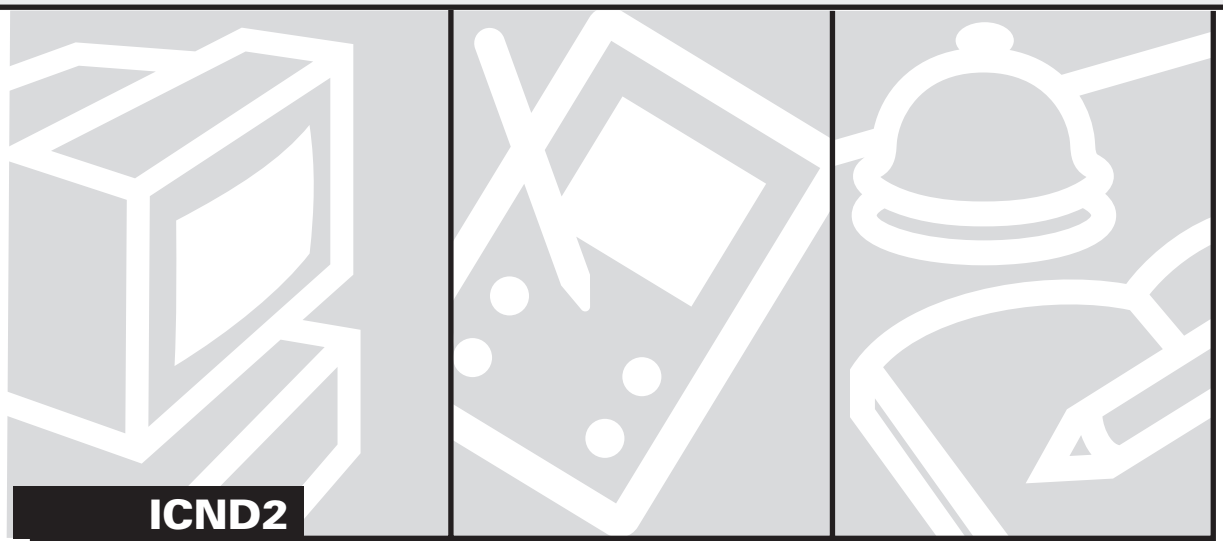

ICND2

Section 6
Implementing EIGRP

Enhanced Interior Gateway Routing Protocol (EIGRP) was introduced in Cisco IOS Release 9.21 as an enhancement to the limitations of IGRP. IGRP was developed by Cisco in the mid-1980s as a solution to Routing Information Protocol's (RIP's) limitations. RIP's limitation was based on its metric of hop count. In RIP, a hop larger than 15 was unreachable. As a result, RIP networks could contain no more than 15 hops. IGRP was developed to support larger networks than RIP. IGRP can support networks up to 255 hops in diameter. Additionally, IGRP could support multiple routed protocols. IGRP was created to support not only IP but also Internetwork Packet Exchange (IPX) and AppleTalk; RIP supports only IP.

Although IGRP provided many solutions to RIP, it still had its limitations. First, IGRP is a distance vector protocol. As such, IGRP was still susceptible to distance vector issues such as slow convergence; additionally, IGRP was classful.

Because of the limitation of IGRP, Cisco developed EIGRP. EIGRP is a classless protocol, meaning that it sends the subnet mask of its interfaces in routing updates. Although EIGRP is an enhancement of IGRP, it is not a distance vector routing protocol. EIGRP has some of the characteristics of a distance vector routing protocol and of a link-state protocol. As such, Cisco considers EIGRP a balanced hybrid protocol.

Like Open Shortest Path First (OSPF), EIGRP maintains awareness of the network through neighbor and topology tables, multicasts routing updates, and is fast to converge. However, EIGRP does not use the SPF algorithm to determine routes; instead, EIGRP uses the Diffusing Update Algorithm (DUAL).

Like IGRP, EIGRP is a Cisco-proprietary protocol and supports IP, IPX, and AppleTalk. EIGRP also uses the same metric as IGRP.

Cisco discontinued IGRP in Cisco IOS Release 12.2(13)T and 12.2(R1s4)S. As a result of IGRP's discontinuation, the ICND2 exam focuses on EIGRP and not IGRP. This section provides flash cards that cover EIGRP concepts, configuring EIGRP, and troubleshooting EIGRP.

Question 1

What are the four components of EIGRP?

Question 2

By default, what does EIGRP use for calculating routes?

Question 3

What are three general steps that EIGRP uses to add routes to the router's routing table?

Question 1 Answer

The four components of EIGRP are as follows:

- Protocol-independent modules
- Reliable Transport Protocol (RTP)
- Neighbor discovery/recovery
- Diffusing Update Algorithm (DUAL)

Question 2 Answer

Bandwidth and delay.

By default, bandwidth and delay are used by EIGRP to calculate its metric. EIGRP can also be configured to use reliability, load, and maximum transmission unit (MTU). The metric of EIGRP is the metric of IGRP multiplied by 256 for improved granularity.

Question 3 Answer

The three general steps that EIGRP uses to add routes to the router's routing table are as follows:

Step 1 Discover other EIGRP routers attached to the same subnet and form a neighbor relationship with the discovered routers. All discovered routers are kept in the router's EIGRP neighbor table.

Step 2 Exchange network topology information with all discovered neighbors. This information is stored in the EIGRP topology table.

Step 3 Run DUAL on all topology information and put the lowest-metric routes in the routing table.

Question 4

How does EIGRP discover neighbors?

Question 5

When routing information changes in the routing table, how does EIGRP send updates?

Question 6

What is the EIGRP neighbor table?

Question 4 Answer

EIGRP neighbors are discovered through Hello messages. On most networks, Hello messages are multicast every 5 seconds to address 224.0.0.10. On Frame Relay and link speeds of T1 (1.544 Mbps) or slower, Hellos are unicast every 60 seconds.

Question 5 Answer

When routing information changes, EIGRP sends update messages to all neighbors, informing them of the change. If EIGRP has to send to multiple neighbors on the same subnet, the update messages are multicast to IP address 224.0.0.10. If sending updates to one router, the messages are unicast to the neighbor.

NOTE All update messages are sent using RTP.

Question 6 Answer

The EIGRP neighbor table lists all adjacent routers. Each EIGRP router maintains a neighbor table.

Question 7

What is the EIGRP topology table?

Question 8

In EIGRP, what is a successor?

Question 9

In EIGRP, what is the feasible successor?

Question 7 Answer

The EIGRP topology table contains all learned routes to a destination. In other words, the topology table holds all feasible routes in its table.

NOTE EIGRP maintains a topology table for each network protocol configured. For example, if the router is configured for IP and IPX, EIGRP would maintain a topology table for IP and IPX.

Question 8 Answer

A successor is a route selected as the primary route used to reach a destination. It is the route kept in the routing table.

Question 9 Answer

The feasible successor is the backup route. These routes are selected at the same time the successors are identified, but they are only kept in the topology table, not the routing table. They are used for fast convergence. If the successor fails, the router can immediately route through the feasible successor. Multiple feasible successors can exist for a destination.

Question 10

In EIGRP, what is the advertised distance (AD)?

Question 11

In EIGRP, what is the feasible distance (FD)?

Question 12

What IOS commands enable EIGRP on a Cisco router and advertise 192.168.3.0 and 192.168.4.0 as its directly connected networks?

Question 10 Answer

The AD is the cost between the next-hop router and the destination.

Question 11 Answer

The FD is the metric from the local router, through the next-hop router, and to the destination.

Question 12 Answer

The **router eigrp** *process-id* command, followed by the **network** command, enables EIGRP on the router. The following commands enable EIGRP using AS 100 and then advertise networks 192.168.3.0 and 192.168.4.0:

```
RouterA(config)#router eigrp 100   (100 is the AS)
RouterA(config-router)#network 192.168.3.0
RouterA(config-router)#network 192.168.4.0
```

NOTE All routers configured for EIGRP must be configured with the same autonomous system (AS) to share information. The AS is a number from 1 to 65,535.

Question 13

What command would you use to see EIGRP adjacencies?

Question 14

What IOS command allows you to view all EIGRP routes in the routing table?

Question 15

How do you view the EIGRP neighbor table?

Question 13 Answer

The **show ip eigrp neighbors** command displays EIGRP adjacencies and directly connected neighbors, as follows:

```
RouterA# show ip eigrp neighbors

IP-EIGRP Neighbors for process 100
Address             Interface    Holdtime Uptime   Q     Seq  SRTT  RTO
                                 (secs)   (h:m:s)  Count  Num  (ms)  (ms)
192.168.10.2          Ethernet1    13       0:02:00  0     11   4     20
192.168.11.2          Ethernet0    14       0:02:01  0     10   12    24
```

Question 14 Answer

The **show ip route eigrp** command allows you to view all EIGRP-learned routes in the routing table.

Question 15 Answer

The **show ip eigrp neighbors** command shows the EIGRP neighbor table.

Question 16

How do you view the EIGRP topology table?

Question 17

What IOS command would you use to view the EIGRP neighbor states?

Question 18

Router A is running EIGRP and has four paths to network 192.168.100.0. All four paths have the same cost. Which path will Router A choose to route to network 192.168.100.0?

Question 16 Answer

The **show ip eigrp topology** command shows the EIGRP topology table, including successors and feasible successors, as follows:

```
RouterB# show ip eigrp topology

IP-EIGRP Topology Table for process 100

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

P 192.168.4.0 255.255.255.0, 1 successors, FD is 2172416
          via 192.168.3.2 (2172416/28160), Serial0
          via 192.168.2.2 (2684416/1794560), Serial1
```

Question 17 Answer

You would use the **debug eigrp neighbors** command to check the EIGRP neighbor states. This command displays the contents of the Hello packet used in EIGRP as well as the neighbors discovered by EIGRP.

Question 18 Answer

All four paths. By default, EIGRP can load-balance up to four equal-cost routes. This is called equal-cost load balancing. Because EIGRP has four equal-cost paths to network 192.168.100.0, all paths are included in Router A's routing table.

NOTE Remember, the default metric that EIGRP uses is bandwidth of the interface and delay.

Question 19

Why would you want to enable EIGRP route authentication on EIGRP routers?

Question 20

What is the default bandwidth of a serial interface on a Cisco router?

Question 21

What type of route authentication does EIGRP support?

Question 19 Answer

EIGRP route authentication causes EIGRP routers to authenticate with each other using an MD5 key digest. This prevents the introduction of unauthorized or false routing messages from unauthorized or unapproved routers.

Question 20 Answer

1544 kbps.

NOTE This is important to know because EIGRP uses bandwidth and delay as its default metric. Say that you have two point-to-point links configured between two routers. The speed on interface serial0 is 1.544 Mbps, and the speed on interface serial1 is 256 kbps. Because the delay on serial interfaces is the same, and the default bandwidth on serial interfaces is set to 1544, EIGRP will perform equal-cost load balancing between the two links. To avoid this, you need to manually set the bandwidth on the interface using the **bandwidth** *bandwidth-in-kbps* interface command.

Question 21 Answer

EIGRP supports message digest algorithm 5 (MD5) route authentication.

Question 22

As a network administrator, you have EIGRP enabled on all network routers. Your company has two locations. The corporate network is 172.16.1.0 255.255.255.0. The branch office network is 192.168.1.0 255.255.255.0. Your company acquires another company, and you connect the newly acquired company through a WAN link at the branch office. The network range of the newly acquired company is 172.16.2.0 255.255.255.0. You enable EIGRP on the router that connects to the new office, but users at the corporate network cannot access devices on the newly acquired company's network. Why?

Question 23

What command allows you to troubleshoot EIGRP authentication?

Question 22 Answer

Users from the corporate network cannot access the newly acquired company's network because the network is discontiguous, and by default, EIGRP summarizes routes across classful boundaries. As a result, the router on network 192.168.1.0 advertises 172.16.0.0 to both the corporate network and the newly acquired network. Thus when a user at the corporate network tries to connect to a device at the newly acquired network, the router drops the packet because it thinks it is local. Additionally, the branch office router thinks it has two equal-cost paths to the 172.16.0.0 network.

To fix the issue, you need to disable auto-summary on the routers. This is done with the **no auto-summary** EIGRP router configuration mode command.

Question 23 Answer

The **debug eigrp packets** command allows you to view the neighbor adjacency process. When authentication is enabled on two routers, it is part of the adjacency process, and you can view whether authentication is the cause of failed neighbor adjacencies.

If the failed neighbor adjacency is due to a misconfiguration in EIGRP authentication, you will see the following debugging output from the router:

```
*Mar 26 12:48:15.749: EIGRP: pkt key id = 2, authentication mismatch
*Mar 26 12:48:15.749: EIGRP: Serial 0: ignored packet from 192.168.1.2,
opc ode = 5 (invalid authentication)
*Mar 26 12:48:15.749: EIGRP: Dropping peer, invalid authentication19
```

Question 24

Router A is connected to Router B through a point-to-point T1 link. Router B is connected to network 192.168.100.0 on its Fast Ethernet interface. EIGRP is running on both routers. You install a second point-to-point link between the two routers for redundancy. The new link has a bandwidth of 256 kbps. Because the new link has a higher cost than the T1 link, the new link is not installed in the routing table and is idle. EIGRP only uses the T1 link to route to network 192.168.100.0. You want to load-balance between the two links. How do you enable EIGRP to load-balance between the two links?

Question 25

How do you enable EIGRP MD5 authentication on Cisco routers?

Question 24 Answer

By default, EIGRP can only load-balance equal-cost links and not load-balance between unequal-cost links. EIGRP needs to be configured to load-balance between unequal-cost links. The goal is to configure EIGRP to spread the traffic load inversely proportionally to the metrics on the two links. EIGRP uses the **variance** command to perform unequal-cost load balancing. The **variance** command defines a multiplier by which a metric can vary from the lowest-cost route. A variance of 1 means that the metrics of multiple routes must be equal.

In this question, the metric of the T1 link is 1,657,856. The composite metric to network 192.168.100.0 (the total of the cost of the T1 link and the Fast Ethernet interface) is 2,172,416. The composite metric of the 256-kbps link is 10,514,432. To find the variance between the two paths to perform unequal-cost load balancing, divide the metric of the 256-kbps link by the T1 link: 10,514,432/2,172,416 = 4.8. Thus to configure unequal-cost load balancing, the variance on Router A needs to be set to 5, as follows:

```
RouterA(config)#router eigrp 100
RouterA(config-router)#variance 5
```

NOTE The variance must be specified in whole numbers.

Question 25 Answer

The steps to configure EIGRP authentication are as follows:

Step 1 Enter the interface that you want to configure authentication on.

Step 2 Enable MD5 authentication using the **ip authentication mode eigrp** *process-id* **md5** interface command.

Step 3 Create an authentication key using the **ip authentication key-chain eigrp** *process-id key-chain* command. The *key-chain* parameter is the name of the key you want to create.

Step 4 Exit interface configuration mode.

Step 5 Identify the key chain that you configured in Step 3 using the **key chain** *name-of-key-chain* command.

Step 6 Create a key number with the **key** *number* command.

Step 7 Identify the key string using the **key-string** *text* command.