IGRP Metric

Document ID: 13678

Contents

Introduction

Interior Gateway Routing Protocol (IGRP) adds together weighted values of different characteristics of the link to the network in question in order to calculate a metric. The link characteristics from which IGRP calculates a composite metric are bandwidth, delay, load, reliability, and maximum transmission unit (MTU). By default, IGRP chooses a route based on bandwidth and delay.

Prerequisites

Requirements

Readers of this document should have knowledge of these topics:

- IGRP and related features

    Note: Refer to An Introduction to IGRP for more information.

Components Used

The information in this document is based on the software and hardware versions:

- Cisco IOS® Software Release 12.2(24a)
- Cisco 2500 Series Routers

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

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

Find the IGRP Metric

This section uses an example in order to illustrate how to find the metric when IGRP is the routing protocol.

Network Diagram

The diagram for the given scenario is provided here:

Here is the formula used to calculate the composite metric for IGRP:

Metric = [K1 * Bandwidth + (K2 * Bandwidth)/(256–load) + K3*Delay] * [K5/(reliability + K4)]

The default constant values are K1 = K3 = 1 and K2 = K4 = K5 = 0.

If K5 = 0, the [K5/(reliability + K4)] term is not used. So, given the default values for K1 through K5, the composite metric calculation used by IGRP reduces to Metric = Bandwidth + Delay.

The K values in these formulas are constants that you are able to define with the router configuration command, **metric weights** *tos k1 k2 k3 k4 k5* .

Note: Cisco strongly suggests that you do not change the default K parameters.

To find the bandwidth, find the smallest of all the bandwidths in Kbps from outgoing interfaces and divide 10,000,000 by that number. (The bandwidth is scaled by 10,000,000 in kilobits per second.)

In order to find the delay, add all of the delays (in microseconds) from the outgoing interfaces and divide this number by 10. (The delay is in tenths of microseconds.)

Remember, the path with the smallest metric is the best path.

The various outputs of the **show** commands for both the routers are as shown here:

```
Venus# show interfaces ethernet 0
Ethernet0 is up, line protocol is up
Hardware is Lance, address is 0060.5cf4.a9a8 (bia 0060.5cf4.a9a8)
Internet address is 12.1.1.1/24
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set

Venus# show interfaces serial 0
Serial0 is up, line protocol is up
Hardware is HD64570
Internet address is 172.16.10.2/24
MTU 1500 bytes, BW 784 Kbit, DLY 20000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation FRAME-RELAY, loopback not set
Keepalive set (10 sec)
LMI enq sent 981, LMI stat recvd 330, LMI upd recvd 0, DTE LMI up
LMI enq recvd 340, LMI stat sent 0, LMI upd sent 0
```

```
                LMI DLCI 1023 LMI type is CISCO frame relay DTE

        Saturn# show interfaces serial 1
        Serial0 is up, line protocol is up
        Hardware is HD64570
        Internet address is 172.16.10.1/24
        MTU 1500 bytes, BW 224 Kbit, DLY 20000 usec,
        reliability 255/255, txload 1/255, rxload 1/255
        Encapsulation FRAME-RELAY, loopback not set
        Keepalive set (10 sec)
        LMI enq sent 167, LMI stat recvd 168, LMI upd recvd 0, DTE LMI up
        LMI enq recvd 0, LMI stat sent 0, LMI upd sent 0
        LMI DLCI 1023 LMI type is CISCO frame relay DTE

        Saturn# show interfaces ethernet 0
        Ethernet0 is up, line protocol is up
        Hardware is Lance, address is 0060.5cf4.a955 (bia 0060.5cf4.a955)
        Internet address is 172.17.10.1/16
        MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
        reliability 255/255, txload 1/255, rxload 1/255
        Encapsulation ARPA, loopback not set
```

You are able to view the metric values calculated by IGRP with the **show ip route** command:

```
        Venus# show ip route 172.17.10.1
        Routing entry for 172.17.0.0/16
        Known via "igrp 100", distance 100, metric 14855
        Redistributing via igrp 100
        Advertised by igrp 100 (self originated)
        Last update from 172.16.10.1 on serial0, 00:00:13 ago
        Routing Descriptor Blocks:
          * 172.16.10.1, from 172.16.10.1, 00:00:13 ago, via Serial0
              Route metric is 14855, traffic share count is 1
              Total delay is 21000 microseconds, minimum bandwidth is 784 Kbit
              Reliability 255/255, minimum MTU 1500 bytes
              Loading 1/255, Hops 0
```

The corresponding calculations are:

Metric = Bandwidth + Delay = 10000000/784 + (20000 + 1000)/10 = **14855**

```
        Saturn# show ip route 12.1.1.1
        Routing entry for 12.0.0.0/8
        Known via "igrp 100", distance 100, metric 46742
        Redistributing via igrp 100
        Advertised by igrp 100 (self originated)
        Last update from 172.16.10.2 on serial1, 00:00:43 ago
        Routing Descriptor Blocks:
          * 172.16.10.2, from 172.16.10.2, 00:00:43 ago, via Serial1
              Route metric is 46742, traffic share count is 1
              Total delay is 21000 microseconds, minimum bandwidth is 224 Kbit
              Reliability 255/255, minimum MTU 1500 bytes
              Loading 1/255, Hops 0
```

The corresponding calculations are:

Metric = Bandwidth + Delay = 10000000/224 + (20000 + 1000)/10 = **46742**

How Often Is Load Calculated?

The constant K2 defaults to zero. If K2 is set to 1, load becomes a variable used in routing. The problem seems to be if the load jumps. If the metric cost jumps at the start of an FTP session, it is possible for the route

go into holddown due to the increase. How often is load calculated?

The load is a five–minute exponentially weighted average that is updated every five seconds.

How Fast Can the Load Value Rise?

Is it possible for the load value to rise fast enough to make the route unstable?

Yes, it is. And worse, when the load falls, the metric decreases. This failure causes a Flash update.

Can IGRP Be Configured to Use the Fastest Path through the Network Cloud?

Since the composite metric cost to a given site is determined by the slowest link in the path and the slowest link is normally the access line into the cloud, how can IGRP be configured to use the fastest path through the network cloud?

Once the slowest link has been determined, the rest of the routing is done on hops (delay) without regard for hop–link speeds. With the large gaps in the bandwidth values, it does not seem practical to try and use delay to bias network cloud routing. One obvious solution is to configure the **bandwidth** command on the access lines to be faster than any network cloud backbone line.

Another solution is to configure the delay on the WAN links to be an accurate measurement of the delay for that particular link. You should not have to tweak the delays at all, and you should have good routing.

It is certainly worthwhile to change the bandwidths on the access line if you have radically different bandwidths within your WAN.

What Metric Should Be Used When Redistributing Routes into IGRP?

Issue the **default–metric** command to set the metric for the redistributed routes. This statement is appropriate for most cases:

```
Venus(config)# router igrp 100
  Venus(config-router)# default-metric 10000 100 255 1 1500
```

where 10000 = Bandwidth, 100 = Delay, 255 = Reliability, 1 = Loading, and 1500 = MTU.

Related Information

- **How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP?**
- **An Introduction to IGRP**
- **IGRP Support Page**
- **IP Routing Technology Support Page**
- **Technical Support – Cisco Systems**