Ivan Pepelnjak
CCIE #1354

The definitive resource for EIGRP design, deployment, and operation

EIGRP Network Design Solutions

Cisco Systems
Cisco Press
www.ciscopress.com

EIGRP Network Design Solutions

Ivan Pepelnjak

Cisco Systems
Cisco Press

Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

EIGRP normally only installs routes to the successors in the main IP routing table. As many routes are installed as are permitted (up to the maximum of six) yielding equal-cost load balancing. EIGRP and IGRP are also the only routing protocols that support unequal-cost load balancing via a mechanism called *variance.*

NOTE The administrative distance is usually the only means of comparing routes coming from different routing sources due to incompatible metrics used in different routing protocols. However, the administrative distance is used as the sole criteria even when comparing routes coming from two EIGRP processes even though the metrics could be compared in this scenario. If the two EIGRP processes have the same administrative distance, the route that was changed most recently (the latest flap) takes precedence and replaces the older (more stable) route. Therefore, it's mandatory to use different administrative distances for different EIGRP processes if they cover overlapping parts of your network. Some hints for proper administrative distance selection in these scenarios are given in Chapter 9, "Integrating EIGRP with Other Enterprise Routing Protocols."

Administrative Distance of EIGRP Routes

The EIGRP route selection process sets different administrative distances for internal and external routes, the defaults being 90 for internal routes and 170 for external routes. (For further explanation about using different distances for internal and external routes, see Chapter 9.) As shown in Table 2-18, you can change these distances in a variety of ways.

Table 2-18 *Different Ways of Setting Nondefault Administrative Distances of EIGRP Routes*

To Change . . .	. . . Use This Command
Default distances for internal and external routes in an EIGRP process	**router eigrp <as>** **distance eigrp <default-internal-distance> <default-external-distance>**
Distance of all internal routes received from a neighbor or a set of neighbors	**router eigrp <as>** **distance <distance> <neighbor-ip-address> <wildcard-bits>**
Distance of a select set of internal routes received from a neighbor or a set of neighbors	**router eigrp <as>** **distance <distance> <neighbor-ip-address> <wildcard-bits> <route-selection-ACL>**

The **distance eigrp** command sets the new defaults for both internal and external routes. Use this command to prefer one EIGRP process over another in the event that the address space of the two processes overlaps.

The **distance** command influences only the distance of internal EIGRP routes (nondefault value for external routes cannot be specified) and enables you to set different administrative distances for routes received from specific neighbors (matched by **neighbor-ip-address** and **wildcard-bits**) or matched by a route filter (standard or extended IP access list specified in **route-selection-ACL**). If you want to ignore internal EIGRP routes received from a neighbor or internal EIGRP routes matching a route filter, specify distance 255 (which means ignore this entry).

Use the **distance** command only in very special circumstances and with extreme care because side effects of this command are not easily evaluated.

EIGRP Variance and Its Influence on Traffic Load Sharing

EIGRP is the only protocol that can load-balance between unequal cost routes. To enable and fine-tune EIGRP load balancing, use the commands in Table 2-19.

Table 2-19 *Configure Unequal-Cost Load-Sharing with EIGRP*

Task	Configure With
Configure unequal-cost load balancing	**router eigrp <as>** **variance <factor>**
Configure proportional load balancing between unequal-cost routes	**router eigrp <as>** **traffic-share balanced**
Use only minimum-cost routes for load balancing	**router eigrp <as>** **traffic-share min across-interfaces**
Configure the maximum number of equal-cost or unequal-cost routes for a given destination	**route eigrp <as>** **maximum-paths <1 to 6>**
Configure per-packet load balancing over an interface on all platforms	**interface <int>** **no ip route-cache**
Configure Cisco Express Forwarding (CEF) per destination load balancing	**interface <int>** **ip route-cache cef** **ip load-sharing per-destination**
Configure CEF per-packet load balancing	**interface <int>** **ip route-cache cef** **ip load-sharing per-packet**

The **variance** command enables unequal-cost balancing under the following conditions:

- The router's own distance from the topology table entry is less than *feasibility distance variance*.
- The alternate path toward the destination goes through a feasible successor.

The number of alternate paths that can be entered in the IP routing table is controlled by the maximum-paths command. (Only the best N entries that match feasibility condition are entered in the routing table.)

You configure the load-balancing mode on unequal-cost routes by using the **traffic-share** command. If you specify **traffic-share balanced**, the traffic is load-balanced inversely proportionally to the EIGRP composite metric. If, on the other hand, you specify **traffic-share min**, the routed traffic is balanced only across the minimum-cost paths, but all the other paths are already entered in the IP routing table to speed up the convergence process in case of a link or neighbor failure.

The load-balancing mechanism used by the router is dependent on the switching path taken for the interface in question as detailed in Table 2-20.

Table 2-20 *Load Sharing Mechanism Used Depending on the Switching Path*

Switching Path	Load Sharing Mechanism
Process switching	Per-packet load sharing
Fast switching, Optimum switching, Autonomous switching, Silicon switching, Netflow switching	Per-prefix load sharing (for example, all traffic for a certain prefix in the routing table flows over one interface)
Cisco Express Forwarding	Per source-destination-pair load sharing (for example, all traffic for a certain source-destination IP address pair flows over one interface)
Cisco Express Forwarding with per-packet load sharing configured	Per-packet load sharing

Designing EIGRP networks for unequal-cost load balancing is a nontrivial task as you'll see in the next section, but you'll be able to design successful networks using unequal-cost load-balancing by following these rules:

Variance Rule 1

Verify that the paths over which you want to load-balance the traffic lead to successors and feasible successors. In many intuitively correct designs, the neighbor you want to balance the load with is not a feasible successor in EIGRP terms.

Variance Rule 2

Always verify that the load-balancing mechanism works in both directions; if the load balancing works for traffic flowing in one direction it may not work for the return traffic.

Variance Rule 3

If there is more than one router connected to a LAN and you want to load-balance outgoing traffic from that LAN, you'll need additional mechanisms like Hot Standby Routing Protocol (HSRP) to select the proper exit point from the LAN.

Variance Rule 4

You can solve some load-balancing problems where load balancing intuitively works but does not work in reality due to feasible successor limitations by introducing another layer of routers to distribute the traffic.

Valid and Invalid Examples of Using Variance

In several commonly used designs, you might expect to get the desired load-sharing functionality, but the design does not work because it does not comply with EIGRP requirements for unequal-cost load sharing. Several examples are presented in this section, some with comments and solutions, the other as exercises.

One of the most common designs is load sharing between two unequal speed links between two adjacent routers. A simple network using this design appears in Figure 2-22.

The design requirement for this network is to provide proportional load balancing on parallel links between San Jose and Chicago in a 2:1 ratio for all traffic running over those two links. To evaluate whether this design yields the required load balancing functionality, you need to verify that all the variance rules are satisfied:

Variance Rule 1—Verification:

The San Jose router has two paths to the Chicago LAN. The Chicago router over the 2-Mbps link is the successor, and the same router over a 1-Mbps link is a feasible successor.

The Chicago router has two paths to the San Jose LAN. The San Jose router over the 2-Mbps link is the successor, and the same router over a 1-Mbps link is the feasible successor. The Chicago router also has two paths to the San Francisco LAN, but they have the same distance; the minimum bandwidth is 64 kbps over both paths and the delays across the links are the same.

Variance Rule 1 is thus satisfied. Both routers that have to do load sharing have paths going to successors and/or feasible successors over links where the traffic should be load-shared.

Figure 2-22 *Simple Network with Unequal Speed Links*

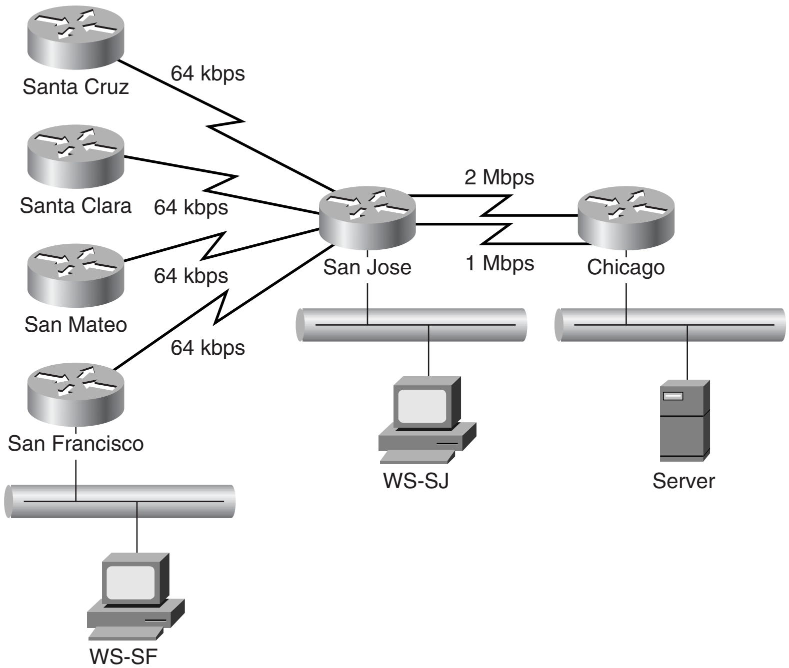

Variance Rule 2—Verification:

The traffic from San Francisco and San Jose toward the Chicago server load-share in the proper ratio; the same is also true for the return traffic from the Chicago server toward the San Jose workstation. The traffic from the Chicago server toward the San Francisco workstation load-share equally across the links because both paths from Chicago toward San Francisco have the same cost. The traffic from Chicago toward San Francisco therefore places too high a load on the lower-speed link. If the number of San Francisco-type offices connected to the San Jose router is high enough and most of the traffic goes from Chicago toward those offices, the 1-Mbps link becomes saturated, while the 2-Mbps link is only 50 percent used.

Variance Rules 3 and 4 do not apply in this particular design.

Based on the previous facts, the design in Figure 2-22 provides load balancing between unequal speed links, but not completely within the required specifications.

Exercise 2-2

How could you modify the EIGRP design in Figure 2-19 to ensure proportional load balancing from Chicago toward all destinations in the San Jose area?

Exercise 2-3

The customer validated load balancing in an environment where the majority of the traffic flowed from the Chicago server toward the remote locations. A new remote backup application was deployed a few months later where the majority of the traffic flowed from remote offices toward the Chicago server. To the customer's surprise, all the traffic from San Jose toward Chicago uses only one of the parallel links. Why?

After implementing the load balancing between Chicago and San Jose, the customer discovered that the slow-speed links in San Jose are saturated over the peak usage periods. The customer decided to upgrade the leased lines to 128 kbps and installed ISDN into all the offices in that area. ISDN is supposed to be used in the dial-backup mode with remote offices calling the San Jose office (dial-in). Unequal-cost load balancing is to be used between the 64-kbps leased line and one ISDN B-channel dial-up connection and the desired traffic ratio between the leased line and ISDN connection is 2:1. The target network diagram is shown in Figure 2-23.

This network conforms to all applicable variance rules for load balancing between remote offices in the Bay area and the San Jose site as follows:

Variance Rule 1—Verification:

The San Jose router has two paths to a remote LAN. The best path to the remote LAN is over the 128-kbps leased line, and the remote router over the ISDN dial-up connection is a feasible successor. The same conclusion can also be reached for paths to the Chicago LAN as seen from the remote office router.