

# Cisco's Interior Gateway Routing Protocol

## Overview

The Interior Gateway Routing Protocol (IGRP) is a routing protocol to provide routing within an autonomous system (AS). In the mid-1980s, the most popular interior 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. The popularity of Cisco routers and the robustness of IGRP encouraged many organizations with large internetworks to replace RIP with IGRP. Cisco developed Enhanced IGRP in the early 1990s to improve the operating efficiency of IGRP.

IGRP is a distance vector Interior Gateway Protocol (IGP). Distance vector routing protocols mathematically compare routes using some measurement of distance. Distance vector routing protocols are often contrasted with link-state routing protocols, which send local connection information to all nodes in the internetwork.

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.

## IGRP Protocol Characteristics

IGRP is a distance-vector interior gateway protocol (IGP). Distance-vector routing protocols call for each router to send all or a portion of its routing table in a routing-update message at regular intervals to each of its neighboring routers. As routing information proliferates through the network, routers can calculate distances to all nodes within the internetwork.

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 "Open Shortest Path First (OSPF)," and "Open System Interconnection (OSI) Protocols," respectively.

IGRP uses a combination (vector) of metrics. Internetwork delay, bandwidth, reliability, and load are all factored into the routing decision. Network administrators can set the weighting factors for each of these metrics. IGRP uses either the administrator-set or the default weightings to automatically calculate optimal routes.

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 1,200 bps to 10 gigabits per second, while delay can take on any value from 1 to 2 to the 24th power. Wide metric ranges allow satisfactory metric setting in internetworks with widely

varying performance characteristics. Most importantly, the metric components are combined in a user-definable algorithm. As a result, network administrators can influence route selection in an intuitive fashion.

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. Also, multiple paths can be used even if the metrics for the paths are different. If, for example, one path is three times better than another because 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 of the best route are used as multiple paths.

## Configuration

A very simple configuration of IGRP can be:

### Router A

```
RouterA# conf t
RouterA(config)# interface eth0
RouterA(config-if)# ip address 70.0.0.1 255.0.0.0
RouterA(config-if)# exit
RouterA(config)# interface serial0
RouterA(config-if)# ip address 20.30.40.2 255.255.255.252
RouterA(config-if)# exit
RouterA(config)# router igrp 1
RouterA(config-router)# redistribute connected
RouterA(config-router)# network 20.0.0.0
RouterA(config-router)# network 70.0.0.0
RouterA(config-router)# network 71.0.0.0
```

### Router B

```
RouterB# conf t
RouterB(config)# interface eth0
RouterB(config-if)# ip address 71.0.0.1 255.0.0.0
RouterB(config-if)# exit
RouterB(config)# interface serial0
RouterA(config-if)# ip address 20.30.40.1 255.255.255.252
RouterA(config-if)# exit
RouterA(config)# router igrp 1
RouterA(config-router)# redistribute connected
RouterA(config-router)# network 20.0.0.0
RouterA(config-router)# network 70.0.0.0
RouterA(config-router)# network 71.0.0.0
```

A few other commands might come in useful. **variance 2** can be used to configure IGRP to load balance between equal cost paths. The command **passive-interface eth0** disables IGRP from sending updates out of eth0.

## Testing

```
router# debug ip igrp events
```

Only shows the sending or receiving of IGRP packets and the number of routes in each update. It does show the routes that are advertised!

```
router# debug ip igrp transactions
```

Sames as debug ip igrp events but also shows the routes that are advertised.

```
router# show ip route
```

As with debugging any routing problem, look at the routing table. Is there a static route that takes precendece?

```
router# show ip interface brief
```

This command is always useful to quickly verify which links are and which aren't.