



## Chapter Goals

- Identify the metrics that IGRP uses to compare routes.
- Explain how an administrator can influence route selection.
- Understand multipath routing.
- Identify IGRP's native stability features.
- Identify IGRP's timing mechanisms and explain their purpose.

# Interior Gateway Routing Protocol

---

## Background

The *Interior Gateway Routing Protocol (IGRP)* is a routing protocol that was developed in the mid-1980s by Cisco Systems, Inc. Cisco's principal goal in creating IGRP was to provide a robust protocol for routing within an autonomous system (AS). Such protocols are known as Interior Gateway Routing Protocols.

In the mid-1980s, the most popular Interior Gateway Routing Protocol was the Routing Information Protocol (RIP). Although RIP was quite useful for routing within small- to moderate-sized, relatively homogeneous internetworks, its limits were being pushed by network growth. In particular, RIP's small hop-count limit (16) restricted the size of internetworks; single metric (hop count) support of only equal-cost load balancing (in all-Cisco networks only!) did not allow for much routing flexibility in complex environments. The popularity of Cisco routers and the robustness of IGRP encouraged many organizations with large internetworks to replace RIP with IGRP.

Cisco's initial IGRP implementation worked in Internet Protocol (IP) networks. IGRP was designed to run in any network environment, however, and Cisco soon ported it to run in OSI Connectionless-Network Protocol (CLNP) networks. Cisco developed Enhanced IGRP in the early 1990s to improve the operating efficiency of IGRP. This chapter discusses IGRP's basic design and implementation. Enhanced IGRP is discussed in Chapter 40, "Enhanced IGRP."

## IGRP Protocol Characteristics

IGRP is a *distance vector* Interior Gateway Protocol (IGP). Distance vector routing protocols mathematically compare routes using some measurement of distance. This measurement is known as the distance vector. Routers using a distance vector protocol must send all or a portion of their routing table

in a routing-update message at regular intervals to each of their neighboring routers. As routing information proliferates through the network, routers can identify new destinations as they are added to the network, learn of failures in the network, and, most importantly, calculate distances to all known destinations.

Distance vector routing protocols are often contrasted with link-state routing protocols, which send local connection information to all nodes in the internetwork. For a discussion of Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS), two popular link-state routing algorithms, see Chapter 46, “Open Shortest Path First,” and Chapter 45, “Open System Interconnection Routing Protocol,” respectively.

IGRP uses a composite metric that is calculated by factoring weighted mathematical values for internetwork delay, bandwidth, reliability, and load. Network administrators can set the weighting factors for each of these metrics, although great care should be taken before any default values are manipulated. IGRP provides a wide range for its metrics. Reliability and load, for example, can take on any value between 1 and 255; bandwidth can take on values reflecting speeds from 1200 bps to 10 Gbps, while delay can take on any value from 1 to  $2^{24}$ . These wide metric ranges are further complemented by a series of user-definable constants that enable a network administrator to influence route selection. These constants are hashed against the metrics, and each other, in an algorithm that yields a single, composite metric. Thus, the network administrator can influence route selection by giving higher or lower weighting to specific metrics. This flexibility allows administrators to fine-tune IGRP’s automatic route selection.

To provide additional flexibility, IGRP permits multipath routing. Dual equal-bandwidth lines can run a single stream of traffic in round-robin fashion, with automatic switchover to the second line if one line goes down. Multiple paths can have unequal metrics yet still be valid multipath routes. For example, if one path is three times better than another path (its metric is three times lower), the better path will be used three times as often. Only routes with metrics that are within a certain range or variance of the best route are used as multiple paths. Variance is another value that can be established by the network administrator.

## Stability Features

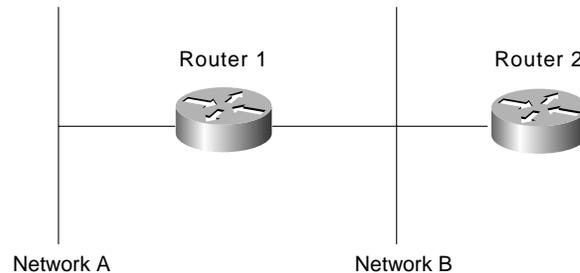
IGRP provides a number of features that are designed to enhance its stability. These include holddowns, split horizons, and poison-reverse updates.

*Holddowns* are used to prevent regular update messages from inappropriately reinstating a route that might have gone bad. When a router goes down, neighboring routers detect this via the lack of regularly scheduled update messages. These routers then calculate new routes and send routing update messages to inform their neighbors of the route change. This activity begins a wave of triggered updates that filter through the network. These triggered updates do not instantly arrive at every network device. Thus, it is possible for a device that has yet to be informed of a network failure to send a regular update message, which advertises a failed route as being valid to a device that has just been notified of the network failure. In this case, the latter device would contain (and potentially advertise) incorrect routing information. Holddowns tell routers to hold down any changes that might affect routes for some period of time. The holddown period usually is calculated to be just greater than the period of time necessary to update the entire network with a routing change.

*Split horizons* derive from the premise that it is never useful to send information about a route back in the direction from which it came. Figure 42-1 illustrates the split-horizon rule. Router 1 (R1) advertises that it has a route to Network A. There is no reason for Router 2 (R2) to include this route in its update back to R1 because R1 is closer to Network A. The split-horizon rule says that R2 should strike this route from any updates that it sends to R1. The split-horizon rule helps prevent routing loops. Consider, for example, the case in which R1’s interface to Network A goes down. Without split horizons, R2 continues to inform R1 that it can get to Network A (through R1). If R1 does not have sufficient intelligence, it

actually might pick up R2's route as an alternative to its failed direct connection, causing a routing loop. Although holddowns should prevent this, split horizons are implemented in IGRP because they provide extra algorithm stability.

*Figure 42-1 The Split-Horizon Rule Helps Protect Against Routing Loops*



Split horizons should prevent routing loops between adjacent routers, but *poison-reverse updates* are necessary to defeat larger routing loops. Increases in routing metrics generally indicate routing loops. Poison-reverse updates then are sent to remove the route and place it in holddown. In Cisco's implementation of IGRP, poison-reverse updates are sent if a route metric has increased by a factor of 1.1 or greater.

## Timers

IGRP maintains a number of timers and variables containing time intervals. These include an update timer, an invalid timer, a hold-time period, and a flush timer. The *update timer* specifies how frequently routing update messages should be sent. The IGRP default for this variable is 90 seconds. The *invalid timer* specifies how long a router should wait in the absence of routing-update messages about a specific route before declaring that route invalid. The IGRP default for this variable is three times the update period. The *hold-time variable* specifies the holddown period. The IGRP default for this variable is three times the update timer period plus 10 seconds. Finally, the *flush timer* indicates how much time should pass before a route should be flushed from the routing table. The IGRP default is seven times the routing update period.

## Summary

IGRP has proven to be one of the most successful routing protocols of all time. No small part of its success has been due to its functional similarity to RIP, a simple yet highly successful and widely deployed routing protocol. Cisco took great pains to carefully preserve many of the effective features of RIP, while greatly expanding its capabilities. Today, IGRP is showing its age; it lacks support for variable-length subnet masks (VLSM). Rather than develop an IGRP version 2 to incorporate that capability, Cisco has built upon IGRP's legacy of success with Enhanced IGRP. Enhanced IGRP is examined in Chapter 40 of this book.

## Review Questions

Q—Name the benefits of using IGRP instead of RIP.

**A**—Despite its enduring success as an Interior Gateway Routing Protocol, RIP suffers from fundamental limitations that are not easily avoided. For example, it (and the subsequent RIPv2) is limited to a maximum of 16 hops per route. This tends to limit the size and complexity of a network that can be effectively routed with RIP. Other RIP limitations include its support for only equal-cost load balancing and its single, simple routing metric (hop count). IGRP was specifically designed to offer an alternative to RIP that was as easy to implement and administer as RIP, yet that did not suffer from these fundamental limitations.

**Q**—*How can an administrator influence route selection?*

**A**—A network administrator can accept IGRP's default settings or can fine-tune network performance by manipulating any of IGRP's four routing metrics or their constant weights. These mathematical components of IGRP's composite routing metric offer remarkable latitude to network administrators by enabling them to emphasize or de-emphasize delay, link speed, historical reliability, or load levels in the selection of optimal routes.

**Q**—*What is variance, and how does it affect multipath routing?*

**A**—Variance is another value that can be established and modified by a network administrator to fine-tune an IGRP network. In essence, variance enables a range of routing costs to be used to select multiple redundant paths of unequal cost to any given destination. Thus, variance is the mechanism by which IGRP supports unequal-cost load balancing.

**Q**—*Identify and explain IGRP's stability features.*

**A**—IGRP uses holddowns, split horizons, and poison-reverse updates to improve operational stability. Holddowns prevent IGRP's interval updates from wrongly reinstating an invalid route. Split horizons help prevent routing loops by preventing a router from updating neighbors of any routing changes that it originally learned from those neighbors. Poison-reverse updates function similarly but are not limited to use between adjacent routers. Thus, they prevent large routing loops from occurring between nonadjacent routers.

**Q**—*What timers are used by IGRP, and what is their function?*

**A**—IGRP features several functionally distinct timers, including an update timer, an invalid timer, a hold-time period, and a flush timer. The *update* timer specifies how frequently routing update messages should be sent. The *invalid* timer specifies how long a router should wait in the absence of routing-update messages about a specific route before declaring that route invalid. The *hold-time variable* specifies the holddown period. Finally, the *flush* timer indicates how much time should pass before a route should be flushed from the routing table.

## For More Information

Sportack, Mark A. *IP Routing Fundamentals*. Indianapolis: Cisco Press, 1999.

[http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed\\_cr/np1\\_c/1cigrp.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/np1_c/1cigrp.htm)