Configuring EIGRP

This chapter describes how to configure Enhanced Interior Gateway Routing Protocol (E IGRP). For a complete description of the EIGRP commands listed in this chapter, refer to the "EIGRP Commands" chapter of the *Network Protocols Command Reference, Part 1*. To locate documentation of other commands that appear in this chapter, use the command reference master index or search online.

Refer to the *Network Protocols Configuration Guide, Part 2* for information on AppleTalk Enhanced IGRP or IPX Enhanced IGRP.

For protocol-independent features that work with EIGRP, see the chapter "Configuring IP Routing Protocol-Independent Features."

EIGRP is an enhanced version of IGRP developed by Cisco Systems, Inc. EIGRP uses the same distance vector algorithm and distance information as IGRP. However, the convergence properties and the operating efficiency of EIGRP have improved significantly over IGRP.

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

Cisco's EIGRP Implementation

EIGRP provides the following features:

- Automatic redistribution—IGRP routes can be automatically redistributed into EIGRP, and EIGRP routes can be automatically redistributed into IGRP. If desired, you can turn off redistribution. You can also completely turn off EIGRP and IGRP on the router or on individual interfaces.
- Increased network width—With IP RIP, the largest possible width of your network is 15 hops. When EIGRP is enabled, the largest possible width is 224 hops. Because the EIGRP metric is large enough to support thousands of hops, the only barrier to expanding the network is the transport layer hop counter. Cisco works around this problem by incrementing the transport control field only when an IP packet has traversed 15 routers and the next hop to the destination was learned by way of EIGRP. When a RIP route is being used as the next hop to the destination, the transport control field is incremented as usual.

Note Redistribution between EIGRP and IGRP differs from normal redistribution in that the metrics of IGRP routes are compared with the metrics of external EIGRP routes. The rules of normal administrative distances are not followed, and routes with the lowest metric are selected.

EIGRP offers the following features:

- Fast convergence—The DUAL algorithm allows routing information to converge as quickly as any currently available routing protocol.
- Partial updates—EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature minimizes the bandwidth required for EIGRP packets.
- Less CPU usage than IGRP—This occurs because full update packets do not have to be processed each time they are received.
- Neighbor discovery mechanism—This is a simple hello mechanism used to learn about neighboring routers. It is protocol-independent.
- Variable-length subnet masks
- Arbitrary route summarization
- Scaling—EIGRP scales to large networks.

EIGRP has the following four basic components:

- Neighbor discovery/recovery
- Reliable transport protocol
- DUAL finite state machine
- Protocol-dependent modules

Neighbor discovery/recovery is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery/recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, the Cisco IOS software can determine that a neighbor is alive and functioning. Once this status is determined, the neighboring routers can exchange routing information.

The reliable transport protocol is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be transmitted reliably and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities (such as Ethernet) it is not necessary to send hellos reliably to all neighbors individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets (such as updates) require acknowledgment, and this is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. Doing so helps ensure that convergence time remains low in the presence of varying speed links.

The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information (known as a metric) to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. 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. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur.

This is the process whereby a new successor is determined. The amount of time it takes to recompute the route affects the convergence time. Recomputation is processor-intensive; it is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, it will use any it finds in order to avoid unnecessary recomputation.

The protocol-dependent modules are responsible for network layer protocol-specific tasks. An example is the EIGRP module, which is responsible for sending and receiving EIGRP packets that are encapsulated in IP. It is also responsible for parsing EIGRP packets and informing DUAL of the new information received. EIGRP asks DUAL to make routing decisions, but the results are stored in the IP routing table. Also, EIGRP is responsible for redistributing routes learned by other IP routing protocols.

EIGRP Configuration Task List

To configure EIGRP, complete the tasks in the following sections. At a minimum, you must enable EIGRP. The remaining tasks are optional.

- Enable EIGRP
- Transition from IGRP to EIGRP
- Log EIGRP Neighbor Adjacency Changes
- Configure the Percentage of Link Bandwidth Used
- Adjust the EIGRP Metric Weights
- Apply Offsets to Routing Metrics
- Disable Route Summarization
- Configure Summary Aggregate Addresses
- Configure EIGRP Route Authentication
- Configure EIGRP's Protocol-Independent Parameters
- Monitor and Maintain EIGRP

See the section "EIGRP Configuration Examples" at the end of this chapter for configuration examples.

Enable EIGRP

To create an EIGRP routing process, use the following commands, beginning in global configuration mode:

Step	Command	Purpose
1	**router eigrp** *autonomous-system*	Enable an EIGRP routing process in global configuration mode.
2	**network** *network-number*	Associate networks with an EIGRP routing process in router configuration mode.

EIGRP sends updates to the interfaces in the specified networks. If you do not specify an interface's network, it will not be advertised in any EIGRP update.

Transition from IGRP to EIGRP

If you have routers on your network that are configured for IGRP, and you want to make a transition to routing EIGRP, you must designate transition routers that have both IGRP and EIGRP configured. In these cases, perform the tasks as noted in the previous section, "Enable EIGRP," and also read the chapter, "Configuring IGRP" in this document. You must use the same autonomous system number in order for routes to be redistributed automatically.

Log EIGRP Neighbor Adjacency Changes

You can enable the logging of neighbor adjacency changes to monitor the stability of the routing system and to help you detect problems. By default, adjacency changes are not logged. To enable such logging, use the following command in global configuration mode:

Command	Purpose
eigrp log-neighbor-changes	Enable logging of EIGRP neighbor adjacency changes.

Configure the Percentage of Link Bandwidth Used

By default, EIGRP packets consume a maximum of 50 percent of the link bandwidth, as configured with the **bandwidth** interface configuration command. You might want to change that value if a different level of link utilization is required or if the configured bandwidth does not match the actual link bandwidth (it may have been configured to influence route metric calculations).

To configure the percentage of bandwidth that may be used by EIGRP on an interface, use the following command in interface configuration mode:

Command	Purpose
ip bandwidth-percent eigrp *percent*	Configure the percentage of bandwidth that may be used by EIGRP on an interface.

Adjust the EIGRP Metric Weights

You can adjust the default behavior of EIGRP routing and metric computations. For example, this adjustment allows you to tune system behavior to allow for satellite transmission. Although EIGRP metric defaults have been carefully selected to provide excellent operation in most networks, you can adjust the EIGRP metric. Adjusting EIGRP metric weights can dramatically affect network performance, so be careful if you adjust them.

To adjust the EIGRP metric weights, use the following command in router configuration mode:

Command	Purpose
metric weights *tos k1 k2 k3 k4 k5*	Adjust the EIGRP metric.

Note Because of the complexity of this task, it is not recommended unless it is done with guidance from an experienced network designer.

By default, the EIGRP composite metric is a 32-bit quantity that is a sum of the segment delays and the lowest segment bandwidth (scaled and inverted) for a given route. For a network of homogeneous media, this metric reduces to a hop count. For a network of mixed media (FDDI, Ethernet, and serial lines running from 9600 bps to T1 rates), the route with the lowest metric reflects the most desirable path to a destination.

Apply Offsets to Routing Metrics

An offset list is the mechanism for increasing incoming and outgoing metrics to routes learned via EIGRP. This is done to provide a local mechanism for increasing the value of routing metrics. Optionally, you can limit the offset list with either an access list or an interface. To increase the value of routing metrics, use the following command in router configuration mode:

Command	Purpose
offset-list [*access-list-number* | *name*] {**in** | **out**} *offset* [*type number*]	Apply an offset to routing metrics.

Disable Route Summarization

You can configure EIGRP to perform automatic summarization of subnet routes into network-level routes. For example, you can configure subnet 131.108.1.0 to be advertised as 131.108.0.0 over interfaces that have subnets of 192.31.7.0 configured. Automatic summarization is performed when there are two or more **network** router configuration commands configured for the EIGRP process. By default, this feature is enabled.

To disable automatic summarization, use the following command in router configuration mode:

Command	Purpose
no auto-summary	Disable automatic summarization.

Route summarization works in conjunction with the **ip summary-address eigrp** interface configuration command, in which additional summarization can be performed. If automatic summarization is in effect, there usually is no need to configure network level summaries using the **ip summary-address eigrp** command.

Configure Summary Aggregate Addresses

You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

To configure a summary aggregate address, use the following command in interface configuration mode:

Command	Purpose
ip summary-address eigrp *autonomous-system-number address mask*	Configure a summary aggregate address.

See the "Route Summarization Example" at the end of this chapter for an example of summarizing aggregate addresses.

Configure EIGRP Route Authentication

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol. The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

Before you can enable EIGRP route authentication, you must enable EIGRP.

To enable authentication of EIGRP packets, use the following commands beginning in interface configuration mode:

Step	Command	Purpose
1	**interface** *type number*	Configure an interface type and enter interface configuration mode
2	**ip authentication mode eigrp** *autonomous-system* **md5**	Enable MD5 authentication in EIGRP packets.
3	**ip authentication key-chain eigrp** *autonomous-system key-chain*	Enable authentication of EIGRP packets.
4	**exit**	Exit to global configuration mode.
5	**key chain** *name-of-chain*	Identify a key chain. (Match the name configured in Step 1.)
6	**key** *number*	In key chain configuration mode, identify the key number.
7	**key-string** *text*	In key chain key configuration mode, identify the key string.
8	**accept-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	Optionally specify the time period during which the key can be received.
9	**send-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	Optionally specify the time period during which the key can be sent.

Each key has its own key identifier (specified with the **key** *number* command), which is stored locally. The combination of the key identifier and the interface associated with the message uniquely identifies the authentication algorithm and MD5 authentication key in use.

You can configure multiple keys with lifetimes. 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 uses the first valid key it encounters. Note that the router needs to know the time. Refer to the NTP and calendar commands in the "Performing Basic System Management" chapter of the *Configuration Fundamentals Configuration Guide*.

For an example of route authentication, see the section "Route Authentication Example" at the end of this chapter.

Configure EIGRP's Protocol-Independent Parameters

EIGRP works with AppleTalk, IP, and IPX. The bulk of this chapter describes EIGRP. However, this section describes EIGRP features that work for AppleTalk, IP, and IPX. To configure such protocol-independent parameters, perform one or more of the tasks in the following sections:

- Adjust the Interval between Hello Packets and the Hold Time
- Disable Split Horizon

For more protocol-independent features that work with EIGRP, see the chapter "Configuring IP Routing Protocol-Independent Features."

Adjust the Interval between Hello Packets and the Hold Time

You can adjust the interval between hello packets and the hold time.

Routing devices periodically send hello packets to each other to dynamically learn of other routers on their directly attached networks. This information is used to discover who their neighbors are, and to learn when their neighbors become unreachable or inoperative.

By default, hello packets are sent every 5 seconds. The exception is on low-speed, nonbroadcast, multiaccess (NBMA) media, where the default hello interval is 60 seconds. Low speed is considered to be a rate of T1 or slower, as specified with the **bandwidth** interface configuration command. The default hello interval remains 5 seconds for high-speed NBMA networks. Note that for the purposes of EIGRP, Frame Relay and SMDS networks may or may not be considered to be NBMA. These networks are considered NBMA if the interface has not been configured to use physical multicasting; otherwise they are not considered NBMA.

You can configure the hold time on a specified interface for a particular EIGRP routing process designated by the autonomous system number. The hold time is advertised in hello packets and indicates to neighbors the length of time they should consider the sender valid. The default hold time is three times the hello interval, or 15 seconds. For slow-speed NBMA networks, the default hold time is 180 seconds.

To change the interval between hello packets, use the following command in interface configuration mode:

Command	Purpose
ip hello-interval eigrp *autonomous-system-number seconds*	Configure the hello interval for an EIGRP routing process.

On very congested and large networks, the default hold time might not be sufficient time for all routers to receive hello packets from their neighbors. In this case, you may want to increase the hold time.

To change the hold time, use the following command in interface configuration mode:

Command	Purpose
ip hold-time eigrp *autonomous-system-number seconds*	Configure the hold time for an EIGRP routing process.

Note Do not adjust the hold time without advising technical support.

Disable Split Horizon

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, these packets are not sent for destinations for which this interface is the next hop. This reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks (such as Frame Relay and SMDS), situations can arise for which this behavior is less than ideal. For these situations, you may want to disable split horizon.

To disable split horizon, use the following command in interface configuration mode:

Command	Purpose
no ip split-horizon eigrp *autonomous-system-number*	Disable split horizon.

Monitor and Maintain EIGRP

To delete neighbors from the neighbor table, use the following command in EXEC mode:

Command	Purpose
clear ip eigrp neighbors [*ip-address* | *interface*]	Delete neighbors from the neighbor table.

To display various routing statistics, use the following commands in EXEC mode:

Command	Purpose
show ip eigrp interfaces [*interface*] [*as-number*]	Display information about interfaces configured for EIGRP.
show ip eigrp neighbors [*type number*]	Display the EIGRP discovered neighbors.
show ip eigrp topology [*autonomous-system-number* | [[*ip-address*] *mask*]]	Display the EIGRP topology table for a given process.
show ip eigrp traffic [*autonomous-system-number*]	Display the number of packets sent and received for all or a specified EIGRP process.

EIGRP Configuration Examples

This section contains the following examples:

- Route Summarization Example
- Route Authentication Example

Route Summarization Example

The following example configures route summarization on the interface and also configures the auto-summary feature. This configuration causes EIGRP to summarize network 10.0.0.0 out Ethernet interface 0 only. In addition, this example disables auto summarization.

```
interface Ethernet 0
 ip summary-address eigrp 1 10.0.0.0 255.0.0.0
!
router eigrp 1
 network 172.16.0.0
 no auto-summary
```

Note You should not use the **ip summary-address eigrp** summarization command to generate the default route (0.0.0.0) from an interface. This causes the creation of an EIGRP summary default route to the null 0 interface with an administrative distance of 5. The low administrative distance of this default route can cause this route to displace default routes learned from other neighbors from the routing table. If the default route learned from the neighbors is displaced by the summary default route, or if the summary route is the only default route present, all traffic destined for the default route will not leave the router, instead, this traffic will be sent to the null 0 interface where it is dropped.

The recommended way to send only the default route out a given interface is to use a **distribute-list** command. You can configure this command to filter all outbound route advertisements sent out the interface with the exception of the default (0.0.0.0).

Route Authentication Example

The following example enables MD5 authentication on EIGRP packets in autonomous system 1. Figure 28 shows the scenario.

Figure 28 EIGRP Route Authentication Scenario

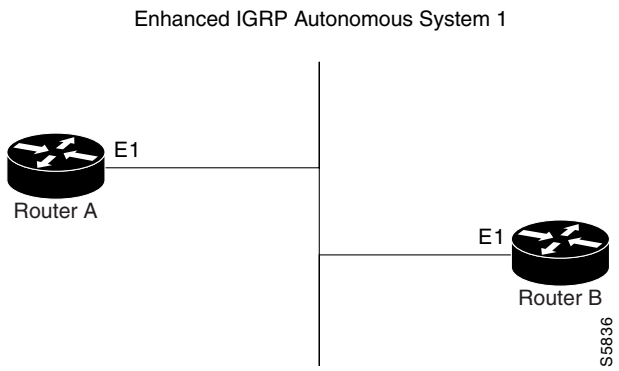

Router A

```
interface ethernet 1
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 holly
key chain holly
 key 1
  key-string 0987654321
  accept-lifetime 04:00:00 Dec 4 2006 infinite
  send-lifetime 04:00:00 Dec 4 2006 04:48:00 Dec 4 1996
 exit
 key 2
  key-string 1234567890
  accept-lifetime 04:00:00 Jan 4 2007 infinite
  send-lifetime 04:45:00 Jan 4 2007 infinite
```

Router B

```
interface ethernet 1
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 mikel
key chain mikel
 key 1
  key-string 0987654321
  accept-lifetime 04:00:00 Dec 4 2006 infinite
  send-lifetime 04:00:00 Dec 4 2006 infinite
 exit
 key 2
  key-string 1234567890
  accept-lifetime 04:00:00 Jan 4 2007 infinite
  send-lifetime 04:45:00 Jan 4 2007 infinite
```

Router A will accept and attempt to verify the MD5 digest of any EIGRP packet with a key equal to 1. It will also accept a packet with a key equal to 2. All other MD5 packets will be dropped. Router A will send all EIGRP packets with key 2.

Router B will accept key 1 or key 2, and will use key 1 to send MD5 authentication, since key 1 is the first the first valid key off the key-chain. Key 1 will no longer be valid to be used for sending after December 4, 2006. After this date, key 2 would be used to send MD5 authentication., since it is valid until January 4, 2007.