EIGRP Successor and Feasible Successor Calculations
Created by Keith Bogart

The Enhanced Interior Gateway Routing Protocol (EIGRP) is an interior gateway protocol (IGP) used to determine forwarding paths and next hops for IP packets. In virtually any network where path redundancy exists, EIGRP commonly receives routing updates from multiple EIGRP neighbors about the same destination IP prefix (IP subnet). By default, EIGRP compares these routing updates to determine which neighbor has the best path to that destination prefix. This document examines that process and how EIGRP makes the best-path determination.

This document is for readers who are already familiar with the following topics:

- The differences between a routing protocol and a routable protocol
- IP version 4 (IPv4) packet creation and delivery between end systems
- The way that EIGRP dynamically discovers its neighbors using hello packets
- Basic configuration of EIGRP

EIGRP Successor and Feasible Successor

When any IGP routing protocol receives two or more updates to the same destination prefix from neighboring routers, the IGP makes use of some kind of metric for comparing those updates against each other to determine the best path. The best path is then installed into the IP routing table.

When a route for a prefix is placed in the routing table there are three basic components to that route:
1. The prefix itself and mask (such as 10.10.0.0 /16)
2. The metric for that prefix, such as hop count for Routing Information Protocol (RIP) or cost for Open Shortest Path First (OSPF)
3. The IP address of the next hop for that route

If an EIGRP router receives multiple updates for the same prefix from various EIGRP neighbors, it must determine which of those neighbors is “closest” to that destination prefix. EIGRP installs all the updates into its EIGRP topology table and then determines how close a neighbor is to a given prefix by looking at the metric that neighbor provided in its routing update. The lower the metric value, the closer that neighbor is to the destination prefix. The neighbor that is the closest is the one whose IP address will be stored as the next hop in the IP routing table.

In EIGRP terminology, this neighbor is called the successor. It is possible that multiple neighbors will advertise the same prefix with exactly the same metric. In that case, EIGRP can do equal-cost load-balancing by installing both neighbors into the routing table. Thus, the prefix will have multiple successors.

By default, EIGRP installs up to four equal-cost successors into the IP routing table. This number can be increased or decreased by means of configuration commands.

If an EIGRP router receives multiple routes to the same destination prefix and those routes do not have the same metric, you can install the router that “lost” (because its metric was higher than that of the successor) into the EIGRP topology table as a backup next hop. In other words, this neighbor can be used as a next hop if the current successor goes away or withdraws the route. In this situation, the backup neighbor is called a feasible successor.

Just as an EIGRP router may have multiple equal-cost successors to a given prefix, it may also contain multiple feasible successors. However, just because a router advertises a route does not mean that the router necessarily qualifies as a feasible successor; it must first meet the feasibility condition (FC). The FC is described later in this paper.

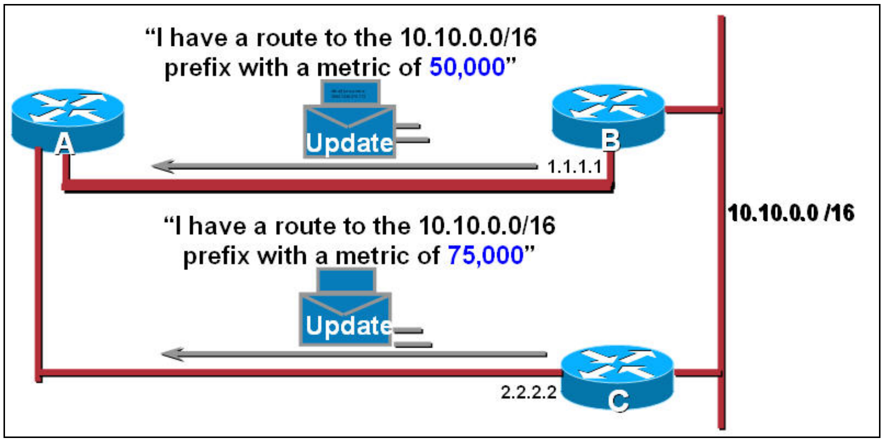

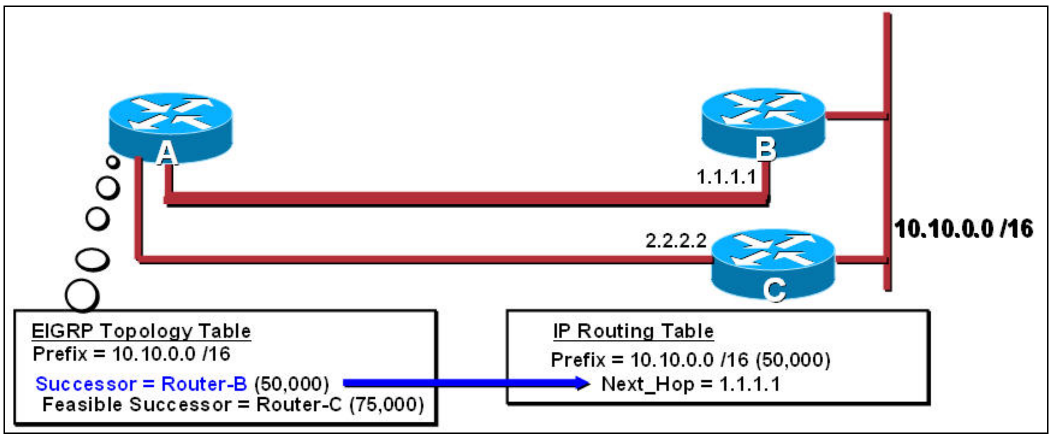

The value used for the IGP metric varies depending on which IGP is in use. RIP uses hop count. OSPF uses cost. EIGRP uses distance, which is a value derived from several vector metrics. Those vector metrics are bandwidth, delay, load, and reliability; however, bandwidth and delay are the only two factors used in the default calculation.

When an interface on a router becomes active, metrics such as bandwidth and delay are automatically assigned to that interface, whether or not a routing protocol is in use. You can see these values by using the Cisco IOS command **show interface**, as shown in the following example.

```
Router#show interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is AmdFE, address is 0011.5c10.8260 (bia 0011.5c10.8260)
  Description: Connection to R7
  Internet address is 200.1.1.129/27
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

It is these same values that are used to derive the EIGRP distance for that directly connected network.

Advertised Distance

When an EIGRP router is directly connected to a network and must advertise that prefix to its neighbors, it must determine the “advertised distance” for that prefix. The advertised distance is the distance the local router perceives that prefix to be from itself. This is also sometimes called the reported distance.

EIGRP uses (by default) the values of bandwidth (BW) and delay from the interface, and plugs those numbers into the following formula to derive the advertised distance of that directly connected network:

$$256 \times \left\{ \left[(K1 \times BW) + \frac{(K2 \times BW)}{(256 - Load)} + (K3 \times Delay) \right] \times \left[\frac{K5}{(Reliability + K4)} \right] \right\}$$

The K values shown in the formula are simple multipliers and by default take on the following values:

$$K1 = 1,\ K2 = 0,\ K3 = 1,\ K4 = K5 = 0$$

With the default K values shown above, the EIGRP formula ignores the last part of the equation (that uses reliability), and the portion of the equation using load drops out as well. This leaves the following as the default formula for determining the EIGRP distance of a prefix:

$$256 \times \left\{ \left[(K1 \times BW) + (K3 \times Delay) \right] \right\}$$

