Configuring IGRP

This chapter describes how to configure the Interior Gateway Routing Protocol (IGRP). For a complete description of the IGRP commands in this chapter, refer to the "IGRP 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.

IGRP is a dynamic distance-vector routing protocol designed by Cisco in the mid-1980s for routing in an autonomous system that contains large, arbitrarily complex networks with diverse bandwidth and delay characteristics.

For protocol-independent features, see the chapter "Configuring IP Routing Protocol-Independent Features" in this document.

Cisco's IGRP Implementation

IGRP uses a combination of user-configurable metrics, including internetwork delay, bandwidth, reliability, and load.

IGRP also advertises three types of routes: interior, system, and exterior, as shown in Figure 17. *Interior routes* are routes between subnets in the network attached to a router interface. If the network attached to a router is not subnetted, IGRP does not advertise interior routes.

System routes are routes to networks within an autonomous system. The Cisco IOS software derives system routes from directly connected network interfaces and system route information provided by other IGRP-speaking routers or access servers. System routes do not include subnet information.

Exterior routes are routes to networks outside the autonomous system that are considered when identifying a *gateway of last resort*. The Cisco IOS software chooses a gateway of last resort from the list of exterior routes that IGRP provides. The software uses the gateway (router) of last resort if it does not have a better route for a packet and the destination is not a connected network. If the autonomous system has more than one connection to an external network, different routers can choose different exterior routers as the gateway of last resort.

Figure 17 Interior, System, and Exterior Routes

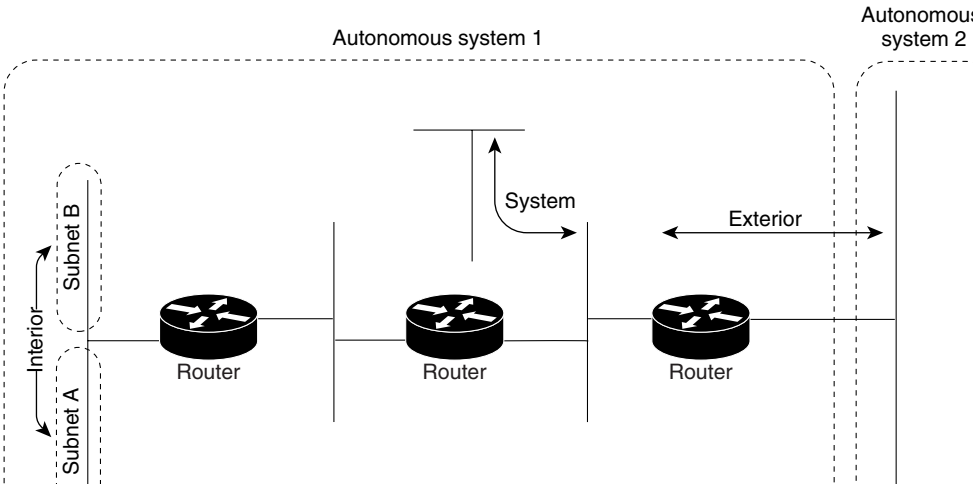

IGRP Updates

By default, a router running IGRP sends an update broadcast every 90 seconds. It declares a route inaccessible if it does not receive an update from the first router in the route within 3 update periods (270 seconds). After 7 update periods (630 seconds), the Cisco IOS software removes the route from the routing table.

IGRP uses *flash update* and *poison reverse updates* to speed up the convergence of the routing algorithm. Flash update is the sending of an update sooner than the standard periodic update interval of notifying other routers of a metric change. Poison reverse updates are intended to defeat larger routing loops caused by increases in routing metrics. The poison reverse updates are sent to remove a route and place it in *holddown*, which keeps new routing information from being used for a certain period of time.

IGRP Configuration Task List

To configure IGRP, perform the tasks in the following sections. Creating the IGRP routing process is mandatory; the other tasks are optional.

- Create the IGRP Routing Process
- Apply Offsets to Routing Metrics
- Allow Unicast Updates for IGRP
- Define Unequal-Cost Load Balancing
- Control Traffic Distribution
- Adjust the IGRP Metric Weights
- Adjust Timers

- Disable Holddown
- Enforce a Maximum Network Diameter
- Validate Source IP Addresses
- Enable or Disable Split Horizon

Also see the examples in the "IGRP Configuration Examples" section at the end of this chapter.

Create the IGRP Routing Process

To create the IGRP routing process, perform the following required tasks starting in global configuration mode:

Task		Command
Step 1	Enable an IGRP routing process, which places you in router configuration mode.	**router igrp** *autonomous-system*
Step 2	Associate networks with an IGRP routing process.	**network** *network-number*

IGRP sends updates to the interfaces in the specified networks. If an interface's network is not specified, it will not be advertised in any IGRP update.

It is not necessary to have a registered autonomous system number to use IGRP. If you do not have a registered number, you are free to create your own. We recommend that if you do have a registered number, you use it to identify the IGRP process.

Apply Offsets to Routing Metrics

An offset list is the mechanism for increasing incoming and outgoing metrics to routes learned via IGRP. 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, perform the following task in router configuration mode:

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

Allow Unicast Updates for IGRP

Because IGRP is normally a broadcast protocol, in order for IGRP routing updates to reach nonbroadcast networks, you must configure the Cisco IOS software to permit this exchange of routing information.

To permit information exchange, perform the following task in router configuration mode:

Task	Command
Define a neighboring router with which to exchange routing information.	**neighbor** *ip-address*

To control the set of interfaces with which you want to exchange routing updates, you can disable the sending of routing updates on specified interfaces by configuring the **passive-interface** command. See the discussion on filtering in the "Filter Routing Information" section in the "Configuring IP Routing Protocol-Independent Features" chapter.

Define Unequal-Cost Load Balancing

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 multiple (up to four) unequal-cost paths to provide greater overall throughput and reliability. Alternate *path variance* (that is, 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 IGRP unequal-cost load balancing:

- 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). You can define how much worse an alternate path can be before that path is disallowed by performing the following task in router configuration mode:

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

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

See the "IGRP Feasible Successor Relationship Example" at the end of this chapter.

Control Traffic Distribution

If variance is configured as described in the preceding section "Define Unequal-Cost Load Balancing," IGRP or Enhanced IGRP will distribute traffic among multiple routes of unequal cost to the same destination. If you want to have faster convergence to alternate routes, but you do not want to send traffic across inferior routes in the normal case, you might prefer to have no traffic flow along routes with higher metrics.

To control how traffic is distributed among multiple routes of unequal cost, perform the following task in router configuration mode:

Task	Command
Distribute traffic proportionately to the ratios of metrics, or by the minimum-cost route.	**traffic-share** {**balanced** | **min**}

Adjust the IGRP Metric Weights

You have the option of altering the default behavior of IGRP routing and metric computations. This allows, for example, tuning system behavior to allow for transmissions via satellite. Although IGRP metric defaults were carefully selected to provide excellent operation in most networks, you can adjust the IGRP metric. Adjusting IGRP metric weights can dramatically affect network performance, however, so ensure that you make all metric adjustments carefully.

To adjust the IGRP metric weights, perform the following task in router configuration mode. Because of the complexity of this task, we recommend that you only perform it with guidance from an experienced system designer.

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

By default, the IGRP composite metric is a 24-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.

Adjust Timers

Routing protocols use several timers that determine such variables as the frequency of routing updates, the length of time before a route becomes invalid, and other parameters. You can adjust these timers to tune routing protocol performance to better suit your internetwork needs. You can make the following timer adjustments:

- The rate (time in seconds between updates) at which routing updates are sent
- The interval of time (in seconds) after which a route is declared invalid
- The interval (in seconds) during which routing information regarding better paths is suppressed
- The amount of time (in seconds) that must pass before a route is removed from the routing table
- The amount of time for which routing updates will be postponed

It also is possible to tune the IP routing support in the software to enable faster convergence of the various IP routing algorithms, and, hence, quicker fallback to redundant routers. The total effect is to minimize disruptions to end users of the network in situations where quick recovery is essential.

To adjust the timers, perform the following task in router configuration mode:

Task	Command
Adjust routing protocol timers.	**timers basic** *update invalid holddown flush* [*sleeptime*]

Disable Holddown

When the Cisco IOS software learns that a network is at a greater distance than was previously known, or it learns the network is down, the route to that network is placed in holddown. During the holddown period, the route is advertised, but incoming advertisements about that network from any router other than the one that originally advertised the network's new metric will be ignored. This mechanism is often used to help avoid routing loops in the network, but has the effect of increasing the topology convergence time.

To disable holddowns with IGRP, perform the following task in router configuration mode. All devices in an IGRP autonomous system must be consistent in their use of holddowns.

Task	Command
Disable the IGRP holddown period.	**no metric holddown**

Enforce a Maximum Network Diameter

The Cisco IOS software enforces a maximum diameter to the IGRP network. Routes whose hop counts exceed this diameter are not advertised. The default maximum diameter is 100 hops. The maximum diameter is 255 hops.

To configure the maximum diameter, perform the following task in router configuration mode:

Task	Command
Configure the maximum network diameter.	**metric maximum-hops** *hops*

Validate Source IP Addresses

To disable the default function that validates the source IP addresses of incoming routing updates, perform the following task in router configuration mode:

Task	Command
Disable the checking and validation of the source IP address of incoming routing updates.	**no validate-update-source**

Enable or Disable Split Horizon

Normally, routers that are connected to broadcast-type IP networks and that use distance-vector routing protocols employ the *split horizon* mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router out of 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 might want to disable split horizon. This applies to IGRP and RIP.

If an interface is configured with secondary IP addresses and split horizon is enabled, updates might not be sourced by every secondary address. One routing update is sourced per network number unless split horizon is disabled.

To enable or disable split horizon, perform the following tasks in interface configuration mode:

Task	Command
Enable split horizon.	**ip split-horizon**

Task	Command
Disable split horizon.	**no ip split-horizon**

Split horizon for Frame Relay and SMDS encapsulation is disabled by default. Split horizon is not disabled by default for interfaces using any of the X.25 encapsulations. For all other encapsulations, split horizon is enabled by default.

See the "Split Horizon Examples" section at the end of this chapter for examples of using split horizon.

Note In general, changing the state of the default is not recommended unless you are certain that your application requires making a change in order to advertise routes properly. Remember: If split horizon is disabled on a serial interface (and that interface is attached to a packet-switched network), you *must* disable split horizon for all routers in any relevant multicast groups on that network.

IGRP Configuration Examples

This section contains the following IGRP configuration examples:

- IGRP Feasible Successor Relationship Example
- Split Horizon Examples

IGRP Feasible Successor Relationship Example

In Figure 18, 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 18 Assigning Metrics for IGRP Path Feasibility

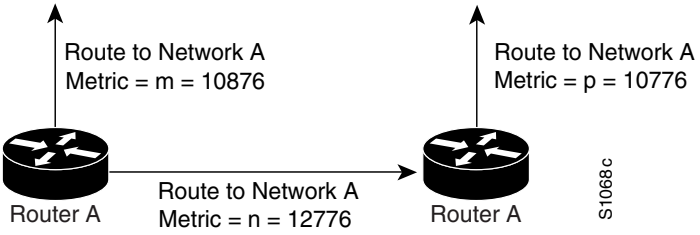

The feasibility test would work as follows:

- 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 both of the following two conditions are met, the route to network A through C2 will be included in C1's routing table:
  - If *m* is greater than *p*.
  - 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 igrp 109
 variance 10
```

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

Split Horizon Examples

The following sample configuration illustrates a simple example of disabling split horizon on a serial link. In this example, the serial link is connected to an X.25 network.

```
interface serial 0
 encapsulation x25
 no ip split-horizon
```

In the next example, Figure 19 illustrates a typical situation in which the **no ip split-horizon** interface configuration command would be useful. This figure depicts two IP subnets that are both accessible via a serial interface on Router C (connected to Frame Relay network). In this example, the serial interface on Router C accommodates one of the subnets via the assignment of a secondary IP address.

The Ethernet interfaces for Router A, Router B, and Router C (connected to IP networks 12.13.50.0, 10.20.40.0, and 20.155.120.0) all have split horizon *enabled* by default, while the serial interfaces connected to networks 128.125.1.0 and 131.108.1.0 all have split horizon *disabled* by default. The partial interface configuration specifications for each router that follow Figure 19 illustrate that the **ip split-horizon** command is *not* explicitly configured under normal conditions for any of the interfaces.

Figure 19 Disabled Split Horizon Example for Frame Relay Network

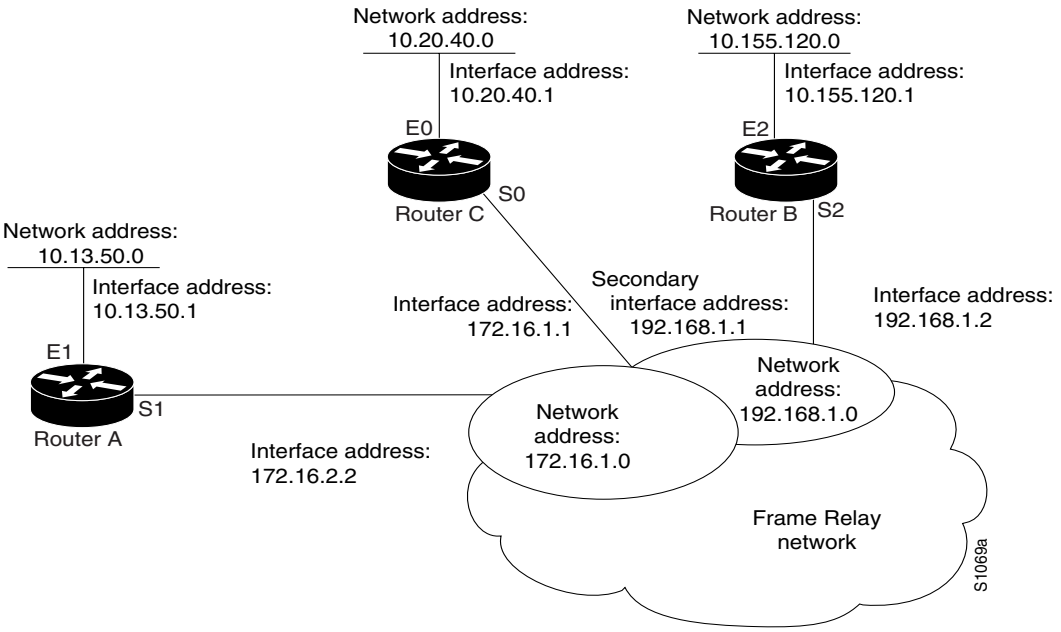

In this example, split horizon must be disabled in order for network 128.125.0.0 to be advertised into network 131.108.0.0, and vice versa. These subnets overlap at Router C, interface S0. If split horizon were enabled on serial interface S0, it would not advertise a route back into the Frame Relay network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

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
!
interface serial 0
 ip address 128.124.1.1
 ip address 131.108.1.1 secondary
 encapsulation frame-relay
```