



**Vendor: Cisco**

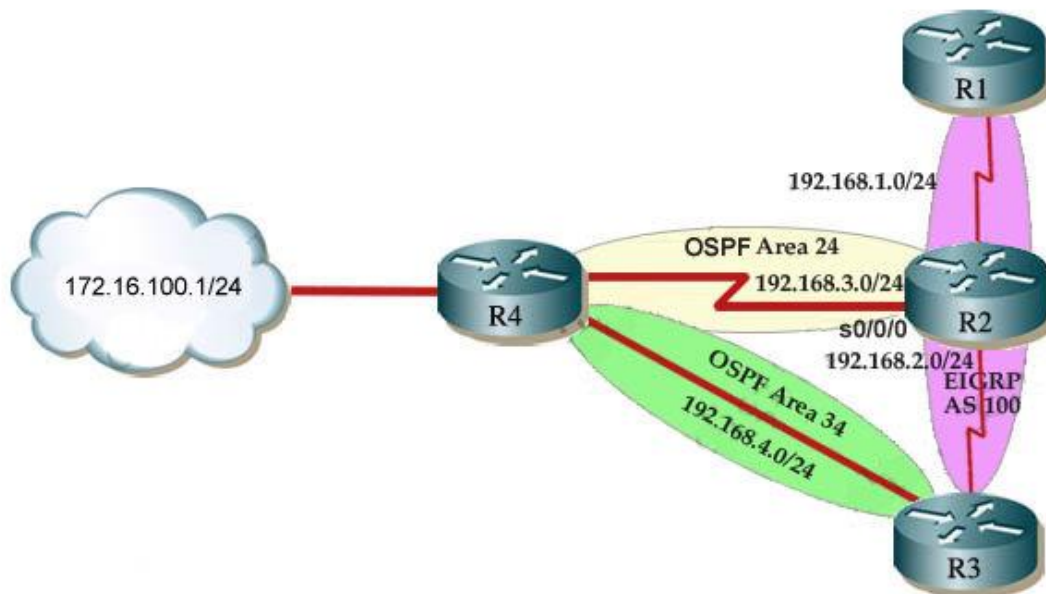
**Exam Code: 642-902**

**Exam Name: Implementing Cisco IP Routing (ROUTE)**

**Version: Demo**

**QUESTION 1**

In this question you need to redistribute between OSPF and EIGRP such that 172.16.100.1 is reachable from router R1.



R2 is an ASBR for EIGRP 100 and OSPF AREA 24

R3 is an ASBR for EIGRP 100 and OSPF AREA 34

Note: There are TWO separate areas on TWO separate ASBRS thus you need to do redistribution on R2 and R3

R1 is ONLY in EIGRP 100, and is THE ONLY router you can ping from. R4 has a loopback interface that must be pinged from R1.

R4 is running OSPF and has redundant links to EIGRP network over R3 router.

Note: You should ping from R1 to 172.16.100.1 network to make sure everything is working correctly.

**Correct Answer:**

First we need to find out 5 K-Values used for EIGRP (Bandwidth, Delay, Reliability, Load, MTU) of the s0/0/0 interface (the interface of R2 connected to R4) for redistribution:

```
R2#show interface s0/0/0
```

Write down these 5 parameters, notice that we have to divide the Delay by 10 because its metric unit is tens of microsecond. For example, we get Bandwidth=1544 Kbit, Delay=20000 us, Reliability=255, Load=1, MTU=1500 bytes then we would redistribute as follows:

```
R2#config terminal
```

```
R2(config)#router ospf 1
R2(config-router)# redistribute eigrp 100 metric-type 1 subnets
R2(config-router)#exit
R2(config-router)#router eigrp 100
R2(config-router)#redistribute ospf 1 metric 1544 2000 255 1 1500
```

(Notice: In fact, these parameters are just used for reference and we can use other parameters with no problem. Also, a candidate said that the simulator didn't accept the Bandwidth of 1544; in that case, we can use a lower value, like 128.)

If the delay is 20000us then we need to divide it by 10, that is  $20000 / 10 = 2000$ )

For R3 we use the show interface fa0/0 to get 5 parameters too

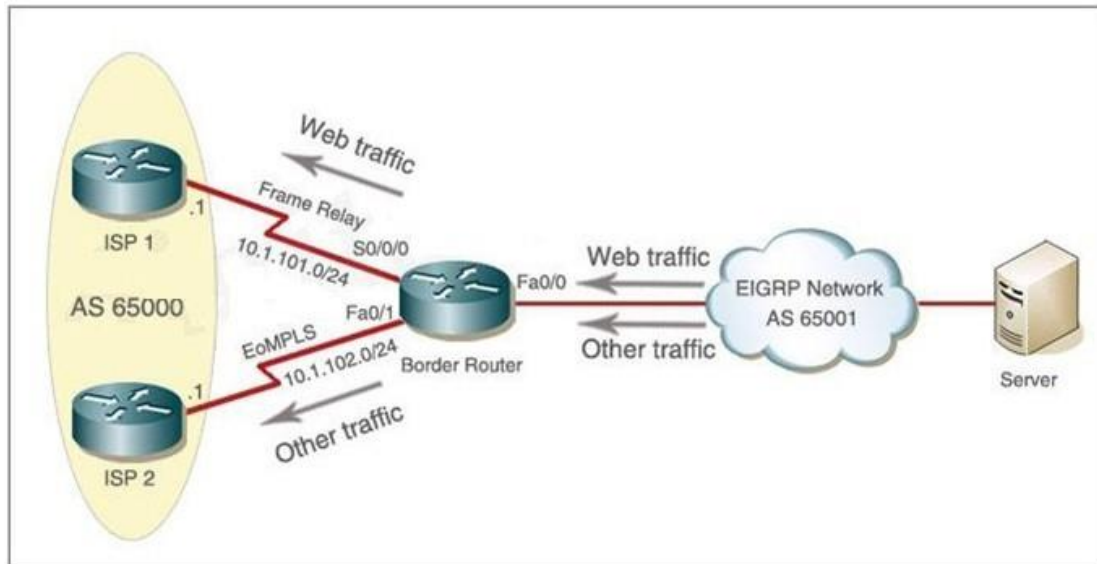
```
R3#show interface fa0/0
```

For example we get Bandwidth=10000 Kbit, Delay=1000 us, Reliability=255, Load=1, MTU=1500 bytes

```
R3#config terminal
R3(config)#router ospf 1
R3(config-router)#redistribute eigrp 100 metric-type 1 subnets
R3(config-router)#exit
R3(config-router)#router eigrp 100
R3(config-router)#redistribute ospf 1 metric 10000 100 255 1 1500
```

## QUESTION 2

Company Acan has two links which can take it to the Internet. The company policy demands that you use web traffic to be forwarded only to Frame Relay link and other traffic can go through any links that are available. Use BGP attributes to solve the requested action.



### Correct Answer:

1) Create an Access list that catches the HTTP traffic:

```
BorderRouter#access-list 101 permit tcp any any eq www
```

Note that the server was not directly connected to the Border Router. There were a lot of EIGRP routes on it. In the real exam you do not know the exact IP address of the server in the EIGRP network so we have to use the source as "any" to catch all the source addresses.

2) Route map that sets the next hop address to be ISP1 and permits the rest of the traffic:

```
BorderRouter(config)#route-map pbr permit 10
BorderRouter(config-route-map)#match ip address 101
BorderRouter(config-route-map)#set ip next-hop 10.1.101.1
BorderRouter(config-route-map)#exit
BorderRouter(config)#route-map pbr permit 20
```

(Notice: the route-map pbr permit 20 line allows other traffic than HTTP to be routed. Otherwise, other traffic will be dropped)

3) Apply the route-map on the interface to the server in the EIGRP Network:

```
BorderRouter(config-route-map)#exit
BorderRouter(config)#int fa0/0
```

```
BorderRouter(config-if)#ip policy route-map pbr
BorderRouter(config-if)#exit
BorderRouter(config)#exit
```

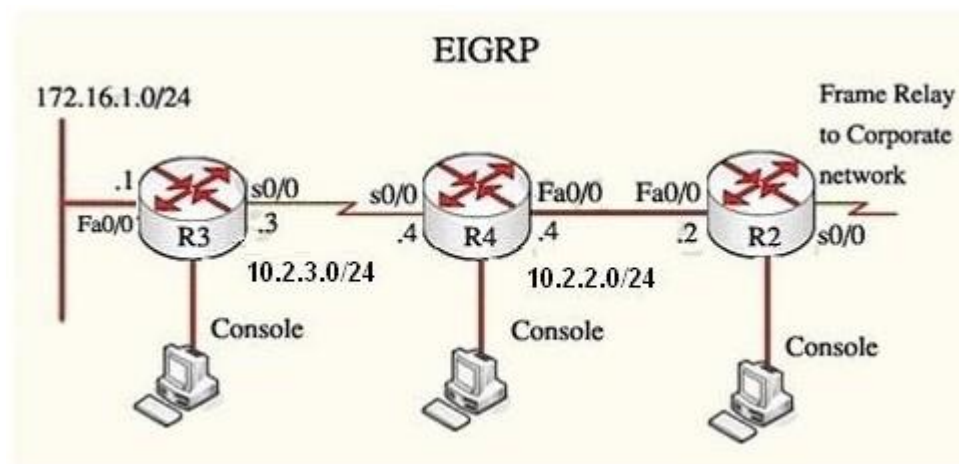
4) There is a "Host for Testing", click on this host to open a box in which there is a button named "Generate HTTP traffic". Click on this button to generate some packets for HTTP traffic. Jump back to the BorderRouter and type the command "show route-map".

```
BorderRouter#show route-map
```

In the output you will see the line "Policy routing matches: 9 packets...". It means that the route-map we configured is working properly.

### QUESTION 3

By increasing the first distant office, JS manufactures has extended their business. They configured the remote office router (R3) from which they can reach all Corporate subnets. In order to raise network stableness and lower the memory usage and CPU utilization to R3, JS manufactures makes use of route summarization together with the EIGRP Stub Routing feature. Another network engineer is responsible for the implementing of this solution. However, in the process of configuring EIGRP stub routing connectivity with the remote network devices off of R3 has been missing.



Presently JS has configured EIGRP on all routers in the network R2, R3, and R4. Your duty is to find and solve the connectivity failure problem with the remote office router R3. You should then configure route summarization only to the distant office router R3 to complete the task after the problem has been solved. The success of pings from R4 to the R3 LAN interface proves that the fault has been corrected and the R3.

IP routing table only contains two 10.0.0.0 subnets.

**Correct Answer:**

Use the show runningconfig command on router R3

```
R3# show running-config
<output omitted>
!
router eigrp 123
  network 10.0.0.0
  network 172.16.0.0
  no auto-summary
  eigrp stub receive-only
!
<output omitted>
```

Notice that R3 is configured as a stub receive-only router. The receive-only keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system. This keyword will also prevent any type of route from being sent. Therefore we will remove this command and replace it with the eigrp stub command:

```
R3#configure terminal
R3(config)#router eigrp 123
R3(config-router)#no eigrp stub receive-only
R3(config-router)#eigrp stub
R3(config-router)#end
```

Now R3 will send updates containing its connected and summary routes to other routers. Notice that the eigrp stub command equals to the eigrp stub connected summary because the connected and summary options are enabled by default. Next we will configure router R3 so that it has only 2 subnets of 10.0.0.0 network. Use the show ip route command on R3 to view its routing table R3#show ip route

```
Router3# show ip route
```

```
10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
```

- D 10.0.0.0/8 is a summary, 02:04:27, Null0
- D 10.2.2.0/24 [90/30720] via 10.2.3.4, 00:00:06, Serial0/0
- C 10.2.3.0/24 is directly connected, Serial0/0
- D 10.2.4.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
- D 10.2.5.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
- D 10.2.6.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
- D 10.2.7.0/24 [90/161280] via 10.2.3.4, 00:00:02, Serial0/0
- D 10.2.8.0/24 [90/161280] via 10.2.3.4, 00:00:02, Serial0/0
- D 10.2.9.0/24 [90/161280] via 10.2.3.4, 00:00:02, Serial0/0

```
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

- D 172.16.0.0/16 is a summary, 02:04:27, Null0
- C 172.16.1.0/24 is directly connected, FastEthernet0/0

Because we want the routing table of R3 only have 2 subnets so we have to summary sub-networks at the interface which is connected with R3, the s0/0 interface of R4. There is one interesting thing about the output of the show ip route shown above. The 10.2.3.0 /24, which is a directly connected network of R3. We can't get rid of it in the routing table no matter what technique we use to summarize the networks. Therefore, to make the routing table of R3 has only 2 subnets we have to summarize the other subnets into one subnet. In conclusion, we will use the ip summary-address eigrp 123 10.0.0.0 255.0.0.0 at the interface s0/0 of R4 to summary.

```
R4>enable
```

```
R4#configure terminal
```

```
R4(config)#interface s0/0
```

```
R4(config-if)#ip summary-address eigrp 123 10.0.0.0 255.0.0.0
```

```
Router3# show ip route
```

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
```

- D 10.0.0.0/8 is a summary, 00:06:22, Null0
  - C 10.2.3.0/24 is directly connected, Serial0/0
- ```
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```
- D 172.16.0.0/16 is a summary, 00:06:22, Null0
  - C 172.16.1.0/24 is directly connected, FastEthernet0/0

Now we jump back to R3 and use the show ip route command to verify the effect, the output is shown below: (But please notice that the ip addresses and the subnet masks in your real exam might be different so you might use different ones to solve this question) Just for your information, notice that if you use another network than 10.0.0.0/8 to summary, for example, if you use the command ip summary-address eigrp 123 10.2.0.0 255.255.0.0 you will leave a /16 network in the output of the show ip route command.

```
Router3# show ip route
```

```
10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
```

```
D    10.0.0.0/8 is a summary, 00:18:43, Null0
```

```
D    10.2.0.0/16 [90/30720] via 10.2.3.4, 00:00:06, Serial0/0
```

```
C    10.2.3.0/24 is directly connected, Serial0/0
```

```
    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
D    172.16.0.0/16 is a summary, 00:18:43, Null0
```

```
C    172.16.1.0/24 is directly connected, FastEthernet0/0
```

But in your real exam, if you don't see the line "10.0.0.0/8 is a summary ...Null0" then you can summary using the network 10.2.0.0/16. This summarization is better because all the pings can work well. Finally don't forget to use the copy running-config startup-config command on routers R3 and R4 to save the configurations.

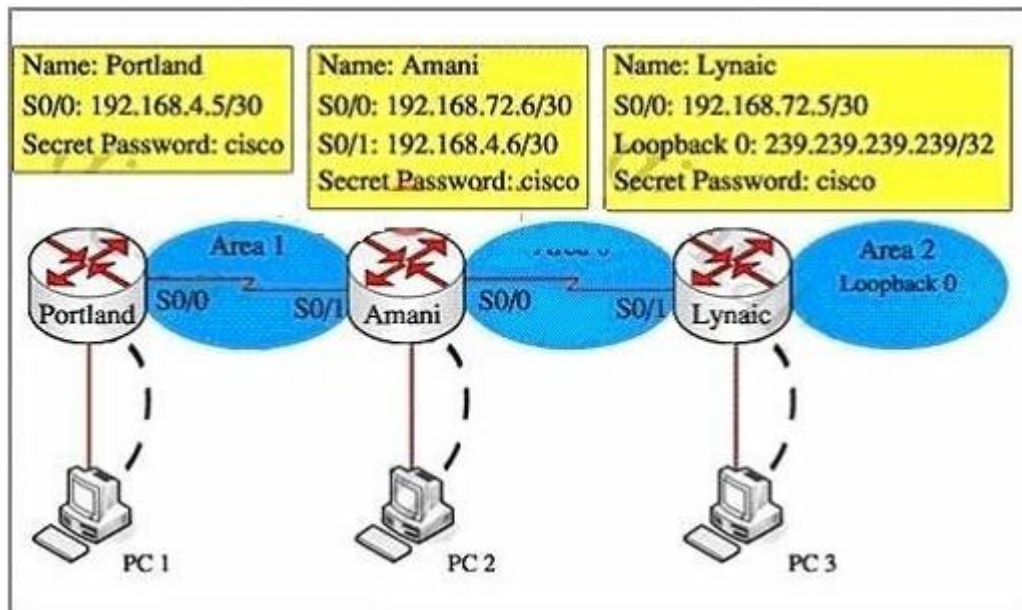
```
R4(config-if)#end
```

```
R4#copy running-config startup-config
```



**QUESTION 4**

OSPF is configured on routers Amani and Lynaic. Amani's S0/0 interface and Lynaic's S0/1 interface are in Area 0. Lynaic's Loopback0 interface is in Area 2.



Your task is to configure the following:

Portland's S0/0 interface in Area 1

Amani's S0/1 interface in Area 1

Use the appropriate mask such that ONLY Portland's S0/0 and Amnani's S0/1 could be in Area 1. Area 1 should not receive any external or inter-area routes (except the default route).

**Correct Answer:**

First, we configure Portland's S0/0 interface so that it belongs to Area 1. So, we have to find out which subnetwork the IP address 192.168.4.5/30 (the IP of interface S0/0 of Portland) belongs to. This address belongs to a subnetwork which has:

Increment: 4 (/30 = 255.255.255.252 or 1111 1111.1111 1111.1111 1111.1111 1100)

Network address: 192.168.4.4 (because 4 = 4 \* 1 and 4 < 5)

Broadcast address: 192.168.4.7 (because 7 = 4 + 4 - 1) (It is not necessary to find out the broadcast address but we should know it)

The question requires that only Portland's S0/0 and Amani's S0/1 could be in Area 1, therefore we must use a wildcard of 0.0.0.3 (this wildcard is equivalent with a subnet mask of /30) so that there are only 2 IP addresses can participate in area 1 (they are 192.168.4.5 & 192.168.4.6). The full command we use here is network 192.168.4.4 0.0.0.3 area 1 The question also requires that "Area 1 should not receive any external or inter-area routes (except the default route)". Recall that if we don't want the router to receive external routes, we have to stop LSA Type 5. And if we don't want to receive inter-area routes, we have to stop LSA Type 3 and Type 4. Therefore we have to configure area 1 as a totally stubby area. For your information, here is the definition of a

totally stubby area. "Totally stubb area - This area does not accept summary LSAs from other areas (types 3 or 4) or external summary LSAs (Type 5). Types 3,4 and 5 LSAs are replaced by the Area Border Router(ABR) with a default router. Totally stubby areas protect internal routers by minimizing the routing table and summarizing everything outside the area with a default route." (CCNP BSCI Official Exam Certification Guide, Fourth Edition) In conclusion, we have to configure area 1 as a totally stubby area. We do that by configuring Portland as stub and configuring Amani (ABR router) as a stub + "no-summary"suffix. + Configure Portland router as a stub:

```
Portland#configure terminal
Portland(config)#router ospf 1
```

Allow network 192.168.4.4/30 to join Area 1, notice that you have to convert subnet mask into wildcard mask:

```
Portland(config-router)#network 192.168.4.4 0.0.0.3 area 1
```

Configure Portland as a stub:

```
Portland(config-router)#area 1 stub
Portland(config-router)#end
Portland#copy running-config startup-config
```

+ Configure Amani router as a "totally stub":

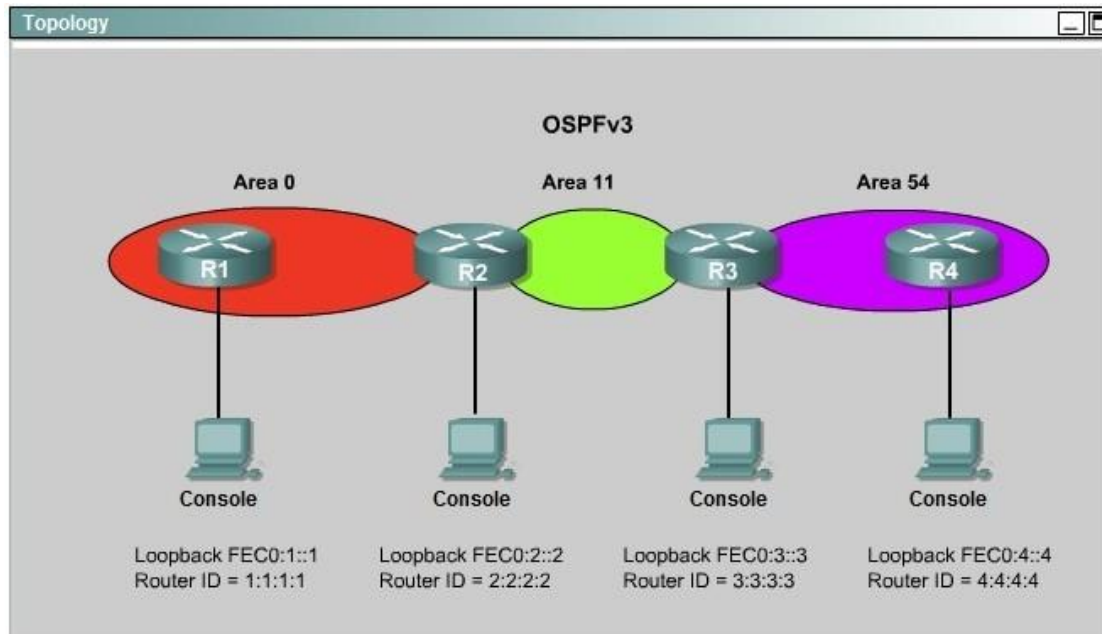
```
Amani#configure terminal
Amani(config)#router ospf 1
Amani(config-router)#network 192.168.4.4 0.0.0.3 area 1
```

Make area 1 become a totally stubby area, notice that we can only use this command on ABR router:

```
Amani(config-router)#area 1 stub no-summary
Amani(config-router)#end
Amani#copy running-config startup-config
```

### QUESTION 5

Acme is a small export company that has an existing enterprise network that is running IPv6 OSPFv3. Currently OSPF is configured on all routers. However, R4's loopback address (FEC0:4:4) cannot be seen in R1's IPv6 routing table. You are tasked with identifying the cause of this fault and implementing the needed corrective actions that uses OSPF features and does no change the current area assignments. You will know that you have corrected the fault when R4's loopback address (FEC0:4:4) can ping from R1 to R4 loopback address.



#### Correct Answer:

To troubleshoot the problem, first issue the show running-config on all of 4 routers. Pay more attention to the outputs of routers R2 and R3. The output of the "show running-config" command of R2:

```
<output omitted>
!
ipv6 router ospf 1
router-id 2.2.2.2
log-adjacency-changes
!
<output omitted>
```

The output of the "show running-config" command of R3:

```
<output omitted>
|
ipv6 router ospf 1
router-id 3.3.3.3
log-adjacency-changes
area 54 virtual-link 4.4.4.4
|
<output omitted>
```

We knew that all areas in an Open Shortest Path First (OSPF) autonomous system must be physically connected to the backbone area (Area 0). In some cases, where this is not possible, we can use a virtual link to connect to the backbone through a non-backbone area. The area through which you configure the virtual link is known as a transit area. In this case, the area 11 will become the transit area. Therefore, routers R2 and R3 must be configured with the area <area id> virtual-link <neighbor router-id>command. + Configure virtual link on R2 (from the first output above, we learned that the OSPF process ID of R2 is 1):

```
R2>enable
R2#configure terminal
R2(config)#ipv6 router ospf 1
R2(config-rtr)#area 11 virtual-link 3.3.3.3
```

Save the configuration:

```
R2(config-rtr)#end
R2#copy running-config startup-config
```

(Notice that we have to use neighbor router-id 3.3.3.3, not R2's router-id 2.2.2.2) + Configure virtual link on R3 (from the second output above, we learned that the OSPF process ID of R3 is 1 and we have to disable the wrong configuration of "area 54 virtual-link 4.4.4.4"):

```
R3>enable
R3#configure terminal
R3(config)#ipv6 router ospf 1
R3(config-rtr)#no area 54 virtual-link 4.4.4.4
R3(config-rtr)#area 11 virtual-link 2.2.2.2
```

Save the configuration:

```
R3(config-rtr)#end
R3#copy running-config startup-config
```

You should check the configuration of R4, too.

```
R4(config)#ipv6 router ospf 1
R4(config-router)#no area 54 virtual-link 3.3.3.3
R4(config-router)#end
```

After finishing the configuration doesn't forget to ping between R1 and R4 to make sure they work well! Note. If you want to check the routing information, use the show ipv6 route command, not "show ip route".

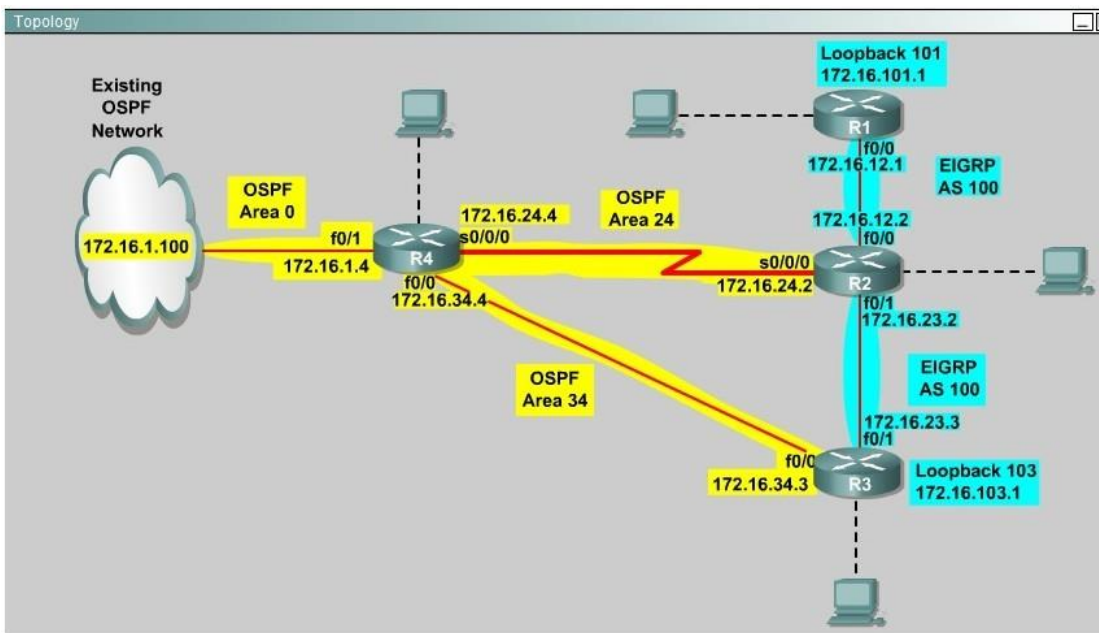
### QUESTION 6

**Scenario**

You are a network engineer with ROUTE.com, a small IT company. They have recently merged two organizations and now need to merge their networks as shown in the topology exhibit. One network is using OSPF as its IGP and the other is using EIGRP as its IGP. R4 has been added to the existing OSPF network to provide the interconnect between the OSPF and EIGRP networks. Two links have been added that will provide redundancy.

The network requirements state that you must be able to ping and telnet from loopback 101 on R1 to the OSPF domain test address of 172.16.1.100. All traffic must use the shortest path that provides the greatest bandwidth. The redundant paths from the OSPF network to the EIGRP network must be available in case of a link failure. No static or default routing is allowed in either network.

A previous network engineer has started the merger implementation and has successfully assigned and verified all IP addressing and basic IGP routing. You have been tasked with completing the implementation and ensuring that the network requirements are met. You may not remove or change any of the configuration commands currently on any of the routers. You may add new commands or change default values.




**Correct Answer:**

```

On router R2:
Config t
Router ospf 1
Redistribute eigrp 100 metric 100 subnets
Router eigrp 100
Redistribute ospf 1 metric 100000 10 200 10 1500
Then from R1 you should be able to ping 172.16.1.100
    
```

**QUESTION 7**

The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.



Examine the following excerpt from the "show ip ospf" command on D1:

**Area BACKBONE (0)**

Number of interfaces in the this area is 1  
Area has no authentication  
SPF algorithm last executed 00:00:31.280 ago  
SPF algorithm executed 5 times  
Area ranges are  
Number of LSA 13. Checksum Sum 0x16F0FD  
Number of opaque link LSA 0. Checksum Sum 0x000000  
Number of DCbitless LSA 0  
Number of indication LSA 0  
Number of DoNotAge LSA 0  
Flood list length 0

**Area 16**

Number of interfaces in this area is 2  
Area has message digest authentication  
SPF algorithm last executed 00:00:34.928 ago  
SPF algorithm executed 7 times  
Area ranges are  
Number of LSA 5. Checksum Sum 0x02FCD3  
Number of opaque link LSA 0. Checksum Sum 0x000000  
Number of DCbitless LSA 0  
Number of indication LSA 0  
Number of DoNotAge LSA 0  
Flood list length 0

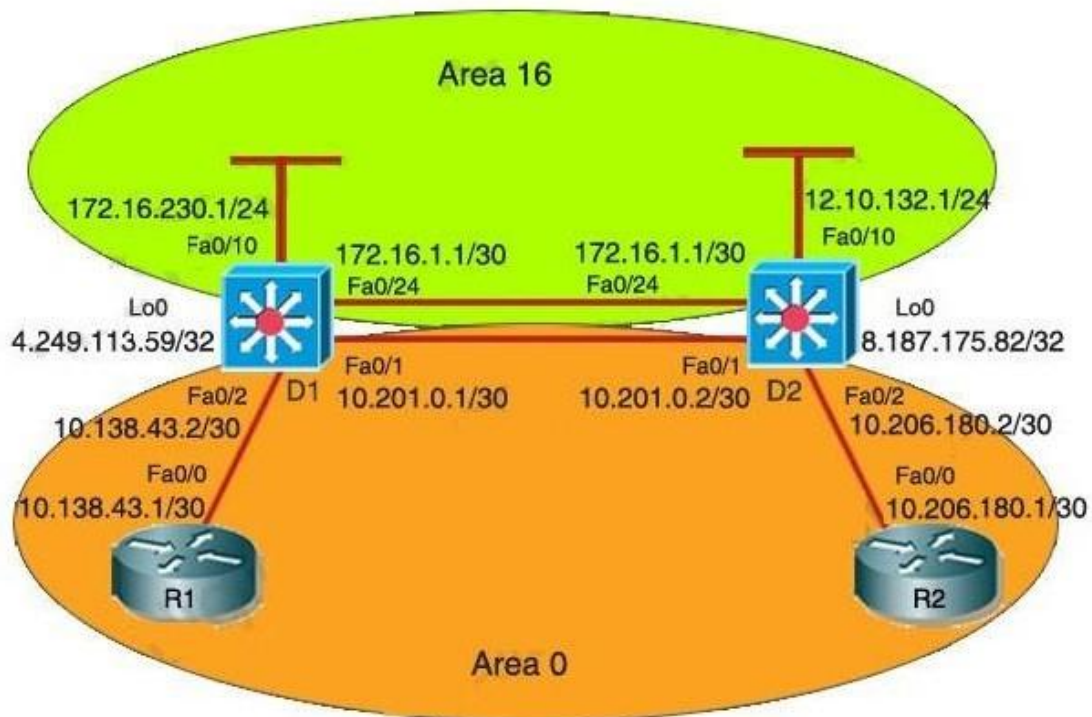
Based on the information shown above, what is most likely causing the different missing routes throughout the network?

- A. Area 16 is configured with authentication.
- B. Area 16 has been configured to use the same interfaces as Area 0.
- C. Area 0 and Area 32 have been configured with mismatched LSA numbers.
- D. Area 16 has been configured as a total stub network
- E. Area 16 has been configured as a stub network
- F. Area 0 is discontinuous.
- G. None of the above

**Correct Answer: F**

**QUESTION 8**

The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.



Which configuration command on D1 (with a similar command on D2) will provide an immediate solution to the missing route problem?


- A. no area 16 stub
- B. no area 16 authentication message-digest
- C. area 16 virtual-link 8.187.175.82
- D. area 16 virtual-link 172.16.4.2
- E. no area 16 stub no-summary
- F. network 172.16.0.0.0.255.255 area 16
- G. None of the above

**Correct Answer: C**



**QUESTION 9**

The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.



The log of d1 reports the following:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
%OSPF-5-ADJCHG: Process 1, Nbr 8.187.175.82/32 on FasstEthernet0/1 from FULL to DOWN,
Neighbor Down: Interface down or detached
```

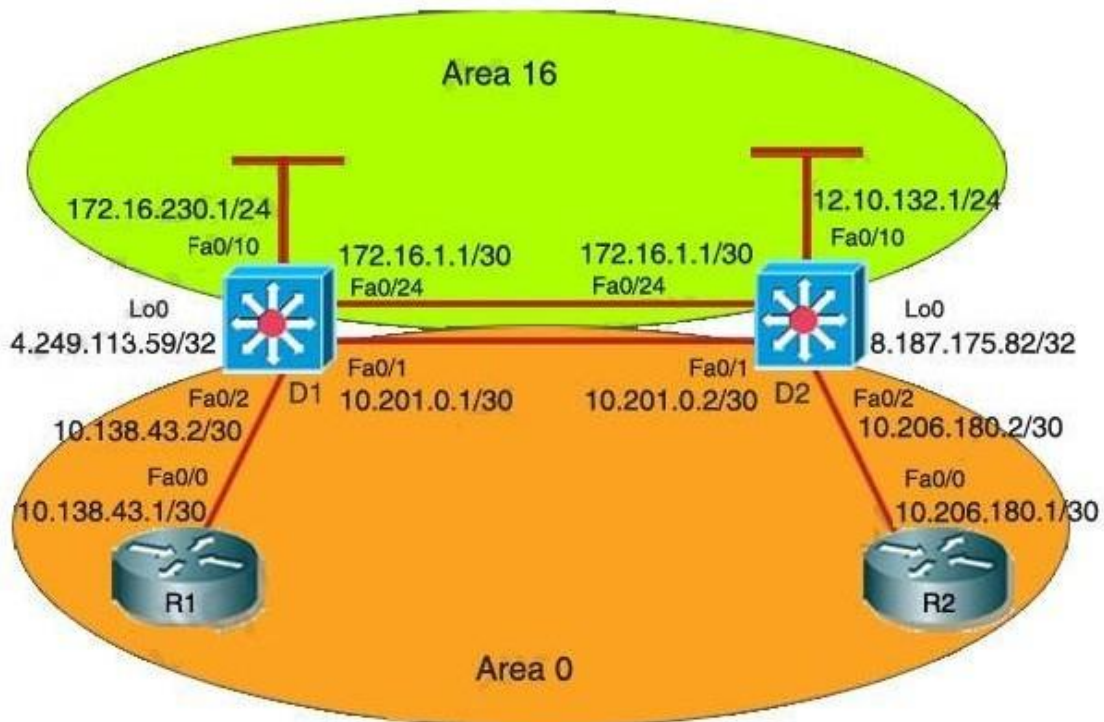
This event was anticipated due to maintenance; however, it resulted in excessive lost routes. Which route should be the only one removed from the routing tables of the routers?

- A. 8.187.175.82/32
- B. 10.138.43.0/30
- C. 10.206.180.0/30
- D. 4.249.113.59/32
- E. 10.201.0.0/30
- F. None of the above

**Correct Answer: E**

**QUESTION 10**

The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.



The R2 router has lost connectivity to R1. The following is R1's current route table:

```

172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O IA 172.16.230.0/24 [110/21] via 10.138.43.1, 00:00:03, FastEthernet0/0
O IA 172.16.209.0/24 [110/22] via 10.138.43.1, 00:00:03, FastEthernet0/0
O IA 172.16.1.0/30 [110/21] via 10.138.43.1, 00:00:03, FastEthernet0/0
10.0.0.0/30 is subnetted, 1 subnets
C 10.138.43.0 is directly connected, FastEthernet0/0
    
```

Which expected route is missing from R1's route table based on the topology during the maintenance period?

- A. O 172.16.0.0 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0
- B. O IA 9.152.105.122 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0
- C. O IA 10.138.0.0 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0
- D. O IA 10.249.0.0 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0
- E. O IA 4.249.113.59 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0

- F.  8.187.175.82 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0
- G.  10.206.180.0/30 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0

**Correct Answer: F**

#### QUESTION 11

Which three statements about the EIGRP routing protocol are true? (Choose three)

- A. EIGRP supports five generic packet types, including Hello, Database Description (DBD), Linkstate Request (LSR), Link-State Update (LSU), and LSAck.
- B. EIGRP sends periodic hello packets to the multicast IP address 224.0.0.10.
- C. EIGRP will not form a neighbor relationship with another peer when their AS number and K values, either or both are mismatched.
- D. EIGRP sends periodic hello packets to the multicast IP address 224.0.0.9.
- E. EIGRP will form a neighbor relationship with another peer even when their K values are mismatched.
- F. EIGRP supports five generic packet types, including Hello, Update, Query, Reply, and ACK packets.

**Correct Answer: BCF**

#### QUESTION 12

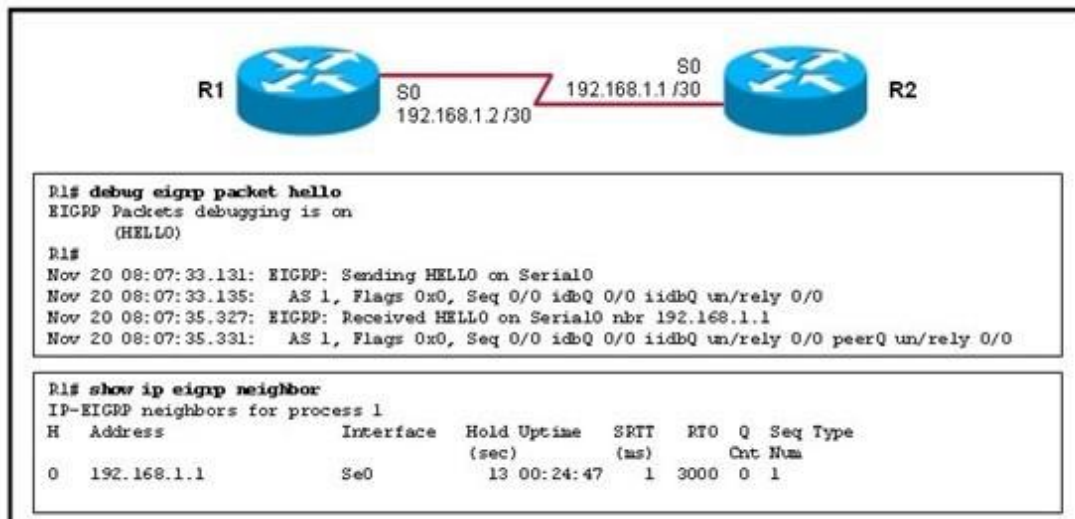
After DUAL calculations, a router has identified a successor route, but no routes have qualified as a feasible successor. In the event that the current successor goes down, what process will EIGRP use in the selection of a new successor?

- A. EIGRP will find the interface with the lowest MAC address
- B. The route will transition to the active state
- C. The route will transition to the passive state
- D. EIGRP will automatically use the route with the lowest feasible distance(FD)
- E. EIGRP will automatically use the route with the lowest advertised distance(AD)

**Correct Answer: B**

**QUESTION 13**

Refer to the exhibit. Routers R1 and R2 have established a neighbor relationship and are exchanging routing information. The network design requires that R1 receive routing updates from R2, but not advertise any routes to R2. Which configuration command sequence will successfully accomplish this task?



- A. R1(config)# router eigrp 1  
R1(config-router)# passive-interface serial 0
- B. R2(config)# router eigrp 1  
R2(config-router)# passive-interface serial 0
- C. R1(config)# access-list 20 deny any  
R1(config)# router eigrp 1  
R1(config-router)# distribute-list 20 out serial 0
- D. R2(config)# access-list 20 deny any  
R2(config)# router eigrp 1  
R2(config-router)# distribute-list 20 out serial 0
- E. R1(config)# access-list 20 permit any  
R1(config)# router eigrp 1  
R1(config-router)# distribute-list 20 in serial 0
- F. R2(config)# access-list 20 permit any  
R2(config)# router eigrp 1  
R2(config-router)# distribute-list 20 in serial 0

**Correct Answer: C**

#### QUESTION 14

EIGRP has been configured to operate over Frame Relay multipoint connections. What should the bandwidth command be set to?

- A. the CIR rate of the lowest speed connection multiplied by the number of circuits
- B. the CIR rate of the lowest speed connection
- C. the CIR rate of the highest speed connection
- D. the sum of all the CIRs divided by the number of connections

**Correct Answer: A**

#### QUESTION 15

Refer to the exhibit. EIGRP is configured on all routes in the network. On a basis of the show ip eigrp topology output provided, what conclusion can be derived?

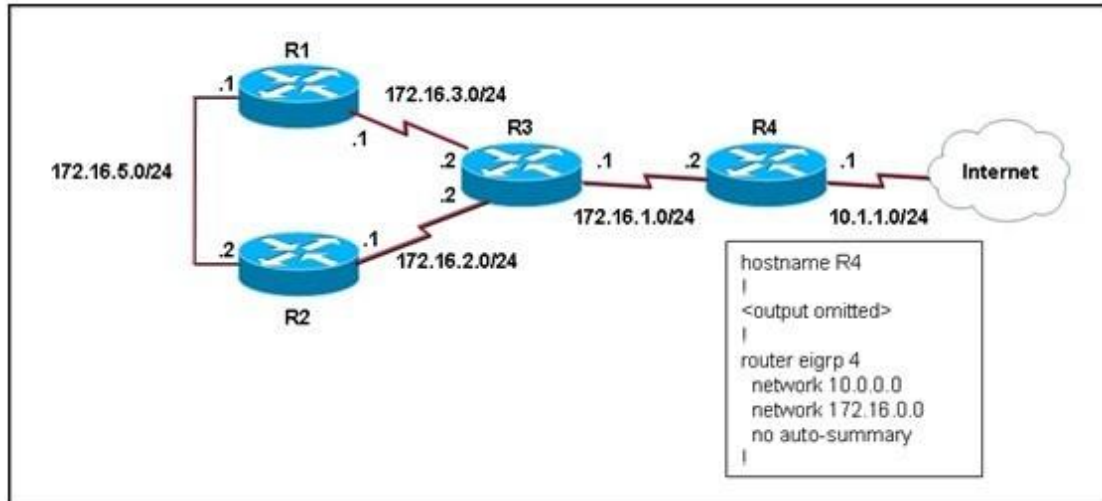
```
R1# show ip eigrp topology
<output omitted>
P 10.1.2.0/24, 1 successors, FD is 281600
   via Connected, FastEthernet0/0
A 10.6.1.0/24, 0 successors, FD is 3385160704, Q
   1 replies, active 00:00:41, query-origin: Local origin
   Remaining replies:
     via 10.1.2.1, r. FastEtherent0/0
```

- A. Router R1 can send traffic destined for network 10.6.1.0/24 out of interface FastEthernet0/0
- B. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out before it declares the neighbor unreachable
- C. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out inquiring for a second successor to network 10.6.1.0/24
- D. Router R1 is waiting for a reply from the neighbor 10.1.2.1 in response to the query sent about network 10.6.1.0/24

**Correct Answer: D**

#### QUESTION 16

Refer to the exhibit. EIGRP has been configured on all routers in the network. What additional configuration statement should be included on router R4 to advertise a default route to its neighbors?




- A. R4(config)# ip default-network 10.0.0.0
- B. R4(config)# ip route 0.0.0.0 0.0.0.0 10.1.1.1
- C. R4(config)# ip route 10.0.0.0 255.0.0.0 10.1.1.1
- D. R4(config-router)# default-information originate

**Correct Answer: A**

#### QUESTION 17

Refer to the exhibit. Router RTA is the hub router for routers RTB and RTC. The Frame Relay network is configured with EIGRP, and the entire network is in autonomous system 1. However, router RTB and RTC are not receiving each other's routes. What is the solution?



- A. Configure the auto summary command under router eigrp 1 on router RTA.
- B. Issue the no ip split horizon command on router RTA.
- C. Configure subinterfaces on the spoke routers and assign different IP address subnets for each subinterface.
- D. Check and change the access lists on router RTA.
- E. Issue the no ip split horizon eigrp 1 command on router RTA.
- F. Configure a distribute list on router RTA that allows it to advertise all routes to the spoke routers.

**Correct Answer: E**

#### **QUESTION 18**

Which two routing protocols require a metric to be configured when redistributing routes from other protocols? (Choose two.)

- A. RIP
- B. BGP
- C. IS-IS
- D. OSPF
- E. EIGRP

**Correct Answer: AE**

#### **QUESTION 19**


When troubleshooting an EIGRP connectivity problem, you notice that two connected EIGRP routers are not becoming EIGRP neighbors. A ping between the two routers was successful. What is the next thing that should be checked?

- A. Verify that the EIGRP hello and hold timers match exactly.
- B. Verify that EIGRP broadcast packets are not being dropped between the two routers with the show ip EIGRP peer command.
- C. Verify that EIGRP broadcast packets are not being dropped between the two routers with the show ip EIGRP traffic command.
- D. Verify that EIGRP is enabled for the appropriate networks on the local and neighboring router.

**Correct Answer: D**

**QUESTION 20**

Refer to the exhibit. You are the network administrator of the Route.com company. You have been tasked to implement a hub and spoke EIGRP topology over Frame Relay to provide connectivity between the networks at headquarters and all 300 spokes. Before you begin the actual implementation, which three pieces of information are more important to know than the others? (Choose three.)




- A. the Committed Information Rate of all the Frame Relay PVCs
- B. the Cisco IOS version running on all the routers
- C. the router model number of all the spoke routers
- D. the number of HQ networks connected behind the headquarter routers
- E. the routing policy, such as whether or not the spokes can be used as backup transient point between the two headquarter routers

**Correct Answer: ABE**



**QUESTION 21**

Refer to the exhibit. The Route.com company is running EIGRP between all the routers. Currently, if one of the LAN links (LAN1 or LAN2) at the headquarters flaps (goes up and down), the HQ-RTR1 and HQ-RTR2 routers will experience high CPU usage and have a long EIGRP convergence time. As the new network administrator, you are asked to investigate this situation and determine if there is a quick way to resolve this issue. Which is the most important thing that you can quickly verify first to resolve this issue?




- A. Verify that the bandwidth setting on all WAN links is correct.
- B. Verify that the HQ-RTR1 and HQ-RTR2 routers are configured to send only a default route to all the spoke routers.
- C. Verify that the HQ-RTR1 and HQ-RTR2 routers are configured for EIGRP Nonstop Forwarding.
- D. Verify that all the spoke routers are configured for auto summarization.
- E. Verify that all the spoke routers are configured as EIGRP stub.

**Correct Answer: E**

**QUESTION 22**

Refer to the exhibit. When you examine the routing tables of R1 and R4, you are not able to see the R1 Ethernet subnet on the R4 routing table. You are also not able to see the R4 Ethernet subnet on the R1 routing table. Which two configuration changes should be made to resolve this issue? Select the routers where the configuration change will be required, and select the required EIGRP configuration command(s). (Choose two.)




- A. R1 and R4
- B. R2 and R3
- C. ip summary-address eigrp 1 10.1.1.0 255.255.255.0 and ip summary-address eigrp 1 10.2.2.0 255.255.255.0
- D. variance 2
- E. eigrp stub connected
- F. no auto-summary

**Correct Answer: BF**

**QUESTION 23**

Refer to the exhibit. The actual speed of the serial links between R2 and R3 are 256 kb/s and 512 kb/s. When configuring EIGRP on routers R2 and R3, the network administrator configured the bandwidth of both serial interfaces to 512 kb/s. What will be the effect?

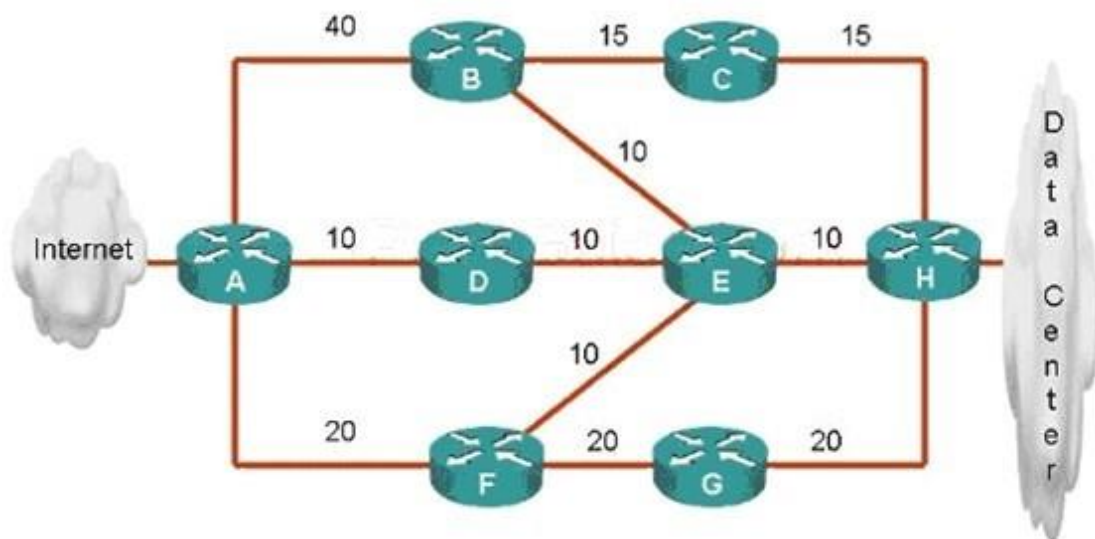


- A. EIGRP will over utilize the 512 kb/s link.
- B. The interface "delay" value used in the EIGRP metric calculation will be inaccurate on the 256 kb/s serial interface.
- C. The amount of bandwidth used for EIGRP routing protocol traffic on the 256 kb/s link can become excessive.
- D. EIGRP can load balance between the two serial links only if the variance is set to 2 or higher.
- E. Unequal cost load balancing will be disabled.

**Correct Answer: C**

**QUESTION 24**

Refer to the exhibit. ROUTE.com has just implemented this EIGRP network. A network administrator came to you for advice while trying to implement load balancing across part of their EIGRP network. If the variance value is configured as 2 on all routers and all other metric and K values are configured to their default values, traffic from the Internet to the data center will be load balanced across how many paths?

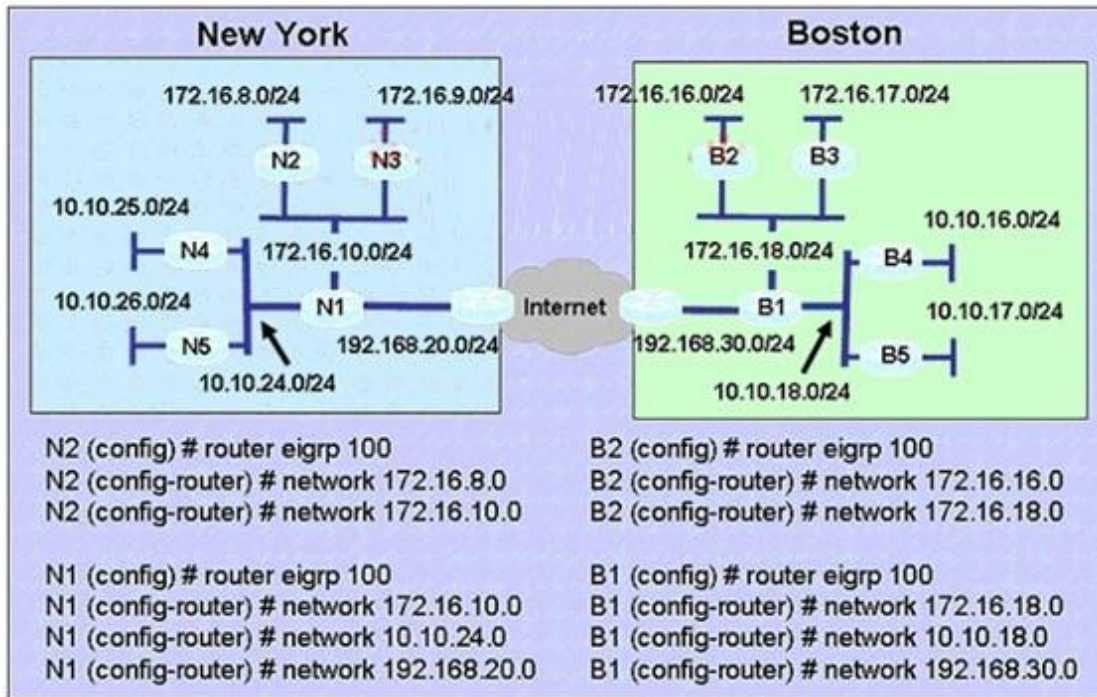


- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

**Correct Answer: C**

**QUESTION 25**

Refer to the exhibit. A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP. The show command output shows that traffic will not flow between the networks. As a network consultant, you were asked to modify the configuration and certify the interoperability of the two networks. For traffic to flow from subnet 172.16.8.0/24 to the 172.16.16.0/24 subnet. Which configuration change do you recommend?

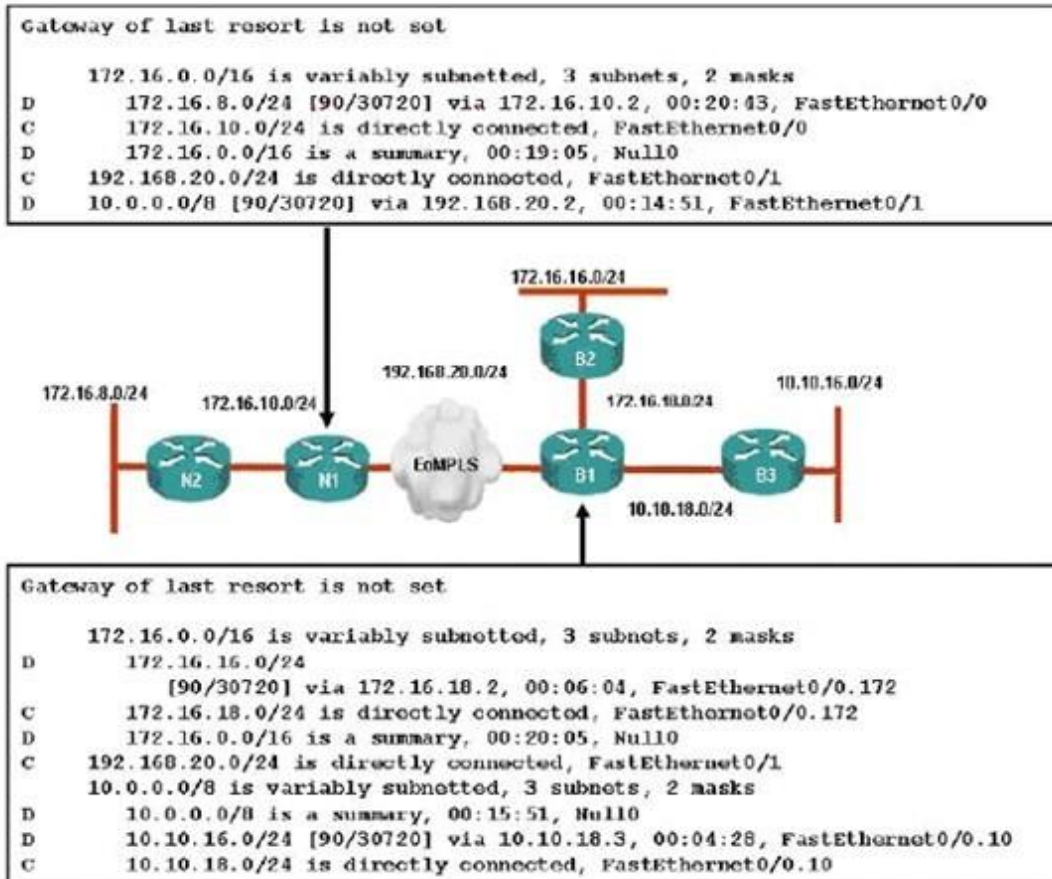


- A. Turn off autosummarization on routers N1 and B1.
- B. Add IP summary addresses to the Internet-pointing interfaces of routers N1 and B1.
- C. Turn off auto summarization on routers N2 and B2.
- D. Add wildcard masks to the network commands on routers N2 and B2.

**Correct Answer: A**

**QUESTION 26**

Refer to the exhibit. A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP over EoMPLS. As a network consultant, you were asked to verify the interoperability of the two networks. From the show ip route command output, what can you tell the customer about the traffic flow between the subnet in New York (172.16.8.0/24) and the subnets in Boston (172.16.16.0/24 and 10.10.16.0/24)?



- A. Traffic is flowing between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0 and no configuration changes are needed.
- B. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0.
- C. Traffic will flow between the 172.16.8.0 subnet and 172.16.16.0 without any further configuration changes. However, auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 10.10.16.0 subnet.
- D. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 172.16.16.0 subnet. However, traffic will flow between the 172.16.8.0 subnet and 10.10.16.0 without any further configuration changes.

**Correct Answer: B**

**QUESTION 27**

Which condition must be satisfied before an EIGRP neighbor can be considered a feasible successor?

- A. The neighbor's advertised distance must be less than or equal to the feasible distance of the current successor.

- B. The neighbor's advertised distance must be less than the feasible distance of the current successor.
- C. The neighbor's advertised distance must be greater than the feasible distance of the current successor.
- D. The neighbor's advertised distance must be equal to the feasible distance of the current successor.
- E. The neighbor's advertised distance must be greater than or equal to the feasible distance of the current successor.

**Correct Answer: B**

#### **QUESTION 28**

Which statement about a non-zero value for the load metric (k2) for EIGRP is true?

- A. A change in the load on an interface will cause EIGRP to recalculate the routing metrics and send a corresponding update out to each of its neighbors.
- B. EIGRP calculates interface load as a 5-minute exponentially weighted average that is updated every 5 minutes.
- C. EIGRP considers the load of an interface only when sending an update for some other reason.
- D. A change in the load on an interface will cause EIGRP to recalculate and update the administrative distance for all routes learned on that interface.

**Correct Answer: C**

#### **QUESTION 29**

Your network consists of a large hub-and-spoke Frame Relay network with a CIR of 56 kb/s for each spoke. Which statement about the selection of a dynamic protocol is true?

- A. EIGRP would be appropriate if LMI type ANSI is NOT used.
- B. EIGRP would be appropriate, because the Frame Relay spokes could be segmented into their own areas.
- C. EIGRP would be appropriate, because by default, queries are not propagated across the slow speed Frame Relay links.
- D. EIGRP would be appropriate, because you can manage how much bandwidth is consumed over the Frame Relay interface.

**Correct Answer: D**

**QUESTION 30**

Refer to the exhibit. Why are the EIGRP neighbors for this router not learning the routes redistributed from OSPF?

```
router eigrp 123
 redistribute ospf 123
 network 116.16.35.0 0.0.0.255
 network 130.130.0.0
 auto-summary
!

router ospf 123
 log-adjacency-changes
 network 116.16.34.0 0.0.0.255 area 0
 neighbor 116.16.34.4
```

- A. Redistribution must be enabled mutually (in both directions) to work correctly.
- B. Auto-summary causes the OSPF routes redistributed into EIGRP to be summarized; thus the OSPF network 116.16.34 is summarized to 116.34.0.0, which is already covered by the EIGRP protocol.
- C. Default metrics are not configured under EIGRP.
- D. Both routing protocols must have unique autonomous system numbers for redistribution to function correctly.

**Correct Answer: C**

## EnsurePass.com Members Features:

1. Verified Answers researched by industry experts.
2. Q&As are downloadable in PDF and VCE format.
3. 98% success Guarantee and **Money Back** Guarantee.
4. Free updates for **180** Days.
5. **Instant Access to download the Items**

View list of All Exam provided:

<http://www.ensurepass.com/certifications?index=A>

To purchase Lifetime Full Access Membership click here:

<http://www.ensurepass.com/user/register>

**Valid Discount Code for 2015: JREH-G1A8-XHC6**

To purchase the HOT Exams:

<u>Cisco</u>		<u>CompTIA</u>		<u>Oracle</u>	<u>VMWare</u>	<u>IBM</u>
<a href="#">100-101</a>	<a href="#">640-554</a>	<a href="#">220-801</a>	<a href="#">LX0-101</a>	<a href="#">1Z0-051</a>	<a href="#">VCAD510</a>	<a href="#">C2170-011</a>
<a href="#">200-120</a>	<a href="#">200-101</a>	<a href="#">220-802</a>	<a href="#">N10-005</a>	<a href="#">1Z0-052</a>	<a href="#">VCP510</a>	<a href="#">C2180-319</a>
<a href="#">300-206</a>	<a href="#">640-911</a>	<a href="#">BR0-002</a>	<a href="#">SG0-001</a>	<a href="#">1Z0-053</a>	<a href="#">VCP550</a>	<a href="#">C4030-670</a>
<a href="#">300-207</a>	<a href="#">640-916</a>	<a href="#">CAS-001</a>	<a href="#">SG1-001</a>	<a href="#">1Z0-060</a>	<a href="#">VCAC510</a>	<a href="#">C4040-221</a>
<a href="#">300-208</a>	<a href="#">640-864</a>	<a href="#">CLO-001</a>	<a href="#">SK0-003</a>	<a href="#">1Z0-474</a>	<a href="#">VCP5-DCV</a>	<a href="#">RedHat</a>
<a href="#">350-018</a>	<a href="#">642-467</a>	<a href="#">ISS-001</a>	<a href="#">SY0-301</a>	<a href="#">1Z0-482</a>	<a href="#">VCP510PSE</a>	<a href="#">EX200</a>
<a href="#">352-001</a>	<a href="#">642-813</a>	<a href="#">JK0-010</a>	<a href="#">SY0-401</a>	<a href="#">1Z0-485</a>		<a href="#">EX300</a>
<a href="#">400-101</a>	<a href="#">642-832</a>	<a href="#">JK0-801</a>	<a href="#">PK0-003</a>	<a href="#">1Z0-580</a>		
<a href="#">640-461</a>	<a href="#">642-902</a>			<a href="#">1Z0-820</a>		

