CHAPTER 3

EIGRP and Troubleshooting Routing Protocols

Objectives

Upon completion of this chapter, you should be able to answer the following questions:

- What are the features of balanced hybrid routing?
- What are the particular features of EIGRP?
- How does EIGRP compare with IGRP?
- How do you configure EIGRP?
- How do you verify the EIGRP configuration?
- What is a general process for troubleshooting routing protocols?
- How are **debug** commands used to troubleshoot a RIP configuration?
- How are **debug** commands used to troubleshoot an EIGRP configuration?
- How are **debug** commands used to troubleshoot an OSPF configuration?

Additional Topics of Interest

Some chapters contain additional coverage of previous topics or related topics that are secondary to the main goals of the chapter. You can find the additional coverage in the "Additional Topics of Interest" section near the end of the chapter. For this chapter, the following additional topic is covered:

- Troubleshooting IGRP

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary:

Diffusing Update Algorithm (DUAL) *page 67*

neighbor table *page 68*

topology table *page 68*

successor *page 68*

feasible successor *page 68*

Reliable Transport Protocol (RTP) *page 69*

hello packets *page 70*

passive state *page 70*

acknowledgment packets *page 70*

reply packets *page 70*

update packets *page 71*

query packets *page 71*

active state *page 71*

feasible distance *page 75*

EIGRP is a Cisco-proprietary routing protocol that is based on IGRP. EIGRP supports CIDR and VLSM, allowing network designers to maximize address space. When compared to IGRP, a classful routing protocol, EIGRP boasts faster convergence times, improved scalability, and superior management of routing loops.

EIGRP is often described as a hybrid routing protocol that offers the best of distance vector and link-state algorithms. EIGRP is an advanced routing protocol that relies on features commonly associated with link-state protocols. Some of the best features of OSPF, such as partial updates and neighbor discovery, are similarly put to use by EIGRP; however, EIGRP is easier to configure than OSPF. EIGRP is an ideal choice for large, multiprotocol networks built primarily on Cisco routers.

This chapter discusses common EIGRP configuration tasks. The emphasis is on ways in which EIGRP establishes relationships with adjacent routers, calculates primary and backup routes, and responds to failures in known routes to a particular destination.

A network is made up of many devices, protocols, and media that allow data communication to occur. When a network component does not work correctly, it can affect the entire network. In any case, network engineers must quickly identify and troubleshoot problems when they arise. The following are some reasons network problems occur:

- Commands are entered incorrectly.
- Access lists are constructed or placed incorrectly.
- Routers, switches, or other network devices are misconfigured.
- Physical connections are bad.

A network engineer should troubleshoot in a methodical manner with the use of a general problem-solving model. It is often useful to check for physical layer problems first and then move up the layers in an organized manner. Although this chapter closes with a focus on how to troubleshoot Layer 3 protocols, it is important to troubleshoot and eliminate any problems that might exist at the lower layers.

EIGRP Concepts

Balanced hybrid routing protocols combine aspects of both distance vector and link-state protocols. The balanced hybrid routing protocol uses distance vectors with more accurate metrics to determine the best paths to destination networks. However, the balanced hybrid routing protocol differs from most distance vector protocols in that it uses topology changes instead of automatic periodic updates to trigger the routing of database updates.

The balanced hybrid routing protocol converges more rapidly than distance vector routing protocols, which is similar to link-state routing protocols. However, the balanced hybrid differs from distance vector and link-state routing protocols in that it emphasizes economy in the use

of required resources, such as bandwidth, memory, and processor overhead. Enhanced Interior Gateway Routing Protocol (EIGRP) is an example of a balanced hybrid routing protocol.

EIGRP has several advantages over Routing Information Protocol (RIP) and Interior Gateway Routing Protocol (IGRP), and even some advantages over Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS). EIGRP's enhancements come with many complexities that take place behind the scenes. Although configuring EIGRP is relatively simple, the underlying protocol and algorithm are not so simple. This section describes EIGRP concepts, terminology, and features.

Comparing EIGRP and IGRP

EIGRP uses metric calculations similar to those that IGRP uses, and EIGRP supports the same unequal-cost path load balancing as IGRP. It is also important to note that Cisco IOS Release 12.2(13)T is the last version to support the legacy IGRP. The convergence properties and the operating efficiency of EIGRP are substantially improved compared with IGRP. EIGRP has a dramatically improved convergence time and reduced network overhead. Although the metric (bandwidth and delay, by default, and the option to use load and reliability) is the same for both IGRP and EIGRP, the weight assigned to the metric is 256 times greater for EIGRP. Automatic redistribution occurs between IGRP and EIGRP if they are using the same autonomous system number. Also of note is that EIGRP has a maximum hop count of 224 and supports route tagging during redistribution.

The convergence technology, which is based on research conducted at SRI International by Dr. J.J. Garcia-Luna-Aceves, employs ***Diffusing Update Algorithm (DUAL)***. This algorithm guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize simultaneously. Routers that are not affected by topology changes are not involved in recomputations. The convergence time with DUAL rivals that of any other existing routing protocol.

EIGRP Features

In a well-designed network, EIGRP scales well and provides extremely quick convergence times with minimal network traffic. Some of the features of EIGRP are as follows:

- EIGRP has rapid convergence times for changes in the network topology. In some situations, convergence can be almost instantaneous. EIGRP uses DUAL to achieve rapid convergence. A router that runs EIGRP stores backup routes for destinations when they are available so that it can quickly adapt to alternate routes. If no appropriate route or backup route exists in the local routing table, EIGRP queries its neighbors to discover an alternate route. These queries are propagated until an alternate route is found.
- EIGRP has low usage of network resources during normal operation; only hello packets are transmitted on a stable network. Like other link-state routing protocols, EIGRP uses EIGRP hello packets to establish relationships with neighboring EIGRP routers. Each

router builds a neighbor table from the hello packets that it receives from adjacent EIGRP routers. EIGRP does not send periodic routing updates like IGRP does. When a change occurs, routing table changes are only propagated, not the entire routing table. When changes are only propagated, the bandwidth required for EIGRP packets is minimized, which reduces the load that the routing protocol itself places on the network.

- EIGRP supports automatic (classful) route summarization at major network boundaries as the default. However, unlike other classful routing protocols, such as IGRP and RIP, manual route summarization can be configured on arbitrary network boundaries to reduce the size of the routing table.

EIGRP Terminology

EIGRP relies on various tables for its computations. These include the neighbor table, the topology table, and the routing table. Table 3-1 summarizes the terms related to EIGRP.

Table 3-1 EIGRP Terminology

Term	Definition
Neighbor table (AppleTalk, Internetwork Packet Exchange [IPX], IPv4, IPv6)	Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the adjacencies database that OSPF uses, and it serves the same purpose (to ensure bidirectional communication between each of the directly connected neighbors). There is a neighbor table for each protocol that EIGRP supports.
Topology table (AppleTalk, IPX, IPv4, IPv6)	Each EIGRP router maintains a topology table for each configured routed protocol. This table includes route entries for all destinations that the router has learned.
Routing table v4, IPv6	EIGRP chooses the best (successor) routes to a destination from the topology table and places these routes in the routing table. The router maintains one routing table for each network protocol.
Successor	A route selected as the primary route to reach a destination. Successors (up to four) are the entries kept in the routing table.
Feasible successor	Considered a backup route. Backup routes are selected when the successors are identified; however, these routes are kept in a topology table. Multiple feasible successors for a destination can be retained.

Figure 3-1 displays the routing protocols supported by EIGRP.

Figure 3-1 Routing Protocols Supported by EIGRP

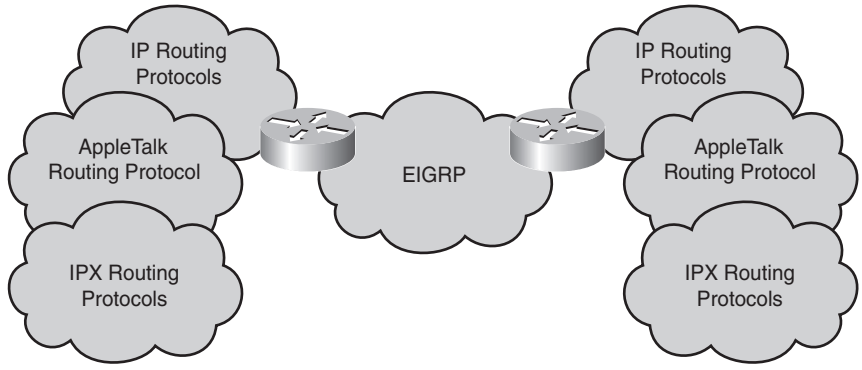

Figure 3-2 illustrates the fundamental contents of each table that EIGRP uses.

Figure 3-2 Contents of the Tables Used by EIGRP

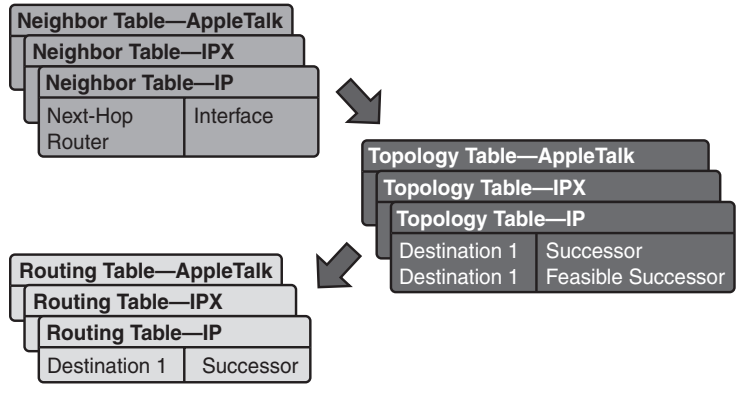

Reliable Transport Protocol (RTP) is a transport layer protocol that guarantees ordered delivery of EIGRP packets to all neighbors. On an IP network, hosts use Transmission Control Protocol (TCP) to sequence packets and ensure their timely delivery. However, EIGRP is protocol-independent, which means that it does not rely on Transmission Control Protocol/Internet Protocol (TCP/IP) to exchange routing information the way that RIP, IGRP, and OSPF do. To stay independent of IP, EIGRP uses RTP as its own proprietary transport layer protocol to guarantee delivery of routing information.

EIGRP can call on RTP to provide reliable or unreliable service as the situation warrants. With RTP, EIGRP can simultaneously multicast and unicast to different peers, which allows for maximum efficiency.

EIGRP Packet Types

Like OSPF, EIGRP relies on different packet types to maintain its tables and establish relationships with neighbor routers. EIGRP uses the following five types of packets:

- Hello
- Acknowledgment
- Update
- Query
- Reply

EIGRP relies on ***hello packets*** to discover, verify, and rediscover neighbor routers. Rediscovery occurs if EIGRP routers do not receive hellos from each other for a hold time interval but then reestablish communication.

Hello packets are always unreliably sent. This means that no acknowledgment is transmitted. EIGRP routers send hello packets at a fixed interval called the hello interval. The default hello interval depends on the interface's bandwidth. On IP networks, EIGRP routers send hello packets to the multicast IP address 224.0.0.10. On low-speed (T1 or slower) NBMA networks, hello packets are sent every 60 seconds; for all other networks, the hello interval is 5 seconds.

The neighbor table includes the Sequence Number field to record the number of the last received EIGRP packet that each neighbor sent. The neighbor table also includes a Hold Time field, which records the time the last packet was received. Packets must be received within the hold time interval period to maintain a ***passive state***, which is a reachable and operational status.

If EIGRP does not receive a packet from a neighbor within the hold time, EIGRP considers that neighbor down. DUAL then steps in to reevaluate the routing table. By default, the hold time is three times the hello interval, but an administrator can configure both timers as desired.

OSPF requires neighbor routers to have the same hello and dead intervals to communicate. EIGRP has no such restriction. Neighbor routers learn about each of the other respective timers through the exchange of hello packets. They then use that information to forge a stable relationship regardless of unlike timers.

EIGRP routers use ***acknowledgment packets*** to indicate receipt of any EIGRP packet during a reliable exchange. RTP provides reliable communication between EIGRP hosts. The recipient must acknowledge a message that is received to make it reliable. Acknowledgment packets, which are hello packets without data, are used for this purpose. Unlike multicast hello packets, acknowledgment packets are unicast. Acknowledgments can be attached to other kinds of EIGRP packets, such as ***reply packets***.

Update packets are used when a router discovers a new neighbor. EIGRP routers send unicast update packets to that new neighbor so that the neighbor can add to its topology table. More than one update packet can be needed to convey all the topology information to the newly discovered neighbor.

Update packets are also used when a router detects a topology change. In this case, the EIGRP router sends a multicast update packet to all neighbors, which alerts them to the change. All update packets are reliably sent.

An EIGRP router uses ***query packets*** whenever it needs specific information from one or all of its neighbors. A reply packet is used to respond to a query.

If an EIGRP router loses its successor and cannot find a feasible successor for a route, DUAL places the route in the ***active state***. A query is then multicasted to all neighbors in an attempt to locate a successor to the destination network. Neighbors must send replies that either provide information on successors or indicate that no information is available. Queries can be multicast or unicast, while replies are always unicast. Both packet types are reliably sent.

EIGRP Configuration

Configuring EIGRP is similar to configuring RIP and IGRP. In fact, EIGRP is most similar to RIP version 2 (RIPv2) in its configuration syntax and configuration options. This section explores basic EIGRP configuration, EIGRP configuration examples, and how to verify EIGRP configurations.

Basic EIGRP Configuration

Use the **router eigrp** and **network** commands to create an EIGRP routing process:

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

autonomous-system-number identifies all routers that belong within the internetwork. The number does not have to be registered, but it must match all routers within the internetwork.

The **network** command assigns a major network number to which the router is directly connected. Indicate which networks belong to the EIGRP autonomous system (AS) on the local router with the *network-number*. The EIGRP routing process associates interface addresses with the advertised network number and begins EIGRP packet processing on the specified interfaces.

Figure 3-3 displays a simple network. Example 3-1 shows the basic EIGRP configuration for the three routers in Figure 3-3.

Figure 3-3 Simple EIGRP Network

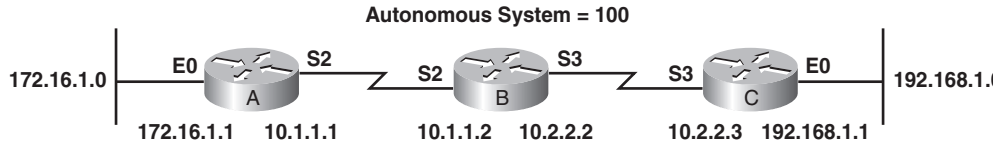

Example 3-1 Enabling EIGRP

```
RouterA(config)#router eigrp 100
RouterA(config-router)#network 172.16.0.0
RouterA(config-router)#network 10.0.0.0

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

RouterC(config)#router eigrp 100
RouterC(config-router)#network 192.168.1.0
RouterC(config-router)#network 10.0.0.0
```

Table 3-2 describes the router A configuration.

Table 3-2 Router A Command Descriptions

Command	Description
router eigrp 100	Enables the EIGRP routing process for AS 100
network 172.16.0.0	Associates network 172.16.0.0 with the EIGRP routing process
network 10.0.0.0	Associates network 10.0.0.0 with the EIGRP routing process

On router A, EIGRP sends updates out the interfaces in networks 10.0.0.0 and 172.16.0.0. The updates include information about networks 10.0.0.0, 172.16.0.0, and any other networks about which EIGRP learns.

When configuring serial links using EIGRP, it is important to configure the bandwidth setting on the interface. If the bandwidth for these interfaces is not changed, EIGRP assumes the default bandwidth on the link instead of the true bandwidth. If the link is slower, the router might not be able to converge, routing updates might become lost, or suboptimal path selection might result. To set the interface bandwidth, use the following syntax:

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

The **bandwidth** command is only used by the routing process and must be set to match the line speed of the interface.

Cisco Systems also recommends adding the following command to all EIGRP configurations:

```
Router(config-router)#eigrp log-neighbor-changes
```

This command enables the logging of neighbor adjacency changes to monitor the stability of the routing system and to help detect problems. By default, this command is enabled.

Lab 3.2.1 Configuring EIGRP Routing

In this lab, you configure EIGRP routing.

Configuring EIGRP Summarization

Prior to Cisco IOS Release 12.2(8)T, EIGRP automatically summarized routes at the classful boundary. The classful boundary is the boundary where the network address ends, as defined by class-based addressing. This means that although router RTC in Figure 3-4 is connected to subnet 2.1.1.0, it advertises that it is connected to the entire Class A network, 2.0.0.0. In some cases, autosummarization is beneficial because it keeps routing tables as compact as possible. However, over time, it has become general consensus that it is best not to have the router automatically summarize at the classful boundary, as evidenced by Cisco Systems move to disable autosummarization as the default behavior for EIGRP.

Figure 3-4 Effect of Autosummarization Is to Summarize at the Classful Boundary

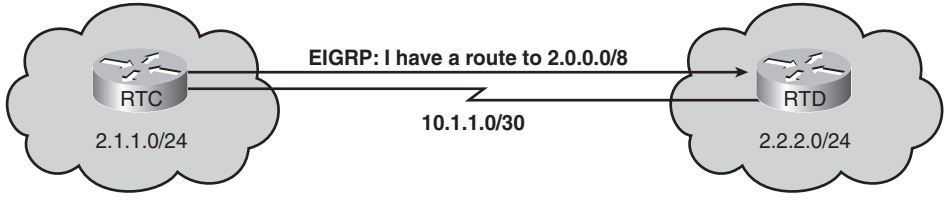

In many instances, autosummarization is not the preferred option. For example, if discontiguous subnetworks exist, autosummarization must be disabled for routing to work properly, Figure 3-5 illustrates. Autosummarization prevents routers from learning about discontinguous subnets; with summarization turned off, EIGRP routers will advertise subnets. To turn off autosummarization, use the following command:

```
Router(config-router)#no auto-summary
```

With EIGRP, a summary address can be manually configured by configuring a network prefix. With EIGRP, manual summary routes are configured on a per-interface basis, so the interface that propagates the route summary must be selected first. Then, the summary address can be defined with the **ip summary-address eigrp** command:

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

Figure 3-5 Discontiguous Networks with and Without Autosummarization

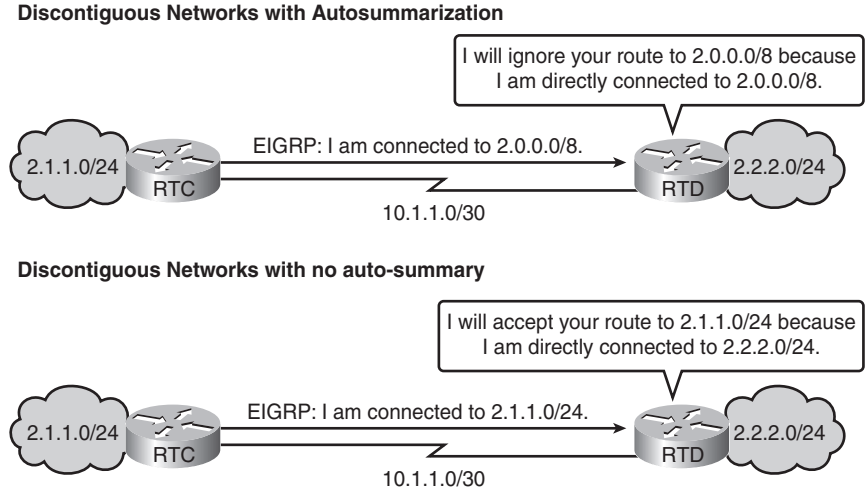

By default, EIGRP summary routes have an administrative distance of 5. Optionally, they can be configured for a value between 1 and 255.

In Figure 3-6, router RTC can be configured by using the commands shown in Example 3-2.

Figure 3-6 Granular Routing Updates with Interface Summarization

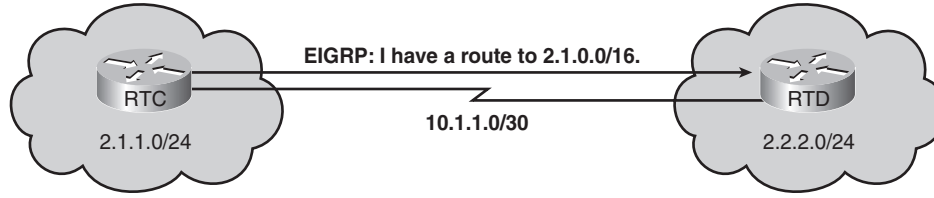

Example 3-2 Using Interface Summarization with EIGRP

```
RTC(config)#router eigrp 2446
RTC(config-router)#no auto-summary
RTC(config-router)#exit
RTC(config)#interface serial 0/0
RTC(config-if)#ip summary-address eigrp 2446 2.1.0.0 255.255.0.0
```

Router RTC adds a route to its table as follows:

```
D 2.1.0.0/16 is a summary, 00:00:22, Null0
```

Notice that the summary route is sourced from Null0 and not from an actual interface. This is because this route is used for advertisement purposes and does not represent a path that router RTC can take to reach that network. On router RTC, this route has an administrative distance of 5.

Router RTD is not aware of the summarization, but it accepts the route. The route is assigned the administrative distance of a normal EIGRP route, which, by default, is 90.

In the configuration for router RTC, autosummarization is turned off with the **no auto-summary** command. If autosummarization was not turned off, router RTD would receive two routes: the manual summary address, which is 2.1.0.0/16; and the automatic, classful summary address, which is 2.0.0.0/8. Normally, when manually summarizing, the **no auto-summary** command needs to be issued.

Verifying the EIGRP Configuration

As with OSPF, numerous **show** commands verify the EIGRP configuration. Table 3-3 summarizes these commands.

Table 3-3 EIGRP **show** Commands

Command	Description
show ip eigrp neighbors	Displays neighbors discovered by EIGRP.
show ip eigrp topology	Displays the EIGRP topology table. This command shows the topology table, the active or passive state of routes, the number of successors, and the feasible distance to the destination. ***Feasible distance*** is the best metric along a path to a destination network, including the metric to the neighbor advertising that path.
show ip route eigrp	Displays the current EIGRP entries in the routing table.
show ip protocols	Displays the parameters and current state of the active routing protocol process. This command shows the EIGRP AS number. It also displays filtering and redistribution numbers and neighbors and distance information.
show ip eigrp traffic	Displays the number of EIGRP packets sent and received. This command displays statistics on hello packets, updates, queries, replies, and acknowledgments.

Many network engineers use the **show ip eigrp neighbors** command when first configuring EIGRP to ensure that neighbor relationships are forming (without which no EIGRP routing can occur).

Lab 3.2.3 Verifying Basic EIGRP Configuration

In this lab, you verify EIGRP routing.

Troubleshooting Routing Protocols

Routing-protocol troubleshooting needs to begin with a logical sequence or process flow. This process flow is not a rigid outline for troubleshooting an internetwork; however, it is a foundation from which a network engineer can build a problem-solving process to suit a particular environment:

Step 1 When analyzing a network failure, make a clear problem statement:

- Define the problem in terms of a set of symptoms and potential causes.
- To properly analyze the problem, identify the general symptoms and then ascertain what kinds of problems or causes might result in these symptoms. For example, hosts might not be responding to service requests from clients, which is a symptom.
- Possible causes might include a misconfigured access host, bad interface cards, or missing router configuration commands.

Step 2 Gather the facts needed to help isolate possible causes:

- Gather the facts needed to help isolate possible causes. Ask questions of affected users, network administrators, managers, and other key people.
- Collect information from sources such as network management systems, protocol analyzer traces, output from router diagnostic commands, or software release notes.

Step 3 Consider possible problems based on the facts that have been gathered:

- Using these facts helps eliminate some of the potential problems from the list.
- Depending on the data, it might be possible to eliminate hardware as a problem, so you can then focus on software problems.
- At every opportunity, try to narrow the number of potential problems to create an efficient action plan.

Step 4 Create an action plan based on the remaining potential problems:

- Begin with the most likely problem and devise a plan in which only one variable is changed.
- Changing only one variable at a time helps to reproduce a given solution to a specific problem. Do not try to alter more than one variable at the same time. Such an action might solve the problem. However, identifying the specific change that eliminated the symptom becomes far more difficult and will not help to solve the same problem if it occurs in the future.

Step 5 Implement the action plan, performing each step carefully while testing to see whether the symptom disappears.

Step 6 Analyze the results to determine whether the problem has been resolved. If it has, the process is complete.

Step 7 If the problem has not been resolved, create an action plan based on the next most likely problem in the list. Return to Step 4, change one variable at a time, and repeat the process until the problem is solved.

Step 8 After the actual cause of the problem is identified, try to solve it:

- At this point, it is important to document the problem and the solution for future reference.
- If all attempts to this point have failed, it might now be necessary to ask for technical support from the manufacturer of the suspect equipment.
- Alternative resources include professional experts or technical engineers to help complete the troubleshooting process.

Cisco routers provide numerous integrated commands to assist you in monitoring and troubleshooting an internetwork:

- **show** commands help monitor installation behavior, normal network behavior, and isolate problem areas.
- **debug** commands assist in the isolation of protocol and configuration problems.
- TCP/IP network tools such as ping, traceroute, and Telnet help to isolate the OSI layer where the problem exists, as well as the location of the problem.

Cisco IOS **show** commands are among the most important tools for understanding the status of a router, detecting neighboring routers, monitoring the network in general, and isolating problems in the network. Chapter 1, "Introduction to Classless Routing," Chapter 2, "Single-Area OSPF," and this chapter describe the various **show** commands used with RIP, OSPF, and EIGRP. (Note that no **show** commands are specific to IGRP.)

Cisco routers provide numerous **debug** commands to assist you in troubleshooting an internetwork. EXEC **debug** commands can provide a wealth of information about interface traffic, internal error messages, protocol-specific diagnostic packets, and other useful troubleshooting data. **debug** commands isolate problems; they do not monitor normal network operation. **debug** commands look for specific types of traffic or problems. Before using a **debug** command, narrow the problems to a likely subset of causes. The **show debugging** command views which debugging features are enabled.

The remainder of this chapter explores the particular troubleshooting techniques and various **debug** commands used when troubleshooting RIP, EIGRP, and OSPF.

Troubleshooting RIP

The most common problem found in RIP that prevents RIP routes from being advertised is variable-length subnet mask (VLSM). This is because RIP version 1 (RIPv1) does not support VLSM. If the RIP routes are not being advertised, check the following:

- Layer 1 or Layer 2 connectivity issues exist.
- VLSM subnetting is configured. VLSM subnetting cannot be used with RIPv1.
- Mismatched RIPv1 and RIPv2 routing configurations exist.
- Network statements are missing or are incorrectly assigned.
- The outgoing interface is down.
- The advertised network interface is down.

Use the **debug ip rip** EXEC command to display information on RIP routing transactions. **no debug ip rip** turns off debugging for RIP. In general, the **no debug all** or **undebug all** command turns off all debugging.

Example 3-3 shows the **debug ip rip** output on router A of Figure 3-7.

Figure 3-7 Sample RIP Network for Troubleshooting

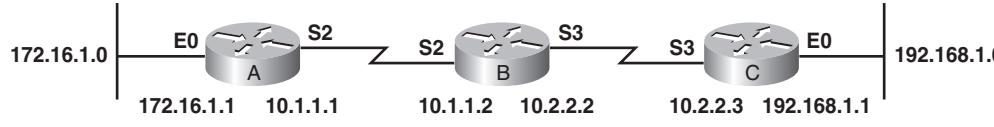

Example 3-3 Troubleshooting with **debug ip rip**

```
Router#debug ip rip
RIP protocol debugging is on
RouterA#
00:06:24: RIP: received v1 update from 10.1.1.2 on Serial2
00:06:24: 10.2.2.0 in 1 hops
00:06:24: 192.168.1.0 in 2 hops
00:06:33: RIP: sending v1 update to 255.255.255.255 via Ethernet0 (172.16.1.1)
00:06:34: network 10.0.0.0, metric 1
00:06:34: network 192.168.1.0, metric 3
00:06:34: RIP: sending v1 update to 255.255.255.255 via Serial2 (10.1.1.1)
00:06:34: network 172.16.0.0, metric 1
```

Example 3-3 shows that the router being debugged has received updates from one router at source address 10.1.1.2. That router sent information about two destinations in the routing table update. The router being debugged also sent updates (in both cases, to broadcast address 255.255.255.255 as the destination). The number in parentheses is the source address that is encapsulated into the IP header.

Other output that you might see from the **debug ip rip** command includes entries such as the following:

```
RIP: broadcasting general request on Ethernet0
RIP: broadcasting general request on Ethernet1
```

Entries like these can appear at startup or when an event occurs, such as when an interface transitions or a user manually clears the routing table.

The following output shows an entry most likely caused by a malformed packet from the transmitter:

```
RIP: bad version 128 from 160.89.80.43.
```

Troubleshooting EIGRP

Normal EIGRP operation is stable, efficient in bandwidth utilization, and relatively simple to monitor and troubleshoot.

Some possible reasons why EIGRP might not work correctly are as follows:

- Layer 1 or Layer 2 connectivity issues exist.
- AS numbers on EIGRP routers are mismatched.
- The link might be congested or down.
- The outgoing interface is down.
- The advertised network interface is down.
- Autosummarization is enabled on routers with discontiguous subnets. Use the **no auto-summary** command to disable automatic network summarization.

One of the most common reasons for a missing neighbor is a failure on the actual link. Another possible cause of missing neighbors is an expired hold-down timer. Because hellos are sent every 5 seconds on most networks, the hold time value in a **show ip eigrp neighbors** command output should normally be a value between 10 and 15.

The **debug ip eigrp** privileged EXEC command helps you analyze the packets pertaining to EIGRP routing that are sent and received on an interface, as Example 3-4 shows.

Example 3-4 Using **debug ip eigrp** to Troubleshoot

```
Router#debug ip eigrp
IP-EIGRP: Processing incoming UPDATE packet
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM 360960 - 256000
  104960
IP-EIGRP: Ext 192.168.0.0 255.255.255.0 M 386560 - 256000 130560 SM 360960 - 256000
  104960
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM 360960 - 256000
  104960
IP-EIGRP: 172.69.43.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.68.43.0 255.255.255.0 metric 371200 - 25600 115200
IP-EIGRP: 192.135.246.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.246.0 255.255.255.0 metric 46310656 - 45714176 596480
IP-EIGRP: 172.69.40.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.68.40.0 255.255.255.0 metric 2272256 - 1657856 614400
IP-EIGRP: 192.135.245.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.245.0 255.255.255.0 metric 40622080 - 40000000 622080
IP-EIGRP: 192.135.244.0 255.255.255.0, - do advertise out Ethernet0/1
```

Because the **debug ip eigrp** command generates a substantial amount of output, use it only when traffic on the network is light. Table 3-4 describes some fields in the output from the **debug ip eigrp** command shown in Example 3-4.

Table 3-4 **debug ip eigrp** Output Fields

Field	Description
IP-EIGRP:	Indicates that this is an IP EIGRP packet.
Ext	Indicates that the following address is an external destination rather than an internal destination, which would be labeled as "Int."
M	Displays the computed metric, which includes SM and the cost between this router and the neighbor. The first number is the composite metric. The next two numbers are the inverse bandwidth and the delay, respectively.
SM	Displays the metric as reported by the neighbor.

The **debug eigrp fsm** command is used for EIGRP debugging. This command displays information on DUAL feasible successor metrics and helps network engineers analyze the packets that are sent and received on an interface.

Troubleshooting OSPF

The majority of problems encountered with OSPF relate to the formation of adjacencies and the synchronization of the link-state databases.

To display information on OSPF-related events, such as adjacencies, flooding information, designated router selection, and SPF calculation, use the **debug ip ospf events** command. Example 3-5 shows output from the **debug ip ospf events** command.

Example 3-5 **debug ip ospf events** Is a Useful Troubleshooting Command

```
Router1#debug ip ospf events
OSPF events debugging is on
OSPF: hello with invalid timers on interface Ethernet0
hello interval received 10 configured 10
net mask received 255.255.255.0 configured 255.255.255.0
dead interval received 40 configured 30
```

The **debug ip ospf events** output that Example 3-5 shows might appear if any of the following situations occur:

- The IP subnet masks for routers on the same network do not match.
- The OSPF hello interval for the router does not match that configured for a neighbor.
- The OSPF dead interval for the router does not match that configured for a neighbor.

If a router configured for OSPF routing is not seeing an OSPF neighbor on an attached network, perform the following tasks:

- Make sure that both routers have been configured with the same IP mask, OSPF hello interval, and OSPF dead interval.
- Make sure that both neighbors are part of the same area type.

In the following line of sample output, the neighbor and this router are not both part of a stub area (stub areas are explored in CCNP); that is, one is a part of a transit area and the other is a part of a stub area, as explained in RFC 1247:

```
OSPF: hello packet with mismatched E bit
```

To display information about each OSPF packet received, use the **debug ip ospf packet** privileged EXEC command. Example 3-6 shows the sample output.

Example 3-6 debug ip ospf packet Provides Detailed Output Relating to OSPF

```
Router#debug ip ospf packet
OSPF: rcv. v:2 t:1 l:48 rid:200.0.0.117
      aid: 0.0.0.0 chk:6AB2 aut:0 auk:
      rcv. v:2 t:1 l:48 rid:200.0.0.116
      aid:0.0.0.0 chk:0 aut:2 keyid:1 seq:0x0
```

The **debug ip ospf packet** command produces one set of information for each packet received. The output varies slightly depending on which authentication is used. Table 3-5 gives a description of the output in Example 3-6.

Table 3-5 debug ip ospf packet Output Fields

Field	Description
v:	OSPF version
t:	OSPF packet type; possible packet types are as follows:
	1—Hello
	2—Data description
	3—Link-state request
	4—Link-state update
	5—Link-state acknowledgment
l:	OSPF packet length in bytes
rid:	OSPF router ID
aid:	OSPF area ID
chk:	OSPF checksum
aut:	OSPF authentication type; possible authentication types are as follows:
	0—No authentication
	1—Simple password
	2—MD5
auk:	OSPF authentication key
keyid:	MD5 key ID
seq:	Sequence number

As you can see from the output in Example 3-6, and referencing Table 3-5, MD5 authentication is in use.

Additional Topics of Interest

Some chapters of this book include additional topics of interest, which typically cover either more details about previous topics or topics that are optional or secondary to the chapter's main goals.

This chapter's "Additional Topics of Interest" section provides additional details of how to troubleshoot IGRP.

Troubleshooting IGRP

IGRP is a distance vector routing protocol that Cisco Systems developed in the 1980s. IGRP has several features that differentiate it from other distance vector routing protocols, such as RIP.

If IGRP does not appear to be working correctly, check the following:

- Layer 1 or Layer 2 connectivity issues exist.
- AS numbers on IGRP routers are mismatched.
- Network statements are missing or are incorrectly assigned.
- The outgoing interface is down.
- The advertised network interface is down.

To view IGRP debugging information, use the following commands:

```
debug ip igrp transactions [ip-address]
debug ip igrp events [ip-address]
```

Use **debug ip igrp transactions** [*ip-address*] to display IGRP transaction information.

The *ip-address* parameter is optional and indicates the IP address of an IGRP neighbor. If this option is used, the output includes only messages describing updates from that neighbor and updates that the router broadcasts toward that neighbor.

Entries such as the following occur on startup or when some event occurs, such as an interface making a transition or a user manually clearing the routing table:

```
IGRP: broadcasting request on Ethernet0
IGRP: broadcasting request on Ethernet1
```

The following type of entry can result when routing updates become corrupted between sending and receiving routers:

```
IGRP: bad checksum from 172.69.64.43
```

Use **debug ip igrp events** [*ip-address*] to display summary information on IGRP routing messages that indicate the source and destination of each update and the number of routes in each update.

To see how these commands are used to troubleshoot, first see Figure 3-8. Example 3-7 provides sample **debug ip igrp transactions** output for router A in Figure 3-8.

Figure 3-8 Sample IGRP Network for Troubleshooting

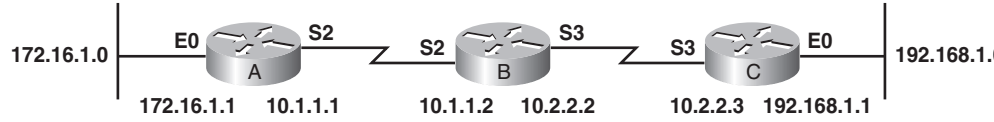

Example 3-7 Troubleshooting with **debug ip igrp transactions**

```
Router#debug ip igrp transactions
RouterA#:
00:21:06: IGRP: sending update to 255.255.255.255 via Ethernet0 (172.16.1.1)
00:21:06: network 10.0.0.0, metric=88956
00:21:06: network 192.168.1.0, metric=91056
00:21:07: IGRP: sending update to 255.255.255.255 via Serial2 (10.1.1.1)
00:21:07: network 172.16.0.0, metric=1100
00:21:16: IGRP: received update from 10.1.1.2 on Serial2
00:21:16: subnet 10.2.2.0, metric 90956 (neighbor 88956)
00:21:16: network 192.168.1.0, metric 91056 (neighbor 89056)
```

The output in Example 3-7 shows that the router being debugged has received an update from the router at source address 10.1.1.2, including information about two destinations (the networks being advertised). The fields are the same as in the sending output, but the metric in parentheses indicates the metric advertised by the neighbor sending the information. "Metric…inaccessible" usually means that the neighbor router has put the destination in a hold-down state.

When many networks exist in your routing table, displaying every update for every route can flood the console and make the router unusable. In this case, use the **debug ip igrp events** command to display a summary of IGRP routing information. The output of this command indicates the source and destination of each update and the number of routes in each update. Messages are not generated for each route. Example 3-8 illustrates typical output of **debug ip igrp events**. This output comes from router A in Figure 3-8.

In Figure 3-8, router A exchanges update IGRP messages with its neighbors. The router that is being debugged has sent two updates (in both cases, to broadcast address 255.255.255.255 as the destination address). The type of route information is categorized as subnet (interior), network (system), or exterior (exterior). The number of each type of route and the total number of routes are also indicated.

Example 3-8 Troubleshooting with **debug ip igrp events**

```
Router#debug ip igrp events
IGRP event debugging is on
RouterA#
00:23:44: IGRP: sending update to 255.255.255.255 via Ethernet0 (172.16.1.1)
00:23:44: IGRP: Update contains 0 interior, 2 system, and 0 exterior routes.
00:23:44: IGRP: Total routes in update: 2
00:23:44: IGRP: sending update to 255.255.255.255 via Serial2 (10.1.1.1)
00:23:45: IGRP: Update contains 0 interior, 1 system, and 0 exterior routes.
00:23:45: IGRP: Total routes in update: 1
00:23:48: IGRP: received update from 10.1.1.2 on Serial2
00:23:48: IGRP: Update contains 1 interior, 1 system, and 0 exterior routes.
00:23:48: IGRP: Total routes in update: 2
```

To delve into more detail about using the **debug ip igrp transactions** command, here is a troubleshooting scenario (see Figure 3-9).

Figure 3-9 IGRP Network Fails

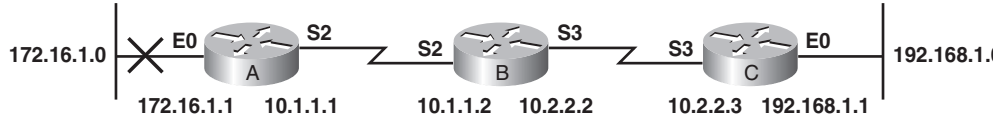

In Figure 3-9, the Ethernet network attached to router A fails. Router A sends a triggered update to router B that indicates that network 172.16.0.0 is inaccessible (with a metric of 4294967295), as Example 3-9 illustrates. Router B sends back a poison reverse update.

Example 3-9 Troubleshooting an IGRP Network (Router A)

```
RouterA#
00:31:15: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0, changed state to
  down
00:31:15: IGRP: edition is now 3
00:31:15: IGRP: sending update to 255.255.255.255 via Serial2 (10.1.1.1)
00:31:15: network 172.16.0.0, metric=4294967295
00:31:16: IGRP: Update contains 0 interior, 1 system, and 0 exterior routes.
00:31:16: IGRP: Total routes in update: 1
00:31:16: IGRP: broadcasting request on Serial2
00:31:16: IGRP: received update from 10.1.1.2 on Serial2
00:31:16: subnet 10.2.2.0, metric 90956 (neighbor 88956)
00:31:16: network 172.16.0.0, metric 4294967295 (inaccessible)
00:31:16: network 192.168.1.0, metric 91056 (neighbor 89506)
00:31:16: IGRP: Update contains 1 interior, 2 system, and 0 exterior routes.
00:31:16: IGRP: Total routes in update: 3
```

In Example 3-10, router B receives the triggered update from router A, sends a poison reverse to router A, and sends a triggered update to router C, thereby notifying both routers that network 176.16.0.0 is "possibly down."

Example 3-10 Troubleshooting an IGRP Network (Router B)

```
RouterB#debug ip igrp transactions
IGRP protocol debugging is on
RouterB#
1d19h: IGRP: sending update to 255.255.255.255 via Serial2 (10.1.1.2)
1d19h:      subnet 10.2.2.0, metric=88956
1d19h:      network 192.168.1.0, metric=89056
1d19h: IGRP: sending update to 255.255.255.255 via Serial3 (10.2.2.2)
1d19h:      subnet 10.1.1.0, metric=88956
1d19h:      network 172.16.0.0, metric=89056
1d19h: IGRP: received update from 10.1.1.1 on Serial2
1d19h:      network 172.16.0.0, metric 4294967295 (inaccessible)
1d19h: IGRP: edition is now 10
1d19h: IGRP: sending update to 255.255.255.255 via Serial2 (10.1.1.2)
1d19h:      subnet 10.2.2.0, metric=88956
1d19h:      network 172.16.0.0, metric=4293967295
1d19h:      network 192.168.1.0, metric=89056
1d19h: IGRP: sending update to 255.255.255.255 via Serial3 (10.2.2.2)
1d19h:      subnet 10.1.1.0, metric=88956
1d19h:      network 172.16.0.0, metric=4294967295
```

Example 3-11 Troubleshooting an IGRP Network 2 (Router B)

```
RouterB#show ip route
Codes:  C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX, - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        I - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
        U - per-user static route, o - ODR
        T - traffic engineered route

Gateway of last resort is not set

I 172.16.0.0/16 is possibly down, routing via 10.1.1.1, Serial2
 10.0.0.0/24 is subnetted, 2 subnets
C    10.1.1.0 is directly connected, Serial2
C    10.2.2.0 is directly connected, Serial3
I 192.168.1.0/24 [100/89506] via 10.2.2.3, 00:00:14, Serial3
```

In addition to sending updates, router B places the route to network 172.16.0.0 in the hold-down state for 280 seconds. While in the hold-down state, the route to network 172.16.0.0 is marked as "possibly down" in the routing table, as Example 3-11 shows. Router B still tries to send traffic to network 172.16.0.0 until the hold-down timer expires.

In Example 3-12, a network engineer unsuccessfully attempts to ping 172.16.1.1.

Example 3-12 Troubleshooting an IGRP Network 3 (Router B)

```
RouterB#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
RouterB#
```

If the Ethernet link on router A comes back up, router A sends another triggered update to router B stating that network 172.16.0.0 is now accessible (with metric 89056), as Example 3-13 shows. Router B receives the triggered update.

Example 3-13 Troubleshooting an IGRP Network 4 (Router B)

```
RouterB#debug ip igrp transactions
RouterB#
1d20h: IGRP: received update from 10.1.1.1 on Serial2
1d20h:      network 172.16.0.0, metric 89056 (neighbor 1100)
RouterB#
```

Although router B receives the update, router B keeps the route in the hold-down state. Router B does not remove the route from the hold-down state and update its routing table until the hold-down timer expires.

In Example 3-14, the hold-down timer has not yet expired, so the route is still "possibly down."

Example 3-14 Troubleshooting an IGRP Network 5 (Router B)

```
RouterB#show ip route
Codes:  C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX, - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        I - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
        U - per-user static route, o - ODR
        T - traffic engineered route

Gateway of last resort is not set

I 172.16.0.0/16 is possibly down, routing via 10.1.1.1, Serial2
 10.0.0.0/24 is subnetted, 2 subnets
C    10.1.1.0 is directly connected, Serial2
C    10.2.2.0 is directly connected, Serial3
I 192.168.1.0/24 [100/89506] via 10.2.2.3, 00:00:14, Serial3
```

However, the administrator at router B can now successfully ping network 172.16.0.0, as Example 3-15 shows.

Example 3-15 Troubleshooting an IGRP Network 6 (Router B)

```
RouterB#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5)
```

Chapter Summary

EIGRP is an IGP that scales well and provides quick convergence times with minimal network traffic. EIGRP is an enhanced version of IGRP, which was developed by Cisco, but EIGRP has improved convergence properties and operating efficiency over IGRP. New versions of the IOS no longer support IGRP.

Although IGRP and EIGRP are compatible with each other, there are some differences. EIGRP offers multiprotocol support, but IGRP does not. The EIGRP metric is the same as the IGRP metric except for a multiplier of 256 (which makes the EIGRP metrics larger).

EIGRP routers keep route and topology information readily available in RAM. Like OSPF, EIGRP saves this information in three tables. The neighbor table lists adjacent routers, the topology table is made up of all the EIGRP routes in the AS, and the routing table holds the best routes to a destination. DUAL, which is the EIGRP distance vector algorithm, takes the information supplied in the neighbor and the topology tables and calculates the lowest cost routes to each destination. The preferred primary route is called the successor route, and the backup route is called the feasible successor.

EIGRP is a balanced hybrid routing protocol (also referred to as an advanced distance vector routing protocol) and acts as a link-state protocol when updating neighbors and maintaining routing information. Advantages include rapid convergence, efficient use of bandwidth, support for VLSM and CIDR, support for multiple network layers, and independence from routed protocols.

DUAL results in the fast convergence of EIGRP. Each router has constructed a topology table that contains information about how to route to specific destinations. Each topology table identifies the routing protocol, or EIGRP; the lowest cost of the route, which is called feasible distance; and the cost of the route as advertised by the neighboring router, called reported distance.

EIGRP configuration commands vary depending on which protocol is used. Some examples of these protocols are IP, IPX, and AppleTalk. The **network** command configures only connected networks. EIGRP automatically summarizes routes at the classful boundary only prior to Cisco IOS Release 12.2(8)T. If discontiguous subnetworks exist, autosummarization must be disabled for routing to work properly. Manual summarization is done at the interface level with the **ip summary-address eigrp** command. The **show ip eigrp** command can verify an EIGRP configuration. The **debug ip eigrp** command can display information on EIGRP packets and troubleshoot EIGRP.

Troubleshooting at Layer 3 can be approached in a systematic fashion by using an eight-step troubleshooting methodology. Network engineers rely on **show** and **debug** commands to troubleshoot routing protocols. RIP, IGRP, EIGRP, and OSPF have their own set of **debug** commands that are tailored for culling important information that is used to troubleshoot issues with the respective routing protocol.

Check Your Understanding

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. Answers are listed in Appendix A, "Answers to Check Your Understanding and Challenge Questions and Activities."

1. In this output line from the **debug ip rip** command, what do the numbers within the parentheses signify?

```
RIP: sending v1 update to 255.255.255.255 via Ethernet1 (10.1.1.2)
```

   A. Source address
   B. Next-hop address
   C. Destination address
   D. Address of the routing table entry

2. What could cause the message "RIP: bad version 128 from 160.89.80.43" to display in the output of the **debug ip rip** command?

   A. Receiving a malformed packet
   B. Sending a routing table update
   C. Receiving a routing table update

3. Which command displays metric information that is contained in an IGRP update?

   A. **debug ip igrp events**
   B. **debug ip igrp transactions**
   C. **debug ip igrp events summary**
   D. **debug ip igrp transactions summary**

4. How is the bandwidth requirement for EIGRP packets minimized?

   A. By propagating only data packets
   B. By propagating only hello packets
   C. By propagating only routing table changes and hello packets
   D. By propagating the entire routing table to only those routers affected by a topology change

5. Which command correctly specifies that network 10.0.0.0 is directly connected to a router that runs EIGRP?

   A. Router(config)#**network 10.0.0.0**
   B. Router(config)#**router eigrp 10.0.0.0**
   C. Router(config-router)#**network 10.0.0.0**
   D. Router(config-router)#**router eigrp 10.0.0.0**

6. Which command displays the amount of time since the router heard from an EIGRP neighbor?

A. **show ip eigrp traffic**
B. **show ip eigrp topology**
C. **show ip eigrp interfaces**
D. **show ip eigrp neighbors**

7. The output from which command includes information about the length of the OSPF packet?

A. **debug ip ospf events**
B. **debug ip ospf packet**
C. **debug ip ospf packet size**
D. **debug ip ospf mpls traffic-eng advertisements**

8. What command(s) advertises the summary route 172.16.0.0/12 in EIGRP AS 1 out of interface Serial 0/0?

A. Router(config)**#ip summary-address 172.16.0.0 255.240.0.0 eigrp 1 serial0/0**
B. Router(config)**#interface serial0/0**
   Router(config-if)**#ip summary-address eigrp 1 172.16.0.0 255.240.0.0**
C. Router(config)**#ip summary-address 172.16.0.0 255.255.0.0 eigrp 1 serial0/0**
D. Router(config)**#interface serial0/0**
   Router(config-if)**#ip summary-address 172.16.0.0 255.240.0.0 eigrp 1 serial0/0**

9. What are the five EIGRP packet types?

A. Reply, query, hello, update, acknowledgment
B. Reply, query, hello, acknowledgment, LSU
C. Query, hello, acknowledgment, LSA, LSU
D. Reply, query, hello, RTP, acknowledgment

10. Which command displays the active or passive state of routes?

A. **show ip eigrp traffic**
B. **show ip eigrp topology**
C. **show ip eigrp interfaces**
D. **show ip eigrp neighbors**

Challenge Questions and Activities

These questions and activities are purposefully designed to be similar to the more complex styles of questions you might see on the CCNA exam. Answers are listed in Appendix A.

1. In Figure 3-10, two routers are configured to use EIGRP. Packets are not being forwarded between the two routers. What could be the problem?

Figure 3-10 EIGRP and VLSM

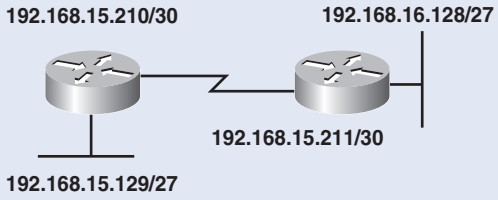

A. EIGRP does not support VLSM.
B. The routers were not configured to monitor neighbor adjacency changes.
C. The default bandwidth was used on the routers.
D. An incorrect IP address was configured on a router interface.

2. In Figure 3-11, routers A and B have EIGRP configured and automatic summarization has been disabled on both routers. Which one of the following router commands summarizes the attached routes and to which interface is the command applied? (Choose two.)

Figure 3-11 Interface Summarization with EIGRP

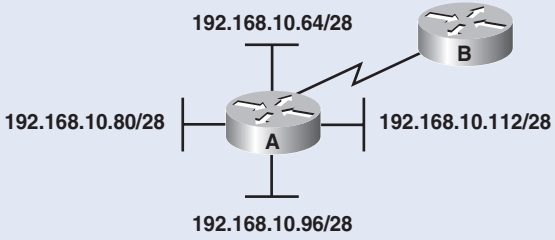

A. **ip summary-address eigrp 1 192.168.10.64 255.255.255.192**
B. **ip area-range eigrp 1 192.168.10.80 255.255.255.224**
C. **summary-address 192.168.10.80 0.0.0.31**
D. **ip summary-address eigrp 1 192.168.10.64 0.0.0.63**
E. Serial interface on Router A
F. Serial interface on Router B

3. When EIGRP is configured on a router, which table of DUAL information calculates the best route to each designated router?

A. Router table

B. Topology table

C. DUAL table

D. CAM table

E. ARP table