



This chapter covers the following topics:

- **Configuring EIGRP**—This section describes the basic configuration commands necessary to support routing.
- **Verifying the EIGRP operation**—This section describes **show** and **debug** commands that will help evaluate the status of EIGRP.
- **Troubleshooting EIGRP**—This section describes the process of diagnosing and repairing EIGRP problems.

Scalable EIGRP

This chapter details steps to configure EIGRP for IP routing. This chapter assumes general knowledge of routing protocols and the terminology, concepts, and operation of EIGRP (from Chapter 3, “EIGRP Principles”). This chapter introduces EIGRP configuration commands by explaining the required commands and then discussing the optional configuration commands that you can use.

“Do I Know This Already?” Quiz

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

The 12-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 4-1 outlines the major topics discussed in this chapter and the “Do I Know This Already?” quiz questions that correspond to those topics.

Table 4-1 “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundation Topics Section	Questions Covered in This Section	Score
Configuring EIGRP	1–7	
Verifying the EIGRP Operation	8–10	
Troubleshooting EIGRP	11–12	
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. What command starts the EIGRP routing process?
 - a. **router eigrp**
 - b. **router eigrp** *autonomous-system-number*
 - c. **ip routing eigrp**
 - d. **ip eigrp routing**

2. What happens if no **network** command is configured?
 - a. The EIGRP process is not activated on any interface.
 - b. The EIGRP defaults to all interfaces.
 - c. The EIGRP process can receive, but cannot send, updates.
 - d. The router prompts you for the networks.

3. What happens immediately after the **network** command is configured?
 - a. Updates are sent.
 - b. The routing table is created.
 - c. Hellos are sent on appropriate interfaces.
 - d. Networks are advertised.

4. Where in the EIGRP network is it possible to summarize IP addresses?
 - a. At a classful network boundary
 - b. At the ASBR
 - c. At the ABR
 - d. At an interface

5. Where is the command **ip summary-address eigrp** *autonomous-system-number address mask* configured?
 - a. At the routing process
 - b. At the interface
 - c. After the network command
 - d. At the EXEC command prompt

6. What command is used to change the hello timer?
 - a. **egrp hello-interval** *seconds*
 - b. **hello interval** *seconds*
 - c. **ip hello-interval eigrp** *autonomous-system-number seconds*
 - d. **ip eigrp hello timer** *seconds*

7. Which of the following commands will show the hold time?
 - a. **show ip eigrp topology**
 - b. **show ip eigrp traffic**
 - c. **show ip eigrp holdtime**
 - d. **show ip eigrp neighbors**

8. Which command is used to display the current status of EIGRP neighbors?
 - a. **show ip eigrp traffic**
 - b. **show ip eigrp neighbors**
 - c. **show ip eigrp adjacencies**
 - d. **show ip eigrp database**

9. Which command is used to see the types of packets sent and received, as well as the statistics on routing decisions?
 - a. **show ip eigrp traffic**
 - b. **show eigrp events**
 - c. **debug ip eigrp packets**
 - d. **debug ip eigrp traffic**

10. What types of information are shown by the **show ip eigrp traffic** command?
 - a. Updates
 - b. Hellos
 - c. Queries and replies
 - d. ACKs

11. What is displayed in the command **debug ip eigrp summary**?
 - a. A summary of EIGRP activity
 - b. A summary of the contents of the neighbor database
 - c. The process taken when a change is made in a summary route
 - d. A summary of the topology database

12. Which command could be used to show that EIGRP is able to contact neighbors?
 - a. **debug ip eigrp traffic**
 - b. **debug ip eigrp packets**
 - c. **debug ip eigrp events**
 - d. **debug ip eigrp**

The answers to the “Do I Know This Already?” quiz are found 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:

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

Foundation Topics

Configuring EIGRP

Routers are able to route IP out of the box, but you need to configure static routes or a routing protocol so they can reach networks that are not directly connected.

EIGRP is classless, meaning that it advertises both a prefix and a mask with each route. Classless routing protocols are also able to arbitrarily summarize. EIGRP summarizes automatically at classful network boundaries, but you will also want to summarize within your network. EIGRP can summarize at any interface on any router, and summarization is extremely important in building a scalable EIGRP network.

The commands for EIGRP are similar to IP routing protocols.

This section covers the following:

- Required commands for configuring EIGRP
- Optional commands for EIGRP
- Optional EIGRP commands specific to WANs

Required Commands for Configuring EIGRP

Like other interior routing protocols, the template for an EIGRP configuration involves turning on the protocol and identifying participating interfaces. EIGRP also uses an autonomous system (AS) to group routers. To run EIGRP, you must enable the following three items:

- **The EIGRP process**—The routing protocol needs to be started on the router.
- **The EIGRP autonomous system number**—Routers in the same administrative domain should be identified and grouped by a common AS number. This 16-bit number is arbitrary. Organizations that have a BGP AS will sometimes use that number; others just make up a number or use “AS 1.” The significance of the AS is that a router will not become a neighbor with a router in a foreign AS.
- **Participating router interfaces**—You can activate EIGRP on some or all interfaces. For instance, a router might run EIGRP on internal interfaces and BGP on external interfaces.

To enable EIGRP as an active routing protocol, use the **router** command. Remember that the AS must be consistent between routers to exchange updates.

```
Router(config)#router eigrp autonomous-system-number
```

At this point, EIGRP is enabled but not active on an interface. EIGRP will not produce hellos or advertise networks until it is activated on particular links.

NOTE Most versions of the IOS software do not offer an error message when the configuration is incomplete, which can make troubleshooting more difficult. Refer to the section titled “Verifying the EIGRP Operation” later in this chapter for more information.

Interfaces are then attached to the EIGRP process with the **network** command. The **network** command is a pattern-matching tool: interfaces with IP addresses that match the pattern are active EIGRP speakers and those subnets are advertised through EIGRP.

There are two ways to utilize the network command. The first way is by classful network number.

```
Router(config-router)#network network-number
```

In this case, the **network** command in EIGRP is similar to the **network** command in RIP. A router with three interfaces—10.0.0.1, 10.1.1.1, and 192.168.0.1—configured with the **network 10.0.0.0** command would speak EIGRP on two interfaces, but not on the 192.168.0.1 port.

In some cases, you might want to leave an interface that is within the classful network out of EIGRP. The command to do this is **passive-interface**. The **passive-interface** command prevents EIGRP from speaking on an interface; it does not send hellos or advertisements. Neighbors are not found on the passive interface and routes are not exchanged; however, the prefix of the passive interface is advertised to EIGRP neighbors on other interfaces.

The **passive-interface** command can be used on interfaces with no neighbors, or on interfaces that run another routing protocol (such as BGP).

For example, a router with two interfaces addressed 172.17.0.1 and 172.17.55.1 is configured in Example 4-1. EIGRP should run only on fastethernet0/0. Configuring the classful network (network 172.17.0.0) enables EIGRP on both interfaces. The **passive-interface** command is then used to suppress fastethernet0/1. Because EIGRP is active but quiet on fastethernet 0/1, the attached subnet (172.17.55.0/24) is advertised to neighbors through fastethernet0/0.

Example 4-1 The EIGRP passive-interface Command

```
Router(config)#interface fastethernet0/0  
Router(config-if)#ip address 172.17.0.1 255.255.255.0  
Router(config-if)#interface fastethernet0/1
```

Example 4-1 *The EIGRP passive-interface Command (Continued)*

```
Router(config-if)#ip address 172.17.55.1 255.255.255.0
Router(config-if)#router eigrp 100
Router(config-router)#network 172.17.0.0
Router(config-router)#passive-interface fastethernet0/1
```

NOTE As shown in Example 4-1, you do not have to exit back to global-configuration mode to move between sub-configuration modes.

The **network** command for EIGRP can also be stated with a mask option. This allows you to identify specific interfaces for EIGRP. Routes that match the network and mask will run EIGRP and those subnets will be advertised.

This alternative syntax (first supported in IOS 12.0(4)T) is

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

To repeat the previous example, a router with interfaces numbered 172.17.0.1 and 172.17.55.1 needs to run EIGRP only on fastethernet0/0. Using the mask option allows the administrator to identify a specific interface, as shown in Example 4-2.

Example 4-2 *EIGRP network mask Command*

```
Router(config)#interface f0/0
Router(config-if)#ip address 172.17.0.1 255.255.255.0
Router(config-if)#interface f0/1
Router(config-if)#ip address 172.17.55.1 255.255.255.0
Router(config)#router eigrp 100
Router(config-router)#network 172.17.0.1 0.0.0.0
```

There is an important difference between the **passive-interface** and **network mask** techniques. The **passive-interface** command forces an interface to be quiet, but advertises the connected network out other interfaces. Interfaces left out of **network mask** do not speak EIGRP and subnets attached to those interfaces are not advertised.

Once interfaces have been selected for EIGRP, the next order of business is to meet and greet neighbors and exchange routes.

At this point, EIGRP is completely configured. It is possible to “tune” it a bit, and the following sections describe how to do so with optional commands. As the discussion turns to these optional commands, do not lose sight of one of the most important features of EIGRP: EIGRP makes smart decisions (such as bandwidth utilization and summarization) automatically and is thus easy to

configure and troubleshoot. Many of the optional tuning commands are unnecessary because EIGRP defaults are effective in most cases.

Optional Commands for Configuring EIGRP

Optional commands described in this section are summarized in Table 4-2.

Table 4-2 *Optional Commands for Configuring EIGRP*

Command	Description
no auto-summary	Turns off automatic classful summarization.
ip summary-address	Manually configures summarization.
eigrp stub	Configures a stub router.
variance	Configures unequal proportional load balancing.
ip hello-interval eigrp <i>autonomous-system-number seconds</i>	Changes the hello frequency.
ip hold-time eigrp <i>autonomous-system-number seconds</i>	Changes the length of time before a neighbor is considered dead.
bandwidth	Changes the bandwidth setting on an interface, which is used to calculate the EIGRP metric and the maximum percentage of bandwidth to which EIGRP traffic will limit itself.
ip bandwidth-percent eigrp <i>as-number</i>	Changes the maximum amount of bandwidth EIGRP traffic will use. The default is 50%.

Summarization with EIGRP

It cannot be said enough: Scalability comes from summarization. Like other routing protocols, summarization is important to EIGRP because it limits advertisements, minimizes the size and optimizes the speed of the routing table, limits route recalculation, and holds down the amount of memory and processor resources consumed. Smaller routing tables make for faster networks.

Summarization in EIGRP also solves a problem unique to EIGRP: *query scoping*. EIGRP queries do not stop until they come to a dead-end; because of this, they can loop. Route convergence cannot proceed until all replies have been received. Summarization limits queries by stating: “only these routes are found past this point.” Summarizing at logical points in the network is important to prevent stuck in active (SIA).

In IOS versions before 12.2(8), summarization automatically happens at classful network boundaries. Automatic summarization is an example of EIGRP trying to be helpful; however, if your network uses more than one classful network—for example, some enterprises use 10.0.0.0/8 and 192.168.x.0/24 networks concurrently—automatic summarization may create problems in some situations and should be disabled.

Summarization may be manually applied at any point in the network. You can configure manual summarization on any router interface. Consider summarization for both upstream and downstream neighbors. Upstream neighbors should receive a consolidated route, and downstream neighbors can receive a default route.

There are two commands for summarization with EIGRP: **no auto-summary** and **ip summary-address eigrp *autonomous-system-number address mask***. The **no auto-summary** command disables automatic summarization. With **no auto-summary** configured, all known subnets are advertised out of each interface, limited only by split-horizon.

The **ip summary-address eigrp** command is used to announce a specific range on a particular interface. This allows bidirectional summarization and allows summarization to be more widely utilized.

Manual summarization is configured at the interface level, as shown here:

```
Router(config)#interface S0
Router(config-if)#ip summary-address eigrp autonomous-system-number address mask
```

For example, imagine a router that connects to the corporate network and to a division with IP subnets such as 172.20.0.0/24, 172.20.1.0/24 . . . 172.20.31.0/24 on serial 1/1. It also has a default route on serial 1/0. Configuration of manual summarization is done at the interface level. The corporation uses AS 100. This is shown in Example 4-3.

Example 4-3 EIGRP Summarization

```
Router(config)#interface S1/0
Router(config-if)#ip summary-address eigrp 100 172.20.0.0 255.255.224.0
Router(config-if)#interface S1/1
Router(config-if)#ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

Thus the division is summarized to the larger network and receives back a simple default route.

CAUTION Summarizing the default route can be a useful technique. However, if used in cases where there are multiple default paths, this will reduce routing options and prevent failover to alternate paths when the primary path fails.

Stub Routers

Cisco IOS software release 12.0 made it possible for you to configure a remote router as an EIGRP stub router. Stub routers in EIGRP networks use EIGRP to send limited information between the stub and the core routers. A stub router is typically used on routers to minimize memory and processor utilization; as such it is a good candidate for slow routers with minimal memory.

The stub router has only one neighbor, a distribution layer router. The remote router only needs a default route pointing to the distribution router—everywhere else can be reached via that router.

Configuring a router as a stub also helps the rest of the network. Queries are responded to much quicker and convergence happens much faster. Sometimes a query can cause delays that result in the path being SIA. If the stub configuration is applied, the router responds to queries as inaccessible, thus limiting the scope of the query range and preventing SIA from occurring.

The following command structure shows the syntax of the **eigrp stub** command:

```
Router(config-router)# eigrp stub [receive-only | connected | redistributed | static | summary]
```

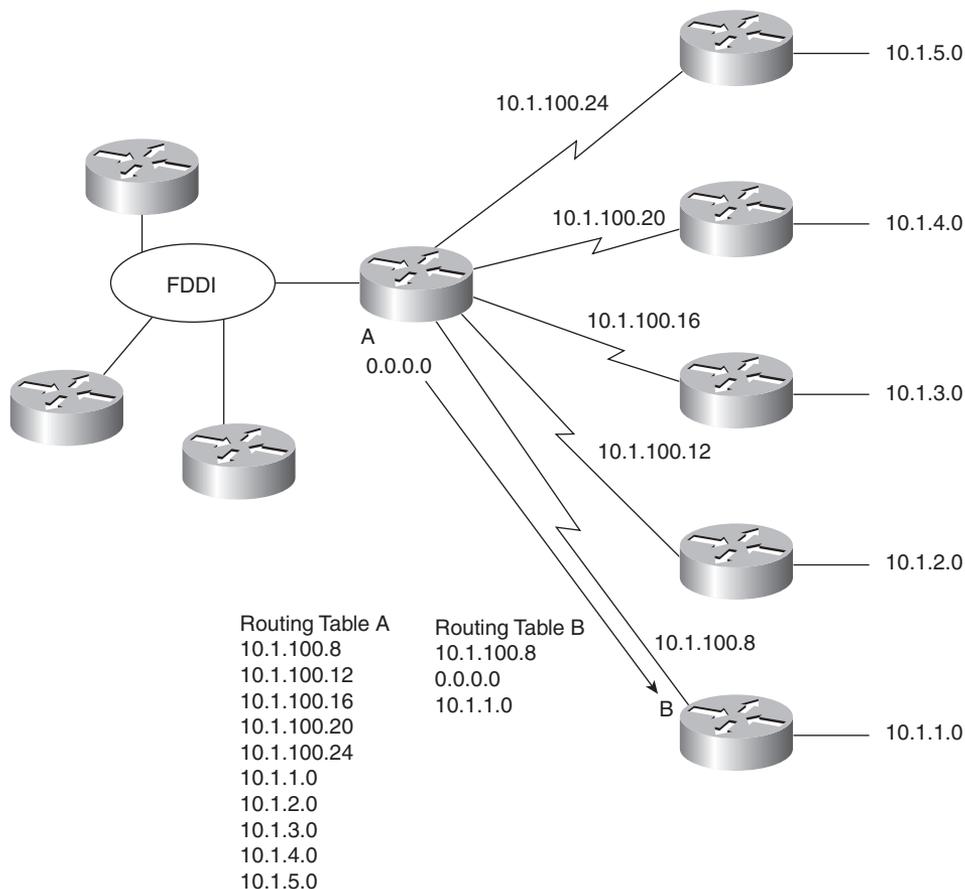
Table 4-3 explains the options further. The third column shows that some optional commands may be combined on a single command line, while others require a separate **eigrp stub** command.

Table 4-3 *The eigrp stub Option Syntax and Description*

Parameter	On by Default	May Be Combined	Description
receive-only	No	No	(optional) Prevents the router from advertising routes.
connected	Yes	Yes	(optional) Permits advertisement of connected routes.
static	No	Yes	(optional) Permits redistribution of static routes.
summary	Yes	Yes	Advertises summary routes.

Figure 4-1 shows a group of routers connected over WAN links. These routers are stub routers because they have no other networks connected to them.

Example 4-4 is the configuration for Router B in Figure 4-1.

Figure 4-1 *Stub Routers*Example 4-4 *The eigrp stub Router Command*

```

RouterB(config)#router eigrp 100
RouterB(config-router)#network 10.0.0.0 255.0.0.0
RouterB(config-router)#eigrp stub

```

This is obviously a simplified example, but it does raise an interesting point. This network could be configured with static routes. By using **eigrp stub**, the on-site technician can simply copy an example configuration into the router and change the IP addresses without having to understand EIGRP. Some organizations use EIGRP in hub-and-spoke topologies because it standardizes the configuration and fits situations where on-site technicians are not as familiar with IP and EIGRP.

Load Balancing in EIGRP

EIGRP automatically load balances across equal-cost links. The exact mechanism depends on the internal switching process of the router. You can also configure EIGRP to load balance proportionally across unequal-cost paths using the **variance** command.

When **variance** is used, the metric of the best path is multiplied by the variance. Alternate paths with a feasible distance lower than that product are used for load balancing. The command syntax is simply the keyword **variance** and the whole number multiplier:

```
Router(config-router)#variance multiplier
```

The *multiplier* argument is a whole number from 1 to 128. The default is 1, which means equal-cost load balancing.

CAUTION When using variance, remember the following important points:

- A variance of two or three is almost always appropriate—use higher values with caution. At a certain point, EIGRP will start activating old 28.8 modems to load-balance with DS1s.
- Variance should be used cautiously with delay-sensitive traffic. A DS1 link takes 8 ms to transmit a 1500 B packet, while a 256-kbps link takes 47 ms. Voice over IP (VoIP) traffic, if shared between the two links, would perceive the difference as 39 ms of jitter.

Cisco Voice over IP (CVoice), Second Edition, by Kevin Wallace, is a great resource for more information about delay and its effect on voice (Cisco Press, 2006).

If the variance number is higher than the default value of 1, the EIGRP process multiplies the best (lowest) metric by the variance. All paths to the same destination that have metrics less than this product are now included in load balancing. The amount of traffic sent over each link is proportional to the metric for the path.

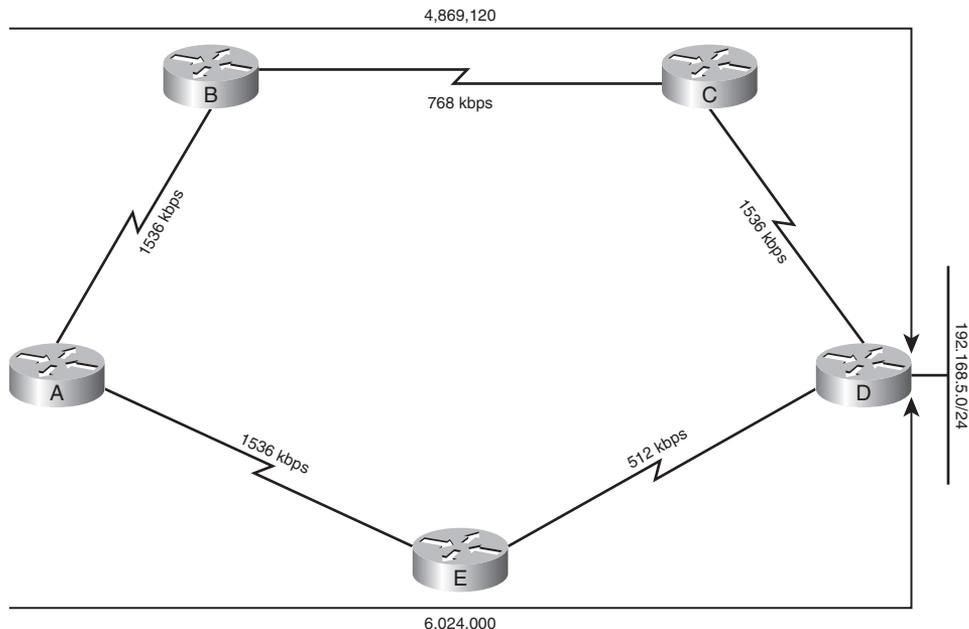
Figure 4-2 shows EIGRP metrics to 192.168.5.0/24. To demonstrate the **variance** command, consider Router B.

Example 4-5 shows the configuration of the **variance** command.

Example 4-5 The **variance** Command

```
RouterB(config)#router eigrp 100
RouterB(config-router)#variance 2
```

Figure 4-2 EIGRP Distance and the **variance** Command



Because the clockwise metric is about five million, a variance of two will load-balance with paths less than ten million. Because the counterclockwise metric is about six million, six packets will go clockwise for every five that go counterclockwise.

NOTE Only those paths that are in the topology table as feasible successors are eligible to be included in the **variance** command.

Tuning the EIGRP Process

You can tune a network in many ways. Load balancing and summarizing routes are probably the two most important ways, but other techniques include reducing the frequency of the hello and hold timers.

There is a trade-off between reducing the hello traffic and maintaining the stability of the network. Fewer hellos, for example, result in the network taking longer to notice a failure, which in turn delays convergence. When the network does not have an accurate understanding of the available routes, the router cannot forward packets correctly.

EIGRP only sends updates when a new route is advertised or an existing route is withdrawn (changes state to down). Link failure causes an interface to change state without delay. But when

a failed neighbor is not directly connected (on the other side of an Ethernet switch, for example), the only way to notice the change is to notice that hellos are no longer received. Neighborship is important because it alerts the router to topology changes and because the router is responsible to the rest of the network to publicize the lost routes.

Change the timers only after careful consideration. In the majority of cases, the default values are reasonable. Instead of improving performance, changing timers can create issues if not done deliberately. For instance, timers are changed per interface and changing the timers on only one side of a link can create a problem where a neighborhood forms and dissolves periodically.

Timer values default based on the speed of the interface. Because the timers are assumed based on speed, they will usually be the same.

Timers are not communicated between neighbors and are not a requirement for neighborship. If Router A has a hello interval of 5 seconds and a hold timer of 15 seconds and Router B has a hello interval of 30 seconds, then the two routers will be neighbors for 15 seconds and then down for 15 seconds. Both the hello timer and the hold timer are discussed in the next sections.

The Hello Timer

Tuning the hello timer directly affects the ability of the network to notice a change in the state of a neighbor. Only after a router's interface is recognized as being down, or the router has failed to hear from a neighbor after a certain amount of time, does the router declare the neighbor dead and take action to update the routing table and neighbors.

For these reasons, the **ip hello-interval eigrp** command is typically used to decrease the time between hellos to ensure that the network converges quickly.

The defaults for the hello timer are as follows:

- High-bandwidth links (every 5 seconds):
 - Broadcast media, such as Ethernet, Token Ring, and FDDI
 - Point-to-point serial links, such as PPP or HDLC leased circuits, Frame Relay point-to-point subinterfaces, and ATM
 - Point-to-point subinterfaces
 - High bandwidth (greater than T1) multipoint circuits, such as ISDN PRI and Frame Relay

- Low-bandwidth links (every 60 seconds):
 - Multipoint circuits T1 bandwidth or slower, such as Frame Relay multipoint interfaces, ATM multipoint interfaces, and ATM
 - Switched virtual circuits and ISDN BRIs

The command to change how often the hellos are sent to neighbors is

```
Router(config-if)#ip hello-interval eigrp autonomous-system-number seconds
```

Notice that the command is applied to an interface and does not affect the default interval on other interfaces. The autonomous system number ties the command to a particular EIGRP process. The last variable then changes the default hello interval. Example 4-6 shows the hello interval on serial0/0 changing to 10 seconds.

Example 4-6 Hello Timer

```
Router(config)#interface Serial 0/0
Router(config-if)#ip hello-interval eigrp 100 10
```

This command is very useful with nonbroadcast multiaccess (NBMA) WAN links. EIGRP uses a 60-second hello timer for Frame Relay by default.

The Hold Timer

The *hold time* is how long the router waits for a hello before pronouncing a neighbor unavailable. By default, the hold time is three times the hello time but changing the hello interval does not automatically change the hold time. The hold timer for an interface must be changed manually using the **ip hold-time eigrp** command, as follows:

```
Router(config-if)#ip hold-time eigrp autonomous-system-number seconds
```

The command is specific to an interface and EIGRP AS. Example 4-7 shows the hold interval on serial0/0 changing to 30 seconds.

Example 4-7 Hold Timer

```
Router(config)#interface serial0/0
Router(config-if)#ip hold-time eigrp 100 30
```

Authentication

EIGRP supports two forms of authentication: simple passwords and MD5 hashes. Simple passwords are sent as plain-text and matched to the key on the receiver. Simple passwords are not secure, because any listener can see this traffic and read the key value. Because simple passwords do not accomplish a change in the security structure in your network, the commands to implement them are not discussed here.

Hashed keys, sent as MD5 values, are secure because a listener cannot use the value in one transmission to compute the key.

Using MD5 Authentication, the router generates a hash value for every EIGRP transmission and checks the hash of every received EIGRP packet.

To specify MD5 authentication, the command is

```
router(config-if)#ip authentication mode eigrp autonomous-system md5
```

Once MD5 is specified, the command to set the key is

```
router(config-if)#ip authentication key-chain eigrp autonomous-system chain-name
```

Finally, the key-chain is configured and the key is specified:

```
router(config-if)#key chain chain-name
router(config-if)#key key-id
router(config-keychain-key)#key-string key
```

An example configuration in AS 100 might therefore look like the following:

```
router(config-if)#ip authentication mode eigrp 100 md5
router(config-if)#ip authentication key-chain eigrp 100 my-chain
router(config)#key chain my-chain
router(config-keychain)#key 1
router(config-keychain-key)#key-string secret
```

Authentication results are not shown under the **show** commands; authentication must be assumed based on the successful establishment of neighborship. However, the **debug eigrp packets** command does show the process.

Optional EIGRP Commands over WANs

There are EIGRP design and configuration issues concerning WANs. Here you must deal with limited capacity to a greater degree than at other points in the network. Furthermore, routing updates and data transmissions compete for that limited space.

EIGRP is unique in that it restricts its use of bandwidth. By default, it uses no more than half the link capacity, which is superior to the scant consideration shown by every other routing protocol. Although the default setting for EIGRP is usually sufficient, you may need to adjust it at times. The following sections describe the commands necessary to make these changes.

EIGRP Defaults in Bandwidth Utilization

Routers usually understand link capacity. Serial interfaces are problematic in this regard, though, because they usually attach through a DSU. The router therefore assumes a default serial speed of 1544 kbps.

If the link is 56 kbps, however, EIGRP would calculate an inaccurate metric and—even limiting itself to 772 kbps—could saturate the link. This could result in dropped EIGRP and data packets because of congestion.

The **show interface** command allows you to check that the interface bandwidth is accurate. The output shows the configured bandwidth of the link.

Example 4-8 shows the bandwidth of serial0.

Example 4-8 *Bandwidth*

```
Router# show interface serial 0
Serial0 is up, line protocol is up
Hardware is HD64570
Description: to Internet
Internet address is 172.25.146.182/30
MTU 1500 bytes, BW 1280 Kbit, DLY 20000 usec, rely 255/255, load 1/255
...
```

Note that, in this case, the bandwidth has been adjusted to 1280 kbps.

The bandwidth does not affect the actual speed of the link. It is used for routing-protocol calculations and load calculations. The following sections provide additional guidelines about bandwidth configuration.

In general, the bandwidth is set using the **bandwidth** command. The following syntax shows the structure of the command:

```
Router#interface S0
Router(config-if)#bandwidth speed-of-line
```

Configuring Bandwidth over an NBMA Cloud

EIGRP works well over all WANs, including point-to-point and NBMA environments such as Frame Relay or ATM. The NBMA topology can include either point-to-point subinterfaces or multipoint interfaces.

Cisco identifies three rules that you should follow when configuring EIGRP over an NBMA cloud:

- EIGRP traffic should not exceed the committed information rate (CIR) capacity of the virtual circuit (VC).
- EIGRP aggregated traffic over all the VCs should not exceed the access line speed of the interface.
- The bandwidth allocated to EIGRP on each VC must be the same in both directions.

If you understand and follow these rules, EIGRP works well over the WAN. If you do not take care in the configuration of the WAN, EIGRP can swamp the network.

NOTE Come back and look through this section after you have read about OSPF support for Frame Relay. OSPF over Frame Relay is complicated and contrasts with the simple implementation of EIGRP over Frame Relay. EIGRP handles whatever topology that is thrown at it without requiring a special configuration.

Configuring Bandwidth over a Multipoint Network

In addition to being used in the EIGRP metric, the **bandwidth** command influences how EIGRP uses NBMA virtual circuits (VCs). If the serial line has many VCs in a multipoint configuration, EIGRP will assume that each VC has an even share of bandwidth. EIGRP will confine itself to using half that share for its traffic. This limitation is important—no other routing protocol does this—but will work best if the bandwidth is correctly defined. A 56-kbps link that has a bandwidth set to 128 kbps will not benefit from restraint shown by EIGRP.

The **bandwidth** command should reflect the access-link speed into the Frame Relay cloud. Your company might have five PVCs from your router's serial interface, each carrying 56 kbps. The access link will need a capacity of $5 * 56$ kbps (280 kbps).

Configuring Bandwidth over a Hybrid Multipoint Network

If the multipoint network has differing speeds allocated to the VCs, a more complex solution is needed. There are two main approaches:

- Take the lowest CIR and multiply it by the number of circuits. Apply the product as the bandwidth of the physical interface. The problem with this configuration is that EIGRP will underutilize the higher-bandwidth links.
- If possible, it is much easier to configure and manage an environment that has used subinterfaces, where a VC is logically treated as if it were a separate interface. The **bandwidth** command can be configured on each subinterface, which allows different speeds on each VC. In this solution, subinterfaces are configured for each VC and the CIR is configured as the bandwidth. Cisco recommends this as the preferred solution.

Configuring the Pure Point-to-Point Network

If there are many VCs, there might not be enough bandwidth at the access speed of the interface to support the aggregate EIGRP traffic. The subinterfaces should be configured with a bandwidth that is much lower than the real speed of the circuit. In this case, it is necessary to use the **bandwidth-percent** command to indicate to the EIGRP process that it can still function.

The **ip bandwidth-percent-igrp** command adjusts the percentage of capacity that EIGRP may use from the default fifty percent. You would use this command because the **bandwidth** command does not reflect the true speed of the link. The **bandwidth** command might have been altered to manipulate the routing metric and path selection of a routing protocol (but it might be better to use other methods of controlling the routing metric and return the bandwidth to a true value). The following shows the structure of the **ip bandwidth-percent igrp** command:

```
Router(config)#interface S0
Router(config-if)#ip bandwidth-percent igrp autonomous-system-number percent
```

Verifying the EIGRP Operation

It is important to understand the output of the **show** commands to prepare for the exam. The ability to interpret these **show** command output examples in conjunction with the physical and logical topology diagrams of your organization will ensure your understanding of the operation of EIGRP.

This section describes the **show** commands shown in Table 4-4.

Table 4-4 *EIGRP show Commands*

Command Option	Description
show ip eigrp neighbors	Gives detailed information about neighbors. This command records the communication between the router and the neighbors in addition to the interface and address by which they communicate.
show ip eigrp topology	Gives details about the routes held in the topology table, detailed information on the networks that the router is aware of and the preferred paths to those networks, and the next hop in the path. Also allows the router to track EIGRP packets that have been sent to neighbors.
show ip eigrp traffic	Gives information about the aggregate traffic sent to and from the EIGRP process.

The EIGRP **show** commands are highly detailed and give a comprehensive understanding of the state of the network. Troubleshooting EIGRP should start with the generic routing commands **show ip route** and **show ip protocols**, before continuing with EIGRP-specific troubleshooting.

The show ip eigrp neighbors Command

The **show ip eigrp neighbors** command shows the neighbor table. The syntax is

```
Router#show ip eigrp neighbors [type number]
```

Example 4-9 shows the output of this command.

Example 4-9 *The show ip eigrp neighbors Command Output*

Table 4-5 explains the meaning of the important fields in Example 4-9.

Table 4-5 *Explanation of the show ip eigrp neighbors Command Output*

Field	Explanation
process 100	Autonomous system number.
Address	IP address of the EIGRP neighbor.
Interface	Interface the router receives hellos on.
Holdtime	Length of time, in seconds, that the router will wait before declaring the link down.
Uptime	Time (in hours, minutes, and seconds) since the router first heard from this neighbor.
Q Count	Number of EIGRP packets (update, query, and reply) that the router has queued.
Seq Num	The sequence number of the last packet received from the neighbor.
SRTT	Smooth round-trip time. The time (in milliseconds) from the sending of the packet to the receipt of an acknowledgment from the neighbor.
RTO	Retransmission timeout, in milliseconds. This shows how long the router will wait for an acknowledgment before it retransmits the packet.

The show ip eigrp topology Command

The **show ip eigrp topology** command shows the topology table. This command shows the effect of DUAL, which is described in Chapter 3. It shows whether the successor or the route is in an active or passive state and if there is a feasible successor. The syntax is

```
Router#show ip eigrp topology [autonomous-system-number | [[ip-address] mask]]
```

Example 4-10 shows the output of this command.

Example 4-10 *The show ip eigrp topology Command Output*

```

Router# show ip eigrp topology
IP-EIGRP Topology Table for process 100
Codes:P - Passive, A - Active, U - Update, Q - Query, R - Reply, r - Reply status
P 140.100.56.0 255.255.255.0, 2 successors, FD is 0
  via 140.100.32.22 (46251776/46226176), Ethernet0
  via 140.100.48.22 (46251776/46226176), Ethernet1
  via 140.100.32.31 (46277376/46251776), Ethernet0
P 140.100.48.0 255.255.255.0, 1 successors, FD is 307200
  via Connected, Ethernet1
  via 140.100.48.22 (307200/281600), Ethernet1
  140.100.32.22 (307200/281600), Ethernet0
  via 140.100.32.31 (332800/307200), Ethernet0

```

Table 4-6 explains the meaning of the important fields in Example 4-10.

Table 4-6 *Explanation of the show ip eigrp topology Command Output*

Field	Explanation
P	Passive—The router has not received any EIGRP input from a neighbor, and the network is assumed to be stable.
A	Active—When a route or successor is down, the router attempts to find an alternative path. The router is currently actively querying for a backup path.
U	Update—A value in this field identifies that the router has sent an update packet to a neighbor.
Q	Query—A value in this field identifies that the router has sent a query packet to a neighbor.
R	Reply—A value here shows that the router has sent a reply to the neighbor.
r	This is used in conjunction with the query counter; the router has sent out a query and is awaiting a reply.
140.100.48.0	Destination IP network number.
255.255.255.0	Destination subnet mask.
successors	This is the number of routes or the next logical hop. The number stated here is the same as the number of applicable routes in the routing table.
via	This is the next-hop address. The first of these entries are the current successors. The remaining entries on the list are feasible successors.
(46251776/46226176)	The first number is the EIGRP metric. After the slash is the EIGRP metric the peer advertised, or the advertised distance.
Ethernet0	The interface through which the EIGRP advertisement was received and the outgoing interface.

The show ip eigrp traffic Command

The **show ip eigrp traffic** command shows the EIGRP traffic received and generated by the router. The command syntax is

```
Router#show ip eigrp traffic [autonomous-system-number]
```

Example 4-11 shows the output of this command.

Example 4-11 The show ip eigrp traffic Command Output

```
Router# show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 100
Hellos sent/received: 218/205
Updates sent/received: 7/23
Queries sent/received: 2/0
Replies sent/received: 0/2
Acks sent/received: 21/14
```

Table 4-7 explains the meaning of the important fields in Example 4-11.

Table 4-7 Explanation of the **show ip eigrp traffic** Command Output

Field	Explanation
process 100	Autonomous system number
Hellos sent/received	Number of hello packets sent and received by the router
Updates sent/received	Number of update packets sent and received by the router
Queries sent/received	Number of query packets sent and received by the router
Replies sent/received	Number of reply packets sent and received by the router
Acks sent/received	Number of acknowledgment packets sent and received by the router

Troubleshooting EIGRP

Having many methods and tools helps in troubleshooting any network. One of the most beneficial tools is documentation, for several reasons. For example, you can see progress and easily eliminate the obvious in a checklist manner, and you can clearly explain the problem and the steps taken thus far in solving it if you need to call in expert help. Cisco provides many tools, both on its web page and in service contracts, to help solve your network problems. One of the mainstays in troubleshooting any routing protocol is the group of debug commands, which provide the ability to see traffic and router processes in real time.

Exercise care when using the **debug** command, because it can be very greedy in terms of the resources that it consumes. You should use it only for a specific option and for a finite time. When debugging is complete, the **undebug all** command will disable debugging.

The options available for monitoring EIGRP are covered in Table 4-8.

Table 4-8 *The debug Command Options for EIGRP*

Command Option	Description
debug ip eigrp packet	Shows the EIGRP packets sent and received by the router. The packet types to be monitored can be selected. Up to 11 types are available.
debug ip eigrp neighbors	Shows the hello packets sent and received by the router and the neighbors discovered by this process.
debug ip eigrp	Shows dynamic changes made to the routing table.
debug ip eigrp summary	Shows the process taken when a summary (manual or auto) is changed on the router.

Foundation Summary

The Foundation Summary provides a convenient review of many key concepts in this chapter. If you are already comfortable with the topics in this chapter, this summary might help you recall a few details. If you just read this chapter, this review should help solidify some key facts. If you are doing your final prep before the exam, the following tables are a convenient way to review the day before the exam.

Table 4-9 summarizes the commands covered in this chapter.

Table 4-9 Summary of EIGRP Commands

Command	Function
Router(config)# router eigrp <i>autonomous-system-number</i>	Starts the EIGRP processes on the router with the specified autonomous system number.
Router(config-router)# network <i>network-number</i> [<i>wildcard-mask</i>]	Advertises the networks that match the <i>network mask</i> .
Router(config-router)# no auto-summary	Disables automatic summarization.
Router(config-if)# ip summary-address eigrp <i>autonomous-system-number</i> <i>address mask</i>	Enables manual summarization.
Router(config-router)# eigrp stub [receive-only connected static summary]	Used on single-neighbor routers in a hub-and-spoke environment. Minimizes local routing table and advertises local subnets to core.
Router(config-if)# bandwidth <i>speed-of-line</i>	Identifies bandwidth on interfaces where router cannot automatically determine them. Used primarily on serial lines.
Router(config-if)# ip bandwidth-percent eigrp <i>autonomous-system-number</i> <i>percent</i>	EIGRP, by default, will take only 50% of bandwidth for routing protocol traffic. Allows adjustment of that percentage.
Router(config-router)# variance <i>multiplier</i>	Allows unequal-cost paths to proportionally load share.
Router(config-if)# ip hello-interval eigrp <i>autonomous-system-number seconds</i>	Changes hello interval. May be used to lengthen interval (reducing overhead) or to shorten the interval (decreasing convergence time).

Table 4-9 Summary of EIGRP Commands (Continued)

Command	Function
Router(config-if)# ip hold-time eigrp <i>autonomous-system-number seconds</i>	Defines how long to wait without hearing a hello from a neighbor before declaring it dead. By default, it is three times the length of the hello-interval.
Router# show ip eigrp neighbors	Displays information drawn from the neighbor table.
Router# show ip eigrp topology	Displays information drawn from the topology table.
Router# show ip eigrp traffic	Shows statistics about the EIGRP traffic reaching the router.
Router# debug eigrp packet	Shows EIGRP packets sent and received.
Router# debug ip eigrp neighbors	Shows the hello packets sent and received by the router and the neighbors discovered by this process.
Router# debug ip eigrp summary	Shows when a summary is changed on the router.

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. Imagine a router that has two interfaces—fastethernet 0/0 is 10.1.1.1 and serial 0/0 is 10.2.2.2. Nominally, EIGRP would be run with the “network 10.0.0.0” command, which would enable both interfaces. There are two ways discussed in this chapter that could be used to specify that one interface run EIGRP and not the other. What are those two techniques?
2. What is the preferred configuration for a hybrid multipoint NBMA network when one VC has a CIR of 56 kbps and the other five VCs each have a CIR of 256 kbps?
3. Four Frame Relay PVCs exist on a multipoint interface. On the multipoint interface, the command **bandwidth 224** is used. How much bandwidth will EIGRP assign to each PVC, and how much of that amount will EIGRP use for advertisements?
4. Explain the purpose of the command **no auto-summary**.
5. Explain the meaning of the command **ip bandwidth-percent eigrp 63 100**.
6. How long is the default hold time?
7. How is the **variance** command used?
8. What command is used to display the passive and active state of the routes?
9. What command is used in EIGRP to perform manual summarization?
10. Which command is used to display all types of EIGRP packets that are both received and sent by a router?
11. What problems can you solve by configuring summarization?
12. Why would you configure an EIGRP router as a stub router?
13. Explain the parameters **receive-only**, **connected**, **static**, and **summary** used in the command **eigrp stub**:

```
router(config-router)# eigrp stub [receive-only | connected | static |
summary]
```
14. When configuring the **variance** command, which routes can be used?

15. Give two reasons why you might wish to change the hello timer.
16. If a multipoint serial interface uses five PVCs with CIRs of 56-, 128-, 128-, 128-, and 256-kbps, how should the bandwidth command be implemented on the interface?
17. On what occasions should you consider configuring the bandwidth on subinterfaces?

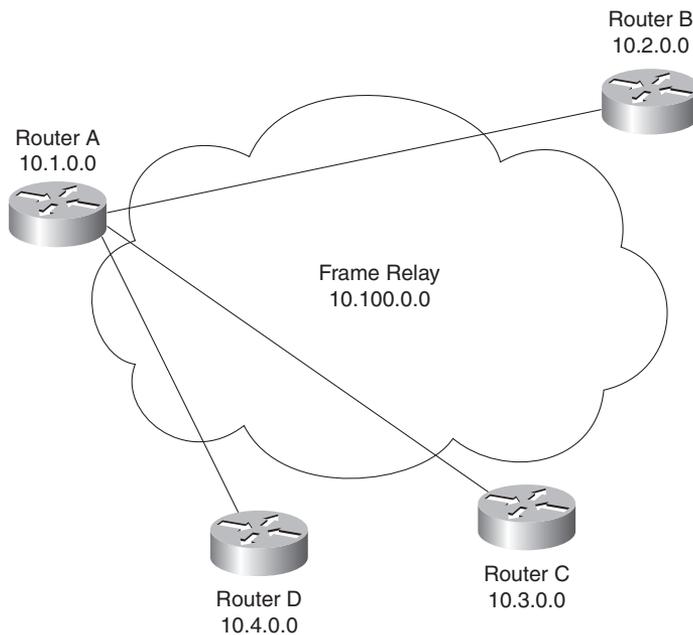
Scenarios

The following scenarios and questions are designed to draw together the content of the chapter and to exercise your understanding of the concepts. There is not necessarily a right answer. The thought process and practice in manipulating the concepts are the goals of this section. The answers to the scenario questions are found at the end of this chapter. The information used in these scenarios was adapted from the Cisco web page, “Cisco Configuration Guidelines.”

Scenario 4-1

The multinational company Gargantuan, Inc., has a private network addressed using the 10.0.0.0/8 space. Figure 4-3 shows the addressing scheme.

Figure 4-3 *Diagram for Scenario 4-1*



The network is experiencing timeouts and network crashes. In addition, EIGRP appears to be losing routes from its routing tables, which is adding to the problem.

Review the Gargantuan, Inc., addressing scheme and answer the following questions:

1. What changes to addressing or EIGRP could affect the route drops and network problems? State the configuration commands necessary to activate this solution on Router A.
2. The WAN is a Frame-Relay cloud, and Router A is the hub in the hub-and-spoke configuration. Each VC is 56 kbps. Give the commands to configure Router A for EIGRP over this NBMA cloud.
3. Give the commands to configure Router B for EIGRP over this NBMA cloud.

Scenario 4-2

Given the configuration of EIGRP in Example 4-12, perform the tasks and answer the questions listed. The WAN has light user traffic and is a hub-and-spoke configuration, as shown in Figure 4-4.

Example 4-12 Scenario 4-2 Configuration for Router A

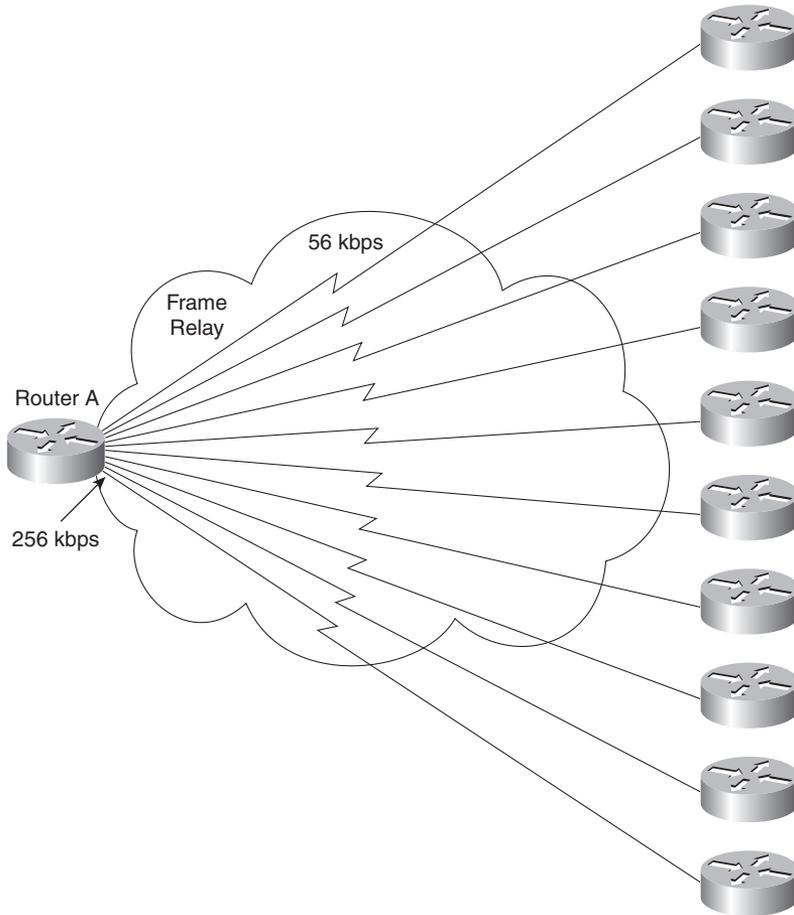
```
Router#show config
interface Serial 0
  encapsulation frame-relay

  interface Serial 0.1 point-to-point
  bandwidth 25
  ip bandwidth-percent eigrp 123 90

  interface Serial 0.2 point-to-point
  bandwidth 25
  ip bandwidth-percent eigrp 123 90

  ...
```

Figure 4-4 Diagram for Scenario 4-2



The 256-kbps access line to the hub has 56-kbps access lines to each of ten spoke sites. Each link has a Frame Relay CIR of 56 kbps. The access line to each router reflects the CIR. The access line to the hub router, Router A, is 256 kbps, but the CIR of the hub is the same as its access line.

From a Frame Relay perspective, a circuit is considered oversubscribed when the sum of CIRs of the remote circuits is higher than the CIR of the hub location. With ten links, each with a CIR of 56 kbps, this circuit is oversubscribed.

1. How much bandwidth has each circuit been allocated? Why was this value chosen by the administrator?

2. What bandwidth utilization is available to EIGRP? Why was this value chosen by the administrator?
3. If Router A fails, what would the effect be on the network?
4. Is summarization possible only on the routers entering the WAN cloud, or is it possible on the networks not shown in the figure that are on the other side of the routers? Give reasons for your answers.

Scenario Answers

The answers provided in this section are not necessarily the only possible answers to the questions. The questions are designed to test your knowledge and to give practical exercise in certain key areas. This section is intended to test and exercise skills and concepts detailed in the body of this chapter.

If your answer is different, ask yourself whether it follows the tenets explained in the answers provided. Your answer is correct not if it matches the solution provided in the book, but rather if it has included the principles of design laid out in the chapter.

If you do not get the correct answer, refer back to the text and review the subject tested. Be certain to also review your notes on the question to ensure that you understand the principles of the subject.

Scenario 4-1 Answers

1. What needs to be done in addition to solve the problems caused by EIGRP? State the configuration commands necessary to activate this solution on Router A.

The WAN is a Frame Relay cloud, and Router A is the hub in the hub-and-spoke configuration. Each VC is 56 kbps.

In addition to readdressing the network, summarization would limit the query range of the routers. This would prevent the routes in the topology table being SIA, which seriously affects the performance of the network.

The commands required are as follows:

```
RouterA(config)# router eigrp 63
RouterA(config)# no auto-summary
RouterA(config)# network 10.0.0.0
RouterA(config)# int serial 0
RouterA(config-if)# ip summary-address eigrp 63 10.1.0.0 255.255.0.0
```

2. Give the commands to configure Router A for EIGRP over this NBMA cloud.

The configuration on Router A is as follows:

```
RouterA(config)# interface serial 0
RouterA(config-if)# frame-relay encapsulation
RouterA(config-if)# bandwidth 178
```

3. Give the commands to configure Router B for EIGRP over this NBMA cloud.

The configuration on Router B is as follows:

```
RouterB(config)# interface serial 0  
RouterB(config-if)# frame-relay encapsulation  
RouterB(config-if)# bandwidth 56
```

Scenario 4-2 Answers

1. How much bandwidth has each circuit been allocated? Why was this value chosen by the administrator?

Because a maximum of 256 kbps is available, you cannot allow any individual PVC to handle more than 25 kbps (256/10). Note that EIGRP will not use more than 22.5 kbps (90 percent of 25 kbps) on this interface, even though its capacity is 56 kbps. This configuration will not affect user data capacity, which will still be able to use the entire 56 kbps, except when EIGRP is using 22 kbps.

2. What bandwidth utilization is available to EIGRP? Why was this value chosen by the administrator?

Because this data rate is low, and because you do not expect very much user data traffic, you can allow EIGRP to use up to 90 percent of the bandwidth.

3. If Router A fails, what would the effect be on the network?

If Router A fails, there would be no communication between the routers in the WAN because Router A is the hub. Each site would function, but they would all be isolated from each other. The neighbor tables would fail to hear the hellos from the other routers connecting to the WAN and would time out all routes that they had heard from these routers. The topology table would be updated, and the routers would send updates to all their other neighbors.

4. Is summarization possible only on the routers entering the WAN cloud, or is it possible on the networks not shown in the figure, that are on the other side of the routers? Give reasons for your answers.

Summarization is possible on all interfaces in EIGRP, as long as the addressing scheme allows for it to be implemented. This is one of the major advantages of EIGRP over OSPF. OSPF can summarize only at ABRs.