CHAPTER 3

Configuring the Enhanced Interior Gateway Routing Protocol

In present-day and future routing environments, Enhanced Interior Gateway Routing Protocol (EIGRP) offers benefits and features over historic distance vector routing protocols, such as Routing Information Protocol Version 1 (RIPv1) and Interior Gateway Routing Protocol (IGRP). These benefits include rapid convergence, lower bandwidth utilization, and multiple-routed protocol support.

NOTE As of Cisco IOS Software Release 12.3, IGRP is no longer supported.

This chapter introduces EIGRP terminology and concepts and EIGRP configuration, verification, and troubleshooting. The chapter also explores topics such as route summarization, load balancing, bandwidth usage, and authentication. The chapter concludes with a discussion of EIGRP design and configuration techniques to implement an effective enterprise network.

EIGRP Overview

This section introduces EIGRP and describes its four underlying technologies.

EIGRP Capabilities and Attributes

EIGRP is a Cisco-proprietary protocol that combines the advantages of link-state and distance vector routing protocols. EIGRP has its roots as a distance vector routing protocol and is predictable in its behavior. Like its predecessor IGRP, EIGRP is easy to configure and is adaptable to a wide variety of network topologies. What makes EIGRP an *advanced* distance vector protocol is the addition of several link-state features, such as dynamic neighbor discovery. EIGRP is an *enhanced* IGRP because of its rapid convergence and the guarantee of a loop-free topology at all times. Features of this hybrid protocol include the following:

- **Fast convergence**—EIGRP uses the Diffusing Update Algorithm (DUAL) to achieve rapid convergence. A router running EIGRP stores its neighbors' routing tables so that it can quickly adapt to changes in the network. If no appropriate route or backup route exists in the local routing table, EIGRP queries its neighbors to discover an alternative route. These queries are propagated until an alternative route is found, or it is determined that no alternative route exists.

- **Variable-length subnet masking (VLSM) support**—EIGRP is a classless routing protocol, which means that it advertises a subnet mask for each destination network; this enables EIGRP to support discontinuous subnetworks and VLSM.
- **Partial updates**—EIGRP sends partial triggered updates instead of periodic updates. These updates are sent only when the path or the metric for a route changes; they contain information about only that changed link rather than the entire routing table. Propagation of these partial updates is automatically bounded so that only those routers that require the information are updated. As a result, EIGRP consumes significantly less bandwidth than IGRP. This behavior is also different than link-state protocol operation, which sends a change update to *all* routers within an area.
- **Multiple network layer support**—EIGRP supports IP, AppleTalk, and Novell NetWare Internetwork Packet Exchange (IPX) using protocol-dependent modules that are responsible for protocol requirements specific to the network layer. EIGRP's rapid convergence and sophisticated metric offer superior performance and stability when implemented in IP, IPX, and AppleTalk networks.

NOTE Only the IP implementation of EIGRP is thoroughly covered in this book. Refer to the Cisco IOS technical documentation at http://www.cisco.com for information about how EIGRP operates, and how to configure it, for AppleTalk and IPX.

Other EIGRP features include the following:

- **Seamless connectivity across all data link layer protocols and topologies**—EIGRP does not require special configuration to work across any Layer 2 protocols. Other routing protocols, such as Open Shortest Path First (OSPF), require different configurations for different Layer 2 protocols, such as Ethernet and Frame Relay (as you will see in Chapter 4, "Configuring the Open Shortest Path First Protocol"). EIGRP was designed to operate effectively in both local-area network (LAN) and wide-area network (WAN) environments. In multiaccess topologies, such as Ethernet, neighbor relationships are formed and maintained using reliable multicasting. EIGRP supports all WAN topologies: dedicated links, point-to-point links, and nonbroadcast multiaccess (NBMA) topologies. EIGRP accommodates differences in media types and speeds when neighbor adjacencies form across WAN links. The amount of bandwidth that EIGRP uses on WAN links can be limited.
- **Sophisticated metric**—EIGRP uses the same algorithm for metric calculation as IGRP, but represents values in a 32-bit format, rather than IGRP's 24-bit format, to give additional granularity (thus, the EIGRP metric is the IGRP metric multiplied by 256). A significant advantage of EIGRP (and IGRP) over other protocols is its support for unequal metric load balancing that allows administrators to better distribute traffic flow in their networks.

- **Use of multicast and unicast**—EIGRP uses multicast and unicast for communication between routers, rather than broadcast. As a result, end stations are unaffected by routing updates or queries. The multicast address used for EIGRP is 224.0.0.10.

Like most IP routing protocols, EIGRP relies on IP packets to deliver routing information (Integrated Intermediate System-to-Intermediate System [IS-IS] is the exception, as you will see in Chapter 6, "Configuring the Integrated Intermediate System-to-Intermediate System Protocol"). The EIGRP routing process is a transport layer function of the Open System Interconnection (OSI) reference model. IP packets carrying EIGRP information have protocol number 88 in their IP header, as illustrated in Figure 3-1.

Figure 3-1 *EIGRP Is a Transport Layer Function*

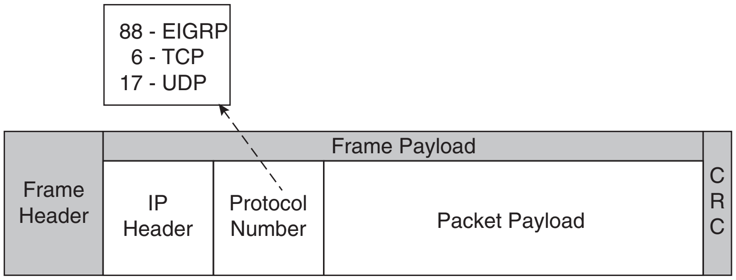

Figure 3-2 illustrates how EIGRP performs automatic route summarization at major network boundaries. Administrators can also configure manual summarization on arbitrary bit boundaries on any router interface (as long as a more-specific route exists in the routing table) to shrink the size of the routing table. EIGRP also supports the creation of supernets or aggregated blocks of addresses (networks).

Figure 3-2 *EIGRP Performs Route Summarization by Default*

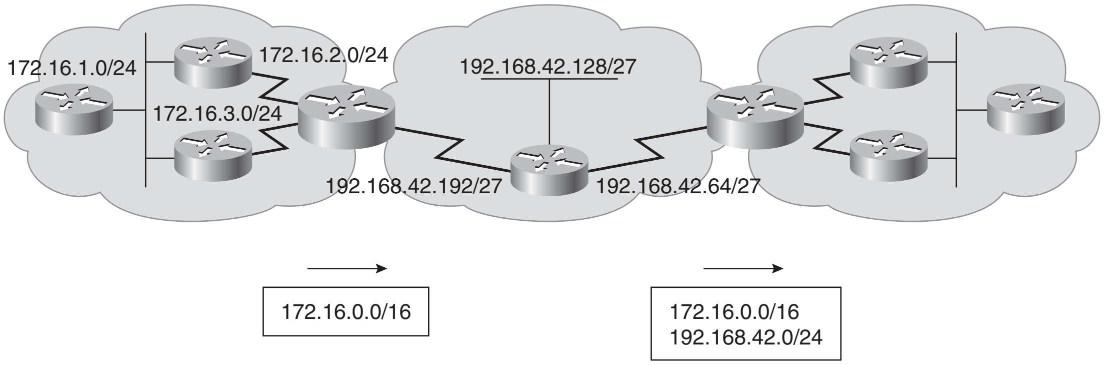

EIGRP supports both hierarchical and nonhierarchical IP addressing.

Underlying Processes and Technologies

EIGRP uses the following four key technologies that combine to differentiate it from other routing technologies:

- **Neighbor discovery/recovery mechanism**—EIGRP's neighbor discovery mechanism enables routers to dynamically learn about other routers on their directly attached networks. Routers also must discover when their neighbors become unreachable or inoperative. This process is achieved with low overhead by periodically sending small hello packets. As long as a router receives hello packets from a neighboring router, it assumes that the neighbor is functioning, and the two can exchange routing information.

- **Reliable Transport Protocol (RTP)**—RTP is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. RTP supports intermixed transmission of multicast or unicast packets. For efficiency, only certain EIGRP packets are transmitted reliably.

  For example, on a multiaccess network that has multicast capabilities, such as Ethernet, it is not necessary to send hello packets reliably to all neighbors individually, so EIGRP sends a single multicast hello packet containing an indicator that informs the receivers that the packet need not be acknowledged. Other types of packets, such as updates, indicate in the packet that acknowledgment is required. RTP contains a provision for sending multicast packets quickly even when unacknowledged packets are pending, which helps ensure that convergence time remains low in the presence of varying speed links.

- **DUAL finite-state machine**—DUAL embodies the decision process for all route computations. DUAL tracks all routes advertised by all neighbors and uses distance information, known as a *metric* or *cost*, to select efficient, loop-free paths to all destinations.

- **Protocol-dependent modules**—EIGRP’s protocol-dependent modules are responsible for network layer protocol-specific requirements. EIGRP supports IP, AppleTalk, and Novell NetWare; each protocol has its own EIGRP module and operates independently from any of the others that might be running. The IP-EIGRP module, for example, is responsible for sending and receiving EIGRP packets that are encapsulated in IP. Likewise, IP-EIGRP is also responsible for parsing EIGRP packets and informing DUAL of the new information that has been received. IP-EIGRP asks DUAL to make routing decisions, the results of which are stored in the IP routing table. IP-EIGRP is also responsible for redistributing routes learned by other IP routing protocols.

EIGRP Terminology and Operation

EIGRP sends out five different types of packets—hello, update, query, reply, and acknowledge (ACK)—that are used to establish the initial adjacency between neighbors and to keep the topology and routing tables current. When troubleshooting an EIGRP network, network administrators must understand what EIGRP packets are used for and how they are exchanged. For example, if routers running EIGRP do not form neighbor relationships, those routers cannot exchange EIGRP updates with each other. Without EIGRP routing updates, users cannot connect to services across the internetwork. This section explains EIGRP terminology, followed by an explanation of the mechanisms for creating the various EIGRP tables and a discussion about the five types of EIGRP packets. This section also explores how EIGRP routers become neighbors, initial route discovery, route selection, and how the DUAL algorithm functions.

EIGRP Terminology

The following terms are related to EIGRP and are used throughout the rest of this chapter:

- **Neighbor table**—EIGRP routers use hello packets to discover neighbors. When a router discovers and forms an adjacency with a new neighbor, it includes the neighbor's address and the interface through which it can be reached in an entry in the neighbor table. This table is comparable to the neighborship (adjacency) database used by link-state routing protocols (as described in Chapter 4). It serves the same purpose—ensuring bidirectional communication between each of the directly connected neighbors. EIGRP keeps a neighbor table for each network protocol supported; in other words, the following tables could exist: an IP neighbor table, an IPX neighbor table, and an AppleTalk neighbor table.

- **Topology table**—When the router dynamically discovers a new neighbor, it sends an update about the routes it knows to its new neighbor and receives the same from the new neighbor. These updates populate the topology table. The topology table contains all destinations advertised by neighboring routers; in other words, each router stores its neighbors' routing tables in its EIGRP topology table. If a neighbor is advertising a destination, it must be using that route to forward packets; this rule must be strictly followed by all distance vector protocols. An EIGRP router maintains a topology table for each network protocol configured (IP, IPX, and AppleTalk).

- **Advertised distance (AD) and feasible distance (FD)**—DUAL uses distance information, known as a *metric* or *cost*, to select efficient, loop-free paths. The lowest-cost route is calculated by adding the cost between the next-hop router and the destination—referred to as the *advertised distance*—to the cost between the local router and the next-hop router. The sum of these costs is referred to as the *feasible distance*.

- **Successor**—A successor, also called a current successor, is a neighboring router that has a least-cost path to a destination (the lowest FD) that is guaranteed not to be part of a routing loop; successors are offered to the routing table to be used for forwarding packets. Multiple successors can exist if they have the same FD.

- **Routing table**—The routing table holds the best routes to each destination and is used for forwarding packets. Successor routes are offered to the routing table. As discussed in Chapter 2, "Routing Principles," if a router learns more than one route to exactly the same destination from different routing sources, it uses the administrative distance to determine which route to keep in the routing table. By default, up to 4 routes to the same destination with the same metric can be added to the routing table (recall that the router can be configured to accept up to 16 per destination). The router maintains one routing table for each network protocol configured.

- **Feasible successor (FS)**—Along with keeping least-cost paths, DUAL keeps backup paths to each destination. The next-hop router for a backup path is called the feasible successor. To qualify as a feasible successor, a next-hop router must have an AD less than the FD of the current successor route; in other words, a feasible successor is a neighbor that is closer to the destination, but it is not the least-cost path and, thus, is not used to forward data. Feasible successors are selected at the same time as successors but are kept only in the topology table. The topology table can maintain multiple feasible successors for a destination.

  If the route via the successor becomes invalid (because of a topology change) or if a neighbor changes the metric, DUAL checks for feasible successors to the destination. If a feasible successor is found, DUAL uses it, thereby avoiding recomputing the route. If no suitable feasible successor exists, a recomputation must occur to determine the new successor. Although recomputation is not processor-intensive, it does affect convergence time, so it is advantageous to avoid unnecessary recomputations.

Populating EIGRP Tables

Figure 3-3 illustrates the three tables that EIGRP uses in its operation:

- The neighbor table lists adjacent routers
- The topology table lists all the learned routes to each destination
- The routing table contains the best route (the successor route) to each destination.

Figure 3-3 *EIGRP Maintains a Neighbor Table, a Topology Table, and a Routing Table*

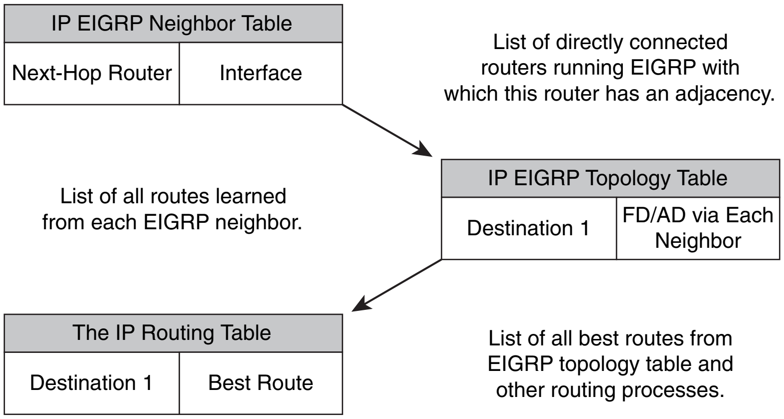
IP EIGRP Neighbor Table
Next-Hop Router
Interface
List of directly connected routers running EIGRP with which this router has an adjacency.
IP EIGRP Topology Table
Destination 1
FD/AD via Each Neighbor
List of all routes learned from each EIGRP neighbor.
The IP Routing Table
Destination 1
Best Route
List of all best routes from EIGRP topology table and other routing processes.

The neighbor table includes the address of each neighbor and the interface through which it can be reached.

The neighbor-table entry also includes information required by RTP. Sequence numbers are employed to match acknowledgments with data packets, and the last sequence number received from the neighbor is recorded, to detect out-of-order packets. A transmission list is used to queue packets for possible retransmission on a per-neighbor basis. Round-trip timers are kept in the neighbor-table entry to estimate an optimal retransmission interval.

Each router forwards a copy of its IP routing table to all its adjacent EIGRP neighbors, as specified in its EIGRP neighbor table. Each router then stores the routing tables of the adjacent neighbors in its EIGRP topology table (database). The topology table also maintains the metric that each neighbor advertises for each destination (the AD) and the metric that this router would use to reach the destination via that neighbor (the FD). The **show ip eigrp topology all-links** command displays all the IP entries in the topology table, while the **show ip eigrp topology** command displays only the successor(s) and feasible successor(s) for IP routes.

The topology table is updated when a directly connected route or interface changes or when a neighboring router reports a change to a route.

A topology-table entry for a destination can exist in one of two states: active or passive.

KEY POINT

Passive Versus Active Routes

A route is considered *passive* when the router is not performing recomputation on that route. A route is *active* when it is undergoing recomputation (in other words, when it is looking for a new successor).

Note that *passive* is the operational, stable state.

If feasible successors are always available, a destination never has to go into the active state, thereby avoiding a recomputation.

A recomputation occurs when the current route to a destination, the successor, goes down and there are no feasible successors for the destination. The router initiates the recomputation by sending a query packet to each of its neighboring routers. If the neighboring router has a route for the destination, it will send a reply packet; if it does not have a route, it sends a query packet to its neighbors. In this case, the route is also in the active state in the neighboring router. While a destination is in the active state, a router cannot change the routing table information for the destination. After a router has received a reply from each neighboring router, the topology table entry for the destination returns to the passive state.

Each router then examines its EIGRP topology table and determines the best route and other feasible routes to every destination network. A router compares all FDs to reach a specific network and then selects the route with the lowest FD and places it in the IP routing table; this is the successor route. The FD for the chosen successor route becomes the EIGRP routing metric to reach that network in the routing table.

EIGRP Packets

EIGRP uses the following five types of packets:

- **Hello**—Hello packets are used for neighbor discovery. They are sent as multicasts and do not require an acknowledgment. (They carry an acknowledgment number of 0.)
- **Update—**Update packets contain route change information. An update is sent to communicate the routes that a particular router has used to converge; an update is sent only to affected routers. These updates are sent as multicasts when a new route is discovered, and when convergence is completed (when the route becomes passive). To synchronize topology tables, updates are sent as unicasts to neighbors during their EIGRP startup sequence. Updates are sent reliably.
- **Query—**When a router is performing route computation and does not have a feasible successor, it sends a query packet to its neighbors, asking if they have a successor to the destination. Queries are normally multicast but can be retransmitted as unicast packets in certain cases; they are sent reliably.
- **Reply—**A reply packet is sent in response to a query packet. Replies are unicast to the originator of the query and are sent reliably.
- **ACK—**The ACK is used to acknowledge updates, queries, and replies. ACK packets are unicast hello packets and contain a nonzero acknowledgment number. (Note that hello and ACK packets do not require acknowledgment.)

The hello packet is the first type exchanged by EIGRP routers. The following section provides details of the hello protocol and how hello packets are used. The details of how the other packet types are used are provided throughout the rest of the chapter.

EIGRP Hello Packets

Through the hello protocol, an EIGRP router dynamically discovers other EIGRP routers directly connected to it. The router sends hello packets out of interfaces configured for EIGRP using the EIGRP multicast address 224.0.0.10. When an EIGRP router receives a hello packet from a router belonging to the same autonomous system (AS), it establishes a neighbor relationship (adjacency).

> **NOTE** The term *autonomous system* as used by EIGRP (and OSPF) is not the same as a Border Gateway Protocol (BGP) autonomous system. For EIGRP, consider the autonomous system to be a group of routers all running the same protocol. You may have more than one EIGRP autonomous system (group) within your network, in which case you might want to redistribute (share) routes between them; redistribution is detailed in Chapter 7, "Manipulating Routing Updates."

The time interval of hello packets varies depending on the medium. Hello packets are released every 5 seconds on a LAN link such as Ethernet, Token Ring, and FDDI. The default interval is also set to 5 seconds for point-to-point links such as PPP, High-Level Data Link Control (HDLC), point-to-point Frame Relay, and Asynchronous Transfer Mode (ATM) subinterfaces, and for multipoint circuits with bandwidth greater than T1, including Integrated Digital Services Network (ISDN) Primary Rate Interface (PRI), ATM, and Frame Relay. Hello packets are sent out less frequently on lower-speed links, such as multipoint circuits with a bandwidth less than or equal to T1, including ISDN Basic Rate Interface (BRI), Frame Relay, ATM, and X.25. Hellos are generated at 60-second intervals on these types of interfaces.

KEY POINT

Hello Packets

By default, hello packets are sent every 60 seconds on T1 or slower multipoint interfaces and every 5 seconds on other serial interfaces and on LANs.

You can adjust the rate at which hello packets are sent, called the *hello interval*, on a per-interface basis with the **ip hello-interval eigrp** *as-number seconds* interface configuration command.

Hello packets include the hold time.

The hold-time interval is set by default to 3 times the hello interval. Therefore, the default hold-time value is 15 seconds on LAN and fast WAN interfaces and 180 seconds on slower WAN interfaces. You can adjust the hold time with the **ip hold-time eigrp** *as-number seconds* interface configuration command.

KEY POINT

Hold Time

The hold time is the amount of time a router considers a neighbor up without receiving a hello or some other EIGRP packet from that neighbor.

NOTE The hold time is not automatically adjusted after a hello interval change. If you change the hello interval, you must manually adjust the hold time to reflect the configured hello interval.

If a packet is not received before the expiration of the hold time, the neighbor adjacency is deleted, and all topology table entries learned from that neighbor are removed, as if the neighbor had sent an update stating that all the routes are unreachable. If the neighbor is a successor for any destination networks, those networks are removed from the routing table, and alternative paths, if available, are computed. This lets the routes quickly reconverge if an alternative feasible route is available.

EIGRP Neighbors

The possibility exists for two routers to become EIGRP neighbors even though the hello and hold time values do not match; this means that the hello interval and hold-time values can be set independently on different routers.

Secondary addresses can be applied to interfaces to solve particular addressing issues, although all routing overhead traffic is generated through the primary interface address. EIGRP will not build peer relationships over secondary addresses, because all EIGRP traffic uses the interface's primary address. To form an EIGRP adjacency, all neighbors use their primary address as the source IP address of their EIGRP packets. Adjacency between EIGRP routers takes place if the primary address of each neighbor is part of the same IP subnet. In addition, peer relationships are not formed if the neighbor resides in a different autonomous system or if the metric-calculation mechanism constants (the K values) are misaligned on that link. (K values are discussed in the "EIGRP Metric Calculation" section later in this chapter.)

Neighbor Table

An EIGRP router multicasts hello packets to discover neighbors; it forms an adjacency with these neighbors so that it can exchange route updates. Only adjacent routers exchange routing information. Each router builds a neighbor table from the hello packets it receives from adjacent EIGRP routers running the same network layer protocol. EIGRP maintains a neighbor table for each configured network-layer protocol. You can display the IP neighbor table with the **show ip eigrp neighbors** command, as shown in Example 3-1.

Example 3-1 *Sample Output for the* **show ip eigrp neighbors** *Command*

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address         Interface   Hold  Uptime    SRTT   RTO  Q  Seq
                                (sec)           (ms)       Cnt Num
0   192.168.1.102   Se0/0/1     10    00:07:22   10    2280  0  5
R1#
```

This table includes the following key elements:

- **H (handle)**—A number used internally by the Cisco IOS to track a neighbor.
- **Address**—The neighbor's network-layer address.
- **Interface**—The interface on this router through which the neighbor can be reached.
- **Hold Time**—The maximum time, in seconds, that the router waits to hear from the neighbor without receiving anything from a neighbor before considering the link unavailable. Originally, the expected packet was a hello packet, but in current Cisco IOS software releases, any EIGRP packets received after the first hello from that neighbor resets the timer.
- **Uptime**—The elapsed time, in hours, minutes, and seconds since the local router first heard from this neighbor.
- **Smooth Round Trip Timer (SRTT)**—The average number of milliseconds it takes for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet. This timer is used to determine the retransmit interval, also known as the retransmit timeout (RTO).

- **RTO**—The amount of time, in milliseconds, that the router waits for an acknowledgment before retransmitting a reliable packet from the retransmission queue to a neighbor.
- **Queue count**—The number of packets waiting in the queue to be sent out. If this value is constantly higher than 0, a congestion problem might exist. A 0 indicates that no EIGRP packets are in the queue.
- **Seq Num**—The sequence number of the last update, query, or reply packet that was received from this neighbor.

EIGRP Reliability

EIGRP's reliability mechanism ensures delivery of critical route information to neighboring routers. This information is required to allow EIGRP to maintain a loop-free topology. For efficiency, only certain EIGRP packets are transmitted reliably.

KEY POINT

Reliable Packets

All packets carrying routing information (update, query, and reply) are sent reliably (because they are not sent periodically), which means that a sequence number is assigned to each reliable packet and an explicit acknowledgment is required for that sequence number.

Recall that RTP is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. RTP supports an intermixed transmission of multicast and unicast packets.

RTP ensures that ongoing communication is maintained between neighboring routers. As such, a retransmission list is maintained for each neighbor. This list indicates packets not yet acknowledged by a neighbor within the RTO. It is used to track all the reliable packets that were sent but not acknowledged.

KEY POINT

RTO Timer

If the RTO expires before an ACK packet is received, the EIGRP process transmits another copy of the reliable packet, up to a maximum of 16 times or until the hold time expires.

The use of reliable multicast packets is efficient. However, a potential delay exists on multiaccess media where multiple neighbors reside. The next reliable multicast packet cannot be transmitted until all peers have acknowledged the previous multicast. If one or more peers are slow to respond, this adversely affects all peers by delaying the next transmission. RTP is designed to handle such exceptions: Neighbors that are slow to respond to multicasts have the unacknowledged multicast packets retransmitted as unicasts. This allows the reliable multicast operation to proceed without delaying communication with other peers, helping to ensure that convergence time remains low in the presence of variable-speed links.

The multicast flow timer determines how long to wait for an ACK packet before switching from multicast to unicast. The RTO determines how long to wait between the subsequent unicasts. The

EIGRP process for each neighbor calculates both the multicast flow timer and RTO, based on the SRTT. The formulas for the SRTT, RTO, and multicast flow timer are Cisco-proprietary. In a steady-state network where no routes are flapping, EIGRP waits the specified hold-time interval before it determines that an EIGRP neighbor adjacency is down. Therefore, by default, EIGRP waits up to 15 seconds on high-speed links and up to 180 seconds on low-speed, multipoint links. When EIGRP determines that a neighbor is down and the router cannot reestablish the adjacency, the routing table removes all networks that could be reached through that neighbor. The router attempts to find alternative routes to those networks so that convergence can occur. The 180-second hold time on low-speed links can seem excessive, but it accommodates the slowest-speed multipoint links, which are generally connected to less-critical remote sites. In some networks with mission-critical or time-sensitive applications (such as IP telephony), even on high-speed links, 15 seconds is too long. The point to remember is that other conditions can override the hold time and allow the network to converge quickly.

For example, if the network is unstable and routes are flapping elsewhere because a remote site is timing out on its adjacency, EIGRP hold timers begin counting down from 180 seconds. When the upstream site sends the remote site an update, and the remote site does not acknowledge the update, the upstream site attempts 16 times to retransmit the update. The retransmission occurs each time the RTO expires. After 16 retries, the router resets the neighbor relationship. This causes the network to converge faster than waiting for the hold time to expire.

Initial Route Discovery

EIGRP combines the process of discovering neighbors and learning routes. Figure 3-4 illustrates the initial route discovery process.

Figure 3-4 *Initial Route Discovery*

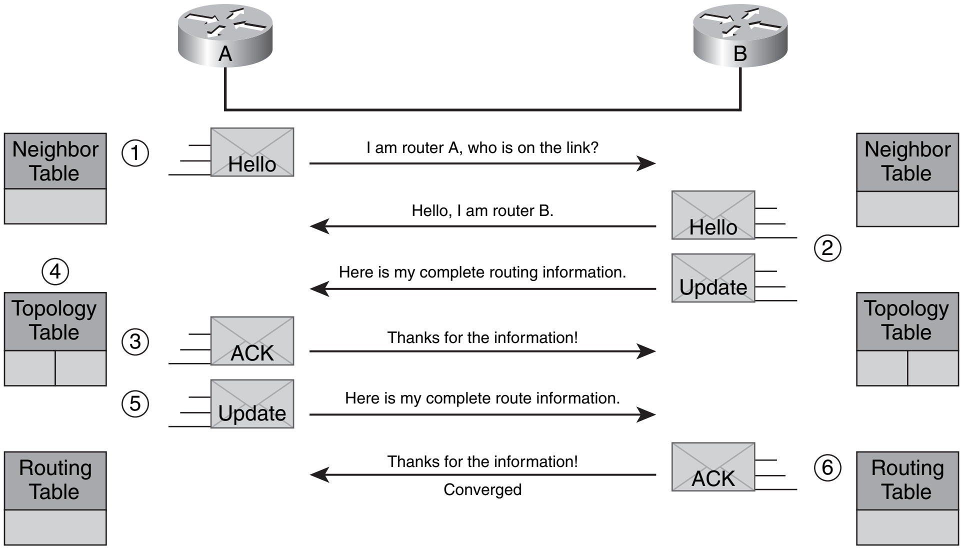
A
B
Neighbor Table
1
Hello
I am router A, who is on the link?
Neighbor Table
Hello, I am router B.
Hello
2
4
Here is my complete routing information.
Update
Topology Table
3
ACK
Thanks for the information!
Topology Table
5
Update
Here is my complete route information.
Routing Table
Thanks for the information!
Converged
ACK
6
Routing Table

The following describes the initial route discovery process:

1. A new router (Router A) comes up on the link and sends out a hello packet through all of its EIGRP-configured interfaces.
2. Routers receiving the hello packet on one interface (Router B in Figure 3-4) reply with update packets that contain all the routes they have in their routing table, except those learned through that interface (because of the split horizon rule). Router B sends an update packet to Router A, but a neighbor relationship is not established until Router B sends a hello packet to Router A. The update packet from Router B has the initial bit set, indicating that this is the initialization process. The update packet contains information about the routes that the neighbor (Router B) is aware of, including the metric that the neighbor is advertising for each destination.
3. After both routers have exchanged hellos and the neighbor adjacency is established, Router A replies to Router B with an ACK packet, indicating that it received the update information.
4. Router A inserts the update packet information in its topology table. The topology table includes all destinations advertised by neighboring (adjacent) routers. It is organized so that each destination is listed, along with all the neighbors that can get to the destination and their associated metrics.
5. Router A then sends an update packet to Router B.
6. Upon receiving the update packet, Router B sends an ACK packet to Router A.

After Router A and Router B successfully receive the update packets from each other, they are ready to chose the successor (best) and feasible successor (backup) routes in the topology table, and offer the successor routes to the routing table.

Split Horizon

Split horizon controls the sending of IP EIGRP update and query packets. When split horizon is enabled on an interface, no update or query packets for destinations for which this interface is the next-hop are sent out of this interface. This reduces the possibility of routing loops. By default, split horizon is enabled on all interfaces.

Split horizon blocks information about a destination from being advertised by a router out of any interface that the router uses to route to that destination. This behavior usually optimizes communications among multiple routers, particularly when links are broken.

When a router changes its topology table in such a way that the interface through which the router reaches a network changes, it turns off split horizon and poison reverses the old route out of all interfaces indicating that the route is unreachable. This ensures that other routers will not try to use the now invalid route.

Route Selection

The EIGRP route selection process is perhaps what most distinguishes it from other routing protocols. EIGRP selects primary (successor) and backup (feasible successor) routes and injects those into the topology table. The primary (successor) routes are then moved to the routing table.

EIGRP supports several types of routes: internal, external, and summary. Internal routes originate within the EIGRP autonomous system. External routes are learned from another routing protocol or another EIGRP autonomous system. Summary routes are routes encompassing multiple subnets.

EIGRP uses DUAL to calculate the best route to a destination. DUAL selects routes based on the composite metric and ensures that the selected routes are loop-free. DUAL also calculates backup routes (feasible successor routes) to a destination that are loop-free. If the best route fails, EIGRP immediately uses a backup route without any need for holddown, because the feasible successor route (if one exists) is loop-free; this results in fast convergence.

EIGRP Metric Calculation

The EIGRP metric calculation can use five variables, but EIGRP uses only two by default:

- **Bandwidth**—The smallest (slowest) bandwidth between the source and destination
- **Delay**—The cumulative interface delay along the path

The following criteria, although available, are not commonly used, because they typically result in frequent recalculation of the topology table:

- **Reliability**—The worst reliability between the source and destination, based on keepalives.
- **Loading**—The worst load on a link between the source and destination based on the packet rate and the interface's configured bandwidth.
- **Maximum transmission unit (MTU)**—The smallest MTU in the path. (MTU is included in the EIGRP update but is actually not used in the metric calculation.)

EIGRP calculates the metric by adding together weighted values of different variables of the path to the network in question. The default constant weight values are K1 = K3 = 1, and K2 = K4 = K5 = 0.

In EIGRP metric calculations, when K5 is 0 (the default), variables (bandwidth, bandwidth divided by load, and delay) are weighted with the constants K1, K2, and K3. The following is the formula used:

$$\text{metric} = (K1 * \text{bandwidth}) + [(K2 * \text{bandwidth}) / (256 - \text{load})] + (K3 * \text{delay})$$

If these K values are equal to their defaults, the formula becomes

$$\text{metric} = (1 * \text{bandwidth}) + [(0 * \text{bandwidth}) / (256 - \text{load})] + (1 * \text{delay})$$
$$\text{metric} = \text{bandwidth} + [0] + \text{delay}$$
$$\text{metric} = \text{bandwidth} + \text{delay}$$

If K5 is not equal to 0, the following additional operation is performed:

$$\text{metric} = \text{metric} * [K5 / (\text{reliability} + K4)]$$

K values are carried in EIGRP hello packets. Mismatched K values can cause a neighbor to be reset (only K1 and K3 are used, by default, in metric compilation). These K values should be modified only after careful planning; changing these values can prevent your network from converging and is generally not recommended.

KEY POINT

Delay and Bandwidth Values

The format of the delay and bandwidth values is different from those displayed by the **show interfaces** command.

The EIGRP delay value is the sum of the delays in the path, in tens of microseconds, multiplied by 256. The **show interfaces** command displays delay in microseconds.

The EIGRP bandwidth is calculated using the minimum bandwidth link along the path, represented in kilobits per second (kbps). 10^7 is divided by this value, and then the result is multiplied by 256.

EIGRP uses the same metric formula as IGRP, but EIGRP represents its metrics in a 32-bit format instead of the 24-bit representation used by IGRP. This representation allows a more granular decision to be made when determining the successor and feasible successor.

The EIGRP metric value ranges from 1 to 4,294,967,296. The IGRP metric value ranges from 1 to 16,777,216. EIGRP metrics are backward compatible with IGRP, as illustrated in Figure 3-5. When integrating IGRP routes into an EIGRP domain using redistribution, the router multiplies the IGRP metric by 256 to compute the EIGRP-equivalent metric. When sending EIGRP routes to an IGRP routing domain, the router divides each EIGRP metric by 256 to achieve the proper 24-bit metric.

Figure 3-5 *Initial Route Discovery*

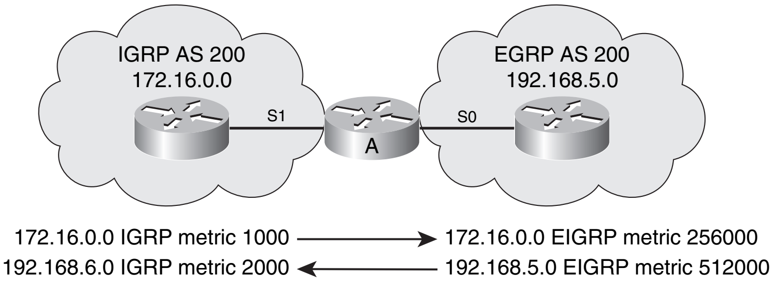
IGRP AS 200
172.16.0.0
EGRP AS 200
192.168.5.0
S1
S0
A
172.16.0.0 IGRP metric 1000
172.16.0.0 EIGRP metric 256000
192.168.6.0 IGRP metric 2000
192.168.5.0 EIGRP metric 512000

EIGRP Metric Calculation Example

In Figure 3-6, Router A has two paths to reach Router D (and thus any networks behind Router D). The bandwidths (in kbps) and the delays (in tens of microseconds) of the various links are also shown.

Figure 3-6 *EIGRP Metric Calculation Example*

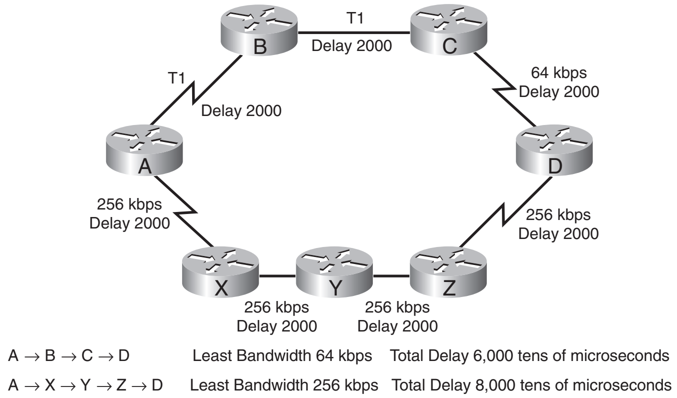
B
C
T1
Delay 2000
T1
Delay 2000
64 kbps
Delay 2000
A
D
256 kbps
Delay 2000
256 kbps
Delay 2000
X
Y
Z
256 kbps
Delay 2000
256 kbps
Delay 2000
A → B → C → D Least Bandwidth 64 kbps Total Delay 6,000 tens of microseconds
A → X → Y → Z → D Least Bandwidth 256 kbps Total Delay 8,000 tens of microseconds

The least bandwidth along the top path (A → B → C → D) is 64 kbps. The EIGRP bandwidth calculation for this path is as follows:

bandwidth = (10^7 / least bandwidth in kbps) * 256
bandwidth = (10,000,000 / 64) * 256 = 156,250 * 256 = 40,000,000

The delay through the top path is as follows:

delay = [(delay A → B) + (delay B → C) + (delay C → D)] * 256
delay = [2000 + 2000 + 2000] * 256
delay = 1,536,000

Therefore, the EIGRP metric calculation for the top path is as follows:

metric = bandwidth + delay
metric = 40,000,000 + 1,536,000
metric = 41,536,000

The least bandwidth along the lower path (A → X → Y → Z → D) is 256 kbps. The EIGRP bandwidth calculation for this path is as follows:

bandwidth = (10^7 / least bandwidth in kbps) * 256
bandwidth = (10,000,000 / 256) * 256 = 10,000,000

The delay through the lower path is as follows:

delay = [(delay A → X) + (delay X → Y) + (delay Y → Z) + (delay Z → D)] * 256
delay = [2000 + 2000 + 2000 + 2000] * 256
delay = 2,048,000

Therefore, the EIGRP metric calculation for the lower path is as follows:

metric = bandwidth + delay
metric = 10,000,000 + 2,048,000
metric = 12,048,000

Router A therefore chooses the lower path, with a metric of 12,048,000, over the top path, with a metric of 41,536,000. Router A installs the lower path with a next-hop router of X and a metric of 12,048,000 in the IP routing table.

The bottleneck along the top path, the 64-kbps link, can explain why the router takes the lower path. This slow link means that the rate of transfer to Router D would be at a maximum of 64 kbps. Along the lower path, the lowest speed is 256 kbps, making the throughput rate up to that speed. Therefore, the lower path represents a better choice, such as to move large files quickly.

Routing Table and EIGRP DUAL

DUAL is the finite-state machine that selects which information is stored in the topology and routing tables. As such, DUAL embodies the decision process for all EIGRP route computations. It tracks all routes advertised by all neighbors; uses the metric to select an efficient, loop-free path to each destination; and inserts that choice in the routing table.

Advertised Distance and Feasible Distance

KEY POINT

Advertised Distance Versus Feasible Distance

The AD is the EIGRP metric for an EIGRP *neighbor router* to reach a particular network. This is the metric between the next-hop neighbor router and the destination network.

The FD is the EIGRP metric for *this router* to reach a particular network. This is the sum of the AD for the particular network learned from an EIGRP neighbor, plus the EIGRP metric to reach that neighbor (the cost between this router and the next-hop router).

A router compares all FDs to reach a specific network in its topology table. The route with the lowest FD is placed in its IP routing table; this is the successor route. The FD for the chosen route becomes the EIGRP routing metric to reach that network in the routing table.

For example, in Figure 3-7, Routers A and B send their routing tables to Router C, whose tables are shown in the figure. Both Routers A and B have paths to network 10.1.1.0/24 (among many others that are not shown).

Figure 3-7 *EIGRP Chooses the Route with the Lowest Feasible Distance*

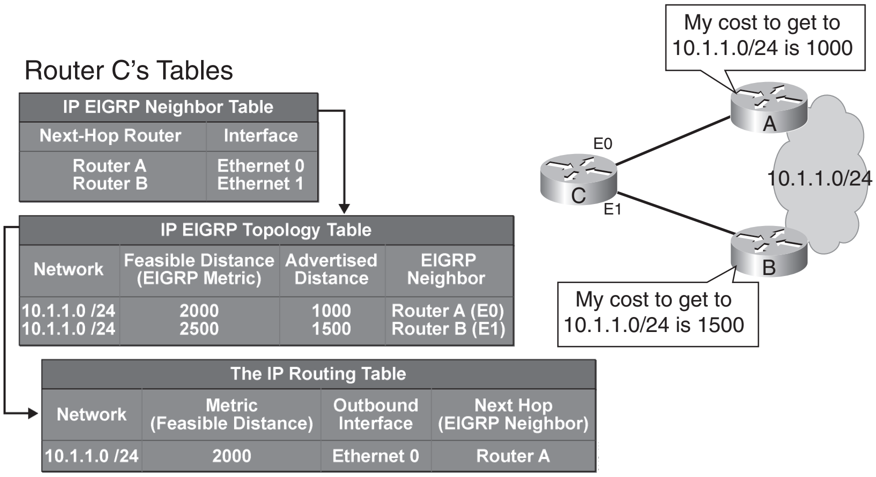

Router C's Tables

IP EIGRP Neighbor Table

Next-Hop Router	Interface
Router A	Ethernet 0
Router B	Ethernet 1

IP EIGRP Topology Table

Network	Feasible Distance (EIGRP Metric)	Advertised Distance	EIGRP Neighbor
10.1.1.0 /24	2000	1000	Router A (E0)
10.1.1.0 /24	2500	1500	Router B (E1)

The IP Routing Table

Network	Metric (Feasible Distance)	Outbound Interface	Next Hop (EIGRP Neighbor)
10.1.1.0 /24	2000	Ethernet 0	Router A

My cost to get to 10.1.1.0/24 is 1000

My cost to get to 10.1.1.0/24 is 1500

The routing table on Router A has an EIGRP metric of 1000 for 10.1.1.0/24. Therefore, Router A advertises 10.1.1.0/24 to Router C with a metric of 1000. Router C installs 10.1.1.0/24 from Router A in its EIGRP topology table with an AD of 1000. Router B has network 10.1.1.0/24 with a metric of 1500 in its IP routing table. Therefore, Router B advertises 10.1.1.0/24 to Router C with an AD of 1500. Router C places the 10.1.1.0/24 network from Router B in the EIGRP topology table with an AD of 1500.

Router C in Figure 3-7 has two entries to reach 10.1.1.0/24 in its topology table. The EIGRP metric for Router C to reach either Router A or B is 1000. This cost (1000) is added to the respective AD from each router, and the results represent the FDs that Router C must travel to reach network 10.1.1.0/24. Router C chooses the least-cost FD (2000) and installs it in its IP routing table as the best route to reach 10.1.1.0/24. The EIGRP metric in the routing table is the best FD from the EIGRP topology table.

Successor and Feasible Successor

KEY POINT

Successor

A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop.

A router is chosen as a successor because it has the lowest FD of all possible paths to that destination network. The successor is the next router in line to reach that destination. In other words, it is the router with the best path to reach that destination network.

An EIGRP router selects the best path to reach a given network and then installs the destination network, the metric to reach that network, the outbound interface to reach the next-hop router, and the IP address of the next-hop router into the IP routing table. If the EIGRP topology table has many entries that have an equal-cost FD to a given destination network, all successors (up to four by default) for that destination network are installed in the routing table.

KEY POINT

FD and AD

Note that it is the FD, not the AD, that affects the selection of the best routes for incorporation in the routing table; the AD is used only to calculate the FD.

All routing protocols can install only the next-hop router information in the routing table; information about the subsequent routers in the path is not put in the routing table. Each router relies on the next-hop router to make a reliable decision to reach a specific destination network. The hop-by-hop path through a network goes from one router to the next. Each router makes a path selection to reach a given network and installs the best next-hop address along the path to reach that destination network. A router trusts a route's successor (the best next-hop router) to send traffic toward that destination address.

The routing table is essentially a subset of the topology table; the topology table contains more detailed information about each route, any backup routes, and information used exclusively by DUAL.

KEY POINT

Feasible Successor

A feasible successor is a router providing a backup route. The route through the feasible successor must be loop free; in other words, it must not loop back to the current successor.

FSs are selected at the same time the successors are identified. These FS routes are kept in the topology table; the topology table can retain multiple FS routes for a destination.

KEY POINT

Feasible Successor Requirements

An FS must be mathematically proven. To qualify as an FS, a next-hop router must have an AD less than the FD of the current successor route for the particular network.

This requirement ensures that the FS cannot use a route through the local router (which would be a routing loop), because the AD through the FS is less than the best route through the local router. For example, as shown in Figure 3-8, Router B is an FS, because the AD through Router B (1500) is less than the FD of the current successor, Router A (2000).

Figure 3-8 *Feasible Successor's AD Must Be Less Than the Successor's FD*

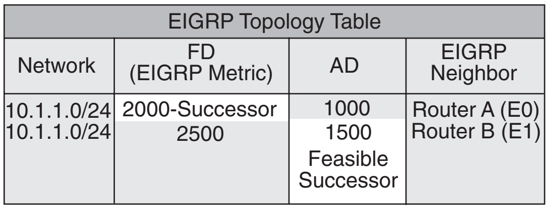

EIGRP Topology Table			
Network	FD (EIGRP Metric)	AD	EIGRP Neighbor
10.1.1.0/24	2000-Successor	1000	Router A (E0)
10.1.1.0/24	2500	1500 Feasible Successor	Router B (E1)

When a router loses a route, it looks at the topology table for an FS. If one is available, the route does not go into an active state; rather, the best FS is promoted as the successor and is installed in the routing table. The FS can be used immediately, without any recalculation. If there are no FSs, a route goes into active state, and route computation occurs. Through this process, a new successor is determined (if there is one). The amount of time it takes to recalculate the route affects the convergence time.

Figure 3-9 illustrates another example. Router C's initial topology table is shown at the top of the figure. Router B is the successor for network 10.1.1.0/24, and Router D is the FS.

Figure 3-9 *With a Feasible Successor, EIGRP Can Recover Immediately from Network Failures*

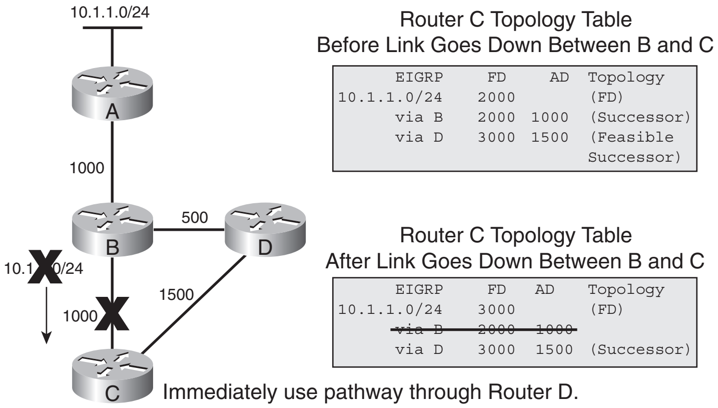

Router C Topology Table
Before Link Goes Down Between B and C

```
        EIGRP     FD     AD  Topology
10.1.1.0/24    2000          (FD)
        via B    2000  1000  (Successor)
        via D    3000  1500  (Feasible
                             Successor)
```

Router C Topology Table
After Link Goes Down Between B and C

```
        EIGRP     FD    AD    Topology
10.1.1.0/24    3000          (FD)
        ~~via B    2000  1000~~
        via D    3000  1500  (Successor)
```

Immediately use pathway through Router D.

In Figure 3-9, the link between Router B and Router C fails. Router C removes the route 10.1.1.0/24 through Router B from its routing table and searches the EIGRP topology table for an FS; Router D is an FS. Because Router D can still reach the network and does not send an update or query packet to inform Router C of the lost route, Router C immediately uses the path through Router D. Router C chooses this path because the AD through Router D (1500) is less than the FD of the best route, through Router B (2000); this path is guaranteed to be loop free.

DUAL Example

The mathematical formula to ensure that the FS is loop free requires that the AD of the backup route be *less than* the FD of the successor. When the AD of the second-best route is greater than or equal to the FD of the successor, an FS cannot be chosen. In this case, a discovery process that uses EIGRP queries and replies must be used to find any alternative paths to the lost networks.

The following example examines partial entries for network 10.1.1.0/24 in the topology tables for Routers C, D, and E in Figure 3-10, to give you a better understanding of EIGRP behavior. The partial topology tables shown in Figure 3-10 indicate the following:

- **AD**—The advertised distance is equal to the cost of the path to network 10.1.1.0/24 as advertised by neighboring routers.
- **FD**—The feasible distance is equal to the sum of the AD for a neighbor to reach 10.1.1.0/24, plus the metric to reach that neighbor.
- **Successor**—The successor is the forwarding path used to reach network 10.1.1.0/24. The cost of this path is equal to the FD.
- **FS**—The feasible successor is an alternative loop-free path to reach network 10.1.1.0/24.

Figure 3-10 *DUAL Example, Step 1*

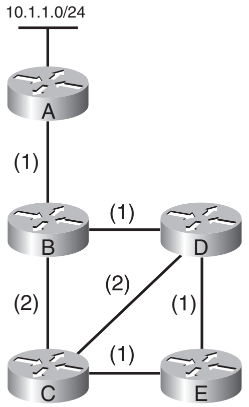

Router C

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   3
      via B   3    1  (Successor)
      via D   4    2  (FS)
      via E   4    3
```

Router D

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   2
      via B   2    1  (Successor)
      via C   5    3
```

Router E

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   3
      via D   3    2  (Successor)
      via C   4    3
```

The network shown in Figure 3-10 is stable and converged.

> **NOTE** As mentioned earlier, EIGRP implements split horizon. For example, Router E does not pass its route for network 10.1.1.0/24 to Router D, because Router E uses Router D as its next hop to network 10.1.1.0/24.

In Figure 3-11, Routers B and D detect a link failure. After being notified of the link failure, DUAL does the following, as shown in Figure 3-11:

- At Router D, it marks the path to network 10.1.1.0/24 through Router B as unusable.

Figure 3-11 *DUAL Example, Step 2*

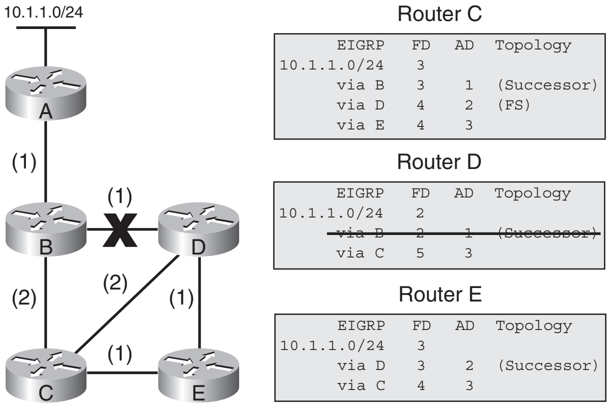

Router C

```
      EIGRP   FD   AD   Topology
10.1.1.0/24    3
      via B    3    1   (Successor)
      via D    4    2   (FS)
      via E    4    3
```

Router D

```
      EIGRP   FD   AD   Topology
10.1.1.0/24    2
~~      via B    2    1   (Successor)~~
      via C    5    3
```

Router E

```
      EIGRP   FD   AD   Topology
10.1.1.0/24    3
      via D    3    2   (Successor)
      via C    4    3
```

The following steps then occur, as shown in Figure 3-12:

- At Router D, there is no FS to network 10.1.1.0/24, because the AD via Router C (3) is greater than the FD via Router B (2). Therefore, DUAL does the following:

  — Sets the metric to network 10.1.1.0/24 as unreachable (–1 is unreachable).

  — Because an FS cannot be found in the topology table, the route changes from the passive state to the active state. In the active state, the router sends out queries to neighboring routers looking for a new successor.

  — Sends a query to Routers C and E for an alternative path to network 10.1.1.0/24.

  — Marks Routers C and E as having a query pending (q).

- At Router E, DUAL marks the path to network 10.1.1.0/24 through Router D as unusable.

- At Router C, DUAL marks the path to network 10.1.1.0/24 through Router D as unusable.

Figure 3-12 *DUAL Example, Step 3*

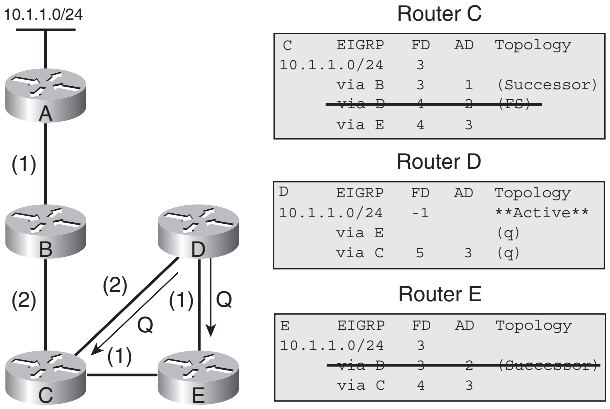

Router C

```
C     EIGRP   FD   AD   Topology
10.1.1.0/24    3
      via B    3    1   (Successor)
~~      via D    4    2   (FS)~~
      via E    4    3
```

Router D

```
D     EIGRP   FD   AD   Topology
10.1.1.0/24   -1        **Active**
      via E             (q)
      via C    5    3   (q)
```

Router E

```
E     EIGRP   FD   AD   Topology
10.1.1.0/24    3
~~      via D    3    2   (Successor)~~
      via C    4    3
```

The following steps then occur, as shown in Figure 3-13:

- At Router D:
  - — DUAL receives a reply from Router C that indicates no change to the path to network 10.1.1.0/24.
  - — DUAL removes the query flag from Router C.
  - — DUAL stays active on network 10.1.1.0/24, awaiting a reply from Router E to its query (q).
- At Router E, there is no FS to network 10.1.1.0/24, because the AD from Router C (3) is not less than the original FD (also 3).
  - — DUAL generates a query to Router C.
  - — DUAL marks Router C as query pending (q).
- At Router C, DUAL marks the path to network 10.1.1.0/24 through Router E as unusable.

Figure 3-13 *DUAL Example, Step 4*

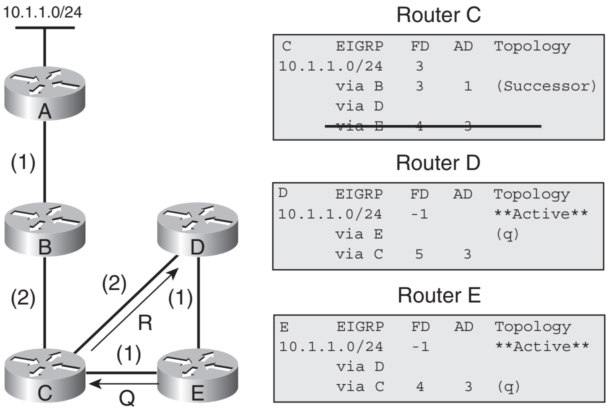

Router C

```
C      EIGRP   FD   AD  Topology
10.1.1.0/24    3
       via B   3     1  (Successor)
       via D
       ~~via E   4     3~~
```

Router D

```
D      EIGRP   FD   AD  Topology
10.1.1.0/24   -1        **Active**
       via E            (q)
       via C   5     3
```

Router E

```
E      EIGRP   FD   AD  Topology
10.1.1.0/24   -1        **Active**
       via D
       via C   4     3  (q)
```

The following steps then occur, as shown in Figure 3-14:

- At Router D, DUAL stays active on network 10.1.1.0/24, awaiting a reply from Router E (q).
- At Router E:
  - — DUAL receives a reply from Router C indicating no change.
  - — It removes the query flag from Router C.
  - — It calculates a new FD and installs a new successor route in the topology table.
  - — It changes the route to network 10.1.1.0/24 from active to passive (converged).

Figure 3-14 *DUAL Example, Step 5*

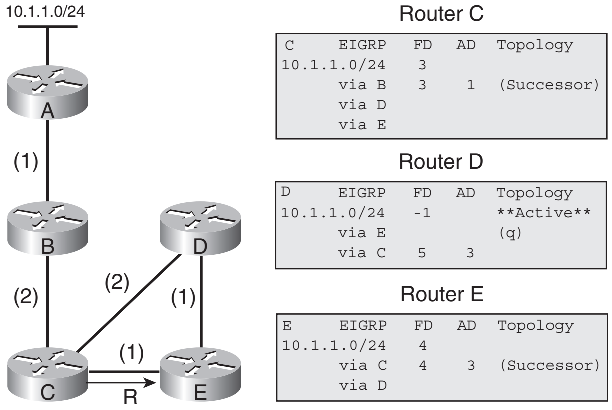

```
C      EIGRP   FD  AD  Topology
10.1.1.0/24    3
       via B   3    1  (Successor)
       via D
       via E
```

```
D      EIGRP   FD  AD  Topology
10.1.1.0/24   -1       **Active**
       via E           (q)
       via C   5    3
```

```
E      EIGRP   FD  AD  Topology
10.1.1.0/24    4
       via C   4    3  (Successor)
       via D
```

The following steps then occur, as shown in Figure 3-15:

- At Router D
  - — DUAL receives a reply from Router E.
  - — It removes the query flag from Router E.
  - — It calculates a new FD.
  - — It installs new successor routes in the topology table. Two routes (through Routers C and E) have the same FD, and both are marked as successors.
  - — It changes the route to network 10.1.1.0/24 from active to passive (converged).

Figure 3-15 *DUAL Example, Step 6*

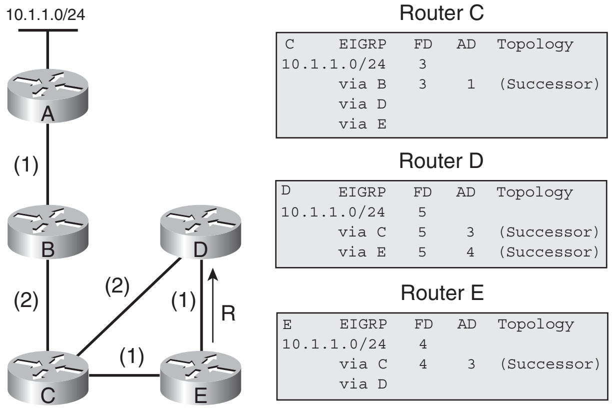

```
C      EIGRP   FD  AD  Topology
10.1.1.0/24    3
       via B   3    1  (Successor)
       via D
       via E
```

```
D      EIGRP   FD  AD  Topology
10.1.1.0/24    5
       via C   5    3  (Successor)
       via E   5    4  (Successor)
```

```
E      EIGRP   FD  AD  Topology
10.1.1.0/24    4
       via C   4    3  (Successor)
       via D
```

The following steps then occur, as shown in Figure 3-16:

- At Router D, two successor routes are in the topology table for network 10.1.1.0/24. Both successor routes are listed in the routing table, and equal-cost load balancing is in effect.

- The network is stable and converged.

Figure 3-16 *DUAL Example, Step 7*

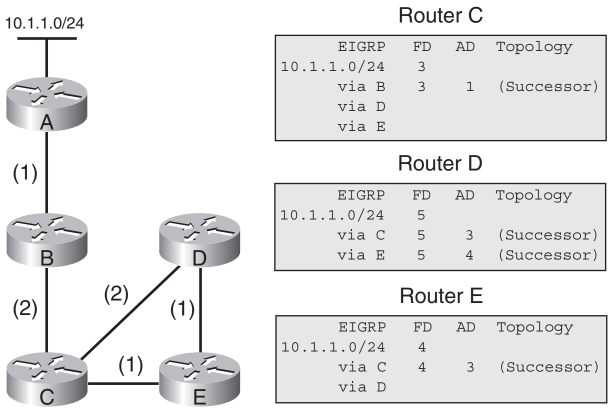

Router C

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   3
     via B   3    1  (Successor)
     via D
     via E
```

Router D

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   5
     via C   5    3  (Successor)
     via E   5    4  (Successor)
```

Router E

```
     EIGRP   FD   AD  Topology
10.1.1.0/24   4
     via C   4    3  (Successor)
     via D
```

Figure 3-10, the original topology before the link failure, shows traffic from Router E passing through Routers D and B. In Figure 3-16, the new topology shows traffic from Routers D and E going through Routers C and B. Notice that throughout the entire convergence process, routes to network 10.1.1.0/24 become active only on Routers D and E. The route to network 10.1.1.0/24 on Router C remains passive because the link failure between Routers B and D does not affect the successor route from Router C to network 10.1.1.0/24.

> **NOTE** When DUAL decides that a packet needs to be transmitted to a neighbor, the packets are not actually generated until the moment of transmission. Instead, the transmit queues contain small, fixed-size structures that indicate which parts of the topology table to include in the packet when it is finally transmitted. This means that the queues do not consume large amounts of memory. It also means that only the latest information is transmitted in each packet. If a route changes state several times, only the last state is transmitted in the packet, thus reducing link utilization.

Configuring and Verifying EIGRP

This section covers the commands used to configure EIGRP features. The following topics are discussed:

- Basic EIGRP configuration
- Configuring the **ip default-network** command for EIGRP
- Route summarization
- EIGRP load balancing
- EIGRP and WAN links

Basic EIGRP Configuration

Follow these steps to configure basic EIGRP for IP:

Step 1 Enable EIGRP and define the autonomous system using the **router eigrp** *autonomous-system-number* global configuration command. In this command, the *autonomous-system-number* identifies the autonomous system and is used to indicate all routers that belong within the internetwork. This value must match on all routers within the internetwork.

Step 2 Indicate which networks are part of the EIGRP autonomous system using the **network** *network-number* [*wildcard-mask*] router configuration command. Table 3-1 summarizes the parameters of this command.

Table 3-1 **network** *Command Parameters*

Parameter	Description
network-number	This parameter can be a network, a subnet, or the address of an interface. It determines which links on the router to advertise to, which links to listen to advertisements on, and which networks are advertised.
wildcard-mask	(Optional) An inverse mask used to determine how to interpret the *network-number*. The mask has wildcard bits, where 0 is a match and 1 is do not care. For example, 0.0.255.255 indicates a match in the first 2 octets.

If you do not use the optional wildcard mask, the EIGRP process assumes that all directly connected networks that are part of the major network will participate in the EIGRP routing process, and EIGRP will attempt to establish EIGRP neighbor relationships from each interface that is part of the overall Class A, B, or C network.

Use the optional wildcard mask to identify a specific IP address, subnet, or network. The router interprets the network number using the wildcard mask to determine which connected networks will participate in the EIGRP routing process. If you want to specify an interface address, use the mask 0.0.0.0 to match all 4 octets of the address. An address and wildcard mask combination of 0.0.0.0 255.255.255.255 matches all interfaces on the router.

Step 3 For serial links, define the link's bandwidth for the purposes of sending routing update traffic on the link. If you do not define the bandwidth value for these interfaces, EIGRP assumes that the bandwidth on the link is the default, which varies with interface type. Recall that EIGRP uses bandwidth as part of its metric calculation. If the link is actually slower than the default, the router might not be able to converge, or routing updates might become lost. The percent of the interface's bandwidth that EIGRP uses can also be limited, as described in the section "EIGRP and WAN Links" later in this chapter. To define the bandwidth, use the **bandwidth** *kilobits* interface configuration command. In this command, *kilobits* indicates the intended bandwidth in kbps.

For generic serial interfaces such as PPP and HDLC, set the bandwidth to match the line speed. For Frame Relay point-to-point interfaces, set the bandwidth to the committed information rate (CIR). For Frame Relay multipoint connections, set the bandwidth to the sum of all CIRs, or if the permanent virtual circuits (PVCs) have different CIRs, set the bandwidth to the lowest CIR multiplied by the number of PVCs on the multipoint connection.

Basic EIGRP Configuration Example

Figure 3-17 shows a sample network, including the configuration of Router A for EIGRP.

All routers in the network are part of autonomous system 109. (For EIGRP to establish a neighbor relationship, all neighbors must be in the same autonomous system.)

Figure 3-17 *Basic EIGRP Configuration Sample Network*

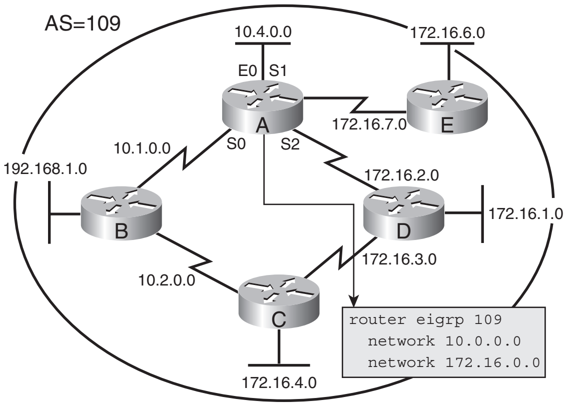

```
router eigrp 109
  network 10.0.0.0
  network 172.16.0.0
```

Because the wildcard mask is not used in Router A's configuration, all interfaces on Router A that are part of network 10.0.0.0/8 and network 172.16.0.0/16 participate in the EIGRP routing process. In this case, this includes all four interfaces. Note that network 192.168.1.0 is not configured in the EIGRP configuration on Router A, because Router A does not have any interfaces in that network.

Instead, suppose that the configuration in Example 3-2 was entered onto Router A.

Example 3-2 *Alternative Configuration of Router A in Figure 3-17*

```
routerA(config)#router eigrp 109
routerA(config-router)#network 10.1.0.0
routerA(config-router)#network 10.4.0.0
routerA(config-router)#network 172.16.7.0
routerA(config-router)#network 172.16.2.0
```

Because no wildcard mask was specified, Router A would automatically change the **network** commands to have classful networks, and the resulting configuration would be as shown in Example 3-3.

Example 3-3 *Router A's Interpretation of the Configuration in Example 3-2*

```
router eigrp 109
 network 10.0.0.0
network 172.16.0.0
```

Alternatively, consider what would happen if the configuration shown in Example 3-4 was entered for Router A.

Example 3-4 *Another Alternative Configuration of Router A in Figure 3-17*

```
routerA(config)#router eigrp 109
routerA(config-router)#network 10.1.0.0 0.0.255.255
routerA(config-router)#network 10.4.0.0 0.0.255.255
routerA(config-router)#network 172.16.2.0 0.0.0.255
routerA(config-router)#network 172.16.7.0 0.0.0.255
```

In this case, Router A uses the wildcard mask to determine which directly connected interfaces participate in the EIGRP routing process for autonomous system 109. All interfaces that are part of networks 10.1.0.0/16, 10.4.0.0/16, 172.16.2.0/24, and 172.16.7.0/24 participate in the EIGRP routing process for autonomous system 109; in this case, all four interfaces participate.

EIGRP Configuration Example Using the Wildcard Mask

Figure 3-18 shows another sample network that runs EIGRP in autonomous system 100. The configuration for Router C uses the wildcard mask, because Router C has subnets of Class B network 172.16.0.0 on all interfaces. Router C connects to a router external to autonomous system 100 on its serial interface, and the administrator does not want to run EIGRP with the same autonomous system number there. Without using the wildcard mask, Router C would send EIGRP packets to the external network. This would waste bandwidth and CPU cycles and would provide unnecessary information to the external network. The wildcard mask tells EIGRP to establish a relationship with EIGRP routers from interfaces that are part of network 172.16.3.0/24 or 172.16.4.0/24, but not 172.16.5.0/24.

Figure 3-18 *EIGRP Configuration with Wildcard Mask Example*

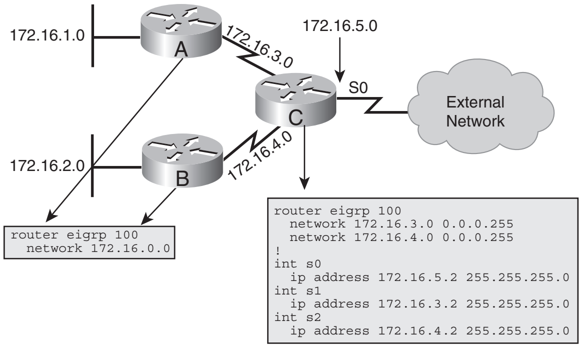

```
router eigrp 100
  network 172.16.0.0
```

```
router eigrp 100
  network 172.16.3.0 0.0.0.255
  network 172.16.4.0 0.0.0.255
!
int s0
  ip address 172.16.5.2 255.255.255.0
int s1
  ip address 172.16.3.2 255.255.255.0
int s2
  ip address 172.16.4.2 255.255.255.0
```

Configuring the ip default-network Command for EIGRP

The EIGRP default route can be created with the **ip default-network** *network-number* global configuration command. A router configured with this command considers the *network-number* the last-resort gateway that it will announce to other routers. The network must be reachable by the router that uses this command before it announces it as a candidate default route to other EIGRP routers. The network number in this command must also be passed to other EIGRP routers so that those routers can use this network as their default network and set their gateway of last resort to this default network. This means that the network must either be an EIGRP-derived network in the routing table, or be generated with a static route and redistributed into EIGRP.

Multiple default networks can be configured; downstream routers then use the EIGRP metric to determine the best default route.

For example, in Figure 3-19, Router A is directly attached to external network 172.31.0.0/16. Router A is configured with the 172.31.0.0 network as a candidate default network using the **ip default-network 172.31.0.0** command. Router A also has that network listed in a **network** command under the EIGRP process and, therefore, passes it to Router B. On Router B, the EIGRP-learned 172.31.0.0 network is flagged as a candidate default network (indicated by the * in the routing table). Router B also sets the gateway of last resort as10.5.1.1 (Router A) to reach the default network of 172.31.0.0.

> **NOTE** In earlier versions of the IOS software, the router on which the **ip default-network** command was configured would not set the gateway of last resort. Figure 3-19 illustrates that it now does set the gateway of last resort to 0.0.0.0, to the network specified in the **ip default-network** command.

> **NOTE** When you configure the **ip default-network** command and specify a subnet, a static route (the **ip route** command) is generated in the router's configuration; however, the IOS does not display a message to indicate that this has been done. The entry appears as a static route in the routing table of the router where the command is configured. This can be confusing when you want to remove the default network; the configuration must be removed with the **no ip route** command, not with the **no ip default-network** command.

Figure 3-19 *EIGRP* **ip default-network** *Sample Network*

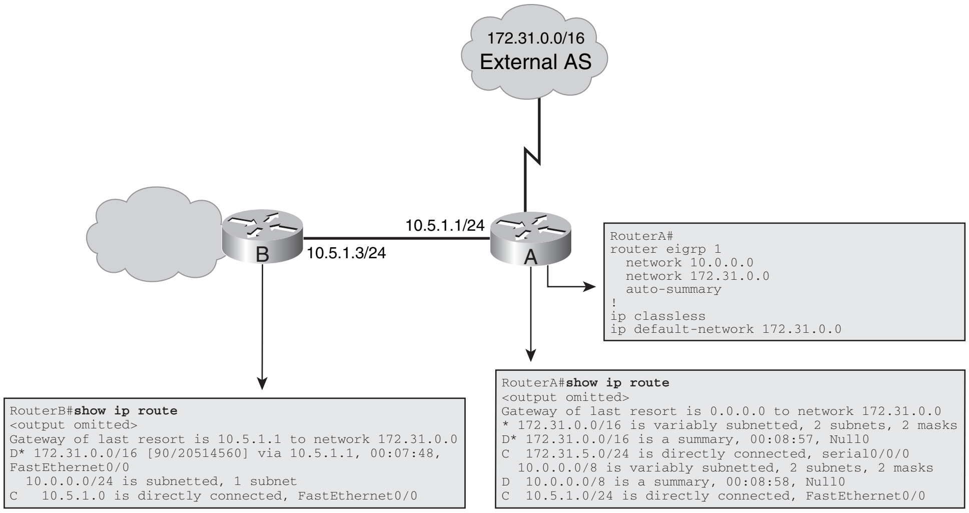

```
RouterA#
router eigrp 1
  network 10.0.0.0
  network 172.31.0.0
  auto-summary
!
ip classless
ip default-network 172.31.0.0
```

```
RouterB#show ip route
<output omitted>
Gateway of last resort is 10.5.1.1 to network 172.31.0.0
D* 172.31.0.0/16 [90/20514560] via 10.5.1.1, 00:07:48,
FastEthernet0/0
  10.0.0.0/24 is subnetted, 1 subnet
C   10.5.1.0 is directly connected, FastEthernet0/0
```

```
RouterA#show ip route
<output omitted>
Gateway of last resort is 0.0.0.0 to network 172.31.0.0
* 172.31.0.0/16 is variably subnetted, 2 subnets, 2 masks
D* 172.31.0.0/16 is a summary, 00:08:57, Null0
C  172.31.5.0/24 is directly connected, serial0/0/0
  10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D  10.0.0.0/8 is a summary, 00:08:58, Null0
C  10.5.1.0/24 is directly connected, FastEthernet0/0
```

> **NOTE** EIGRP (and IGRP) behave differently than RIP when using the **ip route 0.0.0.0 0.0.0.0** command. For example, EIGRP does not redistribute the 0.0.0.0 0.0.0.0 default route by default. However, if the **network 0.0.0.0** command is added to the EIGRP configuration, it redistributes a default route as a result of the **ip route 0.0.0.0 0.0.0.0** *interface* command (but not as a result of the **ip route 0.0.0.0 0.0.0.0** *address* or **ip default-network** command). For example, the partial configuration shown in Example 3-5 illustrates a router with the 0.0.0.0 route passed to the router's EIGRP neighbors.

Example 3-5 *EIGRP Passes a Default Route Only if It Is Configured to Do So*

```
Router#show run
<output omitted>
interface serial 0/0/0
  ip address 10.1.1.1 255.255.255.0
!
```

Example 3-5 *EIGRP Passes a Default Route Only if It Is Configured to Do So (Continued)*

```
ip route 0.0.0.0 0.0.0.0 serial 0/0/0
!
router eigrp 100
  network 0.0.0.0
<output omitted>
```

Route Summarization

Some EIGRP features, such as automatically summarizing routes at a major network boundary, have distance vector characteristics. Traditional distance vector protocols, which are classful routing protocols, must summarize at network boundaries. They cannot presume the mask for networks that are not directly connected, because masks are not exchanged in the routing updates.

KEY POINT

EIGRP Summarization

EIGRP automatically summarizes on the major network boundary by default; this feature can be turned off. In addition, EIGRP summary routes can be configured on any bit boundary within the network as long as a more specific route exists in the routing table.

Summarizing routes at classful major network boundaries creates smaller routing tables. Smaller routing tables, in turn, make the routing update process less bandwidth intensive. Cisco distance vector routing protocols have autosummarization enabled by default. As mentioned earlier, EIGRP has its roots in IGRP and, therefore, summarizes at the network boundary by default. For EIGRP, this feature can be turned off.

The inability to create summary routes at arbitrary boundaries with a major network has been a drawback of distance vector protocols since their inception. EIGRP has added functionality to allow administrators to create one or more summary routes within a network on any bit boundary (as long as a more specific route exists in the routing table). When the last specific route of the summary goes away, the summary route is deleted from the routing table. The minimum metric of the specific routes is used as the metric of the summary route.

When summarization is configured on a router's interface, a summary route is added to that router's routing table, with the route's next-hop interface set to null0—a directly connected, software-only interface. The use of the null0 interface prevents the router from trying to forward traffic to other routers in search of a more precise, longer match, thus preventing traffic from looping within the network. For example, if the summarizing router receives a packet to an unknown subnet that is part of the summarized range, the packet matches the summary route based on the longest match. The packet is forwarded to the null0 interface (in other words, it is dropped or sent to the *bit bucket*). This prevents the router from forwarding the packet to a default route and possibly creating a loop.

For effective summarization, blocks of contiguous addresses (subnets) should funnel back to a common router so that a single summary route can be created and then advertised. The number of subnets that can be represented by a summary route is directly related to the difference in the number of bits between the subnet mask and the summary mask. The formula 2^n, where n equals the difference in the number of bits between the summary and subnet masks, indicates how many subnets can be represented by a single summary route. For example, if the summary mask contains 3 fewer bits than the subnet mask, eight ($2^3 = 8$) subnets can be aggregated into one advertisement.

For example, if network 10.0.0.0 is divided into /24 subnets and some of these subnets are summarized to the summarization block 10.1.8.0/21, the difference between the /24 networks and the /21 summarizations is 3 bits; therefore, $2^3 = 8$ subnets can be aggregated. The summarized subnets range from 10.1.8.0/24 through 10.1.15.0/24.

When creating summary routes, the administrator needs to specify the IP address of the summary route and the summary mask. Cisco IOS handles the details of proper implementation, such as metrics, loop prevention, and removal of the summary route from the routing table if none of the more specific routes are valid.

Configuring Manual Route Summarization

EIGRP automatically summarizes routes at the classful boundary, but, as discussed, in some cases you might want to turn off this feature. For example, if you have discontiguous subnets, you need to disable autosummarization. Note that an EIGRP router does not perform automatic summarization of networks in which it does not participate.

To turn off automatic summarization, use the **no auto-summary** router configuration command. Use the **ip summary-address eigrp** *as-number address mask* [*admin-distance*] interface configuration command to manually create a summary route at an arbitrary bit boundary, as long as a more specific route exists in the routing table. Table 3-2 summarizes the parameters for this command.

Table 3-2 **ip summary-address eigrp** *Command Parameters*

Parameter	Description
as-number	EIGRP autonomous system number.
address	The IP address being advertised as the summary address. This address does not need to be aligned on Class A, B, or C boundaries.
mask	The IP subnet mask being used to create the summary address.
admin-distance	(Optional) Administrative distance. A value from 0 to 255.

For example, Figure 3-20 shows a discontiguous network 172.16.0.0. By default, both Routers A and B summarize routes at the classful boundary; as a result, Router C would have two equally good routes to network 172.16.0.0 and would perform load balancing between Router A and Router B. This would not be correct routing behavior.

Figure 3-20 *Summarizing EIGRP Routes*

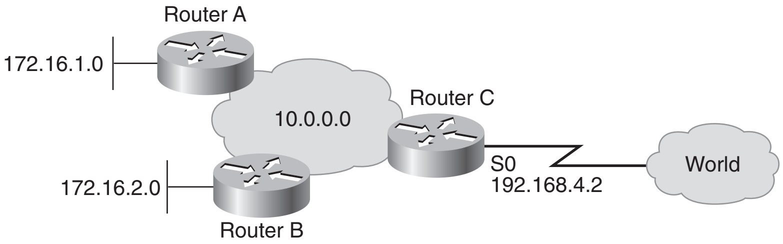

As shown in Example 3-6, you can disable the automatic route summarization on Router A; the same configuration would be done on Router B. With this configuration, Router C knows precisely that 172.16.1.0 is reached via Router A and that 172.16.2.0 is reached only via Router B. The routing tables of the routers in the 10.0.0.0 network, including Router C, now include these discontiguous subnets.

Example 3-6 *Turning Off EIGRP Autosummarization on Router A (and Router B) in Figure 3-20*

```
RouterA(config)#router eigrp 1
RouterA(config-router)#network 10.0.0.0
RouterA(config-router)#network 172.16.0.0
RouterA(config-router)#no auto-summary
```

An EIGRP router autosummarizes routes for only networks to which it is attached. If a network was not autosummarized at the major network boundary, as is the case in this example on Routers A and B because autosummarization is turned off, all the subnet routes are carried into Router C's routing table. Router C will not autosummarize the 172.16.1.0 and 172.16.2.0 subnets because it does not own the 172.16.0.0 network. Therefore, Router C would send routing information about the 172.16.1.0 subnet and the 172.16.2.0 subnet to the WAN.

Forcing a summary route out Router C's interface s0/0/0, as shown in Example 3-7, helps reduce route advertisements about network 172.16.0.0 to the world.

Example 3-7 *Forcing Summarization on Router C in Figure 3-20*

```
RouterC#show run
<output omitted>
router eigrp 1
  network 10.0.0.0
  network 192.168.4.0
```

continues

Example 3-7 *Forcing Summarization on Router C in Figure 3-20 (Continued)*

```
!
<output omitted>
int s0/0/0
  ip address 192.168.4.2 255.255.255.0
  ip summary-address eigrp 1 172.16.0.0 255.255.0.0
<output omitted>
```

Example 3-8 illustrates Router C's routing table. Router C has both 172.16.1.0 and 172.16.2.0, the discontiguous subnets, in its routing table, and the summary route to null 0.

Example 3-8 *Routing Table of Router C in Figure 3-20*

```
RouterC#show ip route
<output omitted>
Gateway of last resort is not set
     172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:00:04, Null0
D       172.16.1.0/24 [90/156160] via 10.1.1.2, 00:00:04, FastEthernet0/0
D       172.16.2.0/24 [90/20640000] via 10.2.2.2, 00:00:04, Serial0/0/1
C    192.168.4.0/24 is directly connected, Serial0/0/0
     10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C       10.2.2.0/24 is directly connected, Serial0/0/1
C       10.1.1.0/24 is directly connected, FastEthernet0/0
D       10.0.0.0/8 is a summary, 00:00:05, Null0
RouterC#
```

KEY POINT

Summary Route

For manual summarization, the summary is advertised only if a component (a more specific entry that is represented in the summary) of the summary route is present in the routing table.

NOTE IP EIGRP summary routes are given an administrative distance value of 5. Standard EIGRP routes receive an administrative distance of 90, and external EIGRP routes receive an administrative distance of 170.

You will notice the EIGRP summary route with an administrative distance of 5 only on the local router that is performing the summarization (with the **ip summary-address eigrp** command), by using the **show ip route** *network* command, where *network* is the specified summarized route.

EIGRP Load Balancing

KEY POINT

Load Balancing

Load balancing is a router's capability to distribute traffic over all of its network ports that are the same metric from the destination address.

Load balancing increases the utilization of network segments, thus increasing effective network bandwidth.

By default, the Cisco IOS balances between a maximum of four equal-cost paths for IP. Using the **maximum-paths** *maximum-path* router configuration command, you can request that up to 16 equally good routes be kept in the routing table (set *maximum-path* to 1 to disable load balancing). When a packet is process-switched, load balancing over equal-cost paths occurs on a per-packet basis. When packets are fast-switched, load balancing over equal-cost paths is on a per-destination basis.

NOTE If you are testing load balancing, do not ping to or from the routers with the fast-switching interfaces, because these locally router-generated packets are process-switched rather than fast-switched and might produce confusing results.

NOTE Load balancing is performed only on traffic that passes *through* the router, not traffic generated by the router.

EIGRP can also balance traffic across multiple routes that have different metrics; this is called unequal-cost load balancing. The degree to which EIGRP performs load balancing is controlled by the **variance** *multiplier* router configuration command. The multiplier is a variance value, between 1 and 128, used for load balancing. The default is 1, which means equal-cost load balancing. The multiplier defines the range of metric values that are accepted for load balancing. For example, in Figure 3-21, a variance of 2 is configured, and the range of the metric values (the FDs) for Router E to get to Network Z is 20 to 45. This range of values is used in the procedure to determine the feasibility of a potential route.

KEY POINT

Feasible Route with Variance

A route is feasible if the next router in the path is closer to the destination than the current router and if the metric for the entire alternate path is within the variance.

Figure 3-21 *EIGRP Load Balancing with a Variance of 2*

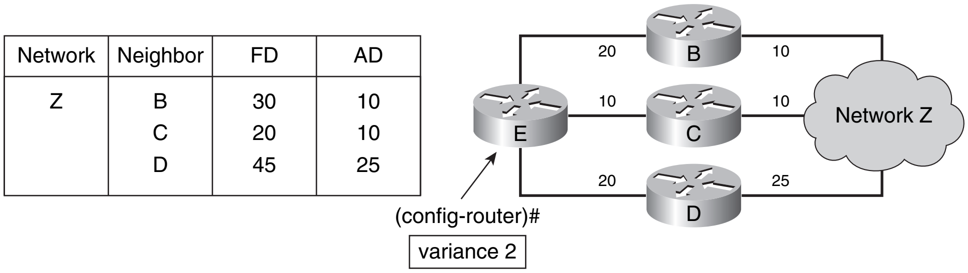

Network	Neighbor	FD	AD
Z	B	30	10
	C	20	10
	D	45	25

Only paths that are feasible can be used for load balancing; the routing table indicates only feasible paths. The two feasibility conditions are as follows:

- The local best metric (the current FD) must be greater than the best metric (the AD) learned from the next router. In other words, the next router in the path must be closer to the destination than the current router; this prevents routing loops.
- The variance multiplied by the local best metric (the current FD) must be greater than the metric through the next router (the alternative FD).

If both of these conditions are met, the route is called feasible and can be added to the routing table.

To control how traffic is distributed among routes when multiple routes exist for the same destination network and they have different metrics, use the **traffic-share** [**balanced** | **min across-interfaces**] router configuration command. With the keyword **balanced**, the router distributes traffic proportionately to the ratios of the metrics associated with the different routes. With the **min across-interfaces** option, the router uses only routes that have minimum costs. (In other words, all routes that are feasible and within the variance are kept in the routing table, but only those with the minimum cost are used.)

In Figure 3-21, Router E has three paths to Network Z, with the following metrics:

- Path 1: 30 (via B)
- Path 2: 20 (via C)
- Path 3: 45 (via D)

Router E uses Router C as the successor because its FD is lowest (20). With the **variance 2** command applied to Router E, the path through Router B meets the criteria for load balancing. In this case, the FD through Router B is less than twice the FD for the successor (Router C). Router D is not considered for load balancing because the FD through Router D is greater than twice the FD for the successor (Router C). In this example, however, Router D would never be a feasible successor, no matter what the variance is. Router D is not a feasible successor because its AD of 25 is greater than Router E's FD of 20; therefore, to avoid a potential routing loop, Router D is not considered closer to the destination than Router E and cannot be a feasible successor.

In another example of unequal load balancing, four paths to a destination have the following different metrics:

- Path 1: 1100
- Path 2: 1100
- Path 3: 2000
- Path 4: 4000

By default, the router routes to the destination using both Paths 1 and 2. Assuming no potential routing loops exist, you would use the **variance 2** command to load balance over Paths 1, 2, and 3, because 1100 * 2 = 2200, which is greater than the metric through Path 3. Similarly, to also include Path 4, you would issue the **variance 4** command.

EIGRP and WAN Links

EIGRP operates efficiently in WAN environments and is scalable on both point-to-point links and NBMA multipoint and point-to-point links. Because of the inherent differences in links' operational characteristics, default configuration of WAN connections might not be optimal. A solid understanding of EIGRP operation coupled with knowledge of link speeds can yield an efficient, reliable, scalable router configuration.

EIGRP Link Utilization

KEY POINT

EIGRP Bandwidth on an Interface

By default, EIGRP uses up to 50 percent of the bandwidth declared on an interface or subinterface. EIGRP uses the bandwidth of the link set by the **bandwidth** command, or the link's default bandwidth if none is configured, when calculating how much bandwidth to use.

You can adjust this percentage on an interface or subinterface with the **ip bandwidth-percent eigrp** *as-number percent* interface configuration command. The *as-number* is the EIGRP autonomous system number. The *percent* parameter is the percentage of the configured bandwidth that EIGRP can use. You can set the percentage to a value greater than 100, which might be useful if the bandwidth is configured artificially low for routing policy reasons. Example 3-9 shows a configuration that allows EIGRP to use 40 kbps (200 percent of the configured bandwidth, 20 kbps) on the interface. It is essential to make sure that the line is provisioned to handle the configured capacity. (The next section, "Examples of EIGRP on WANs," provides more examples of when this command is useful.)

Example 3-9 *Adjusting the EIGRP Link Utilization*

```
Router(config)#interface serial0/0/0
Router(config-if)#bandwidth 20
Router(config-if)#ip bandwidth-percent eigrp 1 200
```

Cisco IOS assumes that point-to-point Frame Relay subinterfaces are operating at the default speed of the interface. In many implementations, however, only fractional T1 speeds are available. Therefore, when configuring these subinterfaces, set the bandwidth to match the contracted CIR.

When configuring multipoint interfaces (especially for Frame Relay, but also for ATM and ISDN PRI), remember that the bandwidth is shared equally by all neighbors. That is, EIGRP uses the **bandwidth** command on the physical interface divided by the number of Frame Relay neighbors connected on that physical interface to get the bandwidth attributed to each neighbor. EIGRP configuration should reflect the correct percentage of the actual available bandwidth on the line.

Each installation has a unique topology, and with that comes unique configurations. Differing CIR values often require a hybrid configuration that blends the characteristics of point-to-point circuits with multipoint circuits. When configuring multipoint interfaces, configure the bandwidth to represent the minimum CIR times the number of circuits. This approach might not fully use the higher-speed circuits, but it ensures that the circuits with the lowest CIR will not be overdriven. If the topology has a small number of very low-speed circuits, these interfaces are typically defined as point-to-point so that their bandwidth can be set to match the provisioned CIR.

Examples of EIGRP on WANs

In Figure 3-22, Router C's interface has been configured for a bandwidth of 224 kbps. Four neighbors exist in this pure multipoint topology, so each circuit is allocated one-quarter of the configured bandwidth on the interface, and this 56-kbps allocation matches the provisioned CIR of each circuit.

Figure 3-22 *Frame Relay Multipoint in Which All VCs Share the Bandwidth Evenly*

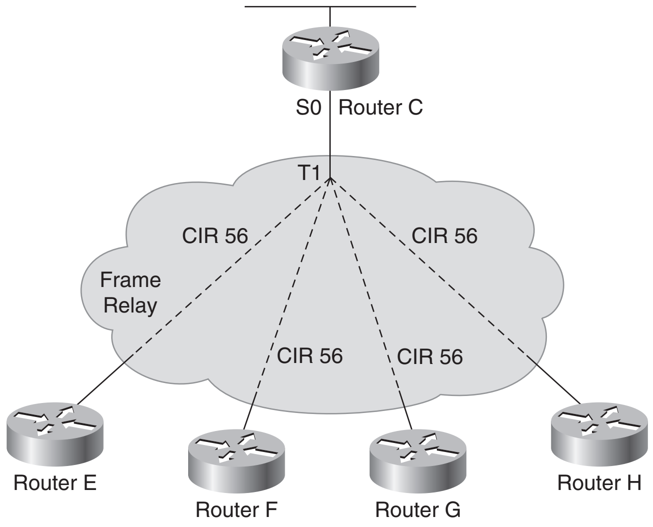

• All VCs share bandwidth evenly: 4 x 56 = 224

Example 3-10 shows the configuration for Router C's Serial 0 interface.

Example 3-10 *Adjusting the* **bandwidth** *Command on an Interface on Router C in Figure 3-22*

```
RouterC(config)#interface serial 0
RouterC(config-if)#encapsulation frame-relay
RouterC(config-if)#bandwidth 224
```

In Figure 3-23, one of the circuits has been provisioned for a 56-kbps CIR, and the other circuits have a higher CIR. This interface has been configured for a bandwidth that represents the lowest CIR multiplied by the number of circuits being supported (56 * 4 = 224). This configuration protects against overwhelming the slowest-speed circuit in the topology.

Figure 3-23 *Frame Relay Multipoint in Which VCs Have Different CIRs*

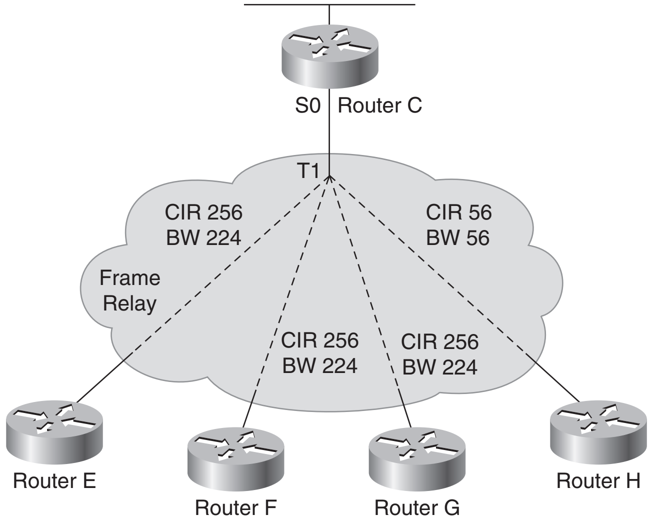

- Lowest CIR x # of VC: 56 x 4 = 224

Figure 3-24 presents a hybrid solution. There is only one low-speed circuit, and other VCs are provisioned for a higher CIR.

Figure 3-24 *Frame Relay Multipoint and Point-to-Point*

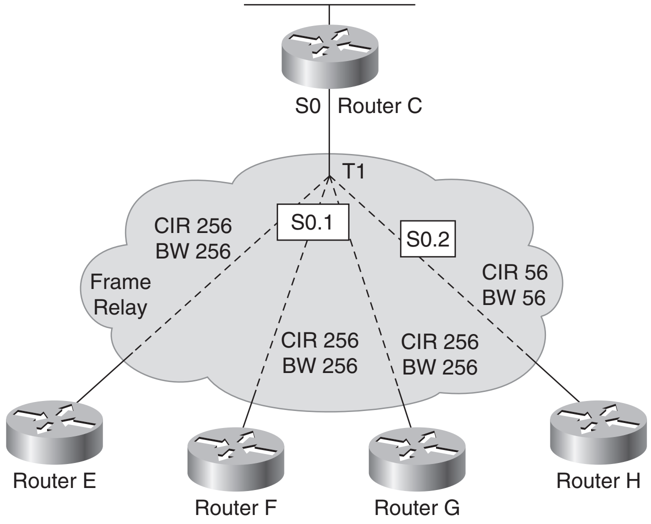

- Configure lowest CIR VC as point-to-point, specify BW = CIR
- Configure higher CIR VCs as multipoint, combine CIRs

Example 3-11 shows the configuration applied to Router C in Figure 3-24.

Example 3-11 *Adjusting the Bandwidth for a Frame Relay Subinterface on Router C in Figure 3-24*

```
RouterC(config)#interface serial 0.1 multipoint
RouterC(config-subif)#bandwidth 768
RouterC(config-subif)#exit
RouterC(config)#interface serial 0.2 point-to-point
RouterC(config-subif)#bandwidth 56
```

Example 3-11 shows the low-speed circuit configured as point-to-point. The remaining circuits are designated as multipoint, and their respective CIRs are added up to set the interface's bandwidth.

Figure 3-25 illustrates a common hub-and-spoke oversubscribed topology with 10 VCs to the remote sites. (Only 4 of the 10 remote sites are shown in the figure.)

Figure 3-25 *Frame Relay Hub-and-Spoke Topology*

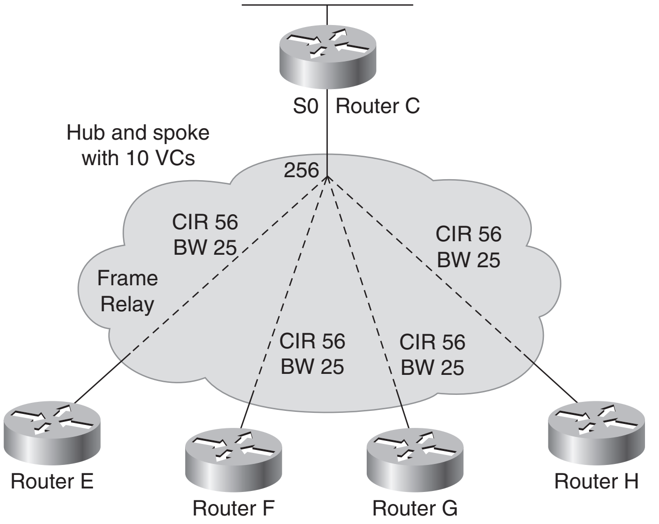

- Configure each VC as point-to-point, specify BW = 1/10 of link capacity
- Increase EIGRP utilization to 50% of actual VC capacity

The circuits are provisioned as 56-kbps links, but there is insufficient bandwidth at the interface to support the allocation. For example, if the hub tries to communicate to all remote sites at the same time, the bandwidth that is required exceeds the available link speed of 256 kbps for the hub—10 times the CIR of 56 kbps equals 560 kbps. In a point-to-point topology, all VCs are treated equally and are therefore configured for exactly one-tenth of the available link speed (25 kbps).

Example 3-12 shows the configuration used on Routers C and G of Figure 3-25.

Example 3-12 *EIGRP WAN Configuration: Point-to-Point Links on Routers C and G in Figure 3-25*

```
RouterC(config)#interface serial 0.1 point-to-point
RouterC(config-subif)#bandwidth 25
RouterC(config-subif)#ip bandwidth-percent eigrp 63 110
<output omitted>
RouterC(config)#interface serial 0.10 point-to-point
RouterC(config-subif)#bandwidth 25
RouterC(config-subif)#ip bandwidth-percent eigrp 63 110

RouterG(config)#interface serial 0
RouterG(config-if)#bandwidth 25
RouterG(config-if)#ip bandwidth-percent eigrp 63 110
```

By default, EIGRP uses 50 percent of a circuit's configured bandwidth. As mentioned, EIGRP configuration should reflect the correct percentage of the actual available bandwidth on the line. Therefore, in an attempt to ensure that EIGRP packets are delivered through the Frame Relay network in Figure 3-25, each subinterface has the EIGRP allocation percentage raised to 110 percent of the specified bandwidth. This adjustment causes EIGRP packets to receive approximately 28 kbps of the provisioned 56 kbps on each circuit. This extra configuration restores the 50-50 ratio that was tampered with when the bandwidth was set to an artificially low value.

NOTE Suppressing ACKs also saves bandwidth. An ACK is not sent if a unicast data packet is ready for transmission. The ACK field in any reliable unicast packet (RTP packet) is sufficient to acknowledge the neighbor's packet, so the ACK packet is suppressed to save bandwidth. This is a significant feature for point-to-point links and NBMA networks, because on those media, all data packets are sent as unicasts and, thus, can carry an acknowledgment themselves (this is also known as a piggyback ACK). In that instance, there is no need for an ACK packet.

Configuring EIGRP Authentication

You can prevent your router from receiving fraudulent route updates by configuring neighbor router authentication. You can configure EIGRP neighbor authentication (also called *neighbor router authentication* or *route authentication*) such that routers can participate in routing based on predefined passwords.

This section first describes router authentication in general, followed by a discussion of how to configure and troubleshoot EIGRP Message Digest 5 (MD5) authentication.

Router Authentication

Neighbor router authentication can be configured such that routers only participate in routing based on predefined passwords.

By default, no authentication is used for routing protocol packets. When neighbor router authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives. This is accomplished by the exchange of an authentication key (also called a *password*) that is known to both the sending and the receiving router.

The router uses two types of authentication:

- **Simple password authentication (also called plain text authentication)**—Supported by Integrated System-Integrated System (IS-IS), Open Shortest Path First (OSPF), and Routing Information Protocol Version 2 (RIPv2)
- **MD5 authentication**—Supported by OSPF, RIPv2, BGP, and EIGRP

NOTE This book covers authentication for EIGRP, OSPF, and BGP.

Both forms of authentication work in the same way, with the exception that MD5 sends a message digest instead of the authenticating key itself. The message digest is created using the key (and a key ID with some protocols) and a message, but the key itself is not sent, preventing it from being read while it is being transmitted. Simple password authentication sends the authenticating key itself over the wire.

NOTE Simple password authentication is not recommended for use as part of your security strategy because it is vulnerable to passive attacks. Anybody with a link analyzer could easily view the password on the wire. The primary use of simple password authentication is to avoid accidental changes to the routing infrastructure. Using MD5 authentication, however, is a recommended security practice.

CAUTION As with all keys, passwords, and other security secrets, it is imperative that you closely guard the keys used in neighbor authentication. The security benefits of this feature are reliant upon keeping all authenticating keys confidential. Also, when performing router management tasks via Simple Network Management Protocol (SNMP), do not ignore the risk associated with sending keys using nonencrypted SNMP.

With simple password authentication, a password (key) is configured on a router; each participating neighbor router must be configured with the same key.

MD5 authentication is a cryptographic authentication. A key (password) and key ID are configured on each router. The router uses an algorithm based on the routing protocol packet, the key, and the key ID to generate a message digest (also called a *hash*) that is appended to the packet. Unlike the simple authentication, the key is not exchanged over the wire—the message digest is sent instead of the key, which ensures that nobody can eavesdrop on the line and learn keys during transmission.

EIGRP MD5 Authentication

By default, no authentication is used for EIGRP packets. You can configure EIGRP to use MD5 authentication.

When EIGRP neighbor authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives. The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

For EIGRP MD5 authentication, you must configure an authenticating *key* and a *key ID* on both the sending router and the receiving router. Each key has its own key ID, which is stored locally. The combination of the key ID and the interface associated with the message uniquely identifies the authentication algorithm and MD5 authentication key in use.

EIGRP allows keys to be managed using *key chains*. Each key definition within the key chain can specify a time interval for which that key will be activated (known as its lifetime). Then, during a given key's lifetime, routing update packets are sent with this activated key. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest, and it uses the first valid key it encounters.

KEY POINT

EIGRP Keys

When configuring EIGRP authentication, you specify the key ID (number), the key (password), and the lifetime of the key. The first (by key ID), valid (by lifetime), key is used.

Keys cannot be used during time periods for which they are not activated. Therefore, it is recommended that for a given key chain, key activation times overlap to avoid any period of time for which no key is activated. If a time period occurs during which no key is activated, neighbor authentication cannot occur, and therefore routing updates will fail.

NOTE The router needs to know the time to be able to rotate through keys in synchronization with the other participating routers, so that all routers are using the same key at the same moment. Refer to the Network Time Protocol (NTP) and calendar commands in the “Performing Basic System Management” chapter of the *Cisco IOS Configuration Fundamentals Configuration Guide* for information about configuring time on your router.

Configuring MD5 Authentication

To configure MD5 authentication for EIGRP, complete the following steps:

Step 1 Enter configuration mode for the interface on which you want to enable authentication.

Step 2 Specify MD5 authentication for EIGRP packets using the **ip authentication mode eigrp** *autonomous-system* **md5** interface configuration command. The *autonomous-system* is the EIGRP autonomous system number in which authentication is to be used.

Step 3 Enable the authentication of EIGRP packets with a key specified in a key chain by using the **ip authentication key-chain eigrp** *autonomous-system name-of-chain* interface configuration command. The *autonomous-system* parameter specifies the EIGRP autonomous system number in which authentication is to be used. The *name-of-chain* parameter specifies the name of the authentication key chain from which a key is to be obtained for this interface.

Step 4 Enter the configuration mode for the key chain using the **key chain** *name-of-chain* global configuration command.

Step 5 Identify a key ID to use and enter configuration mode for that key using the **key** *key-id* key-chain configuration command. The *key-id* is the ID number of an authentication key on a key chain. The range of keys is from 0 to 2147483647; the key ID numbers need not be consecutive.

Step 6 Identify the key string (password) for this key using the **key-string** *key* key-chain-key configuration command. The *key* is the authentication key-string that is to be used to authenticate sent and received EIGRP packets. The key string can contain from 1 to 80 uppercase and lowercase alphanumeric characters, except that the first character cannot be a number.

Step 7 Optionally specify the time period during which this key will be accepted for use on received packets using the **accept-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*} key-chain-key configuration command. Table 3-3 describes the parameters for this command.

Table 3-3 **accept-lifetime** *Command Parameters*

Parameter	Description

start-time	Beginning time that the key specified by the **key** command is valid for use on received packets. The syntax can be either of the following:
	`hh:mm:ss month date year`
	`hh:mm:ss date month year`
	`hh—hours`
	`mm—minutes`
	`ss—seconds`
	`month—first three letters of the month`
	`date—date (1-31)`
	`year—year (four digits)`
	The default start time and the earliest acceptable date is January 1, 1993.

infinite	Indicates the key is valid for use on received packets from the *start-time* value on.
end-time	Indicates the key is valid for use on received packets from the *start-time* value until the *end-time* value. The syntax is the same as that for the *start-time* value. The *end-time* value must be after the *start-time* value. The default end time is an infinite time period.
seconds	Length of time (in seconds) that the key is valid for use on received packets. The range is from 1 to 2147483646.

Step 8 Optionally specify the time period during which this key can be used for sending packets using the **send-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*} key-chain-key configuration command. Table 3-4 describes the parameters for this command.

Table 3-4 **send-lifetime** *Command Parameters*

Parameter	Description

start-time	Beginning time that the key specified by the **key** command is valid to be used for sending packets. The syntax can be either of the following:
	`hh:mm:ss month date year`
	`hh:mm:ss date month year`
	`hh—hours`
	`mm—minutes`
	`ss—seconds`
	`month—first three letters of the month`
	`date—date (1-31)`
	`year—year (four digits)`
	The default start time and the earliest acceptable date is January 1, 1993.

infinite	Indicates the key is valid to be used for sending packets from the *start-time* value on.
end-time	Indicates the key is valid to be used for sending packets from the *start-time* value until the *end-time* value. The syntax is the same as that for the *start-time* value. The *end-time* value must be after the *start-time* value. The default end time is an infinite time period.
seconds	Length of time (in seconds) that the key is valid to be used for sending packets. The range is from 1 to 2147483646.

NOTE If the **service password-encryption** command is not used when implementing EIGRP authentication, the key-string will be stored as plain text in the router configuration. If you configure the **service password-encryption** command, the key-string will be stored and displayed in an encrypted form; when it is displayed, there will be an *encryption-type* of 7 specified before the encrypted key-string.

MD5 Authentication Configuration Example

Figure 3-26 shows the network used to illustrate the configuration, verification, and troubleshooting of MD5 authentication.

Figure 3-26 *Network for EIGRP Authentication Configuration Example*

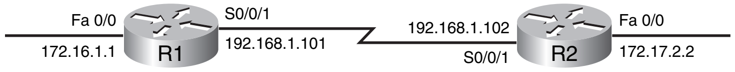

Example 3-13 shows the configuration of the R1 router.

Example 3-13 *Configuration of Router R1 in Figure 3-26*

```
R1#show running-config
<output omitted>
key chain R1chain
 key 1
  key-string firstkey
  accept-lifetime 04:00:00 Jan 1 2006 infinite
  send-lifetime 04:00:00 Jan 1 2006 04:01:00 Jan 1 2006
 key 2
  key-string secondkey
  accept-lifetime 04:00:00 Jan 1 2006 infinite
  send-lifetime 04:00:00 Jan 1 2006 infinite
<output omitted>
interface FastEthernet0/0
 ip address 172.16.1.1 255.255.255.0
!
interface Serial0/0/1
 bandwidth 64
 ip address 192.168.1.101 255.255.255.224
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 R1chain
!
router eigrp 100
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
 auto-summary
```

MD5 authentication is configured on the serial 0/0/1 interface with the **ip authentication mode eigrp 100 md5** command. The **ip authentication key-chain eigrp 100 R1chain** command specifies that the key chain *R1chain* is to be used.

The **key chain R1chain** command enters configuration mode for the *R1chain* key chain. Two keys are defined. Key 1 is set to *firstkey* with the **key-string firstkey** command. This key is acceptable for use on packets received by R1 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. However, the **send-lifetime 04:00:00 Jan 1 2006 04:01:00 Jan 1 2006** command specifies that this key was only valid for use when sending packets for 1 minute on January 1, 2006; it is no longer valid for use in sending packets.

Key 2 is set to *secondkey* with the **key-string secondkey** command. This key is acceptable for use on packets received by R1 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

R1 will therefore accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1 or 2. All other MD5 packets will be dropped. R1 will send all EIGRP packets using key 2, because key 1 is no longer valid for use when sending packets.

Example 3-14 shows the configuration of the R2 router.

Example 3-14 *Configuration of Router R2 in Figure 3-26*

```
R2#show running-config
<output omitted>
key chain R2chain
 key 1
  key-string firstkey
  accept-lifetime 04:00:00 Jan 1 2006 infinite
  send-lifetime 04:00:00 Jan 1 2006 infinite
 key 2
  key-string secondkey
  accept-lifetime 04:00:00 Jan 1 2006 infinite
  send-lifetime 04:00:00 Jan 1 2006 infinite
<output omitted>
interface FastEthernet0/0
 ip address 172.17.2.2 255.255.255.0
!
interface Serial0/0/1
 bandwidth 64
 ip address 192.168.1.102 255.255.255.224
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 R2chain
!
router eigrp 100
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
auto-summary
```

MD5 authentication is configured on the serial 0/0/1 interface with the **ip authentication mode eigrp 100 md5** command. The **ip authentication key-chain eigrp 100 R2chain** command specifies that the key chain *R2chain* is to be used.

The **key chain R2chain** command enters configuration mode for the *R2chain* key chain. Two keys are defined. Key 1 is set to *firstkey* with the **key-string firstkey** command. This key is acceptable for use on packets received by R2 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

Key 2 is set to *secondkey* with the **key-string secondkey** command. This key is acceptable for use on packets received by R2 from January 1, 2006 onward, as specified in the **accept-lifetime**

04:00:00 Jan 1 2006 infinite command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

R2 will therefore accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1 or 2. R2 will send all EIGRP packets using key 1, because it is the first valid key in the key chain.

Verifying MD5 Authentication

Example 3-15 provides the output of the **show ip eigrp neighbors** and **show ip route** commands on the R1 router depicted in the network in Figure 3-26. The neighbor table indicates that the two routers have successfully formed an EIGRP adjacency. The routing table verifies that the 172.17.0.0 network has been learned via EIGRP over the serial connection. Example 3-15 also shows the results of a **ping** to the R2 Fast Ethernet interface address to illustrate that the link is working.

Example 3-15 *Output on Router R1 in Figure 3-26*

```
R1#
*Apr 21 16:23:30.517: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.1.102 (Serial0/
  0/1) is up: new adjacency

R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   192.168.1.102           Se0/0/1           12 00:03:10   17  2280  0  14
R1#show ip route
<output omitted>
Gateway of last resort is not set
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:02:22, Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:31:31, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.96/27 is directly connected, Serial0/0/1
D       192.168.1.0/24 is a summary, 00:31:31, Null0
R1#ping 172.17.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.17.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/15/16 ms
```

Troubleshooting MD5 Authentication

This section provides some examples of how to troubleshoot EIGRP MD5 authentication.

Example of Successful MD5 Authentication

Example 3-16 shows output from the **debug eigrp packets** command on R1, which displays that R1 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 1, from R2.

Example 3-16 **debug eigrp packets** *Command Output on Router R1 in Figure 3-26*

```
R1#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
*Apr 21 16:38:51.745: EIGRP: received packet with MD5 authentication, key id = 1
*Apr 21 16:38:51.745: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.102
*Apr 21 16:38:51.745:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
```

Similarly, the output of the **debug eigrp packets** command on R2 shown Example 3-17 illustrates that R2 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from R1.

Example 3-17 **debug eigrp packets** *Command Output on Router R2 in Figure 3-26*

```
R2#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
R2#
*Apr 21 16:38:38.321: EIGRP: received packet with MD5 authentication, key id = 2
*Apr 21 16:38:38.321: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.101
*Apr 21 16:38:38.321:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
```

Example of Troubleshooting MD5 Authentication Problems

For this example, the key string for router R1's key 2, the one that it uses when sending EIGRP packets, is changed to be different from the key string that router R2 is expecting. Example 3-18 shows the changes to the configuration for R1.

Example 3-18 *Changes to the Configuration of Router R1 in Figure 3-26*

```
R1(config-if)#key chain R1chain
R1(config-keychain)#key 2
R1(config-keychain-key)#key-string wrongkey
```

The output of the **debug eigrp packets** command on R2 shown in Example 3-19 illustrates that R2 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from R1, but that there is an authentication mismatch. The EIGRP packets from R1 are ignored, and the neighbor relationship is declared to be down. The output of the **show ip eigrp neighbors** command confirms that R2 does not have any EIGRP neighbors.

Example 3-19 *Output on Router R2 in Figure 3-26*

```
R2#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
R2#
*Apr 21 16:50:18.749: EIGRP: pkt key id = 2, authentication mismatch
*Apr 21 16:50:18.749: EIGRP: Serial0/0/1: ignored packet from 192.168.1.101, opcode = 5
  (invalid authentication)
*Apr 21 16:50:18.749: EIGRP: Dropping peer, invalid authentication
*Apr 21 16:50:18.749: EIGRP: Sending HELLO on Serial0/0/1
*Apr 21 16:50:18.749:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
*Apr 21 16:50:18.753: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.1.101
 (Serial0/0/1) is down: Auth failure

R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
R2#
```

The two routers keep trying to reestablish their neighbor relationship. Because of the different keys used by each router in this scenario, R1 will authenticate hello messages sent by R2 using key 1. When R1 sends a hello message back to R2 using key 2, however, an authentication mismatch exists. From R1's perspective, the relationship appears to be up for a while, but then it times out, as illustrated by the messages received on R1 shown in Example 3-20. The output of the **show ip eigrp neighbors** command on R1 also illustrates that R1 does have R2 in its neighbor table for a short time.

Example 3-20 *Output on Router R1 in Figure 3-26*

```
R1#
*Apr 21 16:54:09.821: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.1.102 (Serial0/
  0/1) is down: retry limit exceeded
*Apr 21 16:54:11.745: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.1.102 (Serial0/
  0/1) is up: new adjacency
R1#show ip eigrp neighbors
H   Address        Interface   Hold Uptime   SRTT   RTO  Q  Seq
                               (sec)         (ms)       Cnt Num
0   192.168.1.102  Se0/0/1       13 00:00:38    1   5000  1  0
```

Using EIGRP in an Enterprise Network

EIGRP is a scalable routing protocol that ensures that as a network grows larger, it operates efficiently and adjusts rapidly to changes. This section describes practical EIGRP-specific design and configuration techniques to implement an effective, scalable enterprise network.

EIGRP Scalability

The following are some of the many variables that affect network scalability:

- **The amount of information exchanged between neighbors**—If more information than necessary for routing to function correctly is exchanged between EIGRP neighbors, unnecessary work during routing startup and topology changes results.
- **Number of routers**—When a topology change occurs, the amount of resources consumed by EIGRP is directly related to the number of routers that must be involved in the change.
- **The topology's depth**—The topology's depth can affect the convergence time. Depth refers to the number of hops that information must travel to reach all routers. For example, a multinational network without route summarization has a large depth and therefore increased convergence time.

  A three-tiered network design (as described in Chapter 1, "Network Architecture Framework and Design Models") is highly recommended for all IP routing environments. There should never be more than seven hops between any two routing devices on an enterprise internetwork. The propagation delay and the query process across multiple hops when changes occur can slow down the convergence of the network when routes are lost.
- **The number of alternative paths through the network**—A network should provide alternative paths to avoid single points of failure. However, too many alternative paths can create problems with EIGRP convergence, because the EIGRP routing process, using queries, needs to explore all possible paths for lost routes. This complexity creates an ideal condition for a router to become stuck-in-active (described in the later "EIGRP Queries and Stuck-in-Active" section) as it awaits a response to queries that are being propagated through these many alternative paths.

For proper EIGRP operation, you should follow some common design principles. For example, routers located at convergence points within the network need sufficient memory to buffer a large number of packets and to support numerous processes related to routing large volumes of traffic.

On WAN links, and especially with the hub-and-spoke topology, enough bandwidth should be provided to prevent router overhead traffic from interfering with normal user-generated traffic. In this respect, the impact of EIGRP packets being lost because of contention for bandwidth might be greater than application delays experienced by some users.

EIGRP Route Summarization

Route summarization is most effective with a sound address allocation. Having a two- or three-layer hierarchical network design, with routers positioned by function rather than by geography, greatly assists traffic flow and route distribution.

NOTE For a full discussion of internetwork design, refer to *Designing for Cisco Internetwork Solutions (DESGN)* (Cisco Press, 2003).

Figure 3-27 shows the topology of a nonscalable internetwork in which addresses (subnets) are either randomly assigned or assigned by historical requirements. In this example, multiple subnets from different major networks are located in each cloud, requiring many subnet routes to be injected into the core. In addition, because of the random assignment of addresses, query traffic cannot be localized to any portion of the network, thus increasing convergence time. Administration and troubleshooting are also more complex in this scenario.

Figure 3-27 *Nonscalable Internetwork*

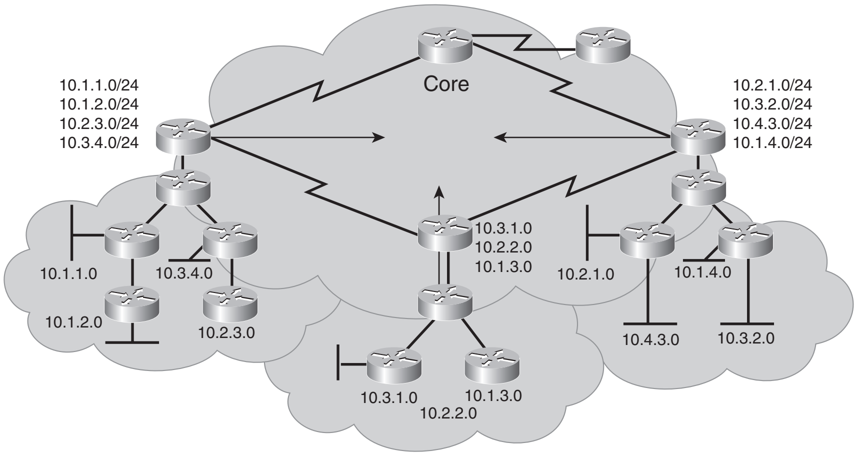

Figure 3-28 illustrates a better-designed network. Subnet addresses from individual major networks are localized within each cloud. This allows summary routes to be injected into the core. As an added benefit, the summary routes act as a boundary for the queries generated by a topology change.

Figure 3-28 *Scalable Internetwork*

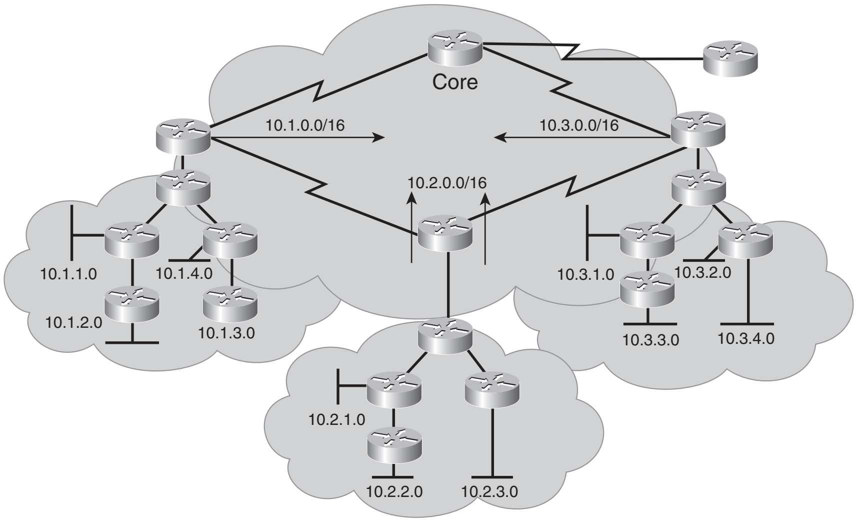

EIGRP Queries and Stuck-in-Active

As an advanced distance vector routing protocol, EIGRP relies on its neighbors to provide routing information. If a route is lost and no FS is available, EIGRP queries its neighbors about the lost route.

Recall that when a router loses a route and does not have an FS in its topology table, it looks for an alternative path to the destination. This is known as *going active* on a route. (A route is considered passive when a router is not recomputing that route.) Recomputing a route involves sending query packets to all neighbors on interfaces other than the one used to reach the previous successor (split horizon), inquiring whether they have a route to the given destination. If a router has an alternative route, it answers the query with a reply packet and does not propagate the query further; the reply includes the alternate route. If a neighbor does not have an alternative route, it queries each of its own neighbors for an alternative path. The queries then propagate through the network, thus creating an expanding tree of queries. When a router answers a query, it stops the spread of the query through that branch of the network; however, the query can still spread through other portions of the network as other routers attempt to find alternative paths, which might not exist.

Because of the reliable multicast approach used by EIGRP when searching for an alternative to a lost route, it is imperative that a reply be received for each query generated in the network. In other words, when a route goes active and queries are initiated, the only way this route can come out of the active state and transition to passive state is by receiving a reply for every generated query.

KEY POINT

Stuck-in-Active

If the router does not receive a reply to all the outstanding queries within 3 minutes (the default time), the route goes to the stuck-in-active (SIA) state.

NOTE You can change the active-state time limit from its default of 3 minutes using the **timers active-time** [*time-limit* | **disabled**] router configuration command. The *time-limit* is in minutes.

When a route goes to SIA state, the router resets the neighbor relationships for the neighbors that failed to reply. This causes the router to recompute all routes known through that neighbor and to re-advertise all the routes it knows about to that neighbor.

The most common reasons for SIA routes are as follows:

- The router is too busy to answer the query—generally as a result of high CPU usage or memory problems—and cannot allocate memory to process the query or build the reply packet.
- The link between the two routers is not good, so some packets are lost between the routers. The router receives an adequate number of packets to maintain the neighbor relationship, but the router does not receive all queries or replies.
- A failure causes traffic on a link to flow in only one direction. This is called a *unidirectional link*.

NOTE Use the **eigrp log-neighbor-changes** command to enable the logging of neighbor adjacency changes to monitor the routing system's stability and to help detect problems related to SIA.

One erroneous approach for decreasing the chances of a stuck-in-active route is to use multiple EIGRP autonomous systems to bound the query range. Many networks have been implemented using multiple EIGRP autonomous systems (to somewhat simulate OSPF areas), with mutual redistribution between the different autonomous systems. Although this approach changes how the network behaves, it does not always achieve the results intended. If a query reaches the edge of the autonomous system (where routes are redistributed into another autonomous system), the original query is answered. However, then the edge router initiates a new query in the other autonomous system. Therefore, the query process has not been stopped; the querying continues in the other autonomous system, where the route can potentially go in SIA.

Another misconception about autonomous system boundaries is that implementing multiple autonomous systems protects one autonomous system from route flaps in another autonomous system. If routes are redistributed between autonomous systems, route transitions from one autonomous system are detected in the other autonomous systems.

Preventing SIA Connections

SIA-Query and SIA-Reply are two new additions to the Type, Length, Value (TLV) triplets in the EIGRP packet header. These packets are generated automatically with no configuration required, from Cisco IOS Software Release 12.1(5) and later, with the *Active Process Enhancement* feature. This feature enables an EIGRP router to monitor the progression of the search for a successor route and ensure that the neighbor is still reachable. The result is improved network reliability by reducing the unintended termination of neighbor adjacency.

The diagram on the left in Figure 3-29 illustrates what would happen before this feature was introduced. Router A sends a query for network 10.1.1.0/24 to Router B. Router B has no entry for this network, so it queries Router C. If problems exist between Router B and C, the reply packet from Router C to Router B might be delayed or lost. Router A has no visibility of downstream progress and assumes that no response indicates problems with Router B. After Router A's 3-minute active timer expires, the neighbor relationship with Router B is reset, along with all known routes from Router B.

In contrast, with the Active Process Enhancement feature, as illustrated in the diagram on the right in Figure 3-29, Router A queries downstream Router B (with an SIA-Query) at the midway point of the active timer (one and a half minutes by default) about the status of the route. Router B responds (with an SIA-Reply) that it is searching for a replacement route. Upon receiving this SIA-Reply response packet, Router A validates the status of Router B and does not terminate the neighbor relationship.

Figure 3-29 *Cisco IOS Active Process Enhancement*

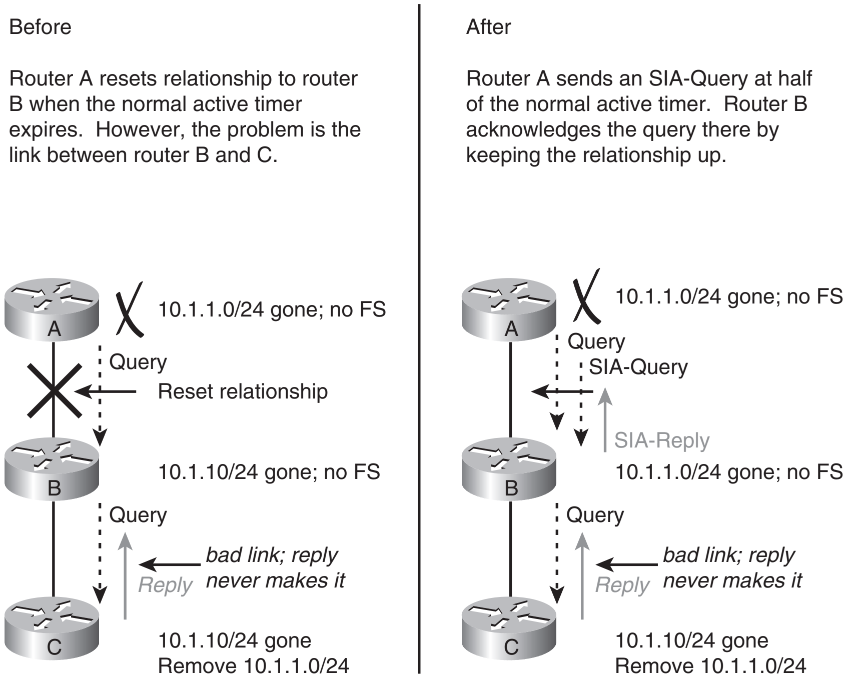

Before
Router A resets relationship to router B when the normal active timer expires. However, the problem is the link between router B and C.
A
10.1.1.0/24 gone; no FS
Query
Reset relationship
B
10.1.10/24 gone; no FS
Query
Reply
bad link; reply never makes it
C
10.1.10/24 gone
Remove 10.1.1.0/24
After
Router A sends an SIA-Query at half of the normal active timer. Router B acknowledges the query there by keeping the relationship up.
A
10.1.1.0/24 gone; no FS
Query
SIA-Query
SIA-Reply
B
10.1.1.0/24 gone; no FS
Query
Reply
bad link; reply never makes it
C
10.1.10/24 gone
Remove 10.1.1.0/24

Meanwhile, Router B will send up to three SIA-Queries to Router C. If they go unanswered, Router B will terminate the neighbor relationship with Router C. Router B will then update Router A with an SIA-Reply indicating that the network 10.1.1.0/24 is unreachable. Routers A and B will remove the active route from their topology tables. The neighbor relationship between Routers A and B remains intact.

EIGRP Query Range

Limiting the scope of query propagation through the network (the query range), also known as *query scoping*, helps reduce incidences of SIA. Keeping the query packets close to the source reduces the chance that an isolated failure in another part of the network will restrict the convergence (query/reply) process. This section introduces an example that examines how to manage the query range.

Remote routers rarely need to know all the routes advertised in an entire network. Therefore, it is the network manager's responsibility to look at what information is necessary to properly route user traffic and to consider the use of a default route.

For example, in Figure 3-30, Router B notices the loss of network 10.1.8.0 and sends a query to Routers A, C, D, and E. In turn, these routers send queries to their neighbors, requesting an FS for 10.1.8.0. When the query process starts, each path receives duplicate queries because of the redundant topology. Therefore, not only are the remote routers required to respond to queries from the regional offices, but they also continue the search by reflecting the queries back toward the other regional office's router. This significantly complicates the convergence process on the network.

Figure 3-30 *Effect of the EIGRP Update and Query Process*

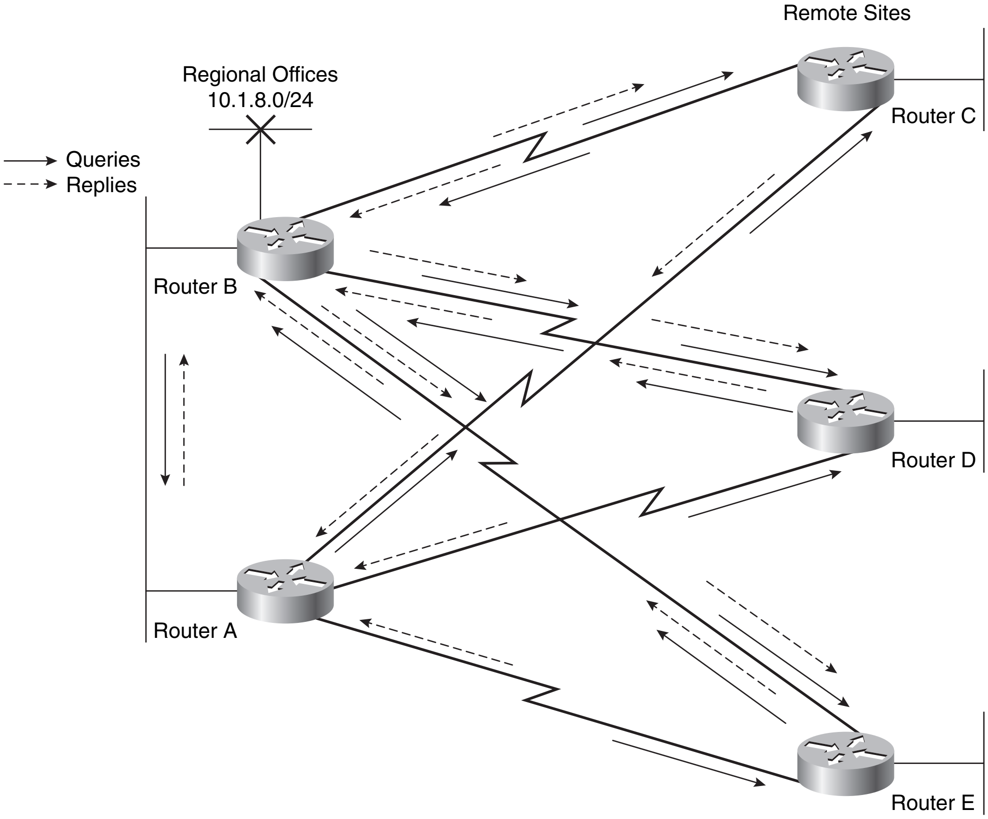

In this sample network with only two regional and three remote routers, the problem might not be very significant. In a network with hundreds of remote offices, the problem can be severe.

Examine the query process for the 10.1.8.0/24 subnet. Router B advertises 10.1.8.0/24 to all other routers. The best path for Router A to reach 10.1.8.0/24 is over the Ethernet link to Router B. The remote routers (C, D, and E) use the serial link to B as their preferred path to reach 10.1.8.0/24 but still learn about an alternative path through Router A. For this example, assume that the EIGRP metric for Ethernet is 1000 and the metric for a serial link is 100,000.

Table 3-5 shows the content of the IP EIGRP topology table on Routers C, D, and E for network 10.1.8.0/24. Table 3-6 shows the content of the IP EIGRP topology table on Router A for network 10.1.8.0/24.

Table 3-5 *IP EIGRP Topology Table for 10.1.8.0/24 on Routers C, D, and E in Figure 3-30*

Neighbor	FD	AD
Router A	102,000	2000
Router B	101,000	1000

Table 3-6 *IP EIGRP Topology Table for 10.1.8.0/24 on Router A in Figure 3-30*

Neighbor	FD	AD
Router B	2000	1000
Router C	201,000	101,000
Router D	201,000	101,000
Router E	201,000	101,000

Note that Routers C, D, and E determine that for network 10.1.8.0/24, Router B is the successor and Router A is an FS (because the AD is 2000 through Router A, which is less than the FD through Router B). Also, note that Router A does not have an FS, because all paths through the remote routers have an AD larger than the FD through Router B.

When Router B loses the path to network 10.1.8.0/24, it queries all four of its neighbors. When the remote sites receive this query, they automatically install the path through Router A in their routing tables and respond to Router B with their supposedly good path through Router A. They also remove the bad path through Router B from their topology tables.

Router B now has responses to three of its four queries, but it must wait until Router A responds as well.

When Router A receives the query from Router B for network 10.1.8.0/24, Router A creates a query and sends it to Routers C, D, and E, because Router A does not have an FS but knows that a path exists through each remote site to reach 10.1.8.0/24.

Routers C, D, and E receive the query from Router A; they now know that their path through Router A is not good, so they check their topology tables for alternative paths. However, none of these routers currently has another path, because Router B has just informed them that it does not have a path to this network. Because the remote routers do not have an answer to the query from Router A, Routers C, D, and E create a query and send it to all neighbors except the neighbor (interface) that these routers received the original query from. In this case, the remote routers send the query only to Router B.

Router B learns from these queries that none of the remote routers has a path to network 10.1.8.0/24, but it cannot respond that it does not know of a path, because Router B is waiting for Router A to reply to a query. Router A is waiting for either Router C, D, or E to reply to its query, and these remote sites are waiting for Router B to reply to their queries. Because Router B sent out the first query, its SIA timer expires first, and Router B reaches the SIA state for network 10.1.8.0/24 first (in 3 minutes by default). Router B resets its neighbor relationship with Router A. As soon as the neighbor relationship goes down, Router B can respond to Routers C, D, and E immediately, saying that Router B does not have a path to 10.1.8.0/24. Routers C, D, and E can then respond to Router A that they do not have a path.

After the EIGRP neighbor relationship between Routers A and B is reestablished (just after the adjacency is reset), Router B, which no longer has a path to 10.1.8.0/24, does not pass the 10.1.8.0/24 network to Router A. Router A learns that the remote sites do not have a path to 10.1.8.0/24, and the new relationship with Router B does not include a path to 10.1.8.0/24, so Router A removes the 10.1.8.0 network from its IP EIGRP topology table.

In Figure 3-30, the network architect provides redundancy with dual links from the regional offices to the remote sites. The architect does not intend for the traffic to go from a regional office to a remote office and back to a regional office, but unfortunately this is the situation. The design of the network shown in Figure 3-30 is sound, but because of EIGRP behavior, remote routers are involved in the convergence process.

If the remote sites are not acting as transit sites between the regional sites, the regional routers can be configured to announce only a default route to the remote routers, and the remote routers can be configured to announce only their directly connected stub network to the regional routers to reduce the complexity and the EIGRP topology table and routing table size.

The following section describes other solutions for limiting the EIGRP query range.

Limiting the EIGRP Query Range

The network manager must determine the information necessary to properly route user traffic to the appropriate destination. The amount of information needed by the remote routers to achieve the desired level of path selection must be balanced against the bandwidth used to propagate this information. To achieve maximum stability and scalability, the remote routers can use a default route to reach the core. If some specific networks need knowledge of more routes to ensure optimum path selection, a business decision is necessary to determine whether the benefits of propagating the additional routing information outweigh the additional bandwidth required to achieve this goal.

In a properly designed network, each remote site has redundant WAN links to separate distribution sites. If both distribution sites pass a default route to the remote sites, the remote sites load balance to all networks behind the distribution site routers. This maximizes bandwidth utilization and allows the remote router to use less CPU and memory, which means that a smaller and less expensive remote router can be used at that site.

If the remote site can see all routes, the router can select a path that is best to reach a given network. However, depending on the number of routes in the internetwork and the amount of bandwidth connecting the remote site to the distribution sites, this approach can mean that higher-bandwidth links or large routers are needed to handle the additional overhead.

After you determine the minimum routing requirements, you can make EIGRP more scalable. Two of the best options are the following:

- Configure route summarization using the **ip summary-address eigrp** command on the outbound interfaces of the appropriate routers.
- Configure the remote routers as stub EIGRP routers.

Other methods to limit query range include route filtering and interface packet filtering. For example, if specific EIGRP routing updates are filtered to a router and that router receives a query about those filtered (blocked) networks, the router indicates that the network is unreachable and does not extend the query any further.

Limiting Query Range with Summarization

One solution to limit the EIGRP query range is to use route summarization.

For example, in Figure 3-31, Router B sends a summary route of 172.30.0.0/16 to Router A. When network 172.30.1.0/24 goes down, Router A receives a query from Router B about that network. Because Router A has received only a summary route, that specific network is not in the routing table. Router A replies to the query with a "network 172.30.1.0/24 unreachable" message and does not extend the query any further.

Figure 3-31 *EIGRP Summarization Can Limit Query Range*

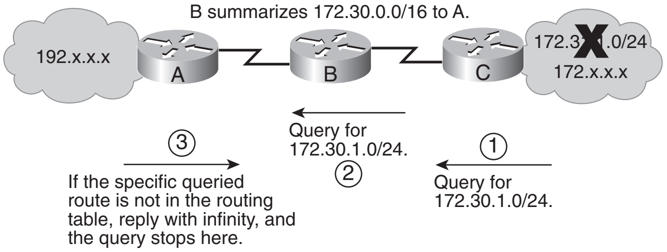

Summarization minimizes the size of the routing table, which means less CPU and memory usage to manage it and less bandwidth to transmit the information. Summarization reduces the chance of networks becoming SIA, because it reduces the number of routers that see each query, so the chance of a query encountering one of these issues is also reduced.

KEY POINT

Query Range

A remote router extends the query about a network only if it has an exact match in the routing table.

Figure 3-32 illustrates how route summarization can affect the network shown in Figure 3-30. The **ip summary-address eigrp** command is configured on the outbound interfaces of Routers A and B so that Routers A and B advertise the 10.0.0.0/8 summary route to the remote Routers C, D, and E.

Figure 3-32 *Limiting Updates and Queries Using Summarization*

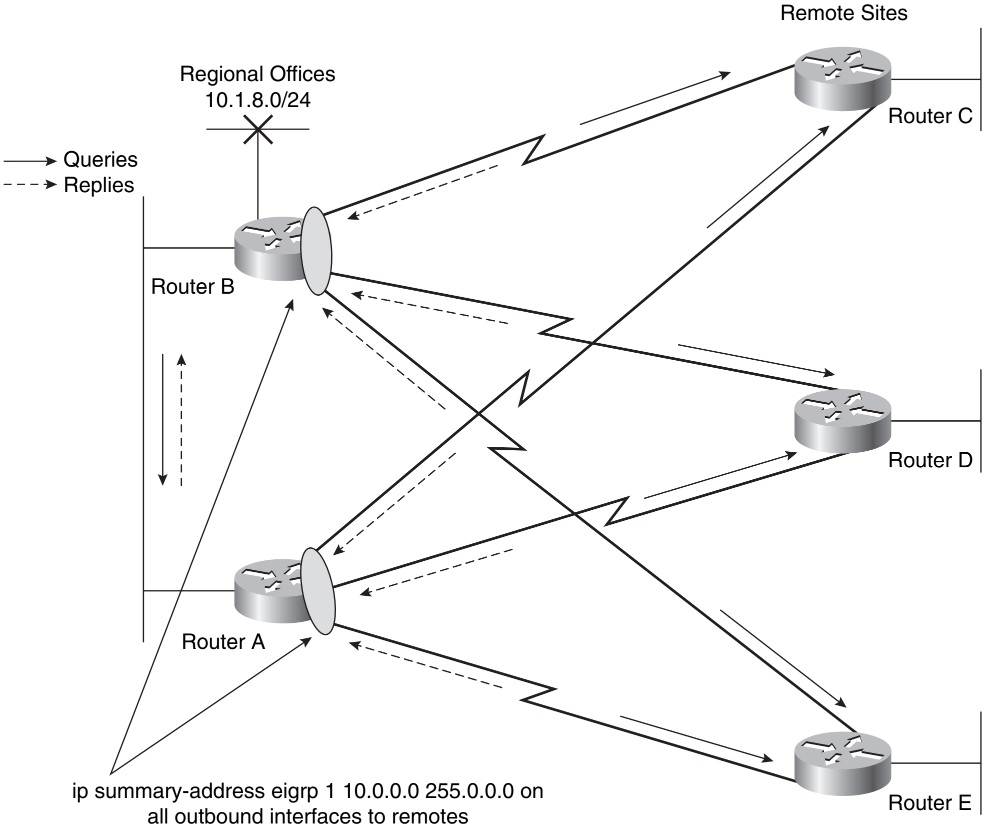

The 10.1.8.0/24 network is not advertised to the remote routers. Therefore, the remote routers (C, D, and E) do not extend the queries about the 10.1.8.0/24 network back to the regional routers (A and B), reducing the convergence traffic (queries and replies) caused by the redundant topology. When Routers A and B send the query for 10.1.8.0/24 to Routers C, D, and E, these routers immediately reply to Routers A and B that the destination is unreachable. Queries for the lost 10.1.8.0/24 networks are not propagated beyond the remote sites, preventing Routers A and B from becoming SIA waiting for the query process to receive all the replies.

Limiting Query Range Using a Stub

Hub-and-spoke network topologies commonly use stub routing. In this topology, the remote router forwards all traffic that is not local to a hub router; the remote router does not need to retain a complete routing table. Generally, the hub router needs to send only a default route to the remote routers.

In a hub-and-spoke topology, having a full routing table on the remote routers serves no functional purpose, because the path to the corporate network and the Internet is always through the hub router. Additionally, having a full routing table at the spoke routers increases the amount of memory required. Route summarization and route filtering can also be used to conserve bandwidth and memory requirements on the spoke routers.

Traffic from a hub router should not use a remote router as a transit path. A typical connection from a hub router to a remote router has significantly less bandwidth than a connection at the network core; attempting to use the connection to a remote router as a transit path typically results in excessive congestion. The EIGRP stub routing feature can prevent this problem by restricting the remote router from advertising the hub router's routes back to other hub routers. For example, routes recognized by the remote router from hub Router A are not advertised to hub Router B. Because the remote router does not advertise the hub routes back to the hub routers, the hub routers do not use the remote routers as a transit path. Using the EIGRP stub routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

KEY POINT

EIGRP Stub

Only the remote routers are configured as stubs. The stub feature does not prevent routes from being advertised to the remote router.

The EIGRP stub feature was first introduced in Cisco IOS Release 12.0(7)T.

A stub router indicates in the hello packet to all neighboring routers its status as a stub router. Any neighbor that receives a packet informing it of the stub status does not query the stub router for any routes. Therefore, a router that has a stub peer does not query that peer.

KEY POINT

EIGRP Stub Routers Are Not Queried

Stub routers are not queried. Instead, hub routers connected to the stub router answer the query on behalf of the stub router.

The EIGRP stub routing feature also simplifies the configuration and maintenance of hub-and-spoke networks. When stub routing is enabled in dual-homed remote configurations, you do not have to configure filtering on remote routers to prevent them from appearing as transit paths to the hub routers.

CAUTION EIGRP stub routing should be used on stub routers only. A stub router is defined as a router connected to the network core or hub layer, and through which core transit traffic should not flow. A stub router should have only hub routers for EIGRP neighbors. Ignoring this restriction causes undesirable behavior.

To configure a router as an EIGRP stub, use the **eigrp stub** [**receive-only** | **connected** | **static** | **summary**] router configuration command. A router configured as a stub with this command shares information about connected and summary routes with all neighbor routers by default. Table 3-7 describes the four optional keywords that can be used with the **eigrp stub** command to modify this behavior.

Table 3-7 **eigrp stub** *Command Parameters*

Parameter	Description
receive-only	The **receive-only** keyword restricts the router from sharing any of its routes with any other router within an EIGRP autonomous system. This keyword does not permit any other keyword to be specified, because it prevents any type of route from being sent. Use this option if there is a single interface on the router.
connected	The **connected** keyword permits the EIGRP stub routing feature to send connected routes. If a **network** command does not include the connected routes, it might be necessary to redistribute connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default and is the most widely practical stub option.
static	The **static** keyword permits the EIGRP stub routing feature to send static routes. Redistributing static routes with the **redistribute static** command is still necessary.
summary	The **summary** keyword permits the EIGRP stub routing feature to send summary routes. You can create summary routes manually with the **ip summary-address eigrp** command or automatically at a major network border router with the **auto-summary** command enabled. This option is enabled by default.

The optional parameters in this command can be used in any combination, with the exception of the **receive-only** keyword. If any of the keywords (except **receive-only**) is used individually, the connected and summary routes are not sent automatically.

In Example 3-21, the **eigrp stub** command is used to configure the router as a stub that advertises connected and summary routes.

Example 3-21 **eigrp stub** *Command to Advertise Connected and Summary Routes*

```
Router(config)#router eigrp 1
Router(config-router)#network 10.0.0.0
Router(config-router)#eigrp stub
```

In Example 3-22, the **eigrp stub receive-only** command is used to configure the router as a stub. Connected, summary, or static routes are not sent.

Example 3-22 **eigrp stub** *Command to Receive Only Routes*

```
Router(config)#router eigrp 1
Router(config-router)#network 10.0.0.0 eigrp
Router(config-router)#eigrp stub receive-only
```

The EIGRP stub feature does not automatically enable route summarization on the hub router. The network administrator should configure route summarization on the hub routers if desired.

If a true stub network is required, the hub router can be configured to send a default route to the spoke routers. This approach is the most simple and conserves the most bandwidth and memory on the spoke routers.

> **NOTE** Although EIGRP is a classless routing protocol, it has classful behavior by default, such as having automatic summarization on by default. When you configure the hub router to send a default route to the remote router, ensure that the **ip classless** command on the remote router. By default, the **ip classless** command is enabled in all Cisco IOS images that support the EIGRP stub routing feature.

Without the stub feature, EIGRP sends a query to the spoke routers if a route is lost somewhere in the network. If there is a communication problem over a WAN link between the hub router and a spoke router, an EIGRP SIA condition can occur and cause instability elsewhere in the network. The EIGRP stub routing feature allows a network administrator to prevent sending queries to the spoke router under any condition. Cisco highly recommends using both EIGRP route summarization and EIGRP stub features to provide the best scalability.

Figure 3-33 illustrates how using the EIGRP stub feature affects the network shown in Figure 3-30. Each of the remote routers is configured as a stub. Queries for network 10.1.8.0/24 are not sent to Routers C, D, or E, thus reducing the bandwidth used and the chance of the routes being stuck-in-active.

Figure 3-33 *Limiting Updates and Queries Using the EIGRP Stub Feature*

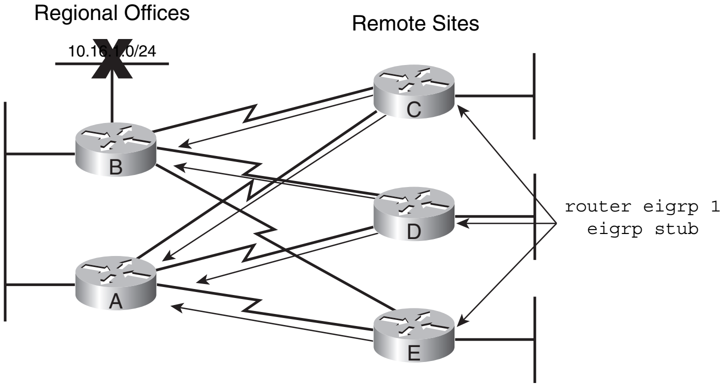

Using the EIGRP stub feature at the remote sites allows the hub (regional offices) sites to immediately answer queries without propagating the queries to the remote sites, saving CPU cycles and bandwidth, and lessening convergence time even when the remote sites are dual-homed to two or more hub sites.

Figure 3-34 illustrates another example network; Example 3-23 shows part of the configuration for Router B.

Figure 3-34 *Network for* **eigrp stub** *Command Example*

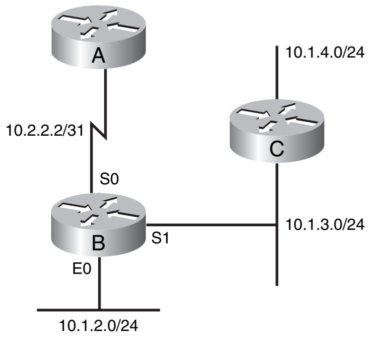

Example 3-23 *Configuration for Router B in Figure 3-34*

```
RouterB#show running-config
<output omitted>
ip route 10.1.4.0 255.255.255.0 10.1.3.10
!
interface ethernet 0
 ip address 10.1.2.1 255.255.255.0
!
interface serial 0
 ip address 10.2.2.3 255.255.255.254
 ip summary-address eigrp 100 10.1.2.0 255.255.254.0
!
interface serial 1
 ip address 10.1.3.1 255.255.255.0
!
router eigrp 100
 redistribute static 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.1
 network 10.1.2.0 0.0.0.255
<output omitted>
```

Using this example network and configuration, consider which networks will be advertised when the various options of the **eigrp stub** command are also configured on Router B:

- With the **eigrp stub connected** command, Router B will advertise only 10.1.2.0/24 to Router A. Notice that although 10.1.3.0/24 is also a connected network, it is not advertised to Router A because it is not advertised in a **network** command, and connected routes are not redistributed.

- With the **eigrp stub summary** command, Router B will advertise only 10.1.2.0/23, the summary route that is configured on the router, to Router A.
- With the **eigrp stub static** command, router B will advertise only 10.1.4.0/24, the static route that is configured on the router, to Router A.
- With the **eigrp stub receive-only** command, Router B will not advertise anything to Router A.

Graceful Shutdown

Graceful shutdown, implemented with the *goodbye message* feature, is designed to improve EIGRP network convergence.

In Figure 3-35, Router A is using Router B as the successor for a number of routes; Router C is the feasible successor for the same routes. Router B normally would not tell Router A if the EIGRP process on Router B was going down, for example, if Router B was being reconfigured. Router A would have to wait for its hold timer to expire before it would discover the change and react to it. Packets sent during this time would be lost.

Figure 3-35 *Graceful Shutdown Causes Router to Say Goodbye*

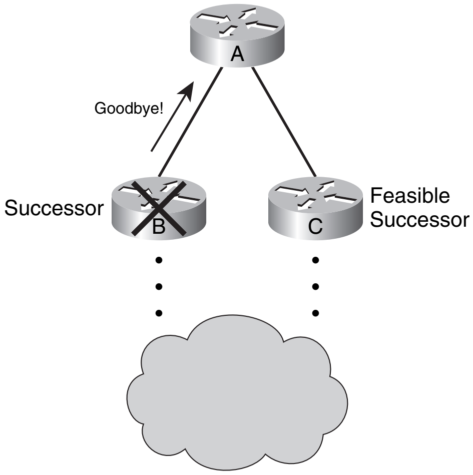

With graceful shutdown, a goodbye message is broadcast when an EIGRP routing process is shut down, to inform adjacent peers about the impending topology change. This feature allows supporting EIGRP peers to synchronize and recalculate neighbor relationships more efficiently than would occur if the peers discovered the topology change after the hold timer expired.

The goodbye message is supported in Cisco IOS Software Release 12.3(2), 12.3(3)B, and 12.3(2)T and later.

KEY POINT

Goodbye Messages

Goodbye messages are sent in hello packets.

EIGRP sends an interface goodbye message with all K values set to 255 when taking down all peers on an interface.

The following message is displayed by routers that support goodbye messages when one is received:

```
*Apr 26 13:48:42.523: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 10.1.1.1 (Ethernet0/0)
  is down: Interface Goodbye received
```

A Cisco router that runs a software release that does not support the goodbye message will misinterpret the message as a K-value mismatch and therefore display the following message:

```
*Apr 26 13:48:41.811: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 10.1.1.1 (Ethernet0/0)
  is down: K-value mismatch
```

NOTE The receipt of a goodbye message by a nonsupporting peer does not disrupt normal network operation. The nonsupporting peer will terminate the session when the hold timer expires. The sending and receiving routers will reconverge normally after the sender reloads.

NOTE An EIGRP router will send a goodbye message on an interface if the **network** command (under the EIGRP process) that encompasses the network on that interface is removed (with the **no network** command). An EIGRP router sends a goodbye message on all interfaces if the EIGRP process is shut down (with the **no router eigrp** command). An EIGRP router will not, however, send a goodbye message if an interface is shut down or the router is reloaded.

Verifying EIGRP Operation

This section discusses commands used to verify EIGRP operation.

Table 3-8 describes some **show** commands used to verify EIGRP operation. Other options might be available with these commands; use the Cisco IOS integrated help feature to see the full-command syntax.

Table 3-8 *EIGRP* **show** *Commands*

Command	Description
show ip eigrp neighbors	Displays neighbors discovered by EIGRP.
show ip route	Displays the current entries in the IP routing table for all configured routing protocols.
show ip route eigrp	Displays the current EIGRP entries in the IP routing table.

Table 3-8 *EIGRP* **show** *Commands (Continued)*

Command	Description
show ip protocols	Displays the parameters and current state of the active routing protocol processes. For EIGRP, this command shows the EIGRP autonomous system number, filtering and redistribution numbers, and neighbors and distance information.
show ip eigrp interfaces	Displays information about interfaces configured for 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 FD to the destination. Note that only successor and feasible successor routes are displayed; add the **all-links** keyword to display all routes, including those not eligible to be successor or feasible successor routes.
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.

Table 3-9 describes **debug** commands used to verify EIGRP operation. Other options might be available with these commands; use the Cisco IOS integrated help feature to see the full command syntax.

Table 3-9 *EIGRP* **debug** *Commands*

Command	Description
debug eigrp packets	Displays the types of EIGRP packets sent and received. A maximum of 11 packet types can be selected for individual or group display.
debug ip eigrp	Displays packets that are sent and received on an interface. Because this command generates large amounts of output, use it only when traffic on the network is light.
debug ip eigrp summary	Displays a summarized version of EIGRP activity. It also displays filtering and redistribution numbers and neighbors and distance information.
debug eigrp neighbors	Displays neighbors discovered by EIGRP and the contents of the hello packets.

The following sections provide sample output from some of these commands, using the network in Figure 3-36 to illustrate the configuration, verification, and troubleshooting of EIGRP. Example 3-24 shows the configuration of the R1 router.

Figure 3-36 *Example Network for EIGRP Verification*

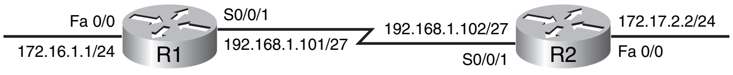

Example 3-24 *Configuration for Router R1 in Figure 3-36*

```
R1#show running-config
<output omitted>
interface FastEthernet0/0
  ip address 172.16.1.1 255.255.255.0

<output omitted>
interface Serial0/0/1
 bandwidth 64
 ip address 192.168.1.101 255.255.255.224

<output omitted>
router eigrp 100
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

On the R1 router, EIGRP is enabled in autonomous system 100. The **network 172.16.1.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R1 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.16.1.0/24 subnet will participate in EIGRP. Note, however, the full Class B network 172.16.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the serial 0/0/1 interface, and allows router R1 to advertise this network.

Example 3-25 shows the configuration of the R2 router.

Example 3-25 *Configuration for Router R2 in Figure 3-36*

```
R2#show running-config
<output omitted>
interface FastEthernet0/0
  ip address 172.17.2.2 255.255.255.0

<output omitted>
interface Serial0/0/1
 bandwidth 64
 ip address 192.168.1.102 255.255.255.224

<output omitted>
router eigrp 100
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
```

EIGRP is also enabled in autonomous system 100 on the R2 router. The **network 172.17.2.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R2 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.17.2.0/24 subnet will participate in EIGRP. Note, however, the full Class B network 172.17.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the serial 0/0/1 interface and allows router R2 to advertise this network.

> **NOTE** The "EIGRP Neighbors" section, earlier in this chapter, provides output from the **show ip eigrp neighbors** command and a description of the output.

show ip route and show ip route eigrp for EIGRP Examples

To verify that the router recognizes EIGRP routes for any neighbors, use the **show ip route eigrp** command, as shown in Example 3-26. Example 3-27 exhibits the **show ip route** command, which displays the full IP routing table, including the EIGRP routes.

Example 3-26 **show ip route eigrp** *Command Output*

```
R1#show ip route eigrp
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:07:01, Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:05:13, Null0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
D       192.168.1.0/24 is a summary, 00:05:13, Null0
```

Example 3-27 **show ip route** *Command Output*

```
R1#show ip route
<output omitted>
Gateway of last resort is not set
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:06:55, Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:05:07, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.96/27 is directly connected, Serial0/0/1
D       192.168.1.0/24 is a summary, 00:05:07, Null0
```

Using the highlighted line in Example 3-26 as an example, the fields in the routing table are interpreted as follows:

- Internal EIGRP routes are identified with a D in the leftmost column. (External EIGRP routes, not shown in this example, are identified with a D EX in the leftmost column.)
- The next column is the network number (172.17.0.0/16 in this example).
- After each network number is a field in brackets (90/40514560 in this example). The second number in brackets is the EIGRP metric. As discussed in the "EIGRP Metric Calculation"

section earlier in this chapter, the default EIGRP metric is the least-cost bandwidth plus the accumulated delays. The EIGRP metric for a network is the same as its FD in the EIGRP topology table.

The first number, 90 in this case, is the administrative distance. Recall from Chapter 2 that administrative distance is used to select the best path when a router learns two or more routes to exactly the same destination from different routing sources. For example, consider that this router uses RIP and EIGRP and that RIP has a route to network 172.17.0.0 that is three hops away. The router, without the administrative distance, cannot compare three hops to an EIGRP metric of 40,514,560. The router does not know the bandwidth associated with hops, and EIGRP does not use hop count as a metric.

To correct this problem, Cisco established an administrative distance for each routing protocol: the lower the value, the more preferred the route is. By default, EIGRP internal routes have an administrative distance of 90, and RIP has an administrative distance of 120. Because EIGRP has a metric based on bandwidth and delays, it is preferred over RIP's hop count metric. As a result, in this example, the EIGRP route would be installed in the routing table.

NOTE Remember that routers use the administrative distance only if the two routes are to the exact same destination (address and mask); for example, a router will choose a RIP route over an EIGRP route if the RIP route is a more specific route than the EIGRP route.

- The next field, via 192.168.1.102 in this example, is the address of the next-hop router to which this router passes packets destined for 172.17.0.0/16. The next-hop address in the routing table is the same as the successor in the EIGRP topology table.
- The route also has a time associated with it (00:07:01 in this example); this is the length of time since EIGRP last advertised this network to this router. EIGRP does not refresh routes periodically; it resends the routing table only when neighbor adjacencies change.
- The interface, serial 0/0/1 in this case, indicates the interface out which packets for 172.17.0.0 are sent.

Notice that the routing table includes routes, to null0, for the advertised (summarized) routes. Cisco IOS Software automatically puts these routes in the table; they are called *summary routes*. Null 0 is a directly connected, software-only interface. The use of the null0 interface prevents the router from trying to forward traffic to other routers in search of a more precise, longer match. For example, if the R1 router in Figure 3-36 receives a packet to an unknown subnet that is part of the summarized range—172.16.3.5 for example—the packet matches the summary route based on the longest match. The packet is forwarded to the null0 interface (in other words, it is dropped, or sent to the *bit bucket*), which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

show ip protocols Example

Use the **show ip protocols** command to provide information about any and all dynamic routing protocols running on the router.

As shown in Example 3-28, the command output displays any route filtering occurring on EIGRP outbound or inbound updates. It also identifies whether EIGRP is generating a default network or receiving a default network in EIGRP updates and provides information about additional settings for EIGRP, such as default K values, hop count, and variance.

Example 3-28 **show ip protocols** *Command Output*

```
R1#show ip protocols
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is in effect
  Automatic address summarization:
    192.168.1.0/24 for FastEthernet0/0
      Summarizing with metric 40512000
    172.16.0.0/16 for Serial0/0/1
      Summarizing with metric 28160
  Maximum path: 4
  Routing for Networks:
    172.16.1.0/24
    192.168.1.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    (this router)         90      00:09:38
    Gateway         Distance      Last Update
    192.168.1.102         90      00:09:40
Distance: internal 90 external 170
```

NOTE Because the routers must have identical K values for EIGRP to establish an adjacency, the **show ip protocols** command helps determine the current K-value setting before an adjacency is attempted.

The output in Example 3-28 also indicates that automatic summarization is enabled (this is the default) and that the router is allowed to load-balance over a maximum of four paths. Cisco IOS

Software allows configuration of up to 16 paths for equal-cost load balancing, using the **maximum-paths** router configuration command.

The networks for which the router is routing are also displayed. As shown in Example 3-28, the format of the output varies, depending on the use of the wildcard mask in the **network** command. If a wildcard mask is used, the network address is displayed with a prefix length. If a wildcard mask is not used, the Class A, B, or C major network is displayed.

The routing information source portion of this command output identifies all other routers that have an EIGRP neighbor relationship with this router. The **show ip eigrp neighbors** command provides a detailed display of EIGRP neighbors.

The **show ip protocols** command output also provides the two administrative distances for EIGRP. An administrative distance of 90 applies to networks from other routers inside the same autonomous system number; these are considered internal networks. An administrative distance of 170 applies to networks introduced to EIGRP for this autonomous system through redistribution; these are called external networks.

show ip eigrp interfaces Example

Example 3-29 demonstrates **show ip eigrp interfaces** command output.

Example 3-29 **show ip eigrp interfaces** *Command Output*

```
R1#show ip eigrp interfaces
IP-EIGRP interfaces for process 100
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0              0        0/0         0       0/10           0          0
Se0/0/1            1        0/0        10      10/380        424          0
```

The **show ip eigrp interfaces** command displays information about interfaces configured for EIGRP. This output includes the following key elements:

- **Interface**—Interface over which EIGRP is configured
- **Peers**—Number of directly connected EIGRP neighbors
- **Xmit Queue Un/Reliable**—Number of packets remaining in the Unreliable and Reliable transmit queues
- **Mean SRTT**—Mean SRTT interval, in milliseconds
- **Pacing Time Un/Reliable**—Pacing time used to determine when EIGRP packets should be sent out the interface (for unreliable and reliable packets)

- **Multicast Flow Timer**—Maximum number of seconds that the router will wait for an ACK packet after sending a multicast EIGRP packet, before switching from multicast to unicast
- **Pending Routes**—Number of routes in the packets in the transmit queue waiting to be sent

show ip eigrp topology Example

Another command used to verify EIGRP operations is the **show ip eigrp topology** command; Example 3-30 demonstrates output generated from this command.

Example 3-30 **show ip eigrp topology** *Command Output*

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(192.168.1.101)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 192.168.1.96/27, 1 successors, FD is 40512000
        via Connected, Serial0/0/1
P 192.168.1.0/24, 1 successors, FD is 40512000
        via Summary (40512000/0), Null0
P 172.16.0.0/16, 1 successors, FD is 28160
        via Summary (28160/0), Null0
P 172.16.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 172.17.0.0/16, 1 successors, FD is 40514560
        via 192.168.1.102 (40514560/28160), Serial0/0/1
```

The command output illustrates that router R1 has an ID of 192.168.1.101 and is in autonomous system 100. The EIGRP ID is the highest IP address on an active interface for this router.

The command output also lists the networks known by this router through the EIGRP routing process. The codes used in the first column of this output are as follows:

- **Passive (P)**—This network is available, and installation can occur in the routing table. Passive is the correct state for a stable network.
- **Active (A)**—This network is currently unavailable, and installation cannot occur in the routing table. Being active means that outstanding queries exist for this network.
- **Update (U)**—This network is being updated (placed in an update packet). This code also applies if the router is waiting for an acknowledgment for this update packet.
- **Query (Q)**—There is an outstanding query packet for this network other than being in the active state. This code also applies if the router is waiting for an acknowledgment for a query packet.
- **Reply (R)**—The router is generating a reply for this network or is waiting for an acknowledgment for the reply packet.
- **Stuck-in-active (S)**—There is an EIGRP convergence problem for this network.

The number of successors available for a route is indicated in the command output. The number of successors corresponds to the number of best routes with equal cost; all networks in Example 3-30 have one successor.

For each network, the FD is listed next, followed by an indication of how the route was learned, such as the next-hop address if the route was learned via another router. Next is a field in brackets. The first number in the brackets is the FD for that network through the next-hop router, and the second number in the brackets is the AD from the next-hop router to the destination network.

show ip eigrp traffic Example

To display the number of various EIGRP packets sent and received, use the **show ip eigrp traffic** command, as illustrated in Example 3-31. For example, in this network, router R1 has sent 429 hello messages and received 192 hello messages.

Example 3-31 **show ip eigrp traffic** *Command Output*

```
R1#show ip eigrp traffic
IP-EIGRP Traffic Statistics for AS 100
  Hellos sent/received: 429/192
  Updates sent/received: 4/4
  Queries sent/received: 1/0
  Replies sent/received: 0/1
  Acks sent/received: 4/3
  Input queue high water mark 1, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 113
PDM Process ID: 73
```

debug eigrp packets Examples

You can use the **debug eigrp packets** command to verify EIGRP connectivity. This command displays the types of EIGRP packets sent and received by the router that this command is executed on. Different packet types can be selected for individual or group display. Example 3-32 shows some output from this command on R2, when an interface on R1 comes up.

Example 3-32 **debug eigrp packets** *Command Output on R2 When a Neighbor's Interface Comes Up*

```
R2#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
*May 11 04:02:55.821: EIGRP: Sending HELLO on Serial0/0/1
*May 11 04:02:55.821:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
R2#
*May 11 04:02:58.309: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.101
```

Example 3-32 **debug eigrp packets** *Command Output on R2 When a Neighbor's Interface Comes Up (Continued)*

```
*May 11 04:02:58.309:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
*May 11 04:02:58.585: EIGRP: Sending HELLO on FastEthernet0/0
*May 11 04:02:58.585:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
*May 11 04:02:59.093: EIGRP: Received UPDATE on Serial0/0/1 nbr 192.168.1.101
*May 11 04:02:59.093:   AS 100, Flags 0x0, Seq 5/4 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
*May 11 04:02:59.093: EIGRP: Enqueueing ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:02:59.093:   Ack seq 5 iidbQ un/rely 0/0 peerQ un/rely 1/0
*May 11 04:02:59.097: EIGRP: Sending ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:02:59.097:   AS 100, Flags 0x0, Seq 0/5 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 1/0
*May 11 04:02:59.109: EIGRP: Enqueueing UPDATE on Serial0/0/1 iidbQ un/rely 0/1 serno 9-9
*May 11 04:02:59.113: EIGRP: Enqueueing UPDATE on Serial0/0/1 nbr 192.168.1.101 iidbQ un/
  rely 0/0 peerQ un/rely 0/0 serno 9-9
*May 11 04:02:59.113: EIGRP: Sending UPDATE on Serial0/0/1 nbr 192.168.1.101
*May 11 04:02:59.113:   AS 100, Flags 0x0, Seq 5/5 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/1 serno 9-9
*May 11 04:02:59.133: EIGRP: Received ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:02:59.133:   AS 100, Flags 0x0, Seq 0/5 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/1
*May 11 04:02:59.133: EIGRP: Serial0/0/1 multicast flow blocking cleared
R2#
*May 11 04:03:00.441: EIGRP: Sending HELLO on Serial0/0/1
*May 11 04:03:00.441:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
R2#
*May 11 04:03:03.209: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.101
*May 11 04:03:03.209:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
```

The **debug eigrp packets** command traces transmission and receipt of EIGRP packets. The output in Example 3-32 shows normal transmission and receipt of EIGRP packets. The serial link is an HDLC point-to-point link; therefore, the default hello time interval is 5 seconds. Hello packets are sent unreliably, so the sequence number (Seq) does not increment.

In this sample output, when R2 receives an update from R1, values appear in the sequence number field. Seq 5/4 indicates that 192.168.1.101 is sending this packet as sequence number 5 to R2 and that sequence number 4 has been received from R2 by neighbor 192.168.1.101. 192.168.1.101 is expecting to receive sequence number 5 in the next reliable packet from R2.

R2 returns an ACK packet with Seq 0/5. The acknowledgment is sent as an unreliable packet, but the neighbor unreliable/reliable flag (un/rel 1/0) is set. This means that the acknowledgment was sent in response to a reliable packet.

The serial number (serno 9-9) reflects the number of changes that the two neighbors register in their EIGRP topology tables. A single update can contain more than 100 networks that all produce an update, because all are now unavailable.

KEY POINT

Sequence Number Versus Serial Number

The sequence number increments each time a query, update, or reply packet is sent, whereas the serial number increments each time the topology table changes. Therefore, if the topology table has more than 100 changes, the serial number increases substantially, but the sequence number may only increase by 1.

When an interface on R1 (R2's EIGRP neighbor 192.168.1.101) is shut down, the resulting output on R2 is shown in Example 3-33. R1 sends a query packet to R2 to determine whether R2 knows a path to the lost network. R2 responds with an ACK packet to acknowledge the query packet; a reliable packet must be explicitly acknowledged with an ACK packet. R2 also responds to the query with a reply packet. The serial number reference (10-12) represents the number of changes to the topology table since the start of the neighbor relationship between these two EIGRP neighbors.

Example 3-33 **debug eigrp packets** *Command Output on R2 When a Neighbor's Interface Is Shut Down*

```
R2#debug eigrp packets
*May 11 04:20:43.361: EIGRP: Received QUERY on Serial0/0/1 nbr 192.168.1.101
*May 11 04:20:43.361:   AS 100, Flags 0x0, Seq 6/5 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/0
*May 11 04:20:43.361: EIGRP: Enqueueing ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:20:43.361:   Ack seq 6 iidbQ un/rely 0/0 peerQ un/rely 1/0
*May 11 04:20:43.365: EIGRP: Sending ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:20:43.365:   AS 100, Flags 0x0, Seq 0/6 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 1/0
*May 11 04:20:43.373: EIGRP: Enqueueing REPLY on Serial0/0/1 nbr 192.168.1.101 iidbQ un/
  rely 0/1 peerQ un/rely 0/0 serno 10-12
*May 11 04:20:43.377: EIGRP: Requeued unicast on Serial0/0/1
R2#
*May 11 04:20:43.381: EIGRP: Sending REPLY on Serial0/0/1 nbr 192.168.1.101
*May 11 04:20:43.381:   AS 100, Flags 0x0, Seq 6/6 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/1 serno 10-12
*May 11 04:20:43.405: EIGRP: Received ACK on Serial0/0/1 nbr 192.168.1.101
*May 11 04:20:43.405:   AS 100, Flags 0x0, Seq 0/6 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/
  rely 0/1
```

debug ip eigrp Examples

You can use the **debug ip eigrp** command to verify EIGRP operation. This command displays EIGRP packets that this router sends and receives. Example 3-34 shows the contents of the updates that are reported when you use the **debug ip eigrp** command on R2 to monitor when an interface on R1 comes up.

Example 3-34 **debug ip eigrp** *Command Output on R2 When a Neighbor's Interface Comes Up*

```
R2#debug ip eigrp
IP-EIGRP Route Events debugging is on
R2#
*May 11 04:24:05.261: IP-EIGRP(Default-IP-Routing-Table:100): Processing incoming UPDATE
  packet
```

Example 3-34 **debug ip eigrp** *Command Output on R2 When a Neighbor's Interface Comes Up (Continued)*

```
*May 11 04:24:05.261: IP-EIGRP(Default-IP-Routing-Table:100): Int 192.168.1.0/24
 M 4294967295 - 40000000 4294967295 SM 4294967295 - 40000000 4294967295
*May 11 04:24:05.261: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.0.0/16
M 40514560 - 40000000 514560 SM 28160 - 25600 2560
*May 11 04:24:05.261: IP-EIGRP(Default-IP-Routing-Table:100): route installed for
172.16.0.0  ()
*May 11 04:24:05.277: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.0.0/16 metric
40514560 - 40000000 514560
```

In this example, an internal route (indicated by Int) for 172.16.0.0/16 is advertised to R2.

Recall that by default the EIGRP metric is equal to the bandwidth plus the delay. The EIGRP process uses the source metric (SM) information in the update to calculate the AD and place it in the EIGRP topology table. In this example, the SM information is SM 28160 – 25600 2560, which means the source metric (the AD) = 28160 = 25600 (the bandwidth) + 2560 (the delay).

The EIGRP metric calculation for the total delay uses the metric (M) information in the update. In this example, the M information is M 40514560 – 40000000 514560, which means the metric (the FD) = 40514560 = 40000000 (the bandwidth) + 514560 (the delay).

The EIGRP metric for this route is equal to the FD and, therefore, is 40,514,560.

Example 3-35 illustrates what occurs when R2 processes an incoming query packet for network 172.16.0.0/16 when the interface on the neighboring router (R1) that leads to that network is shut down. Note that comments (preceded by an exclamation point [!]) have been added to this output for easier understanding.

Example 3-35 **debug ip eigrp** *Command Output*

```
R2#debug ip eigrp
IP-EIGRP Route Events debugging is on
R2#
! An interface on EIGRP neighbor R1 was shutdown
! R2 receives a query looking for a lost path from R1
*May 11 04:35:44.281: IP-EIGRP(Default-IP-Routing-Table:100): Processing incoming QUERY
  packet
*May 11 04:35:44.281: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.1.0/24
M 4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
*May 11 04:35:44.281: IP-EIGRP(Default-IP-Routing-Table:100): Int 192.168.1.0/24
 M 4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
*May 11 04:35:44.281: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.0.0/16
M 4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
! R2 realizes that if it cannot use R1 for this network then
! it does not have an entry in the routing table for this network
```

continues

Example 3-35 **debug ip eigrp** *Command Output (Continued)*

```
*May 11 04:35:44.281: IP-EIGRP(Default-IP-Routing-Table:100): 172.16.0.0/16 routing table
  not updated thru 192.168.1.101
R2#
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): 172.16.1.0/24 - not in IP
  routing table
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.1.0/24 metric
  4294967295 - 0 4294967295
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): 192.168.1.0/24 - poison
  advertise out Serial0/0/1
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): Int 192.168.1.0/24 metric
  40512000 - 40000000 512000
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): 172.16.0.0/16 - not in IP
  routing table
! R2 sends an update to R1 saying it does not know how to reach that network either
*May 11 04:35:44.301: IP-EIGRP(Default-IP-Routing-Table:100): Int 172.16.0.0/16 metric
  4294967295 - 40000000 4294967295
R2#
```

The neighbor previously advertised 172.16.0.0/16 to this router. The query performs the following two functions:

- R2 discovers that its neighbor no longer knows how to get to network 172.16.0.0/16. The metric value (4,294,967,295) is the highest possible value; it indicates that the route is unreachable. R2 removes this entry from the EIGRP topology table and looks for alternative EIGRP routes.

- The debug output indicates that the routing table is not updated; this means that EIGRP did not find an alternative route to the network. The next statement verifies that the EIGRP process has removed the old route and that the route is not in the IP routing table. R2 then informs the neighbor that it does not have a path to this network either.

Summary

In this chapter, you learned about Cisco's own EIGRP, an advanced distance vector routing protocol. The chapter presented the following topics:

- Features of EIGRP, including fast convergence, VLSM support, use of partial updates, multiple network layer support, seamless connectivity across all data link layer protocols and topologies, sophisticated metric, and use of multicast and unicast.

- EIGRP's underlying processes and technologies—neighbor discovery/recovery mechanism, RTP, DUAL finite state machine, and protocol-dependent modules.

- EIGRP terminology, including EIGRP's tables—neighbor table, topology table, and routing table; the advertised distance and the feasible distance; and the successor and feasible successor.

- The five EIGRP packet types: hello, update, query, reply, and acknowledgment.

- Passive and active routes.
- The EIGRP metric calculation, which defaults to bandwidth + delay.
- Basic EIGRP configuration commands.
- EIGRP summarization, EIGRP equal-cost and unequal-cost load balancing, and EIGRP operation in WAN environments.
- Configuring, verifying, and troubleshooting EIGRP MD5 authentication.
- EIGRP scalability factors and EIGRP use in an enterprise network.
- Verifying and troubleshooting EIGRP.

References

For additional information, refer to the following resources:

- The EIGRP protocol home page, http://www.cisco.com/go/eigrp
- The "IGRP Metric" document at http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009405c.shtml (a good reference for the "EIGRP Metric Calculation" section)

Configuration Exercise: Configuring and Tuning EIGRP

In this exercise, you first configure EIGRP and investigate its default behavior. You next configure EIGRP summarization, a stub, and a default route.

Introduction to the Configuration Exercises

This book uses Configuration Exercises to help you practice configuring routers with the commands and topics presented. If you have access to real hardware, you can try these exercises on your routers. See Appendix B, "Configuration Exercise Equipment Requirements and Backbone Configurations," for a list of recommended equipment and initial configuration commands for the backbone routers. However, even if you do not have access to any routers, you can go through the exercises, and keep a log of your own running configurations, or just read through the solution. Commands used and solutions to the Configuration Exercises are provided within the exercises.

In the Configuration Exercises, the network is assumed to consist of two pods, each with four routers. The pods are interconnected to a backbone. You configure pod 1. No interaction between the two pods is required, but you might see some routes from the other pod in your routing tables in some exercises if you have it configured. In most of the exercises, the backbone has only one router; in some cases, another router is added to the backbone. Each Configuration Exercise assumes that you have completed the previous chapters' Configuration Exercises on your pod.

NOTE Throughout this exercise, the pod number is referred to as *x*, and the router number is referred to as *y*. Substitute the appropriate numbers as needed.

Exercise Objectives

The objectives of this exercise are as follows:

- Set up EIGRP
- Investigate the default behavior of EIGRP
- Optimize the EIGRP configuration

Visual Objective

Figure 3-37 illustrates the topology used and what you will accomplish in this exercise.

Figure 3-37 *EIGRP Configuration Exercise Topology*

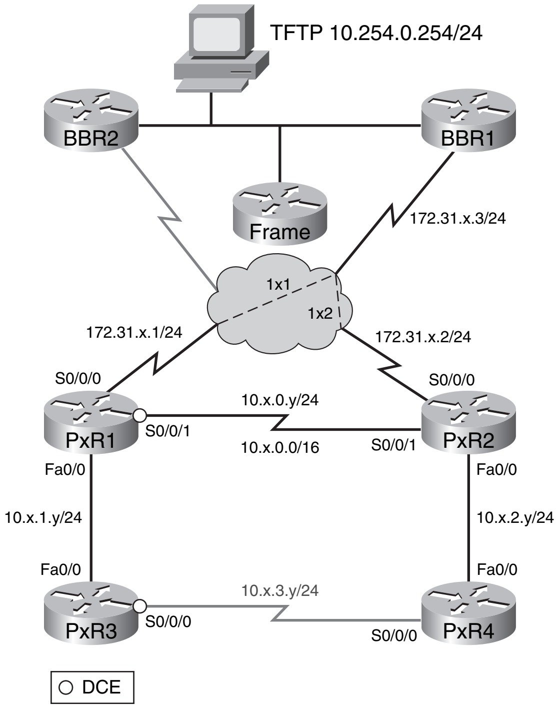

Command List

In this exercise, you use the commands in Table 3-10, listed in logical order. Refer to this list if you need configuration command assistance during the exercise.

CAUTION Although the command syntax is shown in this table, the addresses shown are typically for the P*x*R1 and P*x*R3 routers. Be careful when addressing your routers! Refer to the exercise instructions and the appropriate visual objective diagram for addressing details.

Table 3-10 *EIGRP Configuration Exercise Commands*

Command	Description
(config)#**router eigrp 1**	Enters configuration mode for EIGRP in autonomous system 1
(config-router)#**network 10.*x*.0.0 0.0.255.255**	Specifies that EIGRP should run within network 10.*x*.0.0/16
(config-router)#**no auto-summary**	Turns off automatic summarization at classful network boundaries
#show ip protocols	Displays the parameters and current state of all the active routing protocol processes
#debug ip eigrp	Displays EIGRP updates
(config-if)#**ip summary-address eigrp 1 10.*x*.0.0 255.255.0.0**	Creates and advertises a summary route 10.*x*.0.0/16 for EIGRP autonomous system 1 out of this interface
(config-router)#**eigrp stub**	Specifies that this router should behave as an EIGRP stub router
#show ip eigrp neighbors detail	Displays detailed EIGRP neighbor information
(config-if)#**ip summary-address eigrp 1 0.0.0.0 0.0.0.0**	Creates and advertises a default route for EIGRP autonomous system 1 out of this interface and suppresses all other specific routes
#show ip eigrp topology	Displays the EIGRP topology table
#show ip eigrp traffic	Displays EIGRP traffic statistics
#show ip eigrp interfaces	Displays information about interfaces configured for EIGRP
#show ip eigrp neighbors	Displays EIGRP neighbor information

NOTE The exercise tasks include answers and solutions. Some answers cover multiple steps; the answers are given after the last step to which that answer applies.

Task 1: Configuring Basic EIGRP

In this task, you configure EIGRP on each router in your pod so that there are EIGRP routes from the core, between edge routers, and between the edge and the internal routers. Follow these steps:

Step 1 Shut down the serial interface between the internal routers (s0/0/0 on PxR3 and PxR4); this link is not used in this exercise**.**

Solution:

The following shows the required step on the P1R3 router:

```
P1R3(config)#int s0/0/0
P1R3(config-if)#shutdown
```

Step 2 Configure EIGRP on each router in your pod in autonomous system 1, using the appropriate network and wildcard values to include all interfaces in the EIGRP routing process. Disable autosummarization on the edge routers.

Solution:

The following shows the required steps on the P1R1 and P1R3 routers:

```
P1R1(config)#router eigrp 1
P1R1(config-router)#network 10.1.0.0 0.0.255.255
P1R1(config-router)#network 172.31.1.0 0.0.0.255
P1R1(config-router)#no auto-summary

P1R3(config-if)#router eigrp 1
P1R3(config-router)#network 10.1.0.0 0.0.255.255
```

Step 3 Verify that the routing protocols are set up correctly using the **show ip protocols** command. Make sure that the autonomous system number is correct and that all neighbors are exchanging routes.

Solution:

The following shows example output on the P1R1 router:

```
P1R1#show ip protocols
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.1.0.0/16
    172.31.1.0/24
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.3              90      00:00:37
    10.1.0.2              90      00:00:35
    172.31.1.3            90      00:00:35
  Distance: internal 90 external 170

P1R1#
```

Step 4 Verify that routes from other routers in your pod and from the backbone router BBR1 are being recognized via EIGRP on each router.

Solution:

The following shows example output on the P1R1 router:

```
P1R1#show ip route
<output omitted>
Gateway of last resort is not set

     172.31.0.0/24 is subnetted, 2 subnets
D       172.31.2.0 [90/21024000] via 172.31.1.3, 00:04:41, Serial0/0/0
C       172.31.1.0 is directly connected, Serial0/0/0
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.2.0 [90/20514560] via 10.1.0.2, 00:10:08, Serial0/0/1
C       10.1.1.0 is directly connected, FastEthernet0/0
C       10.1.0.0 is directly connected, Serial0/0/1
D       10.254.0.0 [90/20514560] via 172.31.1.3, 00:04:42, Serial0/0/0
P1R1#
```

The highlighted routes are being learned by EIGRP.

Step 5 Use **debug ip eigrp** on the internal routers in your pod to monitor the EIGRP queries.

Step 6 Shut down the serial interface between the edge routers (the S0/0/1 interface on P*x*R1 and P*x*R2).

Step 7 View the EIGRP queries sent to the internal routers.

Solution:

The following shows the required command on the P1R3 router, the configuration on the P1R1 router, and example output on the P1R3 router:

```
P1R3#debug ip eigrp
IP-EIGRP Route Events debugging is on
P1R3#

P1R1(config)#int s0/0/1
P1R1(config-if)#shutdown

P1R3#
*Mar  6 02:19:11.363: IP-EIGRP(Default-IP-Routing-Table:1): Processing incoming QUERY
  packet
```

```
*Mar  6 02:19:11.367: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.0.0/24 M
   4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
*Mar  6 02:19:11.367: IP-EIGRP(Default-IP-Routing-Table:1): 10.1.0.0/24 routing table
   not updated thru 10.1.1.1
*Mar  6 02:19:11.367: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.2.0/24 M
   4294967295 - 20000000 4294967295 SM 4294967295 - 20000000 4294967295
*Mar  6 02:19:11.367: IP-EIGRP(Default-IP-Routing-Table:1): 10.1.2.0/24 routing table
   not updated thru 10.1.1.1
*Mar  6 02:19:11.387: IP-EIGRP(Default-IP-Routing-Table:1): 10.1.0.0/24- not in IP
   routing table
*Mar  6 02:19:11.387: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.0.0/24 metric
   4294967295 - 20000000 4294967295
*Mar  6 02:19:11.387: IP-EIGRP(Default-IP-Routing-Table:1): 10.1.2.0/24 - not in IP
   routing table
*Mar  6 02:19:11.387: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.2.0/24 metric
   4294967295 - 20000000 4294967295
P1R3#
```

P1R3 receives a query for network 10.1.0.0/24 from P1R1; 10.1.0.0/24 is unreachable, as indicated by the infinite metric 4294967295. P1R3 replies to the query, indicating that 10.1.0.0/24 is unreachable (using the same infinite metric).

Step 8 Turn off all debugging.

Solution:

The following shows the required command on the P1R3 router:

```
P1R3#no debug all
All possible debugging has been turned off
P1R3#
```

Step 9 Reenable the serial interface between the edge routers (the S0/0/1 interface on P*x*R1 and P*x*R2).

Solution:

The following shows the required configuration on the P1R1 router:

```
P1R1(config)#int s0/0/1
P1R1(config-if)#no shutdown
```

Task 2: Configuring EIGRP Summarization

In this task, you configure EIGRP route summarization. This will add stability and speed convergence of the network by controlling the scope of queries, minimizing update traffic, and minimizing routing table size. Follow these steps:

Step 1 Telnet to BBR1 (172.31.*x*.3) and verify that you see the specific subnet routes from your pod.

Solution:

The following shows sample output on the BBR1 router:

```
BBR1>show ip route eigrp
     10.0.0.0/24 is subnetted, 7 subnets
D       10.1.2.0 [90/20514560] via 172.31.1.2, 00:00:28, Serial0/0/0.1
D       10.1.1.0 [90/20514560] via 172.31.1.1, 00:00:29, Serial0/0/0.1
D       10.1.0.0 [90/21024000] via 172.31.1.2, 00:00:32, Serial0/0/0.1
                 [90/21024000] via 172.31.1.1, 00:00:32, Serial0/0/0.1
BBR1>
```

Step 2 Manually configure the edge routers (P*x*R1 and P*x*R2) to summarize the pod EIGRP routes to BBR1 into a single 10.*x*.0.0/16 advertisement (where *x* is your pod number).

Solution:

The following shows the required configuration on the P1R1 router:

```
P1R1(config)#int s0/0/0
P1R1(config-if)#ip summary-address eigrp 1 10.1.0.0 255.255.0.0
P1R1(config-if)#
```

Both edge routers require the same summarization configuration.

Step 3 Telnet to BBR1 (172.31.*x*.3) and verify that you see only the summary route and not the more specific routes from your pod. If both edge routers are configured correctly, you should see two equal-cost paths available to BBR1.

Solution:

The following shows sample output on the BBR1 router:

```
BBR1>show ip route eigrp
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D       10.1.0.0/16 [90/20514560] via 172.31.1.2, 00:00:33, Serial0/0/0.1
                    [90/20514560] via 172.31.1.1, 00:00:33, Serial0/0/0.1
BBR1>
```

Only the summarized 10.1.0.0/16 route is displayed; there are two equal-cost routes to this network, via P1R1 and P1R2.

Task 3: Configuring the EIGRP Stub

Having optimized BBR1's routing table by summarizing the routes from the pod’s edge routers to the core BBR1 router, you now limit the query traffic from the pod’s edge routers to its internal routers. Follow these steps:

Step 1 Configure the internal routers (P*x*R3 and P*x*R4) as EIGRP stubs. Remember that this bounds queries but does not affect the routing table.

Solution:

The following shows the required configuration on the P1R3 router:

```
P1R3(config)#router eigrp 1
P1R3(config-router)#eigrp stub
```

Step 2 Verify that the edge router recognizes its internal EIGRP neighbor as a stub.

Solution:

The following shows sample output on the P1R1 router. The highlighted lines indicate that P1R1 sees P1R3 (10.1.1.3) as a stub:

```
P1R1#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.1.3                Fa0/0             10 00:02:05   12   200  0  12
   Version 12.4/1.2, Retrans: 0, Retries: 0
   Stub Peer Advertising ( CONNECTED SUMMARY ) Routes
   Suppressing queries
0   10.1.0.2                Se0/0/1           12 00:06:46   25  1140  0  40
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 8
2   172.31.1.3              Se0/0/0          159 00:18:03  225  1350  0  4340
   Restart time 00:04:37
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 6
P1R1#
```

Step 3 The stub designation bounds query traffic and helps the router avoid getting into a stuck-in-active state, where EIGRP is unable to resolve routes for long periods. To demonstrate this situation, use the **debug ip eigrp** command on the internal router.

Step 4 Shut down the serial interface between the edge routers (the S0/0/1 interface between P*x*R1 and P*x*R2).

Step 5 Compared to the time before the internal routers were configured as stubs, notice that no queries are now being sent to the internal router. You should *not* see the “processing incoming QUERY” debug message on the internal routers, because they are configured as stub routers.

Solution:

The following shows the required command on the P1R3 router, the configuration on the P1R1 router, and example output on the P1R3 router. Queries are no longer being sent to the internal routers. P1R1 only sends the Update packet to P1R3:

```
P1R3#debug ip eigrp
IP-EIGRP Route Events debugging is on

P1R1(config)#int s0/0/1
P1R1(config-if)#shutdown
```

```
P1R3#
*Mar  6 02:32:34.507: IP-EIGRP(Default-IP-Routing-Table:1): Processing incoming UPDATE
  packet
*Mar  6 02:32:34.507: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.0.0/24 M
  4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
*Mar  6 02:32:34.507: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.2.0/24 M
  4294967295 - 20000000 4294967295 SM 4294967295 - 20000000 4294967295
*Mar  6 02:32:34.523: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.0.0/24 metric
  4294967295 - 0 4294967295
*Mar  6 02:32:34.523: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.2.0/24 metric
  4294967295 - 20000000 4294967295
*Mar  6 02:32:34.543: IP-EIGRP(Default-IP-Routing-Table:1): Processing incoming REPLY
  packet
*Mar  6 02:32:34.543: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.0.0/24 M
  4294967295 - 0 4294967295 SM 4294967295 - 0 4294967295
*Mar  6 02:32:34.543: IP-EIGRP(Default-IP-Routing-Table:1): Int 10.1.2.0/24 M
  4294967295 - 20000000 4294967295 SM 4294967295 - 20000000 4294967295
P1R3#
```

Step 6 Turn off debugging on the internal routers (P*x*R3 and P*x*R4).

Solution:

The following shows the required command on the P1R3 router:

```
P1R3#no debug all
All possible debugging has been turned off
P1R3#
```

Step 7 Reenable the serial interface between the edge routers (the S0/0/1 interface between P*x*R1 and P*x*R2).

Solution:

The following shows the required configuration on the P1R1 router:

```
P1R1(config)#int s0/0/1
P1R1(config-if)#no shutdown
```

Task 4: Configuring an EIGRP Default Route

In this task, you advertise a default route from the edge routers to the internal routers via EIGRP. This change adds stability and speed convergence to the network by minimizing update traffic and routing table size. Follow these steps:

Step 1 Send a default route from the edge routers to the internal routers, and filter all specific routes. You can do this by configuring a summary route of 0.0.0.0 0.0.0.0 on each edge router, on the interface to the internal router.

Solution:

The following shows the required configuration on the P1R1 router:

```
P1R1(config)#int fa0/0
P1R1(config-if)#ip summary-address eigrp 1 0.0.0.0 0.0.0.0
```

Step 2 Examine the routing table on the internal routers. You should see the default routes and the connected routes, but the more specific routes from the edge router should have been filtered.

Solution:

The following shows sample output on the P1R3 router. Notice that the gateway of last resort is also now set on the internal routers:

```
P1R3#show ip route
<output omitted>
Gateway of last resort is 10.1.1.1 to network 0.0.0.0

     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.1.0 is directly connected, FastEthernet0/0
D*   0.0.0.0/0 [90/30720] via 10.1.1.1, 00:01:58, FastEthernet0/0
```

Step 3 Ping the TFTP server (10.254.0.254) from the internal router to verify connectivity.

Solution:

The following shows sample output on the P1R3 router. The ping is successful:

```
P1R3#ping 10.254.0.254

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.254.0.254, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
P1R3#
```

Step 4 Examine the EIGRP topology table, EIGRP traffic statistics, information about interfaces configured for EIGRP, and EIGRP neighbors.

Solution:

The following shows sample output on the P1R1 router:

```
P1R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(1)/ID(172.31.1.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 0.0.0.0/0, 1 successors, FD is 28160
        via Summary (28160/0), Null0
P 10.1.2.0/24, 1 successors, FD is 20514560
        via 10.1.0.2 (20514560/28160), Serial0/0/1
P 10.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 10.1.0.0/16, 1 successors, FD is 28160
        via Summary (28160/0), Null0
P 10.1.0.0/24, 1 successors, FD is 20512000
        via Connected, Serial0/0/1
P 172.31.2.0/24, 1 successors, FD is 21024000
        via 172.31.1.3 (21024000/20512000), Serial0/0/0
```

```
P 172.31.1.0/24, 1 successors, FD is 20512000
        via Connected, Serial0/0/0
P 10.254.0.0/24, 1 successors, FD is 20514560
        via 172.31.1.3 (20514560/28160), Serial0/0/0

P1R1#show ip eigrp traffic
IP-EIGRP Traffic Statistics for AS 1
  Hellos sent/received: 907/905
  Updates sent/received: 341/35
  Queries sent/received: 6/7
  Replies sent/received: 7/6
  Acks sent/received: 33/40
  Input queue high water mark 2, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 150
  PDM Process ID: 88

P1R1#show ip eigrp interfaces
IP-EIGRP interfaces for process 1

                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0                1       0/0         4       0/10          50           0
Se0/0/1              1       0/0        35       5/190        346           0
Se0/0/0              2       0/0        75       5/190        748           0
P1R1#

P1R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.0.2                Se0/0/1           14 00:07:39   35  1140  0  65
1   10.1.1.3                Fa0/0             13 00:14:21    4   200  0  18
2   172.31.1.3              Se0/0/0          139 00:30:19  151  1140  0  4341
P1R1#
```

Step 5 Save your configurations to NVRAM.

Solution:

The following shows how to perform the required step on the P1R1 router:

```
P1R1#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
```

Exercise Verification

You have successfully completed this exercise when you achieve the following results:

- You have successfully implemented EIGRP and have observed EIGRP query traffic.
- You have summarized your pod addresses to the core.
- You have optimized performance on the internal routers.

Review Questions

Answer the following questions, and then refer to Appendix A, "Answers to Review Questions," for the answers.

1. What are some features of EIGRP?
2. Is EIGRP operational traffic multicast or broadcast?
3. What are the four key technologies employed by EIGRP?
4. How do IGRP and EIGRP differ in their metric calculation?
5. Which of the following best describes the EIGRP topology table?
   a. It is populated as a result of receiving hello packets.
   b. It contains all learned routes to a destination.
   c. It contains only the best routes to a destination.
6. Describe the five types of EIGRP packets.
7. True or false: EIGRP hello packets are sent every 5 seconds on LAN links.
8. What is the difference between the hold time and the hello interval?
9. Which of the following statements are true?
   a. A route is considered passive when the router is not performing recomputation on that route.
   b. A route is passive when it is undergoing recomputation.
   c. A route is active when it is undergoing recomputation.
   d. A route is considered active when the router is not performing recomputation on that route.
   e. Passive is the operational state for a route.
   f. Active is the operational state for a route.
10. Which command is used to see the RTO and hold time?
    a. **show ip eigrp traffic**
    b. **show ip eigrp timers**
    c. **show ip eigrp route**
    d. **show ip eigrp neighbors**
11. Why are EIGRP routing updates described as reliable?
12. What units are the bandwidth and delay parameters in the EIGRP metric calculation?
13. Which of the following statements are true regarding AD and FD?
    a. The AD is the EIGRP metric for a *neighbor router* to reach a particular network.
    b. The AD is the EIGRP metric for *this router* to reach a particular network.

c. The FD is the EIGRP metric for *this router* to reach a particular network.

d. The FD is the EIGRP metric for *the neighbor router* to reach a particular network.

14. What does it mean when a route is marked as an FS?

15. In the following table, place the letter of the description next to the term the description describes. The descriptions may be used more than once.

Descriptions:

a. A network protocol that EIGRP supports.

b. A table that contains FS information.

c. The administrative distance determines routing information that is included in this table.

d. A neighbor router that has the best path to a destination.

e. A neighbor router that has a loop-free alternative path to a destination.

f. An algorithm used by EIGRP that ensures fast convergence.

g. A multicast packet used to discover neighbors.

h. A packet sent by EIGRP routers when a new neighbor is discovered and when a change occurs.

Term	Description Letter
Successor	
Feasible successor	
Hello	
Topology table	
IP	
Update	
AppleTalk	
Routing table	
DUAL	
IPX	

16. Answer true or false to the following statements.

a. EIGRP performs autosummarization.

b. EIGRP autosummarization cannot be turned off.

c. EIGRP supports VLSM.

d. EIGRP can maintain three independent routing tables.

e. The EIGRP hello interval is an unchangeable fixed value.

17. Which of the following are true?

- **a.** For Frame Relay point-to-point interfaces, set the **bandwidth** to the CIR.
- **b.** For Frame Relay point-to-point interfaces set the **bandwidth** to the sum of all CIRs.
- **c.** For Frame Relay multipoint connections, set the **bandwidth** to the sum of all CIRs.
- **d.** For generic serial interfaces such as PPP and HDLC, set the **bandwidth** to match the line speed.
- **e.** For Frame Relay multipoint connections, set the **bandwidth** to the CIR.

18. Router A has three interfaces with IP addresses 172.16.1.1/24, 172.16.2.3/24, and 172.16.5.1/24. What commands would be used to configure EIGRP to run in autonomous system 100 on only the interfaces with addresses 172.16.2.3/24 and 172.16.5.1/24?

19. Routers A and B are connected and are running EIGRP on all their interfaces. Router A has four interfaces, with IP addresses 172.16.1.1/24, 172.16.2.3/24, 172.16.5.1/24, and 10.1.1.1/24. Router B has two interfaces, with IP addresses 172.16.1.2/24 and 192.168.1.1/24. There are other routers in the network that are connected on each of the interfaces of these two routers that are also running EIGRP. Which summary routes does Router A generate automatically?

- **a.** 172.16.0.0/16
- **b.** 192.168.1.0/24
- **c.** 10.0.0.0/8
- **d.** 172.16.1.0/22
- **e.** 10.1.1.0/24

20. Router A has four EIGRP paths to a destination with the following EIGRP metrics. Assuming no potential routing loops exist and the command **variance 3** is configured on Router A, which paths are included for load balancing?

- **a.** Path 1: 1100
- **b.** Path 2: 1200
- **c.** Path 3: 2000
- **d.** Path 4: 4000

21. Router A has the following configuration:

```
interface s0
  ip bandwidth-percent eigrp 100 40
  bandwidth 256
router eigrp 100
  network 10.0.0.0
```

What is the maximum bandwidth that EIGRP uses on the S0 interface?

a. 100

b. 40

c. 256

d. 102

e. 10

f. 47

22. What is the default EIGRP authentication?

a. Simple password

b. MD5

c. None

d. IPsec

23. True or false: When configuring EIGRP authentication, each router must have a unique password configured.

24. What does the **accept-lifetime** command do for EIGRP authentication?

25. What command is used to troubleshoot EIGRP authentication?

a. **debug eigrp authentication**

b. **debug ip eigrp packets**

c. **debug eigrp packets**

d. **debug ip eigrp authentication**

26. What is the default EIGRP stuck-in-active timer?

27. With the EIGRP active process enhancement, when does the SIA-Query get sent?

28. How does EIGRP summarization limit the query range?

29. How does the EIGRP stub feature limit the query range?

30. What does the **eigrp stub receive-only** command do?

31. True or false: Goodbye messages are sent in hello packets.

32. The following is part of the output of the **show ip eigrp topology** command:

```
P 10.1.3.0/24, 1 successors, FD is 10514432
        via 10.1.2.2 (10514432/28160), Serial0/0/0
```

What are the two numbers in parentheses?