How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP?

Document ID: 13677

Contents

Introduction

In general, load balancing is the capability of a router to distribute traffic over all the router network ports that are the same distance from the destination address. Load balancing increases the utilization of network segments, and so increases effective network bandwidth. There are two types of load balancing:

- Equal cost path   Applicable when different paths to a destination network report the same routing metric value. The maximum–paths command determines the maximum number of routes that the routing protocol can use.
- Unequal cost path   Applicable when different paths to a destination network report are of different routing metric values. The variance command determines which of these routes is used by the router.

This document explains how unequal cost path load balancing works in Enhanced Interior Gateway Routing Protocol (EIGRP).

Prerequisites

Requirements

This document requires a basic understanding of IP routing protocols and EIGRP routing protocol. In order to learn more about IP routing protocols and EIGRP, refer to these documents:

- Routing Basics
- EIGRP Support Page

Components Used

- EIGRP is supported in Cisco IOS® Software Release 9.21 and later.
- You can configure EIGRP in all routers (such as the Cisco 2500 series and the Cisco 2600 series) and in all Layer 3 switches.

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

Refer to Cisco Technical Tips Conventions for more information on document conventions.

EIGRP Load Balancing

Every routing protocol supports equal cost path load balancing. In addition, Interior Gateway Routing Protocol (IGRP) and EIGRP also support unequal cost path load balancing. Use the **variance** n command in order to instruct the router to include routes with a metric of less than n times the minimum metric route for that destination. The variable n can take a value between 1 and 128. The default is 1, which means equal cost load balancing. Traffic is also distributed among the links with unequal costs, proportionately, with respect to the metric.

Note: If a path is not a feasible successor, the path is not used in load balancing. Refer to the *Feasible Distance, Reported Distance, and Feasible Successor* section of Enhanced Interior Gateway Routing Protocol for more information.

Network Diagram
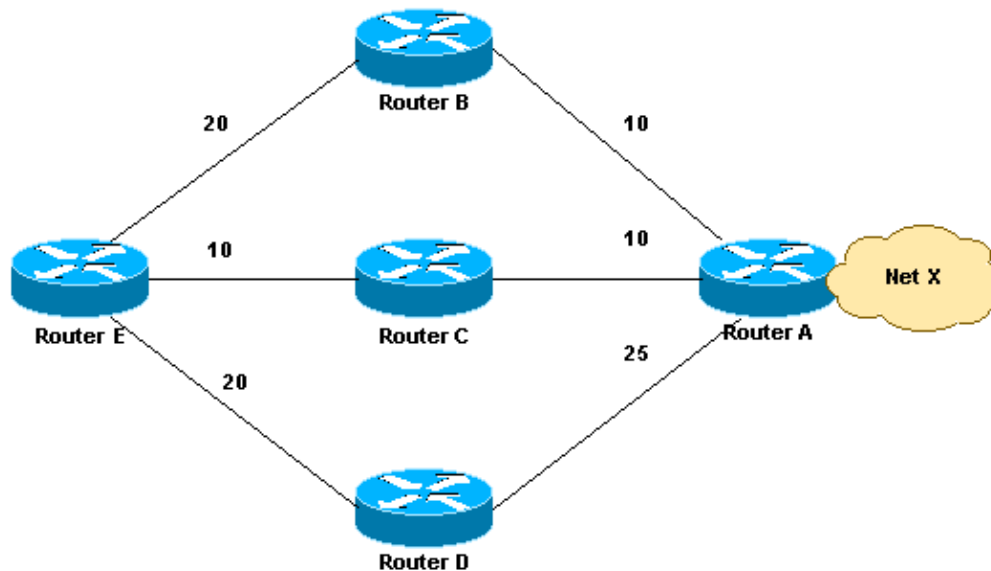
Variance

This section provides an example. In the network diagram, there are three ways to get to Network X:

- E−B−A with a metric of 30
- E−C−A with a metric of 20
- E−D−A with a metric of 45

Router E chooses the path E−C−A with a metric of 20 because 20 is better than 30 and 45. In order to instruct EIGRP to select the path E−B−A as well, configure variance with a multiplier of 2:

```
router eigrp 1
network x.x.x.x
variance 2
```

This configuration increases the minimum metric to 40 (2 * 20 = 40). EIGRP includes all routes that have a metric of less than or equal to 40 and satisfy the feasibility condition. In the configuration in this section, EIGRP now uses two paths to get to Network X, E–C–A and E–B–A, because both paths have a metric of under 40. EIGRP does not use path E–D–A because that path has a metric of 45, which is not less than the value of the minimum metric of 40, because of the variance configuration. Also, the reported distance of neighbor D is 25, which is greater than the feasible distance (FD) of 20 through C. This means that, even if variance is set to 3, the E–D–A path is not selected for load balancing because Router D is not a feasible successor.

Note: For more information on variance, refer to Troubleshooting EIGRP Variance Command.

Traffic Sharing

EIGRP not only provides unequal cost path load balancing, but also intelligent load balancing, such as traffic sharing. In order to control how traffic is distributed among routes when there are multiple routes for the same destination network that have different costs, use the **traffic−share balanced** command. With the keyword **balanced**, the router distributes traffic proportionately to the ratios of the metrics that are associated with different routes. This is the default setting:

```
router eigrp 1
network x.x.x.x
variance 2
traffic-share balanced
```

The traffic share count for this example is:

- For path E–C–A: 30/20 = 3/2 = 1
- For path E–B–A: 30/30 = 1

Because the ratio is not an integer, round down to the nearest integer. In this example, EIGRP sends one packet to E–C–A and one packet to E–B–A.

Now, assume that the metric between E–B is 25 and the metric between B–A is 15. In this case, the E–B–A metric is 40. However, this path will not be selected for load balancing because the cost of this path, 40, is not less than (20 * 2), where 20 is the FD and 2 is the variance. In order to include this path also in load sharing, the variance should be changed to 3. In this case, the traffic share count ratio is:

- For path E–C–A: 40/20 = 2
- For path E–B–A: 40/40 = 1

In this situation, EIGRP sends two packets to E–C–A and one packet to E–B–A. In this way, EIGRP provides both unequal cost path load balancing and intelligent load balancing. Refer to the *Load Balancing* section of Enhanced Interior Gateway Routing Protocol for more information on how EIGRP load balances traffic over unequal cost links.

Similarly, when you use the **traffic−share** command with the keyword **min**, the traffic is sent only across the minimum−cost path, even when there are multiple paths in the routing table.

```
router eigrp 1
network x.x.x.x
variance 3
traffic-share min across-interfaces
```

In this situation, EIGRP sends packets only through E–C–A, which is the best path to the destination network. This is identical to the forwarding behavior without use of the **variance** command. However, if you use the **traffic–share min** command and the **variance** command, even though traffic is sent over the minimum–cost path only, all feasible routes get installed into the routing table, which decreases convergence times.

You have seen how to configure unequal cost path load balancing in EIGRP. Similarly, you can do the same in IGRP, except for the feasibility condition. This condition is not applicable to IGRP.

Load Balancing in CEF

Cisco Express Forwarding (CEF) is an advanced Layer 3 switching technology which can be used for load balancing in routers. By default, CEF uses per–destination load balancing. If it is enabled on an interface, per–destination load balancing forwards packets based on the path to reach the destination. If two or more parallel paths exist for a destination, CEF takes the same path (single path) and avoids the parallel paths. This is a result of the default behavior of CEF. CEF takes the single path in cases when load sharing is done simultaneously on interfaces of different physical types, such as serial and tunnel. The hash algorithm determines the path to be chosen. In order to utilize all the parallel paths in CEF and load balance the traffic, you must enable per–packet load balancing when you have different physical interfaces like serial and tunnel. So, on the basis of the configuration and topology (serial or tunnel), load sharing can fail to work correctly with the default CEF load balancing mode.

Enable these commands for load sharing on a per–packet basis:

```
configure terminal
interface serial 0
ip load-sharing per-packet
```

Related Information

- **Introduction to EIGRP**
- **How Does Load Balancing Work?**
- **EIGRP Support Page**
- **IGRP Support Page**
- **IP Routed Protocols Support Page**
- **IP Routing Support Page**
- **Technical Support & Documentation – Cisco Systems**

Updated: Jun 03, 2009   Document ID: 13677