

EIGRP Setup for IP Routing

Objective

In this exercise you will configure EIGRP routing by enabling EIGRP on all WAN links (Frame Relay sub-interfaces) and all loopback interfaces within your network.

After having configured basic EIGRP routing, you will examine the established EIGRP adjacencies and observe the selection of a successor and a feasible successor for a given Site-1 route on the Site-2 router (see Figure 1). The route via the selected successor is then placed in the IP routing table.

Figure 1 shows the initially selected successor and feasible successor for the Site-1 loopback on the Site-2 router:

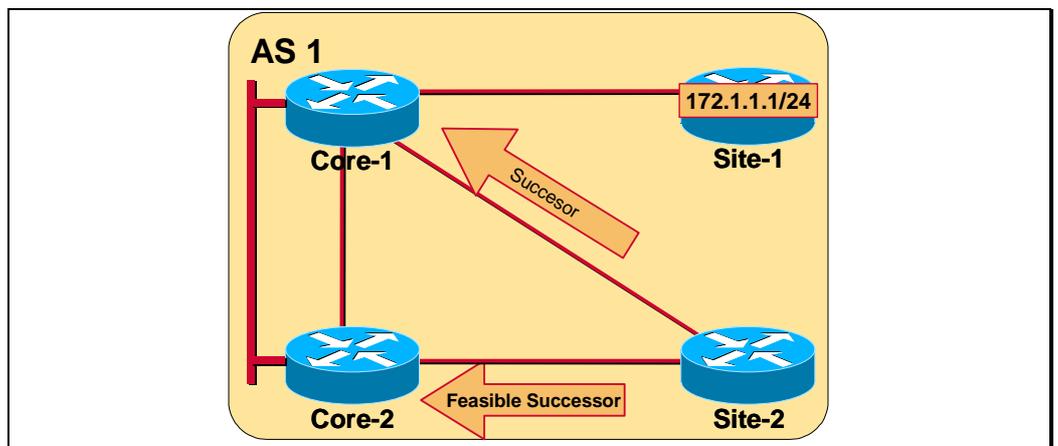


Figure 1: Initial EIGRP successor and feasible successor selection

Since the default bandwidth values do not reflect the actual state on WAN links, you will be asked to reconfigure the bandwidths on all WAN links, which will

result in a re-computation of the successor and the feasible successor. In this way you will optimize EIGRP routing decisions for the inter-site traffic.

Note The default bandwidth value on serial WAN interfaces is 1544 Kbits, on Ethernet interfaces 10000 Kbits, and on FastEthernet interfaces 100000 Kbits. The default delay for serial WAN interfaces is 20000 μ s, for Ethernet interfaces 1000 μ s and for FastEthernet interfaces 100 μ s.

Figure 2 shows the re-election of the successor and the feasible successor after changing the bandwidth on WAN links as to reflect the actual values.

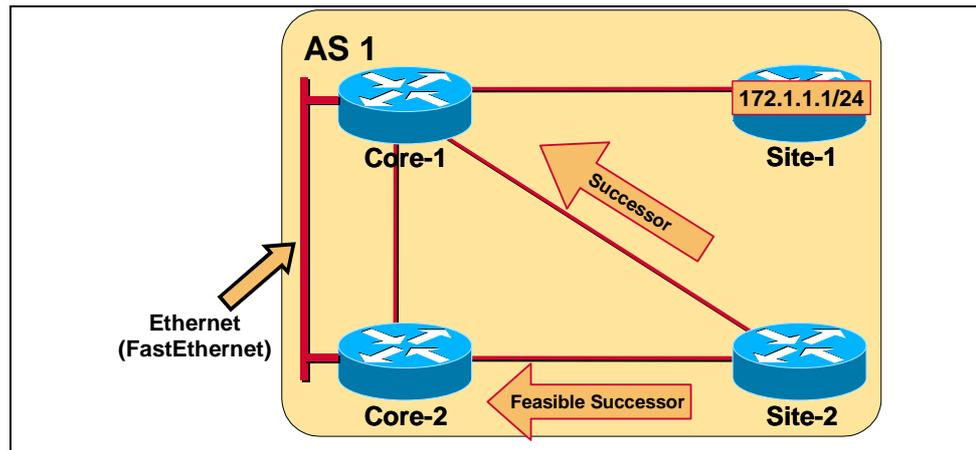


Figure 2: Optimized EIGRP routing decisions

Note Core routers in different lab workgroups may support only Ethernet or only FastEthernet interfaces. As all printouts in the lab were taken on routers with Ethernet interfaces, the default bandwidth and the default delay on FastEthernet interfaces were changed to Ethernet interface default values for compatibility reasons.

After optimizing EIGRP routing decisions, you will take advantage of the EIGRP unequal path load sharing feature, making the Core Serial link collaborate with the Core Ethernet link (the left-most link in the figure) in traffic load sharing.

In a separate, optional exercise, you will improve EIGRP convergence time by changing the values of hello and hold-down timers.

Note Please refer to the *Topology* section for more details on physical connectivity and to the *Addressing* section for detailed IP addressing information.

Command List

Use the following commands to complete this exercise:

Command	Description
(no) auto-summary	Enables (disables) automatic summarization of classless subnets
bandwidth x	Specifies the bandwidth on interfaces (used for routing processes)
clear ip route *	Clears the IP routing table
interface <type> <slot/port>	Enters the interface configuration mode
network x.x.x.x	Enables EIGRP on interfaces belonging to a specified network
Ping	Extended ping
router eigrp <as-number>	Starts an EIGRP routing process with the given AS number
show ip eigrp neighbors	Lists EIGRP neighbors and relevant information on EIGRP neighbor adjacencies
show ip eigrp topology [prefix] [mask]	Displays EIGRP topology table (whole, or only per prefix/mask)
show ip route	Displays the IP routing table
(no) shutdown	Administratively disables (enables) the interface
variance x	Enables unequal path load sharing in the EIGRP routing process

Table 1: Configuration and monitoring commands used in the EIGRP setup for IP routing exercise

Detailed Instructions

All the steps needed to complete the exercise are listed in the following tasks.

Task 1: Enabling EIGRP routing

- Step 1** Start the EIGRP routing process with the autonomous system number 1 on all the routers' active interfaces (including loopback interfaces).

Note Make sure you turn off the automatic summarization in the EIGRP process, to allow the propagation of discontinuous subnets!

Verification

- Step 2** Verify the establishment of EIGRP neighbor relations by using the **show ip eigrp neighbors** command on all routers. Output similar to the one shown below should display.

```
Core-1#sh ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address                Interface    Hold Uptime   SRIT   RTO   Q   Seq Type
      (sec)                (ms)        Cnt  Num
3   192.168.1.6             Se0/0.2     12 00:09:20   44   264  0   6
2   192.168.2.2             Et0/0       11 00:09:54    3   200  0  12
```

1	192.168.1.14	Se0/0.3	11 00:09:57	45	270	0	13
0	192.168.1.2	Se0/0.1	13 00:10:39	62	372	0	8

Printout 1: List of established EIGRP adjacencies

Note The succession of listed EIGRP neighbors depends on the time of adjacency establishment. The most recent neighbor is on top (see "Hold Uptime" column).

Step 3 Check the EIGRP topology table on the Site-2 router for the presence of the 172.1.1.0/24 route (loopback interface on Site-1), and determine the successor, feasible distance, and the feasible successor for this route. Output similar to the one shown below should display.

Note The metric can vary on account of changing lab equipment. The change occurs because some serial interface types are synchronous while others are asynchronous, which causes a change in default configuration.

```

Site-2#sh ip eigrp topology 172.1.1.0 255.255.255.0
IP-EIGRP topology entry for 172.1.1.0 255.255.255.0
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2809856
  Routing Descriptor Blocks:
    192.168.1.5 (Serial0/0.2), from 192.168.1.5, Send flag is 0x0
      Composite metric is (2809856/2297856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 45000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
    192.168.1.9 (Serial0/0.1), from 192.168.1.9, Send flag is 0x0
      Composite metric is (2835456/2323456), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 46000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 3

```

Printout 2: EIGRP topology table contents for the monitored route

Observe there are two possible paths to the 172.1.1.0/24 route, and the 192.168.1.5 router is the current successor. Its feasible distance (FD) is 2809856.

Also, verify the placement of the selected successor in the IP routing table!

```

Site-2#sh ip route 172.1.1.0
Routing entry for 172.1.1.0 255.255.255.0
  Known via "eigrp 1", distance 90, metric 2809856, type internal
  Redistributing via eigrp 1
  Last update from 192.168.1.5 on Serial0/0.2, 00:03:50 ago
  Routing Descriptor Blocks:
    * 192.168.1.5, from 192.168.1.5, 00:03:50 ago, via Serial0/0.2
      Route metric is 2809856, traffic share count is 1
      Total delay is 45000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

```

Printout 3: IP routing table entry for the monitored route

Task 2: Optimizing EIGRP routing decisions

Since the default serial interface bandwidths do not reflect the actual values (frame-relay CIR), resulting in sub-optimal routing, your task is to change bandwidths on routers' **serial** interfaces according to the values in the following table:

Link	Bandwidth
Core-1 to Core-2	2000 Kb
Core-1 to Site-1	256 Kb
Core-2 to Site-2	256 Kb
Core-1 to Site-2	64 Kb

Table 2: Bandwidth values for WAN links

These changes will initiate EIGRP DUAL route re-computation, resulting also in the swapping of the successor and feasible successor for the 172.1.1.0/24 route on Site-2. In that way, Core Ethernet link will be used for inter-site traffic, and Core-1 to Site-2 link for backup purposes.

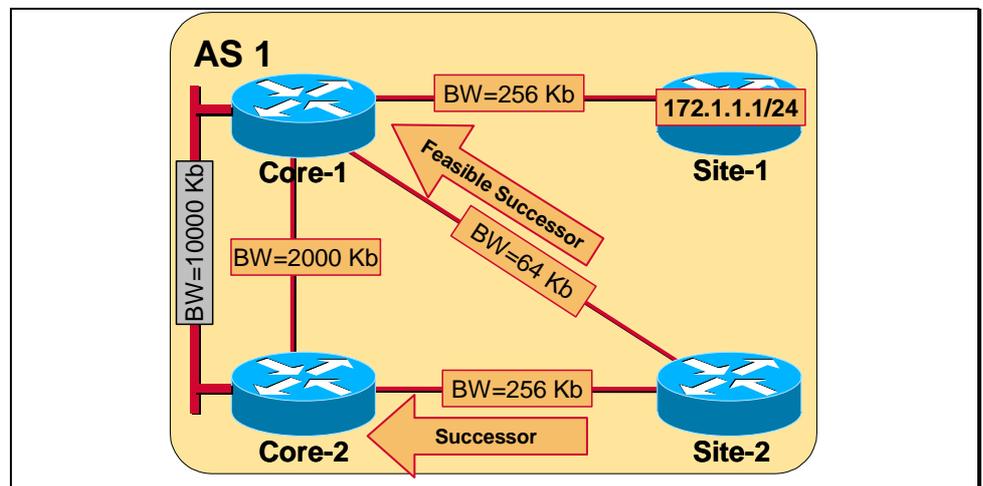


Figure 3: Changes in bandwidth influence routing decisions

- Step 4** Change the bandwidth value on the Core-1 to Core-2 serial link (do not change the bandwidth on the Core Ethernet link).
- Step 5** Change bandwidth values on all the remaining serial links.

Verification

- Step 6** On the Site-2 router, verify the selection of the Core-2 router for successor and Core-1 router for feasible successor for the monitored route (172.1.1.0/24). Since the Core links should also be used for return traffic (Site-1 to Site-2), also verify the selection of Core-1 as successor and Site-2 as feasible successor for the Site-2 loopback interface on the Core-1 router. Output similar to the one shown below should display on Site-2 and Core-1 routers respectively.

```
Site-2#sh ip eigrp topology 172.1.1.0 255.255.255.0
IP-EIGRP topology entry for 172.1.1.0 255.255.255.0
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 11177472
```

```

Routing Descriptor Blocks:
 192.168.1.9 (Serial0/0.1), from 192.168.1.9, Send flag is 0x0
  Composite metric is (11177472/10665472), Route is Internal
  Vector metric:
    Minimum bandwidth is 256 Kbit
    Total delay is 46000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 3
 192.168.1.5 (Serial0/0.2), from 192.168.1.5, Send flag is 0x0
  Composite metric is (41152000/10639872), Route is Internal
  Vector metric:
    Minimum bandwidth is 64 Kbit
    Total delay is 45000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2

```

Printout 4: EIGRP topology table contents for the monitored route, after bandwidth change

Comparing the output with the output from Step 3, you should note the change in the successor. The vector metrics now reflect the changed bandwidths as well.

```

Core-1#sh ip eigrp topology 172.1.3.0 255.255.255.0
IP-EIGRP topology entry for 172.1.3.0 255.255.255.0
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 10665472
  Routing Descriptor Blocks:
    192.168.2.2 (Ethernet0/0), from 192.168.2.2, Send flag is 0x0
    Composite metric is (10665472/10639872), Route is Internal
    Vector metric:
      Minimum bandwidth is 256 Kbit
      Total delay is 26000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 2
  192.168.1.14 (Serial0/0.3), from 192.168.1.14, Send flag is 0x0
  Composite metric is (11151872/10639872), Route is Internal
  Vector metric:
    Minimum bandwidth is 256 Kbit
    Total delay is 45000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
  192.168.1.6 (Serial0/0.2), from 192.168.1.6, Send flag is 0x0
  Composite metric is (40640000/128256), Route is Internal
  Vector metric:
    Minimum bandwidth is 64 Kbit
    Total delay is 25000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1

```

Printout 5: EIGRP topology table contents for the Site-2 loopback interface on the Core-1 router, after bandwidth change

In this way, bi-directional inter-site traffic flow over the Core Ethernet link is enabled, using Site-2 to Core-1 link only for backup purposes.

- Step 7** Start a continuous ping to the Site-2 loopback interface from the Site-1 loopback interface, and perform the shutdown of the Site-2 to Core-2 link on the Site-2 router during pinging. Observe the time needed to switch over from the previous successor to the new one. Since feasible successors for source and destination are

available on Site-2 and Core-1, the switchover should be smooth and relatively quick.

Task 3: Enabling EIGRP unequal paths load sharing

- Step 8** Re-enable the Site-2 to Core-2 link.
- Step 9** Enable EIGRP unequal paths load sharing on Core links. Calculate the minimum needed variance for the effective load sharing of the inter-site traffic (loopback to loopback) by examining the EIGRP topology tables on Core routers.

Note After configuring variance, clear IP routing tables on both Core routers to effectively start the load sharing process!

Verification

- Step 10** Examine the contents of IP routing tables on both Core routers and notice the multiple paths to Site routers (in bold font), resulting in load sharing between the Core Ethernet and Serial link in both directions.

```
Core-1#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       NL - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

172.1.0.0 255.255.255.0 is subnetted, 4 subnets
D       172.1.4.0 [90/409600] via 192.168.2.2, 00:00:03, Ethernet0/0
D       172.1.1.0 [90/10639872] via 192.168.1.2, 00:00:03, Serial0/0.1
D       172.1.3.0 [90/10665472] via 192.168.2.2, 00:00:03, Ethernet0/0
       [90/11151872] via 192.168.1.14, 00:00:03, Serial0/0.3
C       172.1.2.0 is directly connected, Loopback0
192.168.1.0 255.255.255.252 is subnetted, 4 subnets
D       192.168.1.8 [90/10537472] via 192.168.2.2, 00:00:03, Ethernet0/0
       [90/11023872] via 192.168.1.14, 00:00:05, Serial0/0.3
C       192.168.1.12 is directly connected, Serial0/0.3
C       192.168.1.0 is directly connected, Serial0/0.1
C       192.168.1.4 is directly connected, Serial0/0.2
C       192.168.2.0 255.255.255.0 is directly connected, Ethernet0/0
```

Printout 6: Multiple paths for inter-site traffic in the ip routing table of Core-1

```
Core-2#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       NL - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

172.1.0.0 255.255.255.0 is subnetted, 4 subnets
C       172.1.4.0 is directly connected, Loopback0
```

```

D    172.1.1.0 [90/10665472] via 192.168.2.1, 00:17:15, Ethernet0/0
      [90/11151872] via 192.168.1.13, 00:17:15, Serial0/0.2
D    172.1.3.0 [90/10639872] via 192.168.1.10, 00:17:15, Serial0/0.1
D    172.1.2.0 [90/409600] via 192.168.2.1, 00:17:15, Ethernet0/0
      [90/1920000] via 192.168.1.13, 00:17:15, Serial0/0.2
192.168.1.0 255.255.255.252 is subnetted, 4 subnets
C    192.168.1.8 is directly connected, Serial0/0.1
C    192.168.1.12 is directly connected, Serial0/0.2
D    192.168.1.0 [90/10537472] via 192.168.2.1, 00:17:16, Ethernet0/0
      [90/11023872] via 192.168.1.13, 00:17:16, Serial0/0.2
D    192.168.1.4 [90/40537600] via 192.168.2.1, 00:17:16, Ethernet0/0
      [90/41024000] via 192.168.1.10, 00:17:16, Serial0/0.1
      [90/41024000] via 192.168.1.13, 00:17:19, Serial0/0.2
C    192.168.2.0 255.255.255.0 is directly connected, Ethernet0/0

```

Printout 7: Multiple paths for inter-site traffic in the ip routing table of Core-2

Improving EIGRP Convergence Time

Objective

When you shut down the Site-2 to Core-2 link in one of the previous steps, the switchover to the backup route was quick, but not immediate. The reason was that the Core-1 router still believed its successor was Core-2, although Core-2 had no direct connection to Site-2 any more. Such situations, especially if they last for a long period, may potentially result in a temporary routing black hole. After Core-2 realized that the EIGRP adjacency with Site-2 was down (hold-down timer had expired), it sent a query denying all routes learned through Site-2. Core-1 then received this query, realized that Core-2 was no longer a successor, and therefore performed a switchover to the feasible successor (see figures 3 and 4).

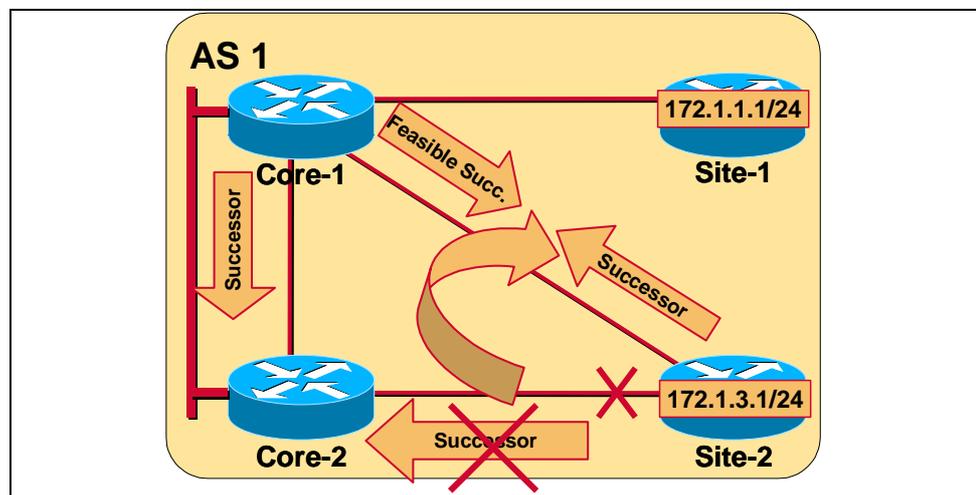


Figure 4: Creating a temporary routing black hole (Core-1 sending Site-2 traffic to Core-2)

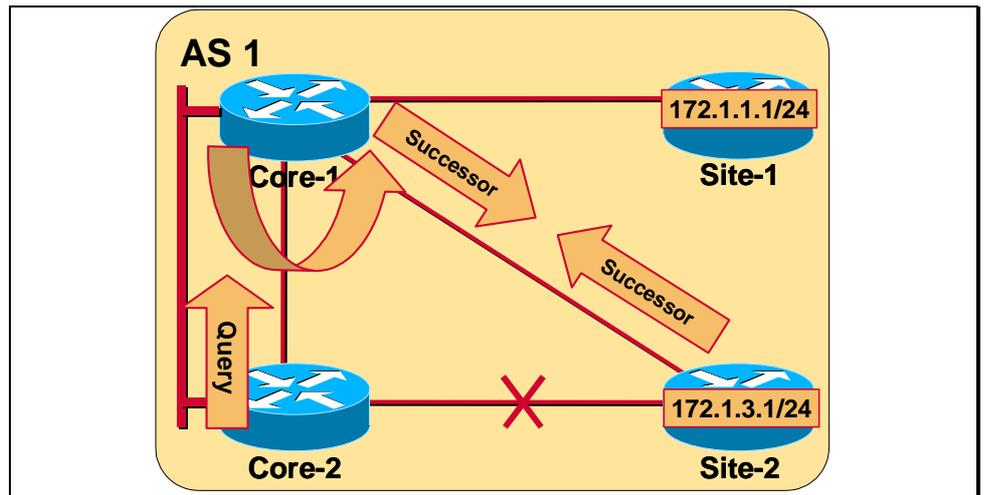


Figure 5: Core-2 originates a Query denying Site-2 routes and initiating the switchover on Core-1

Your objective is to change the default hello and hold-down timers on the Site-2 to Core-2 link to decrease the switchover time.

Command List

Command	Description
ip hello-interval eigrp <as-number> <seconds>	Sets the hello interval on an interface
ip hold-time eigrp <as-number> <seconds>	Sets the hold-time interval on an interface

Table 3: Commands for configuring hello and hold-down timers

If necessary use also commands from Table 1 .

Task 4: Changing hello and hold-down timers

- Step 11** Change the hello and hold-down timers on the Core-2 to Site-2 link to improve EIGRP convergence time. Change the value of the hello timer to 1 second, and the hold-down timer to 3 seconds.

Verification

- Step 12** Start a continuous ping to the Site-2 loopback interface from the Site-1 loopback interface, and perform the shutdown of the Site-2 to Core-2 link on the Site-2 router during pinging. Observe the improvement in EIGRP convergence as compared to the previous measurement.

Review Questions

- When is it necessary to turn off automatic summarization in the EIGRP process?
- When is a router selected as a feasible successor?

- Which two parameters influence EIGRP routing decisions by default? How is the EIGRP metric calculated?
- Which timers affect EIGRP convergence time?