Global Knowledge®

Expert Reference Series of White Papers

Enhanced Interior Gateway Routing Protocol (EIGRP)

Enhanced Interior Gateway Routing Protocol (EIGRP)

Dheeraj (Raj) Tolani, Global Knowledge Instructor

Introduction

Every routing protocol has some pros and cons and its own metric considerations. In this white paper, we will discuss Enhanced Interior Gateway Routing Protocol (EIGRP), a Cisco-proprietary routing protocol.

EIGRP is an advanced distance-vector classless routing protocol. By using the word advanced, we are basically saying that it is better than the other distance vector protocols, like RIP (Routing Information Protocol) for example. By classless, we mean that it will advertise the subnet mask information in its updates.

Before going to EIGRP, Cisco had another protocol called IGRP (Interior Gateway Routing Protocol). Coming from an older protocol, EIGRP still has some of the IGRP limitations, like automatic summarization at network boundary. This limitation can be easily fixed, however, with **no auto-summary** command under router configuration mode. These limitations and fixes are reviewed in detail in Global Knoweldge's BSCI class.

Features of EIGRP

- Since it is advertising its subnet mask information, you could use different subnet masks at different places in your network, based on your host requirements or other business reasons. This means there is a VLSM (Variable Length Subnet Mask) support in EIGRP.
- EIGRP allows you implement unequal load sharing alongside equal load sharing, if you have parallel links between sites.
- EIGRP allows you to specify the different times of the day when different passwords will be effective. (I hope your time servers are keeping things in sync so you can take advantage of this great feature.)
- EIGRP uses IP protocol # 88, so if you are the firewall person or the ACL person, you want to ensure that the protocol number 88 is allowed. Remember, EIGRP rides on IP, not on TCP or UDP like some other protocols do (RIP, for example, rides on UDP port # 520).
- EIGRP allows you to use MD5 authentication between peers. This could be very powerful if you are afraid that someone could be sniffing the network and stealing your routes. With EIGRP authentication, you can validate each update and ensure that you are talking to the correct EIGRP neighbor since they share the same password.

Autonomous Systems and Neighbors

In the early days, when running RIP, one of the issues was the lack of ability to group RIP routers under one administration unit. IGRP gave us the ability to have something called an Autonomous System Number. EIGRP built on that, and now we can group routers under a common administration unit called the Autonomous System.

You can assign any number in the range 1-65,535 to the routers. In order for them to incorporate routes they hear from other EIGRP-speaking routers, the Autonomous System number (AS#) must match. If the AS# doesn't match, then the routes will be discarded.

In order to start speaking EIGRP with other routers, we have to ensure that we are set up as neighbors. In addition to having the same AS#, we must also be on the same L3 network. By default, EIGRP uses the primary interface IP address as the source IP address, so it must be on the same broadcast domain (L3) as the other routers it is trying to talk to.

We said earlier that you can pick any AS number in the range of 1-65,535, but some of my clients use the same number that they use with their external ISP connections running BGP. EIGRP doesn't require this but, if it makes sense to you, then why not.

For metric calculations, EIGRP metric is a 32-bit number. These are big, ugly numbers that are based on a Cisco-proprietary Diffusing Update Algorithm (DUAL). By default, EIGRP uses "K-values" Bandwidth (BW) and Delay (DLY). If you want to, you can also factor in Reliability (REL) and Load, though these are not in use with "out of the box" EIGRP configurations. MTU is also exchanged but not used for metric calculations. So, in order for the EIGRP neighbor relationship to be formed exchanging routes, the K-values must match in the hello messages.

Configuration

In order to configure EIGRP you must be in the router configuration mode. Then, just as with Access Control Lists or OSPF network configurations, you can use the wildcard mask to specify the network/host addresses you want to start EIGRP. If you don't key in the wildcard mask entries, the system defaults to classful assumption.

For example:

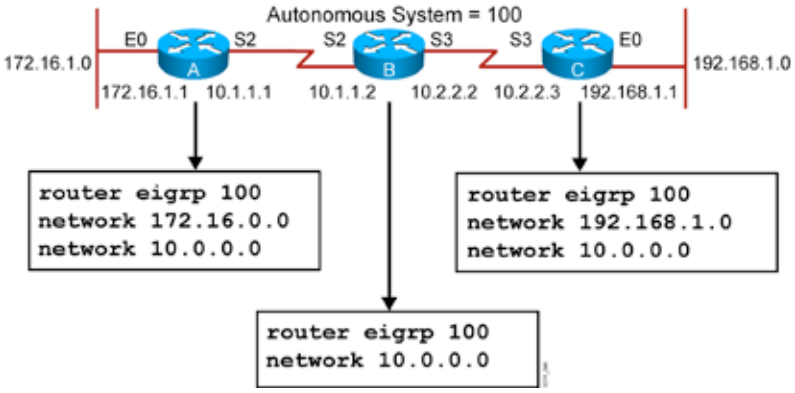

Figure 1. (Courtesy of Cisco Systems curriculum)

In the previous diagram, we have three routers (RouterA, RouterB, and RouterC), and all three routers will be running EIGRP as our routing protocol.

As discussed earlier, they must be in the same Autonomous System in order to run EIGRP. In Figure 1, we see they all have AS# 100, which means they will exchange routes with each other. The first command in Figure 1, **router eigrp 100**, is done in the global configuration mode.

```
RouterA>enable
RouterA#config t
RouterA(config)#router eigrp 100
RouterA(config-router)
```

After typing this command, you end up in router configuration mode. As you might have noticed, you can't tell which routing protocol you are configuring just by looking at the prompt, so you need know what command brought you here.

Now we wish to start EIGRP on Ethernet 0 (E0 in figure 1) and Serial 2 (S2 in figure 1) and advertise these networks in our updates. To do this, we need to key these network commands in the router configuration mode. Even though the E0 and S2 interfaces are subnets of 172.16.0.0 (subnet is 172.16.1.0) and of 10.0.0.0 (subnet 10.1.1.0), we are entering the classful addresses. This used to be the only way of configuring EIGRP in earlier versions of IOS. Today, you can use the wildcard masks and enter the specific matches as previously discussed.

Now, what if users who are using RouterA as their default gateway want to communicate with 192.168.1.0? Shouldn't you key in the network command for that particular class C network under RouterA?

The answer is no, because it has to be an attached network. 192.168.1.0 is not attached to RouterA, only RouterC. Therefore, RouterC has the network statement for 192.168.1.0 and will advertise it to RouterB, and then RouterB will advertise that to Router A. This way, the users for RouterA will be able to get there, but you are still only entering physically attached networks.

You can check to see if there are neighbors discovered after configuring the network commands. Here is a sample output from the **show ip eigrp neighbor** command. The number of neighbors and the interface IP addresses for those neighbors and the interfaces involved will vary, of course.

```
RouterX# show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address         Interface      Hold  Uptime    SRTT  RTO  Q   Seq
                                   (sec)           (ms)       Cnt Num
1   10.23.23.2      Se0        13    00:02:26   29  2280  0  15
0   10.140.1.1      Se1        10    00:28:26   24  2280  0   25
```

Another way to verify your configuration and make sure everything worked out well is to check to see that are you getting any routes from the other routers. In the output of **show ip route,** EIGRP routes show up with the legend **D**.

Also, if your router is running multiple routing protocols, then there is a possibility that the same route will be learned from different protocols. For instance, if we run RIP, EIGRP, and OSPF on a router, we learn the same network from all three protocols. The system then has to decide which one to use, since all three protocols could possibly be sending different ways to get to the same place.

Cisco routers use something called "Administrative distance" to pick one learned route over the other same route from another routing protocol. This rule says that the lower the administrative distance number, the more believable the method. The routing protocol with the lowest number will be the winner and will be installed in the routing table.

EIGRP uses 90 for "Internal EIGRP routes" that are also part of the EIGRP network somewhere in our Autonomous System, or it might originally have been an OSPF route that was translated in the EIGRP network by redistribution. Since we might only be running EIGRP, the other router has to indicate to us that the route it is sending us is not really an EIGRP route but a route translated into EIGRP. This is known as an "External EIGRP route" and, by default, External EIGRP routes have the administrative distance of 170.

You have the option to change the administrative distance numbers to some other number you choose based on network policies. You can see the administrative distance that is assigned to the each network route in the **show ip route** command output. So, let's look at the output of the **show ip route**.

```
RouterX#show ip route
<output omitted>
Gateway of last resort is not set
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:02:22,
Serial0
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:31:31, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.96/27 is directly connected, Serial0/0/1
D       192.168.1.0/24 is a summary, 00:31:31, Null0
```

The point here is to make sure you can see that routes are learned with legend of **D** on the extreme left side of the output and also that the administrative distance is 90 for these EIGRP route (as seen in the table above in the bracket – first number before the "/" - [90/40514560] 90 is the administrative distance and the next big number 40514560 after the slash is the composite metric).

Summary

I hope this gives you basic information on configuring Cisco proprietary routing protocol EIGRP. For more information on EIGRP and to configure other features of EIGRP, such as the unequal load sharing we mentioned earlier in this white paper, you can take the BSCI class at Global Knowledge. Thanks for reading, and I can't wait to see you in the CCNA/ICND2 and BSCI class to explore EIGRP in greater detail.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

Understanding Networking Fundamentals

TCP/IP Networking

ICND1 – Interconnecting Cisco Network Devices 1

ICND2 – Interconnecting Cisco Network Devices 2

CCNA Boot Camp v2.0

BSCI – Building Scalable Cisco Internetworks v3.0

For more information or to register, visit **www.globalknowledge.com** or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs and exercises offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 1,200 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and business training needs.

About the Author

Dheeraj (Raj) Tolani has been working with Global Knowledge as a contract instructor teaching networking courses including CCNP track. He has been in the industry for over 18 years working with various technologies, including Cisco, Banyan Vines, Microsoft, and Novell. Dheeraj has worked as a consultant for various medical, financial, legal, government, and publishing companies. He runs a consulting company based in NYC, providing IP integration solutions (www.rajtolani.com).