



Configuring EIGRP

This chapter describes how to configure Enhanced Interior Gateway Routing Protocol (EIGRP). For a complete description of the EIGRP commands listed in this chapter, refer to the “EIGRP Commands” chapter of the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols* publication. To locate documentation of other commands that appear in this chapter, use the command reference master index, or search online.

Refer to the *Cisco IOS AppleTalk and Novell IPX Configuration Guide* for information on AppleTalk EIGRP or Internetwork Packet Exchange (IPX) EIGRP.

For protocol-independent features that work with EIGRP, see the chapter “Configuring IP Routing Protocol-Independent Features” in this document.

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

To identify the hardware platform or software image information associated with a feature, use the Feature Navigator on Cisco.com to search for information about the feature or refer to the software release notes for a specific release. For more information, see the “Identifying Supported Platforms” section in the “Using Cisco IOS Software” chapter in this book.

The Cisco 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 Routing Information Protocol (RIP), the largest possible width of your network is 15 hops. When EIGRP is enabled, the largest possible width is increased to 100 hops, and the metric is large enough to support thousands of hops.

**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 need not 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 (VLSMs).
- Arbitrary route summarization.
- Scaling—EIGRP scales to large networks.

EIGRP has the following four basic components:

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

Neighbor discovery of neighbor 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 of neighbor 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 sent 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 hello packets 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, which is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when unacknowledged packets are pending. This provision helps to 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 required to recompute the route affects the

convergence time. Recomputation is processor-intensive; it is advantageous to avoid unneeded recomputation. 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, perform the tasks described in the following sections. The tasks in the first section are required; the tasks in the remaining sections are optional:

- [Enabling EIGRP](#) (Required)
- [Making the Transition from IGRP to EIGRP](#) (Optional)
- [Logging EIGRP Neighbor Adjacency Changes](#) (Optional)
- [Configuring the Percentage of Link Bandwidth Used](#) (Optional)
- [Adjusting the EIGRP Metric Weights](#) (Optional)
- [Applying Offsets to Routing Metrics](#) (Optional)
- [Disabling Route Summarization](#) (Optional)
- [Configuring Summary Aggregate Addresses](#) (Optional)
- [Configuring Floating Summary Routes](#) (Optional)
- [Configuring EIGRP Route Authentication](#) (Optional)
- [Configuring EIGRP Protocol-Independent Parameters](#) (Optional)
- [Configuring EIGRP Stub Routing](#) (Optional)
- [Monitoring and Maintaining EIGRP](#)(Optional)

See the section “[EIGRP Configuration Examples](#)” at the end of this chapter for configuration examples.

Enabling EIGRP

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

	Command	Purpose
Step 1	Router(config)# router eigrp <i>autonomous-system</i>	Enables an EIGRP routing process in global configuration mode.
Step 2	Router(config-router)# network <i>network-number</i>	Associates 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 the network of an interface, the interface will not be advertised in any EIGRP update.

Making the 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, “[Enabling EIGRP](#),” and also see the chapter “Configuring IGRP” in this document. You must use the same autonomous system number in order for routes to be redistributed automatically.

Logging EIGRP Neighbor Adjacency Changes

By default, the system logs EIGRP neighbor adjacency changes to help you monitor the stability of the routing system and detect problems. If you disabled logging of such changes and want to reenabling the logging, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# eigrp log-neighbor-changes	Enables logging of EIGRP neighbor adjacency changes.

Configuring 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
Router(config-if)# ip bandwidth-percent eigrp percent	Configures the percentage of bandwidth that may be used by EIGRP on an interface.

Adjusting the EIGRP Metric Weights

EIGRP uses the minimum bandwidth on the path to a destination network and the total delay to compute routing metrics. You can use the **eigrp metric weights** command to 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. EIGRP metric defaults have been carefully selected to provide optimal performance in most networks.



Note

Adjusting EIGRP metric weights can dramatically affect network performance. Because of the complexity of this task, we recommend that you do not change the default values without guidance from an experienced network designer.

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

Command	Purpose
Router(config-router)# metric weights tos k1 k2 k3 k4 k5	Adjusts the EIGRP metric or K value. EIGRP uses the following formula to determine the total metric to the network: metric = [K1*bandwidth + (K2*bandwidth)/(256 - load) + K3*delay] * [K5/(reliability + K4)]

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 bits per second to T1 rates), the route with the lowest metric reflects the most desirable path to a destination.

Mismatched K Values

Mismatched K values (EIGRP metrics) can prevent neighbor relationships from being established and can negatively impact network convergence. The following example explains this behavior between 2 EIGRP peers (ROUTER-A and ROUTER-B).

The following error message is displayed in the console of ROUTER-B because the K values are mismatched:

```
*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
```

There are two scenarios where this error message can be displayed:

- The two routers are connected on the same link and configured to establish a neighbor relationship. However, each router is configured with different K values.

The following configuration is applied to ROUTER-A. The K values are changed with the **metric weights** command. A value of 2 is entered for the *k1* argument to adjust the bandwidth calculation. The value of 1 is entered for the *k3* argument to adjust the delay calculation.

```
hostname ROUTER-A!
interface serial 0
 ip address 10.1.1.1 255.255.255.0
 exit
router eigrp 100
 network 10.1.1.0 0.0.0.255
 metric weights 0 2 0 1 0 0
```

The following configuration is applied to ROUTER-B. However, the **metric weights** command is not applied and the default K values are used. The default K values are 1, 0, 1, 0, and 0.

```
hostname ROUTER-B!
interface serial 0
 ip address 10.1.1.2 255.255.255.0!
 exit
router eigrp 100
 network 10.1.1.0 0.0.0.255
```

The bandwidth calculation is set to 2 on ROUTER-A and set to 1 (by default) on ROUTER-B. This configuration prevents these peers from forming a neighbor relationship.

- The K-value mismatch error message can also be displayed if one of the two peers has transmitted a “goodbye” message, and the receiving router does not support this message. In this case, the receiving router will interpret this message as a K-value mismatch.

The Goodbye Message

The goodbye message is a feature designed to improve EIGRP network convergence. The goodbye message is broadcast when an EIGRP routing process is shutdown 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 Release 12.3(2), 12.3(3)B, and 12.3(2)T and later releases. The following message is displayed by routers that run a supported release when a goodbye message 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 can misinterpret the message as a K-value mismatch and 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 session when the hold timer expires. The sending and receiving routers will reconverge normally after the sender reloads.

Applying Offsets to Routing Metrics

An offset list is the mechanism for increasing incoming and outgoing metrics to routes learned via EIGRP. An offset list provides 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
Router(config-router)# offset-list [<i>access-list-number</i> <i>access-list-name</i>] { in out } <i>offset</i> [<i>interface-type interface-number</i>]	Applies an offset to routing metrics.

Disabling 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
<code>Router(config-router)# no auto-summary</code>	Disables 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.

Configuring Summary Aggregate Addresses

You can configure a summary aggregate address for a specified interface. If any more specific routes are 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
<code>Router(config-if)# ip summary-address eigrp autonomous-system-number ip-address mask</code>	Configures a summary aggregate address.

See the “[Route Summarization Example](#)” at the end of this chapter for an example of summarizing aggregate addresses.

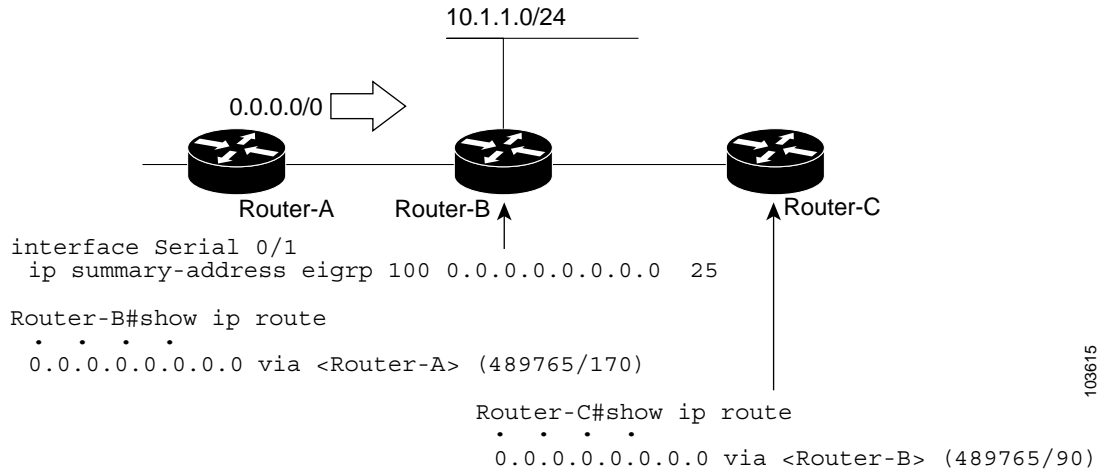
Configuring Floating Summary Routes

You can also use a floating summary route when configuring the **ip summary-address eigrp** command. This enhancement was introduced in Cisco IOS Release 12.2. The floating summary route is created by applying a default route and administrative distance at the interface level. The following scenarios illustrates the behavior of this enhancement.

[Figure 46](#) shows a network with three routers, Router-A, Router-B, and Router-C. Router-A learns a default route from elsewhere in the network and then advertises this route to Router-B. Router-B is configured so that only a default summary route is advertised to Router-C. The default summary route is applied to interface 0/1 on Router-B with the following configuration:

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

Figure 46 Floating Summary Route is Applied to Router-B



103615

The configuration of the default summary route on Router-B sends a 0.0.0.0/0 summary route to Router-C and blocks all other routes, including the 10.1.1.0/24 route, from being advertised to Router-C. However, this also generates a local discard route on Router-B, a route for 0.0.0.0/0 to the null 0 interface with an administrative distance of 5. When this route is created, it overrides the EIGRP learned default route. Router-B will no longer be able to reach destinations that it would normally reach through the 0.0.0.0/0 route.

This problem is resolved by applying a floating summary route to the interface on Router-B that connects to Router-C. The floating summary route is applied by applying an administrative distance to the default summary route on the interface of Router-B with the following statement:

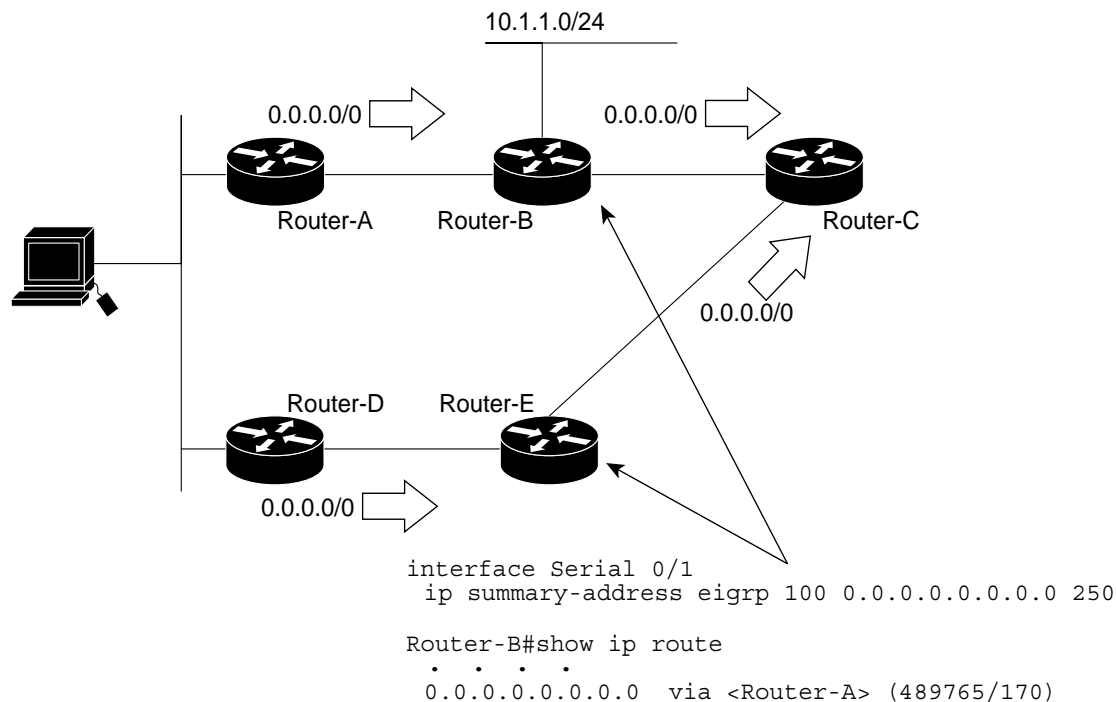
```
Router(config-if)# ip summary-address eigrp 100 0.0.0.0 0.0.0.0 250
```

The administrative distance of 250, applied in the above statement, is now assigned to the discard route generated on Router-B. The 0.0.0.0/0, from Router-A, is learned through EIGRP and installed in the local routing table. Routing to Router-C is restored.

If Router-A loses the connection to Router-B, Router-B will continue to advertise a default route to Router-C, which allows traffic to continue to reach destinations attached to Router-B. However, traffic destined to networks to Router-A or behind Router-A will be dropped when it reaches Router-B.

Figure 47 shows a network with two connections from the core, Router-A and Router-D. Both routers have floating summary routes configured on the interfaces connected to Router-C. If the connection between Router-E and Router-C fails, the network will continue to operate normally. All traffic will flow from Router-C through Router-B to the hosts attached to Router-A and Router-D.

Figure 47 Floating Summary Route Applied for Dual-Homed Remotes



103614

However, if the link between Router-D and Router-E fails, the network may blackhole traffic because Router-E will continue to advertise the default route(0.0.0.0/0) to Router-C, as long as at least one link, (other than the link to Router-C) to Router-E is still active. In this scenario, Router-C still forwards traffic to Router-E, but Router-E drops the traffic creating the black hole. To avoid this problem, you should configure the summary address with an administrative distance on only single-homed remote routers or areas where there is only one exit point between to segments of the network. If two or more exit points exist (from one segment of the network to another), configuring the floating default route can cause a black hole to be formed.

Configuring EIGRP Route Authentication

EIGRP route authentication provides Message Digest 5 (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:

	Command	Purpose
Step 1	Router(config)# interface type number	Configure an interface type and enter interface configuration mode
Step 2	Router(config-if)# ip authentication mode eigrp autonomous-system md5	Enables MD5 authentication in EIGRP packets.

	Command	Purpose
Step 3	<code>Router(config-if)# ip authentication key-chain eigrp autonomous-system key-chain</code>	Enables authentication of EIGRP packets.
Step 4	<code>Router(config-if)# exit Router(config)#</code>	Exits to global configuration mode.
Step 5	<code>Router(config)# key chain name-of-chain</code>	Identifies a key chain. (Match the name configured in Step 1.)
Step 6	<code>Router(config-keychain)# key number</code>	In keychain configuration mode, identifies the key number.
Step 7	<code>Router(config-keychain-key)# key-string text</code>	In keychain key configuration mode, identifies the key string.
Step 8	<code>Router(config-keychain-key)# accept-lifetime start-time {infinite end-time duration seconds}</code>	Optionally specifies the time period during which the key can be received.
Step 9	<code>Router(config-keychain-key)# send-lifetime start-time {infinite end-time duration seconds}</code>	Optionally specifies the time period during which the key can be sent.

Each key has its own key identifier (specified with the **key number** key chain configuration 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 Network Time Protocol (NTP) and calendar commands in the “Performing Basic System Management” chapter of the *Cisco IOS Configuration Fundamentals Configuration Guide*.

For an example of route authentication, see the section “[Route Authentication Example](#)” at the end of this chapter.

Configuring EIGRP 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:

- [Adjusting the Interval Between Hello Packets and the Hold Time](#)
- [Disabling Split Horizon](#)

For more protocol-independent features that work with EIGRP, see the chapter “Configuring IP Routing Protocol-Independent Features” in this document.

Adjusting 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 neighbors and to learn when 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 Switched Multimegabit Data Service (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
<code>Router(config-if)# ip hello-interval eigrp autonomous-system-number seconds</code>	Configures 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
<code>Router(config-if)# ip hold-time eigrp autonomous-system-number seconds</code>	Configures the hold time for an EIGRP routing process.



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

Disabling Split Horizon

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, update and query packets are not sent for destinations for which this interface is the next hop. Controlling update and query packets in this manner 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, including networks in which you have EIGRP configured, you may want to disable split horizon.

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

Command	Purpose
<code>Router(config-if)# no ip split-horizon eigrp autonomous-system-number</code>	Disables split horizon.

Configuring EIGRP Stub Routing

The EIGRP Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

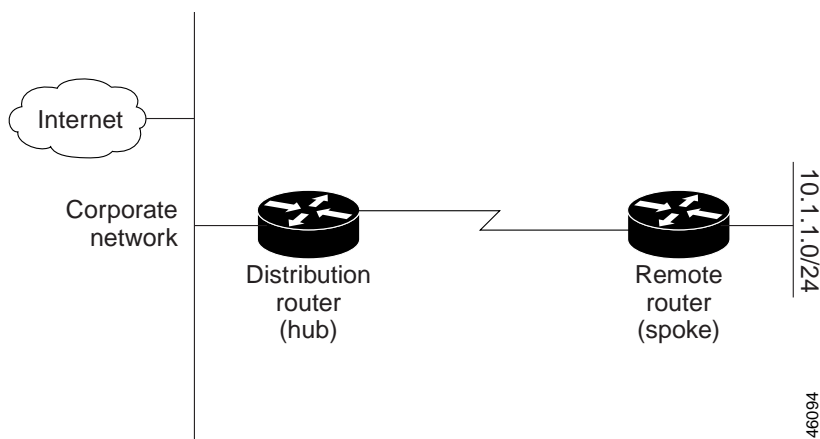
Stub routing is commonly used in a hub-and-spoke network topology. In a hub-and-spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub-and-spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The stub router responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers.

Figure 48 shows a simple hub-and-spoke configuration.

Figure 48 Simple Hub-and-Spoke Network



The stub routing feature by itself does not prevent routes from being advertised to the remote router. In the example in Figure 48, the remote router can access the corporate network and the Internet through the distribution router only. Having a full route table on the remote router, in this example, would serve no functional purpose because the path to the corporate network and the Internet would always be through the distribution router. The larger route table would only reduce the amount of memory required by the remote router. Bandwidth and memory can be conserved by summarizing and filtering routes in the distribution router. The remote router need not receive routes that have been learned from other networks because the remote router must send all nonlocal traffic, regardless of destination, to the distribution router. If a true stub network is desired, the distribution router should be configured to send

only a default route to the remote router. The EIGRP Stub Routing feature does not automatically enable summarization on the distribution router. In most cases, the network administrator will need to configure summarization on the distribution routers.



Note

When configuring the distribution router to send only a default route to the remote router, you must use 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, even after the routes that are sent from the distribution router to the remote router have been filtered or summarized, a problem might occur. If a route is lost somewhere in the corporate network, EIGRP could send a query to the distribution router, which in turn will send a query to the remote router even if routes are being summarized. If there is a problem communicating over the WAN link between the distribution router and the remote router, an EIGRP stuck in active (SIA) condition could occur and cause instability elsewhere in the network. The EIGRP Stub Routing feature allows a network administrator to prevent queries from being sent to the remote router.

Dual-Homed Remote Topology

In addition to a simple hub-and-spoke network where a remote router is connected to a single distribution router, the remote router can be dual-homed to two or more distribution routers. This configuration adds redundancy and introduces unique issues, and the stub feature helps to address some of these issues.

A dual-homed remote router will have two or more distribution (hub) routers. However, the principles of stub routing are the same as they are with a hub-and-spoke topology. [Figure 49](#) shows a common dual-homed remote topology with one remote router, but 100 or more routers could be connected on the same interfaces on distribution router 1 and distribution router 2. The remote router will use the best route to reach its destination. If distribution router 1 experiences a failure, the remote router can still use distribution router 2 to reach the corporate network.


Figure 49 Simple Dual-Homed Remote Topology



[Figure 49](#) shows a simple dual-homed remote with one remote router and two distribution routers. Both distribution routers maintain routes to the corporate network and stub network 10.1.1.0/24.

Dual-homed routing can introduce instability into an EIGRP network. In [Figure 50](#), distribution router 1 is directly connected to network 10.3.1.0/24. If summarization or filtering is applied on distribution router 1, the router will advertise network 10.3.1.0/24 to all of its directly connected EIGRP neighbors (distribution router 2 and the remote router).


Figure 50 *Dual-Homed Remote Topology With Distribution Router 1 Connected to Two Networks*



[Figure 50](#) shows a simple dual-homed remote router where distribution router 1 is connected to both network 10.3.1.0/24 and network 10.2.1.0/24.

If the 10.2.1.0/24 link between distribution router 1 and distribution router 2 has failed, the lowest cost path to network 10.3.1.0/24 from distribution router 2 is through the remote router (see [Figure 51](#)). This route is not desirable because the traffic that was previously traveling across the corporate network 10.2.1.0/24 would now be sent across a much lower bandwidth connection. The over utilization of the lower bandwidth WAN connection can cause a number of problems that might affect the entire corporate network. The use of the lower bandwidth route that passes through the remote router might cause WAN EIGRP distribution routers to be dropped. Serial lines on distribution and remote routers could also be dropped, and EIGRP SIA errors on the distribution and core routers could occur.

Figure 51 Dual-Homed Remote Topology with a Failed Route to a Distribution Router



It is not desirable for traffic from distribution router 2 to travel through any remote router in order to reach network 10.3.1.0/24. If the links are sized to handle the load, it would be acceptable to use one of the backup routes. However, most networks of this type have remote routers located at remote offices with relatively slow links. This problem can be prevented if proper summarization is configured on the distribution router and remote router.

It is typically undesirable for traffic from a distribution router to use a remote router as a transit path. A typical connection from a distribution router to a remote router would have much less bandwidth than a connection at the network core. Attempting to use a remote router with a limited bandwidth connection as a transit path would generally produce excessive congestion to the remote router. The EIGRP Stub Routing feature can prevent this problem by preventing the remote router from advertising core routes back to distribution routers. Routes learned by the remote router from distribution router 1 will not be advertised to distribution router 2. Since the remote router will not advertise core routes to distribution router 2, the distribution router will not use the remote router as a transit for traffic destined for the network core.

The EIGRP Stub Routing feature can help to provide greater network stability. In the event of network instability, this feature prevents EIGRP queries from being sent over limited bandwidth links to nontransit routers. Instead, distribution routers to which the stub router is connected answer the query on behalf of the stub router. This feature greatly reduces the chance of further network instability due to congested or problematic WAN links. 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, it is no longer necessary to configure filtering on remote routers to prevent those remote routers from appearing as transit paths to the hub routers.



Caution

EIGRP Stub Routing should only be used on stub routers. A stub router is defined as a router connected to the network core or distribution layer through which core transit traffic should not flow. A stub router should not have any EIGRP neighbors other than distribution routers. Ignoring this restriction will cause undesirable behavior.



Note

Multi-access interfaces, such as ATM, Ethernet, Frame Relay, ISDN PRI, and X.25, are supported by the EIGRP Stub Routing feature only when all routers on that interface, except the hub, are configured as stub routers.

EIGRP Stub Routing Configuration Task List

To configure EIGRP Stub Routing, perform the tasks described in the following sections. The tasks in the first section are required; the task in the last section is optional.

- [Configuring EIGRP Stub Routing](#) (required)
- [Verifying EIGRP Stub Routing](#) (optional)

Configuring EIGRP Stub Routing

To configure a remote or spoke router for EIGRP stub routing, use the following commands beginning in router configuration mode:

	Command	Purpose
Step 1	<code>router(config)# router eigrp 1</code>	Configures a remote or distribution router to run an EIGRP process.
Step 2	<code>router(config-router)# network network-number</code>	Associates networks with an EIGRP routing process in router configuration mode.
Step 3	<code>router(config-router)# eigrp stub [receive-only connected static summary]</code>	Configures a remote router as an EIGRP stub router.

Verifying EIGRP Stub Routing

To verify that a remote router has been configured as a stub router with EIGRP, use the **show ip eigrp neighbor detail** command from the distribution router in privileged EXEC mode. The last line of the output will show the stub status of the remote or spoke router. The following example shows output is from the **show ip eigrp neighbor detail** command:

```
router# show ip eigrp neighbor detail

IP-EIGRP neighbors for process 1
H   Address                Interface    Hold Uptime    SRTT   RTO  Q  Seq Type
      (sec)                (ms)          Cnt Num
0   10.1.1.2                Se3/1       11 00:00:59    1  4500  0  7
      Version 12.1/1.2, Retrans: 2, Retries: 0
      Stub Peer Advertising ( CONNECTED SUMMARY ) Routes
```

Monitoring and Maintaining EIGRP

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

Command	Purpose
<code>Router# clear ip eigrp neighbors [ip-address interface-type]</code>	Deletes neighbors from the neighbor table.

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

Command	Purpose
Router# <code>show ip eigrp interfaces</code> [<i>interface-type</i> <i>interface-number</i>] [<i>as-number</i>]	Displays information about interfaces configured for EIGRP.
Router# <code>show ip eigrp neighbors</code> [<i>interface-type</i> <i>as-number</i> <code>static</code>]	Displays the EIGRP discovered neighbors.
Router# <code>show ip eigrp topology</code> [<i>as-number</i> [[<i>ip-address</i>] <i>mask</i>]]	Displays the EIGRP topology table for a given process.
Router# <code>show ip eigrp traffic</code> [<i>as-number</i>]	Displays the number of packets sent and received for all or a specified EIGRP process.

To enable EIGRP Stub Routing packet debugging, use the following command in privileged EXEC mode:

Command	Purpose
Router# <code>debug eigrp packet stub</code>	Displays debug information about the stub status of peer routers.

EIGRP Configuration Examples

This section contains the following examples:

- [Route Summarization Example](#)
- [Route Authentication Example](#)
- [Stub Routing Example](#)

Route Summarization Example

The following example configures route summarization on the interface and also configures the automatic summary feature. This configuration causes EIGRP to summarize network 10.0.0.0 out Ethernet interface 0 only. In addition, this example disables automatic 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 52](#) shows the scenario.

Figure 52 EIGRP Route Authentication Scenario



Router A Configuration

```
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 Configuration

```
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.

Stub Routing Example

A router that is configured as a stub with the **eigrp stub** command shares connected and summary routing information with all neighbor routers by default. Four optional keywords can be used with the **eigrp stub** command to modify this behavior:

- **receive-only**
- **connected**
- **static**
- **summary**

This section provides configuration examples for all forms of the **eigrp stub** command. The **eigrp stub** command can be modified with several options, and these options can be used in any combination except for the **receive-only** keyword. The **receive-only** keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system, and the **receive-only** keyword will not permit any other option to be specified because it prevents any type of route from being sent. The three other optional keywords (**connected**, **static**, and **summary**) can be used in any combination but cannot be used with the **receive-only** keyword. If any of these three keywords is used individually with the **eigrp stub** command, connected and summary routes will not be sent automatically.

The **connected** keyword will permit the EIGRP Stub Routing feature to send connected routes. If the connected routes are not covered by a network statement, it may be necessary to redistribute connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default.

The **static** keyword will permit the EIGRP Stub Routing feature to send static routes. Without this option, EIGRP will not send any static routes, including internal static routes that normally would be automatically redistributed. It will still be necessary to redistribute static routes with the **redistribute static** command.

The **summary** keyword will permit the EIGRP Stub Routing feature to send summary routes. Summary routes can be created manually with the **summary address** command or automatically at a major network border router with the **auto-summary** command enabled. This option is enabled by default.

In the following example, the **eigrp stub** command is used to configure the router as a stub that advertises connected and summary routes:

```
router eigrp 1
network 10.0.0.0
eigrp stub
```

In the following example, the **eigrp stub connected static** command is used to configure the router as a stub that advertises connected and static routes (sending summary routes will not be permitted):

```
router eigrp 1
network 10.0.0.0
eigrp stub connected static
```

In the following example, the **eigrp stub receive-only** command is used to configure the router as a stub, and connected, summary, or static routes will not be sent:

```
router eigrp 1
network 10.0.0.0 eigrp
stub receive-only
```