- Vendor: Cisco
- Exam Code: 300-101
- Exam Name: Implementing Cisco IP Routing (ROUTE)
- Question 41 -- Question 60

QUESTION 41
Lab Simulation 2 - EIGRP Stub

Scenario

JS Industries has expanded their business with the addition of their first remote office. The remote office router (R3) was previously configured and all Corporate subnets were reachable from R3. JS Industries is interested in using route summarization along with the EIGRP Stub Routing feature to increase network stability while reducing the memory usage and bandwidth utilization to R3. Another network professional was tasked with implementing this solution. However, in the process of configuring EIGRP stub routing connectivity with the remote network devices off of R3 has been lost.

Currently EIGRP is configured on all routers R2, R3, and R4 in the network. Your task is to identify and resolve the cause of connectivity failure with the remote office router R3. Once the issue has been resolved you should complete the task by configuring route summarization only to the remote office router R3.

You have corrected the fault when pings from R2 to the R3 LAN interface are successful, and the R3 IP routing table only contains 2 10.0.0.0 subnets.

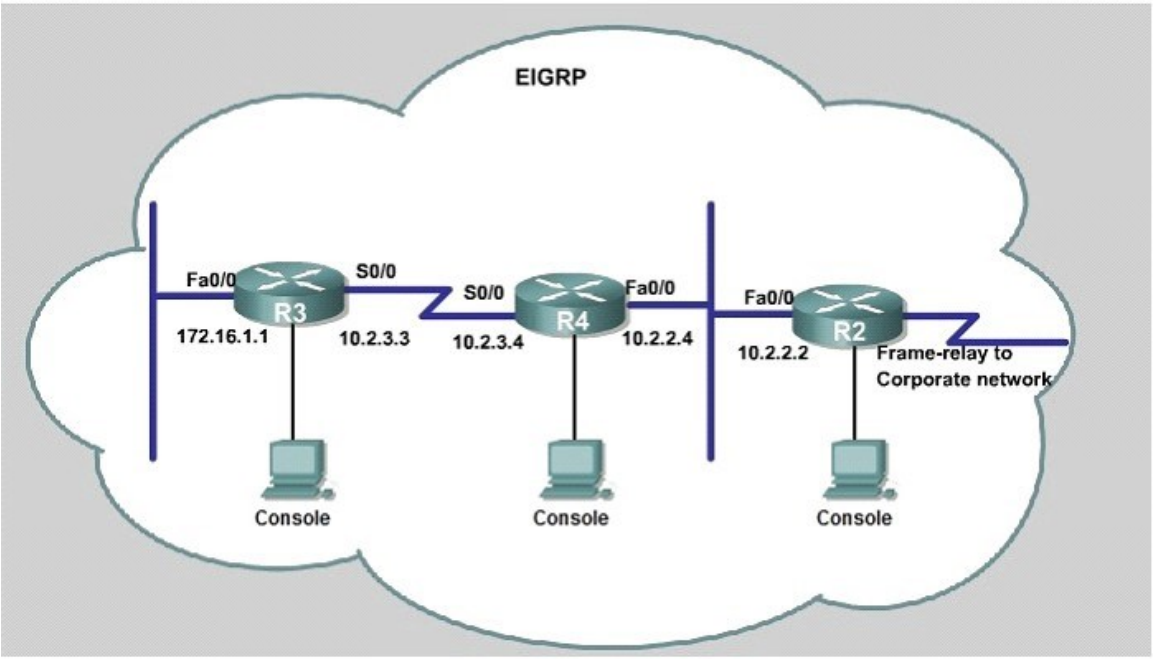

Answer:

First we have to figure out why R3 and R4 can not communicate with each other. Use the "***show running-config***" command on router R3

```
R3#show run

<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

R3#show 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
       N1 - 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

     10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
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, Serial0/1
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, 00:00:03, Serial0/0
D       10.2.8.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D       10.2.9.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 02:04:06, 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 summary the networks. Therefore, to make the routing table of R3 has only 2 subnets we have to summary other subnets into one subnet.

In the output if we don't see the summary line (like 10.0.0.0/8 is a summary…) then we should use the command ip summary-address eigrp 123 10.2.0.0 255.255.0.0 so that all the ping can work well.

In conclusion, we will use the ip summary-address eigrp 123 10.2.0.0 255.255.0.0 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 10.2.0.0 255.255.0.0
```

Now we jump back to R3 and use the show ip route command to verify the effect, the output is shown below:

```
R3#show 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
       N1 - 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

     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 [90/161280] via 10.2.3.4, 00:00:11, Serial0/0
C       10.2.3.0/24 is directly connected, Serial0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 02:04:06, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
```

(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.

```
R3#show 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
       N1 - 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

     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 [90/161280] via 10.2.3.4, 00:00:11, Serial0/0
C       10.2.3.0/24 is directly connected, Serial0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 02:04:06, 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.

```
R3(config-if)# end
R3# copy run start
R4(config-if)# end
R4# copy run start
```

QUESTION 42
Lab Simulation 3 - IPv6 OSPF Virtual Link

Scenario

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 OPSF features and does not change the current area assignments. You will know that you have corrected the fault when R4's loopback address (FEC0:4:4) can be seen in R1's IPv6 routing table and you can ping from R1 to R4 loopback address.

Special Note: To gain the maximum number of points you must remove all incorrect or unneeded configuration statements related to this issue.

Topology

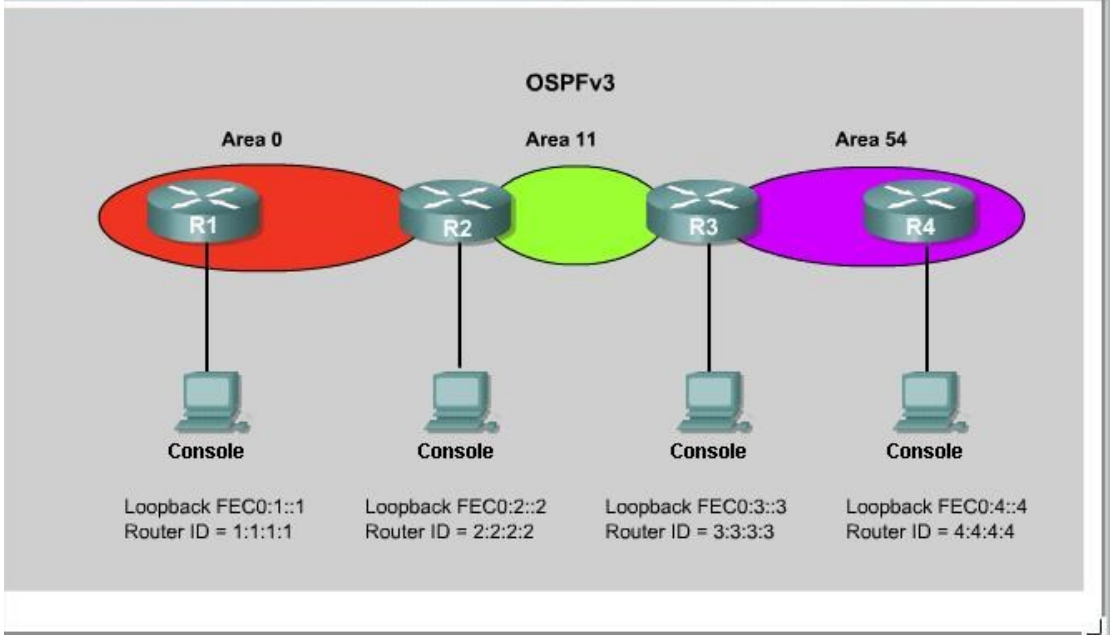

Answer:

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

```
R2#
R2#show run
Building configuration...

Current configuration : 949 bytes
!
version 12.2
!
hostname R2
!
boot-start-marker
boot-end-marker
!
!
no aaa new-model
!
resource policy
memory-size iomem 10
ip subnet-zero
!
!
ip cef
!
!
ipv6 unicast-routing
!
voice-card 0
 no dspfarm
!
!
!
!
!
!
!
!
interface Loopback0
 no ip address
 ipv6 address FEC0:2::2/64
 ipv6 ospf 1 area 11
!
interface Serial0/0/0
 no ip address
 encapsulation
!
interface Serial0/0/1
 description link to R3 (Area 11)
 ip address 10.1.0.1 255.255.255.0
 encapsulation
 clock rate 125000
 ipv6 ospf 1 area 11
!
interface FastEthernet0/0
 description link to R1 (Area 0)
 no ip address
 duplex auto
 speed auto
 ipv6 address FE80::2 link-local
 ipv6 address 2001::2/64
 ipv6 ospf 1 area 0
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
!
!
!
ipv6 router ospf 1
 router-id 2.2.2.2
 log-adjacency-changes
!
!
!
!
!
line con 0
 exec-timeout 0 0
 logging synchronous
line aux 0
line vty 0 15
 transport input ssh
!
end
R2#
```

```
R3#
R3#show run
Building configuration...

Current configuration : 965 bytes
!
version 12.2
!
hostname R3
!
boot-start-marker
boot-end-marker
!
!
no aaa new-model
!
resource policy
memory-size iomem 10
ip subnet-zero
!
!
ip cef
!
!
ipv6 unicast-routing
!
voice-card 0
 no dspfarm
!
!
!
!
!
!
!
!
interface Loopback0
 no ip address
 ipv6 address FEC0:3::3/64
 ipv6 ospf 1 area 11
!
interface Serial0/0/0
 no ip address
 encapsulation
!
interface Serial0/0/1
 description link to R2 (Area 11)
 ip address 10.1.0.2 255.255.255.0
 encapsulation
 ipv6 ospf 1 area 11
!
interface FastEthernet0/0
 description link to R4 (Area 54)
 no ip address
 duplex auto
 speed auto
 ipv6 address FE80::3 link-local
 ipv6 address 2001:54::3/64
 ipv6 ospf 1 area 54
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
!
!
ipv6 router ospf 1
 router-id 3.3.3.3
 log-adjacency-changes
 area 54 virtual-link 4.4.4.4
!
!
!
!
!
line con 0
 exec-timeout 0 0
 logging synchronous
line aux 0
line vty 0 15
 transport input ssh
!
end
R3#
```

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 don'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 43

Lab Simulation 4 - Policy Based Routing

You are a network engineer with ROUTE.com, a small IT company. ROUTE.com has two connections to the Internet; one via a frame relay link and one via an EoMPLS link. IT policy requires that all outbound HTTP traffic use the frame relay link when it is available. All other traffic may use either link. No static or default routing is allowed.

Choose and configure the appropriate path selection feature to accomplish this task.

You may use the Test Workstation to generate HTTP traffic to validate your solution.

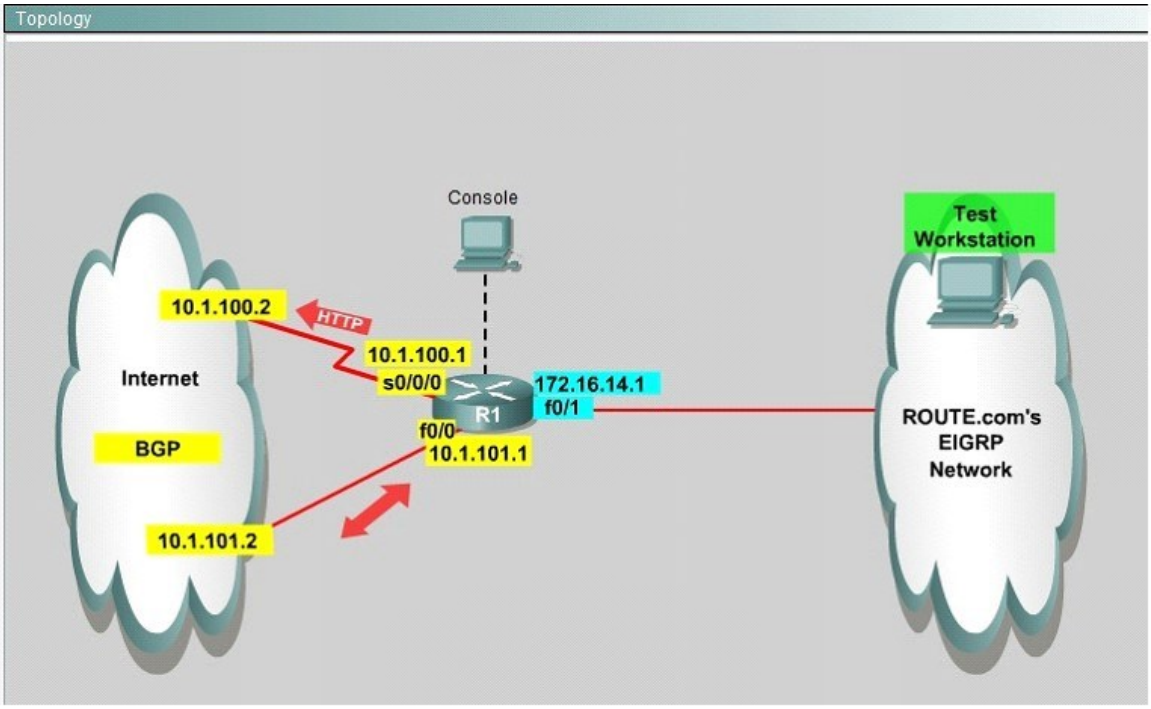

Answer:

(1) First create the access list that catches the HTTP traffic:

```
R1(config)# access-list 101 permit tcp any any eq www
```

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

```
R1(config)# route-map pbr permit 10
R1(config-route-map)# match ip address 101
R1(config-route-map)# set ip next-hop 10.1.100.2
R1(config-route-map)# exit
R1(config)# route-map pbr permit 20
```

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

```
R1(config-route-map)# exit
R1(config)# int fa0/1
R1(config-if)# ip policy route-map pbr
R1(config-if)# exit
R1(config)# exit
```

QUESTION 44

Lab Simulation - EIGRP OSPF Redistribution

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 OPSF 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.

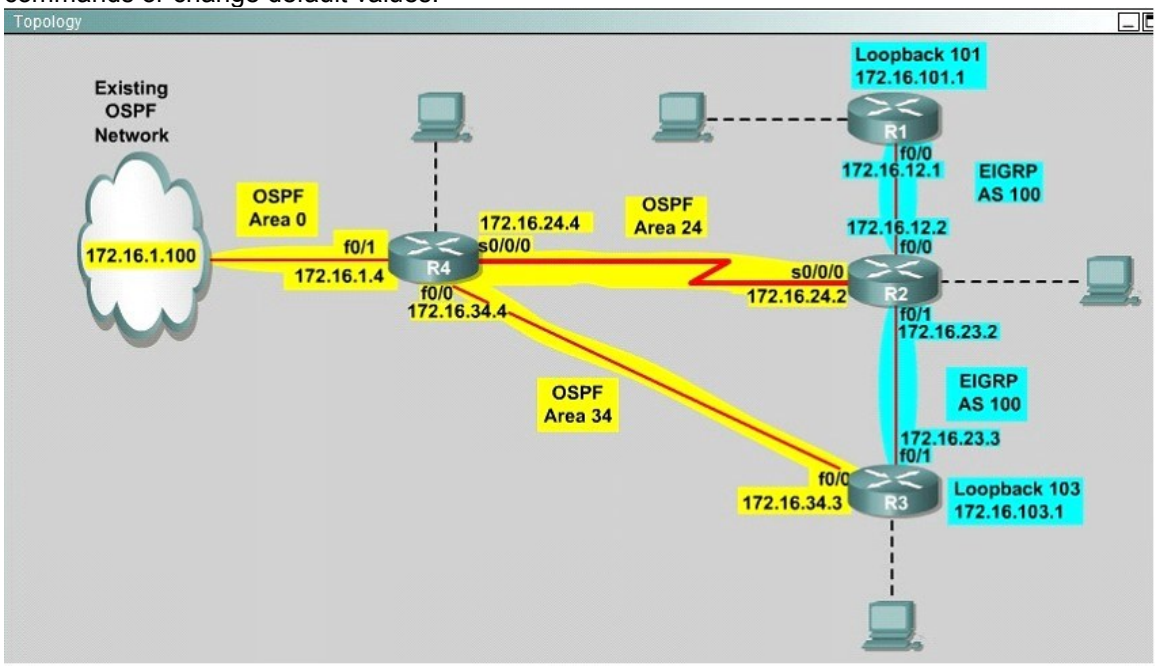

Answer:

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

Write down these 5 parameters, notice that we have to divide the Delay by 10 because the metric unit is in 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
```

Note: In fact, these parameters are just used for reference and we can use other parameters with no problem.

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

Finally you should try to "***show ip route***" to see the 172.16.100.1 network (the network behind R4) in the routing table of R1 and make a ping from R1 to this network.

Note: If the link between R2 and R3 is FastEthernet link, we must put the command below under EIGRP process to make traffic from R1 to go through R3 (R1 -> R2 -> R3 -> R4), which is better than R1 -> R2 -> R4.

```
R2(config-router)# distance eigrp 90 105
```

This command sets the Administrative Distance of all EIGRP internal routes to 90 and all EIGRP external routes to 105, which is smaller than the Administrative Distance of OSPF (110) -> the link between R2 & R3 will be preferred to the serial link between R2 & R4.
Note: The actual OPSF and EIGRP process numbers may change in the actual exam so be sure to use the actual correct values, but the overall solution is the same.

QUESTION 45
You have been asked to evaluate how EIGRP is functioning in a customer network.

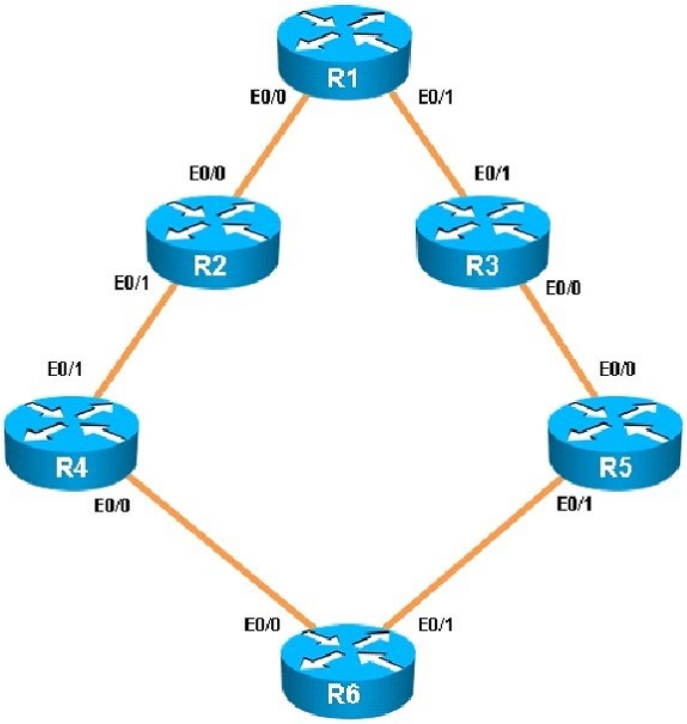

Traffic from R1 to R61 s Loopback address is load shared between R1-R2-R4-R6 and R1- R3-R5-R6 paths. What is the ratio of traffic over each path?

A. 1:1
B. 1:5
C. 6:8
D. 19:80

Answer: D
Explanation:
First, find the IP address of the loopback0 interface on R6:

R6

```
!
!
no ip domain-lookup
no ipv6 cef
ipv6 multicast rpf use-bgp
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 150.1.6.6 255.255.255.255
!
interface Loopback1
 ip address 172.16.6.6 255.255.255.255
!
interface Ethernet0/0
 ip address 192.168.46.6 255.255.255.0

R6#$
```

We see that it is 150.1.6.6, so we issue the "show ip route 150.1.6.6" command from R1 and see this:

```
R1#sh ip route 150.1.6.6
Routing entry for 150.1.6.6/32
  Known via "eigrp 1", distance 90, metric 461056, type internal
  Redistributing via eigrp 1
  Last update from 192.168.13.3 on Ethernet0/1, 00:00:08 ago
  Routing Descriptor Blocks:
  * 192.168.13.3, from 192.168.13.3, 00:00:08 ago, via Ethernet0/1
      Route metric is 1938688, traffic share count is 19
      Total delay is 65730 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3
    192.168.12.2, from 192.168.12.2, 00:00:08 ago, via Ethernet0/0
      Route metric is 461056, traffic share count is 80
      Total delay is 8010 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3

R1#
```

Notice the "traffic share count" shows 19 for the first path, and 80 for the second path.

QUESTION 46

You have been asked to evaluate how EIGRP is functioning in a customer network.

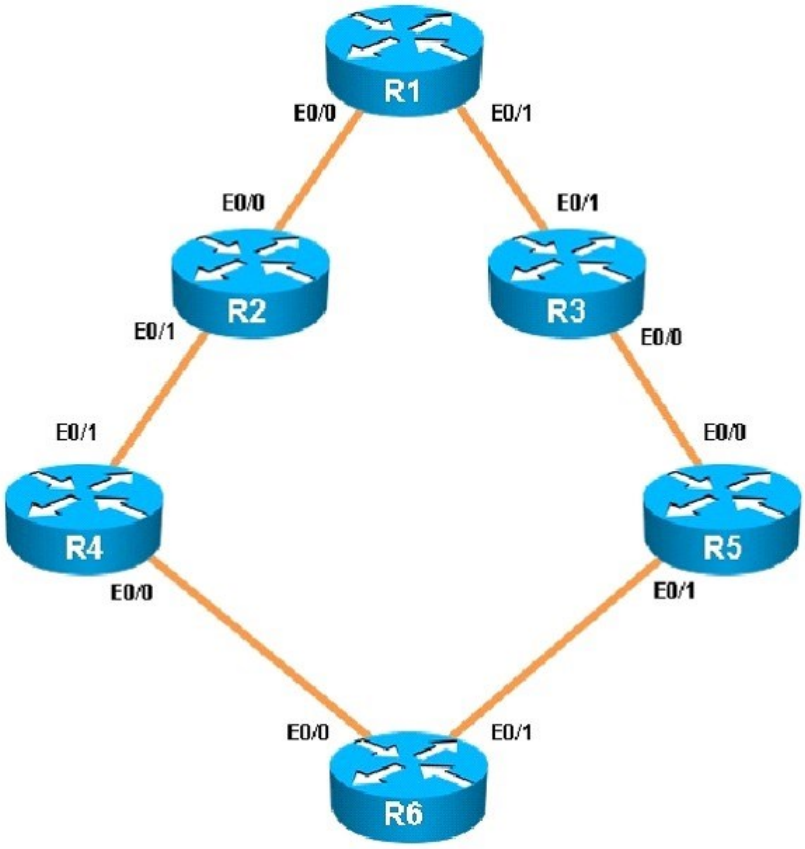

What type of route filtering is occurring on R6

A. Distribute-list using an ACL

B. Distribute-list using a prefix-list
C. Distribute-list using a route-map
D. An ACL using a distance of 255

Answer: A
Explanation:
The configuration on R6 is as follows:

```
router eigrp 1
 distribute-list 1 out
 network 150.1.6.6 0.0.0.0
 network 172.16.6.6 0.0.0.0
 network 192.168.46.0
 network 192.168.56.0
!
!
!
no ip http server
!
access-list 1 permit 192.168.46.0
access-list 1 permit 192.168.56.0
access-list 1 permit 150.1.6.6
access-list 1 deny   172.16.6.6
access-list 2 permit 192.168.47.1
access-list 2 permit 192.168.13.1
access-list 2 permit 192.168.12.1
access-list 2 deny   150.1.1.1
!
```

This is a standard distribute list using access list number 1.

QUESTION 47
You have been asked to evaluate how EIGRP is functioning in a customer network.

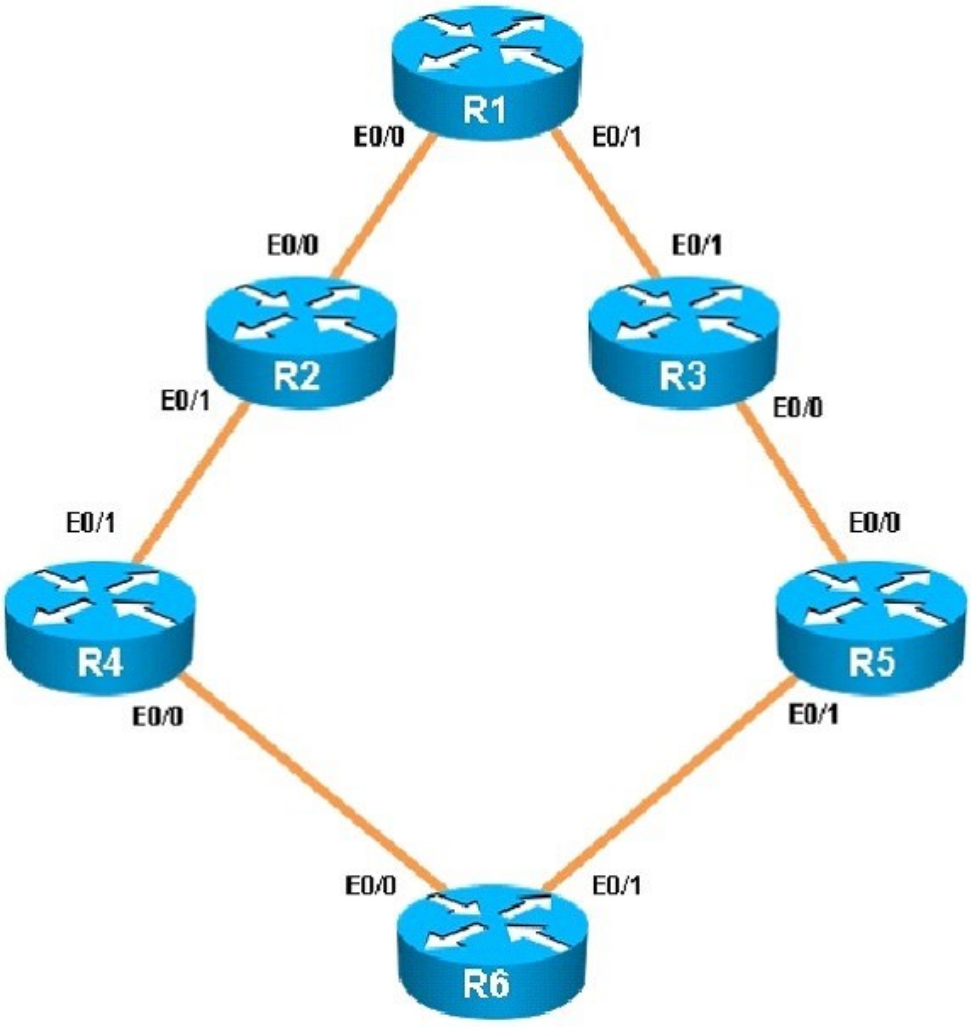

Which key chain is being used for authentication of EIGRP adjacency between R4 and R2?

A. CISCO
B. EIGRP
C. key
D. MD5

Answer: A
Explanation:
R4 and R2 configs are as shown below:

R4

```
!
no ip domain-lookup
no ipv6 cef
ipv6 multicast rpf use-bgp
!
key chain CISCO
 key 1
  key-string firstkey
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 150.1.4.4 255.255.255.255
!
interface Ethernet0/0
 description Link to R6

 ip address 192.168.46.4 255.255.255.0
```

R2

```
!
no aaa new-model
clock timezone PST -8 0
!
ip cef
!
!
no ipv6 cef
ipv6 multicast rpf use-bgp
!
key chain CISCO
 key 1
  key-string firstkey
key chain FIRSTKEY
 key 1
  key-string CISCO
key chain R3
 key 1
  key-string R3
 key 2
  key-string R1
!
!
!

 --- More (164) ---
```

Clearly we see the actual key chain is named CISCO.

QUESTION 48

You have been asked to evaluate how EIGRP is functioning in a customer network.

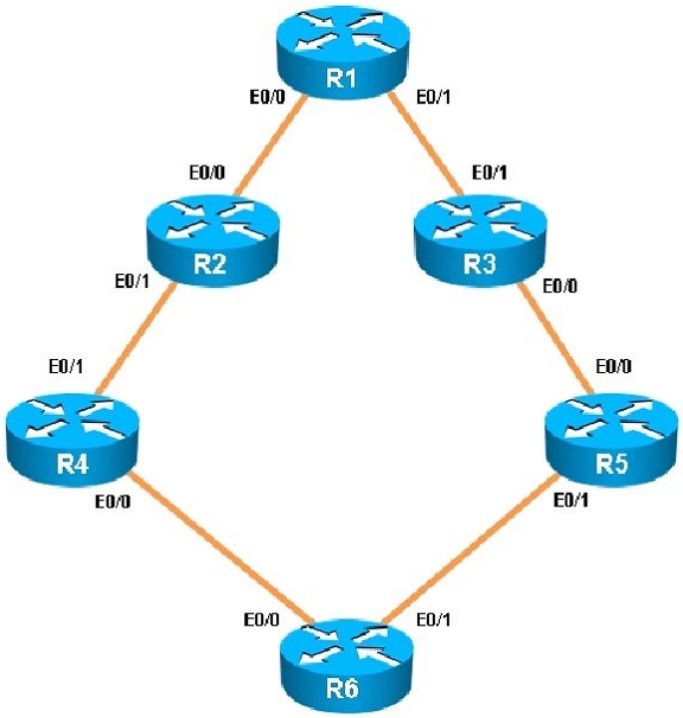

What is the advertised distance for the 192.168.46.0 network on R1?

A. 333056
B. 1938688
C. 1810944
D. 307456

Answer: C
Explanation:
R1's routing table is as follows

R1

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 2 subnets
C        150.1.1.1 is directly connected, Loopback0
D        150.1.6.6 [90/1938688] via 192.168.13.3, 00:13:02, Ethernet0/1
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.12.0/24 is directly connected, Ethernet0/0
L        192.168.12.1/32 is directly connected, Ethernet0/0
      192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.13.0/24 is directly connected, Ethernet0/1
L        192.168.13.1/32 is directly connected, Ethernet0/1
D     192.168.24.0/24 [90/1862144] via 192.168.13.3, 00:13:02, Ethernet0/1
D     192.168.35.0/24 [90/1785088] via 192.168.13.3, 00:13:08, Ethernet0/1
D     192.168.46.0/24 [90/1810944] via 192.168.13.3, 00:13:02, Ethernet0/1
D     192.168.56.0/24 [90/1810688] via 192.168.13.3, 00:13:03, Ethernet0/1

R1#
```

The numbers after the route specify the administrative distance of the route (90 for EIGRP) and the distance metric of that particular route, which is shows as 1810944 for the 192.168.46.0 route.

QUESTION 49
You have been asked to evaluate how EIGRP is functioning in a customer network.

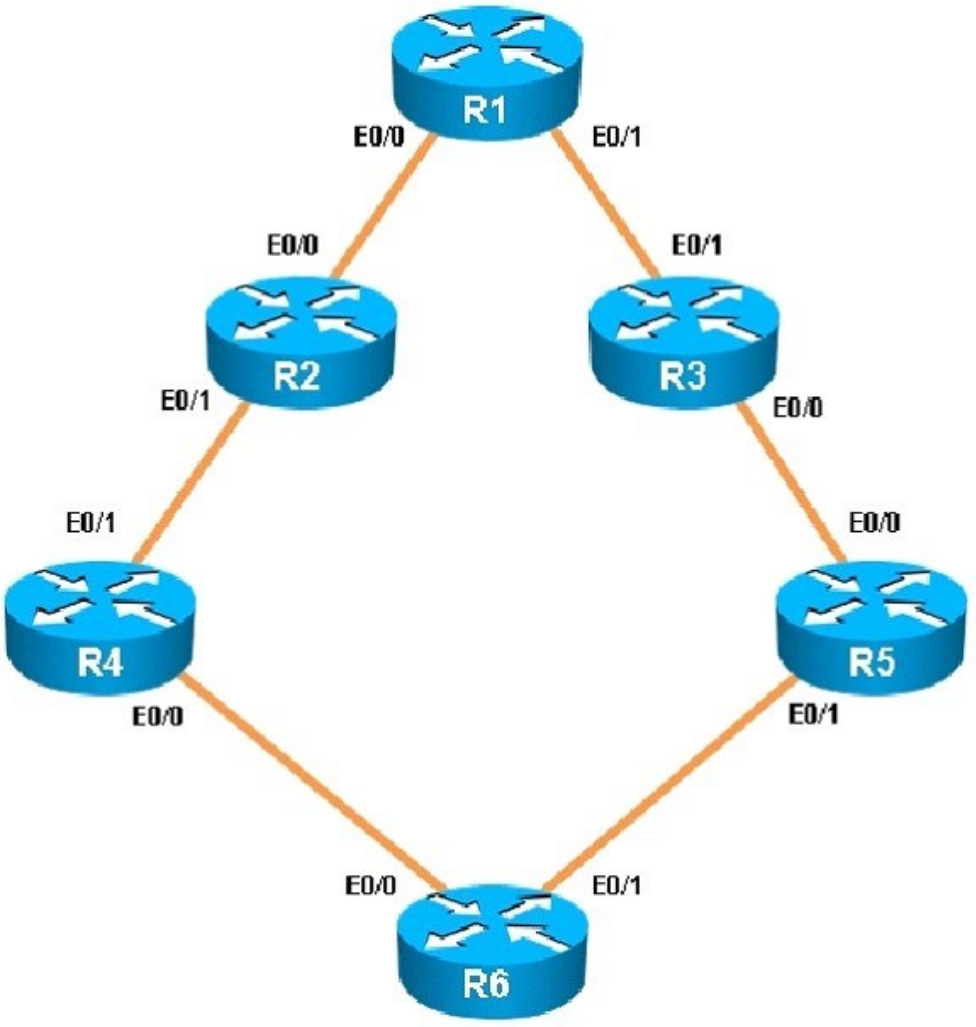

What percent of R1's interfaces bandwidth is EIGRP allowed to use?

A. 10
B. 20
C. 30
D. 40

Answer: B
Explanation:
The relevant configuration of R1 is shown below:

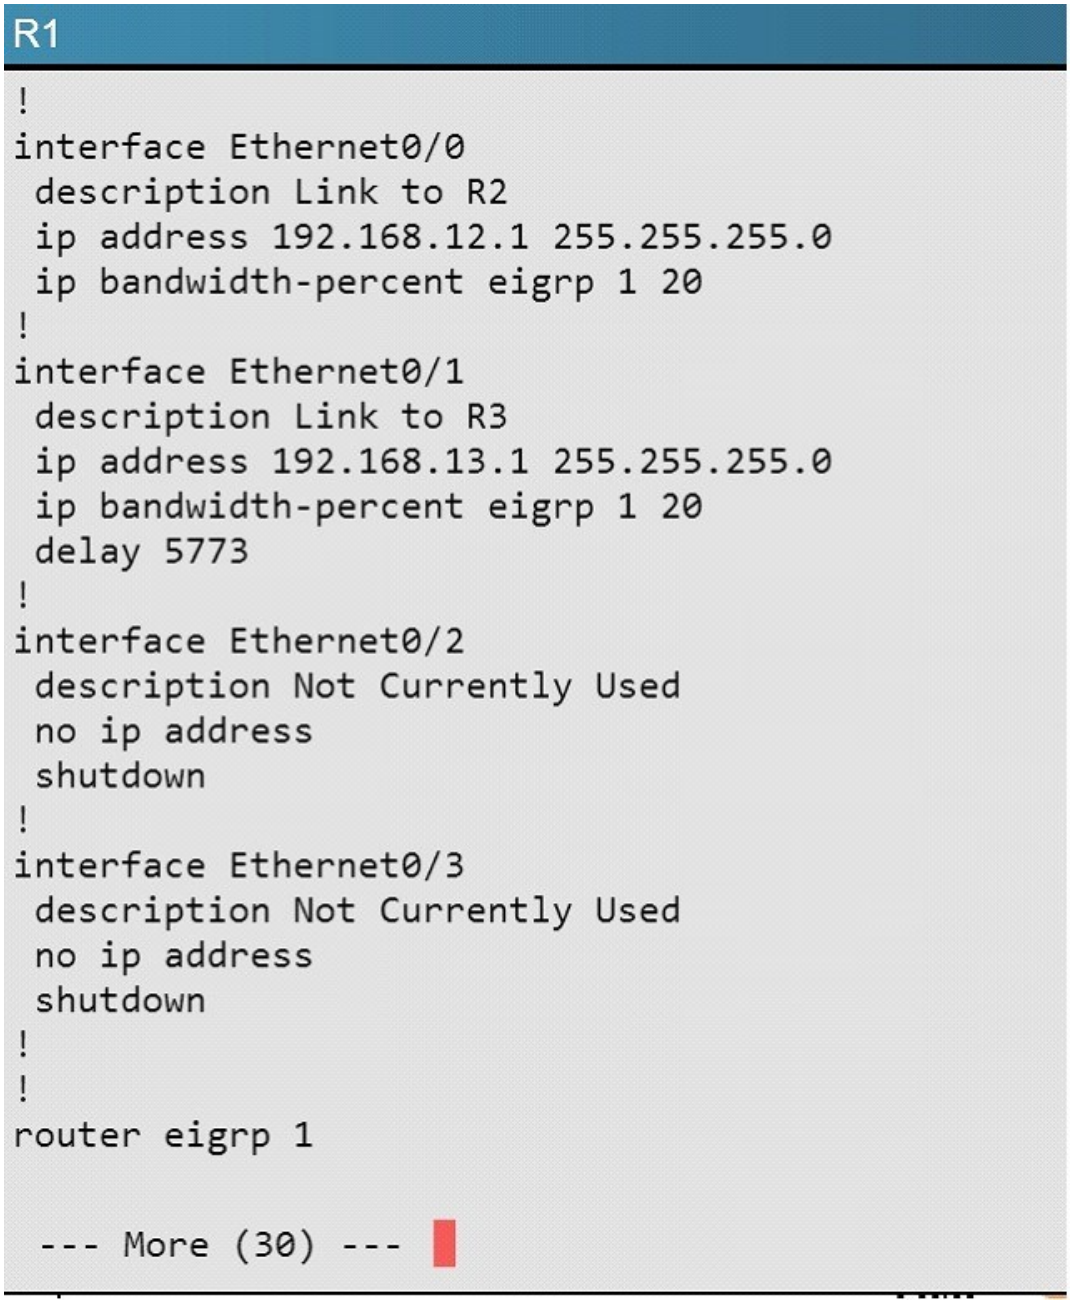

ip bandwidth-percent eigrp 1 20
1 = the EIGRP AS
20 = 20% of the bandwidth

QUESTION 50
Scenario:
You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running-config command.

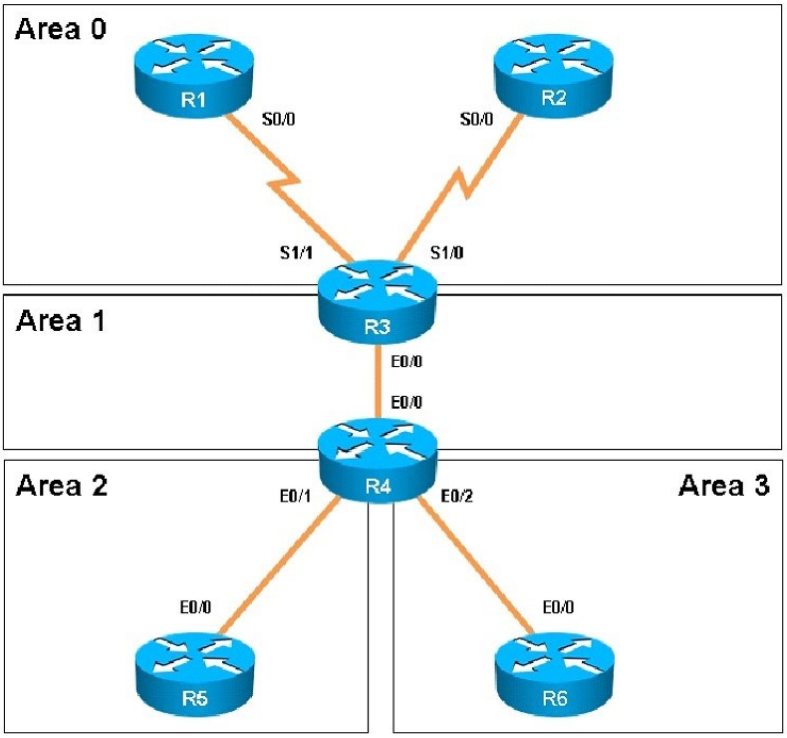

How old is the Type 4 LSA from Router 3 for area 1 on the router R5 based on the output you have examined?

A. 1858
B. 1601
C. 600
D. 1569

Answer: A
Explanation:
Part of the "show ip ospf topology" command on R5 shows this:

```
Link ID         ADV Router      Age         Seq#       Checksum
1.1.1.1         4.4.4.4         600         0x80000002 0x007ED6
2.2.2.2         4.4.4.4         1858        0x80000009 0x004208
3.3.3.3         4.4.4.4         1858        0x80000009 0x00E8FB
4.4.4.4         4.4.4.4         1858        0x80000009 0x00F716
6.6.6.6         4.4.4.4         1601        0x80000009 0x008766
6.6.66.6        4.4.4.4         1601        0x80000009 0x00C7D4
192.168.13.0    4.4.4.4         600         0x80000002 0x006182
192.168.23.0    4.4.4.4         1858        0x80000009 0x00E4ED
192.168.34.0    4.4.4.4         1858        0x80000009 0x004026
192.168.46.0    4.4.4.4         1858        0x80000009 0x00BB9E

R5#
```

The Link ID of R3 (3.3.3.3) shows the age is 1858.

QUESTION 51

Scenario:

You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running-config command.

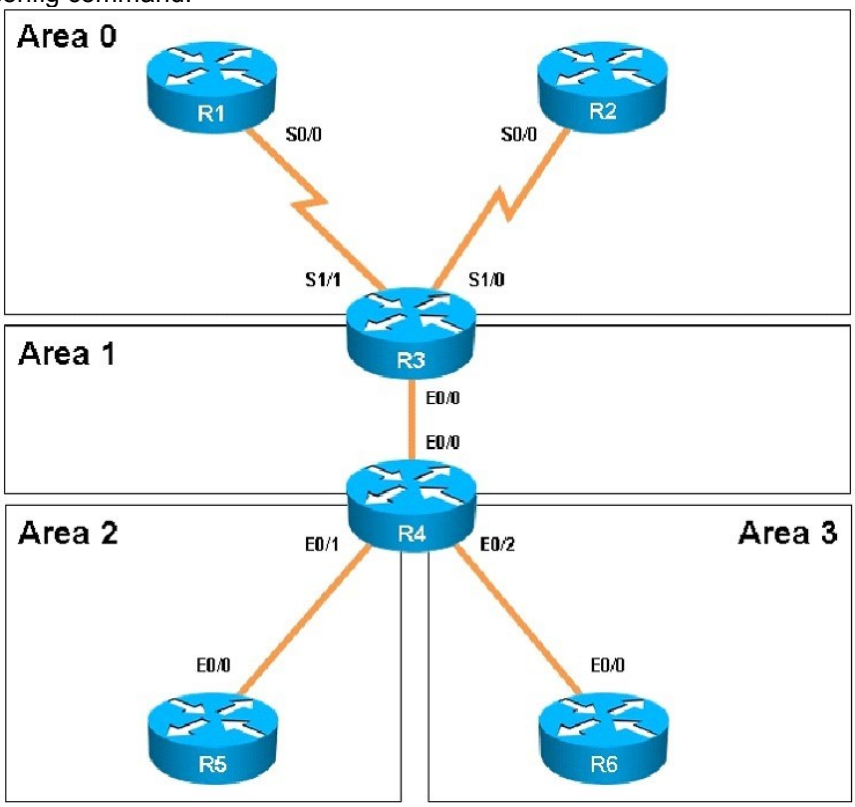

Which of the following statements is true about the serial links that terminate in R3

A. The R1-R3 link needs the neighbor command for the adjacency to stay up
B. The R2-R3 link OSPF timer values are 30, 120, 120
C. The R1-R3 link OSPF timer values should be 10,40,40
D. R3 is responsible for flooding LSUs to all the routers on the network.

Answer: B

Explanation:

We can see the configured timers using the following command:

```
R3#show ip ospf interface serial 1/0
Serial1/0 is up, line protocol is up
  Internet Address 192.168.13.3/24, Area 0, Attached via Network Statement
  Process ID 100, Router ID 3.3.3.3, Network Type NON_BROADCAST, Cost: 1943
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1943      no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 3.3.3.3, Interface address 192.168.13.3
  Backup Designated router (ID) 1.1.1.1, Interface address 192.168.13.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 11
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)

R3#
```

QUESTION 52

Scenario:

You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running-config command.

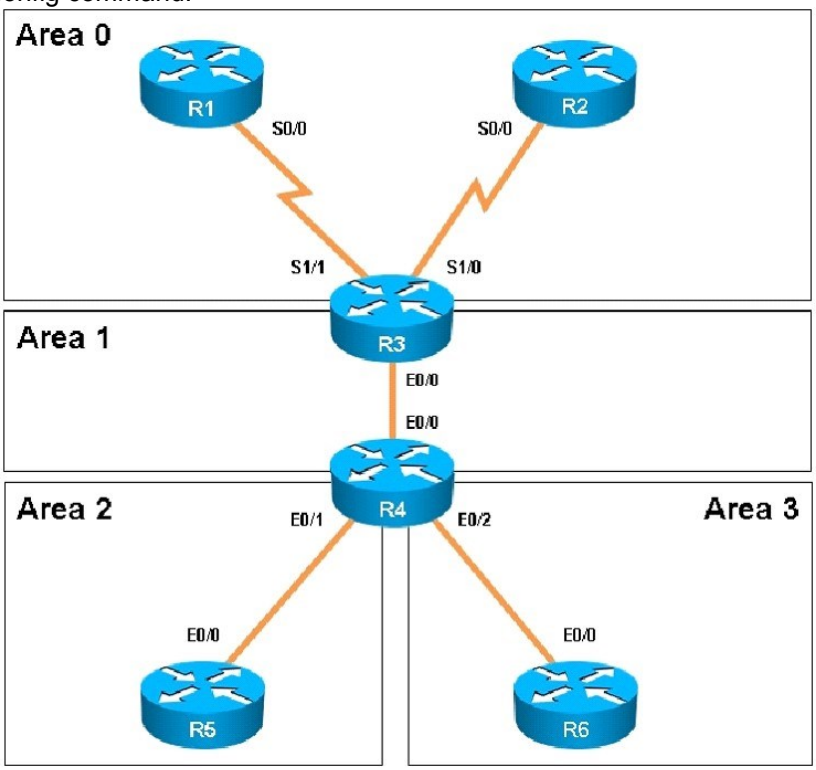

How many times was SPF algorithm executed on R4 for Area 1?

A. 1
B. 5
C. 9
D. 20
E. 54
F. 224

Answer: C
Explanation:
This can be found using the "show ip ospf" command on R4.
Look for the Area 1 stats which shows this:

```
    Flood list length 0
Area 1
    Number of interfaces in this area is 2 (1 loopback)
    This area has transit capability: Virtual Link Endpoint
    Area has no authentication
    SPF algorithm last executed 04:32:05.765 ago
    SPF algorithm executed 9 times
    Area ranges are
    Number of LSA 15. Checksum Sum 0x05538F
    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 2
    Number of interfaces in this area is 1
    It is a NSSA area
    Perform type-7/type-5 LSA translation
    Area has no authentication
```

QUESTION 53
Scenario:
You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running-config command.

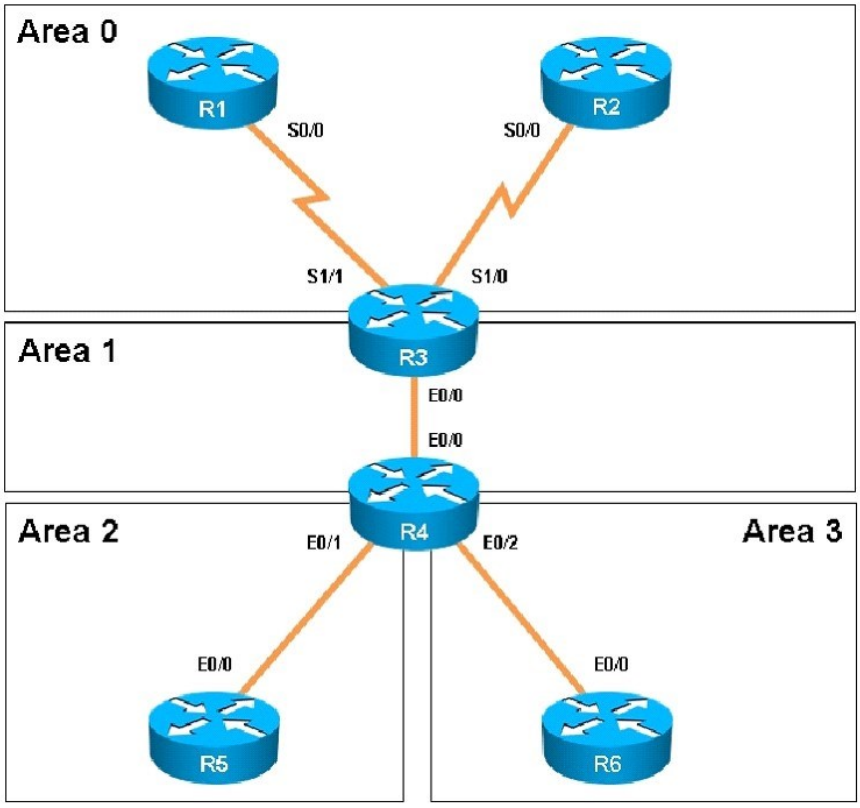

Areas of Router 5 and 6 are not normal areas, inspect their routing tables and determine which statement is true?

A. R5's Loopback and R6's Loopback are both present in R5's Routing table
B. R5's Loopback and R6's Loopback are both present in R6's Routing table
C. Only R5's loopback is present in R5's Routing table
D. Only R6's loopback is present in R5's Routing table
E. Only R5's loopback is present in R6's Routing table

Answer: A
Explanation:
Here are the routing tables of R5 and R6:

R5

```
      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/2544] via 192.168.45.4, 00:46:34, Ethernet0/0
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/2544] via 192.168.45.4, 04:57:48, Ethernet0/0
      3.0.0.0/32 is subnetted, 1 subnets
O IA     3.3.3.3 [110/601] via 192.168.45.4, 04:57:48, Ethernet0/0
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/301] via 192.168.45.4, 04:57:48, Ethernet0/0
      5.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C        5.5.1.0/24 is directly connected, Loopback1
L        5.5.1.1/32 is directly connected, Loopback1
C        5.5.2.0/24 is directly connected, Loopback2
L        5.5.2.1/32 is directly connected, Loopback2
C        5.5.3.0/24 is directly connected, Loopback3
L        5.5.3.1/32 is directly connected, Loopback3
C        5.5.4.0/24 is directly connected, Loopback4
L        5.5.4.1/32 is directly connected, Loopback4
C        5.5.5.5/32 is directly connected, Loopback0
      6.0.0.0/32 is subnetted, 2 subnets
O IA     6.6.6.6 [110/1600] via 192.168.45.4, 04:56:43, Ethernet0/0
O IA     6.6.66.6 [110/601] via 192.168.45.4, 04:56:43, Ethernet0/0
O IA  192.168.13.0/24 [110/2543] via 192.168.45.4, 00:46:44, Ethernet0/0
O IA  192.168.23.0/24 [110/2543] via 192.168.45.4, 04:57:48, Ethernet0/0
O IA  192.168.34.0/24 [110/600] via 192.168.45.4, 04:57:48, Ethernet0/0

      192.168.45.0/24 is variably subnetted, 2 subnets, 2 masks
```

R6

```
R6#show ip route
R6#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 192.168.46.4 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/301] via 192.168.46.4, 05:09:56, Ethernet0/0
      6.0.0.0/32 is subnetted, 2 subnets
C        6.6.6.6 is directly connected, Loopback0
C        6.6.66.6 is directly connected, Loopback1
      192.168.46.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.46.0/24 is directly connected, Ethernet0/0
L        192.168.46.6/32 is directly connected, Ethernet0/0

R6#
```

QUESTION 54

Drag and Drop Question

Drag each OSPF state to the correct definition.

init	No information has been received, but Hello packets can still be sent to the neighbor.
loading	A Hello packet is received, but the ID of the receiving router was not included in the Hello packet.
exstart	Each router see its own Router ID in the neighbor field of the Hello packet; there is a DR / BDR election.
full	The routers and their DR and BDR establish a master-slave relationship.
2-Way	Routers exchange DBD packets that describe the contents of the entire link-state database.
down	Based on the information provided by the DBDs, routers send link-state request packets.
exchange	All the router and network LSAs are exchanged and the router databases are synchronized.

Answer:

QUESTION 55
Drag and Drop Question

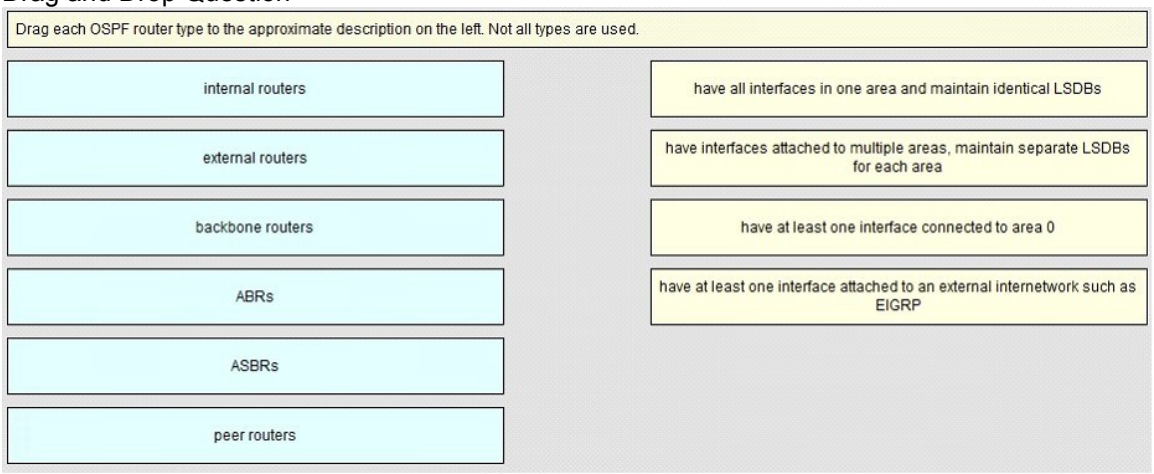

Answer:

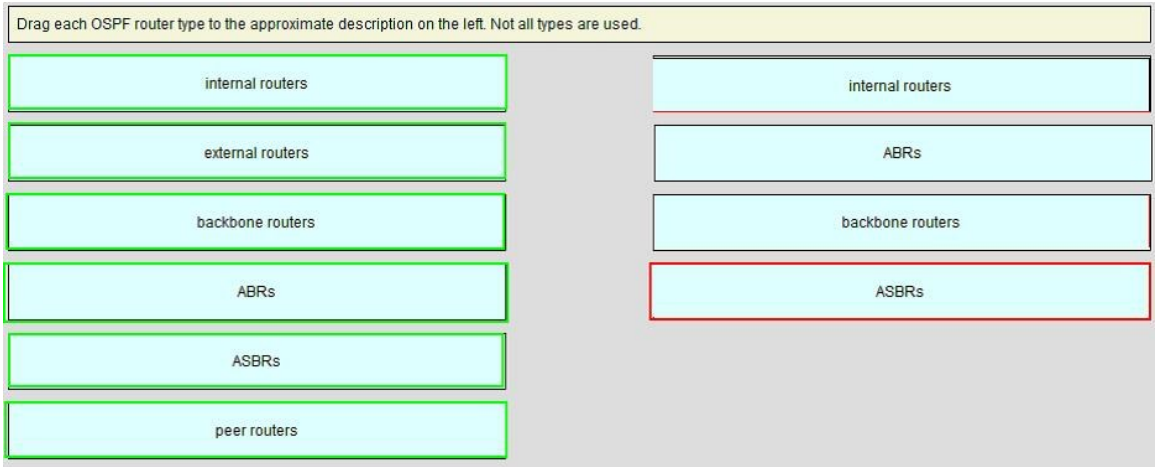

QUESTION 56
Refer to the exhibit. Based on this FIB table, which statement is correct?

```
R2#show ip cef

      Prefix              Next Hop          Interface
0.0.0.0/0             192.168.201.1     FastEthernet0/0
0.0.0.0/32            receive
192.168.201.0/27      attached          FastEthernet0/0
192.168.201.0/32      receive
192.168.201.1/32      192.168.201.1     FastEthernet0/0
192.168.201.2/32      receive
192.168.201.31/32     receive
224.0.0.0/4           drop
224.0.0.0/24          receive
255.255.255.255/32    receive
```

A. There is no default gateway.
B. The IP address of the router on FastEthernet is 209.168.201.1.
C. The gateway of last resort is 192.168.201.1.
D. The router will listen for all multicast traffic.

Answer: C

QUESTION 57
Refer to the exhibit. A network administrator checks this adjacency table on a router.
What is a possible cause for the incomplete marking?

```
Router#show adjacency

 Protocol   Interface              Address
IP          Serial0     192.168.209.130(2)  (incomplete)
IP          Serial0     192.168.209.131(7)
IP          Ethernet0   192.168.201.1(7)
```

A. incomplete ARP information
B. incorrect ACL
C. dynamic routing protocol failure
D. serial link congestion

Answer: A
Explanation:
To display information about the Cisco Express Forwarding adjacency table or the hardware Layer 3-switching adjacency table, use the show adjacency command.
Reasons for Incomplete Adjacencies
There are two known reasons for an incomplete adjacency:
The router cannot use ARP successfully for the next-hop interface.
After a clear ip arp or a clear adjacency command, the router marks the adjacency as incomplete. Then it fails to clear the entry.
In an MPLS environment, IP CEF should be enabeled for Label Switching. Interface level command ip route-cache cef
No ARP Entry
When CEF cannot locate a valid adjacency for a destination prefix, it punts the packets to the CPU for ARP resolution and, in turn, for completion of the adjacency.
http://www.cisco.com/c/en/us/support/docs/ip/express-forwarding-cef/17812-cef-incomp.html#t4

QUESTION 58
A network engineer notices that transmission rates of senders of TCP traffic sharply increase and decrease simultaneously during periods of congestion. Which condition causes this?

A. global synchronization
B. tail drop
C. random early detection
D. queue management algorithm

Answer: A

QUESTION 59
Which three problems result from application mixing of UDP and TCP streams within a network with no QoS? (Choose three.)

A. starvation
B. jitter
C. latency
D. windowing
E. lower throughput

Answer: ACE

QUESTION 60
Which method allows IPv4 and IPv6 to work together without requiring both to be used for a single connection during the migration process?

A. dual-stack method
B. 6to4 tunneling
C. GRE tunneling
D. NAT-PT

Answer: A