

Configuring IP Enhanced IGRP

This chapter describes how to configure Internet Protocol (IP) enhanced IGRP on routers and interfaces configured for IP. For a complete description of the commands mentioned in this chapter, refer to the “IP Enhanced IGRP Commands” chapter in this publication. For a description of other IP configuration commands, refer to the *Router Products Configuration Guide* and *Router Products Command Reference* publications. For historical background and a technical overview of IP routing protocols, see the *Internetworking Technology Overview* publication.

Cisco’s Implementation of IP Enhanced IGRP

IP enhanced IGRP provides the following features:

- Automatic redistribution. IP IGRP routes can be automatically redistributed into enhanced IGRP, and IP enhanced IGRP routes can be automatically redistributed into IGRP. If desired, you can turn off redistribution. You can also completely turn off IP enhanced IGRP and IP 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 IP enhanced IGRP is enabled, the largest possible width is 224 hops. Because the enhanced IGRP 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 via enhanced IGRP. When a RIP route is being used as the next hop to the destination, the transport control field is incremented as usual.

IP Enhanced IGRP Configuration Task List

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

- Enable IP Enhanced IGRP
- Transition from IGRP to Enhanced IGRP
- Configure IP Enhanced IGRP-Specific Parameters
- Configure IP Enhanced IGRP-Specific Parameters
- Monitor IP Enhanced IGRP on an IP Network

See the end of this chapter for configuration examples.

Enable IP Enhanced IGRP

To create an IP enhanced IGRP routing process, perform the following tasks:

Task	Command
Step 1 Enable an IP enhanced IGRP routing process in global configuration mode.	router eigrp <i>autonomous-system-number</i>
Step 2 Associate networks with an IP enhanced IGRP routing process in router configuration mode.	network <i>network-number</i>

IP enhanced IGRP sends updates to the interfaces in the specified network(s). If you do not specify an interface’s network, it will not be advertised in any IP enhanced IGRP update.

Transition from IGRP to Enhanced IGRP

If you have routers on your network that are configured for IGRP and you want to make a transition to routing enhanced IGRP, you need to designate transition routers that have both IGRP and enhanced IGRP configured. In these cases, perform the tasks as noted in the previous section, “Create the IP enhanced IGRP Routing Process,” and also read the section on configuring IGRP in the *Router Products Configuration Guide*. You must use the same autonomous system number in order for routes to be redistributed automatically.

Configure IP Enhanced IGRP-Specific Parameters

To configure IP enhanced IGRP-specific parameters, perform one or more of the following tasks:

- Define Unequal-Cost Load Balancing
- Adjust the IP Enhanced IGRP Metric Weights
- Disable Route Summarization
- Configure Summary Aggregate Addresses

Define Unequal-Cost Load Balancing

IP enhanced IGRP can simultaneously use an asymmetric set of paths for a given destination. This feature is known as *unequal-cost load balancing*. Unequal-cost load balancing allows traffic to be distributed among up to four unequal-cost paths to provide greater overall throughput and reliability. Alternate path variance (the difference in desirability between the primary and alternate paths) is used to determine the feasibility of a potential route. An alternate route is feasible if the next router in the path is closer to the destination (has a lower metric value) than the current router and if the metric for the entire alternate path is within the variance. Only paths that are feasible can be used for load balancing and included in the routing table. These conditions limit the number of cases in which load balancing can occur, but ensure that the dynamics of the network will remain stable.

The following general rules apply to IP enhanced IGRP unequal-cost load balancing:

- IP enhanced IGRP will accept up to four paths for a given destination network.
- The local best metric must be greater than the metric learned from the next router; that is, the next-hop router must be closer (have a smaller metric value) to the destination than the local best metric.

- The alternative path metric must be within the specified *variance* of the local best metric. The multiplier times the local best metric for the destination must be greater than or equal to the metric through the next router.

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

By default, the amount of variance is set to one (equal-cost load balancing). To change the variance to define how much worse an alternate path can be before that path is disallowed, perform the following task in router configuration mode:

Task	Command
Define the variance associated with a particular path.	variance multiplier

See the “IP Enhanced IGRP Configuration Examples” section at the end of this chapter for an example of configuring an IP enhanced IGRP feasible successor.

Note By using the variance feature, the router can balance traffic across all feasible paths and can immediately converge to a new path if one of the paths should fail.

Adjust the IP Enhanced IGRP Metric Weights

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

To adjust the IP enhanced IGRP metric weights, perform the following task in router configuration mode:

Task	Command
Adjust the IP enhanced IGRP metric.	metric weights tos k1 k2 k3 k4 k5

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 IP enhanced IGRP 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.

Disable Route Summarization

You can configure IP enhanced IGRP 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 IP enhanced IGRP process. By default, this feature is enabled.

To disable automatic summarization, perform the following task in router configuration mode:

Task	Command
Disable automatic summarization.	no auto-summary

Route summarization works in conjunction with the **ip summary-address eigrp** interface configuration command, in which additional summarization can be performed. If auto-summary 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, IP enhanced IGRP 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, perform the following task in interface configuration mode:

Task	Command
Configure a summary aggregate address.	ip summary-address eigrp <i>autonomous-system-number</i> <i>address mask</i>

Configure Protocol-Independent Parameters

To configure protocol-independent parameters, perform one or more of the following tasks:

- Redistribute Routing Information
- Set Metrics for Redistributed Routes
- Filter Routing Information
- Adjust the Interval between Hello Packets and the Hold Time
- Disable Split Horizon

Redistribute Routing Information

In addition to running multiple routing protocols simultaneously, the router can redistribute information from one routing protocol to another. For example, you can instruct the router to readvertise IP enhanced IGRP-derived routes using the RIP protocol, or to readvertise static routes using the IP enhanced IGRP protocol. This capability applies to all the IP-based routing protocols.

You may also conditionally control the redistribution of routes between routing domains by defining a method known as *route maps* between the two domains.

To redistribute routes from one protocol into another, perform the following task in router configuration mode:

Task	Command
Redistribute routes from one routing protocol into another.	redistribute <i>protocol autonomous-system-number</i> [route-map <i>map-tag</i>]

To define route maps, perform the following task in global configuration mode:

Task	Command
Define any route maps needed to control redistribution.	route-map <i>map-tag</i> { permit deny } <i>sequence-number</i>

By default, the redistribution of default information between IP enhanced IGRP processes is enabled. To disable the redistribution, perform the following task in router configuration mode:

Task	Command
Disable the redistribution of default information between IP enhanced IGRP processes.	no default-information allowed { in out }

See the “IP Enhanced IGRP Configuration Examples” section at the end of this chapter for examples of configuring redistribution and route maps.

Set Metrics for Redistributed Routes

The metrics of one routing protocol do not necessarily translate into the metrics of another. For example, the RIP metric is a hop count and the IP enhanced IGRP metric is a combination of five quantities. In such situations, an artificial metric is assigned to the redistributed route. Because of this unavoidable tampering with dynamic information, carelessly exchanging routing information between different routing protocols can create routing loops, which can seriously degrade network operation.

To set metrics for redistributed routes, perform the first task when redistributing from IP enhanced IGRP, and perform the second task when redistributing into IP enhanced IGRP. Each task is done in router configuration mode.

Task	Command
Cause the current routing protocol to use the same metric value for all redistributed routes.	default-metric <i>number</i>
Cause the IP enhanced IGRP routing protocol to use the same metric value for all non-IGRP redistributed routes.	default-metric <i>bandwidth delay reliability loading mtu</i>

Filter Routing Information

You can filter routing protocol information by performing the following tasks:

- Suppress the sending of routing updates on a particular router interface. Doing so prevents other systems on an interface from learning about routes dynamically.
- Suppress networks from being advertised in routing updates. Doing so prevents other routers from learning a particular router’s interpretation of one or more routes.
- Suppress a routing protocol from both sending and receiving updates on a particular interface. You usually perform this task when a wildcard command has been used to configure the routing protocol for more router interfaces than is desirable.
- Suppress networks listed in updates from being accepted and acted upon by a routing process. Doing so keeps a router from using certain routes.
- Filter on the source of routing information. You perform this task to prioritize routing information from different sources, because the accuracy of the routing information can vary.
- Apply an offset to routing metrics. Doing so provides a local mechanism for increasing the value of routing metrics.

Use the information in the following sections to perform these tasks.

Prevent Routing Updates through an Interface

To prevent other routers on a local network from learning about routes dynamically, you can keep routing update messages from being sent through a router interface. This feature applies to all IP-based routing protocols except BGP and EGP.

To prevent routing updates through a specified interface, perform the following task in router configuration mode:

Task	Command
Suppress the sending of routing updates through a router interface.	passive-interface <i>type unit</i>

Control the Advertising of Routes in Routing Updates

To control which routers learn about routes, you can control the advertising of routes in routing updates. To do this, perform the following task in router configuration mode:

Task	Command
Control the advertising of routes in routing updates.	distribute-list <i>access-list-number out [interface-name routing-process autonomous-system-number]</i>

Control the Processing of Routing Updates

To control the processing of routes listed in incoming updates, perform the following task in router configuration mode:

Task	Command
Control which incoming route updates are processed.	distribute-list <i>access-list-number in [interface-name]</i>

Apply Offsets to Routing Metrics

To provide a local mechanism for increasing the value of routing metrics, you can apply an offset to routing metrics. To do so, perform the following task in router configuration mode:

Task	Command
Apply an offset to routing metrics.	offset-list {in out} <i>offset</i> [<i>access-list-number</i>]

Filter Sources of Routing Information

An administrative distance is a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers. In a large network, some routing protocols and some routers can be more reliable than others as sources of routing information. Also, when multiple routing processes are running in the same router for IP, the same route may be advertised by more than one routing process. Specifying administrative distance values enables the router to discriminate between sources of routing information. The router always picks the route whose routing protocol has the lowest administrative distance.

There are no general guidelines for assigning administrative distances, because each network has its own requirements. You must determine a reasonable matrix of administrative distances for the network as a whole. Table 3-1 shows the default administrative distance for various routing information sources.

Table 3-1 Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IGRP summary route	5
External BGP	20
Internal enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
External enhanced IGRP	170
Internal BGP	200
Unknown	255

For example, consider a router using IP enhanced IGRP and RIP. Suppose you trust the IP enhanced IGRP-derived routing information more than the RIP-derived routing information. Because the default IP enhanced IGRP administrative distance is lower than that for RIP, the router uses the IP enhanced IGRP-derived information and ignores the RIP-derived information. However, if you lose the source of the IP enhanced IGRP-derived information (for example, because of a power shutdown), the router uses the RIP-derived information until the IP enhanced IGRP-derived information reappears.

Note You can also use administrative distance to rate the routing information from routers running the same routing protocol. This application is generally discouraged if you are unfamiliar with this particular use of administrative distance, since it can result in inconsistent routing information, including forwarding loops.

To filter sources of routing information, perform the following tasks in router configuration mode:

Task	Command
Filter on routing information sources.	distance eigrp <i>internal-distance external-distance</i>

Adjust the Interval between Hello Packets and the Hold Time

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

Routers periodically send hello packets to each other to dynamically learn of other routers on their directly attached networks. The routers use this information 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.

You can configure the hold time on a specified interface for the IP enhanced IGRP 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.

To change the interval between hello packets, perform the following task in interface configuration mode:

Task	Command
Configure the hello interval for an IP enhanced IGRP routing process.	ip hello-interval eigrp <i>autonomous-system-number seconds</i>

On very congested and large networks, 15 seconds may 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, perform the following task in interface configuration mode:

Task	Command
Configure the hold time for an IP enhanced IGRP routing process.	ip hold-time eigrp <i>autonomous-system-number seconds</i>

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

Disable Split Horizon

Split horizon controls the sending of IP enhanced IGRP 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 by default on all interfaces.

Split horizon blocks information about routes from being advertised by a router out any interface from which that information originated. This behavior usually optimizes communications among multiple routers, 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 wish to disable split horizon.

To disable split horizon, perform the following task in interface configuration mode:

Task	Command
Disable split horizon.	no ip split-horizon eigrp <i>autonomous-system-number</i>

See the “IP Enhanced IGRP Configuration Examples” section at the end of this chapter for an example of using split horizon.

Monitor IP Enhanced IGRP on an IP Network

You can display router statistics such as the contents of IP routing tables, caches, and databases. You can use the information displayed to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your router’s packets are taking through the network.

To display various router statistics, perform one or more of the following tasks at the EXEC prompt:

Task	Command
Delete neighbors from the neighbor table.	clear ip eigrp neighbors [<i>ip-address</i> <i>interface</i>]
Display the IP enhanced IGRP discovered neighbors.	show ip eigrp neighbors [<i>interface unit</i>]
Display the IP enhanced IGRP topology table for a given process.	show ip eigrp topology [<i>autonomous-system-number</i> [[<i>ip-address</i>] <i>mask</i>]]
Display the number of packets sent and received for all or a specified IP enhanced IGRP process.	show ip eigrp traffic [<i>autonomous-system-number</i>]
Display the parameters and current state of the active routing protocol process.	show ip protocols
Display the current state of the routing table.	show ip route [<i>ip-address</i> [<i>mask</i>]] [<i>protocol</i> [<i>process-id</i>]]
Display the current state of the routing table in summary form.	show ip route summary

IP Enhanced IGRP Configuration Examples

The following sections provide IP enhanced IGRP configuration examples:

- Overriding Static Routes with Dynamic Protocols Example
- Static Routing Redistribution Example
- IP Enhanced IGRP Redistribution Examples
- RIP and IP Enhanced IGRP Redistribution Example
- Default Metric Values Redistribution Example
- Route Map Example

- Route Summarization Example
- IP Enhanced IGRP Feasible Successor Relationship Example
- Split Horizon Example

Overriding Static Routes with Dynamic Protocols Example

In the following example, packets for network 10.0.0.0 from Router B, where the static route is installed, will be routed through 131.108.3.4 if a route with an administrative distance less than 110 is not available. Figure 3-1 illustrates this point. The route learned by a protocol with an administrative distance less than 110 may cause Router B to send traffic destined for network 10.0.0.0 via the alternate path—through Router D.

```
ip route 10.0.0.0 255.0.0.0 131.108.3.4 110
```




Figure 3-1 Overriding Static Routes

Static Routing Redistribution Example

In the following example, three static routes are specified, two of which are to be advertised. You can advertise them by specifying the **redistribute static** router configuration command, then specifying an access list that allows only those two networks to be passed to the IP enhanced IGRP process. Any redistributed static routes should be sourced by a single router to minimize the likelihood of creating a routing loop.

```
ip route 192.1.2.0 255.255.255.0 192.31.7.65
ip route 193.62.5.0 255.255.255.0 192.31.7.65
ip route 131.108.0.0 255.255.0.0 192.31.7.65
access-list 3 permit 192.1.2.0
access-list 3 permit 193.62.5.0
!
router eigrp 109
network 192.31.7.0
redistribute static route-map static-to-eigrp

route-map static-to-eigrp
match ip address 3
set metric 10000 100 255 1 1500
```

IP Enhanced IGRP Redistribution Examples

Each IP enhanced IGRP routing process can provide routing information to only one autonomous system; the router must run a separate IP enhanced IGRP process and maintain a separate routing database for each autonomous system it services. However, you can transfer routing information between these routing databases.

Suppose the router has one IP enhanced IGRP routing process for network 15.0.0.0 in autonomous system 71 and another for network 192.31.7.0 in autonomous system 109, as the following commands specify:

```
router eigrp 71
network 15.0.0.0
router eigrp 109
network 192.31.7.0
```

To transfer a route from 192.31.7.0 into autonomous system 71 (without passing any other information about autonomous system 109), use the command in the following example:

```
router eigrp 71
redistribute eigrp 109 route-map 109-to-71
route-map 109-to-71 permit
match ip address 3
set metric 10000 100 1 255 1500
access-list 3 permit 192.31.7.0
```

The following example is an alternative way to transfer a route to 192.31.7.0 into autonomous system 71. Unlike the previous configuration, this one does not allow you to arbitrarily set the metric.

```
router eigrp 71
redistribute eigrp 109
distribute-list 3 out eigrp 109
access-list 3 permit 192.31.7.0
```

RIP and IP Enhanced IGRP Redistribution Example

This section provides two examples of RIP and IP enhanced IGRP redistribution, a simple one and a complex one.

Example 1: Simple Redistribution

Consider a wide-area network at a university that uses RIP as an interior routing protocol. Assume that the university wants to connect its wide-area network to a regional network, 128.1.0.0, which uses IP enhanced IGRP as the routing protocol. The goal in this case is to advertise the networks in the university network to the routers on the regional network. The commands for the interconnecting router are listed in the example that follows:

```
router eigrp 109
network 128.1.0.0
redistribute rip
default-metric 10000 100 255 1 1500
distribute-list 10 out rip
```

In this example, the **router** global configuration command starts an IP enhanced IGRP routing process. The **network** router configuration command specifies that network 128.1.0.0 (the regional network) is to send and receive IP enhanced IGRP routing information. The **redistribute** router configuration command specifies that RIP-derived routing information be advertised in the routing updates. The **default-metric** router configuration command assigns an IP enhanced IGRP metric to all RIP-derived routes.

The **distribute-list** router configuration command instructs the router to use access list 10 (not defined in this example) to limit the entries in each outgoing update. The access list prevents unauthorized advertising of university routes to the regional network.

Example 2: Complex Redistribution

The most complex redistribution case is one in which *mutual* redistribution is required between an IGP (in this case IP enhanced IGRP) and BGP.

Suppose that BGP is running on a router somewhere else in AS 1, and that the BGP routes are injected into IP enhanced IGRP routing process 1. You must use filters to ensure that the proper routes are advertised. The example configuration for router R1 illustrates use of access filters and a distribution list to filter routes advertised to BGP neighbors. This example also illustrates configuration commands for redistribution between BGP and IP enhanced IGRP.

```
! Configuration for router R1:
router bgp 1
network 131.108.0.0
neighbor 192.5.10.1 remote-as 2
neighbor 192.5.10.15 remote-as 1
neighbor 192.5.10.24 remote-as 3
redistribute eigrp 1
distribute-list 1 out eigrp 1
!
! All networks that should be advertised from R1 are controlled with access lists:
!
access-list 1 permit 131.108.0.0
access-list 1 permit 150.136.0.0
access-list 1 permit 128.125.0.0
!
router eigrp 1
network 131.108.0.0
network 192.5.10.0
redistribute bgp 1
```

Default Metric Values Redistribution Example

The following example shows a router in autonomous system 1 using both the OSPF and the IP enhanced IGRP routing protocols. The example advertises OSPF-derived routes using the IP enhanced IGRP protocol and assigns the OSPF-derived routes an IP enhanced IGRP metric of 1000 100 255 1 1500.

```
router eigrp 1
network 131.108.0.0
redistribute ospf 1
default-metric 1000 100 255 1 1500

router ospf 1
network 160.89.0.0 0.0.255.255 area 0.0.0.0
```

Route Map Example

The following example redistributes all OSPF routes into IP enhanced IGRP:

```
router eigrp 109
redistribute ospf 110
default-metric 1000 100 255 1 1500
```

Route Summarization Example

The following example configures route summarization on the interface and also configures the auto-summary feature. This configuration causes IP enhanced IGRP to summarize network 10.0.0.0 out interface Ethernet 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 131.108.0.0
no auto-summary
```

IP Enhanced IGRP Feasible Successor Relationship Example

In Figure 3-2, the assigned metrics meet the conditions required for a feasible successor relationship, so the paths in this example can be included in routing tables and used for load balancing.




Figure 3-2 Assigning Metrics for IP Enhanced IGRP Path Feasibility

The feasibility test would work as follows:

Assume that Router C1 already has a route to Network A with metric m and has just received an update about Network A from C2. The best metric at C2 is p . The metric that C1 would use through C2 is n .

If the following two conditions are met, the route to Network A through C2 will be included in C1's routing table:

- 1 If m is greater than p .
- 2 If the *multiplier* (value specified by the **variance** router configuration command) times m is greater than or equal to n .

The configuration for Router C1 would be as follows:

```
router eigrp 109
variance 10
```

A maximum of four paths for a single destination can be present in the routing table. If there are more than four feasible paths, the four best feasible paths are used.

Split Horizon Example

Figure 3-3 illustrates a configuration in which disabling split horizon would be useful. In this configuration, two IP subnets are both accessible via a serial interface on Router C (connected to Frame Relay network). The serial interface on Router C accommodates one of the subnets via the assignment of a secondary IP address. Split horizon must be disabled in order for network

128.125.0.0 to be advertised into network 131.108.0.0, and vice versa. These subnets overlap at Router C, interface S0. If split horizon were enabled on serial interface S0, it would not advertise a route back into the Frame Relay network for either of these networks.

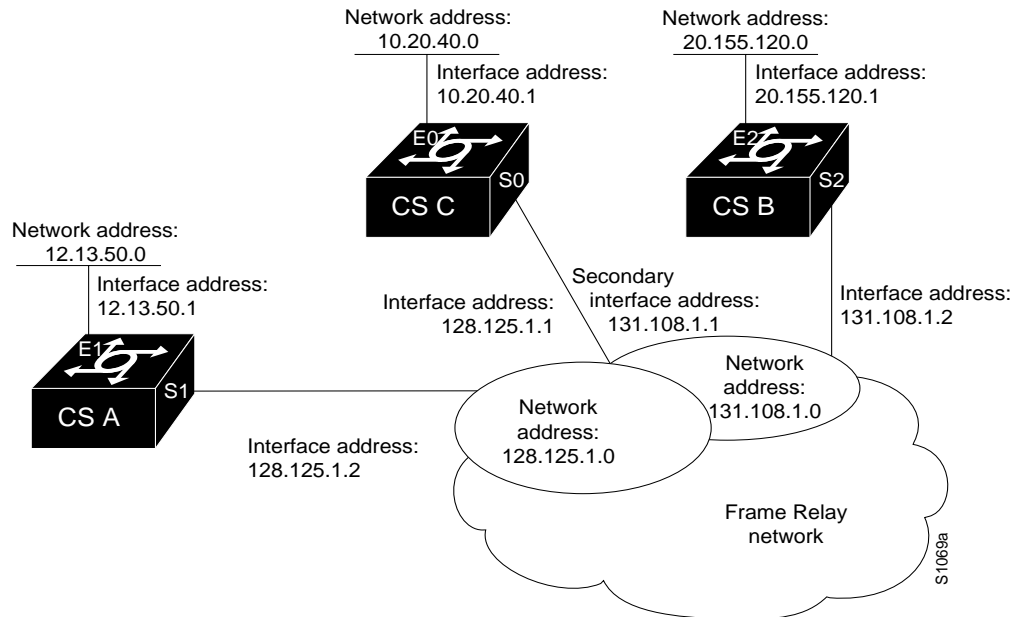


Figure 3-3 Disabled Split Horizon Example

Configuration for Router A

```
interface ethernet 1
ip address 12.13.50.1
!
interface serial 1
ip address 128.125.1.2
encapsulation frame-relay
```

Configuration for Router B

```
interface ethernet 2
ip address 20.155.120.1
!
interface serial 2
ip address 131.108.1.2
encapsulation frame-relay
```

Configuration for Router C

```
interface ethernet 0
ip address 10.20.40.1
no ip split-horizon eigrp
!
interface serial 0
ip address 128.124.1.1
ip address 131.108.1.1 secondary
encapsulation frame-relay
```