

EIGRP White Papers

Cisco System

1.0 EIGRP Theory of Operations

- 1.1 Major Revisions of the Protocol
- 1.2 Basic Theory
- 1.3 Neighbor Discovery and Maintenance
- 1.4 Building the Topology Table
- 1.5 EIGRP Metrics
- 1.6 Feasible Distance, Reported Distance, and the Feasible Successor
- 1.7 Deciding if A Path is Loop-Free

2.0 Split Horizon and Poison Reverse

- 2.1 Startup Mode
- 2.2 Topology Table Change
- 2.3 Queries

3.0 Stuck in Active (SIA) Routes

- 3.1 Troubleshooting SIA Routes

4.0 Redistribution

- 4.1 Redistribution Between Two EIGRP Autonomous Systems
- 4.2 Redistribution Between EIGRP and IGRP in Two Different Autonomous Systems
- 4.3 Redistribution Between EIGRP and IGRP in the Same Autonomous System
- 4.4 Redistribution To and From Other Protocols
- 4.5 Redistribution of Static Routes to Interfaces

5.0 Summarization

- 5.1 Auto-summarization
- 5.2 Manual Summarization
- 5.3 Auto-summarization of External Routes

6.0 Query Processing and Range

- 6.1 How Summarization Points Affect the Query Range
- 6.2 How Autonomous System Boundaries Affect the Query Range
- 6.3 How Distribution Lists Affect the Query Range

7.0 Pacing Packets

8.0 Default Routing

9.0 Load Balancing

10.0 Using the Metrics

11.0 Using Administrative Tags in Redistribution

12.0 Show IP EIGRP Toplogy

1.0 EIGRP Theory of Operation

Enhanced Interior Gateway Routing Protocol (EIGRP) is an interior gateway protocol suited for many different topologies and media. In a well designed network, EIGRP scales well and provides extremely quick convergence times with minimal network traffic. Some of the many advantages of EIGRP are:

- very low usage of network resources during normal operation; only hello packets are transmitted on a stable network
- when a change occurs, only routing table changes are propagated, not the entire routing table; this reduces the load the routing protocol itself places on the network
- rapid convergence times for changes in the network topology (in some situations convergence can be almost instantaneous)

EIGRP is an enhanced distance vector protocol, relying on the Diffused Update Algorithm (DUAL) to calculate the shortest path to a destination within a network.

1.1 Major Revisions of the Protocol

There are two major revisions of EIGRP, versions 0 and 1. Cisco IOS versions earlier than 10.3(11), 11.0(8), and 11.1(3) run the earlier version of EIGRP; some explanations in this paper may not apply to that earlier version. We highly recommend using the later version of EIGRP, as it includes many performance and stability enhancements.

1.2 Basic Theory

A typical distance vector protocol saves the following information when computing the best path to a destination: the distance (total metric or distance, such as hop count) and the vector (the next hop). For instance, all the routers in the network in Figure 1 are running Routing Information Protocol (RIP). Router Two chooses the path to Network A by examining the hop count through each available path.




Figure 1

Since the path through Router Three is three hops, and the path through Router One is two hops, Router Two chooses the path through One and discards the information it learned through Three. If the path between Router One and Network A goes down, Router Two loses all connectivity with this destination until it times out the route of its routing table (three update periods, or 90 seconds), and Router Three re-advertises the route (which occurs every 30 seconds in RIP). Not including any hold-down time, it will take between 90 and 120 seconds for Router Two to switch the path from Router One to Router Three.

EIGRP, instead of counting on full periodic updates to re-converge, builds a topology table from each of its neighbor's advertisements (rather than discarding the data), and converges by either looking for a likely loop-free route in the topology table, or, if it knows of no other route, by querying its neighbors. Router Two saves the information it received from both Routers One and Three. It chooses the path through One as its best path (the successor) and the path through Three as a loop-free path (a feasible successor). When the path through Router One becomes unavailable, Router Two examines its topology table and, finding a feasible successor, begins using the path through Three immediately.

From this brief explanation, it's apparent that EIGRP must provide:

- a system where it sends only the updates needed at a given time; this is accomplished through neighbor discovery and maintenance
- a way of determining which paths a router has learned are loop-free
- a process to clear bad routes from the topology tables of all routers on the network
- a process for querying neighbors to find paths to lost destinations

We'll cover each of these requirements in turn.

1.3 Neighbor Discovery and Maintenance

To distribute routing information throughout a network, EIGRP uses non-periodic incremental routing updates. That is, EIGRP only sends routing updates about paths that have changed when those paths change.

The basic problem with sending only routing updates is that you may not know when a path through a neighboring router is no longer available. You can't time out routes, expecting to receive a new routing table from your neighbors. EIGRP relies on neighbor relationships to reliably propagate routing table changes throughout the network; two routers become neighbors when they see each other's hello packets on a common network.

EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth multipoint links.

- 5-second hello:
 - ✧ broadcast media, such as Ethernet, Token Ring, and FDDI
 - ✧ point-to-point serial links, such as PPP or HDLC leased circuits, Frame Relay point-to-point subinterfaces, and ATM
 - ✧ point-to-point subinterfaces
 - ✧ high bandwidth (greater than T1) multipoint circuits, such as ISDN PRI and Frame Relay
- 60-second hello:
 - ✧ multipoint circuits T1 bandwidth or slower, such as Frame Relay multipoint interfaces, ATM multipoint interfaces, ATM
 - ✧ switched virtual circuits, and ISDN BRIs

The rate at which EIGRP sends hello packets is called the hello interval, and you can adjust it per interface with the **ip eigrp hello-interval** command. The hold time is the amount of time that a router will consider a neighbor alive without receiving a hello packet. The hold time is typically three times the hello interval, by default, 15 seconds and 180 seconds. You can adjust the hold time with the **ip eigrp hold-time interface** command.

Note that if you change the hello interval, the hold time is not automatically adjusted to account for this change - you must manually adjust the hold time to reflect the configured hello interval.

It is possible for two routers to become EIGRP neighbors even though the hello and hold timers don't match. The hold time is included in the hello packets so each neighbor should stay alive even though the hello interval and hold timers don't match.

While there is no direct way of determining what the hello interval is on a router, you can infer it from the output of **show ip eigrp neighbor** on the neighboring router.

```
router#show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H  Address          Interface  Hold Uptime   SRTT  RTO  Q  Seq
                               (sec)           (ms)    Cnt  Num
1  10.1.1.2         Et1        13 12:00:53   12   300  0  620
0  10.1.2.2         S0         174 12:00:56   17   200  0  645
```

```
router#show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H  Address          Interface  Hold Uptime   SRTT  RTO  Q  Seq
                               (sec)           (ms)    Cnt  Num
1  10.1.1.2         Et1        12 12:00:55   12   300  0  620
0  10.1.2.2         S0         173 12:00:57   17   200  0  645
```

The value in the Hold column of the command output should never exceed the hold time, and should never be less than the hold time minus the hello interval (unless, of course, you are

losing hello packets). If the Hold column usually ranges between 10 and 15 seconds, the hello interval is 5 seconds and the hold time is 15 seconds. If the Hold column usually has a wider range - between 120 and 180 seconds - the hello interval is 60 seconds and the hold time is 180 seconds. If the numbers do not seem to fit one of the default timer settings, check the interface in question on the neighboring router - the hello and hold timers may have been configured manually.

Note EIGRP doesn't build peer relationships over secondary addresses. All EIGRP traffic is sourced from the primary address of the interface.

1.4 Building the Topology Table

Now that these routers are talking to each other, what are they talking about? Their topology tables, of course! EIGRP, unlike RIP and IGRP, does not rely on the routing (or forwarding) table in the router to hold all of the information it needs to operate. Instead, it builds a second table, the topology table, from which it installs routes in the routing table.

To see the basic format of the topology table on a router running EIGRP, issue the **show ip eigrp topology** command. The topology table contains the information needed to build a set of distances and vectors to each reachable network, including:

- lowest bandwidth on the path to this destination as reported by the upstream neighbor
- total delay
- path reliability
- path loading
- minimum path maximum transmission unit (MTU)
- feasible distance
- reported distance
- route source (external routes are marked)

Feasible and reported distance are discussed later in this section.

1.5 EIGRP Metrics

EIGRP uses the minimum bandwidth on the path to a destination network and the total delay to compute routing metrics. Although you can configure other metrics, we don't recommend it, as it can cause routing loops in your network. The bandwidth and delay metrics are determined from values configured on the interfaces of routers in the path to the destination network.

For instance, in Figure 2 below, Router One is computing the best path to Network A.




Figure 2

It starts with the two advertisements for this network: one through Router Four, with a minimum bandwidth of 56 and a total delay of 2200; and the other through Router Three, with a minimum bandwidth of 128 and a delay of 1200. Router One chooses the path with the lowest metric.

Let's compute the metrics. EIGRP calculates the total metric by scaling the bandwidth and delay metrics. EIGRP uses the following formula to scale the bandwidth:

$$\text{bandwidth} = (10000000/\text{bandwidth}) * 256$$

EIGRP uses the following formula to scale the delay:

$$\text{delay} = \text{delay} * 256$$

EIGRP uses these scaled values to determine the total metric to the network:

$$\text{metric} = [K1 * \text{bandwidth} + (K2 * \text{bandwidth}) / (256 - \text{load}) + K3 * \text{delay}] * [K5 / (\text{reliability} + K4)]$$

Note These K values should be used after careful planning. Mismatched K values prevent a neighbor relationship from being built, which can cause your network to fail to converge.

The default values for K are:

- K1 = 1
- K2 = 0
- K3 = 1
- K4 = 0
- K5 = 0

For default behavior, you can simplify the formula as follows:

$$\text{metric} = \text{bandwidth} + \text{delay}$$

Combine this formula with the scaling factors and you have:

$$\left[\left(\frac{10^7}{\text{min bandwidth}} \right) + \text{sum of delays} \right] \times 256$$

These formulas assume you are using the delay as configured on the interface, which is in tens of microseconds. The delay as shown in **show ip eigrp topology** or **show interface** is in microseconds, so you must divide by 10 before you use it in this formula. Throughout this paper, we use delay as it is configured and shown on the interface.

Cisco routers do not perform floating point math, so at each stage in the calculation, you need to round down to the nearest integer to properly calculate the metrics.

In this example, the total cost through Router Four is:

$$\text{minimum bandwidth} = 56\text{k}$$

$$\text{total delay} = 100 + 100 + 2000 = 2200$$

$$[(10000000/56) + 2200] \times 256 = (178571 + 2200) \times 256 = 180771 \times 256 = 46021376$$

And the total cost through Router Three is:

$$\text{minimum bandwidth} = 128\text{k}$$

$$\text{total delay} = 100 + 100 + 1000 = 1200$$

$$[(10000000/128) + 1200] \times 256 = (78125 + 1200) \times 256 = 79325 \times 256 = 20307200$$

So to reach Network A, Router One chooses the route through Router Three.

Note the bandwidth and delay values we used are those configured on the interface through which the router reaches its next hop to the destination network. For example, Router Two advertised Network A with the delay configured on its Ethernet interface; Router Four added the delay configured on its Ethernet, and Router One added the delay configured on its serial.

1.6 Feasible Distance, Reported Distance, and Feasible Successor

Feasible distance is the best metric along a path to a destination network, including the metric to the neighbor advertising that path. Reported distance is the total metric along a path to a destination network as advertised by an upstream neighbor. A feasible successor is a path whose reported distance is less than the feasible distance. Figure 3 illustrates this process:




Figure 3

Router One sees that it has two routes to Network A: one through Router Three and another through Router Four.

- The route through Router Four has a cost of 46021376 and a reported distance of 307200.
- The route through Router Three has a cost of 20307200 and a reported distance of 307200.

Note that in each case EIGRP calculates the reported distance from the router advertising the route to the network. In other words, the reported distance from Router Four is the metric to get to Network A from Router Four, and the reported distance from Router Three is the metric to get to Network A from Router Three. EIGRP chooses the route through Router Three as the best path, and uses the metric through Router Three as the feasible distance. Since the reported distance to this network through Router Four is less than the feasible distance, Router One considers the path through Router Four a feasible successor.

When the link between Routers One and Three goes down, Router One examines each path it knows to Network A and finds that it has a feasible successor through Router Four. Router One uses this route, using the metric through Router Four as the new feasible distance. The network converges instantly, and updates to downstream neighbors are the only traffic from the routing protocol.

Let's look at a more complex scenario, shown in Figure 4.

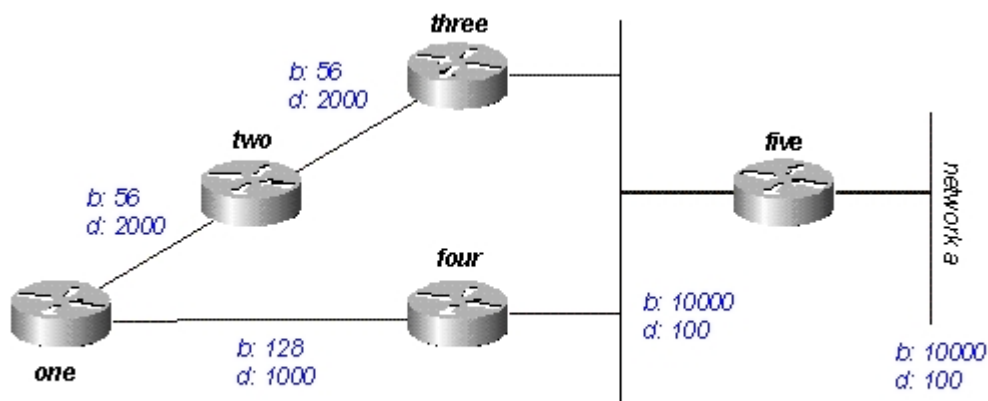


Figure 4

There are two routes to Network A from Router One: one through Router Two with a metric of 46789376 and another through Router Four with a metric of 20307200. Router One chooses the lower of these two metrics as its route to Network A, and this metric becomes the feasible distance. Next, let's look at the path through Router Two to see if it qualifies as a feasible successor. The reported distance from Router Two is 46277376, which is higher than the feasible distance - so this path is not a feasible successor. If you were to look in the topology table of Router One at this point (using **show ip eigrp topology**), you would only see one entry for Network A - through Router Four. (In reality there are two entries in the topology table at Router One, but only one will be a feasible successor, so the other will not be displayed in **show ip eigrp topology**; you can see the routes that are not feasible successors using **show ip eigrp topology all-links**).

Let's suppose that the link between Router One and Router Four goes down. Router One sees that it has lost its only route to Network A, and queries each of its neighbors (in this case, only Router Two) to see if they have a route to Network A. Since Router Two does have a route to Network A, it responds to the query. Since Router One no longer has the better route through Router Four, it accepts this route through Router Two to Network A.

1.7 Deciding if a Path is Loop-Free

How does EIGRP use the concepts of feasible distance, reported distance, and feasible successor to determine if a path is valid, and not a loop? In Figure 4a, Router Three examines routes to Network A. Since split horizon is disabled (for example, if these are multipoint Frame Relay interfaces), Router Three shows three routes to Network A: through Router Four, through Router Two (path is two, one, three, four), and through Router One (path is one, two, three, four).




Figure 4a

If Router Three accepts all of these routes, it results in a routing loop. Router Three thinks it can get to Network A through Router Two, but the path through Router Two passes through Router Three to get to Network A. If the connection between Router Four and Router Three goes down, Router Three believes it can get to Network A through one of the other paths, but because of the rules for determining feasible successors, it will never use these paths as alternates. Let's look at the metrics to see why:

- total metric to Network A through Router Four: 20281600
- total metric to Network A through Router Two: 47019776
- total metric to Network A through Router One: 47019776

Since the path through Router Four has the best metric, Router Three installs this route in the forwarding table and uses 20281600 as its feasible distance to Network A. Router Three then computes the reported distance to Network A through Routers Two and One: 47019776 for the path through Router Two, and 47019776 for the path through Router One. Because both of these metrics are greater than the feasible distance, Router Three doesn't install either route as a feasible successor for Network A.

Suppose that the link between Routers Three and Four goes down. Router Three queries each of its neighbors for an alternative route to Network A. Router Two receives the query and, because the query is from its successor, searches each of the other entries in its topology table to see if there is a feasible successor. The only other entry in the topology table is from Router One, with a reported distance equal to the last known best metric through Router Three. Because the reported distance through Router One is not less than the last known feasible distance, Router Two marks the route as unreachable and queries each of its neighbors - in this case, only Router One - for a path to Network A.

Router Three also sends a query for Network A to Router One. Router One examines its topology table and finds that the only other path to Network A is through Router Two with a reported distance equal to the last known feasible distance through Router Three. Once again, since the reported distance through Router Two is not less than the last known feasible distance, this route is not a feasible successor. Router One marks the route as unreachable and queries its only other neighbor, Router Two, for a path to Network A.

This is the first level of queries. Router Three has queried each of its neighbors in an attempt to find a route to Network A. In turn, Routers One and Two have marked the route unreachable, and queried each of their remaining neighbors in an attempt to find a path to Network A. When Router Two receives Router One's query, it examines its topology table and notes that the destination is marked as unreachable. Router Two replies to Router One that Network A is unreachable. When Router One receives Router Two's query, it also sends back a reply that Network A is unreachable. Now Routers One and Two have both concluded that Network A is unreachable, and they reply to Router Three's original query. The network has converged, and all routes return to the passive state.

2.0 Split Horizon and Poison Reverse

In the previous example, we assumed that split horizon was not in effect to show how EIGRP uses the feasible distance and the reported distance to determine if a route is likely to be a loop. In some circumstances, however, EIGRP uses split horizon to prevent routing loops as well. Before dealing with the details of how EIGRP uses split horizon, let's review what split horizon is and how it works. The split horizon rule states:

- Never advertise a route out of the interface through which you learned it.

For instance, in Figure 4a, if Router One is connected to Routers Two and Three through a single multipoint interface (such as Frame Relay), and Router One learned about Network A from Router Two, it won't advertise the route to Network A back out the same interface to Router Three. Router One assumes that Router Three would learn about Network A directly from Router Two.




Figure 4a

Poison reverse is another way of avoiding routing loops. Its rule states:

- Once you learn of a route through an interface, advertise it as unreachable back through that same interface.

Let's say the routers in Figure 4a have poison reverse enabled. When Router One learns about Network A from Router Two, it advertises Network A as unreachable through its link to Routers Two and Three. Router Three, if it shows any path to Network A through Router One, removes that path because of the unreachable advertisement. EIGRP combines these two rules to help prevent routing loops.

EIGRP uses split horizon or advertises a route as unreachable when:

- two routers are in startup mode (exchanging topology tables for the first time)
- advertising a topology table change
- sending a query

Let's examine each of these situations.

2.1 Startup Mode

When two routers first become neighbors, they exchange topology tables during startup mode. For each table entry a router receives during startup mode, it advertises the same entry back to its new neighbor with a maximum metric (poison route).

2.2 Topology Table Change

In Figure 5, Router One uses variance to balance the traffic destined to Network A between the two serial links - the 56k link between Routers Two and Four, and the 128k link between Routers Three and Four (see the Load Balancing section for a discussion of variance).




Figure 5

Router Two sees the path through Router Three as a feasible successor. If the link between Routers Two and Four goes down, Router Two simply re-converges on the path through Router Three. Since the split horizon rule states that you should never advertise a route out the interface through which you learned about it, Router Two would not normally send an update. However, this leaves Router One with an invalid topology table entry. When a router changes its topology table in such a way that the interface through which the router reaches a network changes, it turns off split horizon and poison reverses the old route out all interfaces. In this case, Router Two turns off split horizon for this route, and advertises Network A as unreachable. Router One hears this advertisement and flushes its route to Network A through Router Two from its routing table.

2.3 Queries

Queries result in a split horizon only when a router receives a query or update from the successor it is using for the destination in the query. Let's take a look at the network in Figure 6.




Figure 6

Router Three receives a query concerning 10.1.2.0/24 (which it reaches through Router One) from Router Four. If Three doesn't have a successor for this destination because a link flap or other temporary network condition, it sends a query to each of its neighbors; in this case, Routers One, Two, and Four. If, however, Router Three receives a query or update (such as a metric change) from Router One for the destination 10.1.2.0/24, it doesn't send a query back to Router One, because Router One is its successor to this network. Instead, it only sends queries to Routers Two and Four.

3.0 Stuck In Active Routes

In some circumstances, it takes a very long time for a query to be answered. So long, in fact, that the router that issued the query gives up and clears its connection to the router that isn't answering, effectively restarting the neighbor session. This is known as a stuck in active (SIA) route. The most basic SIA routes occur when it simply takes too long for a query to reach the other end of the network and for a reply to travel back. For instance, in Figure 7, Router One is recording a large number of SIA routes from Router Two.




Figure 7

After some investigation, the problem is narrowed down to the delay over the satellite link between Routers Two and Three. There are two possible solutions to this type of problem. The first is to increase the amount of time the router waits after sending a query before declaring the route SIA. This setting can be changed using the `timers active-time` command.

The better solution, however, is to redesign the network to reduce the range of queries (so very few queries pass over the satellite link). Query range is covered in the Query Range section. Query range in itself, however, is not a common reason for reported SIA routes. More often, some router on the network can't answer a query for one of the following reasons:

- the router is too busy to answer the query (generally due to high CPU utilization)
- the router is having memory problems, and cannot allocate the memory to process the query or build the reply packet
- the circuit between the two routers is not good - enough packets are getting through to keep the neighbor relationship up, but some queries or replies are getting lost between the routers
- unidirectional links (a link on which traffic can only flow in one direction because of a failure)

3.1 Troubleshooting SIA Routes

Troubleshooting SIA routes is generally a three-step process:

1. Find the routes that are consistently being reported as SIA.
2. Find the router that is consistently failing to answer queries for these routes.
3. Find the reason that router is not receiving or answering queries.

The first step should be fairly easy. If you're logging console messages, a quick perusal of the log indicates which routes are most frequently marked SIA. The second step is more difficult. The command to gather this information is **show ip eigrp topology active**:

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
```

```
A 10.2.4.0/24, 0 successors, FD is 512640000, Q
  1 replies, active 00:00:01, query-origin: Local origin
    via 10.1.2.2 (Infinity/Infinity), Serial1
  1 replies, active 00:00:01, query-origin: Local origin
    via 10.1.3.2 (Infinity/Infinity), r, Serial3
Remaining replies:
  via 10.1.1.2, r, Serial0
```

Any neighbors that show an **R** have yet to reply (the active timer shows how long the route has been active). Note that these neighbors may not show up in the Remaining replies section; they may appear among the other RDBs. Pay particular attention to routes that have outstanding replies and have been active for some time, generally two to three minutes. Run this command several times and you begin to see which neighbors are not responding to queries (or which interfaces seem to have a lot of unanswered queries). Examine this neighbor to see if it's consistently waiting for replies from any of its neighbors. Repeat this process until you find the router that is consistently not answering queries. You can look for problems on the link to this neighbor, memory or CPU utilization, or other problems with this neighbor.

If you run into a situation where it seems that the query range is the problem, it's always best to reduce the query range rather than increasing the SIA timer.

4.0 Redistribution

This section examines different scenarios involving redistribution.

4.1 Redistribution Between Two EIGRP Autonomous Systems

In Figure 8, the routers are configured as follows:




Figure 8

Router One

```
router eigrp 2000 (the 2000 on the end is the autonomous system)
network 172.16.1.0 0.0.0.255
```

Router Two

```
router eigrp 2000
 redistribute eigrp 1000
 network 172.16.1.0 0.0.0.255
!
router eigrp 1000
 redistribute eigrp 2000
 network 10.1.0.0 0.0.255.255
```

Router Three

```
router eigrp 1000
 network 10.1.0.0 0.0.255.255
```

Router Three is advertising the network 10.1.2.0/24 to Router Two through autonomous system 1000; Router Two is redistributing this route into autonomous system 2000 and advertising it to Router One. On Router One, we see:

```

one#show ip eigrp topology 10.1.2.0 255.255.255.0
IP-EIGRP topology entry for 10.1.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 46763776
  Routing Descriptor Blocks:
  20.1.1.1 (Serial0), from 20.1.1.1, Send flag is 0x0
    Composite metric is (46763776/46251776), Route is External
  Vector metric:
    Minimum bandwidth is 56 Kbit
    Total delay is 41000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
  External data:
    Originating router is 10.1.2.1
    AS number of route is 1000
    External protocol is EIGRP, external metric is 46251776
    Administrator tag is 0 (0x00000000)

```

Notice that although the link between Routers One and Two has a bandwidth of 1.544Mb, the minimum bandwidth shown in this topology table entry is 56k. This means that EIGRP preserves all metrics when redistributing between two EIGRP autonomous systems.

4.2 Redistribution Between EIGRP and IGRP in Two Different Autonomous Systems

In Figure 9, we have changed the configurations as follows:




Figure 9

Router One

```

router eigrp 2000
network 20.0.0.0

```

Router Two

```
router eigrp 2000
  redistribute igrp 1000
  network 20.0.0.0
!
router igrp 1000
  network 10.0.0.0
```

Router Three

```
router igrp 1000
  network 10.0.0.0
```

The configuration for Router One is shown below:

```
one#show ip eigrp topology 10.1.2.0 255.255.255.0
IP-EIGRP topology entry for 10.1.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 46763776
  Routing Descriptor Blocks:
  20.1.1.1 (Serial0), from 20.1.1.1, Send flag is 0x0
    Composite metric is (46763776/46251776), Route is External
    Vector metric:
      Minimum bandwidth is 56 Kbit
      Total delay is 41000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
    External data:
      Originating router is 10.1.1.1
      AS number of route is 1000
      External protocol is IGRP, external metric is 180671
      Administrator tag is 0 (0x00000000)
```

IGRP metrics are preserved when routes are redistributed into EIGRP with a different autonomous system, but they are scaled by multiplying the IGRP metric by the constant 256. There is one caveat to redistribution between IGRP and EIGRP that should be noted. If the network is directly connected to the router doing the redistribution, it advertises the route with a metric of 1.

For example, the network 10.1.1.0/24 is directly connected to Router Two, and IGRP is routing for this network (there is a network statement under router IGRP that covers this interface). EIGRP is not routing for this network, but is learning about this directly-connected interface

through redistribution from IGRP. On Router One, the topology table entry for 10.1.1.0/24 shows:

```
one#show ip eigrp topology 10.1.1.0 255.255.255.0
IP-EIGRP topology entry for 10.1.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2169856
  Routing Descriptor Blocks:
  20.1.1.1 (Serial0), from 20.1.1.1, Send flag is 0x0
    Composite metric is (2169856/1), Route is External

  Vector metric:
    Minimum bandwidth is 1544 Kbit
    Total delay is 20000 microseconds
    Reliability is 0/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1

  External data:
    Originating router is 10.1.1.1
    AS number of route is 1000
    External protocol is IGRP, external metric is 0
    Administrator tag is 0 (0x00000000)
```

Note that the reported distance from Router Two, which is highlighted in blue, is 1.

4.3 Redistribution Between EIGRP and IGRP in the Same Autonomous System

The following changes are made to the router configurations in Figure 10:




Figure 10

Router One

```
router eigrp 2000
network 20.0.0.0
```

Router Two

```
router eigrp 2000
network 20.0.0.0
!
router igrp 2000
network 10.0.0.0
```

Router Three

```
router igrp 2000
network 10.0.0.0
```

and Router One, is configured as follows:

```
one#show ip eigrp topology 10.1.2.0 255.255.255.0
IP-EIGRP topology entry for 10.1.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 46763776
  Routing Descriptor Blocks:
  20.1.1.1 (Serial0), from 20.1.1.1, Send flag is 0x0
    Composite metric is (46763776/46251776), Route is External
  Vector metric:
    Minimum bandwidth is 56 Kbit
    Total delay is 41000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
  External data:
    Originating router is 10.1.1.1
    AS number of route is 2000
    External protocol is IGRP, external metric is 180671
    Administrator tag is 0 (0x00000000)
```

This configuration looks amazingly like the earlier output when we were redistributing between two different autonomous systems running IGRP and EIGRP. The directly attached 10.1.1.0/24 network is handled the same way in both scenarios:

```

one#show ip eigrp topology 10.1.1.0 255.255.255.0
IP-EIGRP topology entry for 10.1.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2169856
  Routing Descriptor Blocks:
  20.1.1.1 (Serial0), from 20.1.1.1, Send flag is 0x0
    Composite metric is (2169856/1), Route is External
  Vector metric:
    Minimum bandwidth is 1544 Kbit
    Total delay is 20000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
  External data:
    Originating router is 10.1.1.1
    AS number of route is 2000
    External protocol is IGRP, external metric is 0
    Administrator tag is 0 (0x00000000)

```

So this network, which is directly connected to Router One, is redistributed from IGRP to EIGRP with a metric of 1 – the same metric we see when redistributing between two different autonomous systems.

There are two caveats with EIGRP/IGRP redistribution within the same autonomous system:

- Internal EIGRP routes are always preferred over external EIGRP or IGRP routes.
- External EIGRP route metrics are compared to scaled IGRP metrics (the administrative distance is ignored).

Let's examine these caveats in Figure 11:

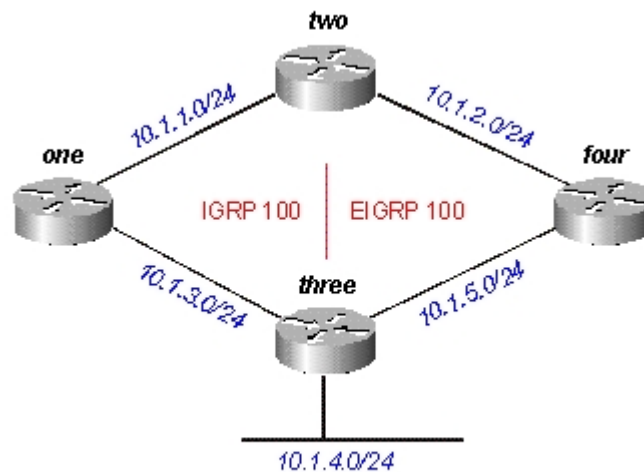


Figure 11

Router One advertises 10.1.4.0/24 in IGRP autonomous system 100; Router Four advertises 10.1.4.0/24 as an external in EIGRP autonomous system 100; Router Two runs both EIGRP and IGRP in autonomous system 100.

If we ignore the EIGRP route advertised by Router Four (by shutting down the link between Routers Two and Four, for instance), Router Two shows:

```
two#show ip route 10.1.4.0
Routing entry for 10.1.4.0/24
  Known via "igrp 100", distance 100, metric 12001
  Redistributing via igrp 100, eigrp 100
  Advertised by igrp 100 (self originated)
    eigrp 100
  Last update from 10.1.2.2 on Serial1, 00:00:42 ago
  Routing Descriptor Blocks:
  * 10.1.2.2, from 10.1.2.2, 00:00:42 ago, via Serial1
    Route metric is 12001, traffic share count is 1
    Total delay is 20010 microseconds, minimum bandwidth is 1000 Kbit
    Reliability 1/255, minimum MTU 1 bytes
    Loading 1/255, Hops 0
```

Note the administrative distance is 100. When we add the EIGRP route, Router Two shows:

```
two#show ip route 10.1.4.0
Routing entry for 10.1.4.0/24
  Known via "eigrp 100", distance 170, metric 3072256, type external
  Redistributing via igrp 100, eigrp 100
  Last update from 10.1.1.2 on Serial0, 00:53:59 ago
  Routing Descriptor Blocks:
  * 10.1.1.2, from 10.1.1.2, 00:53:59 ago, via Serial0
    Route metric is 3072256, traffic share count is 1
    Total delay is 20010 microseconds, minimum bandwidth is 1000 Kbit
    Reliability 1/255, minimum MTU 1 bytes
    Loading 1/255, Hops 1
```

Note the metrics for these two routes are the same after being scaled from IGRP to EIGRP (see the Metrics section):

$$12001 \times 256 = 3072256$$

where 12001, an IGRP metric, is through Router One; and 3072256, an EIGRP metric, is through Router Four.

Router Two prefers the EIGRP external route with the same metric (after scaling) and a higher administrative distance. This is true whenever automatic redistribution occurs between EIGRP and IGRP within the same autonomous system. The router always prefers the path with the lowest cost metric and ignores the administrative distance.

4.4 Redistribution To and From Other Protocols

Redistribution between EIGRP and other protocols - RIP and OSPF, for example - works in the same way as all redistribution. It is always best to use the default metric when redistributing between protocols. You should be aware of the following two issues when redistributing between EIGRP and other protocols:

- Routes redistributed into EIGRP are not always summarized - see the Summarization section for an explanation.
- External EIGRP routes have an administrative distance of 170.

4.5 Redistribution of Static Routes to Interfaces

When you install a static route to an interface, and configure a network statement using `router eigrp`, which includes the static route, EIGRP redistributes this route as if it were a directly connected interface. Let's look at the network in Figure 12.




Figure 12

Router One has a static route to the network 172.16.1.0/24 configured through interface Serial 0:

```
ip route 176.16.1.0 255.255.255.0 Serial0
```

And Router One also has a network statement for the destination of this static route:

```
router eigrp 2000
 network 10.0.0.0
 network 176.16.0.0
 no auto-summary
```

Router One redistributes this route, even though it is not redistributing static routes, because EIGRP considers this a directly attached network. On Router Two, this looks as follows:

```
two#show ip route
....
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C    10.1.1.0/24 is directly connected, Serial0
D    10.1.2.0/24 [90/2169856] via 10.1.1.1, 00:00:47, Serial0
    176.16.0.0/24 is subnetted, 1 subnets
D    172.16.1.0 [90/2169856] via 10.1.1.1, 00:00:47, Serial0
```

Note the route to 176.16.1.0/24 appears as an internal EIGRP route on Router Two.

5.0 Summarization

There are two forms of summarization in EIGRP: auto-summaries and manual summaries.

5.1 Auto-summarization

EIGRP performs an auto-summarization each time it crosses a border between two different major networks. For example, in Figure 13, Router Two advertises only the 10.0.0.0/8 network to Router One, because the interface Router Two uses to reach Router One is in a different major network.




Figure 13

On Router One, this looks like the following:

```
one#show ip eigrp topology 10.0.0.0
IP-EIGRP topology entry for 10.0.0.0/8
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 11023872
  Routing Descriptor Blocks:
  172.16.1.1 (Serial0), from 172.16.1.2, Send flag is 0x0
    Composite metric is (11023872/10511872), Route is Internal
  Vector metric:
    Minimum bandwidth is 256 Kbit
    Total delay is 40000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
```

This route is not marked as a summary route in any way - it looks like an internal route. The metric is the best metric from among the summarized routes. Note that the minimum bandwidth on this route is 256k, although there are links in the 10.0.0.0 network that have a bandwidth of 56k.

On the router doing the summarization, a route is built to null0 for the summarized address:

```
two#show ip route 10.0.0.0
Routing entry for 10.0.0.0/8, 4 known subnets
  Attached (2 connections)
  Variably subnetted with 2 masks
  Redistributing via eigrp 2000

C    10.1.3.0/24 is directly connected, Serial2
D    10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D    10.0.0.0/8 is a summary, 00:23:20, Null0
C    10.1.1.0/24 is directly connected, Serial1
```

The route to 10.0.0.0/8 is marked as a summary through Null0. The topology table entry for this summary route looks like the following:

```
two#show ip eigrp topology 10.0.0.0
IP-EIGRP topology entry for 10.0.0.0/8
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 10511872
  Routing Descriptor Blocks:
  0.0.0.0 (Null0), from 0.0.0.0, Send flag is 0x0
    (note: the 0.0.0.0 here means this route is originated by this router)
  Composite metric is (10511872/0), Route is Internal
  Vector metric:
    Minimum bandwidth is 256 Kbit
    Total delay is 20000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 0
```

To make Router Two advertise the components of the 10.0.0.0 network instead of a summary, configure **no auto-summary** on the EIGRP process on Router Two:

On Router Two

```
router eigrp 2000
  network 20.0.0.0
  network 10.0.0.0
  no auto-summary
```

With auto-summary turned off, Router One now sees all of the components of the 10.0.0.0 network:

```

one#show ip eigrp topology
IP-EIGRP Topology Table for process 2000
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
P 10.1.3.0/24, 1 successors, FD is 46354176
    via 20.1.1.1 (46354176/45842176), Serial0
P 10.1.2.0/24, 1 successors, FD is 11049472
    via 20.1.1.1 (11049472/10537472), Serial0
P 10.1.1.0/24, 1 successors, FD is 11023872
    via 20.1.1.1 (11023872/10511872), Serial0
P 172.16.1.0/24, 1 successors, FD is 2169856
    via Connected, Serial0

```

There are some caveats when dealing with the summarization of external routes that are covered later in the Summary of Externals section.

5.2 Manual Summarization

EIGRP allows you to summarize internal and external routes on virtually any bit boundary using manual summarization. For example, in Figure 14, Router Two is summarizing the 192.1.1.0/24, 192.1.2.0/24, and 192.1.3.0/24 into the CIDR block 192.1.0.0/22.




Figure 14

The configuration on Router Two is shown below:

```

two#show run
....
!
interface Serial0
 ip address 10.1.50.1 255.255.255.0
 ip summary-address eigrp 2000 192.1.0.0 255.255.252.0
 no ip mroute-cache
!
....

```

```
two#show ip eigrp topology
```

```
IP-EIGRP Topology Table for process 2000
```

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
       r - Reply status
```

```
P 10.1.10.0/24, 1 successors, FD is 45842176  
   via Connected, Loopback0  
P 10.1.50.0/24, 1 successors, FD is 2169856  
   via Connected, Serial0  
P 192.1.1.0/24, 1 successors, FD is 10511872  
   via Connected, Serial1  
P 192.1.0.0/22, 1 successors, FD is 10511872  
   via Summary (10511872/0), Null0  
P 192.1.3.0/24, 1 successors, FD is 10639872  
   via 192.1.1.1 (10639872/128256), Serial1  
P 192.1.2.0/24, 1 successors, FD is 10537472  
   via 192.1.1.1 (10537472/281600), Serial1
```

Note the **ip summary-address** command under interface Serial0, and the summary route via Null0. On Router One, we see this as an internal route:

```
one#show ip eigrp topology
```

```
IP-EIGRP Topology Table for process 2000
```

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
       r - Reply status
```

```
P 10.1.10.0/24, 1 successors, FD is 46354176  
   via 10.1.50.1 (46354176/45842176), Serial0  
P 10.1.50.0/24, 1 successors, FD is 2169856  
   via Connected, Serial0  
P 192.1.0.0/22, 1 successors, FD is 11023872  
   via 10.1.50.1 (11023872/10511872), Serial0
```

5.3 Auto-summarization of External Routes

EIGRP will not auto-summarize external routes unless there is a component of the same major network that is an internal route. To illustrate, let's look at Figure 15.




Figure 15

Router Three is injecting an external route to 192.1.1.0/25 into EIGRP using **redistribute connected**, as shown in the configurations below.

Router Three

```
interface Ethernet0
 ip address 192.1.2.1 255.255.255.192
 !
interface Ethernet1
 ip address 192.1.2.65 255.255.255.192
 !
router eigrp 2000
 redistribute connected
 network 10.0.0.0
 default-metric 10000 1 255 1 1500
```

With this configuration on Router Three, the routing table on Router One shows:

```
one#show ip route
....
 10.0.0.0/8 is subnetted, 2 subnets
D    10.1.2.0 [90/11023872] via 10.1.50.2, 00:02:03, Serial0
C    10.1.1.0 is directly connected, Serial0
    192.1.2.0/24 is subnetted, 1 subnets
D EX  192.1.1.0 [170/11049472] via 10.1.50.2, 00:00:53, Serial0
D EX  192.1.2.64 [170/11049472] via 10.1.50.2, 00:00:53, Serial0
```

(Although the network mask is not shown here, it is 26 bits - /26)

Although auto-summary normally causes Router Two to summarize the 192.1.2.0/26 and 192.1.2.128/26 routes into one major net destination (192.1.1.0/24), it doesn't do this because both routes are external. However, if you reconfigure the link between Routers Two and Three to 192.1.1.128/26, and add network statements for this network on Routers Two and Three, the 192.1.1.0/24 auto-summary is then generated on Router Two.

Router Three

```
interface Ethernet0
 ip address 192.1.2.1 255.255.255.192
!
interface Ethernet1
 ip address 192.1.2.65 255.255.255.192
!
interface Serial0
 ip address 192.1.1.130 255.255.255.192
!
router eigrp 2000
 network 192.1.1.0
```

Now Router Two generates the summary for 192.1.1.0/24:

```
two#show ip route
....
D    192.1.1.0/24 is a summary, 00:06:48, Null0
....
```

And Router One shows only the summary route:

```
one#show ip route
....
    10.0.0.0/8 is subnetted, 1 subnets
C    10.1.1.0 is directly connected, Serial0
D    192.1.1.0/24 [90/11023872] via 10.1.1.2, 00:00:36, Serial0
```

6.0 Query Processing and Range

When a router processes a query from a neighbor, the following rules apply:

| Query from | Route state | Action |
|--------------------------------------|--|---|
| neighbor (not the current successor) | passive | reply with current successor information |
| successor | passive | attempt to find new successor; if successful, reply with new information; if not successful, mark destination unreachable and query all neighbors except the previous successor |
| any neighbor | no path through this neighbor before query | reply with best path currently known |
| any neighbor | not known before query | reply that the destination is unreachable |
| neighbor (not the current successor) | active | if there is no current successor to this destinations (normally this would be true), reply with an unreachable |
| | | if there is a good successor, reply with the current path information |
| successor | active | attempt to find new successor; if successful, reply with new information; if not successful, mark destination unreachable and query all neighbors except the previous successor |

The actions in the table above impact the range of the query in the network by determining how many routers receive and reply to the query before the network converges on the new topology. To see how these rules affect the way queries are handled, let's look at the network in Figure 16, which is running under normal conditions.

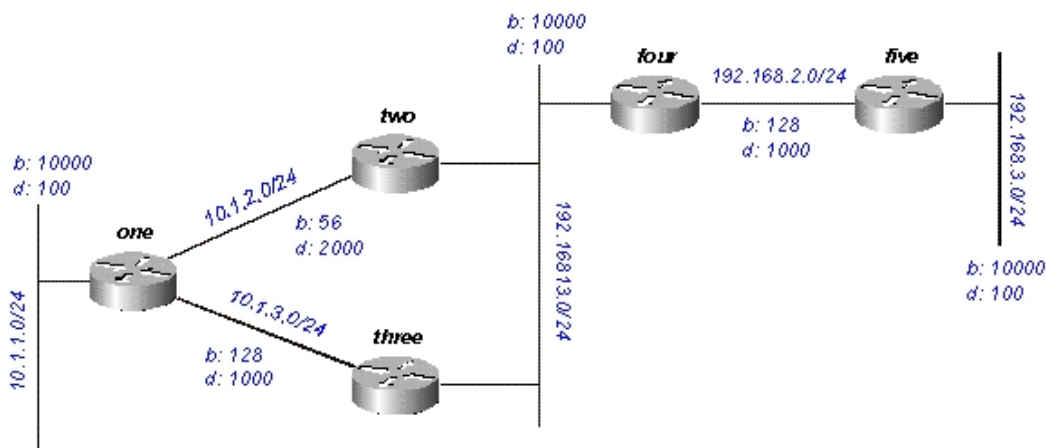


Figure 16

We can expect the following to happen regarding network 192.168.3.0/24 (far right side):

- Router One has two paths to 192.168.3.0/24:
 - ✧ through Router Two with a distance of 46533485 and a reported distance of 20307200
 - ✧ through Router Three with a distance of 20563200 and a reported distance of 20307200

- Router One chooses the path through Router Three and keeps the path through Router Two as a feasible successor
- Routers Two and Three show one path to 192.168.3.0/24 through Router Four

Suppose that 192.168.3.0/24 fails. What activity can we expect to see on this network? Figures 16a through 16h illustrate the process.

Router Five marks 192.168.3.0/24 as unreachable, and queries Router Four:




Figure 16a

Router Four, upon receiving a query from its successor, attempts to find a new feasible successor to this network. It doesn't find one, so it marks 192.168.3.0/24 as unreachable and query Routers Two and Three:




Figure 16b

Routers Two and Three, in turn, see that they have lost their only feasible route to 192.168.3.0/24, and mark it as unreachable; they both send queries to Router One:




Figure 16c

For simplicity, let's assume that Router One receives the query from Router Three first, and marks the route as unreachable. Router One then receives the query from Router Two. Although another order is possible, they will all have the same final result.




Figure 16d

Router One replies to both queries with unreachables; Router One is now passive for 192.168.3.0/24:




Figure 16e

Routers Two and Three reply to the query from Router Four; Routers Two and Three are now passive for 192.168.3.0/24:




Figure 16f

Router Five, upon receiving the reply from Router Four, removes network 192.168.3.0/24 from its routing table; Router Five is now passive for network 192.168.3.0/24. Router Five sends updates back to Router Four so the route is removed from the topology and routing tables of the remaining routers.




Figure 16g

It's important to understand that although there may be other query paths or processing orders, all routers in the network process a query for network 192.168.3.0/24 when that link goes down. Some routers may end up processing more than one query (Router One in this example). In fact, if the queries were to reach the routers in a different order, some would end up processing three or four queries. This is a good example of an unbounded query in an EIGRP network.

6.1 How Summarization Points Affect the Query Range

Now let's look at the paths to 10.1.1.0/24 in the same network:

- Router Two has a topology table entry for the 10.1.1.0/24 network with a cost of 46251885 through Router One.
- Router Three has a topology table entry for the 10.1.1.0/24 network with a cost of 20281600 through Router One.
- Router Four has a topology table entry for the 10.0.0.0/8 network (because Routers Two and Three are autosummarizing to the major network boundary) through Router Three with a metric of 20307200 (the reported distance through Router Two is higher than the total metric through Router Three, so the path through Router Two is not a feasible successor).

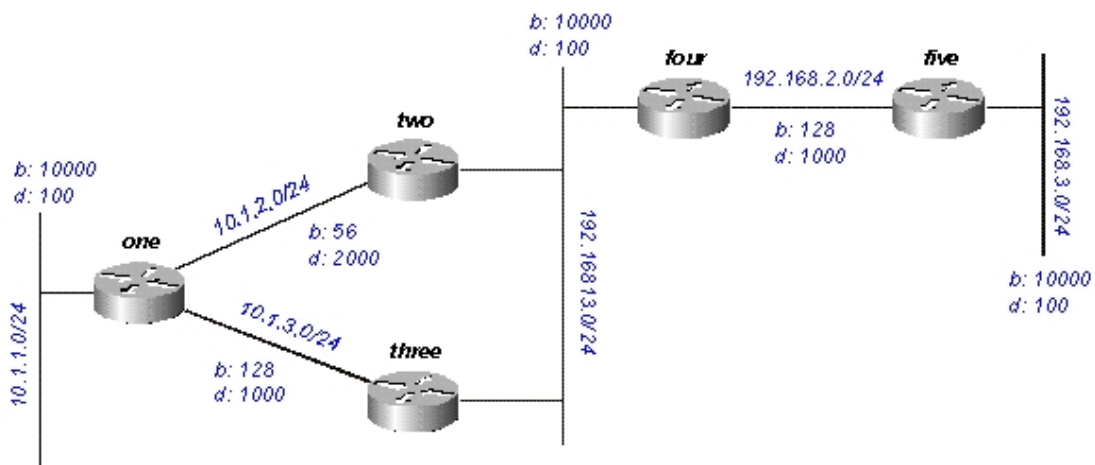


Figure 17

If 10.1.1.0/24 goes down, Router One marks it as unreachable, and then queries each of its neighbors (Routers Two and Three) for a new path to that network:




Figure 17a

Router Two, on receiving the query from Router One, marks the route as unreachable (because the query is from its successor) and then queries Routers Four and Three:

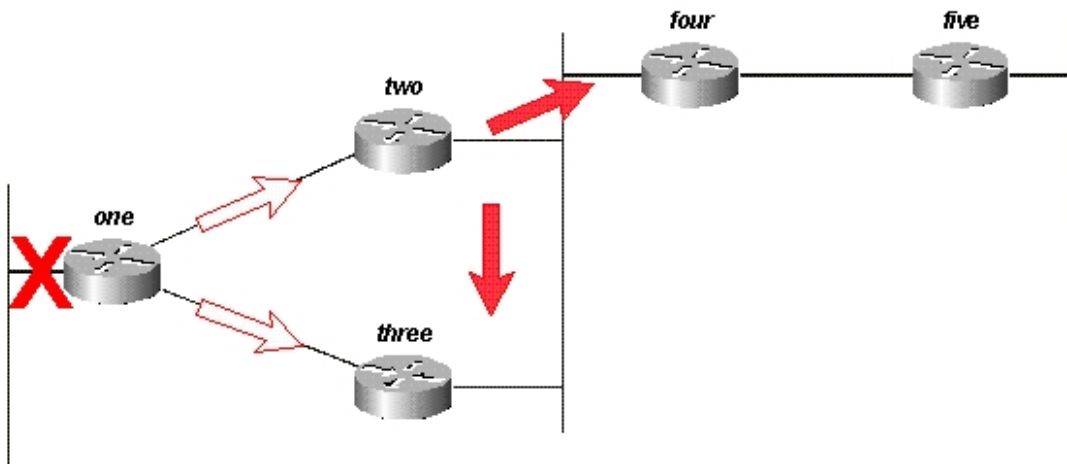


Figure 17b

Router Three, when it receives the query from Router One, marks the destination as unreachable and queries Routers Two and Four:




Figure 17c

Router Four, when it receives the queries from Routers Two and Three, replies that 10.1.10/24 is unreachable (note that Router Four has no knowledge of the subnet in question, since it only has the 10.0.0.0/8 route):




Figure 17d

Routers Two and Three reply to each other that 10.1.10/24 is unreachable:




Figure 17e

Since Routers Two and Three now have no outstanding queries, they both reply to Router One that 10.1.1.0/24 is unreachable:




Figure 17f

The query, in this case, is bounded by the autosummarization at Routers Two and Three. Router Five does not participate in the query process, and is not involved in the re-convergence of the network. Queries can also be bound by manual summarization, autonomous system borders, and distribution lists.

6.2 How Autonomous System Boundaries Affect the Query Range

If a router is redistributing routes between two EIGRP autonomous systems, it replies to the query within the normal processing rules and launches a new query into the other autonomous system. For example, if the link to the network attached to Router Three goes down, Router Three marks the route unreachable and queries Router Two for a new path:




Figure 18a

Router Two replies that this network is unreachable and launches a query into autonomous system 200 toward Router One. Once Router Three receives the reply to its original query, it removes the route from its table. Router Three is now passive for this network:




Figure 18b

Router One replies to Router Two, and the route goes passive:




Figure 18c

While the original query didn't propagate throughout the network (it was bound by the autonomous system border), the original query leaks into the second autonomous system in the form of a new query. This technique may help to prevent stuck in active (SIA) problems in a network by limiting the number of routers a query must pass through before being answered, but it does not solve the overall problem that each router must process the query. In fact, this method of bounding a query may worsen the problem by preventing the auto-summarization of routes that would otherwise be summarized (external routes are not summarized unless there is an external component in that major network).

6.3 How Distribution Lists Affect the Query Range

Rather than block the propagation of a query, distribution lists in EIGRP mark any query reply as unreachable. Let's use Figure 19 as an example.




Figure 19

In the figure above:

- Router Three has a distribute-list applied against its serial interfaces that only permits it to advertise Network B.
- Routers One and Two don't know that Network A is reachable through Router Three (Router Three is not used as a transit point between Routers One and Two).
- Router Three uses Router One as its preferred path to Network A, and does not use Router Two as a feasible successor.

When Router One loses its connection to Network A, it marks the route as unreachable and sends a query to Router Three. Router Three doesn't advertise a path to Network A because of the distribution list on its serial ports.




Figure 19a

Router Three marks the route as unreachable, then queries Router Two:




Figure 19b

Router Two examines its topology table and finds that it has a valid connection to Network A. Note the query was not affected by the distribution list in Router Three:




Figure 19c

Router Two replies that Network A is reachable; Router Three now has a valid route:




Figure 19d

Router Three builds the reply to the query from Router One, but the distribution list causes Router Three to send a reply that Network A is unreachable, even though Router Three has a valid route to Network A:




Figure 19e

7.0 Pacing Packets

Some routing protocols consume all of the available bandwidth on a low bandwidth link while they are converging (adapting to a change in the network). EIGRP avoids this congestion by pacing the speed at which packets are transmitted on a network, thereby using only a portion of the available bandwidth. The default configuration for EIGRP is to use up to 50 percent of the available bandwidth, but this can be changed with the following command:

```
router(config-if)#ip bandwidth-percent eigrp 2 ?
<1-999999> Maximum bandwidth percentage that EIGRP may use
```

Essentially, each time EIGRP queues a packet to be transmitted on an interface, it uses the following formula to determine how long to wait before sending the packet:

$$((1/\text{bandwidth}) \times (\text{packet size in bytes} \times 8)) \times \text{percent of bandwidth allowed}$$

For instance, if EIGRP queues a packet to be sent over a serial interface that has a bandwidth of 56k, and the packet is 512 bytes, EIGRP waits:

$$\begin{aligned} &((1/56000) \times (512 \times 8)) \times .5 \text{ (50\% of the bandwidth)} \\ &0.0731 \text{ seconds} \times .5 \text{ (50\% of the bandwidth)} \\ &0.0366 \text{ seconds} \end{aligned}$$

This allows a packet (or groups of packets) of at least 512 bytes to be transmitted on this link before EIGRP sends its packet. The pacing timer determines when the packet is sent, and is typically expressed in milliseconds. The pacing time for the packet in the above example is 36.6. There is a field in show ip eigrp interface that displays the pacing timer, as shown below:

```
router#show ip eigrp interface
IP-EIGRP interfaces for process 2
```

| Xmit Queue | Mean | Pacing Time | Multicast | Pending | | |
|------------|-------|-------------|-----------|-------------|------------|--------|
| Interface | Peers | Un/Reliable | SRTT | Un/Reliable | Flow Timer | Routes |
| Se0 | 1 | 0/0 | 28 | 0/15 | 127 | 0 |
| Se1 | 1 | 0/0 | 44 | 0/15 | 211 | 0 |

The time displayed is the pacing interval for the maximum transmission unit (MTU), the largest packet that can be sent over the interface.

8.0 Default Routing

There are two ways to inject a default route into EIGRP: redistribute a static route or summarize to 0.0.0.0/0. Use the first method when you want to draw all traffic to unknown destinations to a default route at the core of the network. This method is effective for advertising connections to the Internet. For example:

```
ip route 0.0.0.0 0.0.0.0 x.x.x.x (next hop to the internet)
!
router eigrp 100
 redistribute static
 default-metric 10000 1 255 1 1500
```

Summarizing to a default route is effective only when you want to provide remote sites with a default route. Since summaries are configured per interface, you don't need to worry about using distribute-lists or other mechanisms to prevent the default route from being propagated toward the core of your network. Note that a summary to 0.0.0.0/0 overrides a default route learned from any other routing protocol. The only way to configure a default route on a router using this method is to configure a static route to 0.0.0.0/0. (Beginning in Cisco IOS Software 12.0(4)T, you can also configure an administrative distance on the end of the **summary-address** command, so the local summary doesn't override the 0.0.0.0/0 route).

```
router eigrp 100
 network 10.0.0.0
!
interface serial 0
 encapsulation frame-relay
 no ip address
!
interface serial 0.1 point-to-point
 ip address 10.1.1.1
 frame-relay interface-dlci 10
 ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

9.0 Load Balancing

EIGRP puts up to four routes of equal cost in the routing table, which the router then load-balances. The type of load balancing (per packet or per destination) depends on the type of switching being done in the router. EIGRP, however, can also load-balance over unequal cost links.

Note Using **max-paths**, you can configure EIGRP to use up to six routes of equal cost.

Let's say there are four paths to a given destination, and the metrics for these paths are:

- path 1: 1100
- path 2: 1100
- path 3: 2000
- path 4: 4000

The router, by default, places traffic on both path 1 and 2. Using EIGRP, you can use the variance command to instruct the router to also place traffic on paths 3 and 4. The variance is a multiplier: traffic will be placed on any link that has a metric less than the best path multiplied by the variance. To load balance over paths 1, 2, and 3, use variance 2, because $1100 \times 2 = 2200$, which is greater than the metric through path 3. Similarly, to also add path 4, issue variance 4 under the **router eigrp** command.

How does the router divide the traffic between these paths? It divides the metric through each path into the largest metric, rounds down to the nearest integer, and uses this number as the traffic share count.

```
router#show ip route 10.1.4.0
Routing entry for 10.1.4.0/24
  Known via "igrp 100", distance 100, metric 12001
  Redistributing via igrp 100, eigrp 100
  Advertised by igrp 100 (self originated)
    eigrp 100
  Last update from 10.1.2.2 on Serial1, 00:00:42 ago
  Routing Descriptor Blocks:
  * 10.1.2.2, from 10.1.2.2, 00:00:42 ago, via Serial1
    Route metric is 12001, traffic share count is 1
    Total delay is 20010 microseconds, minimum bandwidth is 1000 Kbit
    Reliability 1/255, minimum MTU 1 bytes
    Loading 1/255, Hops 0
```

For this example, the traffic share counts are:

- for paths 1 and 2: $4000/1100 = 3$
- for path 3: $4000/2000 = 2$
- for path 4: $4000/4000 = 1$

The router sends the first three packets over path 1, the next three packets over path 2, the next two packets over path 3, and the next packet over path 4. The router then restarts by sending the next three packets over path 1, and so on.

10.0 Using the Metrics

When you initially configure EIGRP, remember these two basic rules if you're attempting to influence EIGRP metrics:

- The bandwidth should always be set to the real bandwidth of the interface; multipoint serial links and other mismatched media speed situations are the exceptions to this rule.
- The delay should always be used to influence EIGRP's routing decisions.

Because EIGRP uses the interface bandwidth to determine the rate at which to send packets, it is important that these be set correctly. If it's necessary to influence the path EIGRP chooses, always use delay to do so.

At lower bandwidths, the bandwidth has more influence over the total metric; at higher bandwidths, the delay has more influence over the total metric.

11.0 Using Administrative Tags in Redistribution

External administrative tags are useful for breaking redistribution routing loops between EIGRP and other protocols. By tagging the route when it is redistributed into EIGRP, you can block redistribution from EIGRP into the external protocol. A basic example of configuring these tags follows, but this example doesn't show the entire configuration used for breaking redistribution loops.




Figure 20

Router Three, which is redistributing routes connected into EIGRP, shows:

Router Three, which is redistributing routes connected into EIGRP, shows:

three#**show run**

```
....
interface Loopback0
 ip address 172.19.1.1 255.255.255.0
!
interface Ethernet0
 ip address 172.16.1.1 255.255.255.0
 loopback
 no keepalive
!
interface Serial0
 ip address 172.17.1.1 255.255.255.0

....
router eigrp 444
 redistribute connected route-map foo
 redistribute rip
 network 172.17.0.0
 default-metric 10000 1 255 1 1500

....
access-list 10 permit 172.19.0.0 0.0.255.255
route-map foo permit 10
 match ip address 10
 set tag 1

....
```

three#**show ip eigrp topo**

IP-EIGRP Topology Table for process 444

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - Reply status

```
P 172.17.1.0/24, 1 successors, FD is 2169856
   via Connected, Serial0
   via Redistributed (2169856/0)
P 172.16.1.0/24, 1 successors, FD is 281600
   via Redistributed (281600/0)
P 172.19.1.0/24, 1 successors, FD is 128256, tag is 1
   via Redistributed (128256/0)
```


Router Two, which is redistributing routes from EIGRP into RIP, shows:

```
two#show run
```

```
....  
interface Serial0  
 ip address 172.17.1.2 255.255.255.0  
!  
interface Serial1  
 ip address 172.18.1.3 255.255.255.0  
  
....  
router eigrp 444  
 network 172.17.0.0  
!  
router rip  
 redistribute eigrp 444 route-map foo  
 network 10.0.0.0  
 network 172.18.0.0  
 default-metric 1  
!  
no ip classless  
ip route 1.1.1.1 255.255.255.255 Serial0  
route-map foo deny 10  
 match tag 1  
!  
route-map foo permit 20  
  
....
```

```
two#show ip eigrp topo
```

```
IP-EIGRP Topology Table for process 444
```

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
       r - Reply status
```

```
P 172.17.1.0/24, 1 successors, FD is 2169856  
   via Connected, Serial0  
P 172.16.1.0/24, 1 successors, FD is 2195456  
   via 172.17.1.1 (2195456/281600), Serial0  
P 172.19.1.0/24, 1 successors, FD is 2297856, tag is 1  
   via 172.17.1.1 (2297856/128256), Serial0
```

Note the tag 1 on 172.19.1.0/24.

Router One, which is receiving the RIP routes redistributed by Router B, shows:

```
one#show run
```

```
....  
interface Serial0  
 ip address 172.18.1.2 255.255.255.0  
 no fair-queue  
 clockrate 1000000  
  
router rip  
 network 172.18.0.0
```

```
....
```

```
one#show ip route
```

```
Gateway of last resort is not set
```

```
R   172.16.0.0/16 [120/1] via 172.18.1.3, 00:00:15, Serial0  
R   172.17.0.0/16 [120/1] via 172.18.1.3, 00:00:15, Serial0  
    172.18.0.0/24 is subnetted, 1 subnets  
C    172.18.1.0 is directly connected, Serial0
```

Note that 172.19.1.0/24 is missing.

12.0 Using the **show ip eigrp topology** command

Using the **show ip eigrp topology** command

This command only displays feasible successors. To display all entries in the topology table, use the **show ip eigrp topology all-links** command. Click on each output field below to jump to an explanation.

| Clickable output for show ip eigrp topology |
|--|
| <pre>Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply, r - Reply status A 10.2.4.0/24, 0 successors, FD is 512640000, tag is 0x0, Q 1 replies, active 00:00:01, query-origin: Local origin via 10.1.2.2 (Infinity/Infinity), r, Q, Serial1 Remaining replies: via 10.1.1.2, r, Serial0 P 10.3.9.0/24, 1 successors, FD is 512640000, U * via 10.1.2.2 (512640000/128256), Serial1</pre> |

A means active. This could also show a **P**, meaning passive.

10.2.4.0/24 is the destination or mask.

0 successors shows how many successors (or paths) are available for this destination; if successors is capitalized, the route is in transition.

FD is 512640000 shows the feasible distance, which is the best metric to reach this destination or the best metric known when the route went active.

tag is 0x0 can be set and/or filtered using route maps with the **set tag** and **match tag** commands.

Q means a query is pending. This field can also be: **U**, for update pending; or **R**, for reply pending.

1 replies shows the number of outstanding replies.

active 00:00:01 shows how long this route has been active.

query origin: Local origin shows this route originated the query. This field can also be: Multiple origins, meaning that multiple neighbors have sent queries on this destination, but not the successor; or Successor origin, meaning the successor originated the query.

via 10.1.2.2 shows that we learned of this route from a neighbor whose IP address is 10.1.2.2. This field can also be: Connected, if the network is directly connected to this router; Redistributed, if this route is being redistributed into EIGRP on this router; or Summary, if this is a summary route generated on this router.

(Infinity/Infinity) shows the metric to reach this path through this neighbor in the first field, and the reported distance through this neighbor in the second field.

r shows that we've queried this neighbor and are waiting for a reply.

Q is the send flag for this route, meaning there is a query pending. This field can also be: U, meaning there is an update pending; or R, meaning there is a reply pending.

Serial1 is the interface through which this neighbor is reachable.

Via 10.1.1.2 shows the neighbor from which we are waiting for a reply.

r shows that we have queried this neighbor about the route and have not yet received a reply.

Serial0 is the interface through which this neighbor is reachable.

Via 10.1.2.2 (512640000/128256), Serial1 shows we are using this route (indicates which path the next path/destination will take when there are multiple routes of equal cost).

show ip eigrp topology <network>

This command displays all entries in the topology table for this destination, not just feasible successors. Click on each output field below to jump to an explanation.

Clickable output of show ip eigrp topology <network>

```
IP-EIGRP topology entry for 20.0.0.0/8
State is Passive, Query origin flag is 1, 2 Successor(s), FD is 307200
Routing Descriptor Blocks:
 10.1.1.2 (Ethernet1), from 10.1.1.2, Send flag is 0x0
   Composite metric is (307200/281600), Route is Internal
   Vector metric:
     Minimum bandwidth is 10000 Kbit
     Total delay is 2000 microseconds
     Reliability is 0/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 2
 10.1.2.2 (Ethernet0), from 10.1.2.2, Send flag is 0x0
   Composite metric is (307200/281600), Route is Internal
   Vector metric:
     Minimum bandwidth is 10000 Kbit
     Total delay is 2000 microseconds
     Reliability is 0/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 2
```

If the route is external, the following information is included:

Clickable output for external route info

```
External data:
  Originating router is 10.1.2.2
  AS number of route is 0
  External protocol is Static, external metric is 0
  Administrator tag is 0 (0x00000000)
```

State is Passive means the network is in passive state, or, in other words, we are not looking for a path to this network. Routes are almost always in a passive state in stable networks.

Query origin flag is 1 If this route is active, this field provides information on who originated the query.

0: This route is active but no query has been originated for it (we are searching for a feasible successor locally).

1: This router originated the query for this route (or the route is passive).

2: Multiple diffusing computations for this query. This router has received more than one query for this route from more than one source.

3: The router that we learned the path to this network from is querying for another route.

4: Multiple query sources for this route, including the router through which we learned this route. Similar to 2, but it also means there is a query origin string which describes the queries outstanding for this path.

2 Successor(s) means there are two feasible paths to this network.

FD is 307200 shows the best current metric to this network. If the route is active, this indicates the metric of the path we were previously using to route packets to this network.

Routing Descriptor Blocks Each of the following entries describes one path to the network.

10.1.1.2 (Ethernet1) is the next hop to the network and the interface that next hop is reached through.

from 10.1.2.2 is the source of this path information.

Send flag is:

0x0 If there are packets that need to be sent in relation to this entry, this indicates the type of packet.

0x1: This router has received a query for this network, and needs to send a unicast reply.

0x2: This route is active, and a multicast query should be sent.

0x3: This route has changed, and a multicast update should be sent.

Composite metric is (307200/281600) shows the total calculated costs to the network. The first number in the parentheses is the total cost to the network through this path, including the cost to the next hop. The second number in the parentheses is the reported distance, or, in other words, the cost the next hop router uses.

Route is Internal means this route was originated within this EIGRP autonomous system (AS). If the route was redistributed into this EIGRP AS, this field would indicate that the route is External.

Vector metric shows the individual metrics used by EIGRP to calculate the cost to a network. EIGRP does not propagate total cost information throughout the network; the vector metrics

are propagated, and each router computes the cost and reported distance individually.

Minimum bandwidth is 10000 Kbit shows the lowest bandwidth on the path to this network.

Total delay is 2000 microseconds shows the sum of the delays on the path to this network.

Reliability is 0/255 shows a reliability factor. This number is calculated dynamically, but is not used by default in metric calculations.

Load is 1/255 indicates the amount of load the link is carrying. This number is calculated dynamically, and is not used by default when EIGRP calculates the cost to use this path.

Minimum MTU is 1500 This field is not used in metric calculations.

Hop count is 2 This is not used in metric calculations, but does limit the maximum size of an EIGRP AS. The maximum number of hops that EIGRP will accept is 100 by default, although the maximum can be configured to 220 with metric maximum hops.

External data:

Originating Router shows that this is the router that injected this route into the EIGRP AS.

External AS shows the Autonomous System this route came from (if there is one).

External Protocol shows the protocol this route came from (if there is one).

external metric shows the internal metric in the external protocol.

Administrator Tag can be set and/or filtered using route maps with the **set tag** and **match tag** commands.

show ip eigrp topology [active | pending | zero-successors]

Same output format as **show ip eigrp topology**, but it also shows some portion of the topology table.

show ip eigrp topology all-links

Same output format as **show ip eigrp topology**, but it also shows all links in the topology table, rather than just feasible successors.