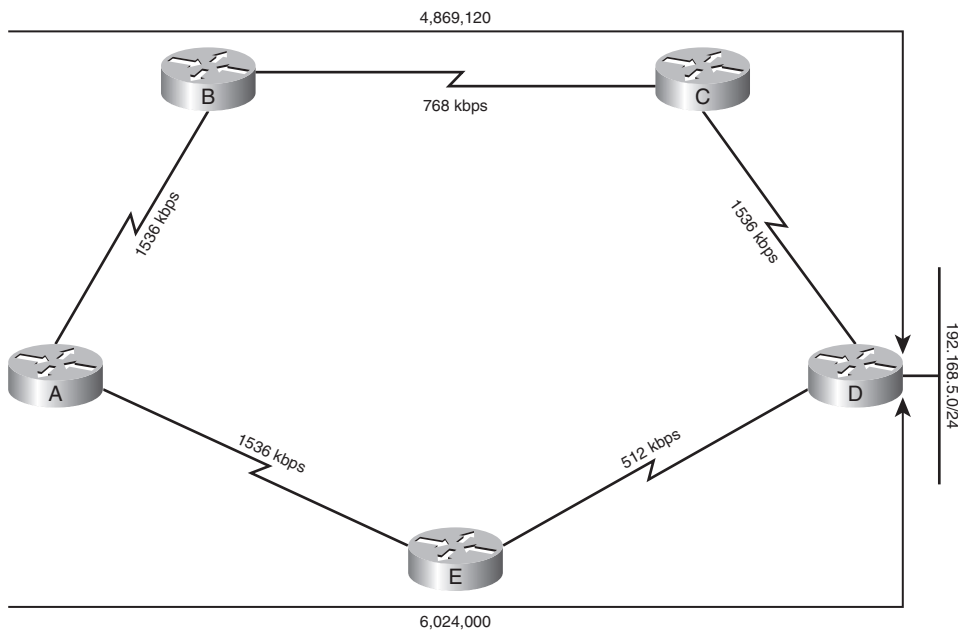


Table 3-3 EIGRP DUAL Results to 192.168.5.0/24

Router	Clockwise	Counter-Clockwise
A	Through B—Successor Min(BW)=768, Delay=60000 FD=4,869,120 AD=4,357,120	Through E Min(BW)=512, Delay=40000 FD=6,023,936 AD=5,511,936
B	Through C—Successor Min(BW)=768, Delay=40000 FD=4,357,120 AD=2,178,560	Through A Min(BW)=512, Delay=60000 FD=6,535,936 AD=6,023,936
C	To D—Successor Min(BW)=1536, Delay=20000 FD=2,178,560 AD=0	Through B Min(BW)=512, Delay=80000 FD=7,047,936 AD=6,535,936
E	To A—Successor Min(BW)=768, Delay=80000 FD=5,381,120 AD=4,869,120	To D—Feasible Successor Min(BW)=512, Delay=20000 FD=5,511,936 AD=0

Notice that Router E has a feasible successor because the advertised distance of the alternative path is less than the Feasible Distance of the best path ($0 < 5,381,120$). On the other hand, A, B, and C do not have feasible successors because their counterclockwise path does not meet the feasibility requirement.

In practical terms, what does this mean for the network? If Router E loses its path through A, it can fail-over to the D path instantly and without consulting neighbors because it knows this will not create a problem. If Router A loses its path through B, it has to query its remaining neighbor (E) and ask whether a path is still available. The query process results in the counterclockwise path being used as a backup but with a very short additional delay.

Queries

Having a feasible successor provides the best convergence. A feasible successor is a backup path, and it can be substituted for a lost path at any point. When a path is lost and no feasible successor exists, the router will send queries to its remaining neighbors. If a neighbor does not know of an alternative path, it will recursively ask its neighbors.

Recursive queries can loop without being resolved, forcing the router to time-out the query. This situation is known as stuck in active (SIA). Fortunately, it is uncommon; understanding its causes can prevent it entirely.

EIGRP uses split-horizon, which says that a router should not advertise a network on the link from which it learned about the network. As shown in Figure 3-3, because Router A learned about 192.168.5.0 from Router B, A does not advertise the network to B.

If the link between B and C goes down, B loses its only path to 192.168.5.0. The query process allows B to actively search remaining neighbors for a replacement route. When B asks A, “Do you have a route to 192.168.5.0/24?,” Router A realizes that its route (through B) is down and recursively queries E. Router E has a route and so replies to A, which passes the news on to B.

Queries continue propagating until an answer is found or until no one is left to ask. When queries are produced, the router changes to an active state and sets a timer (typically three minutes). If the timer expires before an answer comes back then the router is considered stuck in active. SIA typically occurs because queries loop or are not properly limited to an area. The primary way to limit how far queries travel (called *query scoping*) is to summarize.

Queries will not cross summarization because the summary answers the query—the route is either behind the summary or not. For this reason, summarization is extremely important on EIGRP networks. The concepts of summarization are discussed in more detail in Chapter 2, while the EIGRP commands to summarize are found in Chapter 4.

The query process allows EIGRP to avoid periodically transmitting entire routing tables. When new networks are added or advertisements are withdrawn, routers may ask each other for additional information, allowing EIGRP to converge quickly even when there is not a feasible successor.

Additional EIGRP Advantages

The preceding section discussed the key advantages of EIGRP, but some smaller issues relating to network efficiency remain.

Incremental Updates

EIGRP periodically sends hellos to maintain neighborhood, but only sends updates when a change occurs. When a route is added or withdrawn, an incremental update is sent (that includes only those changes). This is an important feature because it prevents EIGRP from monopolizing link access, which was occasionally a problem with older protocols.

Multicast Addressing for Updates

EIGRP uses both multicast and unicast addressing. Some packets are sent reliably using Real-Time Protocol (RTP), a Cisco proprietary protocol that oversees the communication of EIGRP packets. These packets are sent with sequence numbers to make the transmission of data reliable. Hellos and ACKs do not require acknowledgement.

Incremental updates cannot be anticipated; therefore, update, query, and reply packets must be acknowledged by the receiving neighbor.

Updates are sent using a reliable multicast. The address is the reserved class D address, 224.0.0.10. When the neighbor receives a multicast, it acknowledges receipt of the packet with an unreliable unicast.

The use of multicast by EIGRP to send updates is also important because it represents an improvement over other protocols. Older protocols used broadcast, which created issues. Multicast allows hosts to filter out the traffic while preserving the “one to all” aspect of broadcasts.

NOTE A broadcast domain identifies devices that are within the same Layer 2 domain. Although they might not be directly connected to the same physical cable, if they are in a switched environment, from a logical Layer 2 or Layer 3 perspective, they are on the same link. If a broadcast is sent out, all the devices within the broadcast domain will hear the message and will expend resources determining whether it is addressed to them. A Layer 3 device is a broadcast firewall in that a router does not forward broadcasts.

Unequal-Cost Load Sharing

All IP routing protocols on Cisco routers support equal-cost load sharing. EIGRP is unique in its support for unequal-cost load sharing.

Unequal-cost load balancing takes the best FD and multiplies it by variance. Any other path with an FD less than this product is used for load sharing. That is exciting because now a 256 kbps link and a 384 kbps link can work together—but EIGRP actually goes one better than that.

EIGRP does proportional unequal-cost load sharing. EIGRP will pass a relative portion of the traffic to each interface. The 384 kbps link would get 60 percent and the 256 kbps link would handle 40 percent of the traffic. This allows all links to a destination to be used to carry data without saturating the slower links or limiting the faster links.

Understanding EIGRP Tables

EIGRP builds and maintains three tables. A neighbor table is used to make sure all acknowledgements are received. A topology table is used to understand paths through the network. Finally, the best paths from the topology table are fed into the IP routing table. The following sections describe how EIGRP creates and maintains each table.

Creating the Neighbor Table

The neighbor table is maintained by means of Hello packets. Hello packets are multicast announcements that the router is alive. Hello packets place the router into adjacent routers' neighbor tables. Reciprocal hellos build the local neighbor table. Once neighbor tables are built, hellos continue periodically to maintain neighborhood ("I'm still here!").

Each Layer 3 protocol supported by EIGRP (IPv4, IPv6, IPX, and AppleTalk) has its own neighbor table. Information about neighbors, routes, or costs is not shared between protocols.

Contents of the Neighbor Table

The neighbor table includes the following information:

- The Layer 3 address of the neighbor.
- The interface through which the neighbor's Hello was heard.
- The holdtime, or how long the neighbor table waits without hearing a Hello from a neighbor, before declaring the neighbor unavailable and purging the database. Holdtime is three times the value of the Hello timer by default.
- The uptime, or period since the router first heard from the neighbor.
- The sequence number. The neighbor table tracks all the packets sent between the neighbors. It tracks both the last sequence number sent to the neighbor and the last sequence number received from the neighbor.
- Retransmission timeout (RTO), which is the time the router will wait on a connection-oriented protocol without an acknowledgment before retransmitting the packet.
- Smooth Round Trip Time (SRTT), which calculates the RTO. SRTT is the time (in milliseconds) that it takes a packet to be sent to a neighbor and a reply to be received.
- The number of packets in a queue, which is a means by which administrators can monitor congestion on the network.

Becoming a Neighbor

All EIGRP routers periodically announce themselves with the Hello protocol using a multicast address of 224.0.0.10. On hearing Hellos, the receiving routers add an entry in their neighbor table. The continued receipt of these packets maintains the neighbor table. If a Hello from a known neighbor is not heard within the holdtime, the neighbor is treated as no longer operational and removed from the table. The holdtime, by default, is three times the Hello timer. Therefore, if the router misses three Hellos, the neighbor is declared dead. The Hello timer on a LAN is set to 5 seconds; the holdtime, therefore, is 15 seconds. On DS1 (1.5 Mbps) or slower WAN links, the Hello timer is 60 seconds and the holdtime is 180 seconds.

To become a neighbor, the following conditions must be met:

- The router must hear a Hello packet from a neighbor.
- The EIGRP autonomous system number in the Hello must be the same as that of the receiving router.
- K-values used to calculate the metric must be the same.

Example 3-1 demonstrates a neighbor table.

Example 3-1 The `show ip eigrp neighbors` Command

Router#show ip eigrp neighbors							
IP-EIGRP neighbors for process 1							
H	Address	Interface (sec)	Hold (ms)	Uptime Cnt	SRTT	RT0 Q	Seq Num
172.16.0.1	S0/0	14	01:16:26	149	894 0	291	

Creating the Topology Table

After the router knows who its neighbors are, it is able to create a topological database and assign successors and feasible successors for each route. The topology table has a record not only of feasible successors and successors but also of all received routes. The other routes are referred to as *possibilities*. The topology table in EIGRP manages the selection of routes to be added to the routing table.

The topology table includes the following information:

- Whether the route is passive or active.
- Whether an update has been sent to the neighbor.

- Whether a query packet has been sent to the neighbor. If this field is positive, at least one route will be marked as active.
- Whether a query packet has been sent; if so, another field will track whether any replies have been received from the neighbor.
- That a reply packet has been sent in response to a query packet received from a neighbor.
- Prefixes, masks, interface, next-hop, and feasible and advertised distances for remote networks.

The topology table is built from the update packets that are exchanged by the neighbors and by replies to queries sent by the router.

The queries and responses used by EIGRP for DUAL are sent reliably as multicasts using RTP, which is a Cisco proprietary protocol. If a router does not hear an acknowledgment within the allotted time, it retransmits the packet as a unicast. If there is no response after 16 attempts, the router marks the neighbor as dead. Each time the router sends a packet, RTP increments the sequence number by one. The router must hear an acknowledgment from every router before it can send the next packet. The capability to send unicast retransmissions decreases the time that it takes to build the tables.


When the router has an understanding of the topology, it runs DUAL to determine the best path to the remote network. The result is entered into the routing table.

Maintaining the Topology Table

The topology table may be recalculated because a new network is added to the network, successors change, or because a network is lost. Figure 3-4 illustrates the traffic flow seen when a router loses a connection.

Just as the neighbor table tracks the receipt of the EIGRP packets, the topology table records the packets that have been sent by the router to the neighbors. It also identifies the status of the networks in the table. Like a Sunday afternoon, passive is good and active is bad. A healthy network is marked passive; a lost route is active because the router is actively attempting to find an alternative path to the remote network.

Because the routing table is built from the topology table, the topology table must have all the information required by the routing table. This includes the next hop or the address of the neighbor that sent the update, and the metric (which is taken from the feasible distance).

Figure 3-4 *Traffic Flow Used to Maintain the Topology Table*

Adding a Network to the Topology Table

Imagine an access layer router (Router A) that connects to a new network via an Ethernet interface. The administrator has configured another interface to service a department that has moved into the building. At the start of this process, the old interface has converged routing. The following list describes how the new network is propagated to all the routers in the EIGRP autonomous system:

1. As soon as Router A becomes aware of the new network, it starts to send Hello packets out of the new interface. No one answers—no other routers are on the segment.

There are no new entries in the neighbor table because no neighbors have responded to the Hello protocol. There is a new entry in the topology table, however, because it is attached to a new network.

2. EIGRP, sensing a change, is obliged to send an update to all its neighbors on the old interface, informing them of the new network. These updates are tracked in the topology table and the neighbor table because the updates are connection-oriented and the acknowledgments from the neighbors must be received within a set timeframe.

3. Router A has completed its work. However, its neighbors on the old network still have work to do. On hearing the update from Router A, they will update the sequence number in their neighbor table and add the new network to the topology table. They calculate the FD and the successor to place in the routing table.

The next section describes the process for removing a router or path from the topology table.

Removing a Path or Router from the Topology Table

The process of removing a path or router from the topology table is far more complex and gets to the crux of EIGRP.

1. If a network connected to Router A is disconnected, Router A updates its topology and routing table and sends an update to its neighbors.
2. When a neighbor receives the update, it updates the neighbor table and the topology table.
3. The neighbor searches for an alternative route to the remote network. It examines the topology table for alternatives. Because there is only one path to the remote network, no alternatives are found.
4. The neighbor then sends out a query to its neighbors requesting that they look in their tables for paths to the remote network. The route is marked active in the topology table at this time.
5. The query number is tracked, and when all the replies are in, the neighbor and topology tables are updated.
6. DUAL, which starts to compute as soon as a network change is registered, runs to determine the best path, which is placed in the routing table.
7. Before they respond, they query their own neighbors; in this way, the search for an alternative path extends or diffuses throughout the organization.
8. If no alternative route is available, the neighbors reply to the query stating that they have no path.
9. When no router can supply a path to the network, all the routers remove the network from their routing and topology tables.

The next section describes what happens if a neighbor does have an alternative route.

Finding an Alternative Path to a Remote Network

When the path to a network is lost, EIGRP goes to a lot of trouble to find an alternative path. The method that EIGRP uses to find alternative paths is very reliable and very fast.

The following list describes the process after a neighbor goes offline:

1. The router marks the routes that were reached by sending the traffic to that neighbor.

2. The router looks in the topology table, which has every advertisement received, to determine whether there is an alternative route. It is looking for a FS.
3. If a successor is found, the router adds the FS route to its routing table. If the router did not have a FS, it would have placed the route into an active state while it actively queried other routers for an alternative path.
4. After interrogating the topology table, if a feasible route is found, the neighbor replies with the alternative path. This alternative path is then added to the topology table.
5. If no answer is heard, the messages are propagated.

When the router sends a query packet, it is recorded in the topology table. This is to ensure a timely reply. If the router does not hear a reply, the neighbor is removed from the neighbor table and all networks held in the topology table from that neighbor are removed from the topology table. Occasionally, problems can occur because of slow links and burdened routers in a large network. In particular, a router might not receive a reply from all the queries that it sent out. This leads to the route being declared SIA; the neighbor that failed to reply is removed from the neighbor table.

Creating the Routing Table

The routing table is built from the topology table using DUAL. The topology table holds all routing information known to the router and from this information successors and feasible successors are selected. Successor paths are then transferred to the routing table and used as the basis for routing decisions.

EIGRP Path Selection

The DUAL algorithm uses the metric to select the best path or paths to a destination. Up to sixteen paths can be held for one destination. There are three different types of paths. These three path types are described in Table 3-4.

Table 3-4 *EIGRP Route Types*

Route Type	Description
Internal	Paths placed directly into EIGRP
Summary	Internal paths that have been summarized
External	Routes redistributed into EIGRP


The metric is discussed in greater detail earlier in the chapter in the section “Sophisticated Metric” and is calculated from bandwidth and delay.

Updating the Routing Table in Passive Mode with DUAL

When a path is lost, DUAL first looks in the topology table for a feasible successor. If one is found, the router stays in passive mode (passive, in this sense, means that the router is not actively querying for an alternative path).

Figure 3-5 provides an example network.

Figure 3-5 *The Use of Feasible and Advertised Distance—Passive Mode*



The following list explains Figure 3-5 with the metrics and actions that EIGRP takes in determining the path:

- The feasible distance from Router A to Router G is 10 (A to D to G).
- The advertised distance from Router A to Router G is 5 (advertised from Neighbor D).
- Because $10 > 5$, feasible distance $>$ advertised distance. This means that the feasible distance meets the feasibility condition, allowing it to become a feasible distance. If you follow the diagram, it is very straightforward and less algebraic.
- If the link between Router D and Router G were down, Router A would look in its topology table.
- The alternative routes through Router A to D to H to E to G have an advertised distance of 19 ($7 + 5 + 7$).

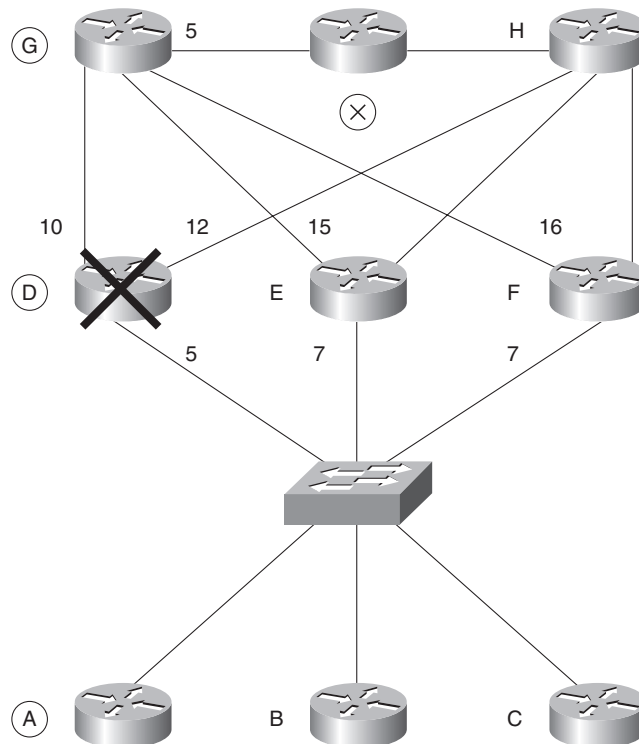
- Because 19 is greater than the original feasible distance of 10, it does not qualify as a feasible successor.
- The path through Router D to H to F to G has an advertised distance of 20 and cannot be a feasible successor.
- The path through Router A to E to G has an advertised distance of 7, however, which is less than the original 10. Therefore, this is a feasible successor and can be replaced as a route without Router A changing from passive to active mode.
- The original topology table would show that the primary route (successor) is Router D, while the backup route (feasible successor) is Router E. After the link between D and G dies, the routing table would be updated from the topology table while the route remains passive.

The following section illustrates what happens when the topology table is interrogated and no feasible route is found.

Updating the Routing Table in Active Mode with DUAL

Figure 3-6 shows feasible and advertised distance.

Figure 3-6 Feasible and Advertised Distance—Active Mode



When no alternative route is found in the routing table, the following actions are taken (using the network in Figure 3-6 as an example). The topology table of Router A starts with a path (successor) of A to D to G to X. The feasible distance is 20, and the advertised distance from Router D is 15. When Router D dies, Router A must find an alternative path to X.

- The router rejects neighbors B, C, E, and F as feasible successors. These neighbors have advertised distances of 27, 27, 20, and 21, respectively. Because all the neighbors have an advertised distance that is the same or greater than the successor feasible distance, they do not meet the feasibility requirement.
- Router A goes into active mode and sends queries to the neighbors.
- Both Routers E and F reply with a feasible successor because both have an advertised distance of 5 from Router G. Remember that equation $\text{feasible distance} > \text{advertised distance}$; the Routers E and F have a feasible distance of 21, and $21 > 5$.
- The network returns to passive mode. The feasible distance is acceptable, the topology and routing tables will be updated, and there is no need for further convergence.
- From this information received from Routers E and F, the router selects the path through E as the best route because it has the lower feasible distance.
- The result is placed in the routing table as the valid neighboring router. EIGRP refers to this neighboring router as a successor.
- Router F will be stored as an FS in the topology table.

The details on how EIGRP computes successors are complex, but the concept is simple, as described in the next section.

EIGRP Network Design

EIGRP is designed to work in very large networks. However, EIGRP is design-sensitive. Scaling a network is a major concern in organizations. New demands are constantly driving the networks to use applications that require more bandwidth and less delay; at the same time networks are becoming larger and more complex. Mergers and acquisitions, for instance, do nothing for good design.

The factors that can affect the scaling of EIGRP include

- The amount of information sent between neighbors
- The number of routers that receive updates
- The distance between neighboring routers
- The number of alternative paths to remote networks

Poorly scaled EIGRP networks can result in

- A stuck-in-active route
- Network congestion
- Lost routing information
- Flapping routes
- Retransmission
- Low router memory
- Overutilized router CPU

Other factors, such as poor design, cause some of these symptoms because resources are overwhelmed by the tasks assigned. Often, a route flagged as SIA characterizes many of these symptoms, because the router waits for a reply from a neighbor across a network that cannot handle the demands made upon it. Careful design and placement of network devices can remedy many of the problems seen in a network.

EIGRP Design Issues

The major concerns in scaling an organizational network are controlling advertisements and limiting query range. These are particularly important over slow WAN links. By sending less information about the network, the capacity available for the data between clients and servers increases. Although sending less routing information relieves the network and speeds convergence, it provides less information for finding alternate routes. A balance between summarization and full information must be struck, but generally this balance will tilt toward more summarization and not less.

EIGRP automatically summarizes at classful network boundaries because summarization is generally helpful and the EIGRP process is built to recognize opportunities such as this to optimize the network. However, most administrators find that automatic summarization sometimes does not match what their network needs and choose to disable it. Instead, they choose to manually configure summarization at the interface level.

Certain topologies pose problems for the EIGRP network. This is true in particular for the hub-and-spoke design often seen implemented between remote sites and regional offices. The popular dual-hub configuration provides redundancy and allows the potential for routers to reflect queries back to one another. Summarization and filters make this network design work well while also allowing queries to be managed effectively.

Guideline to Address Scaling Issues

The design of the network is very important to the ability to scale any network. The following actions are critical to a well-designed network:

- Assign addresses and organize links so that natural points for summarization exist. A hierarchical network design meets this criteria.
- Provide sufficient hardware resources (memory and CPU) on network devices.
- Use sufficient bandwidth on WAN links.
- Use filters to limit advertisements.
- Monitor the network.

Foundation Summary

The “Foundation Summary” section of each chapter lists the most important facts from the chapter. Although this section does not list every fact from the chapter that will be on your exam, a well-prepared candidate should, at a minimum, know all the details in each “Foundation Summary” before going to take the exam.

The key features of EIGRP include

- Neighborhood and reliable incremental updates
- Neighbor discovery and recovery
- Sophisticated metric
- DUAL
- Queries

EIGRP uses five types of packets to communicate:

- Hello
- Ack
- Update
- Query
- Reply

EIGRP uses a sophisticated metric that considers bandwidth, load, reliability, and delay. When default constant values are taken into consideration, the equation simplifies to

$$metric = 256 * \left(\frac{10^7}{\min(\text{bandwidth})} + \sum \frac{\text{delays}}{10} \right)$$

Additional EIGRP advantages include

- Incremental Updates
- Multicast Addressing for Updates
- Easy Configuration
- Unequal-Cost Load Sharing

EIGRP builds and maintains three tables:

- Neighbor table
- Topology table
- Routing table

To become a neighbor, the following conditions must be met:

- The router must hear a Hello packet from a neighbor.
- The EIGRP autonomous system must be the same.
- K-values must be the same.

Q&A

The questions and scenarios in this book are designed to be challenging and to make sure that you know the answer. Rather than allowing you to derive the answers from clues hidden inside the questions themselves, the questions challenge your understanding and recall of the subject.

You can find the answers to these questions in Appendix A. For more practice with exam-like question formats, use the exam engine on the CD-ROM.

1. If a router does not have a feasible successor, what action will it take?
2. Which timers are tracked in the neighbor table?
3. What is the difference between an update and a query?
4. When does EIGRP recalculate the topology table?
5. State two rules for designing a scalable EIGRP network.
6. Which EIGRP packets are sent reliably?
7. How long is the holdtime by default?
8. What is an EIGRP topology table, and what does it contain?
9. What is the advertised distance in EIGRP, and how is it distinguished from the feasible distance?
10. What EIGRP algorithm is run to create entries for the routing table?
11. When does EIGRP place a network in active mode?
12. By default, EIGRP summarizes at which boundary?
13. What is Stuck in Active?
14. State two factors that influence EIGRP scalability.
15. What are reply packets in EIGRP?
16. What conditions must be met for a router to become a neighbor?