

Fast IGRP Details

Document ID: 13680

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Implementation

- The timers basic Command

- The no metric holddown Command

- Fast EIGRP Configuration Example

- Concerns

Related Information

Introduction

Typically, it can take several minutes for Interior Gateway Routing Protocol (IGRP) to converge. This happens since the total convergence time is the sum of the time it takes to detect the change in the state of the network, the delay due to the configured holddown timer, and the time it takes for the information to propagate from one end of the network to the other. However, in certain setups, a smaller convergence time might be desirable.

Fast IGRP is used to speed up the IGRP convergence process. Fast IGRP is able to detect a downed network faster. Also, fast IGRP does not go into holddown when a route disappears. Instead, it accepts new routing information as soon as it hears about it. However, in rare cases, this might mean ending up in a "count to infinity" condition, a situation which can be handled by reducing the maximum hop count.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

Implementation

There are two commands that you can use in order to speed up the IGRP process running on a network:

timers basic and **no metric holddown**. Both these commands work by reducing the delay before a route is accepted by the router.

The timers basic Command

The **timers basic** command lets you control how often IGRP sends updates. The default is once every 90 seconds. In order to allow for dropped packets, three consecutive packets must be lost before the route can be declared to be invalid. Thus, IGRP can not timeout expired routes until several minutes have elapsed.

The first thing to do is to speed up the time constants. Use 15 seconds instead of 90 for the basic time constant. This allows routes to expire after 45 seconds. All the other times decrease proportionally.

It should be noted that the expiration time turns out not to be as important as you might expect. Normally routes don't just expire. They are killed because keepalive fails on some interface, or you lose carrier. Keepalives are normally done every 10 seconds, so it takes 30 seconds to detect an interface down that way. Use a keepalive of 4 on T1 lines where you care about speed of routing adjustment. This lets you detect a failure in 12 seconds.

Refer to IGRP Commands for more information.

The no metric holddown Command

The other critical parameter change is **no metric holddown**. This disables holddowns, meaning that after a route has been removed, a new one will be accepted immediately. There were good theoretical reasons for using holddowns. There could be cases where, without holddowns, an old route will never get out of the system. However, Cisco IOS® has a couple of checks to prevent bad routes from surviving indefinitely, and we have not seen any old routes staying around. The router does keep a hop count, specifically to get rid of old routes that somehow manage to avoid the other tests.

If you do **show ip protocol**, you'll notice an "IGRP maximum hopcount" of 100. If all else fails, IGRP will go into "count to infinity," and stop at 100. Since IGRP uses triggered updates, counting to 100 may not take too long. However, you should set the maximum hopcount to something smaller, unless you have an enormous network. It should be a number at least as large as the "diameter" of your network, such as the maximum number of routers a route might ever have to go through. If you exchange IGRP routing with an external network, the diameter must include your network plus that external network. When you compute diameter, take into account what the configuration would look like if a few lines go down.

Fast EIGRP Configuration Example

Below is an example router configuration that uses all of the features listed above. The configuration is built based on the assumption that network 128.6.0.0 belongs to the organization for which the following is configured. Use your own network number in place of 128.6.0.0.

```
router igrp 46
timers basic 15 45 0 60

!--- Command to set the update, invalid, holddown, and flush timers

network 128.6.0.0
no metric holddown

!--- Disables the holddown state

metric maximum-hop 50

!--- Command reduces maximum hop count to 50 to better handle
```

possible routing loops due to the removal of the holddown timer

With this configuration, routing will generally adapt to change within 30 seconds (assuming that keepalive has been set to 4).

Concerns

Fast EIGRP should not be used in large networks. You wouldn't want to run IGRP quite this fast if you were circulating the whole Internet routing table. These adjustments should be made only after consulting with Cisco support personnel.

Related Information

- [IGRP Support Page](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 13680
