


Cisco 642-902



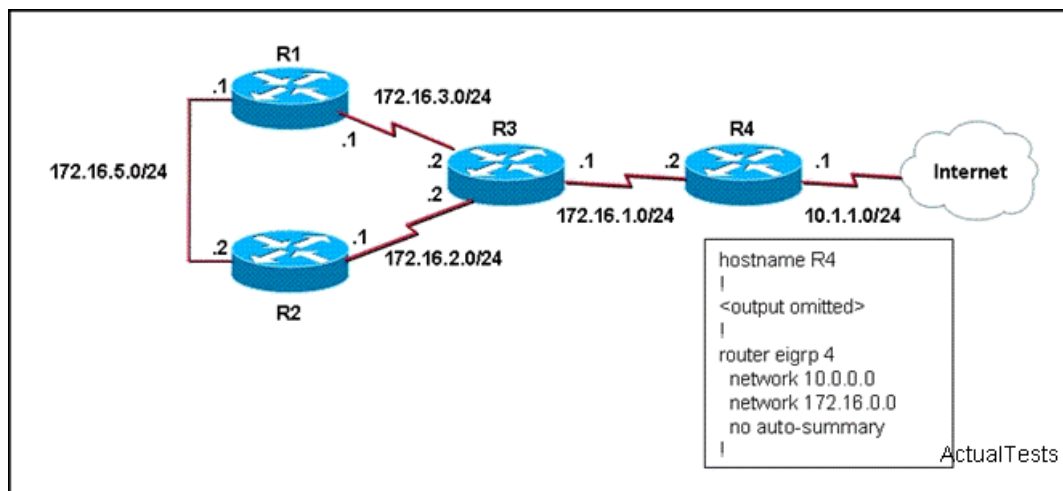
642-902 Implementing Cisco IP Routing (ROUTE)

Practice Test

Version 5.1

QUESTION NO: 1

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

Answer: B

Explanation:

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router.

For every network configured with ip default-network, if a router has a route to that network, that route is flagged as a candidate default route.

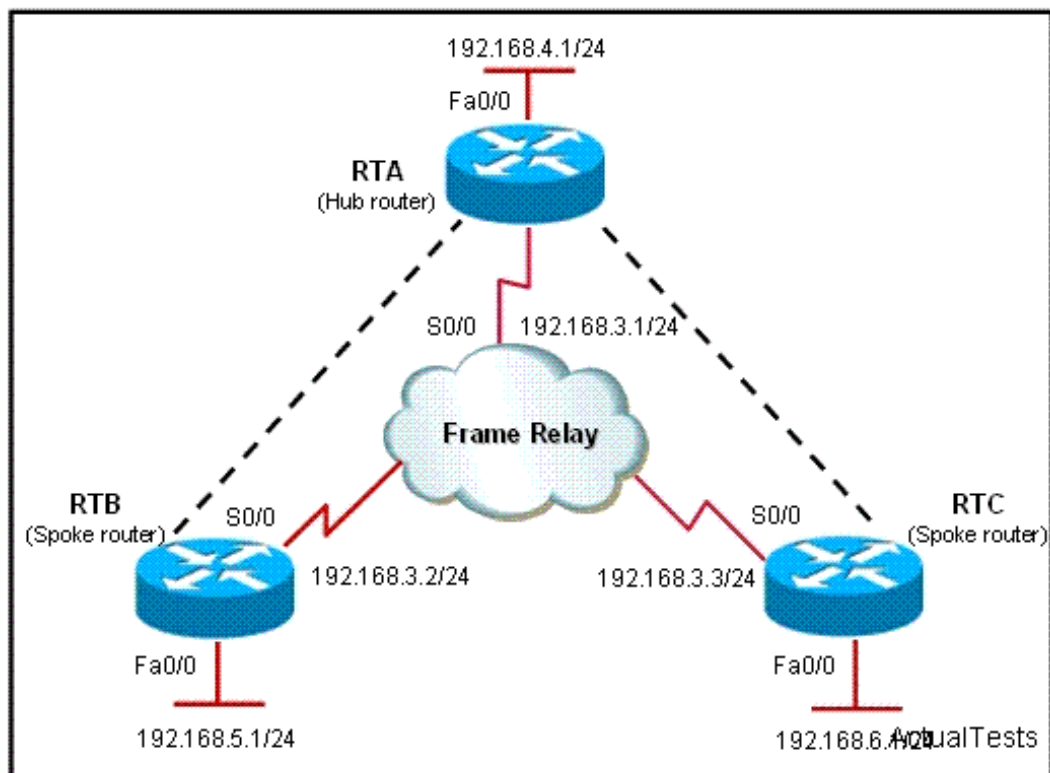
Gateways of last resort selected using the ip default-network command are propagated differently depending on which routing protocol is propagating the default route. For IGRP and EIGRP to propagate the route, the network specified by the ip default-network command must be known to IGRP or EIGRP. This means the network must be an IGRP- or EIGRP-derived network in the routing table, or the static route used to generate the route to the network must be redistributed into IGRP or EIGRP, or advertised into these protocols using the network command. In this case, the 10.0.0.0 network is indeed being advertised via EIGRP.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094374.shtml#ipnetwork

QUESTION NO: 2

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. Check and change the access lists on router RTA.
- B. Configure the auto summary command under router eigrp 1 on router RTA.
- C. Configure subinterfaces on the spoke routers and assign different IP address subnets for each subinterface.
- D. Issue the no ip split horizon command 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.

Answer: E

Explanation:

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, these packets are not sent for destinations for which this interface is the next hop. This reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks (such as Frame Relay and SMDS), situations can arise for which this behavior is less than ideal. For these situations, you may want to disable split horizon. In this example, routes received by RTB and RTC are not being sent back out the same serial interface on RTA, so they are not receiving

each other's routes. Disabling Split horizons on interface S0/0 on RTA will fix this issue.

QUESTION NO: 3

Refer to the exhibit. EIGRP is configured on all routers 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. FastEthernet0/0
```

ActualTests

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

Answer: D

Explanation:

The "show ip eigrp topology" command lists all routes that EIGRP is aware of and shows whether EIGRP is actively processing information on that route. Under most normal conditions, the routes should all be in a passive state and no EIGRP process are running for that route. If the routes are active, this could indicate the dreaded stuck in active, or SIA, state.

The fields to note in this output are as follows: P- Passive; no EIGRP computation is being performed. This is the ideal state. A- Active; EIGRP computations are "actively" being performed for this destination. Routes constantly appearing in an active state indicate a neighbor or query problem. Both are symptoms of the SIA problem. U- Update; an update packet was sent to this destination. Q- Query; a query packet was sent to this destination. R- Reply; a reply packet was sent to this destination. Route information- IP address of the route or network, its subnet mask, and the successor, or next hop to that network, or the feasible successor.

QUESTION NO: 4

Refer to the exhibit. Which three statements are true? (Choose three.)

S 62.99.153.0/24 [1/0] via 209.177.64.130
 172.209.12.0/32 is subnetted, 1 subnets
 D EX 172.209.1
 [170/2590720] via 209.179.2.114, 06:47:28, Serial0/0/0.1239
 62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
 D EX 99.3.215.0/24
 [170/27316] via 209.180.96.45, 09:52:10, FastEthernet11/0/0
 [170/27316] via 209.180.96.44, 09:52:10, FastEthernet11/0/0
 25.248.17.0/24
 [90/1512111] via 209.179.66.25, 10:33:13, Serial0/0/0.1400001
 [90/1512111] via 209.179.66.41, 10:33:13, Serial0/0/0.1402001
 62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
 D 62.113.1.227/32
 [90/2611727] via 209.180.96.45, 10:33:13, FastEthernet1/0/0
 [90/2611727] via 209.180.96.44, 10:33:13, FastEthernet1/0/0
 S* 0.0.0.0/0 [1/0] via 209.180.96.14

- A. On the routing table of R4, the 10.1.1.0/24 route appears as an O E2 route.
- B. On R4, the 172.16.1.0/24 route has a metric of 20.
- C. The R3 S0/0 interface should not need the no ip split-horizon eigrp 1 configuration command for the 172.16.1.0/24 route to appear in the routing table of R2 as an D EX route.
- D. The administrative distance of the 172.16.1.0/24 route in the routing table of R3 is 170.
- E. On R5, the 4.0.0.0/8 route will have an administrative distance of 120 and a hop count of 6.

Answer: A,B,D

QUESTION NO: 5

Which command will display EIGRP packets sent and received, as well as statistics on hello packets, updates, queries, replies, and acknowledgments?

- A. debug eigrp packets
- B. show ip eigrp traffic
- C. debug ip eigrp
- D. show ip eigrp interfaces

Answer: B

Explanation:

The show ip eigrp traffic command displays the number of Enhanced IGRP (EIGRP) packets sent and received.

Example:

The following is sample output from the show ip eigrp traffic command:

Router# show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 77

Hellos sent/received: 218/205

Updates sent/received: 7/23

Queries sent/received: 2/0

Replies sent/received: 0/2

Acks sent/received: 21/14

Reference

:http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800ca5a9.html#wp1018815

QUESTION NO: 6

Which three statements are true about EIGRP operation? (Choose three.) Select 3 response(s).

- A. When summarization is configured, the router will also create a route to null 0.
- B. The summary route remains in the route table, even if there are no more specific routes to the network.
- C. Summarization is configured on a per-interface level.
- D. The maximum metric for the specific routes is used as the metric for the summary route.
- E. Automatic summarization across major network boundaries is enabled by default.

Answer: A,C,E

QUESTION NO: 7

Which two statements about the EIGRP DUAL process are correct? (Choose two.) Select 2 response(s).

- A. An EIGRP route will go active if there are no successors or feasible successors in the EIGRP topology table.
- B. An EIGRP route will go passive if there are no successors in the EIGRP topology table.
- C. DUAL will trigger an EIGRP query process while placing the flapping routes in the holddown state.
- D. A feasible successor in the EIGRP topology table can become the successor only after all the query requests have been replied to.
- E. The stuck in active state is caused when the wait for the query replies have timed out.
- F. EIGRP queries are sent during the loading state in the EIGRP neighbor establishment process.

Answer: A,E

QUESTION NO: 8

What are three key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Choose three.) Select 3 response(s).

- A. A hub router prevents routes from being advertised to the remote router.
- B. Only remote routers are configured as stubs.
- C. Stub routers are not queried for routes.
- D. Spoke routers connected to hub routers answer the route queries for the stub router.
- E. A stub router should have only EIGRP hub routers as neighbors.
- F. EIGRP stub routing should be used on hub routers only.

Answer: B,C,E

QUESTION NO: 9

Based on the exhibited output, which three statements are true? (Choose three.)

```
R1# show ip eigrp topology
IP-EIGRP Topology Table for process 200
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 192.168.1.64/28 1 successors, FD is 281600
   via Connected, Ethernet0
P 192.168.1.32/28 1 successors, FD is 40512000
   via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
   via 192.168.1.66 (40537600/40512000), Ethernet0
   via 192.168.1.17 (41024000/40512000), Serial0
   via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28 1 successors, FD is 40512000
   via Connected, Serial0
```

ActualTests

- A. All the routes are in the passive mode because R1 is in the query process for those routes.
- B. R1 is in AS 200.
- C. R1 will load balance between three paths to reach the 192.168.1.48/28 prefix because all three paths have the same advertised distance (AD) of 40512000.
- D. 40512000 is the advertised distance (AD) via 192.168.1.66 to reach the 192.168.1.48/28 prefix.
- E. All the routes are in the passive mode because these routes are in the hold-down state.
- F. The best path for R1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.

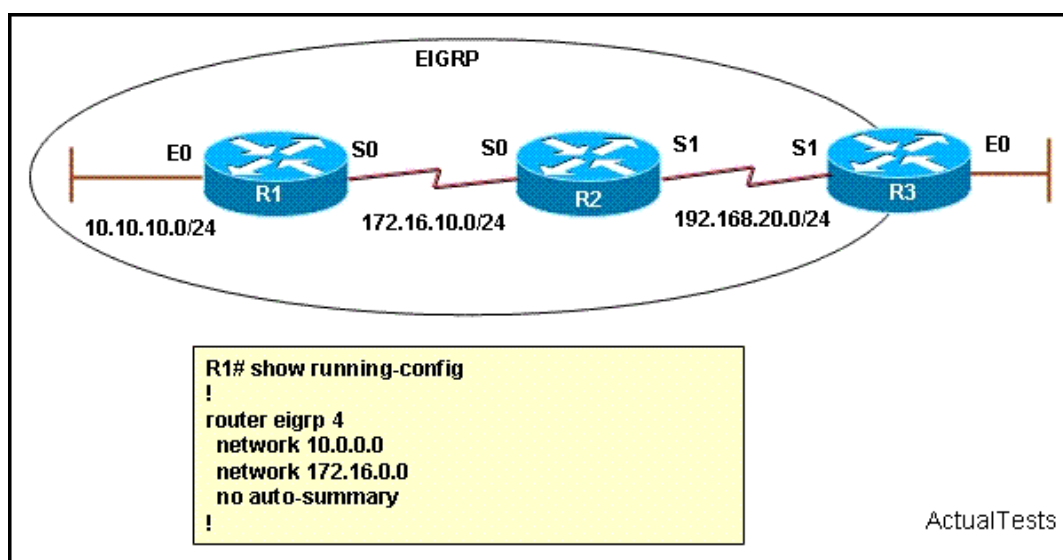
Answer: B,D,F

Explanation:

It can be determined that AS 200 is used, from the fact that the IS-IS process ID is labeled as 200. The best path to reach the network 192.168.1.48/28 is the first one displayed in the routing table. This can be further demonstrated by the fact that the metric is less than the alternative route, via serial 0. Finally, the AD can be found by viewing the second number within the parentheses, which in this case is 40512000.

QUESTION NO: 10

Refer to the exhibit. EIGRP is configured with the default configuration on all routers. Auto summarization is enabled on routers R2 and R3, but it is disabled on router R1. Which two EIGRP routes will be seen in the routing table of router R3? (Choose two.)

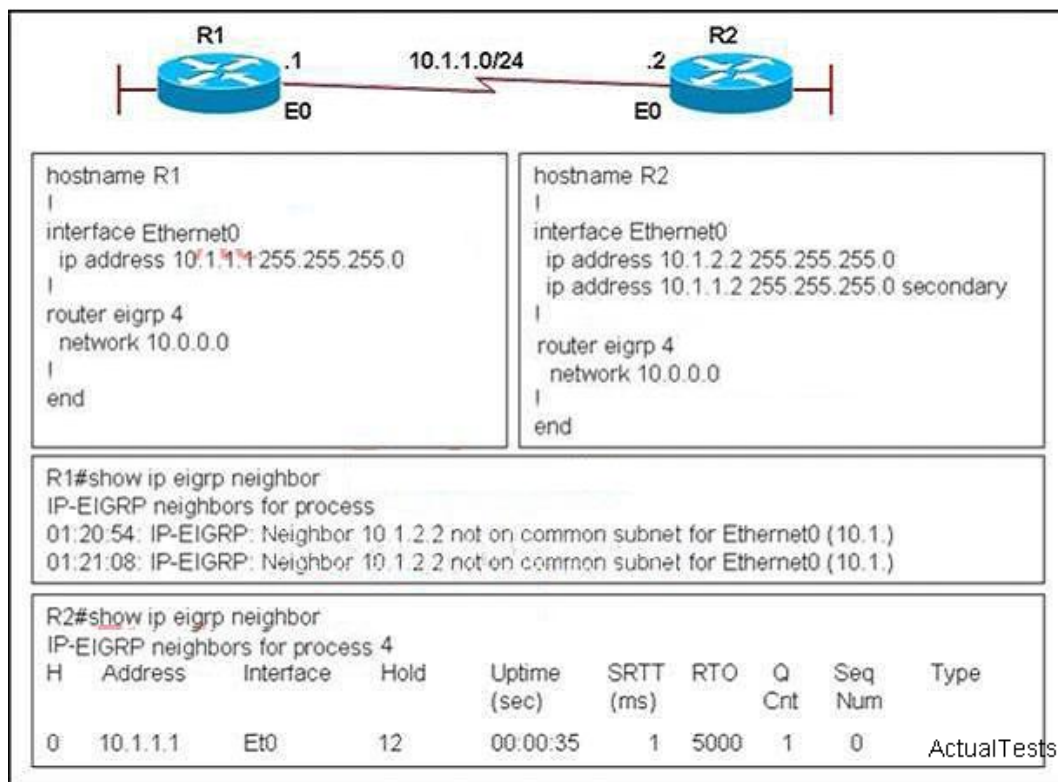


- A. 10.0.0.0/8
- B. 172.16.10.0/24
- C. 172.16.0.0/24
- D. 10.10.10.0/24
- E. 172.16.0.0/16
- F. 10.10.0.0/16

Answer: D,E

QUESTION NO: 11

Refer to the exhibit. EIGRP has been configured on routers R1 and R2. However, R1 does not show R2 as a neighbor and does not accept routing updates from R2. What could be the cause of the problem?



- A. The no auto-summary command has not been issued under the EIGRP process on both routers.
- B. Interface E0 on router R1 has not been configured with a secondary IP address of 10.1.2.1/24.
- C. EIGRP cannot exchange routing updates with a neighbor's router interface that is configured with two IP addresses.
- D. EIGRP cannot form neighbor relationship and exchange routing updates with a secondary address.

Answer: D

QUESTION NO: 12

Which EIGRP packet statement is true?

- A. Update packets route reliable change information only to the affected routers.
- B. Reply packets are multicast to IP address 224.0.0.10 using RTP.
- C. On high-speed links, hello packets are broadcast every 5 seconds for neighbor discovery.
- D. Reply packets are used to send routing updates.
- E. On low-speed links, hello packets are broadcast every 15 seconds for neighbor discovery.

Answer: A

QUESTION NO: 13

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

Answer: A

Explanation:

If the multipoint network has different speeds allocated to the VCs, take the lowest CIR and simply multiply it by the number of circuits. This is because in Frame-relay all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC, so we have to get the lowest speed CIR rate and multiply it by the number of circuits. This result will be applied on the main interface (or multipoint connection interface).

QUESTION NO: 14

Refer to the exhibit. Router R1 is connected to networks 172.16.1.0 /26 and 172.16.1.64 /27. On the basis of the partial output in the exhibit, which statement is true?

```
R1#show running-config
<Output omitted>
!
router eigrp 100
 network 172.16.0.0
 distribute-list prefix TEST out
 auto-summary
 no eigrp log-neighbor-changes
!
ip prefix-list TEST seq 5 permit 172.16.1.0/26
!
<Output omitted>
```

ActualTests

- A. Router R1 will deny the 172.16.1.0/26 route while permitting the 172.16.1.64/27 route to be advertised.
- B. Router R1 will advertise both routes.
- C. Router R1 should be reconfigured with an ACL instead of an ip prefix-list command.
- D. Router R1 will deny the 172.16.1.0/27 route while permitting the 172.16.1.0/26 route to be advertised.

Answer: D

QUESTION NO: 15

Based on the exhibited output, which three statements are true? (Choose three.)

```

R1#show ip route

Gateway of last resort is 10.1.1.2 to network 0.0.0.0

C    1.0.0.0/8 is directly connected, Loopback0
    172.17.0.0/24 is subnetted, 1 subnets
D    172.17.1.0 [90/25632000] via 10.1.1.2, 00:05:20, Serial0/0
    172.16.0.0/24 is subnetted, 1 subnets
D    172.16.1.0 [90/23072000] via 10.1.1.2, 00:05:20, Serial0/0
        [90/20640000] via 10.1.1.3, 00:00:13, Serial0/0
D    172.19.0.0/16 [90/391248640] via 10.1.1.3, 00:05:20, Serial0/0
D    172.22.0.0/16 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D EX 172.25.0.0/16 [170/32032000] via 10.1.1.2, 00:00:10, Serial0/0
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D    10.2.0.0/16 is a summary, 00:06:18, Null0
C    10.2.1.0/24 is directly connected, FastEthernet0/0
C    10.1.1.0/24 is directly connected, Serial0/0
D*EX 0.0.0.0/0 [170/20514560] via 10.1.1.2, 00:00:11, Serial0/0
R1#

```

- A. The route to 10.2.0.0/16 was redistributed into EIGRP.
- B. A default route has been redistributed into the EIGRP autonomous system.
- C. R1 is configured with the ip summary-address command.
- D. R1 is sourcing an external EIGRP route from Null0.
- E. The router at 10.1.1.2 is configured with the ip default-network 0.0.0.0 command.
- F. R1 is configured with the variance command.

Answer: B,C,F

QUESTION NO: 16

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

Answer: B

QUESTION NO: 17

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

- A. EIGRP supports five generic packet types, including hello, database description (DBD), link-state 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 K values 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.

Answer: B,C,F

QUESTION NO: 18

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

Answer: A

Explanation:

If the multipoint network has different speeds allocated to the VCs, take the lowest CIR and simply multiply it by the number of circuits. This is because in Frame-relay all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC, so we have to get the lowest speed CIR rate and multiply it by the number of circuits. This result will be applied on the main interface (or multipoint connection interface).

QUESTION NO: 19

Refer to the exhibit. EIGRP is configured on all routers 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
ActualTests

```

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

Answer: D

Explanation:

The "show ip eigrp topology" command lists all routes that EIGRP is aware of and shows whether EIGRP is actively processing information on that route. Under most normal conditions, the routes should all be in a passive state and no EIGRP process are running for that route. If the routes are active, this could indicate the dreaded stuck in active, or SIA, state.

The fields to note in this output are as follows: P- Passive; no EIGRP computation is being performed. This is the ideal state. A- Active; EIGRP computations are "actively" being performed for this destination. Routes constantly appearing in an active state indicate a neighbor or query problem. Both are symptoms of the SIA problem. U- Update; an update packet was sent to this destination. Q- Query; a query packet was sent to this destination. R- Reply; a reply packet was sent to this destination. Route information- IP address of the route or network, its subnet mask, and the successor, or next hop to that network, or the feasible successor.

QUESTION NO: 20

Which three statements are true about EIGRP route summarization? (Choose three.)

- A. When manual summarization is configured, the router immediately creates a route that points to null0 interface.
- B. Manual route summarization is configured in router configuration mode when the router is configured for EIGRP routing.
- C. Manual route summarization is configured on the interface.
- D. The ip summary-address eigrp command generates a default route with an administrative distance of 5.

- E. When manual summarization is configured, the summary route will use the metric of the largest specific metric of the summary routes.
- F. The ip summary-address eigrp command generates a default route with an administrative distance of 90.

Answer: A,C,D

Explanation:

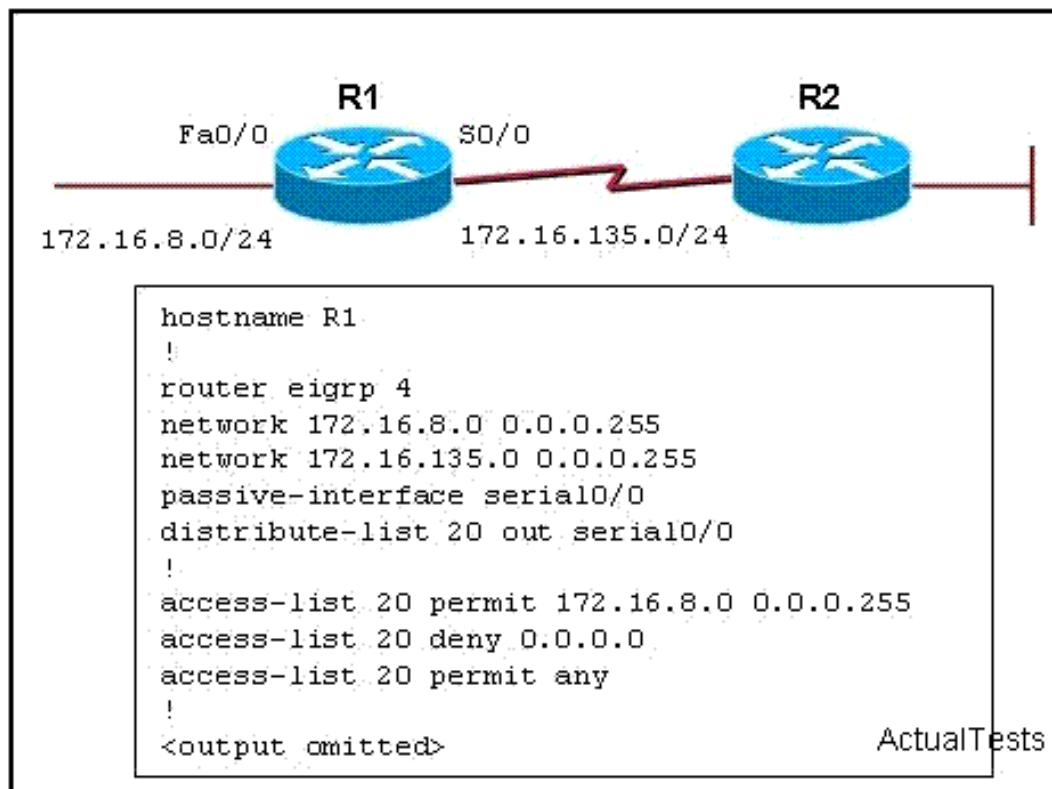
The purpose of route summarization is small routing tables, smaller updates. On major network boundaries, subnetworks are summarized to a single classful network and automatic route summarization is enabled by default. Manual route summarization can be configured on per interface basis. When summarization is configured on an interface, the router immediately creates a route pointing to null0.

Route summarization works in conjunction with the ip summary-address eigrp interface configuration command, in which additional summarization can be performed. If automatic summarization is in effect, there usually is no need to configure network level summaries using the ip summary-address eigrp command. You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Reference: http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1ceigrp.html

QUESTION NO: 21

Refer to the exhibit. Routers R1 and R2 are running EIGRP and have converged. On the basis of the information that is presented, which statement is true?



- A. Both outgoing and incoming routing updates on R1 will be permitted because the distribute-list 20 out Serial0/0 command cannot be used with association with the outgoing interface.
- B. All incoming routing updates from R2 will be suppressed, but the outgoing updates will continue to be sent.
- C. Both outgoing and incoming routing updates on R1 will be stopped because of the passive-interface Serial0/0 configuration statement.
- D. All outgoing routing updates from router R1 to router R2 will be suppressed, but the inbound updates will continue to be received.

Answer: C

Explanation:

You can use the passive-interface command to control the advertisement of routing information. The command enables the suppression of routing updates over some interfaces while it allows updates to be exchanged normally over other interfaces.

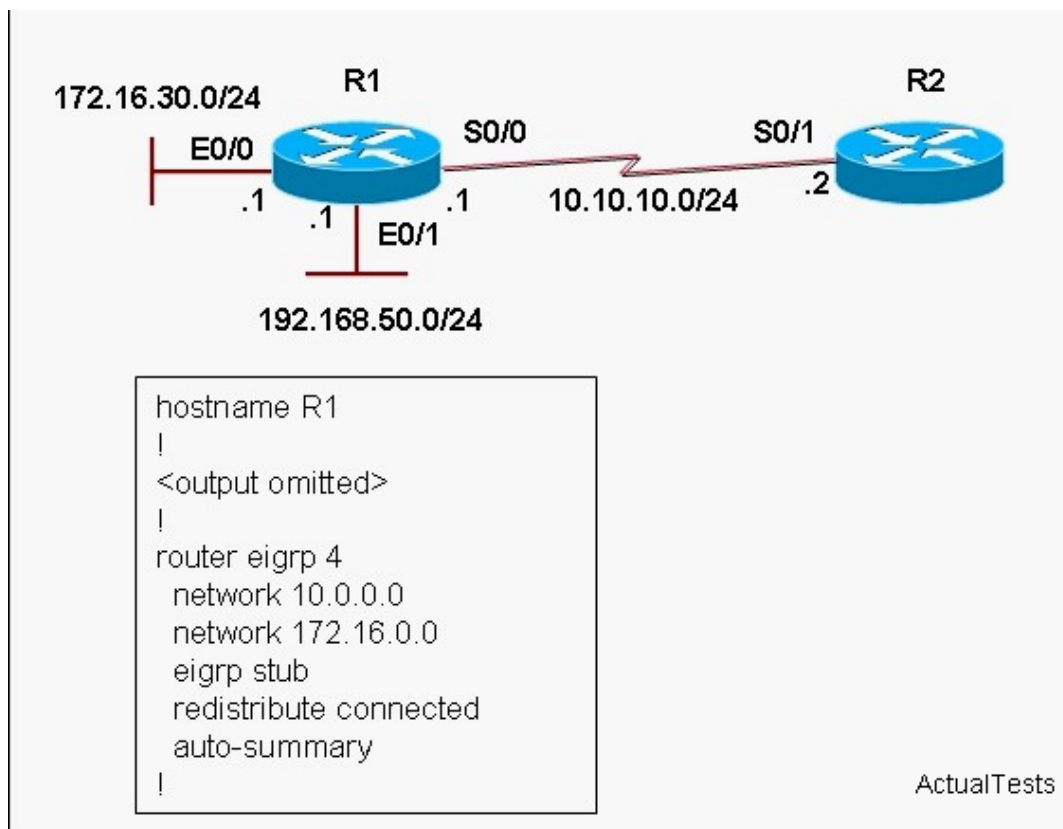
With most routing protocols, the passive-interface command restricts outgoing advertisements only. However, when used with Enhanced Interior Gateway Routing Protocol (EIGRP), the effect is slightly different. With EIGRP running on a network, the passive-interface command stops both outgoing and incoming routing updates, since the effect of the command causes the router to stop sending and receiving hello packets over an interface.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0a.shtml

QUESTION NO: 22

Refer to the exhibit. Which two statements are true? (Choose two.)



- A. Router R1 will advertise connected and summary routes only.
- B. The eigrp stub command will automatically enable summarization of routes on R2.
- C. Router R1 is configured as a receive-only neighbor and will not send any connected, static, or summary routes.
- D. The eigrp stub command prevents all routes except a default route from being advertised to R1.
- E. The eigrp stub command prevents queries from being sent from R2 to R1.
- F. Router R1 will advertise connected and static routes. The sending of summary routes will not be permitted.

Answer: A,E

QUESTION NO: 23

A router has been configured with the commands `router eigrp 9` and `network 172.16.1.0 0.0.0.255`. No other EIGRP-related commands have been configured. The answers list the IP addresses that could be assigned to this router's Fa0/0 interface. Which answers list an IP address/prefix length that would cause the router to enable EIGRP on Fa0/0?

- A. 172.16.0.1/23
- B. 172.16.1.1/26
- C. 172.16.1.1/24

- D. 172.16.0.255/23
- E. None of the other answers is correct.

Answer: B,C

Explanation:

The network 172.16.1.0 0.0.0.255 command tells IOS to match the first three octets when comparing the interface IP addresses to the configured "172.16.1.0" value. Only two answers match in the first three octets. The other two answers have a 0 in the 3rd octet, making the addresses not match the network command.

QUESTION NO: 24

Router R1 has working interfaces S0/0, S0/1, and S0/2, with IP address/prefix combinations of 10.10.10.1/24, 10.10.11.2/24, and 10.10.12.3/22. R1's configuration includes the commands router eigrp 9 and network 10.0.0.0. The show ip eigrp interfaces command lists S0/0 and S0/1 in the command output, but not S0/2. Which answer gives the reason for the omission? (Choose two answers.)

- A. R1 has EIGRP neighbors reachable via S0/0 and S0/1, but not via S0/2, so it is not included.
- B. S0/2 may currently be in a state other than up/up.
- C. The network 10.0.0.0 command requires the use of mask 255.0.0.0 due to EIGRP being classful by default.
- D. S0/2 may be configured as a passive interface.

Answer: B

Explanation:

The show ip eigrp interfaces command displays working (up/up) interfaces on which EIGRP has been enabled but omits passive interfaces. A failure of the interface, or making the interface passive, would omit the interface from the output of this command.

QUESTION NO: 25

Routers R1 and R2 are EIGRP neighbors using their Fa0/0 interfaces, respectively. An engineer adds the ip hello-interval eigrp 9 6 command to R1's Fa0/0 configuration. Which of the following is true regarding the results from this change?

- A. The show ip eigrp neighbors command on R1 lists the revised Hello timer.
- B. The show ip eigrp interfaces command on R1 lists the revised Hello timer.
- C. The R1-R2 neighborship fails due to Hello timer mismatch.
- D. The show ip eigrp interfaces detail command on R1 lists the revised Hello timer.

Answer: D

Explanation:

The show ip eigrp interfaces detail command does display a router's EIGRP Hello timer setting for each enabled interface. The other listed commands do not display the timer. Also, EIGRP routers do not have to have matching Hello timers to become neighbors.

QUESTION NO: 26

Routers R1 and R2, currently EIGRP neighbors over their Fa0/0 interfaces (respectively), both use EIGRP authentication. Tuesday at 8 p.m. the neighborship fails.

Which of the following would not be useful when investigating whether authentication had anything to do with the failure?

- A. debug eigrp packet
- B. show key chain
- C. show ip eigrp neighbor failure
- D. show clock

Answer: C

Explanation:

The show ip eigrp neighbors failure command is not a valid IOS command. The debug displays messages that state when a neighborship fails due to authentication each time Hellos are exchanged. The show key chain command lists the specific keys that are currently valid, allowing you to determine if both routers use the same key values in currently valid keys. The show clock command displays the current time-of-day clock setting on a router, allowing you to check the valid times of the various keys versus the two router's clocks.

QUESTION NO: 27

Router R1 has been configured with the commands router eigrp 9 and network 172.16.2.0 0.0.0.255, with no other current EIGRP configuration. R1's (working) Fa0/0 interface has been configured with IP address 172.16.2.2/26. R1 has found three EIGRP neighbors reachable via interface Fa0/0, including the router with IP address 172.16.2.20. When the engineer attempts to add the neighbor 172.16.2.20 fa0/0 command in EIGRP configuration mode, which of the following occurs?

- A. Fa0/0 fails.
- B. The command is rejected.
- C. The existing three neighbors fail.

- D. The neighborship with 172.16.2.20 fails and then reestablishes.
- E. None of the other answers is correct.

Answer: C

Explanation:

The neighbor 172.16.2.20 fa0/0 command would only be rejected if the IP address (172.16.2.20) is not inside the range of addresses in the subnet (172.16.2.0/26, range 172.16.2.0-172.16.2.63). This command does not impact interface state. The command does disable all EIGRP multicasts, and because the three dynamically discovered neighbors require the EIGRP multicasts, all three neighbors fail. Although 172.16.2.20 is a valid potential neighbor, both routers must be configured with static neighbor commands, and we know that 172.16.2.20 was not previously configured with a static neighbor command; otherwise, it could not have been a neighbor with R1.

QUESTION NO: 28

Which of the following settings could prevent two potential EIGRP neighbors from becoming neighbors? (Choose two answers.)

- A. The interface used by one router to connect to the other router is passive in the EIGRP process.
- B. Duplicate EIGRP router IDs.
- C. Mismatched Hold Timers.
- D. IP addresses of 10.1.1.1/24 and 10.2.2.2/24, respectively.

Answer: A,D

QUESTION NO: 29

An engineer has added the following configuration snippet to an implementation planning document. The configuration will be added to Router R1, whose Fa0/0 interface connects to a LAN to which Routers R2 and R3 also connect. R2 and R3 are already EIGRP neighbors with each other. Assuming the snippet shows all commands on R1 related to EIGRP authentication, which answer lists an appropriate comment to be made during the implementation plan peer review?

```
key chain fred
key 3
key-string whehew
interface fa0/0
ip authentication key-chain eigrp 9 fred
```


- A. The configuration is missing one authentication-related configuration command.
- B. The configuration is missing two authentication-related configuration commands.
- C. Authentication type 9 is not supported; type 5 should be used instead.
- D. The key numbers must begin with key 1, so change the key 3 command to key 1.

Answer: A

Explanation:

The configuration requires the `ip authentication mode eigrp asn md5` command, which is currently missing. This command enables MD5-style authentication, rather than the default of no authentication. Adding this one command completes the configuration. Any valid key numbers can be used. Also, the 9 in the `ip authentication key-chain eigrp 9 fred` command refers to the EIGRP ASN, not an authentication type.

QUESTION NO: 30

Which of the following are methods EIGRP uses to initially populate (seed) its EIGRP topology table, before learning topology data from neighbors? (Choose two.)

- A. By adding all subnets listed by the `show ip route connected` command
- B. By adding the subnets of working interfaces over which static neighbors have been defined
- C. By adding subnets redistributed on the local router from another routing source
- D. By adding all subnets listed by the `show ip route static` command

Answer: B,C

Explanation:

Other than the two listed correct answers, the local router also adds connected routes for which the network command matches the corresponding interfaces, so it may not add all connected routes. Also, EIGRP does not add static routes to the EIGRP topology table, unless those routes are redistributed.

QUESTION NO: 31

Which of the following are both advertised by EIGRP in the Update message and included in the formula for calculating the integer EIGRP metric? (Choose two.)

- A. Jitter
- B. Delay
- C. MTU
- D. Reliability

Answer: B,D

Explanation:

EIGRP sends bandwidth, delay, reliability, load, MTU, and hop-count in the message. The formula to calculate the metric includes bandwidth, delay, reliability, and load.

QUESTION NO: 32

Router R1 uses S0/0 to connect via a T/1 to the Frame Relay service. Five PVCs terminate on the serial link. Three PVCs (101, 102, and 103) are configured on subinterface S0/0.1, and one each (104 and 105) are on S0/0.2 and S0/0.3. The configuration shows no configuration related to EIGRP WAN bandwidth control, and the bandwidth command is not configured at all. Which of the following is true about how IOS tries to limit EIGRP's use of bandwidth on S0/0?

- A. R1 limits EIGRP to around 250Kbps on DLCI 102.
- B. R1 limits EIGRP to around 250Kbps on DLCI 104.
- C. R1 limits EIGRP to around 150Kbps on every DLCI.
- D. R1 does not limit EIGRP because no WAN bandwidth control has been configured.

Answer: A

Explanation:

EIGRP performs WAN bandwidth control without any explicit configuration, using default settings. Because no bandwidth commands have been configured, each subinterface uses the default 1544 Kbps setting. For S0/0.1, WAN bandwidth control divides the 1544 by 3 (515 Kbps), and then takes the (default) WAN bandwidth of 50 percent, meaning about 250 Kbps for each of the three DLCIs. For the two subinterfaces with one PVC, the default 1544 is multiplied by the 50 percent default WAN bandwidth, meaning that each could use about 750 Kbps.

QUESTION NO: 33

The output of show ip eigrp topology on Router R1 shows the following output, which is all the output related to subnet 10.11.1.0/24. How many feasible successor routes does R1 have for 10.11.1.0/24?

```
P 10.11.1.0/24, 2 successors, FD is 2172419
via 10.1.1.2 (2172423/28167), Serial0/0/0.1
via 10.1.1.6 (2172423/28167), Serial0/0/0.2
```

A. 0

- B. 1
- C. 2
- D. 3

Answer: A

Explanation:

This command lists all successor and feasible successor routes. The output states that two successors exist, and only two routes (listed with the "via..." text) exist. So, no feasible successor routes exist.

QUESTION NO: 34

A network design shows that R1 has four different possible paths from itself to the Data Center subnets. Which of the following can influence which of those routes become feasible successor routes, assuming that you follow the Cisco recommended practice of not changing metric weights? (Choose two.)

- A. The configuration of EIGRP offset lists
- B. Current link loads
- C. Changing interface delay settings
- D. Configuration of variance

Answer: A,C

Explanation:

By default, the metric weights cause EIGRP to consider bandwidth and delay in the metric calculation, so changing either bandwidth or delay impacts the calculation of the feasible distance and reported distance, and impacts choice of feasible successor routes. Offset lists also change the metric, which in turn can change whether a route is an FS route. Link loading would impact the metrics, but not without changing the metric weights to nonrecommended values. Finally, variance impacts which routes end up in the IP routing table, but it is not considered by EIGRP when determining which routes are FS routes.

QUESTION NO: 35

Router R1 is three router hops away from subnet 10.1.1.0/24. According to various show interfaces commands, all three links between R1 and 10.1.1.0/24 use the following settings: bandwidth: 1000, 500, 100000 and delay: 12000, 8000, 100. Which of the following answers correctly identifies a value that feeds into the EIGRP metric calculation? (Choose two correct answers.)

- A. Bandwidth of 101,500
- B. Bandwidth of about 34,000
- C. Bandwidth of 500
- D. Delay of 1200
- E. Delay of 2010
- F. Delay of 20100

Answer: C,E

Explanation:

The EIGRP metric calculation treats bandwidth and delay differently. For bandwidth, EIGRP takes the lowest bandwidth, in Kbps, which is in this case 500 Kbps. For delay, EIGRP takes the cumulative delay, which is 20100 per the various show interfaces commands. However, the show interfaces command uses a unit of microseconds, and the interface delay command, and the EIGRP metric formula uses a unit of tens-of-microseconds, making the delay that feeds into the formula be 2010.

QUESTION NO: 36

Routers R1 and R2 are EIGRP neighbors. R1 has been configured with the eigrp stub connected command. Which of the following is true as a result? (Choose two correct answers.)

- A. R1 can learn EIGRP routes from R2, but R2 cannot learn EIGRP routes from R1.
- B. R1 can send IP packets to R2, but R2 cannot send IP packets to R1.
- C. R2 no longer learns EIGRP routes from R1 for routes not connected to R1.
- D. R1 no longer replies to R2's Query messages.
- E. R2 no longer sends to R1 Query messages.

Answer: C,E

Explanation:

R1, as a stub router with the connected option, still advertises routes, but only routes for connected subnets. R1 announces its stub attribute to R2, so R2 chooses to not send Query messages to R1, knowing that R1 cannot be a transit router for other subnets anyway.

QUESTION NO: 37

EIGRP performs automatic summarization at network boundaries. What administrative distance is given to these EIGRP summary routes?

- A. 1
- B. 90

- C. 95
- D. 0
- E. 5
- F. 170
- G. 255
- H. None of the other alternatives apply

Answer: E

QUESTION NO: 38

While troubleshooting an EIGRP routing problem you notice that one of the routers have generated a large number of SIA messages. What are two possible causes for EIGRP Stuck-In-Active routes? (Select two)

- A. Some query or reply packets are lost between the routers.
- B. The neighboring router starts receiving route updates from this router.
- C. A failure causes traffic on a link between two neighboring routers to flow in only one direction (unidirectional link).
- D. The neighboring router stops receiving ACK packets from this router.

Answer: A,C

QUESTION NO: 39

Router R1 is configured as shown below:

```
router eigrp 100
network 10.0.0.0
eigrp stub
```

Based on the information shown above, which two types of routes will be advertised? (Select two)

- A. Receive-only
- B. Stub
- C. Static
- D. Summary
- E. Connected
- F. Dynamic

Answer: D,E

QUESTION NO: 40

Into which two types of areas would an area border router (ABR) inject a default route? (Choose two.)

- A. stub
- B. the autonomous system of an exterior gateway protocol (EGP)
- C. NSSA
- D. totally stubby
- E. the autonomous system of a different interior gateway protocol (IGP)
- F. area 0

Answer: A,D

QUESTION NO: 41

Refer to the output. What IOS command produces this output?

Select the best response

```
Routing Process "ospfv3 1" with ID 172.16.3.3
It is an autonomous system boundary router
Redistributing External Routes from,
  static
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Number of external LSA 1. Checksum Sum 0x218D
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Area 1
    Number of interfaces in this area is 2
    SPF algorithm executed 9 times
    Number of LSA 15. Checksum Sum 0x67581
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

ActualTests

- A. show ip ospf
- B. show ip ospf interface
- C. show ipv6 ospf interface
- D. show ipv6 ospf

Answer: D

QUESTION NO: 42

Which command displays the number of times that the OSPF Shortest Path First (SPF) algorithm has been executed?

- A. show ip protocol
- B. show ip ospf
- C. show ip ospf database
- D. show ip ospf interface

Answer: B

Explanation:

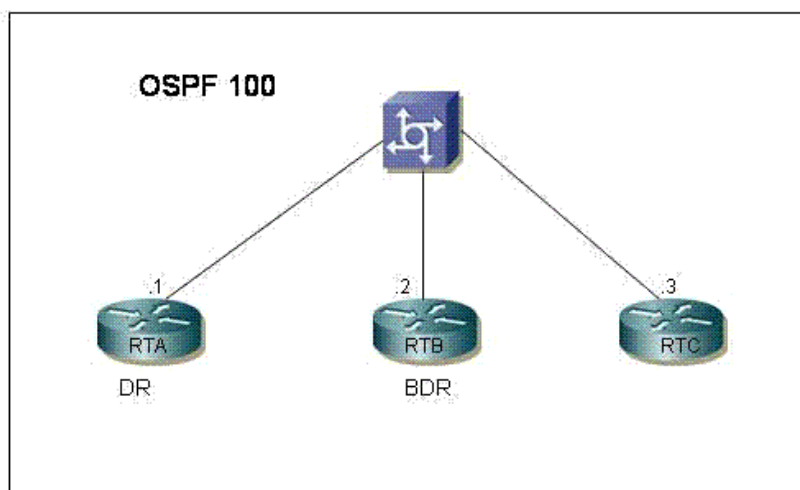
The following table describes the output of the "show ip ospf" command and their meanings:

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469

QUESTION NO: 43

During a recent OSPF election among three routers, RTA was elected the DR and RTB was elected the BDR, as seen in the graphic. Assume that RTA fails, and that RTB takes the place of the DR while RTC becomes the new BDR. What will happen when RTA comes back online?



ActualTests

- A. A new election will take place establishing an all new DR and BDR based on configured priority levels and MAC addresses.
- B. RTA will take the place of DR immediately upon establishing its adjacencies.
- C. RTA will take the place of DR only if RTB fails.
- D. RTA will take the place of DR only if both RTB and RTC fail.

Answer: D

Explanation:

If a router with a higher priority value gets added to the network, it does not preempt the DR and BDR. The only time a DR and BDR changes is if one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new BDR is elected. In a multi-access network, the router that is powered on first will generally become the DR, since the DR/BDR process is not pre-emptive.

Reference: : CCNP Self-Study Second Edition P.243

QUESTION NO: 44

The Dev-1 and Dev-3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Dev-1 and Dev-3 routers, which statement is true?

Dev-1#sh ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
10.200.200.13	1	FULL/BDR	00:00:33	10.1.1.3	Ethernet0/0

Dev-3#sh ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
172.31.1.1	2	FULL/DR	00:00:31	10.1.1.1	Ethernet0/0

- A. Dev-1 is the DR because it has a lower OSPF router ID.
- B. Dev-3 is the DR because it has a higher OSPF router priority.
- C. Both Dev-1 and Dev-3 are using the default OSPF router priority.
- D. Dev-3 is the DR because it has a lower OSPF router ID.
- E. Dev-1 is the DR because it has a higher OSPF router priority.


Answer: E

Explanation:

The output shown above displays information about the neighbors, so from Dev-3 we see that the priority of Dev-1 is 2, and that Dev-1 is the DR. From Dev-1 we can see that its neighbor (Dev-3) is the BDR and has an OSPF priority of 1.

QUESTION NO: 45

Refer to the diagram. Which OSPF configuration command is required to configure Area 3 as a totally stubby area?



- A. On the ABR
router ospf 10
area 3 stub
- B. On all Area 3 routers
router ospf 10
area 3 stub
- C. On all Area 3 routers
router ospf 10
area 3 stub no-summary
- D. On the ABR
router ospf 10
area 3 nssa
- E. On the ABR
router ospf 10
area 3 stub no-summary

Answer: E

QUESTION NO: 46

Which show command will display only the Type 5 LSAs in the OSPF topology database?

- A. show ip ospf database external
- B. show ip ospf database nssa-external
- C. show ip route ospf
- D. show ip ospf database summary
- E. show ip route

Answer: A

QUESTION NO: 47

Which three are advantages to creating multiple areas in OSPF? (Choose three.) Select 3 response(s).

- A. less frequent SPF calculations
- B. fewer hello packets
- C. smaller routing tables
- D. reduced LSU overhead
- E. fewer adjacencies needed

Answer: A,C,D

QUESTION NO: 48

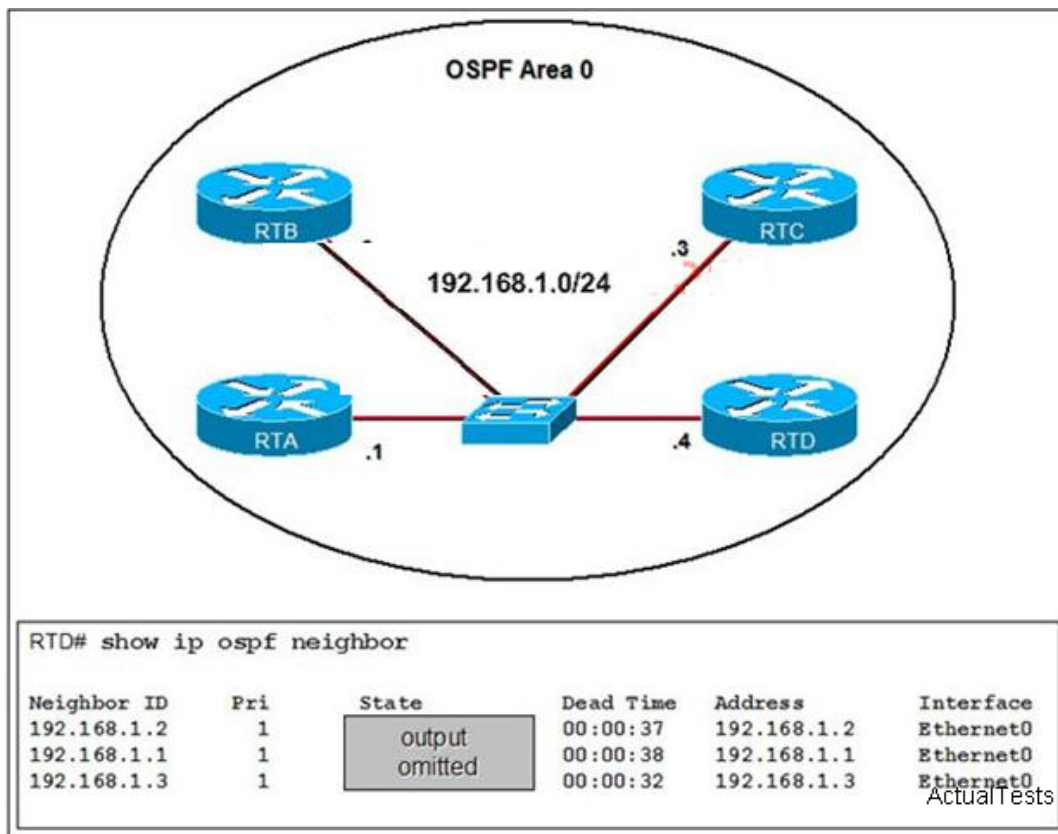
What are two Cisco IOS commands that can be used to view neighbor adjacencies? (Choose two.)

- A. show ip ospf interfaces
- B. show ip ospf protocols
- C. show ip ospf neighbors
- D. show ip ospf database

Answer: A,C

QUESTION NO: 49

DR (Designated Router) is for environments where many routers on the same network such as Ethernet. In the following presented network, all routers are reloaded simultaneously, and DR is selected as expected. What is the RTC status?



- A. 2WAY/DR
- B. FULL/DR
- C. FULL/DROTHER
- D. FULL/BDR

Answer: D

Explanation:

The point of this question is about the select principles of DR and BDR.

DR and BDR election is done via the Hello protocol. Hello packets are exchanged via IP multicast packets (Appendix B) on each segment. The router with the highest OSPF priority on a segment will become the DR for that segment. The same process is repeated for the BDR. In case of a tie, the router with the highest RID will win. The default for the interface OSPF priority is one.

Remember that the DR and BDR concepts are per multiaccess segment. Setting the ospf priority on an interface is done using the ip ospf priority <value> interface command. In this case, all routers have the same priority, but RTD has the highest RID, and RTC was followed, so RTC was the BDR.


QUESTION NO: 50

Refer to the exhibit. During the process of configuring a virtual link to connect area 2 with the backbone area, the network administrator received this console message on R3:

*Mar 1 00:25:01.084: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from

backbone area must be virtual link but not found from 20.20.20.1, Serial 0

How should the virtual link be configured on the OSPF routers to establish full connectivity between the areas?

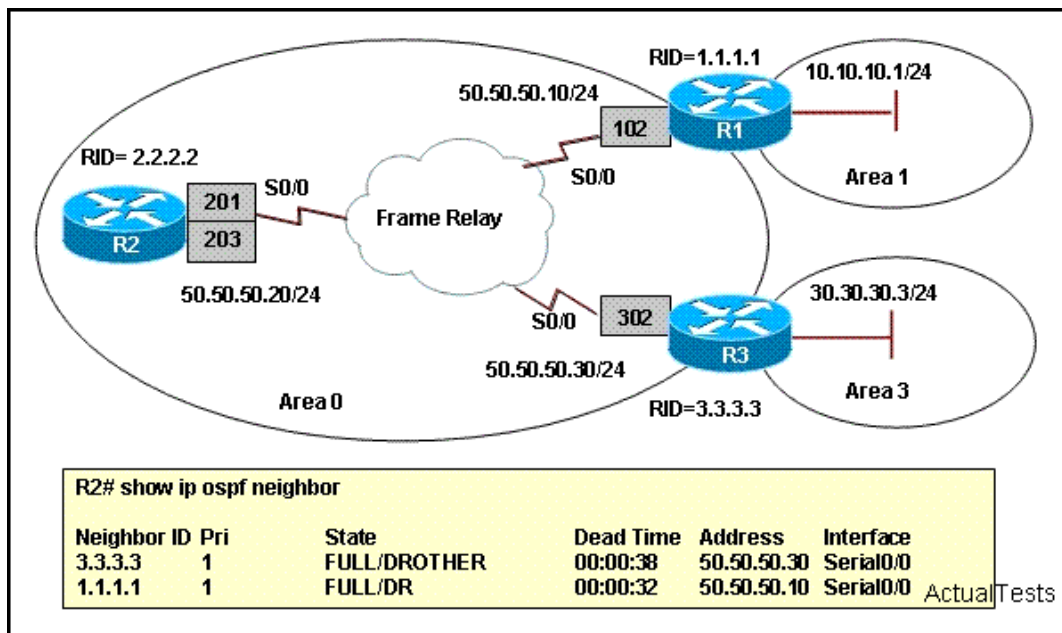


- A. R1(config-router)# area 0 virtual-link 1.1.1.1
R3(config-router)# area 2 virtual-link 3.3.3.3
- B. R1(config-router)# area 1 virtual-link 2.2.2.2
R3(config-router)# area 1 virtual-link 2.2.2.2
- C. R1(config-router)# area 1 virtual-link 20.20.20.2
R3(config-router)# area 1 virtual-link 30.30.30.2
- D. R1(config-router)# area 1 virtual-link 30.30.30.3
R3(config-router)# area 1 virtual-link 20.20.20.1
- E. R1(config-router)# area 1 virtual-link 3.3.3.3
R3(config-router)# area 1 virtual-link 1.1.1.1

Answer: E

QUESTION NO: 51

OSPF is configured over a Frame Relay network as shown in the exhibit. All PVCs are active. However, R1 and R3 fail to see all OSPF routes in their routing tables. The show ip ospf neighbor command executed on R2 shows the state of the neighbors. What should be done to fix the problem?



- A. The ip ospf network broadcast command should be configured on each Frame Relay interface.
- B. The ip ospf network non-broadcast command should be configured on each Frame Relay interface.
- C. The ip ospf priority value on the spoke routers should be set to 0.
- D. The ip ospf priority value on the hub router should be set to 0.
- E. The neighbor command should be configured under the OSPF routing process on all routers.

Answer: C

QUESTION NO: 52

Open Shortest Path First (OSPF) is a routing protocol developed for Internet Protocol (IP) networks by the Interior Gateway Protocol (IGP) working group of the Internet Engineering Task Force (IETF). Which two statements best describe the OSPF link-state routing protocol? (Choose two.)

- A. OSPF sends updates every 10 seconds.
- B. OSPF sends triggered updates when a network change occurs.
- C. When a link state is changed, the router that detected the change creates a link-state advertisement (LSA) and propagates it to all OSPF devices using the 224.0.0.6 multicast address.
- D. OSPF sends summaries of individual link-state entries every 30 minutes to ensure LSDB synchronization.

Answer: B,D

Explanation:

The point of this question is the basis of OSPF.

Incorrect answer A. OSPF send hello packets every 10 seconds, not the updates, OSPF sends triggered updates when a network change occurs. For OSPF, D Rother use the multicast address

224.0.0.6 to send packets to DR and BDR, only DR and BDR can get the information from this multicast address.

QUESTION NO: 53

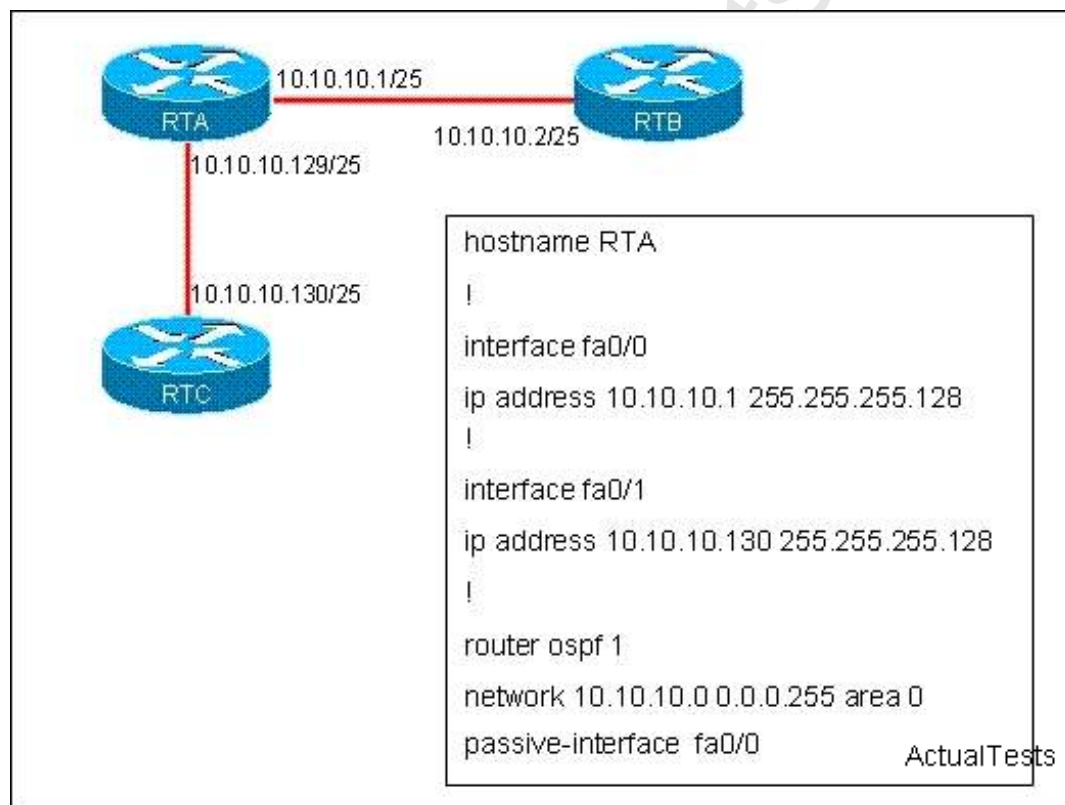
An administrator types in the command `router ospf 1` and receives the error message: "OSPF process 1 cannot start." (Output is omitted.) What should be done to correctly set up OSPF?

- A. Ensure that an interface has been configured with an IP address.
- B. Ensure that IP classless is enabled.
- C. Ensure that the interfaces can ping their directly connected neighbors.
- D. Ensure that an interface has been configured with an IP address and is up.

Answer: D

QUESTION NO: 54

Refer to the exhibit. Which statement is true about the configuration?



- A. RTA will not establish an OSPF adjacency with RTB.
- B. RTA will send OSPF hello packets, but will not send OSPF updates.
- C. RTA will not accept OSPF hello packets from RTB.
- D. RTA will send OSPF updates, but will not establish an adjacency with RTB.

Answer: A

QUESTION NO: 55

Refer to the exhibit. Which statement is true?

RTA# **show ip ospf neighbor**

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.45.1	1	FULL/DR	00:00:36	10.0.0.2	Ethernet0
192.168.45.2	0	2-WAY/DROTHER	00:00:12	10.0.0.1	Ethernet1


ActualTests

- A. Router RTA is the BDR.
- B. Router RTA is directly connected to interface 192.168.45.1.
- C. Neighbor 192.168.45.1 has changed its OSPF priority number.
- D. Router RTA and neighbor 192.168.45.2 are exchanging OSPF LSAs.

Answer: A

QUESTION NO: 56

Refer to the exhibit. OSPF has been configured on all routers in the network and Area 1 has been configured as a NSSA. Which statement is true about the NSSA Area 1?



- A. Only redistributed RIP routes will appear in Area 1. They will be advertised via type 7 LSAs.
- B. No redistributed routes can appear in Area 1, only summary routes.
- C. Redistributed RIP and IGRP routes will appear in Area 1. They will be advertised via type 5 LSAs.
- D. Only redistributed IGRP routes will appear in Area 1. They will be advertised via type 7 LSAs.

Answer: D

QUESTION NO: 57

According to RFC 2328, what is the stateful order in which an OSPF router transitions to a full adjacency with a neighbor router?

- A. Down, Init, 2-Way, Exstart, Exchange, Loading, and Full
- B. Down, Init, 2-Way, Exchange, Exstart, Loading, and Full
- C. Down, 2-Way, Init, Loading, Exstart, Exchange, and Full
- D. Down, 2-Way, Init, Exchange, Exstart, Loading, and Full
- E. Down, Init, 2-Way, Loading, Exstart, Exchange, and Full
- F. Down, 2-Way, Init, Exstart, Exchange, Loading, and Full

Answer: A

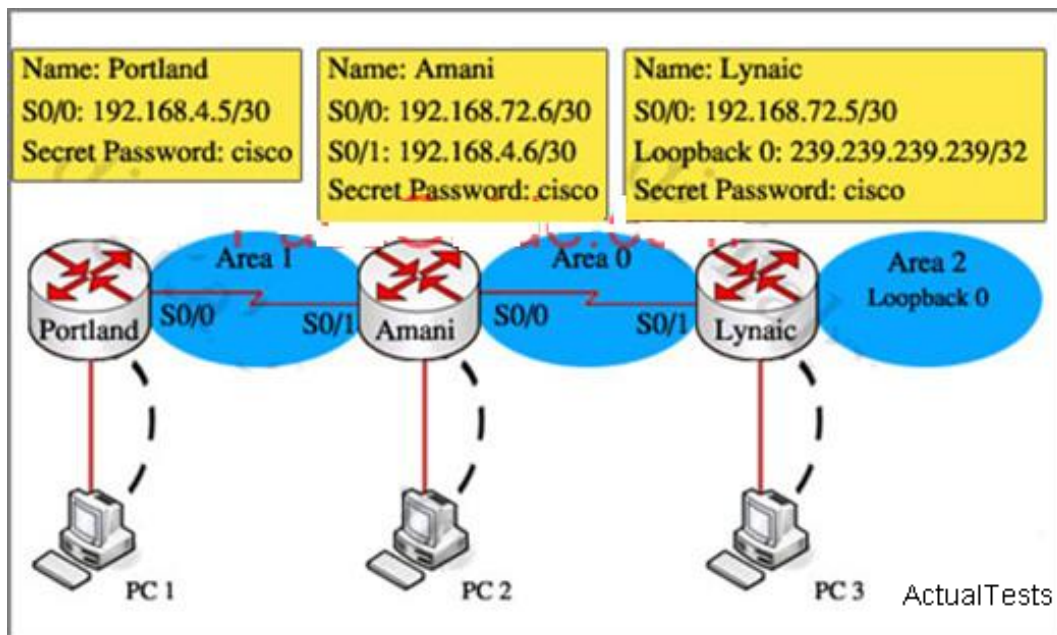
QUESTION NO: 58

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).

**Explanation:**

First, we configure Portland's S0/0 interface so that it belongs to Area 1. So, we have to find out which sub-network 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 NO: 59

A network administrator has enabled OSPF across an NBMA network and has issued the command `ip ospf network nonbroadcast`. Given those facts, which two statements are true? (Choose two.)

- A. Interfaces will automatically detect and build adjacencies with neighbor routers.
- B. DR and BDR elections will not occur.
- C. DR and BDR elections will occur.
- D. The neighbor command is required to build adjacencies.
- E. All routers must be configured in a fully meshed topology with all other routers.

Answer: C,D

Explanation:

Even if there is only one router, broadcast multiaccess networks elect a DR and a BDR to serve as focal points for routing information. In contrast, point-to-point OSPF networks do not elect a DR because they can never include more than two nodes.

Another type of OSPF network, Nonbroadcast Multiaccess (NBMA), can include more than two nodes. Therefore, NBMA will try to elect a DR and a BDR. Common NBMA implementations include Frame Relay, X.25, and SMDS. NBMA networks follow rules at Layer 2 that prevent the delivery of broadcasts and multicasts.

A non-broadcast environment requires that all OSPF neighbors be manually configured. This is the default setting for physical interfaces with Frame Relay encapsulation, as well as for their point-to-multipoint subinterfaces. By manually configuring each neighbor, OSPF knows exactly which neighbors need to participate and which neighbor is identified as the DR. Also, communication between neighbors is done via unicast instead of multicast. This configuration also requires a full mesh and has the same weakness as the broadcast environment.

For non-broadcast networks the default Hello interval is 30 seconds and the Dead interval is four times the Hello interval, 120 seconds. Non-broadcast multi-access networks do elect a DR and BDR, due to their multi-access nature. In order to set which router you want as the DR, you must set the priority in the neighbor statement to elect the neighbor as the DR. In order to manually configure who your neighbors are, the following command must be entered in router configuration mode for the selected OSPF process: `neighbor ip_address ip_address` = the ip address of the neighbor.

If you would like to set the priority of this router to become the DR, you will need to append the priority of the neighbor:

```
neighbor ip_address priority value
```

QUESTION NO: 60

Which of the following is an OSPF configuration parameter that is used on an ABR, but not on an internal router?

- A. A virtual link to area 0.
- B. OSPF summarization command.
- C. default-cost extension to the area command.
- D. no-summary extension to the area stub command.
- E. None of the other alternatives apply

Answer: D

QUESTION NO: 61

Refer to the exhibit. Which two statements are true about the partial configuration that is provided. (Choose two.)

```
router bgp 100
 neighbor internal peer-group
 neighbor internal remote-as 100
 neighbor internal update-source loopback 0
 neighbor internal route-map set-med out
 neighbor internal filter-list 1 out
 neighbor internal filter-list 2 in
 neighbor 171.69.232.53 peer-group internal
 neighbor 171.69.232.54 peer-group internal
 neighbor 171.69.232.55 peer-group internal
 neighbor 171.69.232.55 filter-list 3 in
```

- A. The peer group shortens the EBGP configuration.
- B. The peer group shortens the IBGP configuration.
- C. All the configured neighbors are in autonomous system 100.
- D. Three AS-path filters are applied to each BGP neighbor.
- E. Only the outgoing filters are applied to BGP updates.

Answer: B,C

QUESTION NO: 62

Refer to the exhibit. Which two statements are correct? (Choose two.)

```
Router# show ip bgp

BGP table version is 5, local router ID is 10.0.33.34
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.1.0.0         0.0.0.0           0             32768 ?
* 10.2.0.0         10.0.33.35        10            0 35 ?
*>                 0.0.0.0           0             32768 ?
* 10.0.0.0         10.0.33.35        10            0 35 ?
*>                 0.0.0.0           0             32768 ?
*> 192.168.0.0/16  10.0.33.35        10            0 35 ?
```

- A. All the routes were redistributed into BGP from an IGP.
- B. All the routes were originated by BGP with the network command.
- C. All six routes will be installed in the routing table.
- D. Four routes will be installed in the routing table.
- E. Two routes will be installed in the routing table.

Answer: A,D

QUESTION NO: 63

A router has two paths to reach another network in a different autonomous system. Neither route was generated by the local router and both routes have the same default weight and local preference values. Which statement is true about how BGP would select the best path?

- A. To influence one route to be preferred, its default local preference value will be changed via the use of the command `bgp default local-preference 50`.
- B. The router will prefer the shortest autonomous system path.
- C. If the command `bgp always-compare-med` has been given, then the router will prefer the route with the highest MED.
- D. The router will prefer the route with the lower MED.

Answer: B

Explanation:

The shortest AS path is the fourth item checked by BGP routers, while the MED values are one of the last items looked at by BGP routers. The following is a partial list of rules used by BGP to determine the best path. Prefer the path with the highest WEIGHT.

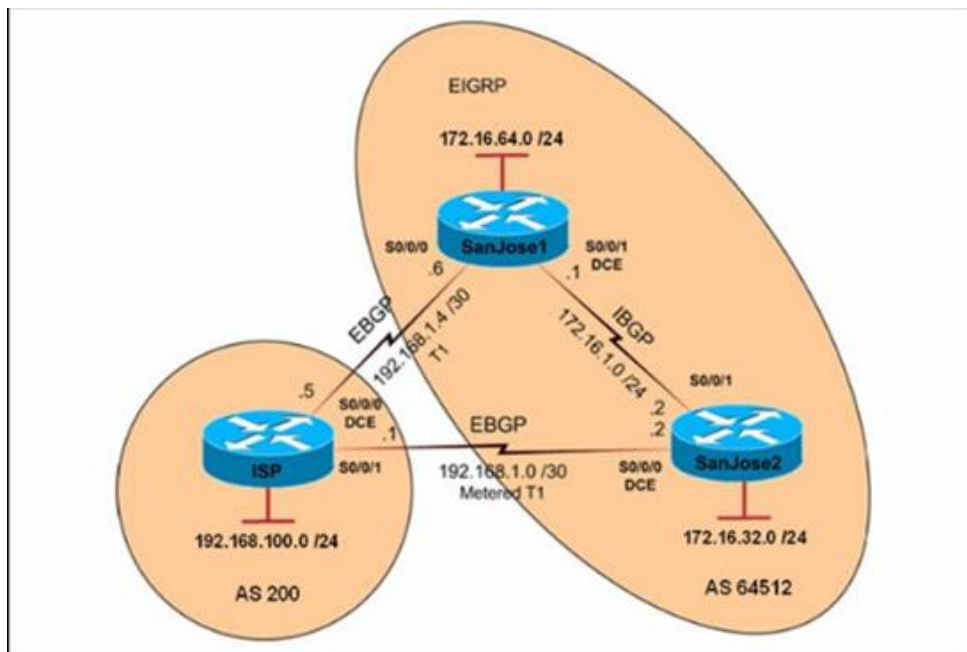
Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured. Prefer the path with the highest LOCAL_PREF. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command. Prefer the path with the shortest AS_PATH. Note the following: Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE. Prefer the path with the lowest multi-exit discriminator (MED).

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml

QUESTION NO: 64

Refer to the exhibit. On the basis of the information in the exhibit, which two statements are true? (Choose two.)



```

BGP table version is 12, local router ID is 192.168.100.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal Origin codes: i - IGP, e - EGP, ? - incomplete
  Network          Next Hop          Metric LocPrf Weight Path
* 172.16.0.0       192.168.1.2       75      0 64512 i
*> 172.16.0.0       192.168.1.6       50      0 64512 i
*> 192.168.1.0/30 0.0.0.0           0          32768 i
*> 192.168.1.4/30 0.0.0.0           0          32768 i
*> 192.168.100.0  0.0.0.0           0          32768 i

```

- A. When traffic is sent from the ISP to autonomous system 64512, the traffic will be forwarded to SanJose2 because of the higher MED value of SanJose2.
- B. The serial 0/0/1 interface on the ISP router has been configured with the set metric 50 command.
- C. The output was generated by entering the show ip bgp command on the SanJose1 router.
- D. The output was generated by entering the show ip bgp command on the ISP router.
- E. The serial 0/0/1 interface on the ISP router has been configured with the set metric 75 command.
- F. When traffic is sent from the ISP to autonomous system 64512, the traffic will be forwarded to SanJose1 because of the lower MED value of SanJose1.

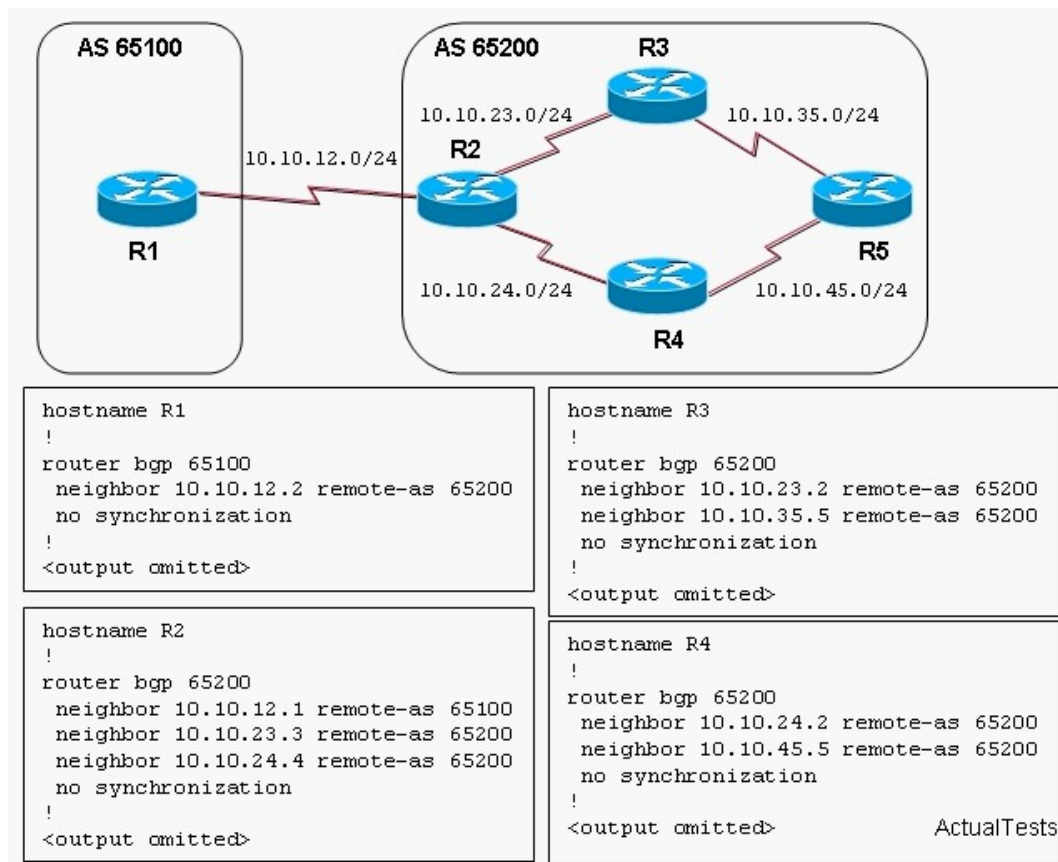
Answer: D,F

Explanation:

The "show ip route bgp" command will display any BGP-learned routes that make it into the IP routing table, the command "show ip bgp" is required to display the contents of the actual BGP routing table. This output was seen on ISP because the local router ID is 192.168.100.1 (ISP). Since we know that this output must have been seen by ISP, we know the serial 0/0/1 interface has been configured with a metric of 75, as this is the metric to the peer with IP address 192.168.1.2 (the other side of the serial 0/0/1 interface).

QUESTION NO: 65

Refer to the exhibit. On the basis of the configuration that is provided, how would the BGP updates that come from router R1 be replicated inside autonomous system 65200?




- A. All BGP updates that are received on router R2 will be sent to routers R3 and R4. Routers R3 and R4 will then forward those BGP updates to router R5.
- B. None of the BGP updates that are received on router R2 will ever be received by router R5.
- C. All BGP updates that are received on router R2 will not be sent to routers R3 and R4.
- D. All BGP updates that are received on router R2 will be sent directly to router R5.

Answer: B

QUESTION NO: 66

Refer to the exhibit. Router RT-1 chooses one path to network 198.133.219.0/24. Indicate the reason Router RT-1 chooses this "best" path.



```
RT-1# show ip bgp 198.133.219.0
BGP routing table entry for 198.133.219.0/24, version 75
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
    1
  5005 7007 31250
    128.107.254.2 (metric 782) from 128.107.2.2 (128.107.2.2)
      Origin IGP, metric 2100, localpref 100, valid, internal, best
  3003 6006 31250
    128.107.255.2 from 128.107.255.2 (128.107.0.3)
      Origin incomplete, metric 2500, localpref 100, valid, external
```

- A. In making its decision about the best path, RT-1 gives precedence to the BGP MED values.
- B. RT-1 prefers internal BGP routes.
- C. In making its decision about the best path, RT-1 gives precedence to the origin code.
- D. In making its decision about the best path, RT-1 prefers the IGP metrics.
- E. IP address 128.107.254.2 is lower than 128.107.255.2.
- F. IP address 128.107.2.2 is lower than 128.107.255.2.

Answer: C

Explanation:


After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination.

The following process summarizes how BGP chooses the best route on a Cisco router. Prefer the route with the highest weight. (The weight attribute is proprietary to Cisco and is local to the router only.) If multiple routes have the same weight, prefer the route with the highest local preference value. (The local preference is used within an autonomous system.) If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table. If none of the routes were locally originated, prefer the route with the shortest autonomous system path. If the autonomous system path length is the same, prefer the lowest origin code (IGP < EGP < incomplete). If all origin codes are the same, prefer the path with the lowest MED. (The MED is exchanged between autonomous systems.) The MED comparison is made only if the neighboring autonomous system is the same for all routes considered, unless the `bgp always-compare-med` command is enabled. If the routes have the same MED, prefer external paths to internal paths. If synchronization is disabled and only internal

paths remain, prefer the path through the closest IGP neighbor, which means that the router prefers the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop). For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping). Prefer the route with the lowest neighbor BGP router ID value. If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.

QUESTION NO: 67

Refer to the exhibit. Router RT-1 and router RT-2 both advertise network 131.25.0.0/16 to router RT-3 via internal BGP. What is the reason that router RT-3 chose router RT-1 as its best path to network 131.25.0.0/16.



```
RT-3# show ip bgp 131.25.0.0
BGP routing table entry for 131.25.0.0/16, version 14
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Not advertised to any peer
5005 7007 31250
  165.105.254.2 (metric 782) from 165.105.3.2 (165.2.2.2)
    Origin incomplete, metric 0, localpref 100, valid, internal
3003 6006 31250
  165.105.255.2 (metric 782) from 165.105.11.1 (165.1.1.1)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
```

ActualTests

- A. It has a better router ID.
- B. It advertises a lower autonomous system.
- C. It advertises the best AS-path.
- D. It advertises the best local preference.
- E. It advertises the best origin code.
- F. It advertises the best MED.

Answer: A

Explanation:

After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination.

The following process summarizes how BGP chooses the best route on a Cisco router. Prefer the route with the highest weight. (The weight attribute is proprietary to Cisco and is local to the router only.) If multiple routes have the same weight, prefer the route with the highest local preference

value. (The local preference is used within an autonomous system.) If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table. If none of the routes were locally originated, prefer the route with the shortest autonomous system path. If the autonomous system path length is the same, prefer the lowest origin code (IGP < EGP < incomplete). If all origin codes are the same, prefer the path with the lowest MED. (The MED is exchanged between autonomous systems.) The MED comparison is made only if the neighboring autonomous system is the same for all routes considered, unless the `bgp always-compare-med` command is enabled. If the routes have the same MED, prefer external paths to internal paths. If synchronization is disabled and only internal paths remain, prefer the path through the closest IGP neighbor, which means that the router prefers the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop). For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping). Prefer the route with the lowest neighbor BGP router ID value. If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.

In this case, since everything else is equal the router ID will be used, and the specific router ID that is chosen is 165.1.1.1 (lowest one), which is RT- 3.

QUESTION NO: 68

During BGP configuration on a router that has peered with other BGP speakers, the BGP command `aggregate-address 172.32.0.0 255.255.252.0` is issued. However, the peers do not receive this aggregate network in BGP advertisements. Also, the router does not have this aggregate network in its BGP table. Which option indicates a possible reason this command did not cause the router to advertise the aggregate network to its peers?

- A. Interface NULL 0 is likely shutdown.
- B. The next hop IP address must be a loopback address.
- C. The BGP command `no auto-summary` is missing.
- D. The BGP command `no synchronization` is missing.
- E. Subnets of 172.32.0.0/22 do not exist in the BGP table.
- F. The IGP running on this router does not have network 172.32.0.0/22 installed.

Answer: E

Explanation:

BGP allows the aggregation of specific routes into one route using the "aggregate-address address mask" command. Aggregation applies to routes that exist in the BGP routing table. This is in contrast to the network command, which applies to the routes that exist in IP routing table. Aggregation can be performed if at least one or more of the specific routes of the aggregate address exists in the BGP routing table.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800945ff.shtml

QUESTION NO: 69

For the accompanying router output, which of the following statements describes the state that neighbor 172.16.254.3 is in?

```
Router# show ip bgp neighbors 172.16.254.3
```

```
BGP neighbor is 172.16.254.3, remote AS 150, internal link
BGP version 4, remote router ID 172.16.254.3
BGP state = Active, table version 0
Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds
Default minimum time between advertisement runs is 5 seconds
Received 3 messages, 0 notifications, 0 in queue
Connections established 1; dropped 1
Last user reset 00:00:54, due to User reset request
No active TCP connection
```

ActualTests

- A. The router will not accept connections from the peer.
- B. BGP can exchange routing information in this state.
- C. The router is listening on its server port for connection requests from the peer.
- D. The router has sent out an active TCP connection request to the peer.

Answer: C

QUESTION NO: 70

Above is the output from show ip bgp neighbors command. What is line 21 stating about the BGP connection?

BGP neighbor is 172.16.254.3, remote AS 150, internal link

1. BGP version 4, remote router ID 172.16.254.3
2. BGP state = Established, up for 19:24:07
3. Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds
4. Neighbor capabilities:
5. Route refresh:advertised and received(new)
6. Address family IPv4 Unicast:advertised and received
7. Graceful Restart Capability:advertised and received
8. Remote Restart timer is 120 seconds
9. Address families preserved by peer:
10. IPv4 Unicast
11. Received 4231 messages, 0 notifications, 0 in queue
12. Sent 4167 messages, 0 notifications, 0 in queue
13. Default minimum time between advertisement runs is 5 seconds

14. For address family:IPv4 Unicast
15. BGP table version 159559, neighbor version 159559
16. Index 90, Offset 11, Mask 0x4
17. Route refresh request:received 0, sent 0
18. 10031 accepted prefixes consume 441364 bytes
19. Prefix advertised 29403, suppressed 0, withdrawn 9801
20. Number of NLRI in the update sent:max 242, min 0

21. Connections established 2; dropped 1
22. Last reset 19:26:54, due to NSF peer closed the session
23. Connection state is ESTAB, I/O status:1, unread input bytes:0
24. Local host:150.254.254.2, Local port:11005
25. Foreign host:172.16.254.3, Foreign port:179

ActualTests

- A. the number of consecutive TCP connections to the specified remote neighbor
- B. the number of total TCP connections that the router has
- C. the number of times the router has established a TCP connection
- D. the number of neighbors that the router has

Answer: C

QUESTION NO: 71

The 192.168.0.0 network is not being propagated throughout the network. Observe the BGP configuration commands from the advertising router. What is the reason the 192.168.0.0 route is not being advertised?

```
router bgp 65111
```

```
neighbor 172.16.1.1 remote-as 65111
```

```
neighbor 172.16.2.1 remote-as 65112
```

```
network 192.168.0.0
```

```
network 10.0.0.0
```

```
!
```

```
ip route 192.168.0.0 255.255.0.0 null0
```

Select the best response.

- A. The network 192.168.0.0 statement is missing mask 255.255.0.0
- B. The network 192.168.0.0 statement is missing mask 0.0.255.255.
- C. The network 10.0.0.0 statement is missing mask 255.0.0.0.
- D. The network 10.0.0.0 statement is missing mask 0.255.255.255.

E. The auto-summary configuration is missing.

Answer: A

QUESTION NO: 72

Which BGP attribute is used by BGP to prevent routing loops? Select the best response.

- A. AS-path
- B. next-hop
- C. MED
- D. weight
- E. local preference
- F. origin

Answer: A

QUESTION NO: 73

Which BGP path attribute is Cisco proprietary? Select the best response.

- A. weight
- B. MED
- C. local preference
- D. origin
- E. next-hop
- F. AS-path

Answer: A

QUESTION NO: 74

Which attribute must exist in the BGP update packet?
Select the best response.

- A. LOCAL_PREF
- B. AGGREGATOR
- C. AS_Path
- D. Weight

Answer: C

QUESTION NO: 75

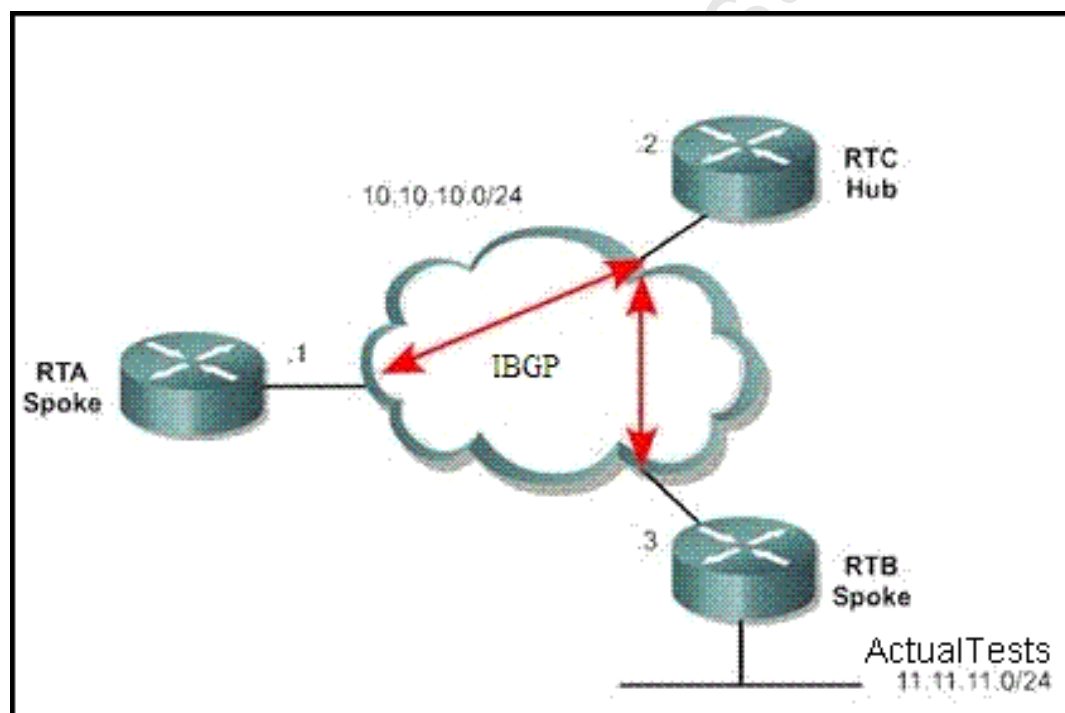
Which BGP attribute will not be advertised in routing updates to its neighboring routers? Select the best response.

- A. weight
- B. local preference
- C. origin
- D. AS_path
- E. next hop

Answer: A

QUESTION NO: 76

Observe the diagram. RTC is the hub router and RTA and RTB are the spokes. There are no virtual circuits between the spoke locations. What is needed to successfully route traffic to the 11.11.11.0/24 network from RTA?



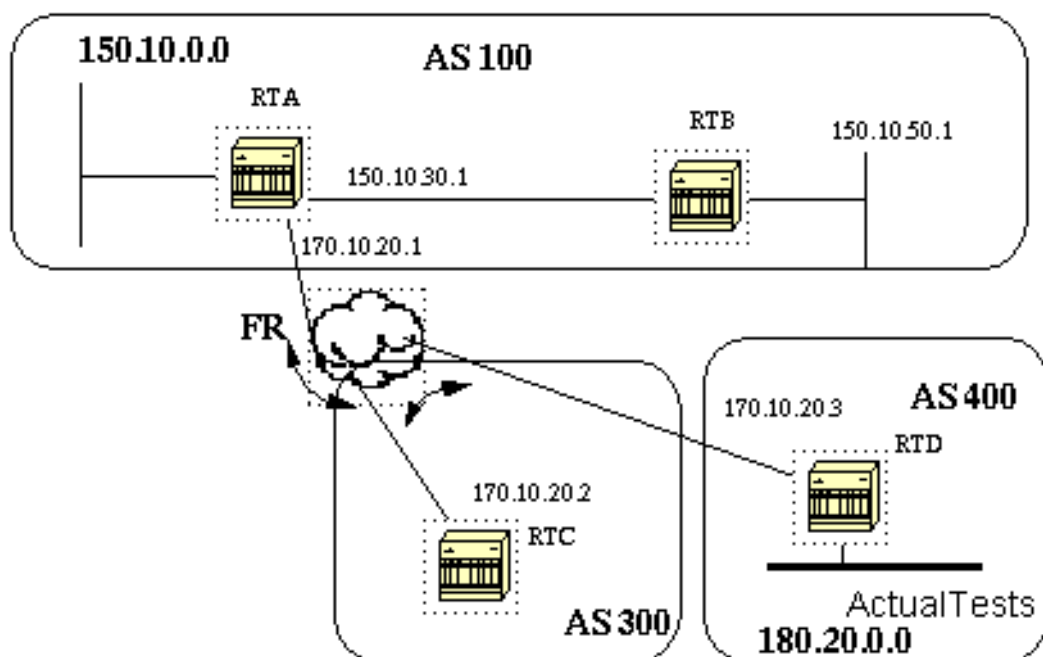
- A. The neighbor 10.10.10.1 next-hop-self command on RTA.
- B. Nothing is required. This is the default behavior on this topology.
- C. The neighbor 10.10.10.1 next-hop-self command on RTB.
- D. The neighbor 10.10.10.1 next-hop-self command on RTC.

Answer: D

Explanation:

The following example illustrates the issue:

BGP Next Hop (NBMA)



If the common media as you see in the shaded area above is a frame relay or any NBMA cloud then the exact behavior will occur as if we were connected via Ethernet. RTC will advertise 180.20.0.0 to RTA with a next hop of 170.10.20.3.

The problem is that RTA does not have a direct PVC to RTD, and cannot reach the next hop. In this case routing will fail.

In order to remedy this situation a command called next-hop-self is created.

The next-hop-self Command

Because of certain situations with the next hop as we saw in the previous example, a command called next-hop-self is created. The syntax is:

```
neighbor { ip-address|peer-group-name } next-hop-self
```

The next-hop-self command allows us to force BGP to use a specified IP address as the next hop rather than letting the protocol choose the next hop.

In the previous example, the following configuration solves our problem:

RTC#

```
router bgp 300
```

```
neighbor 170.10.20.1 remote-as 100
```

```
neighbor 170.10.20.1 next-hop-self
```

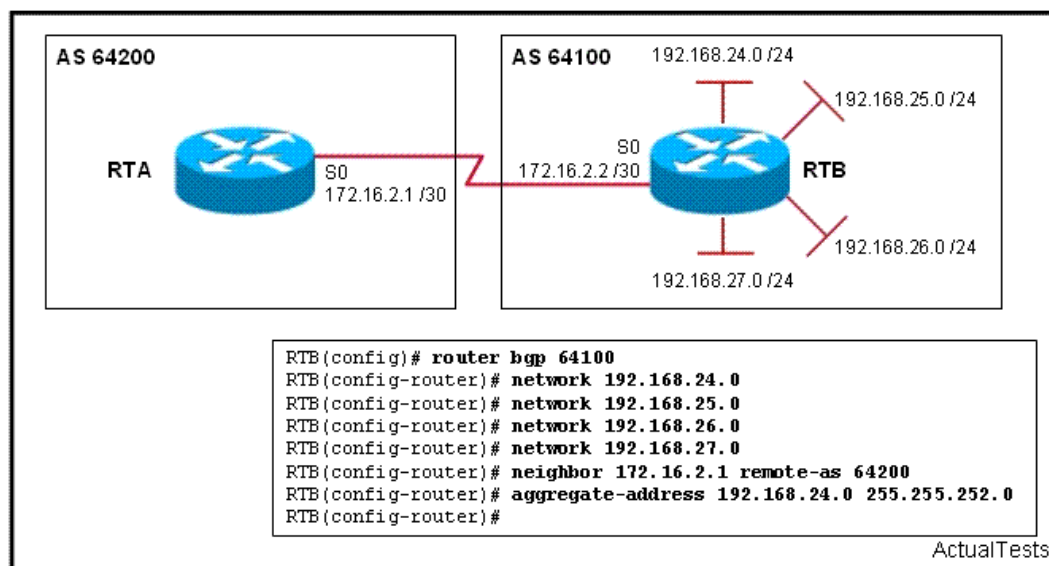
RTC advertises 180.20.0.0 with a next hop = 170.10.20.2

Reference: <http://www.cisco.com/warp/public/459/bgp-toc.html#bgpnexthop>

QUESTION NO: 77

Refer to the exhibit diagram and configuration. RTB is summarizing its networks from AS 64100 with the aggregate-address command. However, the show ip route command on RTA reveals the

RTB individual networks as well as its summary route. Which option would ensure that only the summary route would appear in the routing table of RTA?



- A. Add a static route with a prefix of 192.168.24.0 255.255.252.0 pointing to the null0 interface.
- B. Create a route map permitting only the summary address.
- C. Delete the four network statements and leave only the aggregate-address statement in the BGP configuration.
- D. Add the keyword summary-only to the aggregate-address command.

Answer: D

QUESTION NO: 78

Which statement is true about IBGP routers?

- A. They must be fully meshed.
- B. They can be in a different AS.
- C. They do not need to be directly connected.
- D. They must be directly connected.

Answer: C

Explanation:

When two routers establish a TCP enabled BGP connection, they are called neighbors or peers. Each router running BGP is called a BGP speaker. Peer routers exchange multiple messages to open and confirm the connection parameters, such as the version of BGP to be used. If there are any disagreements between the peers, notification errors are sent and the connection fails. To be a neighbor of BGP no need to be directly connected.

QUESTION NO: 79

When an IPv6 enabled host boots, it sends a router solicitation (RS) message. An IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA? (Choose two.)

- A. any route advertisements it has received
- B. prefixes for the link
- C. keepalive timers
- D. IPv6 address for the host
- E. request for the local host IP address
- F. lifetime of the prefix

Answer: B,F

Explanation:

In IP Version 6, Router Advertisements have the following attributes:

Routers advertise periodically

- Max. time between advertisements can be in the range from 4 and 1800 seconds

- The advertisement has a lifetime (= 0 if not a default router)

Advertisement contains one or more prefixes

- Prefixes have a lifetime

Preferred lifetime

Valid lifetime

Specifies if stateful or stateless autoconfiguration is to be used

Plays a key role in site renumbering

QUESTION NO: 80

Which statement is true about EBGP?

- A. An internal routing protocol can be used to reach an EBGP neighbor.
- B. The next hop does not change when BGP updates are exchanged between EBGP neighbors.
- C. EBGP requires a full mesh.
- D. A static route can be used to form an adjacency between neighbors.

Answer: D

Explanation:

When BGP is running between routers in different autonomous systems, it is called External BGP (EBGP). When BGP is running between routers in the same AS, it is called Internal BGP (IBGP). BGP allows the path that packets take to be manipulated by the AS, as described in this module. It

is important to understand how BGP works to avoid creating problems for your AS as a result of running BGP.

QUESTION NO: 81

Based on the show ip bgp summary output, which two statements are true? (Choose two.)

```
R1# show ip bgp summary
BGP table version is 8, main routing table version 8
4 network entries (8/12 paths) using 832 bytes of memory
5 BGP path attribute entries using 576 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
2 received paths for inbound soft reconfiguration
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.1	4	50001	80	81	8	0	0	04:15:10	2
10.2.2.2	4	50002	79	81	0	0	0	00:00:15	Active
10.3.3.3	4	50003	82	82	0	0	0	02:00:00	Idle

- A. The BGP session to the 10.3.3.3 neighbor is established, but the router has not received any BGP routing updates from the 10.3.3.3 neighbor.
- B. The router is attempting to establish a BGP peering session with the 10.2.2.2 neighbor.
- C. The router is attempting to establish a BGP peering session with the 10.1.1.1 neighbor.
- D. The BGP session to the 10.3.3.3 neighbor is established.
- E. The BGP session to the 10.2.2.2 neighbor is established.
- F. The BGP session to the 10.1.1.1 neighbor is established.


Answer: B,F

Explanation:

In BGP, established valid neighbors will have a number under the state/pfxRcd heading. Neighbors that show an active state mean that the router is actively trying to establish a peering relationship. Idle peers are not active, and not actively trying to peer. In this case, only the peer to 10.1.1.1 is currently operational with two route prefixes received, while the peer to 10.2.2.2 is trying to become established.

QUESTION NO: 82

Refer to the exhibit. By default, when RTB passes BGP advertisements from RTA about network 192.168.2.0 to RTC, what address will be listed as the next-hop address?



- A. 10.1.1.2
- B. 192.168.1.49
- C. 10.1.1.1
- D. 192.168.1.50

Answer: D


Explanation:

The next hop attribute is a well-known mandatory attribute, type code 3. In terms of an IGP, such as RIP, the "next hop" to reach a route is the IP address of the router that has announced the route.

The next hop concept with BGP is more complex and takes one of the following three forms: For EBGP sessions, the next hop is the IP address of the neighbor that announced the route. For IBGP sessions, where routes originated inside the AS, the next hop is the IP address of the neighbor that announced the route. For routes injected into the AS by way of EBGP, the next hop learned from EBGP is carried unaltered into IBGP. The next hop is the IP address of the EBGP neighbor from which the route was learned. When the route is advertised on a multiaccess medium, such as Ethernet or Frame Relay, the next hop is usually the IP address of the interface of the router. This will be the interface connected to the media that originated the route.

QUESTION NO: 83

Refer to the exhibit. Autonomous systems 200 and 300 have EBGP sessions established with their directly connected routers in autonomous system 100. IGP has been configured on all routers in autonomous system 100 and they successfully exchange routing updates. Traffic originated in autonomous system 200 cannot reach the destination autonomous system 300. What configuration should be done on the routers in autonomous system 100 in order for the traffic coming from autonomous system 200 to be forwarded to autonomous system 300?



- A. IBGP session must be established between routers R1 ?R2 and R2 ?R3, and the synchronization must be turned off.
- B. IBGP speakers within autonomous 100 must be fully meshed, and the synchronization must be turned on.
- C. IBGP session must be established between routers R1 ?R2 and R2 ?R3, and the synchronization must be turned on.
- D. IBGP session must be established between routers R1 and R3, and the synchronization must be turned off.
- E. IBGP speakers within autonomous 100 must be fully meshed, and the synchronization must be turned off.
- F. IBGP session must be established between routers R1 and R3, and the synchronization must be turned on.

Answer: E

QUESTION NO: 84

Which statement is true about IPv6?

- A. Only one IPv6 address can be assigned to each interface.
- B. IPv6 hosts use anycast addresses to assign IP addresses to interfaces.
- C. Each host can autoconfigure its address without the aid of a DHCP server.
- D. Only one IPv6 address is assigned per node.

Answer: C

QUESTION NO: 85

Refer to the exhibit. Which statement is true about the 6.6.6.0/24 prefix?

```
R2# show ip bgp 6.6.6.0
BGP routing table entry for 6.6.6.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    10.10.23.3
    30
    10.10.23.3 from 10.10.23.3 (6.6.6.1)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
```

- A. The command neighbor send-community is configured on BGP neighbor 10.10.23.3.
- B. If another path advertises the 6.6.6.0/24 path and has the default local preference, that path is more preferred.
- C. Route 6.6.6.0/24 is learned by an IBGP peer.
- D. The route 10.10.23.3 is not being advertised to other BGP neighbors.

Answer: A

QUESTION NO: 86

Which one of the following statements about BGP is FALSE?

- A. BGP uses TCP port 179.
- B. BGP ensures reliability of updates by using the reliable transport services of TCP.
- C. The network command with the mask option never installs a prefix into the BGP table unless there is a matching prefix exists in the IP route table.
- D. A TCP connection is required before exchanging updates.
- E. BGP uses notification and the update messages to establish and maintain the BGP neighbor relationship.

Answer: E

QUESTION NO: 87 DRAG DROP

Drop

Place the BGP attributes in the correct order used for determining a route.

originate route	1st
AS_Path	2nd
weight	3rd
local preference	4th
MED	5th

ActualTests

Answer:

Place the BGP attributes in the correct order used for determining a route.

originate route	weight
AS_Path	local preference
weight	originate route
local preference	AS_Path
MED	MED

ActualTests

Explanation:

Place the BGP attributes in the correct order used for determining a route.

originate route	weight
AS_Path	local preference
weight	originate route
local preference	AS_Path
MED	MED

ActualTests

QUESTION NO: 88

Examine the above configuration. What does the route map named test accomplish?


```
router bgp 50001
network 10.0.0.0
network 172.16.0.0
neighbor 10.1.1.1 remote-as 5003
neighbor 10.1.1.1 route-map test out
!
access-list 1 permit 10.0.0.0
access-list 2 permit any
!
route-map test permit 10
match ip address 1
set metric 200
!
route-map test permit 20
match ip address 2
```

ActualTests

- A. permits only the 10.0.0.0/8 prefix to be advertised to the 10.1.1.1 neighbor
- B. marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200
- C. marks the 10.0.0.0/8 prefix advertised to the 10.1.1.1 neighbor with a MED of 200
- D. permits only the 10.0.0.0/8 prefix to be received from the 10.1.1.1 neighbor
- E. marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200
- F. marks the 10.0.0.0/8 prefix received from the 10.1.1.1 neighbor with a MED of 200

Answer: C

Explanation:

The Multiple-exit-discriminator (MED) attribute is an optional nontransitive attribute, type code 4. MED informs external neighbors about the preferred path into an AS that has multiple entry points. A lower MED is preferred over a higher MED.

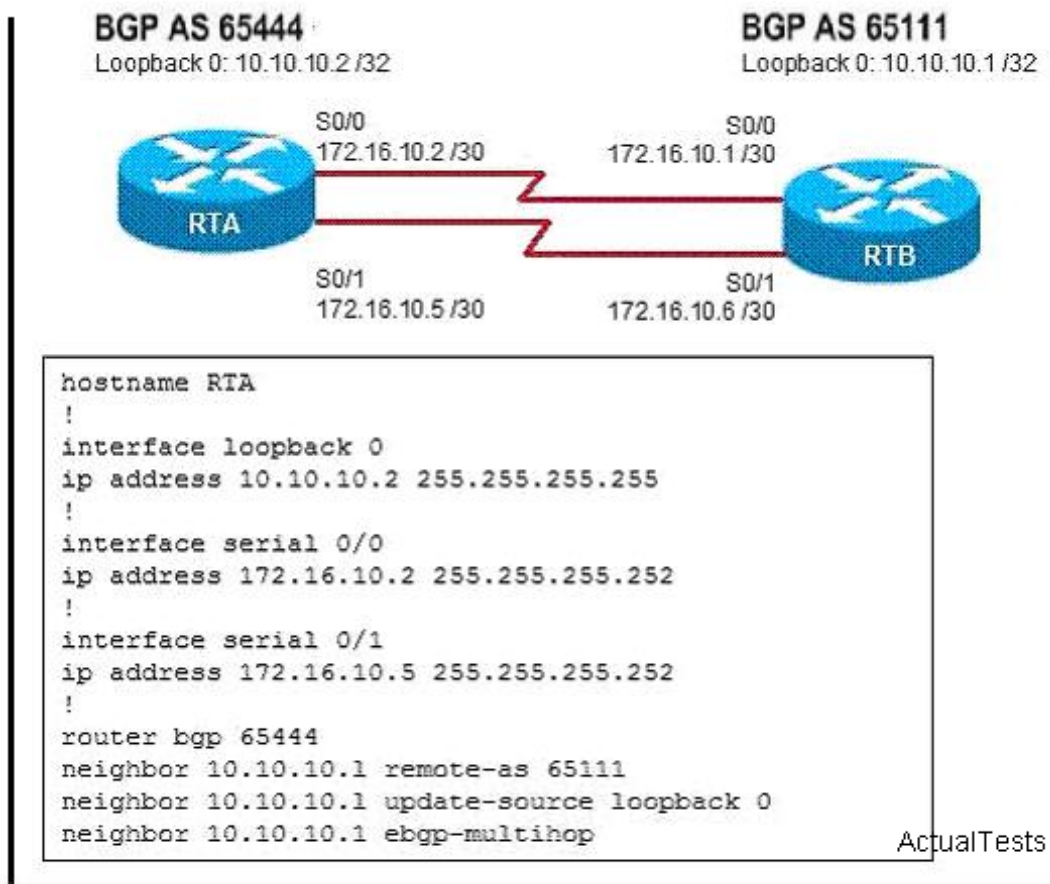
Unlike Local Preference, the MED attribute is exchanged between autonomous systems, but a MED attribute that comes into an AS does not leave the AS. When an update enters the AS with a certain MED value, that value is used for decision making within the AS. When BGP forwards the routing update to another AS, the MED is reset to zero. This is true unless the outgoing MED is set to a specific value.

Example:

```
RTB(config)# route-map setmedout permit 10
RTB(config-route-map)# set metric 50
RTB(config)#
router bgp 400
```

QUESTION NO: 89

Refer to the exhibit. Routers RTA and RTB are running BGP but the session is active. What command needs to be added to establish the BGP session?



- A. ip route 10.10.10.1 255.255.255.255 s0/0
ip route 10.10.10.1 255.255.255.255 s0/1
- B. network 10.10.10.0
- C. no synchronization
- D. neighbor 10.10.10.1 next-hop-self

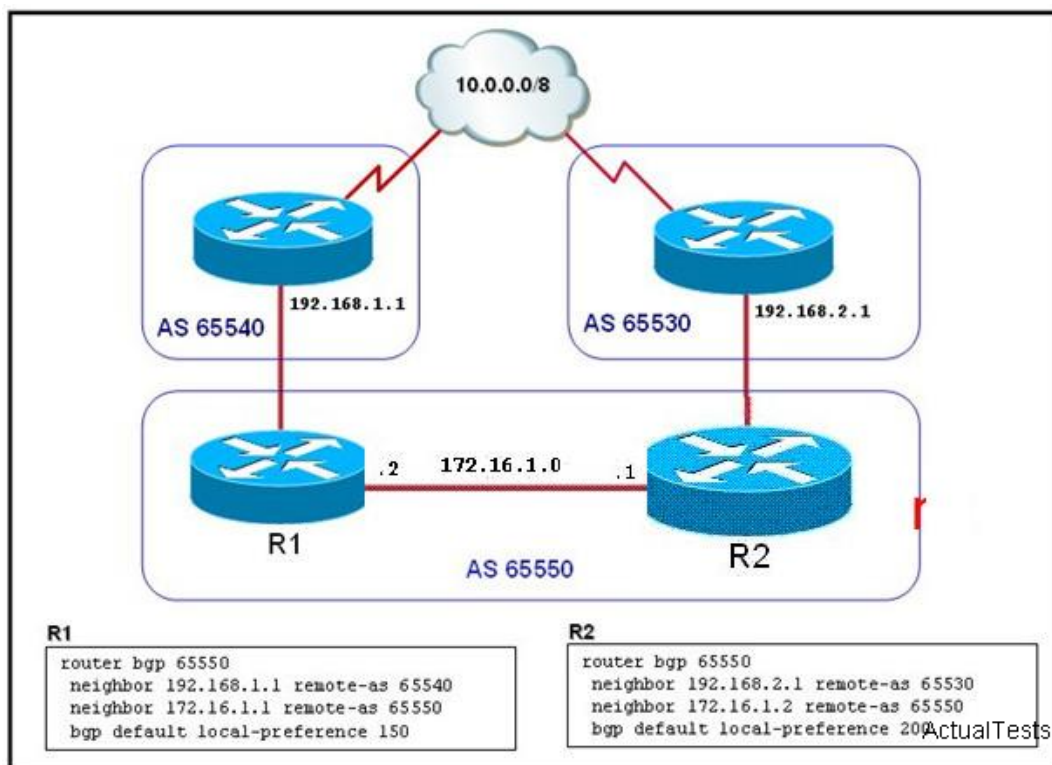
Answer: A

Explanation:

When BGP is running between routers in different autonomous systems, it is called External BGP (EBGP). When BGP is running between routers in the same AS, it is called Internal BGP (IBGP). BGP allows the path that packets take to be manipulated by the AS, as described in this module. It is important to understand how BGP works to avoid creating problems for your AS as a result of running BGP. A static route can be used to form an adjacency between EBGP neighbors.

QUESTION NO: 90

Refer to the exhibit. Network 10.0.0.0/8 is being advertised to autonomous system 65550 via both external links. Which statement about the preferred path to the 10.0.0.0/8 network is true?



- A. Router R2 will be preferred because it has the higher neighbor IP address.
- B. Router R2 will be preferred because its neighbor has a lower autonomous system number.
- C. Router R1 will be preferred because it has a lower local preference.
- D. Router R2 will be preferred because it has a higher local preference.
- E. Router R1 will be preferred because its neighbor has the higher autonomous system number.
- F. Router R1 will be preferred because it has the lower neighbor IP address.

Answer: D

Explanation:

The preferred path to 10.0.0.0/8 network is R2 because it has a higher local preference.

The following process summarizes how BGP chooses the best route on a Cisco router. Prefer the route with the highest weight. (The weight attribute is proprietary to Cisco and is local to the router only.) If multiple routes have the same weight, prefer the route with the highest local preference value. (The local preference is used within an autonomous system.) If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table. If none of the routes were locally originated, prefer the route with the shortest autonomous system path. If the autonomous system path length is the same, prefer the lowest origin code (IGP < EGP < incomplete). If all origin codes are the same, prefer the path with the lowest MED. (The MED is exchanged between autonomous systems.) The MED comparison is made only if the neighboring autonomous system is the same for all routes

considered, unless the `bgp always-compare-med` command is enabled. If the routes have the same MED, prefer external paths to internal paths. If synchronization is disabled and only internal paths remain, prefer the path through the closest IGP neighbor, which means that the router prefers the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop). For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping). Prefer the route with the lowest neighbor BGP router ID value. If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.

In this example, since the weights remained the same (default) value the next thing that is looked at is the highest local preference.

QUESTION NO: 91

What technique should be used on BGP edge routers to prevent a multi-homed autonomous system from becoming a transit system?

- A. Advertise with a high MED value all networks that are discovered via external BGP.
- B. Set the no-export community attribute on all networks that are advertised externally.
- C. Remove the AS-Path information on all routes in the BGP table prior to advertising externally.
- D. Set the origin code to incomplete for all networks that are discovered via external BGP.
- E. Only advertise networks externally if they have been discovered via internal BGP.
- F. Use an outgoing distribution list to filter all networks not originating from inside the autonomous system.

Answer: F

Explanation:

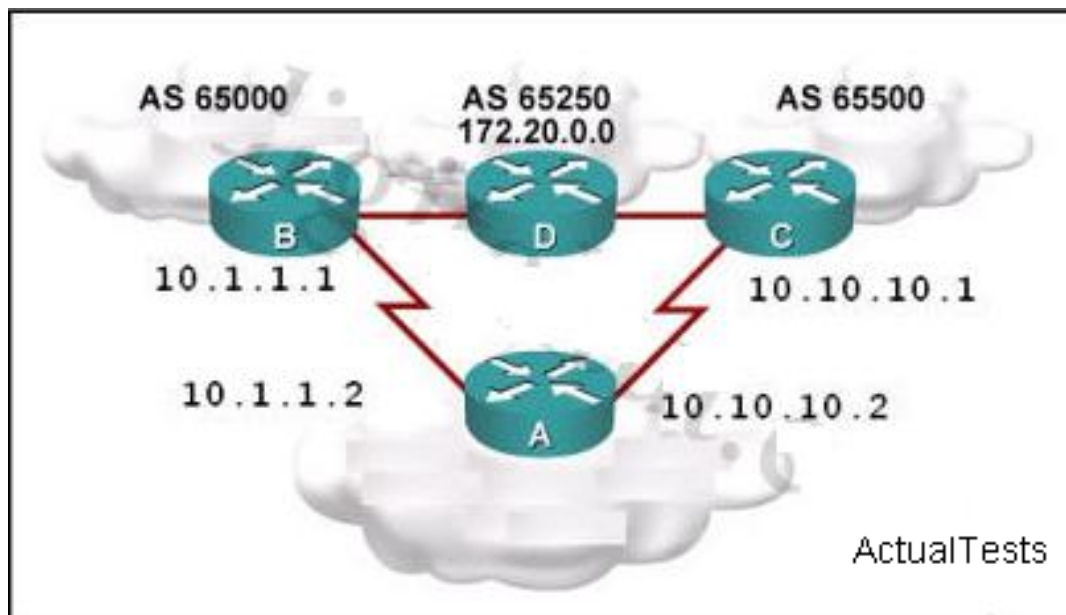
The minimum configuration that guarantees you won't become a transit AS is shown in the following example:

```
router bgp 65000 neighbor 10.1.1.1 filter-list 1 out ! ip as-path access-list 1 permit ^$
```

In this example, the outgoing filter list says that all traffic not sourced from the local AS, should not be advertised. This will prevent the announcement of routes that originated from other Autonomous Systems, and prevent traffic destined for networks in another AS from going through your AS. This is useful real-world information to know and if you configure your network with BGP in a multi-homed environment, odds are good that you will want to implement this kind of filter.

QUESTION NO: 92

Refer to the exhibit. The `neighbor 10.1.1.1 weight 200` BGP configuration command has been configured on router A. What will be the result of this configuration?



- A. Router A will prefer the path through router B for network 172.20.0.0
- B. Router A will prefer the path through router C for network 172.20.0.0
- C. Packets from router D will prefer the path through router B for networks advertised by router A
- D. Packets from router D will prefer the path through router C for networks advertised by router A

Answer: A

Explanation:

The weight attribute is a special Cisco attribute that is used in the path selection process when there is more than one route to the same destination. The higher the weight value, the better the path. The default weight is 0. Therefore, by configuring weight 200 to the neighbor 10.1.1.1, router A will prefer the path through router B for network 172.20.0.0 then the path through router C. The weight attribute is local to the router and not propagated to other routers. In this case the weight is local to router A so it has no effect on the decision of transferring packets from router D.

QUESTION NO: 93

Which two methods advertise internal networks to external ISPs via BGP? (Choose two.)

- A. using aggregate routes
- B. defining routes via the network statement
- C. disabling synchronization
- D. forcing the next-hop address

Answer: A,B

Explanation:

In BGP, the network command tells the BGP process what locally learned networks to advertise. The networks can be connected routes, static routes, or routes learned by way of a dynamic routing protocol, such as RIP. These networks must also exist in the routing table of the local

router or they will not be sent out in updates. The mask keyword can be used with the network command to specify individual subnets. Routes learned by the BGP process are propagated by default but are often filtered by a routing policy.

```
Router(config-router)# network network-number [ mask network-mask ]
```

QUESTION NO: 94

When the BGP path selection process is being performed on a Cisco router, which BGP attribute is used first when determining the best path?

- A. local preference
- B. MED
- C. weight
- D. origin
- E. next-hop
- F. AS-path

Answer: C

QUESTION NO: 95

Which command displays the IBGP and EBGP neighbors that are configured?

- A. show ip bgp paths
- B. show ip bgp peers
- C. show ip bgp
- D. show ip bgp summary

Answer: D

Explanation:

Show ip bgp summary command displays the summary of all BGP connections.

QUESTION NO: 96

Which three conditions can cause BGP neighbor establishment to fail? (Choose three.)

- A. The IBGP neighbor is not directly connected.
- B. BGP synchronization is enabled in a transit autonomous system with fully-meshed IBGP neighbors.
- C. The EBGP neighbor ebgp-multihop option is set to the default value.

- D. The BGP update interval is different between the two BGP neighbors.
- E. The BGP neighbor is referencing an incorrect autonomous system number in its neighbor statement.
- F. There is an access list blocking all TCP traffic between the two BGP neighbors.

Answer: C,E,F

Explanation:

BGP uses TCP port 179 to establish and maintain neighbor relationships, so any access lists or firewalls must permit this port for BGP to function.

By default, EBGP multi-hop is not enabled, so the EBGP peer must be directly connected in order for the local BGP router to know how to reach the EBGP peer.

BGP uses neighbor statements that specify the AS number of the BGP peer. If it is the same number as the local BGP router process, then the BGP router knows that IBGP is used. If the AS number for the specified peer is different, then EBGP is used. Either way, the specified neighbor must be configured correctly, or the BGP peers will not become neighbors.

QUESTION NO: 97

Which three of the statements below correctly describe the characteristics of Autonomous Systems in routed networks? (Select three)

- A. Within an AS, all routers must run either BGP or IBGP.
- B. An AS uses exterior gateway protocols (EGPs) to exchange information with other autonomous systems.
- C. An AS is a group of routers under the same technical administration.
- D. Within an AS, routes learned through BGP can be redistributed using interior gateway protocols.
- E. Within an AS, routes learned through an interior protocol cannot be redistributed using BGP to other autonomous systems.

Answer: B,C,D

QUESTION NO: 98

What are the two reasons for the appearance of 0.0.0.0 as the next hop for a network when using the "show ip bgp" command? (Choose two)

- A. The network was originated via redistribution of an interior gateway protocol into BGP.
- B. The network was defined by a static route.
- C. The network was learned via IBGP.

- D. The network was learned via EBGp.
- E. The network was originated via a network or aggregate command.

Answer: A,E

QUESTION NO: 99

What are two rules for compacting IPv6 addresses? (Choose two.)

- A. Every 16-bit segment that consists of all zeroes can be represented with a single colon.
- B. The trailing zeroes in any 16-bit segment do not have to be written.
- C. The leading zeroes in any 16-bit segment do not have to be written.
- D. Any single, continuous string of one or more 16-bit segments that consists of all zeroes can be represented with a double colon.
- E. The maximum number of times a double colon can replace a 16-bit segment that consists of all zeroes is two.
- F. Two zeroes in the middle of any 16-bit segment do not have to be written.

Answer: C,D

QUESTION NO: 100

What is the difference between the IPv6 addresses ::/0 and ::/128?

- A. ::/0 is the default route, and ::/128 is the unspecified address.
- B. ::/0 is the unicast address, and ::/128 is the anycast address.
- C. ::/0 is the anycast address, and ::/128 is the multicast address.
- D. ::/0 is the unicast address, and ::/128 is the multicast address.
- E. ::/0 is the unspecified address, and ::/128 is the multicast address.
- F. ::/0 is the anycast address, and ::/128 is the default address.

Answer: A

QUESTION NO: 101

How is authentication handled with OSPFv3?

- A. OSPFv3 for IPv6 authentication is supported by IPv6 IPsec.
- B. OSPFv3 for IPv6 authentication is supported by MD5 authentication.
- C. OSPFv3 for IPv6 authentication is supported by IPv4 IPsec.

D. OSPFv3 for IPv6 authentication is supported by SHA-1 authentication.

Answer: A

QUESTION NO: 102

Which three are characteristics of IPv6? (Choose three.)

- A. An IPv6 address is 128 bits long.
- B. An IPv6 header contains the header checksum field.
- C. An IPv6 header is 20 bits long.
- D. An IPv6 header contains the next header field.
- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the protocol field.

Answer: A,D,E

Explanation:

All IPv6 addresses are 128 bits long to accommodate a far larger number of stations than what was possible with the 32 bit IPv4 addresses.

The following displays the IPv6 header field in detail:

IPv6 header:

Version . 4 bits. IPv6 version number.

Traffic Class . 8 bits. Internet traffic priority delivery value.

Flow Label . 20 bits. Used for specifying special router handling from source to destination(s) for a sequence of packets.

Payload Length . 16 bits unsigned. Specifies the length of the data in the packet. When cleared to zero, the option is a hop-by-hop Jumbo payload.

Next Header . 8 bits. Specifies the next encapsulated protocol. The values are compatible with those specified for the IPv4 protocol field.

Hop Limit . 8 bits unsigned. For each router that forwards the packet, the hop limit is decremented by 1. When the hop limit field reaches zero, the packet is discarded. This replaces the TTL field in the IPv4 header that was originally intended to be used as a time based hop limit.

Source address . 16 bytes. The IPv6 address of the sending node.

Destination address . 16 bytes. The IPv6 address of the destination node.

Reference: <http://www.networksorcery.com/enp/protocol/ipv6.htm>

QUESTION NO: 103

Which IPv6 address correctly compresses the IPv6 unicast address 2001:0:0:0:0DB8:0:0:417A?
Select the best response.

- A. 2001:0DB8:417A
- B. 2001::0DB8::417A
- C. 2001:::0DB8::417A
- D. 2001:0DB8:0:0:417A
- E. 2001::DB8:0:0:417A
- F. 2001:::0DB8:0:0:417A

Answer: E

Explanation:

The point of this question is the about the different form of Ipv6 address.

The IPv6 address is 128 bits long, written as eight 16-bit pieces, separated by colons. Each piece is represented by four hexadecimal digits. You can compact multiple contiguous fields of zero even further. This is the exception to the rule that at least one digit must be present in every field. You can replace multiple fields of zeros with double colons (::).

Note that :: can replace only one set of contiguous zero fields. Multiple ::s would make the address ambiguous.

QUESTION NO: 104

Refer to the exhibit. What two statements are true? (Choose two.)

```

<output omitted>
!
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::205:5FFF:FED3:5808, Interface ID 3
  Area 1, Process ID 1, Instance ID 0, Router ID 172.16.3.3
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 172.16.6.6, local address
FE80::205:5FFF:FED3:6408
  Backup Designated router (ID) 172.16.3.3, local address
FE80::205:5FFF:FED3:5808
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 12, maximum is 12
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.6.6 (Designated Router)
  Suppress hello for 0 neighbor(s)

```

ActualTests

- A. This is the designated router (DR) on the FastEthernet 0/0 link.
- B. The output was generated by the show ip interface command.
- C. The router was configured with the commands:
router ospf 1
network 172.16.6.0 0.0.0.255 area 1
- D. The IP address of the backup designated router (BDR) is FE80::205:5FFF:FED3:5808.
- E. Interface FastEthernet 0/0 was configured with the ipv6 ospf 1 area 1 command.

F. OSPF version 2 has been enabled to support IPv6.

Answer: D,E

Explanation:

OSPFv3 supports IPv6. The configuration of OSPFv3 is not a subcommand mode of the router ospf command as it is in OSPFv2 configuration. For example, instead of using the network area command to identify networks that are part of the OSPFv3 network, the interfaces are directly configured to specify that IPv6 networks are part of the OSPFv3 network.

The following describes the steps to configure OSPF for IPv6:

Step 1	Complete the OSPF network strategy and planning for your IPv6 network. For example, you must decide whether multiple areas are required.
Step 2	Enable IPv6 <u>unicast</u> routing using the ipv6 <u>unicast-routing</u> command.
Step 3	Enable IPv6 on the interface using the ipv6 <u>ospf area</u> command.
Step 4	(Optional) Configure OSPFv3 interface specific settings, including area, router priority, and OSPFv3 path cost.
Step 5	(Optional) Configure routing specifics from router configuration mode, including router priority, route summarization, and so on. ActualTests

There are several commonly used OSPFv3 show commands, including the show ipv6 ospf [process-id] [area-id] interface [interface] command.

QUESTION NO: 105

Refer to the exhibit. On the basis of the output from the show ipv6 interface command, what two statements must be true? (Choose two.) Select 2 response(s).

```

RTA# show ipv6 ospf interface
<Output omitted>
!
Ethernet0/0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:6E00, Interface ID 2
  Area 0, Process ID 1, Instance ID 0, Router ID 10.10.10.1
  Network Type BROADCAST, Cost:10
  MD5 Authentication SPI 500, secure socket state UP (errors:0)
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 11.11.11.1, local address
FE80::A8BB:CCFF:FE00:6F00
  Backup Designated router (ID) 10.10.10.1, local address
FE80::A8BB:CCFF:FE00:6E00
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:01
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 11.11.11.1 (Designated Router)
  Suppress hello for 0 neighbor(s)

```

ActualTests

- A. Interface Ethernet 0/0 was configured with the ipv6 ospf 1 area 1 command.
- B. Interface Ethernet 0/0 has been configured with the ipv6 ospf authentication ipsec spi 500 md5 command.
- C. OSPF version 3 is enabled to support IPv6.
- D. The IP address of the designated router (DR) is FE80::A8BB:CCFF:FE00:6E00
- E. This is the designated router (DR) on the Ethernet 0/0 link.

Answer: B,C

QUESTION NO: 106

What is the IPv6 address FF02::2 used for?

- A. all routers in a local segment
- B. all routers in an autonomous system
- C. all hosts in a local segment
- D. all hosts in a particular multicast group

Answer: A

Explanation:

To identify all nodes for the node-local and link-local scopes, the following multicast addresses are defined:

FF01::1 (node-local scope all-nodes address)

FF02::1 (link-local scope all-nodes address)

To identify all routers for the node-local, link-local, and site-local scopes, the following multicast addresses are defined:

FF01::2 (node-local scope all-routers address)

FF02::2 (link-local scope all-routers address)

FF05::2 (site-local scope all-routers address)

Reference: <http://technet2.microsoft.com/windowsserver/en/library/8c07faf2-35b3-4f2a-920f-d61ad76581ad1033.mspx?mfr=true>

QUESTION NO: 107

What does the command `clear ipv6 ospf process` accomplish?

- A. The shortest path first (SPF) algorithm is performed on the LSA database.
- B. The OSPF adjacencies are cleared and initiated again.
- C. The OSPF database is repopulated. Then the shortest path first (SPF) algorithm is performed.
- D. The route table is cleared. Then the OSPF neighbors are reformed.

Answer: C

QUESTION NO: 108

In IPv6, the interfaces running OSPF can be configured with multiple address prefixes. Which statement is true about the IPv6 addresses that can be included into the OSPF process?"

- A. Specific addresses can be selected using a prefix list.
- B. Specific addresses can be selected using a route map.
- C. Specific addresses cannot be selected for importation into the OSPF process.
- D. Specific addresses can be selected using an ACL.

Answer: C

QUESTION NO: 109

Which statement is true about the command `ipv6 ospf 1 area 0`?

- A. It must be issued before the network command to enable the OSPF process for IPv6.
- B. It must be issued after the network command to enable the OSPF process for IPv6.
- C. It must be issued in router global configuration mode to enable the OSPF process for IPv6.
- D. It must be issued in interface configuration mode to enable the OSPF process for IPv6.

Answer: D

QUESTION NO: 110

Refer to the exhibit. Based on the output from the show command on RT1 which statement is true? Select the best response.

```

RT1# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::218:B9FF:FE12:2CD1, Interface ID 4
  Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State Brother, Priority 1
  Designated Router (ID) 10.1.3.1, local address FE80::218:B9FF:FECD:BEF0
  Backup Designated router (ID) 10.1.2.1, local address FE80::218:B9FF:FE92:28D8
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:04
  Index 1/3/3, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 10.1.3.1 (Designated Router)
    Adjacent with neighbor 10.1.2.1 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)

```

ActualTests

- A. OSPFv3 uses global IPv6 addresses to establish neighbor adjacencies.
- B. RT1 has a subnet mask of 64 bits.
- C. RT1 has FastEthernet0/0 set as a DR for network type broadcast.
- D. OSPFv3 uses Link-local addresses to establish neighbor adjacencies.
- E. RT1 does not have a global IPv6 address set on FastEthernet0/0.
- F. OSPFv3 uses IPv4 addresses to establish neighbor adjacencies.

Answer: D

QUESTION NO: 111

To enable BGP tunneling over an IPv4 backbone, the IPv4 address 192.168.30.1 is converted into a valid IPv6 address. Which three IPv6 addresses are acceptable formats for the IPv4 address? (Choose three.)

- A. 192.168.30.1:0:0:0:0:0
- B. 0:0:0:0:0:0:192.168.30.1
- C. ::192.168.30.1

- D. C0A8:1E01::
- E. 192.168.30.1::
- F. ::C0A8:1E01

Answer: B,C,F

QUESTION NO: 112

What is IPv6 router solicitation?

- A. a request made by a node to join a specified multicast group
- B. a request made by a node for its IP address
- C. a request made by a node for the IP address of the DHCP server
- D. a request made by a node for the IP address of the local router

Answer: D

QUESTION NO: 113

Your network trainee asks you, in the context of IPv6 and OSPF, what best describes a type 9 LSA? What should you tell her?

- A. Link LSA
- B. Interarea prefix LSA for ABRs
- C. Router LSA
- D. Switch LSA
- E. Intra-area prefix LSA
- F. None of the above

Answer: E

QUESTION NO: 114

Which command enables OSPF for IPv6?

- A. router ospf ipv6 process-id
- B. ipv6 ospf process-id
- C. router ospf process-id
- D. ipv6 router ospf process-id

Answer: D

Explanation:

To enable Open Shortest Path First (OSPF) for IPv6 router configuration mode, use the `ipv6 router ospf process-id` command in global configuration mode.

`ipv6 router ospf process-id`

Syntax Description

Reference:

http://www.cisco.com/en/US/products/ps6922/products_command_reference_chapter09186a00801d6615.html#wp2027608

QUESTION NO: 115

Your company plans on migrating their network from IPv4 to IPv6 in the near future. Which three techniques can be used to transition from IPv4 to IPv6? (Select three)

- A. Dual stack
- B. NAT
- C. Flow label
- D. Mobile IP
- E. 6to4 tunneling
- F. Anycast
- G. MBGP

Answer: A,B,E

QUESTION NO: 116

Which command must be globally enabled on a Cisco router to support IPv6?

- A. `ip routing ipv6`
- B. `ipv6 unicast-routing`
- C. `ipv6 routing`
- D. `ip classless`
- E. `ipv6 cef`

Answer: B

QUESTION NO: 117

The company network is in the process of migrating the IP address scheme to use IPv6. Which of the following address types are associated with IPv6? (Select three)

- A. Unicast
- B. Private
- C. Broadcast
- D. Public
- E. Multicast
- F. Anycast

Answer: A,E,F

QUESTION NO: 118

Your Company has migrated to IPv6 in their network. Which three IPv6 notations represent the same address? (Select three)

- A. 2031::130F::9C0:876A:130B
- B. 2031:0000:130F:0000:0000:09C0:876A:130B
- C. 2031:0:130F:::9C0:876A:130B
- D. 2031::130F:0::9C0:876A:130B
- E. 2031:0:130F:0:0:09C0:876A:130B
- F. 2031:0:130F::9C0:876A:130B

Answer: B,E,F

QUESTION NO: 119

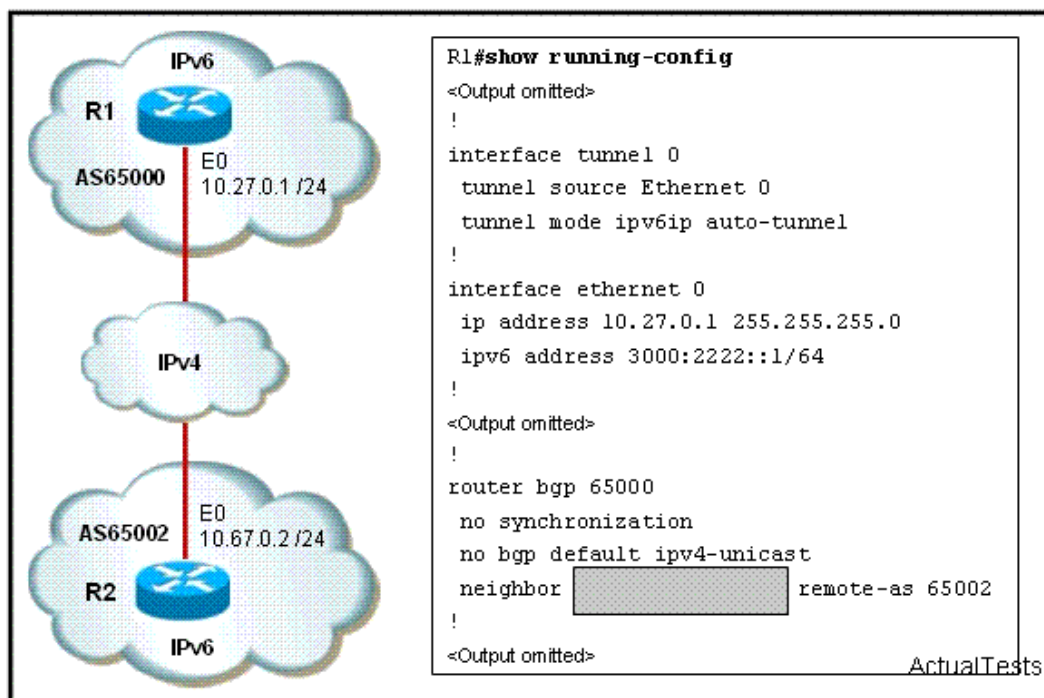
Which statement is true concerning 6to4 tunneling?

- A. The edge routers can use any locally configured IPv6 address.
- B. Hosts and routers inside a 6to4 site will need a special code.
- C. An edge router must use IPv6 address of 2002::/16 in its prefix.
- D. IPv4 traffic is encapsulated with an IPv6 header.

Answer: C

QUESTION NO: 120

Refer to the exhibit. Routers R1 and R2 are IPv6 BGP peers that have been configured to support a neighbor relationship over an IPv4 internetwork. Which three neighbor IP addresses are valid choices to use in the highlighted section of the exhibit? (Choose three.)



- A. ::0A43:0002
- B. 0A43:0002::
- C. 10.67.0.2::
- D. 10.67.0.2::0:0:0:0:0:0
- E. ::10.67.0.2
- F. 0:0:0:0:0:0:10.67.0.2

Answer: A,E,F

QUESTION NO: 121

Which two statements are true about using IPv4 and IPv6 simultaneously on a network segment? (Choose two.)

- A. IPv6 allows a host to create its own IPv6 address that will allow it to communicate to other devices on a network configured via DHCP. IPv4 does not provide a similar capability for hosts.
- B. IPv6 provides for more host IP addresses but IPv4 provides for more network addresses.
- C. Hosts can be configured to receive both IPv4 and IPv6 addresses via DHCP.
- D. Host configuration options for IPv4 can be either statically assigned or assigned via DHCP. Host configuration options for IPv6 can be statically assigned only.
- E. IPv4 and IPv6 addresses can be simultaneously assigned to a host but not to a router interface.

Answer: A,C

QUESTION NO: 122

Which statement describes the difference between a manually configured IPv6 in IPv4 tunnel versus an automatic 6to4 tunnel?

- A. An automatic 6to4 tunnel does not require dual-stack (IPv4 and IPv6) routers at the tunnel endpoints.
- B. An automatic 6to4 tunnel allows multiple IPv4 destinations.
- C. A manually configured IPv6 in IPv4 tunnel allows multiple IPv4 destinations.
- D. A manually configured IPv6 in IPv4 tunnel does not require dual-stack (IPv4 and IPv6) routers at the tunnel endpoints.

Answer: B

QUESTION NO: 123

In IPv6 network, an IPv6 overlay tunnel is required to communicate with isolated IPv6 networks across the existing IPv4 infrastructure. There are currently five IPv6 overlay tunnel types. Which three IPv6 overlay tunnel statements are true? (Choose three)

- A. Overlay tunnels can only be configured between border routers capable of supporting IPv4 and IPv6.
- B. An manual overlay tunnel supports point-to-multipoint tunnels capable of carrying IPv6 and Connectionless Network Service (CLNS) packets.
- C. Overlay tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure.
- D. Overlay tunnels can be configured between border routers or between a border router and a host capable of supporting IPv4 and IPv6.
- E. Cisco IOS supports manual, generic routing encapsulation (GRE), IPv4-compatible, 6to4, and Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) overlay tunneling mechanisms.
- F. Cisco IOS supports manual, generic routing encapsulation (GRE), IPv6-compatible, 4to6, and Multiprotocol Label Switching (MPLS) overlay tunneling mechanisms.

Answer: C,D,E

QUESTION NO: 124

When building an IP multicast domain using PIM which mode assumes that other routers do not want to forward multicast packets for the group? Select the best response.

- A. PIM-DM
- B. PIM-SM

- C. PIM-RP
- D. CGMP
- E. IGMP snooping

Answer: B

QUESTION NO: 125

A network administrator assigns a multicast address of 239.255.8.5 to an application running on a device with an Ethernet MAC address of 01.b2.7d.05.f1.80. Which Layer 2 multicast address will this device use?

- A. 01.00.5e.05.f1.80
- B. 01.b2.7d.05.f1.80
- C. ff.ff.ff.ff.ff
- D. 01.00.5e.7F.08.05
- E. 01.b2.7d.0a.08.05

Answer: D

QUESTION NO: 126

Which enhancement was added to IGMP version 3?

Select the best response.

- A. membership query message
- B. membership report message
- C. leave group message
- D. source filtering
- E. destination filtering

Answer: D

QUESTION NO: 127

When IP multicast is enabled via PIM, which mode uses the flood and prune method? Select the best response.

- A. PIM sparse-dense
- B. Bidir-PIM

- C. PIM-RP
- D. PIM-DM
- E. PIM-SM

Answer: D

QUESTION NO: 128

Which three statements are true about the Internet Group Management Protocol (IGMP)? (Choose three.)

- A. There are three IGMP modes: dense mode, sparse mode, and sparse-dense mode.
- B. IGMP is used to register individual hosts with a multicast group.
- C. IGMP version 3 enables a multicast receiving host to specify to the router which sources it should forward traffic from.
- D. IGMP messages are IP datagrams with a protocol value of 2, destination address of 224.0.0.2, and a TTL value of 1.
- E. IGMP is a multicast routing protocol that makes packet-forwarding decisions independent of other routing protocols such as EIGRP.
- F. IGMP snooping runs on Layer 3 routers.

Answer: B,C,D

QUESTION NO: 129

Which two statements are true about Internet Group Management Protocol (IGMP) snooping? (Choose two.) Select 2 response(s).

- A. IGMP snooping and Cisco Group Membership Protocol (CGMP) can be used simultaneously on a switch.
- B. IGMP snooping and Cisco Group Membership Protocol (CGMP) were developed to help Layer 3 switches make intelligent forwarding decisions on their own.
- C. IGMP snooping examines IGMP join/leave messages so that multicast traffic is forwarded only to hosts that sent an IGMP message toward the router.
- D. IGMP snooping is an IP multicast constraining mechanism for Layer 2 switches.
- E. IGMP snooping is enabled with the ip multicast-routing global configuration command.

Answer: C,D

QUESTION NO: 130

Refer to the exhibit. Given the output of a debug ip mrouting command, which two statements are true? (Choose two.)

```
Router# debug ip mrouting 224.2.0.1
MRT: Create (*, 224.2.0.1), if_input NULL
MRT: Create (224.69.15.0/24, 225.2.2.4), if_input Ethernet0, RPF nbr 224.69.61.15
MRT: Create (224.69.39.0/24, 225.2.2.4), if_input Ethernet1, RPF nbr 224.0.0.0
MRT: Create (10.9.0.0/16, 224.2.0.1), if_input Ethernet1, RPF nbr 0.0.0.0
MRT: Create (10.16.0.0/16, 224.2.0.1), if_input Ethernet1, RPF nbr 0.0.0.0
```

- A. The route to 224.69.15.0/24 will be out Ethernet 0.
- B. The reverse path forwarding (RPF) for the route 224.2.0.1 failed to find the interface on which the multicast packet was received.
- C. Multicast route to 224.69.15.0/24 was added to the mroute table and created by a source directly connected to the router.
- D. Multicast route to 10.16.0.0/16 was added to the mroute table and created by a source directly connected to the router.
- E. This router received an IGMP host report from a group member or a PIM join message.

Answer: D,E

QUESTION NO: 131

What is the purpose of a rendezvous point (RP)? Select the best response.

- A. acts as a meeting place for sources and receivers of multicast traffic
- B. used in PIM dense mode to create a database of all multicast sources
- C. used in PIM dense and sparse mode to create a database of all multicast sources
- D. acts as the designated router for a broadcast segment when multicast routing is enabled

Answer: A

QUESTION NO: 132

Which three IP addresses are valid multicast addresses? (Choose three.) Select 3 response(s).

- A. 169.254.23.59
- B. 223.254.255.254
- C. 225.1.1.1
- D. 227.43.34.2
- E. 238.3.3.3

F. 249.1.2.3

Answer: C,D,E**QUESTION NO: 133**

Which two multicast IP addresses can be represented by the multicast MAC address 0100.5e0A.0A07? (Choose two.)

- A. 228.10.10.7
- B. 229.11.10.7
- C. 229.138.10.7
- D. 229.138.10.8
- E. 228.10.10.8
- F. 228.10.138.7

Answer: A,C**QUESTION NO: 134**

Refer to the exhibit. What are two important facts in interpreting the output of the show ip pim interface command? (Choose two.)

```

R1# show ip pim interface
Address          Interface      Version/Mode  Nbr   Query   DR
                  Count Intvl
192.168.10.1     Ethernet0     v2/Sparse-Dense  1    30     192.168.10.2
192.168.9.3      Ethernet1     v2/Sparse-Dense  1    30     192.168.9.5
  
```

- A. Multiaccess, multicast segments do not elect a DR.
- B. Point-to-Point links do not display DR information.
- C. Multiaccess segments elect a DR based on highest IP address.
- D. Multiaccess segments elect a DR based on lowest IP address.
- E. The RP is only seen in version 2 of Sparse-Dense mode.

Answer: B,C**Explanation:**

show ip pim interface: Use this command to display information about interfaces configured for PIM. In addition, you can use this command to verify that the correct PIM mode (dense or sparse) is configured on the interface, the neighbor count is correct, and the designated router (DR) is correct (which is critical for PIM sparse mode). Multi-access segments (such as Ethernet, Token

Ring, FDDI) elect a DR based on highest IP address. Point-to-Point links do not display DR information.

Reference:

http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080093f21.shtml

QUESTION NO: 135

Which three statements are true about IP multicast configuration? (Choose three.)

- A. PIM sparse-dense mode and PIM dense mode require an RP on the network.
- B. PIM sparse mode and PIM dense mode require an RP on the network.
- C. PIM sparse mode interfaces are always added to the multicast routing table in a router.
- D. PIM sparse mode and PIM sparse-dense mode require an RP on the network.
- E. PIM dense mode interfaces are always added to the multicast routing table in a router.
- F. PIM sparse-dense mode acts as PIM dense mode if an RP is not known.

Answer: D,E,F

Explanation:

If you configure PIM to operate in sparse mode, you must also choose one or more routers to be rendezvous points (RPs). RPs are used by senders to a multicast group to announce their existence and by receivers of multicast packets to learn about new senders.

Dense mode interfaces are always added to the table when the multicast routing table is populated. Sparse mode interfaces are added to the table only when periodic join messages are received from downstream routers, or there is a directly connected member on the interface.

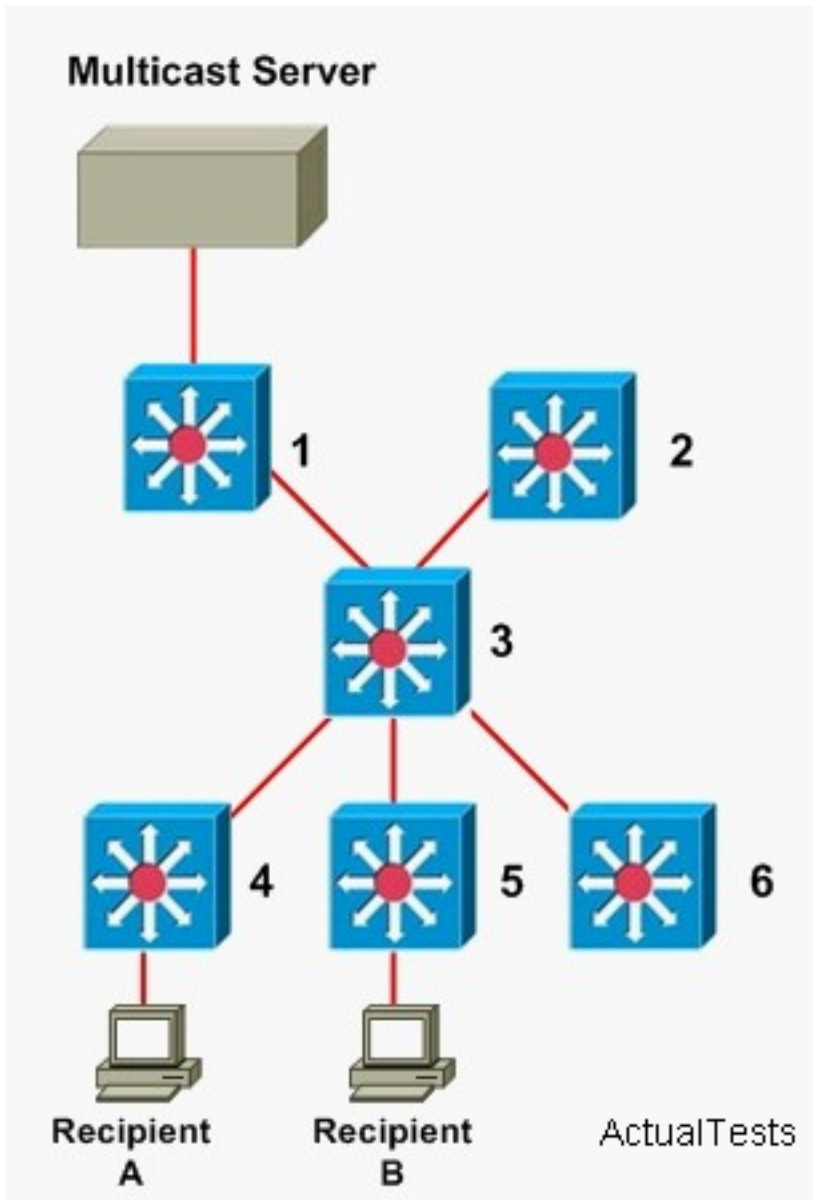
In PIM sparse-dense mode, if an RP is not known for a group, the router sends data using PIM dense mode. However, if the router discovers an RP or you configure an RP statically, PIM sparse mode takes over.

References: http://www.cisco.com/en/US/docs/ios/12_3/ipmulti/command/reference/ip3_i2g.html

<http://www.juniper.net/techpubs/software/erx/junose53/swconfig-routing-vol1/html/ip-multicast-config15.html>

QUESTION NO: 136

Refer to the exhibit. All multilayer switches are running PIM dense mode. Recipient A and Recipient B are sending IGMPv2 join messages to their respective multilayer switches. Which statement is true?




- A. The Multicast Server is the root of the multicast tree. Switches 1, 3, 4, and 5 will participate in the multicast tree once pruning has taken place.
- B. The Multicast Server is the root of the multicast tree. Switches 4 and 5 will participate in the multicast tree once pruning has taken place.
- C. Switch 3 is the root of the multicast tree. Switches 3,4, and 5 will participate in the multicast tree once pruning has taken place.
- D. Switch 3 is the root of the multicast tree. Switches 1,3,4, and 5 will participate in the multicast tree once pruning has taken place.
- E. Switch 1 is the root of the multicast tree. Switches 1,4, and 5 will participate in the multicast tree once pruning has taken place.
- F. Switch 1 is the root of the multicast tree. Switches 1,3,4, and 5 will participate in the multicast tree once pruning has taken place.

Answer: A

QUESTION NO: 137

Refer to the network of Layer 3 switches in the exhibit. The RPI Multicast Server only multicasts to hosts connected to multilayer switches 5 and 6. The CMU Multicast Server multicasts to hosts on multilayer switches 1-6. Given the number of configuration steps involved, what is the most efficient way to configure the network while meeting the requirements for multicast data flow?



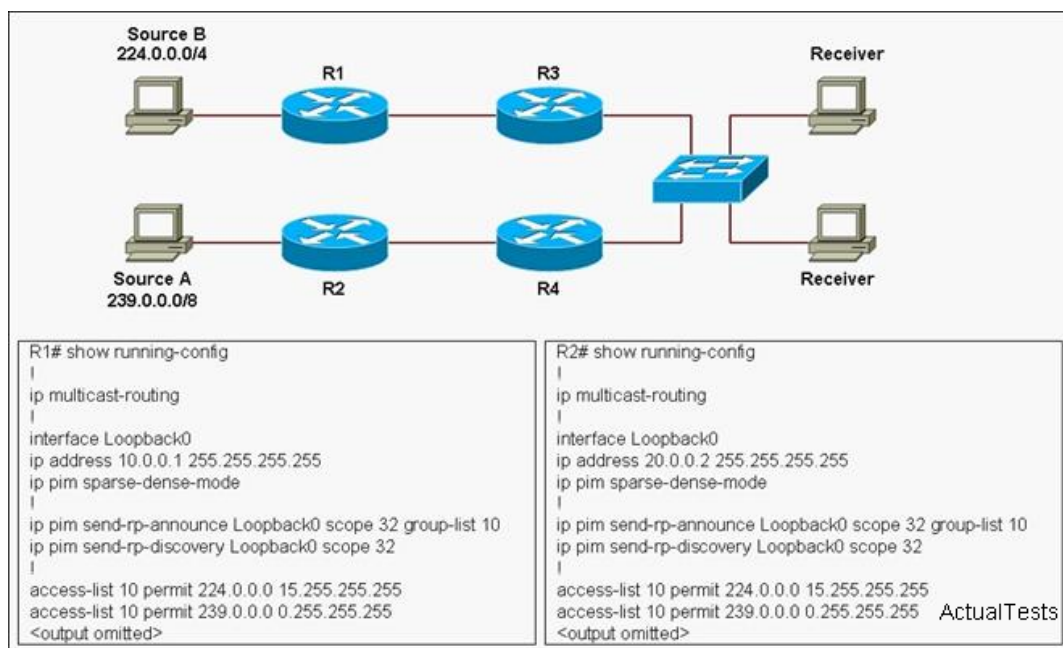
- Configure each switch with PIM sparse mode and a separate instance of PIM dense mode. Specify switch 1 as the root for the RPI Multicast Server. Specify switch 2 as the root for the CMU Multicast Server.
- Configure each switch with PIM sparse mode and a separate instance of PIM dense mode. Leave each multicast server as the root of its own multicast tree.
- Configure each switch with PIM sparse-dense mode. Configure switch 3 as a rendezvous point for the RPI multicast stream.
- Configure each switch with PIM sparse mode and a separate instance of PIM dense mode. Allow the switches to elect their own root for each multicast tree.
- Configure each switch with PIM sparse mode.

F. Configure each switch with PIM dense mode.

Answer: C

QUESTION NO: 138

Refer to the exhibit. All routers have Protocol Independent Multicast (PIM) enabled interfaces. On the basis of the configuration provided on routers R1 and R2, which router will take on the function of rendezvous point (RP) for the multicast network?



- A. both routers R1 and R2
- B. router R1
- C. router R2
- D. none of the routers since they are not configured with static RP

Answer: C

QUESTION NO: 139

Which two statements are true about the rendezvous point (RP) in a multicast network? (Choose two.)

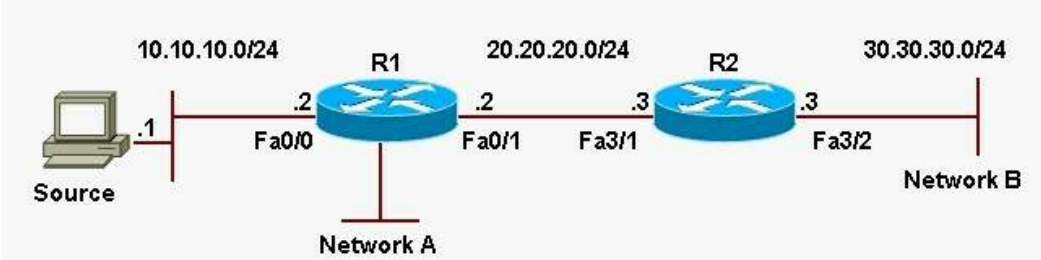
- A. The multicast sources must register with the RP to form the multicast distribution tree.
- B. An RP is required only in networks running Protocol Independent Multicast dense mode (PIM DM).
- C. The multicast receivers must register with the RP to form the multicast distribution tree.
- D. An RP is required only in networks running Protocol Independent Multicast sparse mode (PIM SM).

- E. An RP is required only in networks running Protocol Independent Multicast sparse-dense mode (PIM-SDM).
- F. To form the multicast distribution tree, the multicast sources register with and the receivers join the RP.

Answer: D,F

QUESTION NO: 140

Refer to the exhibit. IP multicast for group address 224.1.1.1 has been enabled on all routers in the network. Hosts on Network A receive the multicast traffic. However, hosts on Network B do not. On the basis of outputs provided, what could be the cause of the problem?



```

R2# show ip pim neighbor
PIM Neighbor Table
Neighbor Address  Interface          Uptime  Expires  Ver  Mode
20.20.20.2       FastEthernet3/1    2d00h   00:01:15 v2
  
```

```

R2# show ip mroute 224.1.1.1
IP Multicast Routing Table
<output omitted>
(*, 224.1.1.1), 00:10:42/stopped, RP 0.0.0.0, flags: DC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet3/1, Forward/Dense, 00:10:42/00:00:00
    FastEthernet3/2, Forward/Dense, 00:10:42/00:00:00
(10.10.10.1, 224.1.1.1), 00:01:10/00:02:48, flags:
  Incoming interface: FastEthernet3/2, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet3/1, Forward/Dense, 00:01:10/00:00:00
    FastEthernet3/2, Forward/Dense, 00:00:16/00:00:00
  
```

ActualTests

- A. Because of RPF failure, Router R2 does not forward multicast packets to Network B.
- B. The multicast packets are sourced from a server with an unspecified IP address.
- C. Router R2 does not see the upstream router R1 as a PIM neighbor.
- D. Router R2 does not have an RP configured on the multicast network.

Answer: A

Explanation:

Multicast routing uses a mechanism called Reverse Path Forwarding (RPF) to prevent forwarding loops and to ensure the shortest path from the source to the receivers. If there is RPF failure, the

router does not forward the multicast packets to neighbor.

QUESTION NO: 141

Refer to the exhibit. IP multicast configuration changes have been made on several routers. However, the IP multicast table displayed in the exhibit still does not reflect the changes. What should be done to display the new routing table information?

```

SW1# show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
  Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
  Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

<Output omitted>
ActualTests

```

- A. Issue the clear ip route * privileged EXEC command.
- B. Issue the clear ip igmp group privileged EXEC command.
- C. Issue the no ip mroute-cache privileged EXEC command.
- D. Issue the clear ip mroute privileged EXEC command.

Answer: D

QUESTION NO: 142

Which three IP multicast group concepts are true? (Choose three.)

- A. If a packet is sent to a multicast group address, the multicast frame contains the source multicast address.
- B. A router does not have to be a member of a multicast group to receive multicast data.
- C. A router must be a member of a multicast group to send to the group.
- D. A router does not have to be a member of a multicast group to send to the group.
- E. If a packet is sent to a multicast group address, all members of the multicast group will receive it.
- F. A router must be a member of a multicast group to receive multicast data.

Answer: D,E,F

QUESTION NO: 143

Which three IP multicast address related statements are true? (Choose three.)

- A. Multicast addresses 224.0.0.13 and 224.0.0.22 are reserved link-local addresses used by PIMv2 and IGMPv3.
- B. Because they would map to overlapping IP multicast MAC addresses, multicast addresses 224.0.1.1 and 238.1.1.1 could not be used together.
- C. Multicast addresses 224.0.0.0 through 224.0.0.255 are always forwarded because they are transmitted with Time to Live (TTL) greater than 1.
- D. Multicast addresses 224.0.0.5 and 224.0.0.6 are source multicast addresses for OSPF routers.
- E. The administratively scoped multicast addresses 239.0.0.0 through 239.255.255.255 are similar in purpose to RFC 1918 private unicast addresses.
- F. Multicast address 224.0.1.1 has been reserved for the Network Time Protocol (NTP) by the IANA.

Answer: A,E,F

QUESTION NO: 144

Which two multicast IP addresses can be represented by the multicast MAC address 0100.5e0A.0A07? (Choose two.)

- A. 228.10.10.7
- B. 229.11.10.7
- C. 229.138.10.7
- D. 229.138.10.8
- E. 228.10.10.8
- F. 228.10.138.7

Answer: A,C

QUESTION NO: 145

A network administrator assigns a multicast address of 239.255.8.5 to an application running on a device with an Ethernet MAC address of 01.b2.7d.05.f1.80. Which Layer 2 multicast address will this device use?

- A. 01.00.5e.05.f1.80

- B. 01.b2.7d.05.f1.80
- C. ff.ff.ff.ff.ff
- D. 01.00.5e.7F.08.05
- E. 01.b2.7d.0a.08.05

Answer: D

QUESTION NO: 146

Which two statements are true about the rendezvous point (RP) in a multicast network? (Choose two.)

- A. The multicast sources must register with the RP to form the multicast distribution tree.
- B. An RP is required only in networks running Protocol Independent Multicast dense mode (PIM DM).
- C. The multicast receivers must register with the RP to form the multicast distribution tree.
- D. An RP is required only in networks running Protocol Independent Multicast sparse mode (PIM SM).
- E. An RP is required only in networks running Protocol Independent Multicast sparse-dense mode (PIM-SDM).
- F. To form the multicast distribution tree, the multicast sources register with and the receivers join the RP.

Answer: D,F

QUESTION NO: 147

Which two multicast protocol statements are true? (Choose two.)

- A. Sparse mode multicast uses a pull model to send multicast traffic to where it is requested.
- B. Dense mode multicast uses a push model to flood traffic throughout the network and then prunes the unwanted traffic.
- C. Dense mode multicast requires explicit join messages from their members.
- D. Sparse mode uses reverse path forwarding (RPF) to prune off redundant flows.
- E. The primary use of sparse mode multicast is for test labs and router performance testing.

Answer: A,B

QUESTION NO: 148

What are three methods that may be used to allow a switch to determine which ports to forward IP multicast messages to? (Choose three.)

- A. IGMP
- B. IGMP snooping
- C. PIM
- D. CGMP
- E. static assignment
- F. CDP

Answer: B,D,E

QUESTION NO: 149

Refer to the exhibit. Which two conclusions can be drawn from the debug ip igmp output? (Choose two.)

```
RTA# debug ip igmp
12:32:51:065: IGMP: Send v2 Query on Ethernet1 to 224.0.0.1
12:32:51:069: IGMP: Set report delay time to 9.4 seconds for 224.0.1.40 on Ethernet1
12:32:56:909: IGMP: Received v1 Report from 192.168.9.1 (Ethernet1) for 239.255.0.1
12:32:56:917: IGMP: Starting old host present timer for 239.255.0.1 on Ethernet1
12:33:01:065: IGMP: Send v2 Report for 224.0.1.40 on Ethernet1
12:33:01:069: IGMP: Received v2 Report from 192.168.9.4 (Ethernet1) for 224.0.1.40
12:33:51:065: IGMP: Send v2 Query on Ethernet1 to 224.0.0.1
```

- A. The router sent an IGMP version 2 query out interface Ethernet1 at multicast address 224.0.0.1.
- B. Router RTA received an IGMP report version 1 from host 192.168.9.1.
- C. Reverse Path Forwarding (RPF) is enabled for 192.168.9.4.
- D. Reverse Path Forwarding (RPF) is enabled for 224.0.1.40.
- E. The IP multicast groups are 224.0.0.1, 224.0.1.40, and 239.255.0.1.
- F. IP PIM RP mapping is static.

Answer: A,B

QUESTION NO: 150

Refer to the exhibit. What statement is true based upon the configuration in the exhibit?

```
Sw1#show run
<output omitted>
ip multicast-routing

interface GigabitEthernet0/0
 ip address 172.16.4.16 255.255.255.0
 ip pim sparse-dense-mode ActualTests
```

- A. CGMP version 2 is being used.
- B. The switch will use sparse mode first, dense mode second.
- C. A rendezvous point must be configured in order for the interface to operate in sparse mode.
- D. IGMP version 2 is being used.
- E. The rendezvous point is IP address 172.16.4.16.
- F. If a rendezvous point is configured, the interface cannot operate in dense mode.

Answer: C

QUESTION NO: 151

In computer networking a multicast address is an identifier for a group of hosts that have joined a multicast group. Multicast addressing can be used in the Link Layer (OSI Layer 2), such as Ethernet Multicast, as well as at the Internet Layer (OSI Layer 3) as IPv4 or IPv6 Multicast. Which two descriptions are correct regarding multicast addressing?

- A. The first 23 bits of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
- B. The last 3 bytes (24 bits) of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
- C. The first 3 bytes (24 bits) of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
- D. To calculate the Layer 2 multicast address, the host maps the last 23 bits of the IP address into the last 24 bits of the MAC address. The high-order bit is set to 0.

Answer: C,D

Explanation:

The point of this question is the form of multicast MAC address, and the conversion between the multicast MAC address and IP address.

The multicast MAC address is 6 bytes(48 bits), the first 3 bytes (24 bits) of the multicast MAC address are 0x01-00-5E, the last 3 bytes(24 bits) of the multicast MAC address =0 + 23 bit(the last

23 bit of the IP address). "0x01-00-5E" is a reserved value that indicates a multicast application. So option B and D are correct.

QUESTION NO: 152

Refer to the exhibit. Which statement is true?

```
hostname RAR1
!  
<output omitted>
!  
router bgp 100
  neighbor 172.16.1.2 remote-as 200
  neighbor 172.16.1.2 distribute-list 101 in
!  
access-list 101 permit ip 10.10.0.0 0.0.0.0 255.255.224.0 0.0.0.0
```

- A. Only traffic with a destination from 10.10.0.0/19 will be permitted.
- B. Router RAR1 will accept only route 10.10.0.0/19 from its BGP neighbor.
- C. Only traffic going to 10.10.0.0/19 will be permitted.
- D. Router RAR1 will send only route 10.10.0.0/19 to its BGP neighbor.

Answer: B

QUESTION NO: 153

If no metric is specified for the routes being redistributed into IS-IS, what metric value is assigned to the routes?

- A. 20
- B. 10
- C. 0
- D. 1

Answer: C

QUESTION NO: 154

Refer to the exhibit. Which three statements accurately describe the result of applying the exhibited route map? (Choose three.)

```

router eigrp 1
 redistribute ospf 1 route-map ospf-to-eigrp
 default-metric 20000 2000 255 1 1500
 !
 !
 route-map ospf-to-eigrp deny 10
  match tag 6
  match route-type external type-2
 !
 route-map ospf-to-eigrp permit 20
  match ip address prefix-list pfx
  set metric 40000 1000 255 1 1500
 !
 route-map ospf-to-eigrp permit 30
  set tag 8

```

ActualTests

- A. The map prohibits the redistribution of all type 2 external OSPF routes.
- B. All routes that do not match clauses 10 and 20 of the route map are redistributed with their tags set to 8.
- C. The map prohibits the redistribution of all external OSPF routes with tag 6 set.
- D. The map permits the redistribution of all type 1 external OSPF routes.
- E. The map redistributes into EIGRP all routes that match the pfx prefix list and the five metric values?0000, 1000, 255, 1, and 1500.
- F. The map prohibits the redistribution of all type 2 external OSPF routes with tag 6 set.

Answer: B,D,F

Explanation:

The route-map command is used to configure policy routing, which is often a complicated task. A route map is defined using the syntax shown in the figure.

Syntax:

```
RouterA(Config)#route-map map-tag [permit | deny ] <Sequence Number>
```

```
RouterA(Config-map-router)#
```

The map-tag is the name, or ID, of the route map. This map-tag can be set to something easily recognizable name. The route-map command changes the mode on the router to the route-map configuration mode, from there conditions can be configured for the route map.

In this example, the first route-map entry (10) will deny all type 2 routes with a tag of 6 (both must be true).

The second route map (20) will set the metrics for all routes matching the pfx prefix list. Choice A is incorrect because the metrics shown do not need to match. These metrics will be assigned to the routes that match.

The last route-map (30) will set all other routes that did not already match 10 or 20 with a tag of 8.

QUESTION NO: 155

Which three restrictions apply to OSPF stub areas? (Choose three.)

- A. The area cannot be a backbone area.
- B. Inter area routes are suppressed.
- C. The area has no more than 10 routers.
- D. No virtual links are allowed.
- E. No autonomous system border routers are allowed.
- F. Redistribution is not allowed unless the packet is changed to a type 7 packet.

Answer: A,D,E

QUESTION NO: 156

Which three IP multicast address related statements are true? (Choose three.)

- A. Multicast addresses 224.0.0.13 and 224.0.0.22 are reserved link-local addresses used by PIMv2 and IGMPv3.
- B. Because they would map to overlapping IP multicast MAC addresses, multicast addresses 224.0.1.1 and 238.1.1.1 could not be used together.
- C. Multicast addresses 224.0.0.0 through 224.0.0.255 are always forwarded because they are transmitted with Time to Live (TTL) greater than 1.
- D. Multicast addresses 224.0.0.5 and 224.0.0.6 are source multicast addresses for OSPF routers.
- E. The administratively scoped multicast addresses 239.0.0.0 through 239.255.255.255 are similar in purpose to RFC 1918 private unicast addresses.
- F. Multicast address 224.0.1.1 has been reserved for the Network Time Protocol (NTP) by the IANA.

Answer: A,E,F

QUESTION NO: 157

Which three IP multicast group concepts are true? (Choose three.)

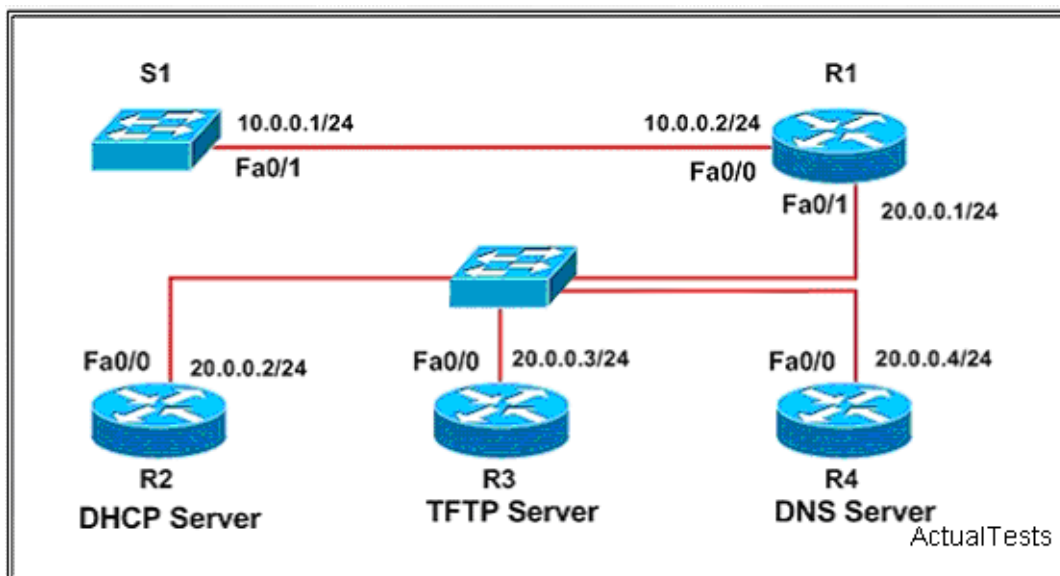
- A. If a packet is sent to a multicast group address, the multicast frame contains the source multicast address.
- B. A router does not have to be a member of a multicast group to receive multicast data.
- C. A router must be a member of a multicast group to send to the group.
- D. A router does not have to be a member of a multicast group to send to the group.
- E. If a packet is sent to a multicast group address, all members of the multicast group will receive it.

F. A router must be a member of a multicast group to receive multicast data.

Answer: D,E,F

QUESTION NO: 158

Refer to the exhibit. Router R1 is being used as a relay device for autoconfiguration of switch S1. Which configuration will accomplish this?



Select the best response.

- A. S1(config)# interface fastethernet 0/1
S1(config-if)# ip helper-address 10.0.0.2
- B. S1(config)# interface fastethernet 0/1
S1(config-if)# ip helper-address 20.0.0.1
- C. R1(config)# interface fastethernet 0/0
R1(config-if)# ip helper-address 20.0.0.2
R1(config-if)# ip helper-address 20.0.0.3
R1(config-if)# ip helper-address 20.0.0.4
R1(config-if)# exit
R1(config)# interface fastethernet 0/1
R1(config-if)# ip helper-address 10.0.0.1
- D. R1(config)# interface fastethernet 0/0
R1(config-if)# ip helper-address 20.0.0.1
R2(config)# interface fastethernet 0/0
R2(config-if)# ip helper-address 20.0.0.1
R3(config)# interface fastethernet 0/0
R3(config-if)# ip helper-address 20.0.0.1
R4(config)# interface fastethernet 0/0
R4(config-if)# ip helper-address 20.0.0.1


```
E. S1(config)# interface fastethernet 0/1
S1(config-if)# ip helper-address 10.0.0.2
R2(config)# interface fastethernet 0/0
R2(config-if)# ip helper-address 20.0.0.1
R3(config)# interface fastethernet 0/0
R3(config-if)# ip helper-address 20.0.0.1
R4(config)# interface fastethernet 0/0
R4(config-if)# ip helper-address 20.0.0.1
```

Answer: C

Explanation:

The point of this question is about the command "ip helper-address".

The ip helper-address command allows you to forward broadcast to unicast. One common application that requires helper addresses is Dynamic Host Configuration Protocol (DHCP).


All of the following conditions must be met in order for a UDP or IP packet to be helpered by the ip helper-address command:

1. The MAC address of the received frame must be all-ones broadcast address (ffff.ffff.ffff).
2. The IP destination address must be one of the following: all-ones broadcast (255.255.255.255), subnet broadcast for the receiving interface; or major-net broadcast for the receiving interface if the no ip classless command is also configured.
3. The IP time-to-live (TTL) value must be at least 2.
4. The IP protocol must be UDP (17).
5. The UDP destination port must be for TFTP, Domain Name System (DNS), Time, NetBIOS, ND, BOOTP or DHCP packet, or a UDP port specified by the ip forward-protocol udp global configuration command.

From this question, we will find the configuration must be done in CK-R1 which is the relay.

QUESTION NO: 159

Refer to the exhibit. This network is running IS-IS. Router RTC is inside Cloud#1. From the output on RTD, which two statements are true? (Choose two.)



```
RTD#show clns
```

```
Global CLNS Information:
```

```
3 Interfaces Enabled for CLNS
```

```
NET: 49.0150.1921.6800.4004.00
```

```
Configuration Timer: 60, Default Holding Timer: 300, Packet Lifetime 64
```

```
ERPDU's requested on locally generated packets
```

```
Running IS-IS in IP-only mode (CLNS forwarding not allowed)
```

```
RTD#show clns neighbors
```

System Id	Interface	SNPA	State	Holdtime	Type	Protocol	
RTE	Fa0/1	0019.06b7.fd5f	Up	9	L2	IS-IS	ActualTests
RTC	Se0/1/1	*HDLC*	Up	22	L1	IS-IS	

- A. Because it is running IS-IS in IP-only mode, RTD will not transmit CLNS packets.
- B. IP routing updates between RTC and RTD will be IP datagrams that are encapsulated by HDLC.
- C. The SNPA for RTE is the source data link address that is used to transmit frames to RTD.
- D. IP routing updates between RTC and RTD will be CLNS datagrams that are encapsulated by HDLC.
- E. The NET for RTE will be 49.0150.0019.06b7.fd5f.00.
- F. IP routing updates between RTC and RTD will be IP datagrams that are encapsulated by CLNS.

Answer: C,D

QUESTION NO: 160

Refer to the exhibit. What three statements are true about the IS-IS configuration? (Choose three.)

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 ip router isis
!
interface Serial 0/0/1
 ip address 10.2.2.2 255.255.255.0
 ip router isis
!
<output omitted>
router isis
 net 49.0001.0002.0003.0004.0005
```

- A. The router is in area 49.0001.0002.
- B. The router acts as a Level 1-2 router.
- C. CLNS routing is enabled for the router.
- D. The router has a system ID of 0003.0004.
- E. The interfaces have a functional IS-IS configuration.
- F. The network service access point selector (NSEL) byte has a value of 0.

Answer: B,E,F

Explanation:

This router is using the default settings, which is a Level 1-2 router. The "ip router isis " command is the only IS-IS command needed for an interface.

The area is 49.001A. The router's system ID is 0002.0003.0004. The easiest way to figure this out is to start from the right and work towards the left. The last two numbers of the NET are the NSEL; they are always 00 on a router. The next 12 numbers (separated into 3 groups of 4 numbers) are the system ID. On Cisco routers, the system ID is always this length-6 bytes. Anything to the left of the system ID is the area ID.

Reference: <http://www.ciscopress.com/articles/article.asp?p=101756&rl=1>

QUESTION NO: 161

Refer to the exhibit. On the basis of the output, what two conclusions can be reached? (Choose two.)

```

R2# show ip protocols
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    FastEthernet0/0
    Loopback0
    Serial0/0/1
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.10.10.10      115          00:00:02
    10.30.30.30      115          00:00:03
  Distance: (default is 115)
R2#

```

ActualTests

- A. The default administrative distance has been changed.
- B. IS-IS is not enabled.
- C. Three physical interfaces are taking part in integrated IS-IS.
- D. IS-IS is not redistributing any other routing protocols.

E. There are two neighboring routers sending IS-IS routing information.

F. Address summarization is configured.

Answer: D,E

Explanation:

Verifying IS-IS Configuration

To verify the IS-IS configuration and IP functionality of an Integrated IS-IS network, use the following commands; these commands can also be useful for troubleshooting problems with the IS-IS network: show ip protocols- Displays the active IP routing protocols, the interfaces on which they are active, and the networks for which they are routing. show ip route [address [mask]] | [protocol [process-id]]- Displays the IP routing table. You can specify the details for a particular route or a list of all routes in the routing table from a particular routing protocol process.

Example 6-3 is sample output from the show ip protocols command that displays information about IP routing being done by Integrated IS-IS. IS-IS is running, it is not redistributing any other protocols, and address summarization has not been configured. Example 6-3 also shows that interfaces FastEthernet 0/0, Loopback 0, and Serial 0/0/1 are taking part in Integrated IS-IS, that there are two sources of routing information (the neighboring routers), and that the administrative distance of Integrated IS-IS is 115.

Example 6-3, "show ip protocols" Command to Examine IS-IS:

Reference : Authorized Self-Study Guide Building Scalable Cisco Internetworks (BSCI), Third Edition

<http://safari.ibmpressbooks.com/1587052237/ch06lev1sec4>

QUESTION NO: 162

Which three options are supported as address allocation mechanisms for DHCP on Cisco routers? (Choose three.)

Select 3 response(s).

A. The IP address can be automatically assigned to a host.

B. The IP address can be assigned as a random hash value of the burned-in-address of the lowest-numbered LAN interface on the router.

C. The network administrator can assign a specific IP address to a specific host MAC address.

D. The IP address can be assigned from configured pools in a reverse lexicographical order.

E. The IP address can be assigned to a host for a limited time or until the host explicitly releases the address.

F. The IP address can be assigned to a host until the host usurps the assigned value using its own dynamic override mechanism.

Answer: A,C,E

QUESTION NO: 163

Refer to the exhibit. Which statement is true about the configuration?

```
<output omitted>

ip dhcp pool 0
  network 10.10.0.0/16
  domain-name cisco.com
  dns-server 10.10.10.50 10.10.20.50
  netbios-name-server 10.10.10.51 10.10.20.51
!
ip dhcp pool 1
  network 10.10.10.0/24
  default-router 10.10.10.100 10.10.10.101
  lease 30
!
ip dhcp pool 2
  network 10.10.20.0/24
  default-router 10.10.20.100 10.10.20.101
  lease 30
```

ActualTests

- A. Hosts belonging to DHCP pool 1 and pool 2 will retain their IP settings for 30 hours before they must renew.
- B. Hosts will receive IP settings from pool 1 until the addresses run out, and then hosts will receive the settings from pool 2.
- C. Hosts in the 10.10.20.0/24 subnet will use 10.10.20.50 as its DNS server.
- D. DHCP pool 0 needs to have the ip dhcp excluded-address command to exclude the default router and DNS servers.

Answer: C

QUESTION NO: 164

Refer to the exhibit. Which statement is true about the router configuration that is shown?
Select the best response

```
ip dhcp pool 1
  network 172.16.1.0/24
  domain-name cisco.com
  dns-server 172.16.1.102
  netbios-name-server 172.16.1.103
  default-router 172.16.1.100 172.16.1.101
  lease 30
!
ip dhcp pool 2
  network 172.16.2.0/24
  domain-name cisco.com
  dns-server 172.16.2.102
  netbios-name-server 172.16.2.103
  default-router 172.16.2.100 172.16.2.101
  lease 30
```

ActualTests

- A. Additional DHCP options will be imported from another DHCP server.
- B. The DHCP server pools need to be bound to an interface to operate.
- C. This configuration will provide IP configuration information to two different subnets.
- D. Additional DHCP option information needs to be imported from another DHCP server.
- E. If the router hands out all the addresses in pool 1, then it will supply addresses from pool 2.

Answer: C

QUESTION NO: 165

Using the rules for IPv6 addressing, how can the address 2031:0000:240F:0000:0000:09C0:123A:121B be rewritten? Select the best response.


- A. 2031:0:240F::09C0:123A:121B
- B. 2031::240F::09C0:123A:121B
- C. 2031::240F:9C0::123A:121B
- D. 2031::240F:::09C0:123A:121B

Answer: A

QUESTION NO: 166

Refer to the exhibit. Which configuration would correctly enable Host A to acquire an IP address from the DHCP server that is located at 10.1.2.10/24?

Select the best response

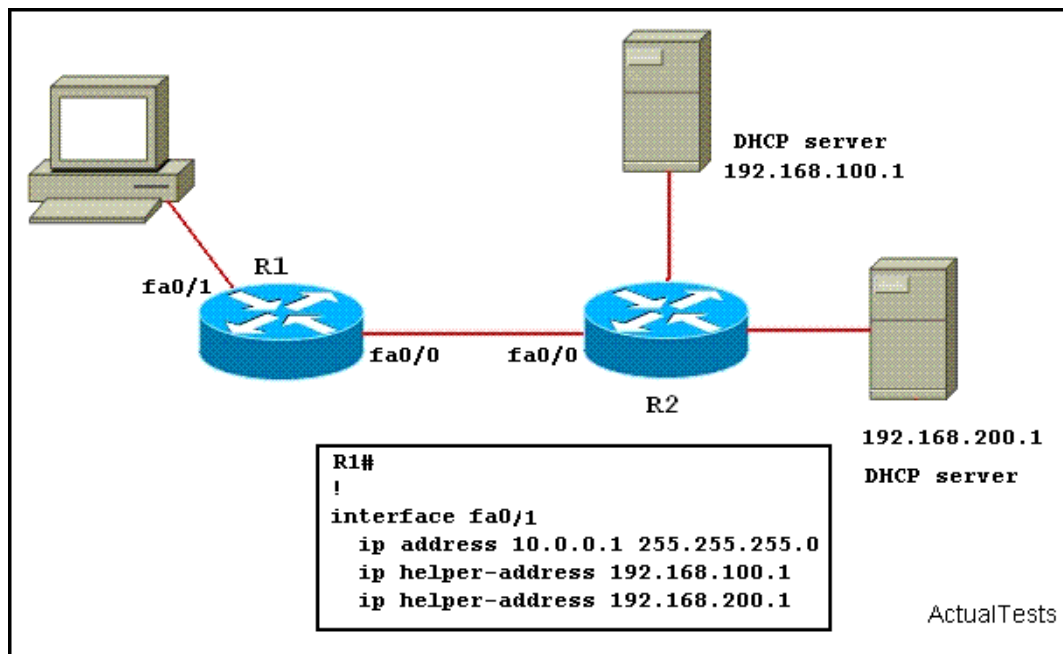


- A. ASw1(config)# interface range FastEthernet 0/1 - 24
ASw1(config-if-range)# ip forward-protocol udp 67
ASw1(config-if-range)# ip forward-protocol udp 68
- B. ASw1(config)# interface range FastEthernet 0/1 - 24
ASw1(config-if-range)# ip helper-address 10.1.2.10
- C. RTA(config)# interface fastethernet0/0
RTA(config-if)# ip forward-protocol udp 67
RTA(config-if)# ip forward-protocol udp 68
- D. RTA(config)# interface fastethernet0/1
RTA(config-if)# ip forward-protocol udp 67
RTA(config-if)# ip forward-protocol udp 68
- E. RTA(config)# interface fastethernet0/0
RTA(config-if)# ip helper-address 10.1.2.10
- F. RTA(config)# interface fastethernet0/1
RTA(config-if)# ip helper-address 10.1.2.10

Answer: E

QUESTION NO: 167

Refer to the exhibit. Based upon the information in the exhibit, which statement is true?



- A. DHCP requests from the host will be rebroadcasted to R2.
- B. To complete this configuration, the R1 fa0/0 interface must be configured with the ip helper-addresses command.
- C. To complete this configuration, the R2 fa0/0 interface must be configured with the ip helper-addresses command.
- D. R1 will forward all DHCP requests to both 192.168.100.1 and 192.168.200.1 as unicast messages.
- E. R1 will forward DHCP requests to 192.168.100.1. If there is no response, R1 will then forward the requests to 192.168.200.1.

Answer: D

Explanation:

A DHCP relay agent is any host that forwards DHCP packets between clients and servers. Relay agents receive DHCP messages and then generate a new DHCP message to send out on another interface. The agents forward requests and replies between clients and servers when they are not on the same physical subnet.

The Cisco IOS DHCP relay agent is enabled on an interface only when the ip helper-address is configured.

If multiple helper-addresses are configured, it tries to get response from first, if no response got from the first helper address then sends the request to second one.

QUESTION NO: 168

Which two statements are true about the rendezvous point (RP) in a multicast network? (Choose two.)

- A. The multicast sources must register with the RP to form the multicast distribution tree.
- B. An RP is required only in networks running Protocol Independent Multicast dense mode (PIM DM).
- C. The multicast receivers must register with the RP to form the multicast distribution tree.
- D. An RP is required only in networks running Protocol Independent Multicast sparse mode (PIM SM).
- E. An RP is required only in networks running Protocol Independent Multicast sparse-dense mode (PIM-SDM).
- F. To form the multicast distribution tree, the multicast sources register with and the receivers join the RP.

Answer: D,F

QUESTION NO: 169

What two situations could require the use of multiple routing protocols? (Choose two)

- A. when all equipment is manufactured by Cisco
- B. when there are multiple paths to destination networks
- C. because having multiple routing protocols confuses hackers
- D. when migrating from an older Interior Gateway Protocol (IGP) to a new IGP
- E. when using UNIX host-based routers
- F. when smaller broadcast domains are desired

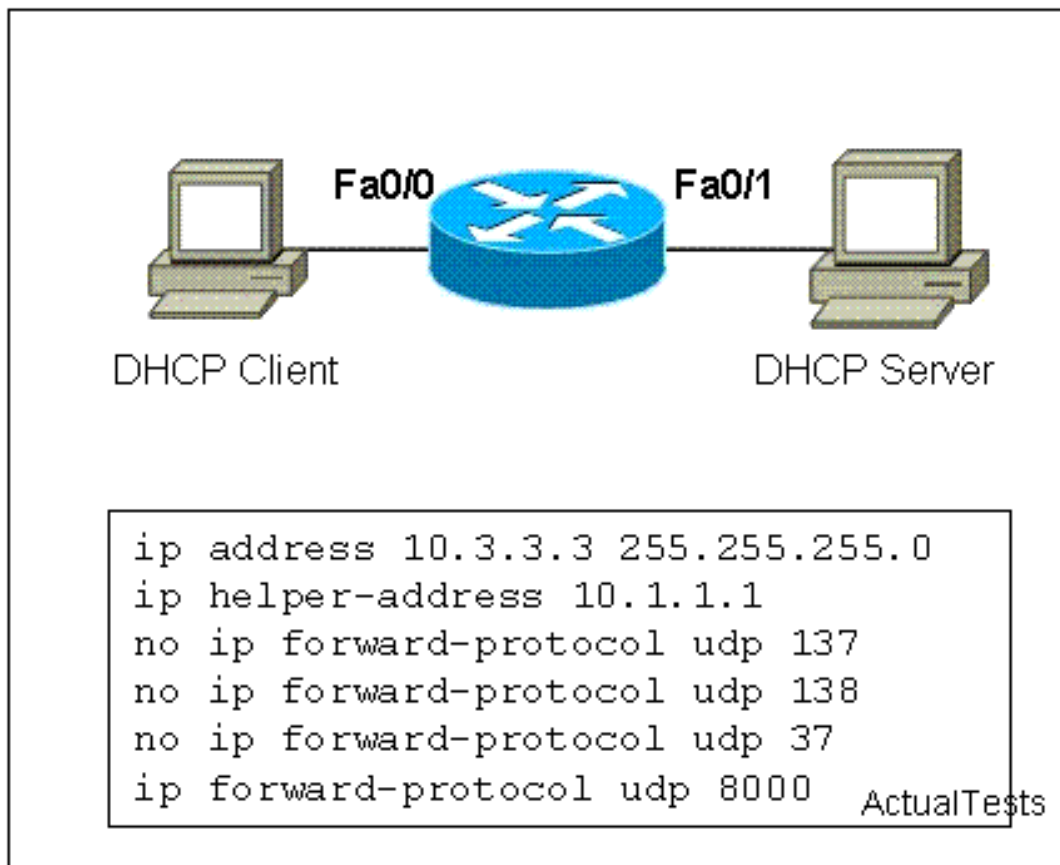
Answer: D,E

Explanation:

Simple routing protocols work well for simple networks, but networks grow and become more complex. While running a single routing protocol throughout your entire IP internetwork is desirable, multiprotocol routing is common for a number of reasons, including company mergers, multiple departments managed by multiple network administrators, multivendor environments, or simply because the original routing protocol is no longer the best choice. Often, the multiple protocols are redistributed into each other during a migration period from one protocol to the other.

QUESTION NO: 170

Refer to the exhibit. Which two statements are true? (Choose two)

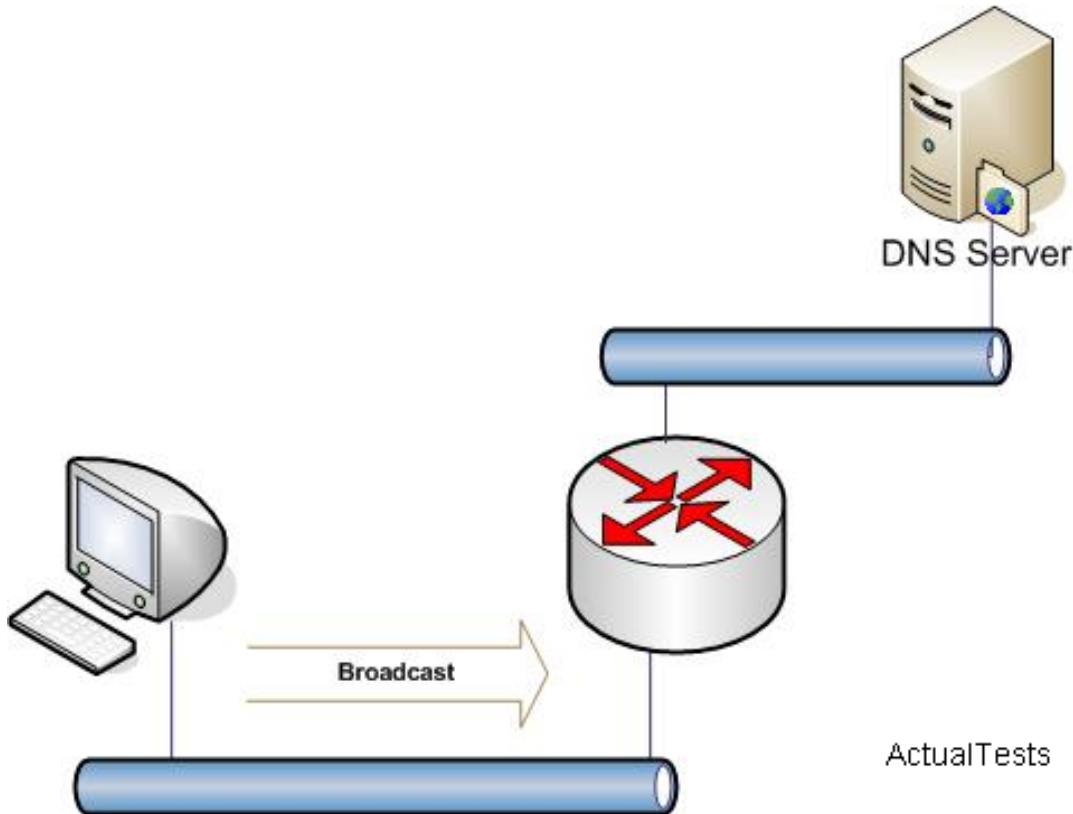


- A. DHCPDISCOVER packets will reach the DHCP server.
- B. The router will not forward DHCPDISCOVER packets because it has not been configured to do so.
- C. This configuration is applied to interface Fa0/1.
- D. DHCPDISCOVER packets will not reach the DHCP server because DHCPDISCOVER packets are broadcasts.
- E. DHCPDISCOVER packets will not reach the DHCP server because ports 67 and 68 have not been explicitly allowed by the ip forward-protocol command.
- F. This configuration is applied to interface Fa0/0.

Answer: A,E

Explanation:

While routers accept and generate broadcasts, they do not forward them. This can be quite a problem when a broadcast needs to get to a device such as a DHCP or TFTP server that's on one side of a router with other subnets on the other side.



If this PC attempts to locate a DNS server with a broadcast, the broadcast will be stopped by the router and will never get to the DNS server. By configuring the `ip helper-address` command on the router, UDP broadcasts such as this will be translated into a unicast by the router, making the communication possible. The command should be configured on the interface that will be receiving the broadcasts.

```
R1(config)#int e0 R1(config-if)#ip helper-address ? A.B.C.D IP destination address
R1(config-if)#ip helper-address 10.1.1.1
```

This command does forward eight common UDP service broadcasts by default. TIME, port 37
TACACS, port 49 DNS, port 53 BOOTP/DHCP Server, port 67 BOOTP/DHCP Client, port 68
TFTP, port 69 NetBIOS name service, port 137 NetBIOS datagram service, port 138

That's going to cover most scenarios where the `ip helper-address` command will be useful, but what about those situations where the broadcast you need forwarded is not on this list? You can use the `ip forward-protocol` command to add any UDP port number to the list. In this particular case, ports 67 and 68 were not included, so the BOOTP packets will not be sent to the DHCP server.

QUESTION NO: 171

Intermediate system to intermediate system (IS-IS), is a protocol used by network devices (routers) to determine the best way to forward datagrams through a packet-switched network, a process called routing. Which three attributes are of the IS-IS routing protocol?

- A. supports VLSM
- B. link-state routing protocol
- C. operation is similar to BGP
- D. supports two routing levels within an autonomous system

Answer: A,B,D

Explanation:

The point of this question is basic knowledge of IS-IS.

IS-IS is a link-state routing protocol, it supports VLSM and CIDR, it supports classless routing, IS-IS supports level-1 and level-2 within an autonomous system. So option A C and D are correct. incorrect Answer: Option B. BGP is path-vector protocol, the operation of BGP is AS-by-AS , not router-by-router.

QUESTION NO: 172

Which two statements are true about IS-IS? (Choose two)
Select 2 response(s).

- A. Interfaces have a default metric of 1.
- B. Level 1 is intra-area and level 2 is interarea.
- C. Level 1 routing is between an end system (ES) and IS while Level 2 routing is between an IS and IS.
- D. Hosts send ISHs while routers send ESHs.
- E. IS-IS is an open standard.
- F. Metrics depend on interface speed.

Answer: B,E

QUESTION NO: 173

Refer to the partial configurations in the exhibit. What address is utilized for DR and BDR identification on Router1? Select the best response.

```
Router1#show run

**** output omitted ****

interface serial1/1
  ipv6 address 2001:410:FFFF: 1::64/64
  ipv6 ospf 100 area 0
!
interface serial2/0
  ipv6 address 3FFF:B00:FFFF:1::2/64
  ipv6 ospf 100 area 0
!
ipv6 router ospf
  router-id 10.1.1.3
```

ActualTests

- A. the serial 1/1 address
- B. the serial 2/0 address
- C. a randomly generated internal address
- D. the configured router-id address

Answer: D

Explanation:

The point of this question is DR and BDR election of OSPFv3.

OSPF test the title of the DR and BDR election.

DR and BDR election there are two standards:

1. Interface priority, the default is 1.
2. Router-id: manually specify> Highest loopback interface> Highest active IP address

This title, there is no priority list, it is the default, manually specify the Router-id: 10.1.1.3, so hand-Router-id specified in the election has become the standard DR and BDR.

The selection of DR and BDR is the same between OSPFv2 and OSPFv3.

QUESTION NO: 174

Refer to the exhibit. What is required to complete the IPv6 routing configurations shown?
Select the best response.

```
Router1#  
interface S1/1  
  ipv6 address  
  2001:410:FFFF:1::1/64  
  ipv6 ospf 100 area 0  
  
interface S2/0  
  ipv6 address  
  3FFE:B00:FFFF:1::2/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.3  
  
Router2#  
interface S3/0  
  ipv6 address  
  3FFE:B00:FFFF:1::1/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.4
```

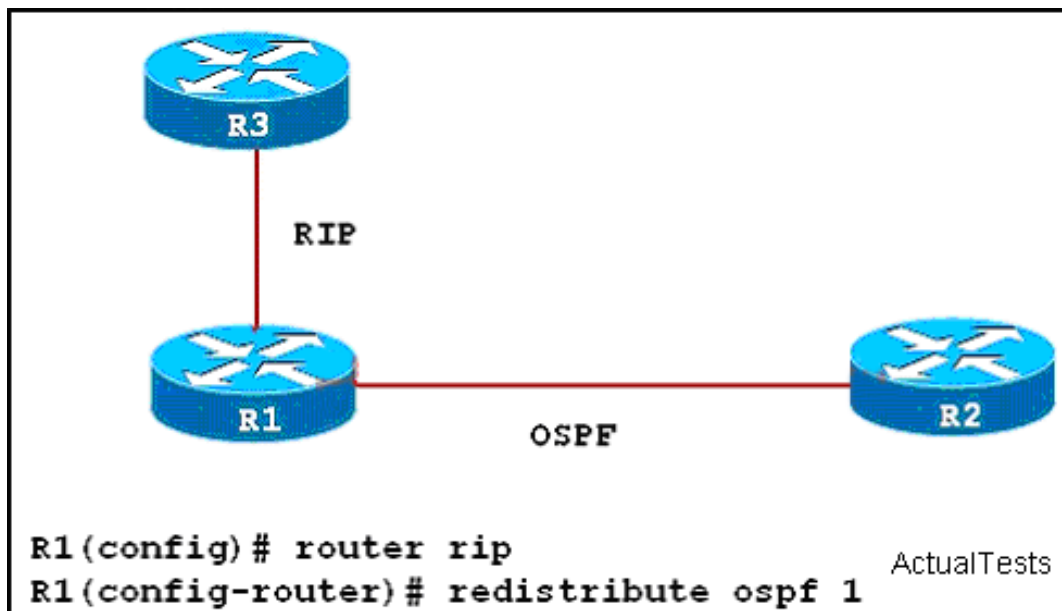
ActualTests

- A. Interface authentication must be configured.
- B. The routing processes must be configured with an area ID.
- C. IP unicast routing must be enabled.
- D. IPv4 addresses must be applied to the interfaces.

Answer: C

QUESTION NO: 175

Refer to the exhibit. Routers R1 and R2 have been configured to operate with OSPF. Routers R1 and R3 have been configured to operate with RIP. After configuring the redistribution between OSPF and RIP on R1, no OSPF routes are distributed into RIP. What should be done to correct this problem?



Select the best response.

- A. The redistribution command should be reentered with the match route-type parameter included.
- B. The redistribution command should be reentered with the route-map map-tag parameter included.
- C. The redistribution command should be reentered with the metric metric-value parameter included.
- D. Routes will first need to be distributed into another protocol, and then into RIP.

Answer: C

QUESTION NO: 176

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

Answer: A,E

Explanation:

Metrics must be set manually via configuration when redistributing into RIP and EIGRP, whereas OSPF uses a default value of 20.

Example:

EIGRP

```
router eigrp 1
redistribute ospf 1 metric 1544 5 255 1 1500
redistribute rip metric 1544 5 255 1 1500
network 15.0.0.0
```

RIP

```
router rip
version 2
redistribute eigrp 1 metric 2
redistribute ospf 1 metric 3
network 16.0.0.0
```

QUESTION NO: 177

Route redistribution (RR) has become an integral part of IP network design as the result of a growing need for disseminating certain routes across routing protocol boundaries. By default, which description is true when redistributing routes from other routing protocols into OSPF?

- A. Summarized routes are unacceptable.
- B. They will appear in the OSPF routing table as type E2 routes.
- c. Only routes with lower administrative distances will be imported.
- D. All imported routes are to be automatically summarized if possible.

Answer: B

Explanation:

The point of this question is OSPF route redistribution.

When redistributing routes from other routing protocols into OSPF, there is a default cost value 20 displaying in the OSPF routing table, it's type E2 route, it's unaltered. In opposite, type E1 route is alterable. So option B is correct answer.

QUESTION NO: 178

The IP multicast global configuration command `ip pirn send-rp-announce loopback0 scope 31 group-list 5` issued on multicast router RTA. What are the two results? (Choose two)

- A. RTA will forward RP announcements provided they are within the scope of 31.
- B. RTA will originate RP announcements with TTL set to 31.
- C. RTA will originate RP announcements for multicast groups that match access-list 5.
- D. RTA will drop all RP announcements it receives if the TTL field is greater than 31.

Answer: B,C

Explanation:

The TTL is perform a limit for the RP to control, so the TP announcements 'TTL is set 31 so limit the range of topology .and the ACL is limit the source address of the multicast packets.

QUESTION NO: 179


When implementing OSPFv3, which statement describes the configuration of OSPF areas? Select the best response.

- A. In interface configuration mode, the OSPFv3 area ID combination assigns interfaces to OSPFv3 areas.
- B. In router configuration mode, the network wildcard area ID combination assigns networks to OSPFv3 areas.
- C. In interface configuration mode, the IPv6 OSPF process area ID combination assigns interfaces to OSPFv3 areas.
- D. In router configuration mode, the IPv6 OSPF interface area ID combination assigns interfaces to OSPFv3 areas.

Answer: C

QUESTION NO: 180

Refer to the exhibit. It is desired to set up a BGP neighbor relationship between routers R1 and R4. BGP packets between them could travel through R2 or R3. What is the simplest configuration that will allow for failover? Select the best response



- A. Configure BGP neighbor relationships between all interfaces on R1 and R4.
- B. Install a direct connection between R1 and R4.
- C. Configure loopback interfaces on R1 and R4 to provide the update source address for BGP packets.
- D. Configure only one neighbor relationship between R1's 192.168.1.2 interface and R4's 172.16.10.2 interface.

Answer: C

Explanation:

The point of this question is about using loopback interface establish neighbor.

You can use two ways to configure IBGP neighbor. First is use physical address, and the second is use loopback interface. In this picture, you find that there are two way to each other, if you use physical address for failover, you need to establish two connection, and you will receive two route, this is no good. If you use loopback interface, there are only one connection, but if one way is down, it also can failover. So answer B is the best.

QUESTION NO: 181

Refer to the exhibit. Routing updates for the 192.168.1.0 network are being received from all three neighbors. Which statement is correct regarding the result of the configuration shown?

```

!
router bgp 65001
  neighbor 172.16.1.1 remote-as 65001
  neighbor 172.26.1.1 remote-as 65555
  neighbor 172.26.1.1 route-map local_pref in
  neighbor 172.30.1.1 remote-as 65510
!
route-map local_pref permit 10
  match ip address 20
  set local-preference 200
!
access-list 65 permit 192.168.1.0 0.0.0.255
!

```

ActualTests

Select the best response.

- A. The router will prefer the next hop of 172.16.1.1 for packets destined for the 192.168.1.0 network.
- B. The router will prefer the next hop of 172.26.1.1 for packets destined for the 192.168.1.0 network.
- C. The router will advertise the 192.168.1.0 network only to 172.30.1.1.

- D. The router will advertise the 192.168.1.0 network only to 172.26.1.1.
- E. The router will prefer the next hop of 172.26.1.1 for packets except those destined for the 192.168.1.0 network.

Answer: B

QUESTION NO: 182

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?

Select the best response.

- 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.

Answer: D

Explanation:

The point of this question is about the condition of establish EIGRP neighbor.

You can use these ways to troubleshoot the EIGRP connectivity problem.

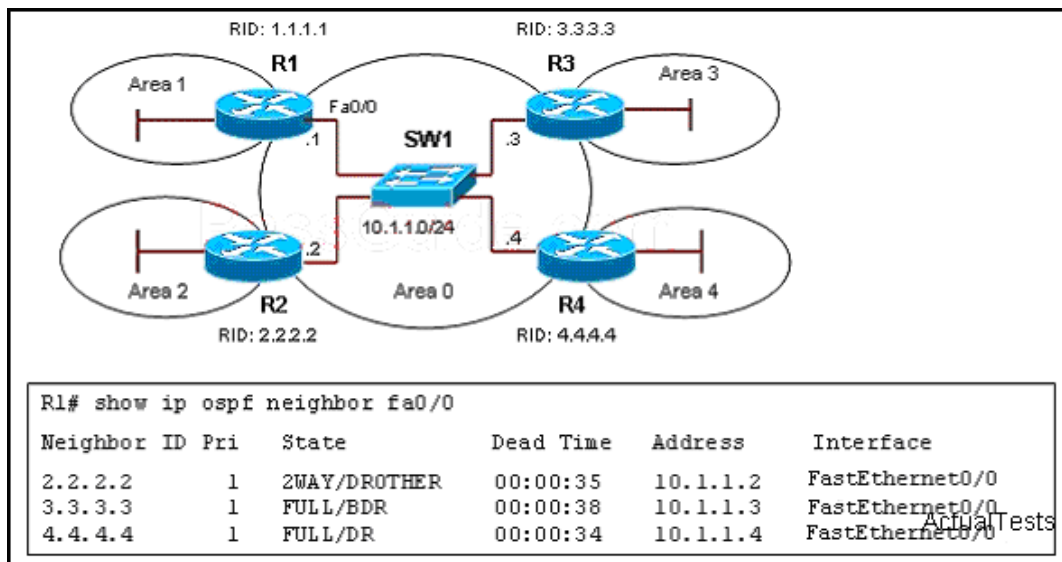
1. Whether EIGRP is enabled for the proper networks.
2. Whether the K values of EIGRP neighbors is the same.
3. Whether EIGRP autonomous number is the same.

Incorrect answers:

- *. EIGRP use multicast, not broadcast.
- *. EIGRP use multicast, not broadcast.
- *. Hello and hold timers match is the condition of establish OSPF neighbor,not EIGRP.

QUESTION NO: 183

Refer to the exhibit. OSPF is configured on all routers in the network. On the basis of the show ip ospf neighbor output, what prevents R1 from establishing a full adjacency with R2?



Select the best response.

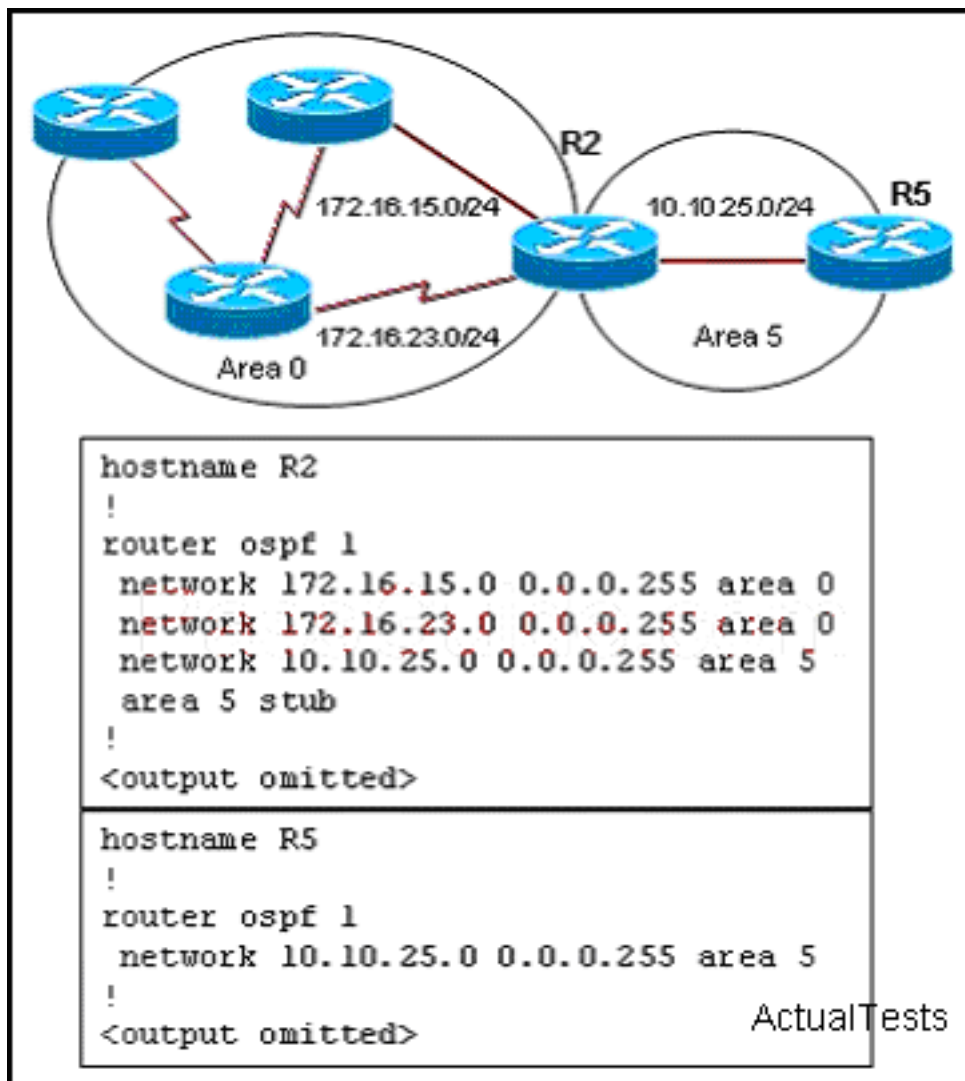
- A. Router R1 will only establish full adjacency with the DR and BDR on broadcast multi-access networks.
- B. Router R2 has been elected as a DR for the broadcast multiaccess network in OSPF area 1.
- C. Routers R1 and R2 are configured as stub routers for OSPF area 1 and OSPF area 2.
- D. Router R1 and R2 are configured for a virtual link between OSPF area 1 and OSPF area 2.
- E. The Hello parameters on routers R1 and R2 do not match.

Answer: A

QUESTION NO: 184

Refer to the exhibit. On the basis of the configuration provided, how are the Hello packets sent by R2 handled by R5 in OSPF area 5?

Select the best response.



- A. The Hello packets will be exchanged and adjacency will be established between routers R2 and R5.
- B. The Hello packets will be exchanged but the routers R2 and R5 will become neighbors only.
- C. The Hello packets will be dropped and no adjacency will be established between routers R2 and R5.
- D. The Hello packets will be dropped but the routers R2 and R5 will become neighbors.

Answer: C

Explanation:

The point of this question is the conditions of OSPF establish adjacency relationship.

For ospf, the optional capabilities must set the same between neighbors, but from the exhibit, R5 was configured as a stub area while R2 in area 0 is a normal area. So there will be no adjacent relationship established between routers R2 and R5.

QUESTION NO: 185

Which description regarding OSPF Network LSAs is correct?

- A. They are originated by Area Border Router and are sent into a single area to advertise an Autonomous System Border Router.
- B. They are originated by Area Border Routers and are sent into a single area to advertise destinations outside that area.
- C. They are originated by the DR on every multi-access network. They include all attached routers including the DR itself.
- D. They are originated by every router in the OPSF network. They include all routers on the link, interfaces, the cost of the link, and any known neighbor on the link.

Answer: C

Explanation:

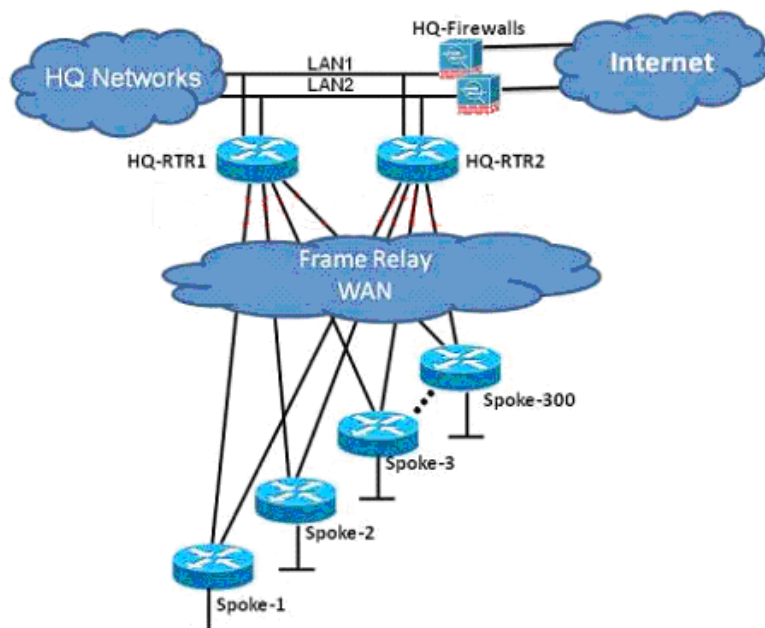
The point of this question is OSPF Network LSAs

The feature of OSPF Network LSAs is that they are generated by DR, and DR only exist on multi-access network, the use of OSPF Network LSAs is that it list all neighbors around and send it to every router which run OSPF.

incorrect Answer: OSPF Network LSAs are not originated by Area Border Routers.

QUESTION NO: 186

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.)




Select 3 response(s).

- 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

Answer: A,B,E

QUESTION NO: 187

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?



ActualTests

Select the best response.

- 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.

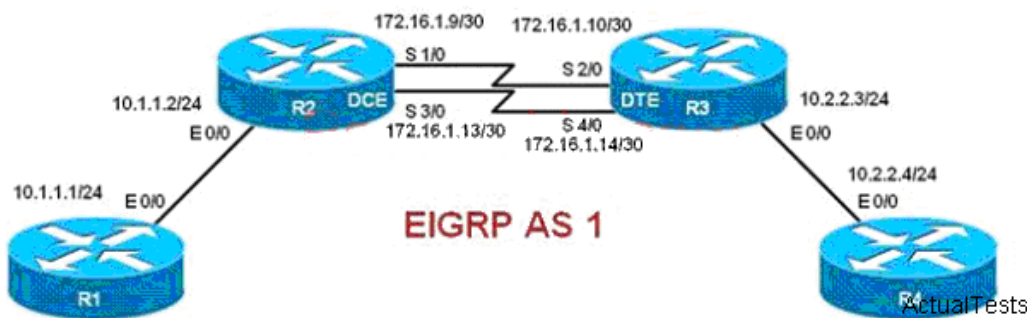
- C. Verify that the HQ-RTR1 and HQ-RTR2 routers are configured for EIGRP Nonstop Forwarding.
- D. Verify that all the spoke routers are configured for autosummarization.
- E. Verify that all the spoke routers are configured as EIGRP stub.

Answer: E

QUESTION NO: 188

Refer to the exhibit. When you examine the routing table 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 configuration change 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 answers. (Choose two.) Select 2 response(s).



- 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

Answer: B,F

QUESTION NO: 189

Refer to the exhibit. The actual speed of the serial links between R2 and R3 are 256 kb/s and 512 kb/s respectively. 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? Select the best response.



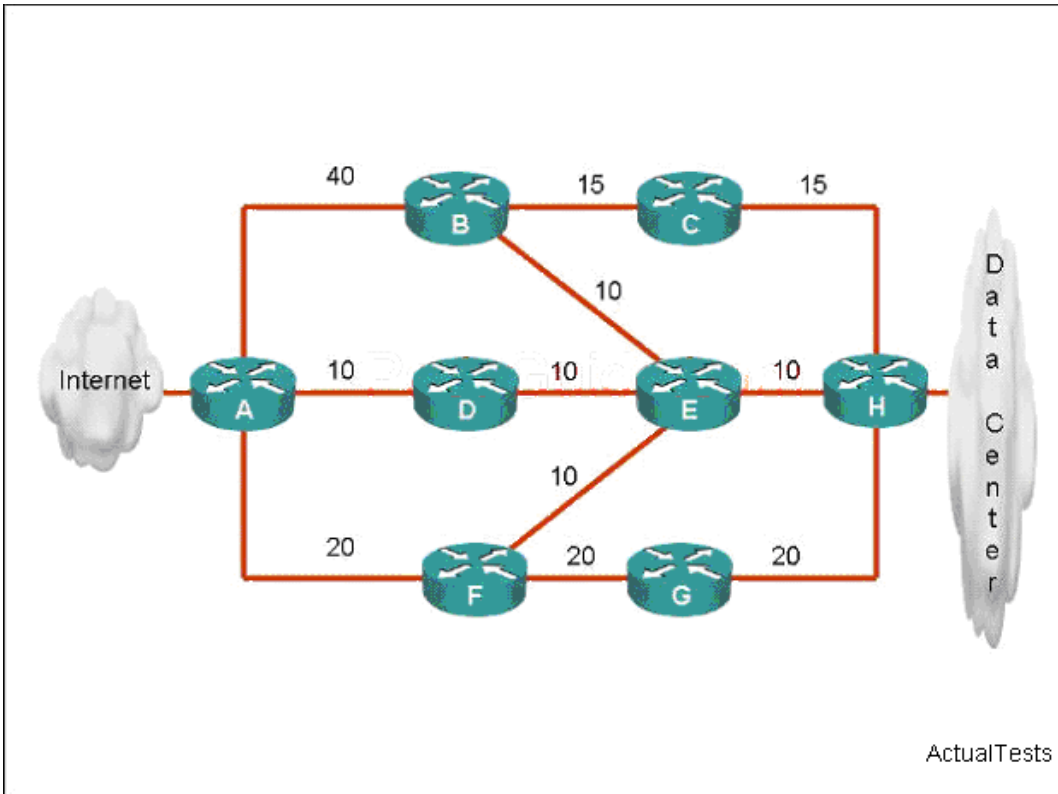
- A. EIGRP will overutilize the 512 kb/s link.
- B. The interface "delay" value used in the EIGRP metric calculation will be inaccurate on the 256 kb/s serial interface.
- C. The amount of bandwidth used for EIGRP routing protocol traffic on the 256 kb/s link can become excessive.
- D. EIGRP can load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

Answer: C

QUESTION NO: 190

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?



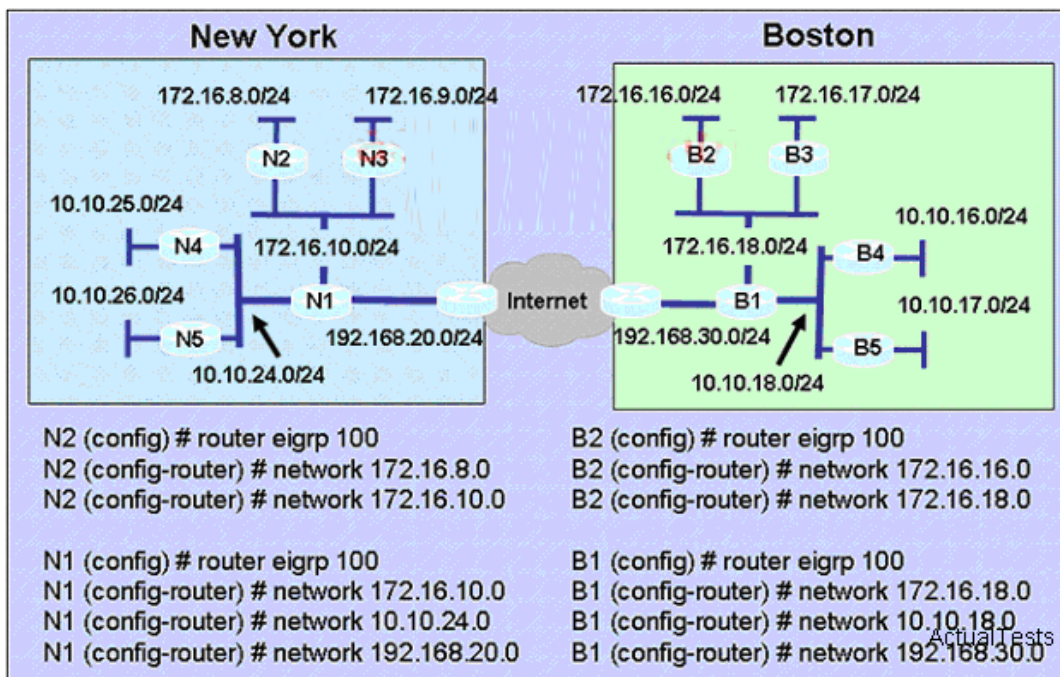
Select the best response.

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

Answer: C

QUESTION NO: 191

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?



Select the best response.

- 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 autosummarization on routers N2 and B2.
- D. Add wildcard masks to the network commands on routers N2 and B2.

Answer: A

QUESTION NO: 192

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)?


Select the best response.


```
Gateway of last resort is not set
```

```

172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D   172.16.8.0/24 [90/30720] via 172.16.10.2, 00:20:43, FastEthernet0/0
C   172.16.10.0/24 is directly connected, FastEthernet0/0
D   172.16.0.0/16 is a summary, 00:19:05, Null0
C   192.168.20.0/24 is directly connected, FastEthernet0/1
D   10.0.0.0/8 [90/30720] via 192.168.20.2, 00:14:51, FastEthernet0/1

```



```
Gateway of last resort is not set
```

```

172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D   172.16.16.0/24
    [90/30720] via 172.16.18.2, 00:06:04, FastEthernet0/0.172
C   172.16.18.0/24 is directly connected, FastEthernet0/0.172
D   172.16.0.0/16 is a summary, 00:20:05, Null0
C   192.168.20.0/24 is directly connected, FastEthernet0/1
D   10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D   10.0.0.0/8 is a summary, 00:15:51, Null0
D   10.10.16.0/24 [90/30720] via 10.10.18.3, 00:04:28, FastEthernet0/0.10
C   10.10.18.0/24 is directly connected, FastEthernet0/0.10

```

- 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.

Answer: B

QUESTION NO: 193

Refer to the exhibit. You are the network administrator responsible for the NProuter, the 10.1.1.1 router, and the 10.1.1.2 router. What can you determine about the OSPF operations from the debug output? Select the best response.


```
NProuter#debug ip ospf events
```

```
OSPF events debugging is on
```

```
NProuter#
```

```
00:02:03: OSPF: Rcv hello from 172.16.1.1 area 0 from Serial0/0 10.1.1.1
```

```
00:02:03: OSPF: Mismatched hello parameters from 10.1.1.1
```

```
00:02:03: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

```
00:02:26: OSPF: Rcv hello from 192.168.1.2 area 0 from Serial0/0 10.1.1.2
```

```
00:02:26: OSPF: Mismatched hello parameters from 10.1.1.2
```

```
00:02:26: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

ActualTests

- A. The NProuter has two OSPF neighbors in the "Full" adjacency state.
- B. The NProuter serial0/0 interface has the OSPF dead timer set to 10 seconds.
- C. The NProuter serial0/0 interface has been configured with an OSPF network type of "point-to-point".
- D. The 10.1.1.1 and 10.1.1.2 routers are not using the default OSPF dead and hello timers setting.
- E. The "Mismatched" error is caused by the expiration of the OSPF timers.

Answer: B

QUESTION NO: 194

You have just completed an OSPF implementation. While executing your verification plan, you determine that R1 is not able to establish full OSPF adjacency with R2. The show ip ospf neighbor command output on R1 shows that R2 is stuck in the INIT state. What could be the cause of this problem?

Select the best response.

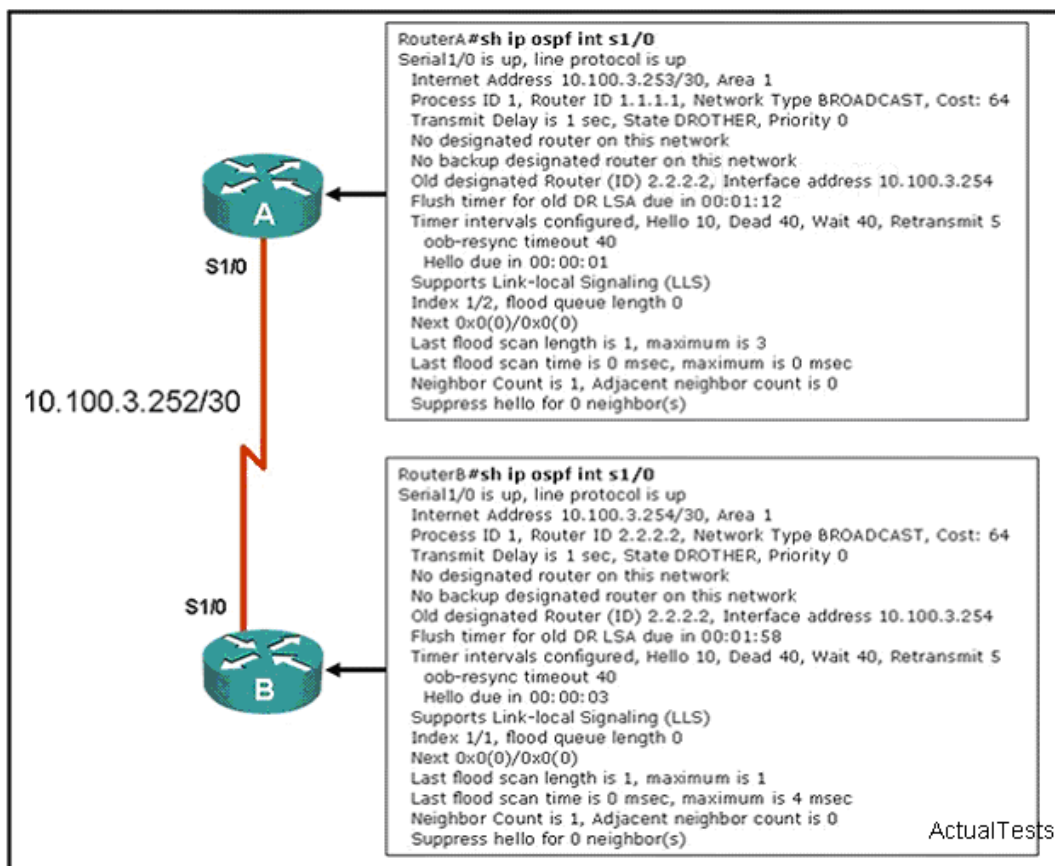
- A. DR and BDR election errors between R1 and R2.
- B. The R2 router has not received the OSPF hello packets from the R1 router.
- C. Mismatched interface maximum transmission unit (MTU) configuration between the R1 and R2.
- D. Mismatched OSPF hello interval configuration between the R1 and R2
- E. Corrupted LSAs exchanges between the R1 and R2.

Answer: B

QUESTION NO: 195

Refer to the exhibit. You have completed an OSPF implementation, and you are verifying OSPF operation. You notice that router A and router B are stuck in the two-way state. From the show ip

ospf interface command output, what is the cause of this issue?



Select the best response.

- A. All OSPF implementations must have at least one interface in area 0.
- B. You are attempting to run in the broadcast mode over an NBMA interface.
- C. Both routers are configured to function as a BDR; therefore, there is no DR router.
- D. Someone has changed the OSPF router ID; therefore you must clear the OSPF process.
- E. The OSPF priority is set to 0 on both routers; therefore neither can become the DR.

Answer: E

QUESTION NO: 196

You have completed an OSPF implementation, and you are verifying OSPF operation. During this verification, you notice that the OSPF route of 172.16.10.0 is repeatedly appearing and disappearing from the routing table. Further investigation finds that the OSPF CPU utilization is very high and the routers are constantly performing SPF calculations. You determine that 172.16.20.2 is the source of the 172.16.10.0 route. Using the show ip ospf database router 172.16.20.1 command, you notice that when this show command is performed repeatedly, the contents of the LSA change every few seconds.

What could be the cause of this problem? Select the best response.

- A. OSPF authentication errors between some of the routers.
- B. Two routers have the same OSPF router ID.
- C. Issues with mistuned OSPF timers.
- D. OSPF LSA pacing issues between some of the routers.
- E. OSPF neighbor adjacency problems between some of the routers.

Answer: B

QUESTION NO: 197

When an OSPF design is planned, which implementation can help a router not have memory resource issues? Select the best response.

- A. Have a backbone area (area 0) with 40 routers and use default routes to reach external destinations.
- B. Have a backbone area (area 0) with 4 routers and 30,000 external routes injected into OSPF.
- C. Have less OSPF areas to reduce the need for interarea route summarizations.
- D. Have multiple OSPF processes on each OSPF router. Example, router ospf 1, router ospf 2.

Answer: A

QUESTION NO: 198

The maximum number of routers per OSPF area typically depends on which three factors? (Choose three.) Select 3 response(s).

- A. the kind of OSPF areas being implemented
- B. the number of external LSAs in the network
- C. the number of DRs and BDRs in the areas
- D. the number of virtual links in the areas
- E. how well the areas can be summarized
- F. the use of LSA filters

Answer: A,B,E

QUESTION NO: 199

You are troubleshooting an OSPF problem where external routes are not showing up in the OSPF database. Which two options are valid checks that should be performed first to verify proper OSPF operation? (Choose two.) Select 2 response(s).

- A. Are the ASBRs trying to redistribute the external routes into a totally stubby area?
- B. Are the ABRs configured with stubby areas?
- C. Is the subnets keyword being used with the redistribution command?
- D. Is backbone area (area 0) contiguous?
- E. Is the CPU utilization of the routers high?

Answer: A,C

QUESTION NO: 200

When verifying the OSPF link state database, which type of LSAs should you expect to see within the different OSPF area types? (Choose three.) Select 3 response(s).

- A. All OSPF routers in stubby areas can have type 3 LSAs in their database.
- B. All OSPF routers in stubby areas can have type 7 LSAs in their database.
- C. All OSPF routers in totally stubby areas can have type 3 LSAs in their database.
- D. All OSPF routers in totally stubby areas can have type 7 LSAs in their database.
- E. All OSPF routers in NSSA areas can have type 3 LSAs in their database.
- F. All OSPF routers in NSSA areas can have type 7 LSAs in their database.

Answer: A,E,F

QUESTION NO: 201

When verifying OSPF virtual link problems, which is an important item to check on the two transit OSPF routers? Select the best response.

- A. OSPF process ID
- B. OSPF router ID
- C. OSPF network type
- D. OSPF memory usage
- E. OSPF CPU utilization
- F. OSPF stub area configurations

Answer: B

QUESTION NO: 202

You are developing a verification plan for an upcoming OSPF implementation. Part of this plan is to verify the status of type 3 LSAs within the network. Which routers should you verify first to ensure that the configurations are correct for generating type 3 LSAs?

Select the best response.

- A. Internal routers within the backbone area (area 0)
- B. Internal routers within the NSSAs
- C. Internal routers within the stubby areas
- D. ASBRs
- E. ABRs
- F. DRs and BDRs

Answer: E

QUESTION NO: 203

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

Select the best response.

- 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.

Answer: A

QUESTION NO: 204

Based on the need to limit processing and bandwidth utilization due to dynamic routing protocol operation, the following routing requirements have been specified for your network.

- partial and incremental routing updates
- only the devices affected by a topology change perform route recomputation
- route recomputation only occurs for routes that were affected

Which dynamic routing protocol should be deployed in your network to best meet these requirements?

Select the best response.

- A. BGP
- B. OSPF
- C. IS-IS
- D. EIGRP
- E. RIPv2

Answer: D

QUESTION NO: 205

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

- 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.

Answer: C

QUESTION NO: 206

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? Select the best response.

- 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.

Answer: D

QUESTION NO: 207

```

R3#show run | include default-ip default-network 140.140.0.0
ip default-network 130.130.0.0
R3#show ip route | begin Gateway
Gateway of last resort is 0.0.0.0 to network 130.130.0.0
116.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
C 116.16.37.0/30 is directly connected, Serial1/0.2
C 116.16.32.0/30 is directly connected, Serial2/0.2
C 116.16.34.0/28 is directly connected, Serial1/0.1
C 116.16.35.0/28 is directly connected, Serial2/0.1
S 116.0.0.0/8 [1/0] via 116.16.34.0
* 140.140.0.0/32 is subnetted, 3 subnets
O 140.140.1.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
O 140.140.3.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
O 140.140.2.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
* 130.130.0.0/16 is variably subnetted, 4 subnets, 2 masks
D* 130.130.0.0/16 is a summary, 00:30:04, Null0
C 130.130.1.0/24 is directly connected, Ethernet0/0
C 130.130.2.0/24 is directly connected, Ethernet0/1
C 130.130.3.0/24 is directly connected, Ethernet1/0
D 150.150.0.0/16 [90/679936] via 116.16.35.5, 00:02:58, Serial2/0.1

```

Refer to the above exhibit. Why is the 140.140.0.0 network not used as the gateway of last resort even though it is configured first?

Select the best response.

- A. The last default-network statement will always be preferred.
- B. A route to the 140.140.0.0 network does not exist in the routing table.
- C. Default-network selection will always prefer the statement with the lowest IP address.
- D. A router will load balance across multiple default-networks; repeatedly issuing the show ip route command would show the gateway of last resort changing between the two networks.

Answer: B

QUESTION NO: 208

```

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
```

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

- 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.

Answer: C

QUESTION NO: 209

Which BGP option is required when load sharing over multiple equal-bandwidth parallel links from a single CE router to a single ISP router over eBGP? Select the best response.

- A. eBGP Multipath
- B. eBGP Multihop
- C. BGP Synchronization
- D. Public AS numbers

Answer: B

QUESTION NO: 210

Which BGP feature should be used to avoid high memory utilization on a router?
Select the best response.

- A. soft-reconfiguration
- B. route refresh
- C. BGP communities
- D. full-mesh BGP peering

Answer: B

QUESTION NO: 211

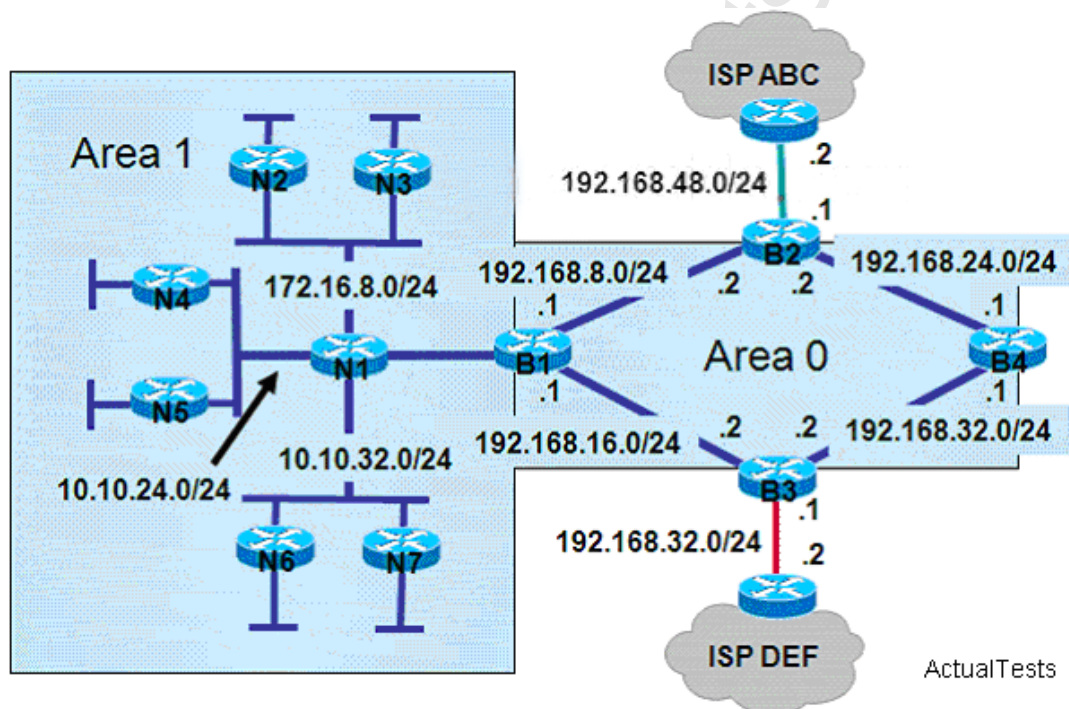
Which functionality is required within an IP router that is situated at the boundary of an IPv4 network and an IPv6 network to allow communication between IPv6-only and IPv4-only nodes? Select the best response.

- A. Autoconfiguration
- B. Automatic 6to4 Tunnel
- C. Automatic 6to4 Relay
- D. Network Address Translator-Protocol Translator (NAT-PT)
- E. Intrasite Automatic Tunnel Address Protocol (ISATAP)

Answer: D

QUESTION NO: 212

Refer to the exhibit. A company would prefer all Internet-bound OSPF routed traffic to use ISP ABC with ISP DEF as a backup. As the network consultant, what three configuration changes might you make? (Choose three.) Select 3 response(s).




- A. The default-information originate command should be configured on router B1 and B4.
- B. The default-information originate command should be configured on router B2 and B3.
- C. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 1.
- D. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 25.
- E. The metric type value should be set to type 1.

F. The metric type value should be set to type 2.

Answer: B,D,F

QUESTION NO: 213

Refer to the exhibit. A network administrator wants to reduce the number of OSPF routes advertised from Area 1 into Area 2. As the router configuration specialist, what two things would you do to accomplish this goal? (Choose two.)



- A. Enter the configuration on router B1.
- B. Enter the configuration on router B4.
- C. On the same router, enter the Summary-address 10.10.1.0 255.255.255.128 subcommand.
- D. On the same router, enter the Area 1 range 10.10.1.0 255.255.255.128 subcommand.
- E. On the same router, enter the Area 2 range 10.10.1.0 255.255.255.128 subcommand.

Answer: A,D

QUESTION NO: 214

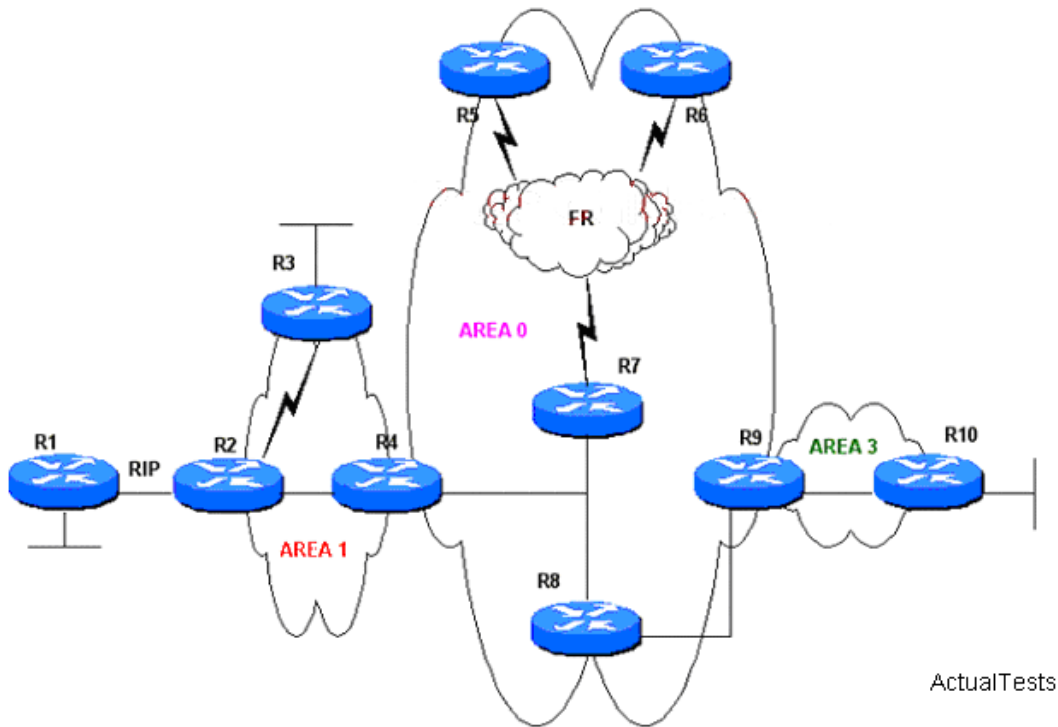
The administrator wants to verify the current state of the OSPF database loading process. Which show command should the administrator use? Select the best response.

- A. show ip ospf [process-id] interface
- B. show ip ospf neighbor
- C. show ip ospf [process-id]
- D. show ip ospf [process-id area-id] database

Answer: B

QUESTION NO: 215

Refer to the exhibit. OSPF is running throughout the network. You want to minimize the propagation of LSAs into and out of Area 1. Which OSPF feature would best achieve this goal?



Select the best response.

- A. stub
- B. totally stubby
- C. NSSA
- D. totally NSSA

Answer: D

QUESTION NO: 216

Refer to the exhibit. Based on this command output, what can we conclude about R3?

```
R3#show ip ospf database
```

```

      OSPF Router with ID (172.16.1.1) (Process ID 123)
      Router Link States (Area 0.0.0.0)
Link ID      ADV Router   Age         Seq#         Checksum Link count
172.16.1.1   172.16.1.1   128        0x80000002  0x000748  1
192.168.0.4  192.168.0.4  128        0x8000000C  0x000EA9  1
      Net Link States (Area 0.0.0.0)
Link ID      ADV Router   Age         Seq#         Checksum
116.16.34.4  192.168.0.4  128        0x80000003  0x00CC36
      Summary Net Link States (Area 0.0.0.0)
Link ID      ADV Router   Age         Seq#         Checksum
116.16.35.0  172.16.1.1   245        0x80000001  0x00E4C0
140.140.1.1  192.168.0.4  728        0x80000001  0x00A013
140.140.2.1  192.168.0.4  728        0x80000001  0x00951D
140.140.3.1  192.168.0.4  728        0x80000001  0x008A27
192.168.0.4  192.168.0.4  728        0x80000001  0x0095CB
      Router Link States (Area 2)
Link ID      ADV Router   Age         Seq#         Checksum Link count
172.16.1.1   172.16.1.1   127        0x80000002  0x00F725  1
192.168.0.5  192.168.0.5  128        0x8000000E  0x00258B  1
      Net Link States (Area 2)
Link ID      ADV Router   Age         Seq#         Checksum
116.16.35.5  192.168.0.5  128        0x80000003  0x00BB43
      Summary Net Link States (Area 2)
Link ID      ADV Router   Age         Seq#         Checksum
116.16.34.0  172.16.1.1   245        0x80000001  0x00EFB6
140.140.1.1  172.16.1.1   124        0x80000001  0x00A77A
140.140.2.1  172.16.1.1   124        0x80000001  0x009C84
140.140.3.1  172.16.1.1   124        0x80000001  0x00918E
192.168.0.4  172.16.1.1   124        0x80000001  0x009C33

```

Select the best response.

- A. R3 is an ABR.
- B. R3 is not connected to the backbone.
- C. R3 has four neighbors.
- D. R3 is the DR for area 2.

Answer: A

QUESTION NO: 217

During the IPv6 autoconfiguration, what does the device append to the 64-bit prefix that it receives from the router to create its IPv6 address?

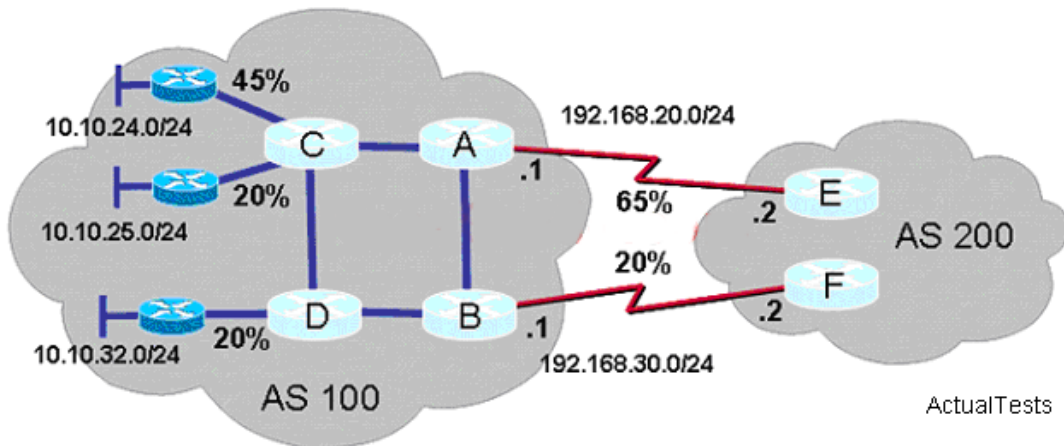
Select the best response.

- A. a pseudorandom generated number
- B. its locally configured IPv4 address
- C. the DHCP-supplied device ID
- D. its MAC address

Answer: D

QUESTION NO: 218

Refer to the exhibit. A client has asked you to consult on an eBGP loading question. Currently the AS 100 eBGP links have an average outbound load of 65% and 20% respectively. On further investigation, traffic from 10.10.24.0 accounts for 45%, and 10.10.25.0 and 10.10.32.0 accounts for 20% each of the outbound load. The customer wants to spread the load between the two eBGP links more evenly. The BGP attributes are currently set at their default values. If you are located at AS 100 and want to influence how AS 100 sends traffic to AS 200, what BGP attribute could you configure to cause AS 100 outbound traffic to load the eBGP links more evenly?



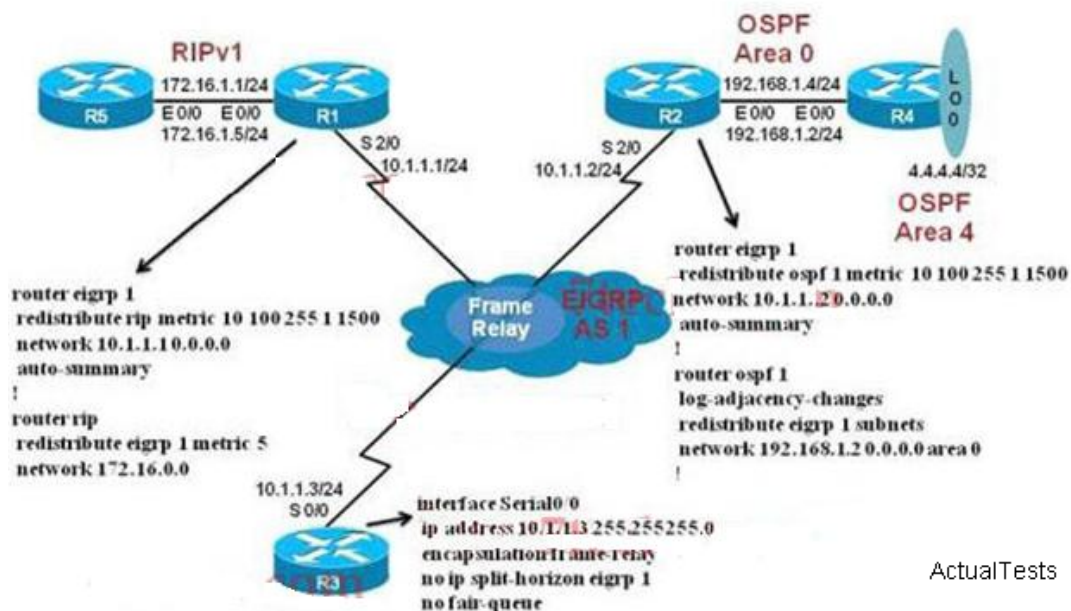
Select the best response.

- A. On router A, set the default local-preference to 50.
- B. On router B, set the default metric to 150.
- C. On router B, configure a route map for 10.10.25.0/24 with a local preference of 150 linked to neighbor 192.168.30.2.
- D. On router B, set the default local-preference to 150.

Answer: C

QUESTION NO: 219

Refer to the exhibit. Looking at the topology diagram and the partial router configurations shown, which statement is true?



ActualTests

Select the best response.

- A. A routing loop will occur due to mutual route redistribution occurring on R1 and R2.
- B. Suboptimal routing will occur due to mutual route redistribution occurring on R1 and R2.
- C. Additional route filtering configurations using route maps and ACLs are required on the R1 and R2 routers to prevent routing loops.
- D. R2 will not be able to redistribute the EIGRP subnets into OSPF, because R2 is missing the default seed metric for OSPF.
- E. The 10.1.1.0/24 subnet will appear as 10.0.0.0/8 in the R5 routing table.

Answer: E

QUESTION NO: 220

Which two methods use IPsec to provide secure connectivity from the branch office to the headquarters office? (Choose two.) Select 2 response(s).

- A. DMVPN
- B. MPLS VPN
- C. Virtual Tunnel Interface (VTI)
- D. SSL VPN
- E. PPPoE

Answer: A,C

QUESTION NO: 221

What is a key benefit of using a GRE tunnel to provide connectivity between branch offices and headquarters? Select the best response.

- A. authentication, integrity checking, and confidentiality
- B. less overhead
- C. dynamic routing over the tunnel
- D. granular QoS support
- E. open standard
- F. scalability

Answer: C

QUESTION NO: 222

You have implemented mutual route redistribution between OSPF and EIGRP on a border router. When checking the routing table on one of the EIGRP routers within the EIGRP routing domain, you are seeing some, but not all of the expected routes. What should you verify to troubleshoot this problem? Select the best response.

- A. The border router is using a proper seed metric for OSPF.
- B. The border router is using a proper seed metric for EIGRP.
- C. The administrative distance is set for OSPF and EIGRP.
- D. The missing OSPF routes are present in the routing table of the border router.
- E. The subnet keyword on the border router in the redistribute OSPF command.

Answer: D

QUESTION NO: 223

You have implemented mutual route redistribution between OSPF and EIGRP on a border router. When checking the routing table on one of the OSPF routers within the OSPF routing domain, you are seeing some, but not all of the expected routes. Which two things should you verify to troubleshoot this problem? (Choose two.) Select 2 response(s).

- A. The border router is using a proper seed metric for OSPF.
- B. The border router is using a proper seed metric for EIGRP.
- C. The administrative distance is set for OSPF and EIGRP.
- D. The missing EIGRP routes are present in the routing table of the border router.
- E. The subnet keyword on the border router in the redistribute EIGRP command.

Answer: D,E

QUESTION NO: 224

Which DSL encapsulation method requires client software running on the end-user PC that is directly connected to a DSL modem?

Select the best response.

- A. PPPoA
- B. PPPoE
- C. PPP
- D. L2TP
- E. ATM

Answer: B

QUESTION NO: 225

What is the purpose of configuring the router as a PPPoE client? Select the best response.

- A. to provide VPN access over L2TP
- B. to enable PPP session from the router to the termination device at the headend for metro Ethernet connectivity
- C. for DSL connectivity and removing the need for the end-user PC to run the PPPoE client software
- D. for connecting the router to a cable modem, which bridges the Ethernet frames from the router to the cable modem termination system

Answer: C

QUESTION NO: 226

What is the international standard for transmitting data over a cable system? Select the best response.

- A. PPPoE
- B. DOCSIS
- C. CMTS
- D. AAL5

Answer: B

QUESTION NO: 227

Under which circumstance will a branch ISR router contain interface vlan configurations?

Select the best response.

- A. performing inter-VLAN routing
- B. performing 802.1Q trunking
- C. performing ISL trunking
- D. Ethernet Switch Module installed
- E. ADSL WIC installed
- F. running Call Manager Express

Answer: D

QUESTION NO: 228

Refer to the exhibit. You want to use all the routes in the EIGRP topology for IP load balancing. Which two EIGRP subcommands would you use to accomplish this goal? (Choose two.) Select 2 response(s).

```
R1#show ip eigrp topology | section 0.0.0.0
P 0.0.0.0/0, 2 successors, FD is 2174976
  via 212.50.185.125 (2174976/2169856), Ethernet0/0
  via 212.50.185.126 (2174976/2169856), Ethernet0/0
  via 212.50.185.65 (2178816/2172416), Ethernet1/0
  via 212.50.185.66 (2178816/2172416), Ethernet1/0
  via 212.50.185.33 (2180096/2172416), Ethernet2/0
  via 212.50.185.34 (2180096/2172416), Ethernet2/0
R1#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "eigrp 212", distance 170, metric 2174976, candidate default path, type external
  Redistributing via eigrp 212
  Last update from 212.50.185.126 on Ethernet0/0, 00:00:32 ago
  Routing Descriptor Blocks:
    * 212.50.185.126, from 212.50.185.126, 00:00:32 ago, via Ethernet0/0
      Route metric is 2174976, traffic share count is 1
      Total delay is 20200 microseconds, minimum bandwidth is 1544 kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 3/255, Hops 1
    212.50.185.125, from 212.50.185.125, 00:00:32 ago, via Ethernet0/0
      Route metric is 2174976, traffic share count is 1
      Total delay is 20200 microseconds, minimum bandwidth is 1544 kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 3/255, Hops 1
```

ActualTests

- A. traffic-share balanced
- B. distance
- C. maximum-paths
- D. default-network
- E. variance

Answer: C,E

QUESTION NO: 229

Refer to the exhibit. R1 accesses the Internet using E0/0. You have been asked to configure R1 so that a default route is generated to its downstream devices (191.0.0.1 and 192.0.0.1). Which commands would create this configuration?

```
R1#show ip route
 1.0.0.0/24 is subnetted, 1 subnets
C    1.1.1.0 is directly connected, Loopback9
D EX 212.50.167.0/24 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 191.0.0.0/30 is subnetted, 1 subnets
C    191.0.0.0 is directly connected, Serial2/0
D EX 212.50.166.0/24 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 20.0.0.0/24 is subnetted, 1 subnets
C    20.20.20.0 is directly connected, Ethernet0/0
 212.50.185.0/27 is subnetted, 3 subnets
D EX 212.50.185.64 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
D EX 212.50.185.96 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
D EX 212.50.185.32 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 192.0.0.0/30 is subnetted, 1 subnets
C    192.0.0.0 is directly connected, Serial1/0
```

ActualTests

Select the best response.

- A. router eigrp 190
redistribute static
!
ip route 0.0.0.0 0.0.0.0 Null0
- B. ip default-network 20.0.0.0
- C. router eigrp 190
redistribute static
!
ip route 0.0.0.0 255.255.255.255 Null0
- D. ip default-network 20.20.20.0

Answer: A

QUESTION NO: 230

Refer to the exhibit. How would you confirm on R1 that load balancing is actually occurring on the default-network (0.0.0.0)? Select the best response.

```

R1#show ip route
Codes: 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

Gateway of last resort is 212.50.185.126 to network 0.0.0.0

D    212.50.167.0/24 [90/160000] via 212.50.185.82, 00:05:55, Ethernet1/0
    212.50.166.0/24 is variably subnetted, 4 subnets, 2 masks
D    212.50.166.0/24 is a summary, 00:05:55, Null0
C    212.50.166.1/32 is directly connected, Loopback1
C    212.50.166.2/32 is directly connected, Loopback2
C    212.50.166.20/32 is directly connected, Loopback20
    212.50.185.0/27 is subnetted, 3 subnets
C    212.50.185.64 is directly connected, Ethernet1/0
C    212.50.185.96 is directly connected, Ethernet0/0
C    212.50.185.32 is directly connected, Ethernet2/0
D*EX 0.0.0.0/0 [170/2174976] via 212.50.185.126, 00:05:55, Ethernet0/0
    [170/2174976] via 212.50.185.125, 00:05:55, Ethernet0/0

```

- A. Use ping and the show ip route command to confirm the timers for each default network resets to 0.
- B. Load balancing does not occur over default networks; the second route will only be used for failover.
- C. Use an extended ping along with repeated show ip route commands to confirm the gateway of last resort address toggles back and forth.
- D. Use the traceroute command to an address that is not explicitly in the routing table.

Answer: D

QUESTION NO: 231

Refer to the exhibit. Based upon the configuration, you need to understand why the policy routing match counts are not increasing. Which would be the first logical step to take? Select the best response.

```

R1#show route-map divert
route-map divert, permit, sequence 1
  Match clauses:
    ip address (access-lists): 101
  Set clauses:
    ip next-hop 212.50.185.126
    ip next-hop recursive 192.0.0.1
    ip default next-hop 212.50.185.125
  Policy routing matches: 0 packets, 0 bytes

```

- A. Confirm if there are other problematic route-map statements that precede divert.
- B. Check the access list for log hits.
- C. Check the routing table for 212.50.185.126.
- D. Remove any two of the set clauses. (Multiple set clause entries will cause PBR to use the routing table.)

Answer: B

QUESTION NO: 232

Refer to the exhibit. A partial routing configuration is shown. Complete the configuration so that only the default-network is redistributed from EIGRP 190 into EIGRP 212. Which ACL statement completes the configuration correctly?

```
router eigrp 190
 redistribute eigrp 212
 network 192.0.0.0 0.0.0.3
!
router eigrp 212
 redistribute eigrp 190 route-map default_route
 network 212.50.185.96 0.0.0.31
!
route-map default_route permit 10
 match ip address 100
```

Select the best response.

- A. access-list 100 permit ip 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
- B. access-list 100 permit ip host 0.0.0.0 any
- C. access-list 100 permit ip any host 0.0.0.0
- D. A default-network cannot be redistributed between routing processes.

Answer: C

QUESTION NO: 233

Refer to the exhibit. The partial configuration for an OSPF ASBR and an Area 0 ABR is shown. Assume the OSPF configurations throughout the network are operable. Which statement about these configurations is true?

```

ASBR
router ospf 123
 redistribute eigrp 1 route-map eigrp-to-ospf
 !
 route-map eigrp-to-ospf permit 10
  match ip address prefix-list private
  set tag 255
 route-map eigrp-to-ospf permit 20
 !
 ip prefix-list private permit 10.0.0.0/8 ge 8 le 30
 ip prefix-list private permit 172.16.0.0/11 ge 11 le 30
 ip prefix-list private permit 192.168.0.0/16 ge 16 le 30

```

```

Area 0 ABR
router ospf 123
 distribute-list route-map private-filter in
 !
 route-map private-filter deny 10
  match tag 255
 route-map private-filter permit 20

```

ActualTests


Select the best response.

- A. The ASBR route-maps are basically useless, because there are no deny prefix-lists.
- B. LSA Type 5s will not be received by the ABR from the ASBR.
- C. The OSPF backbone will not learn any RFC 1918 addresses.
- D. The matched prefix-list addresses will be given a metric of 255, which is essentially unreachable.

Answer: C

QUESTION NO: 234

Refer to the exhibit. Which three commands should be used on router B1 to redistribute the EIGRP AS 10 routes into RIP? (Choose three.)



Select 3 response(s).

- A. router rip
- B. router eigrp 10
- C. redistribute eigrp 10
- D. redistribute rip
- E. default-metric 10000 100 255 1 1500
- F. default-metric 5

Answer: A,C,F

QUESTION NO: 235

Refer to the exhibit. Examine the partial configuration and the routing table excerpt. Which routes would be redistributed into OSPF area 1?

```

Router eigrp 10
  Network 10.0.0.0
Router ospf 100
  redistribute eigrp 10 subnets
  network 172.16.0.0 0.0.255.255 area 1
.....
A partial routing table
  C   10.10.10.16/28 is directly connected, Ethernet1/0
  D   10.10.10.64/26 [90/284160] via 10.10.24.2, 01:54:39, Ethernet1/0
  C   172.16.10.0/24 is directly connected, FastEthernet0/1

```

Select the best response.

- A. 10.10.10.16/28 only
- B. 10.10.10.16/28 and 10.10.10.64/26
- C. 10.10.10.16/28, 10.10.10.64/26, and 172.16.10.0/24
- D. 10.10.10.64/26 only

Answer: D

QUESTION NO: 236

Refer to the exhibit. Examine the partial configuration and the routing table excerpt. Which routes would be redistributed into OSPF area 1?

```

Router eigrp 10
  Network 10.0.0.0
Router ospf 100
  redistribute eigrp 10 subnets
  network 172.16.0.0 0.0.255.255 area 1
.....
A partial routing table
  C   10.10.10.16/28 is directly connected, Ethernet1/0
  D   10.10.10.64/26 [90/284160] via 10.10.24.2, 01:54:39, Ethernet1/0
  C   172.16.10.0/24 is directly connected, FastEthernet0/1

```

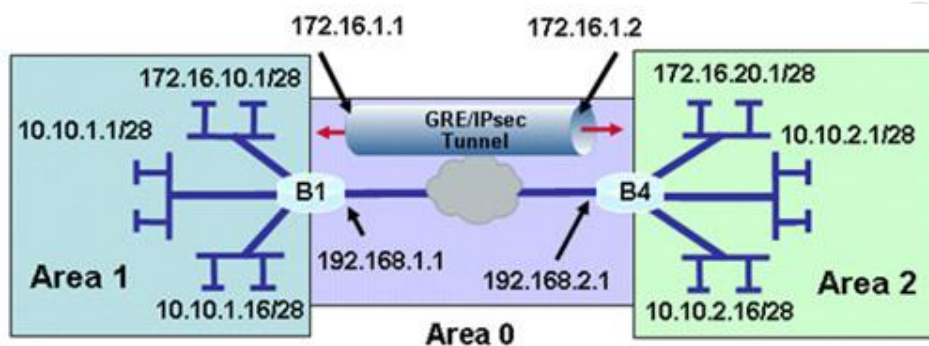
Select the best response.

- A. 10.10.10.16/28 only
- B. 10.10.10.16/28 and 10.10.10.64/26
- C. 10.10.10.16/28, 10.10.10.64/26, and 172.16.10.0/24
- D. 10.10.10.64/26 only

Answer: D

QUESTION NO: 237

Refer to the exhibit. A new TAC engineer came to you for advice. A GRE over IPsec tunnel was configured, but the tunnel is not coming up. What did the TAC engineer configure incorrectly?



Router B1 Configuration

```
crypto ipsec transform-set 10 esp-sha-hmac esp-3des
crypto map tunnel 1 ipsec-isakmp
 set transform-set 10
 set peer 192.168.2.1
 match address 102
crypto isakmp policy 1
 authentication pre-share
 group 2
crypto isakmp key ***** address 172.16.1.2
access-list 102 permit gre host 192.168.1.1 host
192.168.2.1
access-list 102 permit esp host 192.168.1.1 host
192.168.2.1
access-list 102 permit udp host 192.168.1.1 eq isakmp
host 192.168.2.1
```

Router B1 Configuration (con't)

```
Interface f0/0
 Ip address 192.168.1.1 255.255.255.0
interface Tunnel0
 ip address 172.16.1.1 255.255.255.0
 crypto map tunnel
 tunnel source F0/0
 tunnel destination 192.168.2.1
 tunnel path-mtu-discovery
 ip ospf mtu-ignore
router ospf 200
 network 10.10.1.1 0.0.0.224 area 1
 network 172.16.10.1 0.0.0.240 area 1
 network 192.168.1.0 0.0.0.255 area 0
```

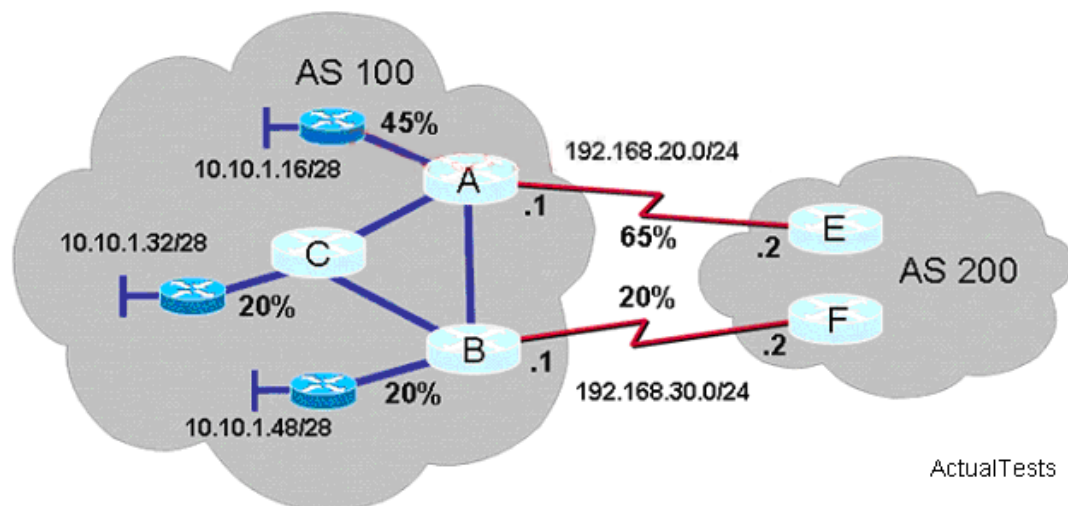
- A. The crypto isakmp configuration is not correct.
- B. The crypto map configuration is not correct.
- C. The interface tunnel configuration is not correct.
- D. The network configuration is not correct; network 172.16.1.0 is missing.

Answer: A

QUESTION NO: 238

Refer to the exhibit. Currently the two eBGP links between AS100 and AS200 have an average inbound load of 65% and 20% respectively. After further investigation, traffic to 10.10.1.16/28 accounts for 45%, and traffic to 10.10.1.32/28 and to 10.10.1.48/28 each account for 20% of the inbound load. The BGP attributes are currently set at their default values in both autonomous systems.

If you want to influence how AS200 sends traffic to AS100, which eBGP configurations would you configure in AS100 to influence AS200 to use the eBGP links more evenly? (Choose two.)



ActualTests

Select 2 response(s).

A. neighbor 192.168.30.2 route-map as_50 out

B. neighbor 192.168.20.2 route-map as_50 out

C. route-map as_50 permit 10

match ip address 50

set metric 150

access-list 50 permit 10.10.1.16 0.0.0.240

D. route-map as_50 permit 10

match ip address 50

set metric 150

access-list 50 permit 10.10.1.32 0.0.0.240

Answer: B,C

QUESTION NO: 239

Which two statements about route redistribution when implementing OSPF are true? (Choose two.)

Select 2 response(s).

A. Routes learned using any IP routing protocol can only be redistributed into non IP routing protocols.

- B. OSPF can import routes learned using EIGRP, RIP, and IS-IS.
- C. OSPF routes cannot be exported into EIGRP, RIP, and IS-IS.
- D. At the interdomain level, OSPF cannot import routes learned using BGP.
- E. OSPF routes can be exported into BGP.

Answer: B,E

QUESTION NO: 240

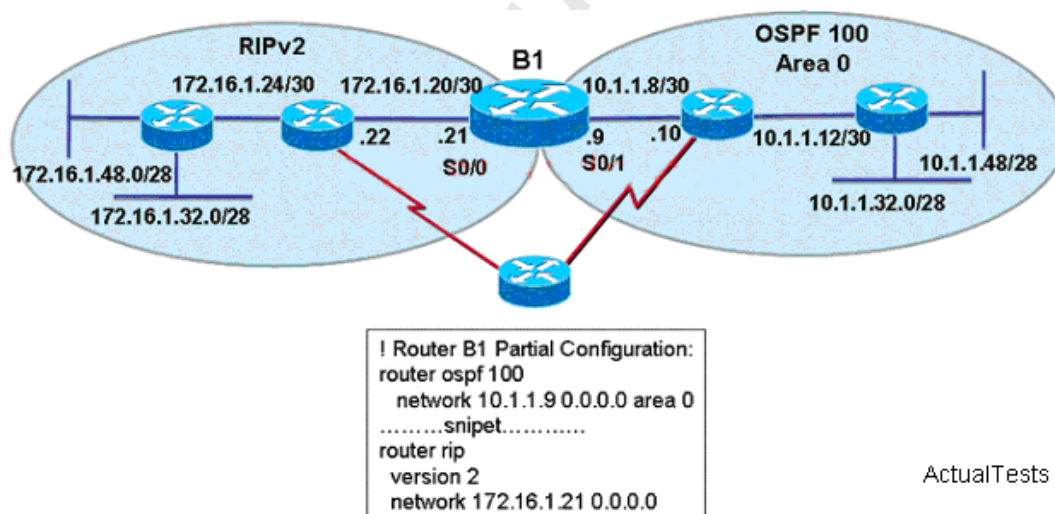
Which two routing interface parameters are supported in OSPF implementations? (Choose two.)
Select 2 response(s).

- A. retransmit-interval
- B. dead-interval
- C. stub area
- D. virtual link
- E. NSSA area

Answer: A,B

QUESTION NO: 241

Refer to the exhibit. A new TAC engineer comes to you for advice. The engineer wants to configure RIPv2-OSPF two-way redistribution while avoiding routing loops. Which two additions to the router B1 configuration should the engineer make? (Choose two.)



ActualTests

Select 2 response(s).

- A. access-list 40 deny 172.16.1.0 0.0.0.255
- access-list 40 permit any
- router rip

```
redistribute ospf 100 metric 5
```

```
distribute-list 40 out ospf 100
```

```
B. ip prefix-list rip_routes permit 172.16.1.16/25 ge 26 le 28
```

```
route-map redis-ospf deny 10
```

```
match ip address prefix-list rip_routes
```

```
router rip
```

```
redistribute ospf 10 route-map redis-ospf subnets
```

```
C. ip prefix-list rip-to-ospf permit 10.1.1.8/25 ge 26 le 28
```

```
route-map redis-rip deny 20
```

```
match ip address prefix-list rip-to-ospf
```

```
router ospf 100
```

```
redistribute rip route-map redis-rip subnets
```

```
D. access-list 15 deny 10.1.1.0 0.0.0.63 access-list 15 permit any
```

```
route-map redis-rip deny 10
```

```
match ip address 15
```

```
route-map redis-rip permit 20
```


```
router ospf 100
```

```
redistribute rip route-map redis-rip subnets
```

Answer: A,D

QUESTION NO: 242

You want the redistributed EIGRP AS 10 routes to have an administrative distance of 121 when they appear as RIP routes in the routing table of A1.



Which command should you use on a router to accomplish this goal? Select the best response.

A. redistribute eigrp 10 metric 121

B. redistribute rip metric 121


C. default-metric 121

D. distance 121 10.1.1.6 0.0.0.0

Answer: D

QUESTION NO: 243

Refer to the exhibit. In a redundant hub-and-spoke deployment using EIGRP, what feature can be used to ensure that routers C through F are not used as transit routers for data traveling from router B to network 10.1.1.0? Select the best response



- A. Use address summarization at routers C, D, E, and F.
- B. Use the EIGRP Stub feature on routers C, D, E, and F.
- C. Use passive-interface on the spoke links in routers A and B.
- D. Change the administrative distance in routers A and B for routes learned from routers C, D, E, and F.

Answer: B

QUESTION NO: 244

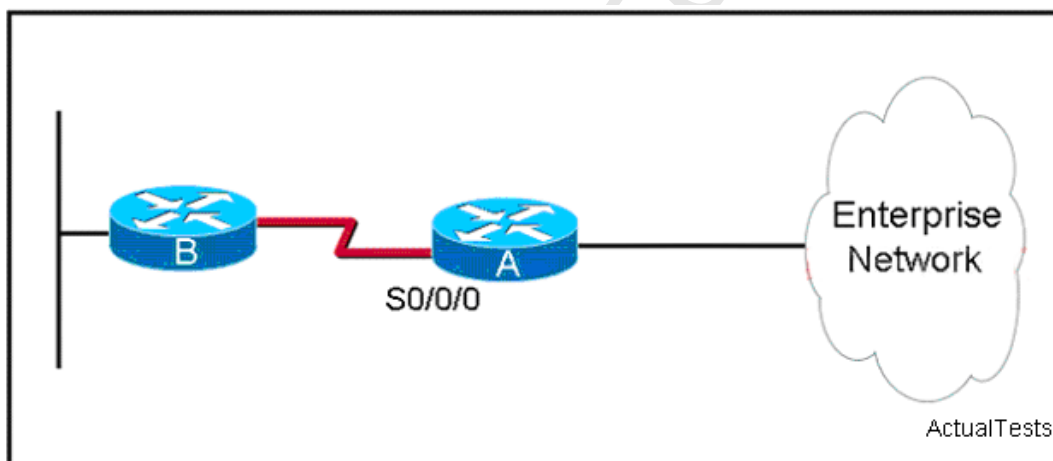
ACME Rocket Sleds is growing, and so is their network. They have determined that they can no longer continue using static routes and must implement a dynamic routing protocol. They want to have data use multiple paths to the destinations, even if the paths are not equal cost. Which routing protocol has the ability to do this? Select the best response.

- A. EIGRP
- B. OSPF
- C. RIPv1
- D. RIPv2
- E. BGP
- F. IS-IS

Answer: A

QUESTION NO: 245

Refer to the exhibit. ROUTE Enterprises has many stub networks in their enterprise network, such as router B and its associated network. EIGRP is to be implemented on router A so that neither the prefix for the S/0/0/0 interface nor the prefixes from router B appear in the routing tables for the router in the enterprise network. Which action will accomplish this goal?




Select the best response.

- A. Declare router B a stub router using the `eigrp stub` command.
- B. Use the `passive-interface` command for interface `Serial0/0/0`.
- C. Use a mask with the `network` command to exclude interface `Serial0/0/0`.
- D. Implement a distribute list to exclude the link prefix from the routing updates.

Answer: C

QUESTION NO: 246

Refer to the exhibit. Which statement about dynamic routing protocols for this network is true?



Select the best response.

- A. No dynamic interior routing protocol can summarize as shown.
- B. Unless configured otherwise, EIGRP would automatically summarize the prefixes as shown in the exhibit.
- C. With this IP addressing scheme, EIGRP can be manually configured to summarize prefixes at the specified summarization points.
- D. The IP address design lends itself to OSPF. Each summarizing router would be an ABR, summarizing to the next area in the address hierarchy.

Answer: C

QUESTION NO: 247

After implementing EIGRP on your network, you issue the show ip eigrp traffic command on router C. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 481/444
Updates sent/received: 41/32
```

Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0

Approximately 25 minutes later, you issue the same command again. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 1057/1020
Updates sent/received: 41/32
Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

Approximately 25 minutes later, you issue the same command a third time. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 1754/1717
Updates sent/received: 41/32
Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

What can you conclude about this network?

Select the best response.

- A. The network has been stable for at least the last 45 minutes.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router A does not know an alternate path to the network.

- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.

Answer: A

QUESTION NO: 248

After implementing EIGRP on your network, you issue the show ip eigrp traffic command on router C. The following output is shown: RouterC#show ip eigrp traffic IP-EIGRP Traffic Statistics for process 1

Hello sent/received: 2112/2076

Updates sent/received: 47/38

Queries sent/received: 5/3

Replies sent/received: 3/4

Acks sent/received: 29/33

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Moments later, you issue the same command a second time and the following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hello sent/received: 2139/2104

Updates sent/received: 50/39

Queries sent/received: 5/4

Replies sent/received: 4/4

Acks sent/received: 31/37

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Moments later, you issue the same command a third time and the following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hello sent/received: 2162/2126

Updates sent/received: 53/42

Queries sent/received: 5/5

Replies sent/received: 5/4

Acks sent/received: 35/41

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0


What information can you determine about this network? Select the best response.

- A. The network is stable.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router C does not know an alternate path to the network.
- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.

Answer: B

QUESTION NO: 249

Refer to the exhibits. Router B should advertise the network connected to the E0/0/0 interface to router A and block all other network advertisements. The IP routing table on router A indicates that it is not receiving this prefix from router B. What is the probable cause of the problem? Select the best response.



```
RouterB#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H  Address                Interface  Hold Uptime  SRTT  RTO  Q  Seq Type
 0  10.255.1.1              Se0/0     138 00:15:26  16   582  0  3
RouterB#

RouterB#debug ip eigrp
IP-EIGRP Route Events debugging is on
02:17:54: IP-EIGRP: 192.168.3.16/28 - denied by distribute list
02:17:54: IP-EIGRP: 192.168.3.32/28 - denied by distribute list
02:17:54: IP-EIGRP: 10.255.1.0/30 - denied by distribute list
02:17:54: IP-EIGRP: 192.168.3.0/24 - denied by distribute list
02:17:54: IP-EIGRP: 10.0.0.0/8 - denied by distribute list
```

ActualTests

- A. An access list on router B is causing the 192.168.3.16/28 network to be denied.
- B. An access list on router B is causing the 192.168.3.32/28 network to be denied.
- C. The distribute list on router B is referencing a numbered access list that does not exist on router B.
- D. The distribute list on router B is referencing the wrong interface.

Answer: A

QUESTION NO: 250

RouterA#

~~~~~

!

router ospf 1

log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 1

network 172.16.1.0 0.0.0.255 area 1

!

~~~~~

RouterB#

~~~~~

router ospf 1

log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 2

network 172.16.2.0 0.0.0.255 area 2

!

~~~~~

RouterC#

~~~~~

!

router ospf 1


log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 0

!

~~~~~

Refer to the exhibits. You are verifying your OSPF implementation, and it does not seem to be functioning properly. What can you conclude from the exhibit and the show running-configuration command output?



Select the best response.

- A. The OSPF areas are not configured correctly.
- B. The wildcard masks for the 10.x.x.x networks are incorrect.
- C. The 172.16.x.x networks need to be connected to area 0 using virtual links.
- D. The 172.16.x.x networks are discontinuous. OSPF is automatically summarizing them to 172.16.0.0/16 and data is being "black holed."
- E. There is not enough information to make a determination.

Answer: A

QUESTION NO: 251

OSPF is enabled on router A. You execute the following command on router A and receive the accompanying output:

```
RouterA#ping 224.0.0.5 repeat 1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 224.0.0.5, timeout is 2 seconds:
Reply to request 0 from 10.100.100.1, 4 ms
10.100.100.1 is the IP address of a loopback interface on router A.
```

What can you conclude about router A?

Select the best response.

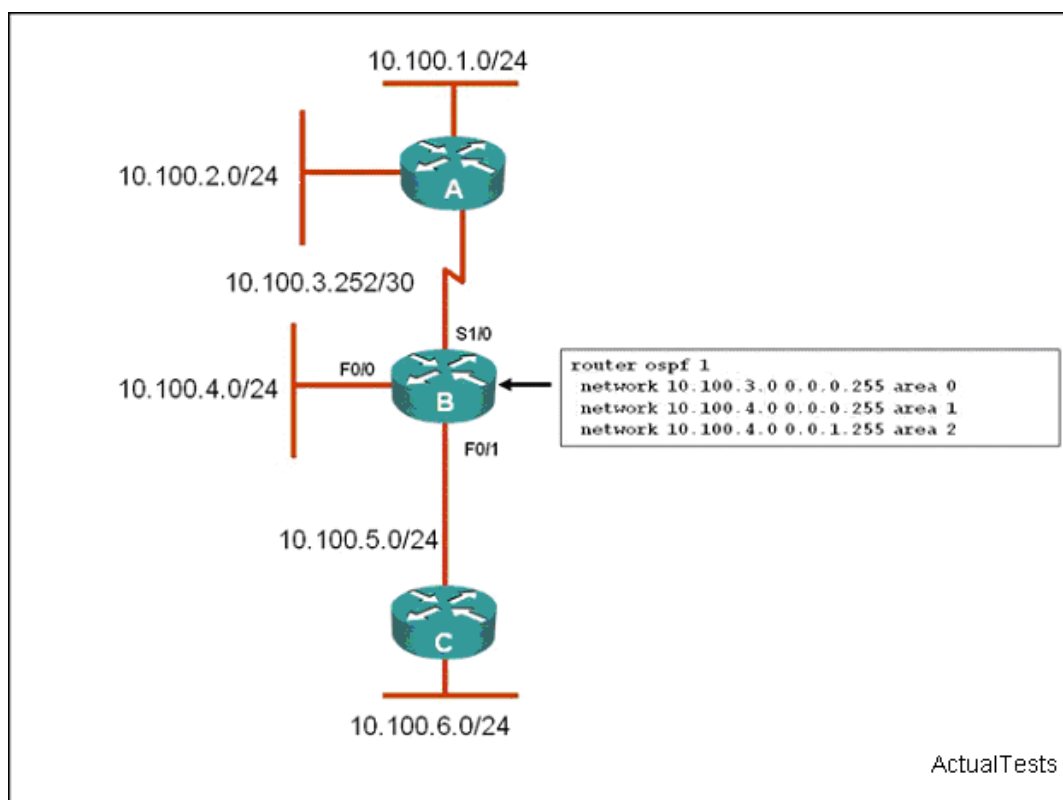
- A. Only the router A loopback interface is participating in the OSPF routing process.
- B. None of the router A interfaces are participating in the OSPF routing process.
- C. Router A is using the loopback interface IP address as its OSPF router ID.

D. Router A does not have any reachable OSPF neighbors.

Answer: D

QUESTION NO: 252

Refer to the exhibit. What is the effect of the OSPF configuration on router B? Select the best response.




- A. All interfaces will be in area 0.
- B. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 1.
- C. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 2.
- D. The router will be an ABR with s1/0 in area 0, f0/0 in area 1, and f0/1 in area 2.

Answer: D

QUESTION NO: 253

Refer to the exhibit. Two routers are connected by Frame Relay and are running OSPF between them. Each router has been configured with the appropriate network statements under router ospf 1, but the routers are not forming an adjacency. Which of the following three commands could be configured on each router to correct this problem? (Choose three.)



Select 3 response(s).

- A. RouterC(config-if)#ip ospf network broadcast
RouterD(config-if)#ip ospf network broadcast
- B. RouterC(config-if)#ip ospf network point-to-point
RouterD(config-if)#ip ospf network point-to-point
- C. RouterC(config-router)#neighbor 10.100.100.4
RouterD(config-router)#neighbor 10.100.100.3
- D. RouterC(config-router)#neighbor 10.255.255.254
RouterD(config-router)#neighbor 10.255.255.253

Answer: A,B,D

QUESTION NO: 254

Your network has a mixture of Fast Ethernet and Gigabit Ethernet links. What needs to be done to ensure optimal data routing when using OSPF? Select the best response.

- A. Nothing. OSPF will determine the most optimal path for routing data by default.
- B. Adjust the hello and dead timers for more rapid detection of link failures.
- C. Increase the reference-bandwidth used to calculate the interface default metrics, on all routers in your network.
- D. Set the priority values on every broadcast interface to ensure that the designated and backup designated routers are the routers with the most processor and memory resources.

Answer: C

QUESTION NO: 255

To make OSPF area 4 totally stubby, the following command was issued on the ABR in router configuration mode: area 4 stub

Which two things need to be done to finish making area 4 a totally stubby area? (Choose two.)
Select 2 response(s).

- A. Apply the area 4 stub command to all routers in the area.
- B. On the ABR, use the area 4 stub command with the no-summary keyword.
- C. On the ABR, specify a default cost for the area with the area default-cost command.
- D. On the ABR, use the default-information originate command to inject a default route into area 4.
- E. Use the auto-cost command on each router in the area to automatically determine the cost to other OSPF areas.

Answer: A,B

QUESTION NO: 256

Refer to the exhibit. Router C was configured so that it could form an adjacency with three OSPF neighbors, one connected to each of its three physical interfaces. Which statement is correct about router C?

```

RouterC#show ip interface brief
Interface                IP-Address      OK? Method Status  Protocol
Ethernet0/0              10.1.0.1        YES NVRAM  up      up
Serial0/0                10.255.255.253 YES NVRAM  up      up
Ethernet0/1              10.2.0.1        YES NVRAM  up      up
Loopback0                10.100.100.3   YES NVRAM  up      up

RouterC#show ip ospf
Routing Process "ospf 1" with ID 10.100.100.3
Supports only single TOS(TOS0) routes
Supports opaque LSA
It is an area border router
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 1 normal 1 stub 0 nssa
External flood list length 0
  Area BACKBONE(0)
    Number of interfaces in this area is 3
    Area has no authentication
    SPF algorithm executed 10 times
    Area ranges are
    Number of LSA 14. Checksum Sum 0x053D21
    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 4
    Number of interfaces in this area is 1
    It is a stub area, no summary LSA in this area
    generates stub default route with cost 1
    Area has no authentication
    SPF algorithm executed 9 times
    Area ranges are
    Number of LSA 4. Checksum Sum 0x01F200
    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

```

ActualTests

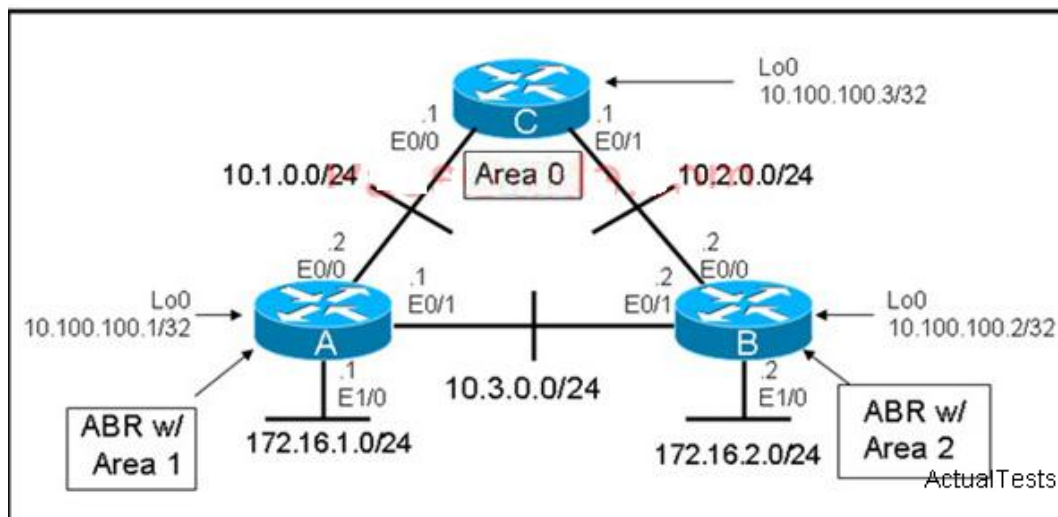
Select the best response.

- A. It is configured and functioning correctly as an OSPF internal router.
- B. It is configured and functioning correctly as an ABR attached to stub area 4.
- C. It is configured and functioning correctly as an ASBR attached to external area 4.
- D. It is configured and functioning correctly as an ABR attached to totally stubby area 4.
- E. It is not configured correctly to function as specified.

Answer: D

QUESTION NO: 257

RouterA#debug ip ospf events OSPF events debugging is on RouterA# 04:43:16: OSPF: Rcv pkt from 10.3.0.2, Ethernet0/1, area 0.0.0.1 mismatch area 0.0.0.2 in the header 04:43:19: OSPF: Rcv hello from 10.100.100.3 area 0 from Ethernet0/0 10.1.0.1 04:43:19: OSPF: End of hello processing Refer to the exhibits. What can be done to fix the problem? Select the best response.




- A. Change router B E0/1 interface to area 0.
- B. Change router A interface E0/1 to area 0.0.0.2.
- C. Configure the E0/1 interfaces of router A and router B to be in area 0.
- D. Shut down the E0/1 interfaces in router A and router B as OSPF does not allow "back doors" between areas.
- E. Remove the E0/1 interfaces in router A and router B from the OSPF process. Use static routes to route data directly from router A to router B to avoid passing data through router C in area 0.

Answer: C

QUESTION NO: 258

Refer to the exhibit. AS 65500 is not advertising the prefix 192.168.12.0/22 to its provider. AS 65500 is running OSPF as its IGP. Which of the following additions to the configuration is most likely to solve the problem? Select the best response.

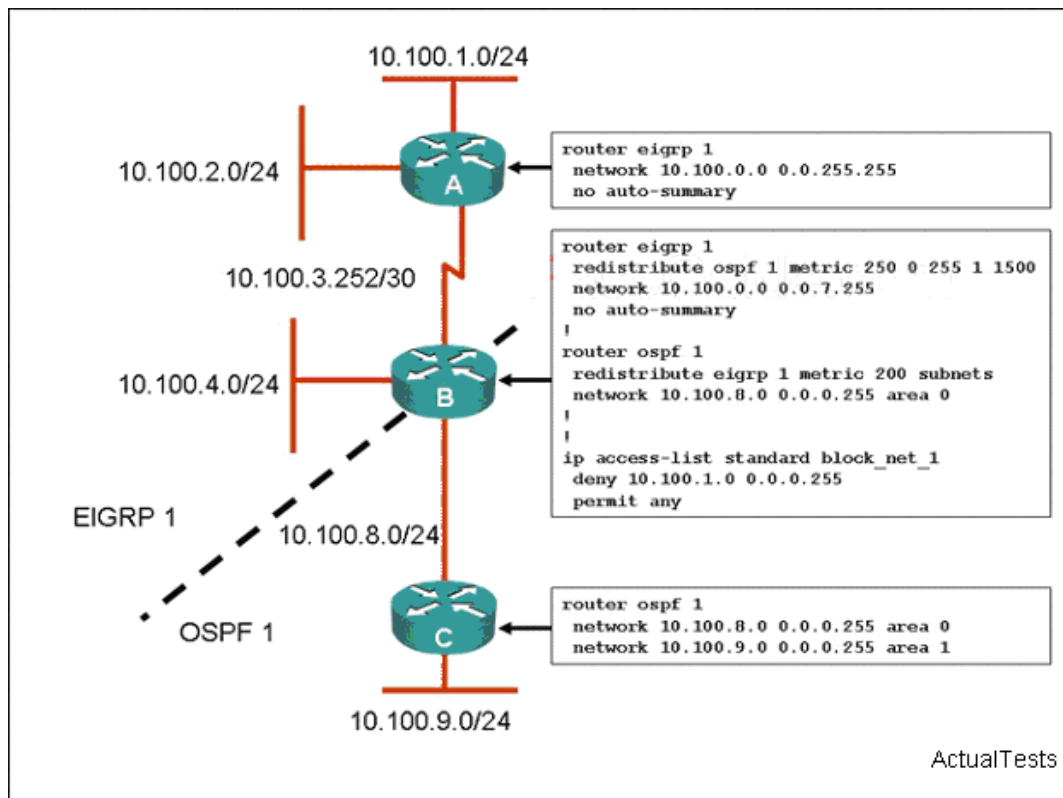


- A. RouterA(config)#ip cef
- B. RouterA(config)#ip route 192.168.12.0 255.255.255.252 null 0
- C. RouterA(config-router)#ebgp multihop 1
- D. RouterA(config-router)#redistribute ospf 1
- E. RouterA(config-router)#neighbor 192.168.14.253 next-hop-self
- F. RouterA(config-router)#neighbor 192.168.14.253 local-as 65500

Answer: B

QUESTION NO: 259

Refer to the exhibit. Router B is performing bidirectional redistribution between EIGRP and OSPF. The network 10.100.1.0/24 should not be reachable from the 10.100.9.0/24 network. However, it needs to be reachable from any network within the EIGRP domain. All other networks should be seen in both domains. Which change to router B would accomplish these goals?




Select the best response.

- A. Under the EIGRP process, insert the distribute-list block_net_1 out ospf 1 command.
- B. Under the OSPF process, insert the distribute-list block_net_1 in serial1/0 command.
- C. Under the EIGRP process, insert the distribute-list block_net_1 in serial1/0 command.
- D. Under the OSPF process, insert the distribute-list block_net_1 out eigrp 1 command.

Answer: D

QUESTION NO: 260

Refer to the exhibit. EIGRP has been configured on router D. Router C is performing mutual redistribution between EIGRP and OSPF. While verifying that the redistribution is functioning properly, you discover that while router C has all of the EIGRP routes in its routing table, router A does not have any routes from the EIGRP domain. What on router C may be the cause of the problem? Select the best response.

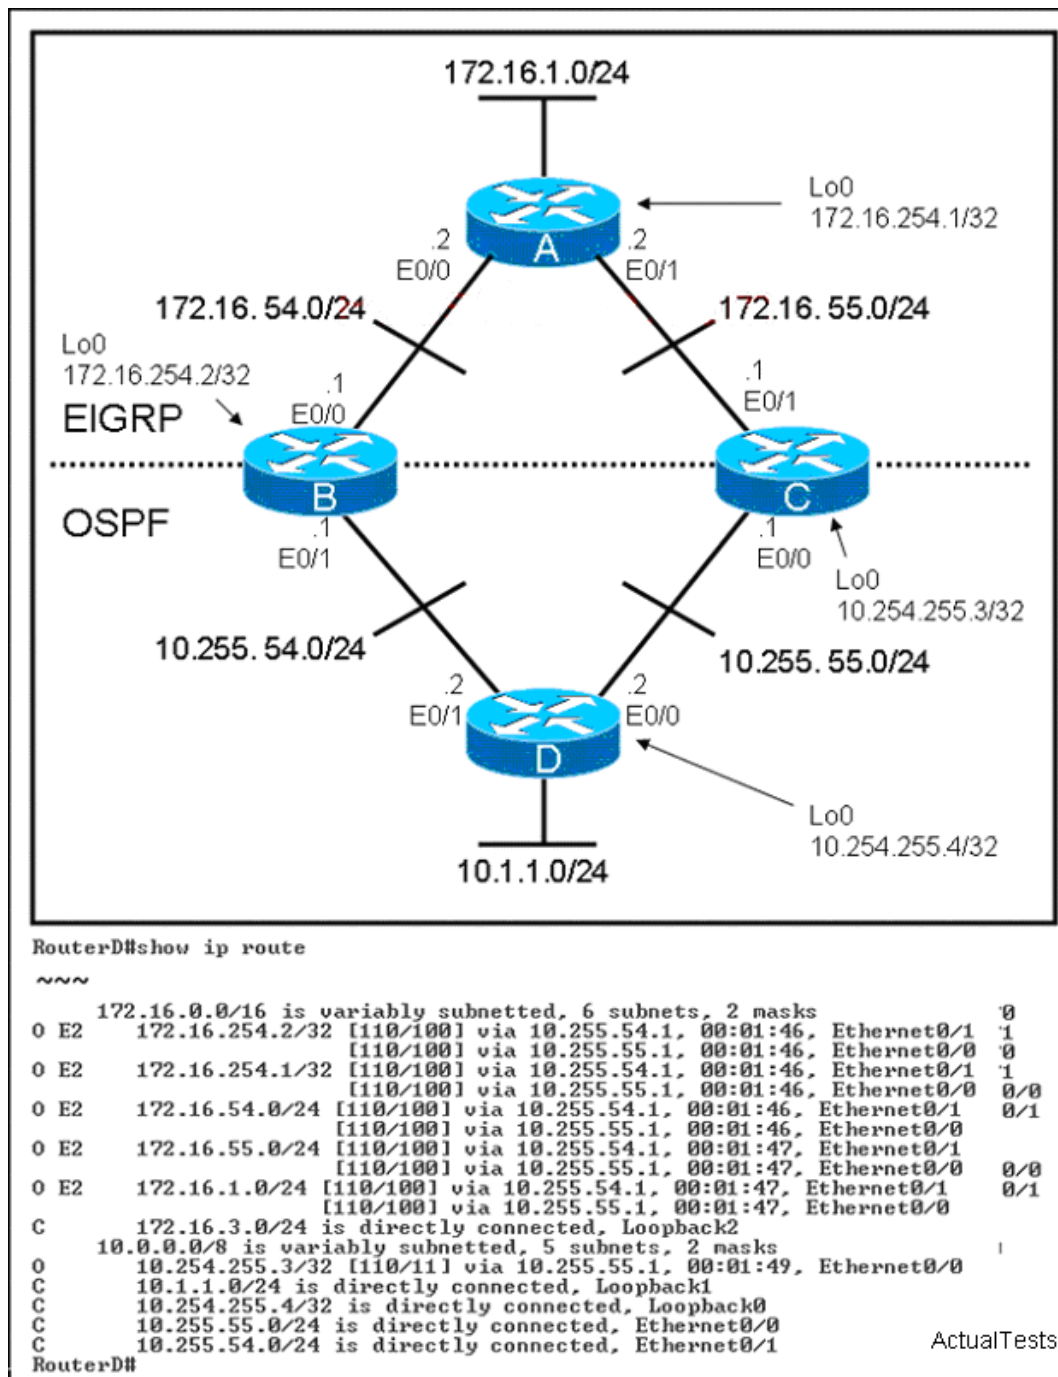


- A. The no auto-summary command needs to be added under router eigrp 1.
- B. The subnets keyword was not included in the redistribute command under router ospf 1.
- C. The metric specified for the redistributed EIGRP routes is too large; making the EIGRP routes unreachable by router A.
- D. The default-information originate command needs to be added under router ospf 1.
- E. The administrative distance of either OSPF or EIGRP must be changed so that EIGRP has a higher administrative distance than OSPF.

Answer: B

QUESTION NO: 261

Refer to the exhibit. Router B and router C are performing mutual redistribution between OSPF and EIGRP, and their default metrics are configured the same. Router D has equal cost paths to networks where both paths are not really equal cost. For example, network 172.16.54.0 shows equal cost through both router B and router C, though in reality the cost is greater using router C. Other routers, though not shown, are connected to the 172.16.54.0 and 172.16.55.0 networks, and the same issues exist to those routers and the networks connected to them. What can be done so that data will be routed along the most optimal path in the network?




Select the best response.

- A. Redistribute connected interfaces on router B and router C.
- B. Set the maximum number of equal cost paths to 1 in all routers.
- C. When redistributing EIGRP into OSPF, set the external metric type to type E1.
- D. Adjust the default metrics in router B and router C so that the values are different in each router.
- E. None of these solutions will fix the problem. Migrate to a single dynamic routing protocol.

Answer: E

QUESTION NO: 262

Refer to the exhibit. ROUTE.com is planning to implement load balancing for traffic between hosts on the 172.16.10.0/24 and 172.16.20.0/24 networks. You have been asked to review the implementation plan for this project. Which statement about the plan is true? Select the best response.

Implementation Plan


1. Configure variance on R1 and R4.
2. Use traceroute to validate load balancing has been activated.
3. Document configuration changes.

ActualTests

- A. It is complete as written.
- B. It should include a task to configure EIGRP multipath equal to 2 on R1 and R4.
- C. It should include a task to implement OSPF because it handles unequal cost load balancing most efficiently using variance.
- D. It should include a task that establishes a baseline before and after the configuration has been changed.

Answer: D**QUESTION NO: 263**

Refer to the exhibit. ROUTE.com is planning to implement load balancing for traffic between host on the 172.16.10.0/24 and 172.16.20.0/24 networks. You have been asked to review the implementation plan for this project. Which statement about the plan is true? Select the best response.



Implementation Plan

1. Establish a traffic throughput baseline.
2. Configure variance on R1 and R4.
3. Use traceroute to validate load balancing has been activated.
4. Establish a new traffic throughput baseline.
5. Compare the new and old baselines and verify that load balancing is implemented as desired.

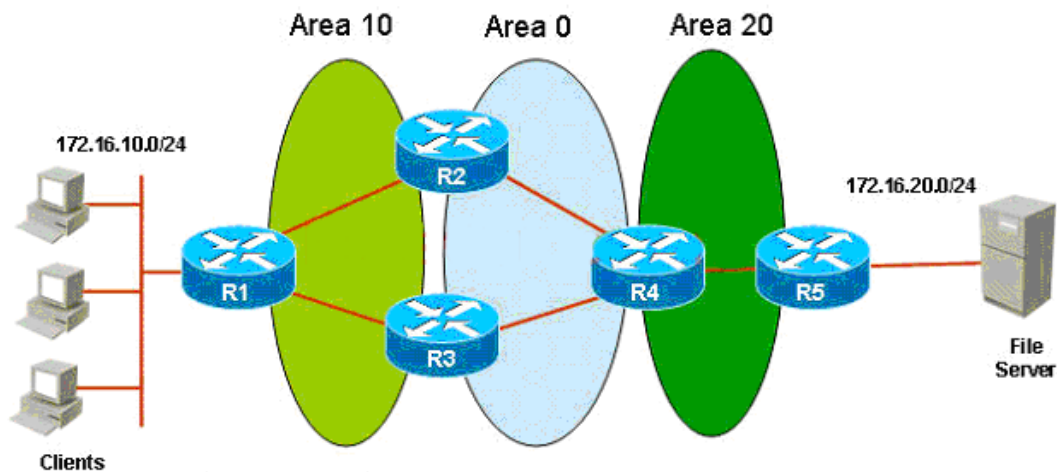
ActualTests

- A. It is complete as written.
- B. It should include a task to configure multipath to equal a value of 2 on R1 and R4.
- C. It should use a ping instead of a traceroute to validate that load balancing has been activated.
- D. It should contain a task that documents the changes made to the configurations.

Answer: D

QUESTION NO: 264

Refer to the exhibit. ROUTE.com is planning to implement a new secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation plan for the OSPF project. Which statement about the plan is true?

Implementation Plan

1. Enable OSPF process 1 on all routers.
2. Enable area 0 on R2, R3, and R4.
3. Enable area 10 on R1, R2 and R3.
4. Enable area 20 on R4 and R5.
5. Verify that all routers contain a complete routing table.
6. Verify that the clients can successfully access the file server.
7. Document configuration changes.

ActualTests

Select the best response.


- A. It is complete as written.
- B. It should include a task that shuts down all unused interfaces.
- C. It should include tasks that enable and verify OSPF authentication.
- D. It should include a task that establishes a file transfer baseline before and after the configuration is changed.

Answer: C

QUESTION NO: 265

Refer to the exhibit. ROUTE.com is planning to implement a secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation and verification plans.

Which statement about the plan is true? Select the best response.



Clients

Verification Plan


1. On R1 and R5, verify end-to-end connectivity between prefixes 172.16.10.0 and 172.16.20.0 using the extended ping command.
2. Verify that all routers have established a full neighbor relationship with the appropriate neighbors.
3. Verify the proper authentication method is active for each neighbor relationship.
4. Verify that each prefix is assigned to the appropriate OSPF area.
5. Verify the appropriate routes and summaries are in the routing table of each router.
6. Verify end-to-end connectivity between the clients and the file server.

Actual Tests

- A. It is complete as written.
- B. It should include a task that verifies that the interarea routes are using the proper MED.
- C. The plan should include a task that verifies that load sharing is active on the appropriate links.
- D. The plan should include a task that verifies end-to-end connectivity between the clients and the file server.

Answer: D**QUESTION NO: 266**

Refer to the exhibit. ROUTE.com is planning to implement a secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation and verification plans for this OSPF project. Which statement about the plan is true?



Clients

Verification Plan

1. On R1 and R5, verify end-to-end connectivity between prefixes 172.16.10.0 and 172.16.20.0 using the extended ping command.
2. Verify that all routers have established a full neighbor relationship with the appropriate neighbors.
3. Verify the proper authentication method is active for each neighbor relationship.
4. Verify that each prefix is assigned to the appropriate OSPF area.
5. Verify the appropriate routes and summaries are in the routing table of each router.
6. Verify end-to-end connectivity between the clients and the file server.

Actual Tests

Select the best response.

- A. It is complete as written.
- B. It should include a task that verifies that the interarea routes are using the proper MED.
- C. It should include a task that verifies that load sharing is active on R1 and R4.
- D. It should include a task that verifies that all redundant links will become active when the primary links are shut down.

Answer: D

QUESTION NO: 267

Which BGP command provides the router ID, local preference, next hop, and BGP path in its output?

Select the best response.

- A. show ip route bgp
- B. show ip bgp
- C. show ip bgp neighbors
- D. show ip bgp summary

Answer: B

QUESTION NO: 268

Refer to the exhibit. Which statement about this neighbor of R1 is true?

R1#show ipv6 neighbor

IPv6 Address	Age	Link-layer Addr	State	Interface
FE80::21E:79FF:FEAB:3141	2	001e.79ab.3141	STALE	Gig0/1


Select the best response.

- A. OSPFv3 adjacency has been lost, which causes the neighbor to be considered Stale.
- B. Aggregate global addresses are always used between IPv6 neighbors.
- C. OSPFv3 adjacency will not work between link-local addresses.
- D. R1 used ICMP to learn about this neighbor.

Answer: D

QUESTION NO: 269

Refer to the exhibit. Which two Cisco IOS commands on R2 would verify its OSPF neighbor relationships? (Choose two.)



Select 2 response(s).

- A. show ip ospf
- B. show ip ospf interface
- C. show ip ospf neighbor
- D. show ip ospf database
- E. show ip ospf statistics

F. show running-config | begin router ospf

Answer: B,C

QUESTION NO: 270

Which three statements about OSPF areas are true? (Choose three.) Select 3 response(s).

- A. Areas introduce a boundary on the link-state updates.
- B. Areas are logical definitions specific to any given router.
- C. All routers within an area have the exact link-state database.
- D. The calculation of the Dijkstra algorithm on a router is limited to changes within an area.
- E. The area designated router will always have a priority of 0.

Answer: A,C,D

QUESTION NO: 271

Refer to the exhibit. Which command would verify if PBR reacts to packets sourced from 172.16.0.0/16?

```
access-list 101 permit ip host 172.16.0.0 0.0.255.255 any
!
route-map divert permit 10
 match ip address 101
  set ip next-hop 212.50.185.126
  set ip next-hop recursive 192.0.0.1
  set ip default next-hop 212.50.185.125
!
interface GigabitEthernet0/1
 ip address 172.16.10.1 255.255.255.0
 ip policy route-map divert
```

ActualTests

Select the best response.

- A. show ip route
- B. show policy-map
- C. show access-lists
- D. show route-map

Answer: D

QUESTION NO: 272

Which technology manages multicast traffic at Layer 2 by configuring Layer 2 LAN interfaces dynamically to forward multicast traffic only to those interfaces that want to receive it? Select the

best response.

- A. IGMP
- B. IGMP snooping
- C. PIM-DM
- D. DVMRP
- E. MOSPF

Answer: B

QUESTION NO: 273

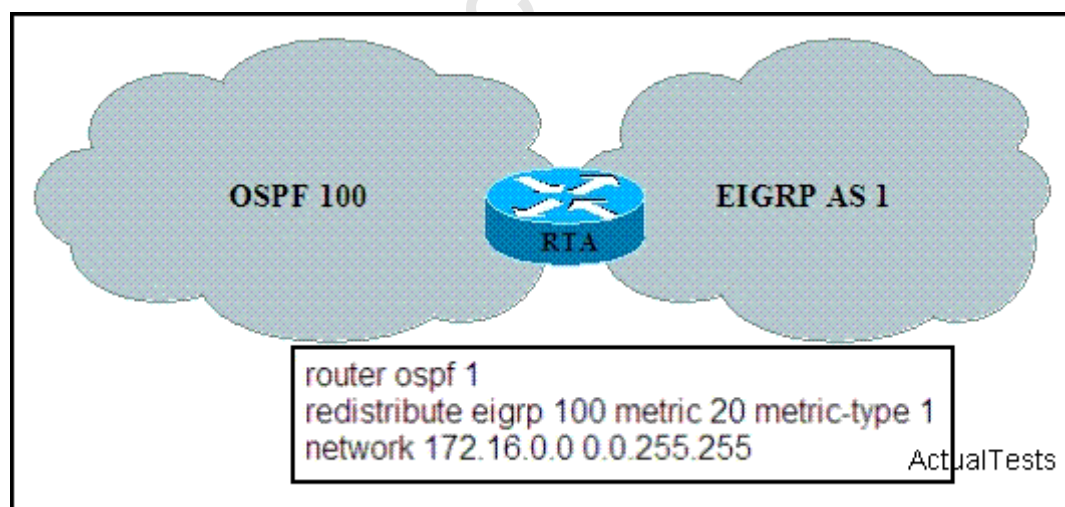
Which of the following NSAP addresses is a private, locally administered address? Select the best response.

- A. 39.0f01.0002.0000.0c00.1111.00
- B. 48.0f01.0002.0000.0c00.1111.00
- C. 49.0004.30ac.0000.3090.c7df.00
- D. 52.0f01.0002.0000.0c00.1111.00

Answer: C

QUESTION NO: 274

During the redistribution process configured on RTA, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24, are not being redistributed into the OSPF routing domain. Which two items could be a solution to this problem? (Choose two.)



- A. Change the EIGRP AS number from 100 to 1 in the redistribute command.
- B. Add the subnets option to the redistribute command.

- C. Change the metric-type to 2 in the redistribute command.
- D. Configure the redistribute command under router eigrp 1 instead.
- E. Change the metric to an EIGRP compatible metric value (bandwidth, delay, reliability, load, MTUs) in the redistribute command.

Answer: A,B

Explanation:

In this example, the router is configured for EIGRP AS 1, but EIGRP AS 100 is being redistributed into OSPF so the EIGRP AS needs to be changed from 100 to 1.

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

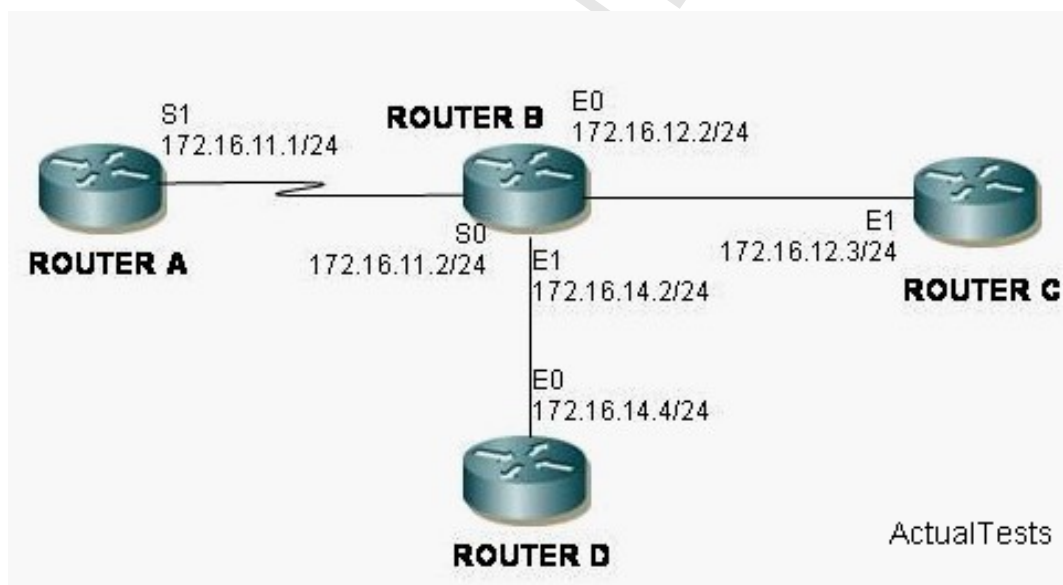
Example:

```
Router A(config)# router ospf 109 Router A(config-router)# redistribute rip subnets Router A(config-router)# network 130.10.62.0 0.0.0.255 area 0 Router A(config-router)# network 130.10.63.0 0.0.0.255 area 0
```

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

QUESTION NO: 275

A policy needs to be implemented on Router B so that any traffic sourced from 172.16.10.0/24 will be forwarded to Router C. Which configuration on Router B will achieve the desired effect?



A. access-list 1 permit 172.16.10.0 0.0.0.255

!

interface e0

ip policy route-map policy

!

route-map policy permit 10

```
match ip address 1
set ip next-hop 172.16.14.4
B. access-list 1 permit 172.16.10.0 0.0.0.255
!
interface s0
ip policy route-map policy
!
route-map policy permit 10
match ip address 1
set ip next-hop 172.16.12.3
C. access-list 1 permit 172.16.10.0 0.0.0.255
!
interface e0
ip policy route-map policy
!
route-map policy permit 10
match ip address 1
set ip next-hop 172.16.12.2
D. access-list 1 deny 172.16.10.0 0.0.0.255
!
interface s0
ip policy route-map policy
!
route-map policy permit 10
match ip address 1
set ip next-hop 172.16.12.2
```

Answer: B

QUESTION NO: 276

According to the following output, can you tell me which command can redistribute IGRP into EIGRP?

```
Router eigrp 123
Network 10.10.10.0
No auto-summary
!
Router igrp 123
Network 172.16.0.0
Network 172.17.0.0
```

- A. Under the router eigrp mode add redistribute igrp 123 subnets
- B. Under the router igrp mode add redistribute eigrp 123
- C. None, EIGRP and IGRP are automatically redistributed in this instance.
- D. Under the router eigrp mode add redistribute igrp 123

Answer: C

Explanation:

The point of this question is redistribute IGRP into EIGRP.

When redistributing IGRP into EIGRP, there is a feature that they are automatically redistributed if they have same autonomous system number; in opposite, they need to manually redistributed if they have different autonomous system number. So option D is correct.

QUESTION NO: 277


IS-IS is often considered an alternative to OSPF in the IP world. Which two statements identify similarities between IS-IS and OSPF? (Choose two.) Select 2 response(s).

- A. support for designated intermediate systems and backup designated intermediate systems
- B. support for multiple areas per router
- C. support for classless routing
- D. support for address summarization between areas
- E. support for both DIS adjacencies and neighbor adjacencies

Answer: C,D

QUESTION NO: 278

Given the network diagram, which routers currently make up the IS-IS backbone?



- A. R3,R4,R6
- B. R2,R3,R6,R7
- C. R1 through R8
- D. R3,R4,R5,R6
- E. R2,R3,R4,R6,R7

Answer: E

Explanation:

An intermediate system can be a level 1 (L1) router, a level 2 (L2) router, or both (L1/L2). L1 routers are analogous to OSPF nonbackbone Internal Routers, L2 routers are analogous to OSPF backbone routers, and L1/L2 routers are analogous to OSPF ABRs. The L1/L2 routers are connected to L1 routers and to L2 routers. These L1/L2 routers must maintain both a level 1 link-state database and a level 2 link-state database, in much the same way that an OSPF ABR must maintain a separate database for each area to which it is attached. Cisco routers are configured as L1-only, L2-only, or L1/L2 with the command `is-type`. By default, they are L1/L2.

Note: In actuality, routers R4 and R5 would also make up the part of the backbone, but since they are not given as choices the best answer is E.

QUESTION NO: 279

Based on the show ip route isis output on R1, which statement is true?

```

R1#show ip route isis
  10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
i L2  10.200.200.14/32 [115/30] via 10.1.0.2, Serial1/0
i L1  10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet0/0
i L1  10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet0/0
i L2  10.1.2.0/23 [115/20] via 10.1.0.2, Serial1/0
i su   10.1.0.0/23 [115/10] via 0.0.0.0, Null0

```

ActualTests

- A. The i su 10.1.0.0/23 route is a suppressed route.
- B. The i su 10.1.0.0/23 route is a summary route.
- C. The i su 10.1.0.0/23 route is an IS-IS external route.
- D. The R1 IS-IS router is an ASBR.
- E. The R1 IS-IS router is an ABR that belongs to multiple IS-IS areas.
- F. The R1 IS-IS router is performing route aggregation and is suppressing the more specific 10.1.0.0/23 prefix.

Answer: B

Explanation:

When viewing the "show IP route ISIS " output, the "su" entry represents a summarized route. This is the direct result of the "summary-address" router configuration command. When creating the summary route. The Cisco IOS automatically creates the summary route and point it to interface Null 0.

QUESTION NO: 280

What will happen after redistributing other routing protocol routes into OSPF?

- A. miss the metric option in the redistribute command.
- B. misconfigure the metric-type option in the redistribute command to type-1.
- C. miss the subnet option in the redistribute command.
- D. miss the tag option in the redistribute command.

Answer: C

Explanation:

The point of this question is about the redistribution of OSPF.

To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [as-number] [metric {metric-value | transparent}] [metric-type type-value] [match {internal | external 1 | external 2}] [tag tag-value] [route-map map-tag] [subnets]

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword,

only networks that are not subnetted are redistributed by OSPF.

QUESTION NO: 281

In reference to the P1R3 show isis route output, which statement is true?

```
P1R3#show isis route

IS-IS Level-1 Routing Table - version 4
System Id      Next-Hop      Interface    SNPA          Metric  State
P1R1           P1R1         Et0/0       aabb.cc00.3300 10      Up      L2-IS
P1R3           --

Default route out of area - (via 1 L2-attached IS)
System Id      Next-Hop      Interface    SNPA          Metric  State
P1R1           P1R1         Et0/0       aabb.cc00.3300 10      Up
```

Select the best response.

- A. P1R1 is the exit point out of the area for P1R3.
- B. P1R1 is a level-1 only IS-IS router.
- C. P1R1 has been configured with a nondefault IS-IS metric.
- D. P1R3 is a level-2 only IS-IS router.
- E. P1R3 has been configured with a nondefault IS-IS metric.
- F. P1R3 routing table should contain i L2 entries.

Answer: A

QUESTION NO: 282

What are the basic configuration steps to enable IS-IS?

- A. Configure the net system-id command under routerisis and enable IS-IS on each interface with the ip router isis command.
- B. Configure the net system-id and the network net-id commands under routerisis.
- C. Configure the network net-id command(s) under routerisis and enable IS-IS on each interface with the ip router isis command.
- D. Configure the network net-id command(s) and the is-type level-1-2 command under routerisis.
- E. Configure the net system-id and the network net-id commands under routerisis and enable IS-IS on each interface with the ip router isis command.

Answer: A

Explanation:

To enable Integrated IS-IS on a router for IP routing is easy. There are many more commands used to tune the IS-IS processes. However, only the following three commands are required to start Integrated IS-IS: Enable IS-IS as an IP routing protocol, using the command `router isis`, and

assign a tag if there are multiple IS-IS processes. If the tag is omitted, a tag of zero (0) is assumed. Identify the router for IS-IS by assigning a NET to the router with the net command. Enable IS-IS on the interfaces participating in IS-IS, using the command `ip router isis`. This is slightly different to most other IP routing protocols where the participating interfaces are specified by network statements. There is no network statement under the IS-IS process. If there are multiple IS-IS processes, interfaces must state which process they belong to by specifying the appropriate tag.

QUESTION NO: 283

What is periodically sent by a DIS on a LAN to ensure that all adjacent neighbors' IS-IS link-state databases are synchronized? Select the best response.

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. database query
- D. database description packet (DDP)
- E. link-state summary
- F. hello

Answer: A

QUESTION NO: 284

Which packet type is used to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link-state database? Select the best response.

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. hello
- D. database query
- E. database description packet (DDP)

Answer: B

QUESTION NO: 285

What is the default metric used on IS-IS routers for each interface? Select the best response.

- A. The cost is set to 10 for all interfaces.
- B. The cost is set to 10 for LAN interfaces and 20 for WAN interfaces.

- C. The cost is based on the speed of the interface.
- D. The cost is based on a composite of bandwidth and delay of the interface.

Answer: A

QUESTION NO: 286

How is network layer addressing accomplished in the OSI protocol suite? Select the best response.

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address

Answer: D

QUESTION NO: 287

Which two are characteristics of the IS-IS protocol but not OSPF? (Choose two.)

- A. forms adjacencies with all neighbors
- B. supports demand circuit routing
- C. provides routing support for multiple network layer protocols
- D. utilizes SPF algorithm
- E. provides for network scalability by allowing the network to be separated into areas
- F. three layers of hierarchical routing

Answer: A,C

Explanation:

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.

Differences between IS-IS and OSPF.

Although IS-IS and OSPF share many common features, they do have quite a few differences:

- * Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.
- * In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
- * IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers instead of

using IP packets, like OSPF does.

- * IS-IS allows for the preempting of DRs, where OSPF does not.
- * OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.

The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION NO: 288

Which three are benefits of IS-IS over OSPF? (Choose three.)

- A. supports route tags
- B. supports more routers in an area
- C. does not require Hello packets to establish neighbor relationships
- D. produces fewer link state advertisements for a given network
- E. supports network layer protocols other than IP
- F. requires fewer neighbor relationships in a broadcast multiaccess network

Answer: B,D,E

Explanation:

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.

Differences between IS-IS and OSPF.

Although IS-IS and OSPF share many common features, they do have quite a few differences:

- * Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.
- * In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
- * IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers instead of using IP packets, like OSPF does.
- * IS-IS allows for the preempting of DRs, where OSPF does not.
- * OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.

The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION NO: 289

Given the NSAP, 39.0100.0102.0001.0c00.1211.00, which portion is interpreted by IS-IS as the area?

- A. 39
- B. 39.0100
- C. 39.0100.0102
- D. 0001
- E. 0001.0c00
- F. 0001.0c00.1211

Answer: C

QUESTION NO: 290

A network administrator is troubleshooting a redistribution of RIP routes into OSPF. Given the exhibited configuration commands, which statement is true?

```
router rip
  network 10.0.0.0
!
router ospf 5
  network 172.10.0.0 0.0.255.255 area 0
  redistribute rip
```


Select the best response.

- A. Redistributed routes will be tagged as external type 1 (E1) with a metric of 30.
- B. Redistributed routes will be tagged as external type 2 (E2) with a metric of 20.
- C. Redistributed routes will maintain their original RIP routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as unreachable and not advertised.
- E. Redistributed routes will have a default metric of 0 but will not be treated as reachable and will be advertised.

Answer: B

QUESTION NO: 291

Refer to the exhibit. Will redistributed RIP routes from OSPF Area 2 be allowed in Area 1?



- A. RIP routes will be allowed in Area 1 only if they are first redistributed into EIGRP.
- B. Because NSSA will discard type 7 LSAs, redistributed RIP routes will not be allowed in Area 1.
- C. Redistributed RIP routes will be allowed in Area 1 because they will be changed into type 5 LSAs in Area 0 and passed on into Area 1.
- D. Redistributed RIP routes will be allowed in Area 1 because they will be changed into type 7 LSAs in Area 0 and passed on into Area 1.
- E. Because Area 1 is an NSSA, redistributed RIP routes will not be allowed.

Answer: E

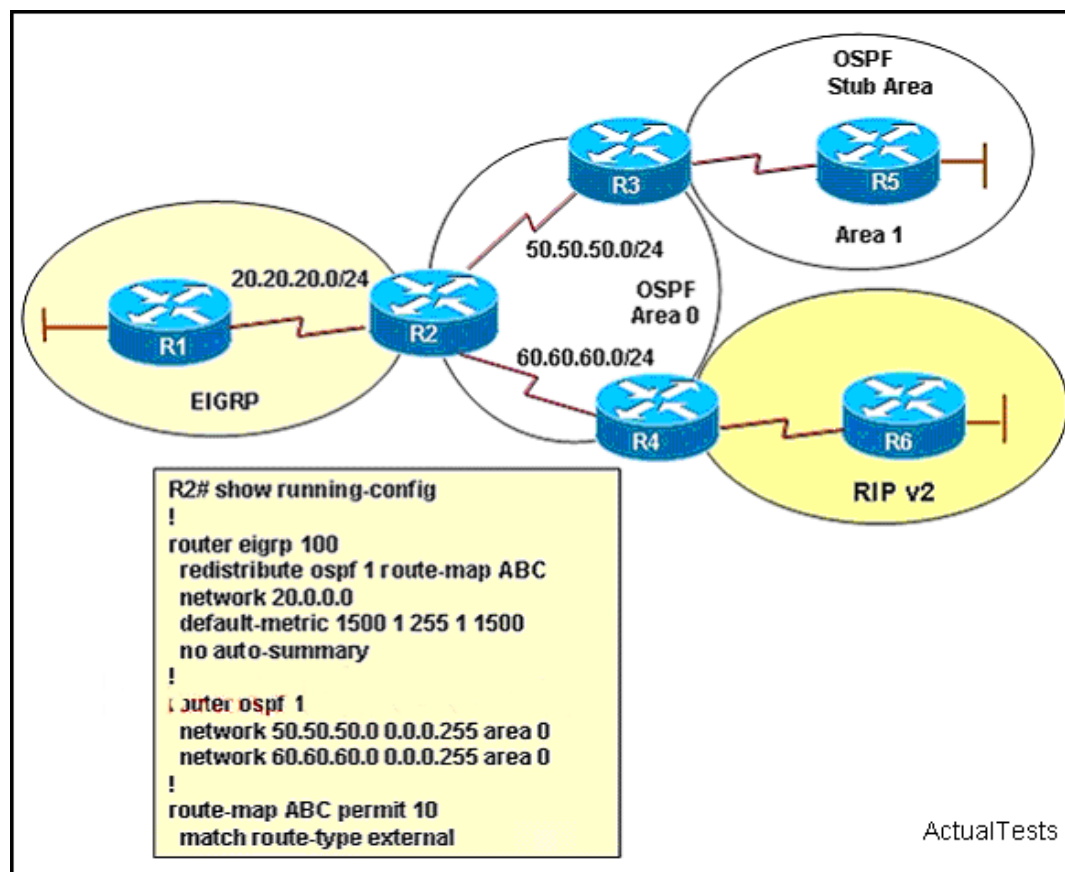
Explanation:

The following are several area types that are possible on OSPF: Standard area - A standard area can accept link updates and route summaries. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area. Stub area - A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0. Totally stubby area - A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature. Not-so-stubby area (NSSA) - An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

QUESTION NO: 292

Refer to the exhibit and the partial configuration on router R2. On router R4 all RIP routes are redistributed into the OSPF domain. A second redistribution is configured on router R2 using a route map. Based on the configuration on router R2, which EIGRP external routes will be present

in the routing table of R1? Select the best response.



- A. the routes originating from the RIP routing domain
- B. the routes originating from the OSPF stub area
- C. all OSPF inter and intra-area routes
- D. all routes originating from RIP and OSPF routing domains

Answer: A

QUESTION NO: 293

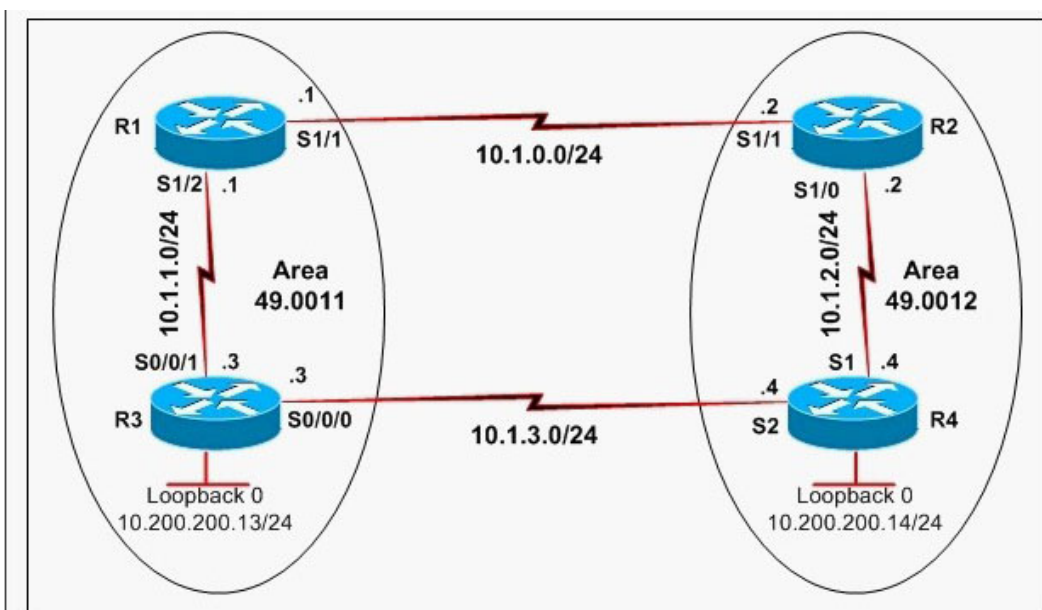
Which two statements are true about IBGP neighbor relationships? (Choose two.)

- A. The BGP split-horizon rule specifies that routes learned via EBGP are never propagated to other IBGP peers.
- B. The BGP split horizon rule specifies that routes learned via IBGP are never propagated to other IBGP peers.
- C. A full-mesh IBGP requires that neighbor relationships be established between all BGP enabled routers in the autonomous system.
- D. IBGP neighbors must be in different autonomous systems.
- E. An EGP or static routing is required between IBGP neighbors.

Answer: B,C

QUESTION NO: 294

Refer to the exhibit. What is the correct output of the command show ip route on router R2?



```
hostname R1
interface Serial1/1
ip address 10.1.0.1 255.255.255.0
ip router isis
isis circuit-type level-2-only
interface Serial1/2
ip address 10.1.1.1 255.255.255.0
ip router isis
isis circuit-type level-1
router isis
net 49.0011.1111.1111.1111.00
```

```
hostname R2
interface Serial1/0
ip address 10.1.2.1 255.255.255.0
ip router isis
isis circuit-type level-1
interface Serial1/1
ip address 10.1.0.2 255.255.255.0
ip router isis
isis circuit-type level-2-only
router isis
net 49.0012.2222.2222.2222.00
```

```
hostname R3
interface Loopback0
ip address 10.200.200.13 255.255.255.255
ip router isis
interface Serial0/0/0
ip address 10.1.3.3 255.255.255.0
ip router isis
interface Serial0/0/1
ip address 10.1.1.3 255.255.255.0
ip router isis
router isis
net 49.0011.3333.3333.3333.00
is-type level-1
```

```
hostname R4
interface Loopback0
ip address 10.200.200.14 255.255.255.255
ip router isis
interface Serial1
ip address 10.1.2.4 255.255.255.0
ip router isis
interface Serial2
ip address 10.1.3.4 255.255.255.0
ip router isis
router isis
net 49.0012.4444.4444.4444.00
is-type level-1
```

ActualTests

A. R2# show ip route

<output omitted>

10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks

i L2 10.200.200.14/32 [115/20] via 10.1.2.4, Serial1/0

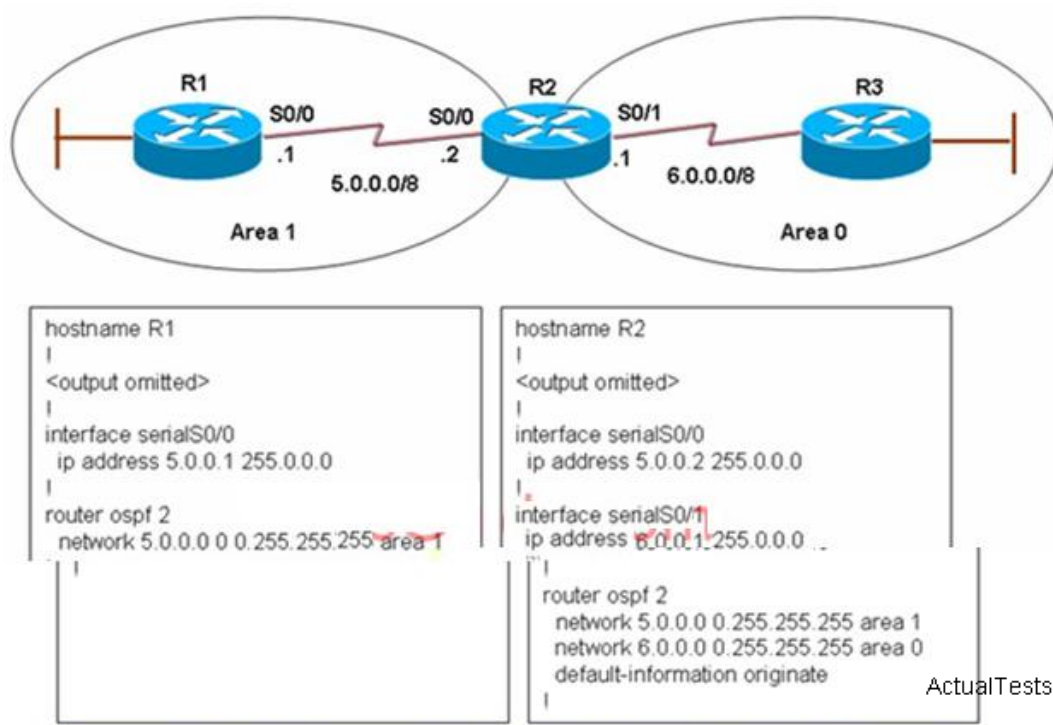
i L2 10.200.200.13/32 [115/30] via 10.1.0.1, Serial1/1

i L1 10.1.3.0/24 [115/20] via 10.1.2.4, Serial1/0

C 10.1.2.0/24 is directly connected, Serial1/0
i L2 10.1.1.0/24 [115/20] via 10.1.0.1, Serial1/1
C 10.1.0.0/24 is directly connected, Serial1/1
B. R2# show ip route
<output omitted>
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
i L1 10.200.200.14/32 [115/20] via 10.1.2.4, Serial1/0
i L2 10.200.200.13/32 [115/30] via 10.1.0.1, Serial1/1
i L1 10.1.3.0/24 [115/20] via 10.1.2.4, Serial1/0
C 10.1.2.0/24 is directly connected, Serial1/0
i su 10.1.2.0/23 [115/10] via 0.0.0.0, Null0
C 10.1.0.0/24 is directly connected, Serial1/1
i L2 10.1.0.0/23 [115/20] via 10.1.0.1, Serial1/1
C. R2# show ip route
<output omitted>
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
i L1 10.200.200.14/32 [115/20] via 10.1.2.4, Serial1/0
i L1 10.200.200.13/32 [115/30] via 10.1.2.4, Serial1/0
[115/30] via 10.1.0.1, Serial1/1
i L1 10.1.3.0/24 [115/20] via 10.1.2.4, Serial1/0
C 10.1.2.0/24 is directly connected, Serial1/0
i L1 10.1.1.0/24 [115/20] via 10.1.0.1, Serial1/1
C 10.1.0.0/24 is directly connected, Serial1/1
D. R2# show ip route
<output omitted>
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
i L1 10.200.200.14/32 [115/20] via 10.1.2.4, Serial1/0
i L2 10.200.200.13/32 [115/30] via 10.1.0.1, Serial1/1
i L1 10.1.3.0/24 [115/20] via 10.1.2.4, Serial1/0
i L1 10.1.2.0/24 is directly connected, Serial1/0
i L2 10.1.1.0/24 [115/20] via 10.1.0.1, Serial1/1
i L2 10.1.0.0/24 is directly connected, Serial1/1
E. R2# show ip route
<output omitted>
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
i L1 10.200.200.14/32 [115/20] via 10.1.2.4, Serial1/0
i L2 10.200.200.13/32 [115/30] via 10.1.0.1, Serial1/1
i L1 10.1.3.0/24 [115/20] via 10.1.2.4, Serial1/0
C 10.1.2.0/24 is directly connected, Serial1/0
i L2 10.1.1.0/24 [115/20] via 10.1.0.1, Serial1/1
C 10.1.0.0/24 is directly connected, Serial1/1

Answer: E**QUESTION NO: 295**

Refer to the exhibit. OSPF has been configured on all routers in the network. However, router R1 does not receive a default route to router R2 as intended. Which configuration change would ensure that R1 would receive a default route from R2?



Select the best response.

- A. Add the area 1 stub command on routers R1.
- B. Add the always keyword to the default-information originate configuration command on router R2.
- C. Remove the default information originate configuration command from router R2 and place it on router R1.
- D. Add the ip route 5.0.0.0 255.255.255.0 0.0.0.0 command to router R2.

Answer: B**Explanation:**

The point of this question is about advertise a default route into an ospf area.

There are two ways to advertise a default route into an ospf area. The first is to advertise 0.0.0.0 into the OSPF domain, provided the advertising router already has a default route. The second is to advertise 0.0.0.0 regardless of whether the advertising router already has a default route. The second method can be accomplished by adding the keyword always to the default-information originate command. In this example, the "always" keyword would be required in order to propagate this default route.

QUESTION NO: 296

Which two statements about the IS-IS routing protocol are true? (Choose two.) Select 2 response(s).

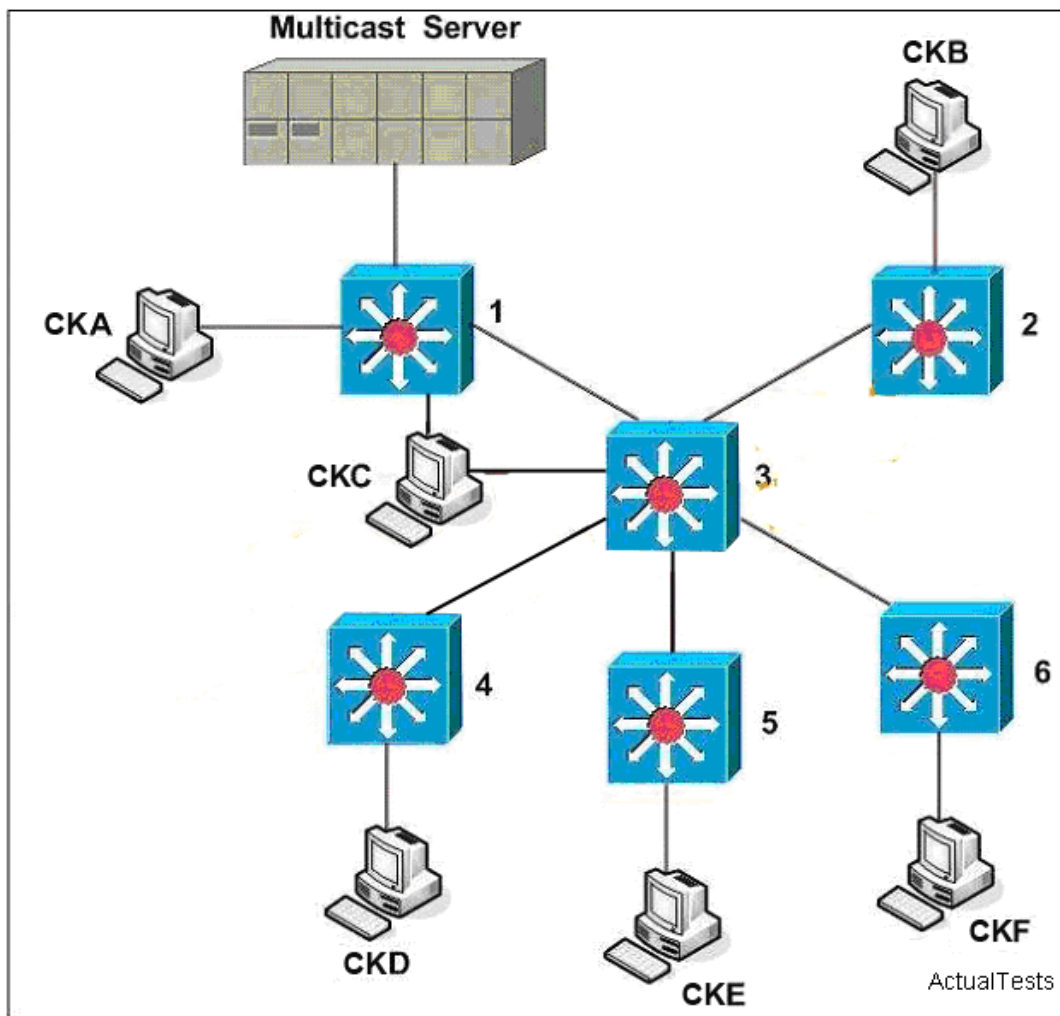
- A. IS-IS is capable of supporting IPv4 and IPv6.
- B. IS-IS is only capable of supporting IPv4 and CLNS.
- C. IS-IS routers use ES-IS hellos (ESH) to establish and to maintain neighbor relationships.
- D. IS-IS routers run the Bellman-Ford algorithm against their LSDBs to pick the best paths.
- E. Level 1 routers learn about paths within the area of which they are a part.
- F. Level 2 routers learn about paths both within areas and between areas.

Answer: A,E

QUESTION NO: 297

You are a network technician, study the exhibit carefully. All multilayer switches are running PIM sparse mode. Host CKB and Host CKF are sending IGMPv2 join messages to their respective multilayer switches. Which description is correct?

ActualTests.com



- A. Switches 1, 2, 3, and 6 will participate in the multicast tree once pruning has taken place.
- B. Switches 2 and 6 will participate in the multicast tree once pruning has taken place.
- C. Switches 1, 2, 3, and 6 will participate in the multicast tree.
- D. The multicast server is the rendezvous point of the multicast tree.

Answer: C

Explanation:

The point of this question is multicast tree

The host B and F sent IGMPv2 join message to the last hop switch, which causes the last hop switch to send join message to RP so that a shared tree is formed. The multicast source sends register message to RP to form the spt. So multicast tree exists on switches 1, 2, 3 and 6.

QUESTION NO: 298

A router is running BGP and receives more than one route for a particular prefix. Assume all the routes for this prefix have the same attributes. Which three path features would be reasons for the router to ignore some of the routes and not consider them as candidates for the best path? (Choose three.)

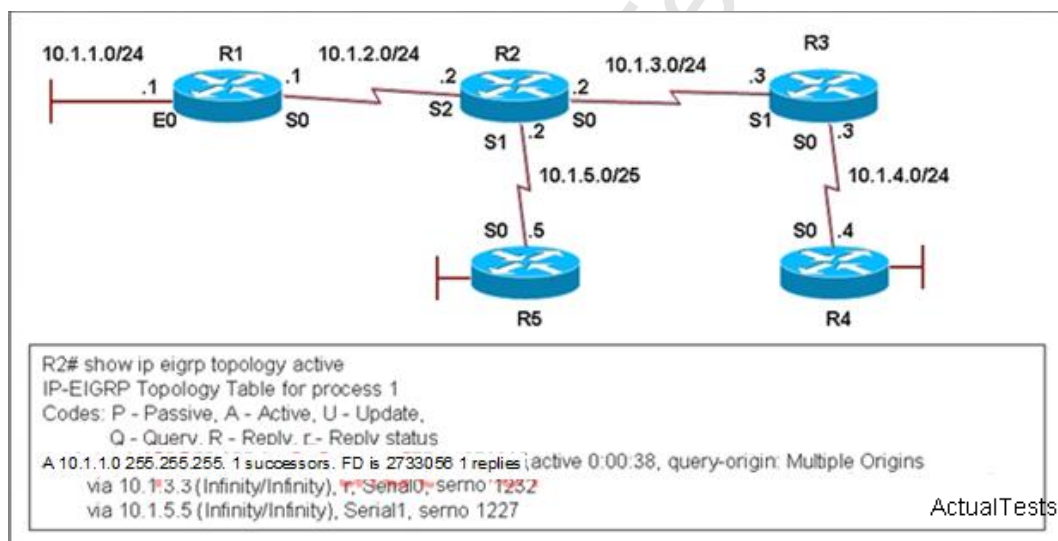
- A. paths from an internal BGP (iBGP) neighbor if the local autonomous system (AS) appears in the AS_PATH
- B. paths for which the NEXT_HOP is inaccessible
- C. paths from an external BGP (eBGP) neighbor if the local autonomous system (AS) appears in the AS_PATH
- D. paths for which the NEXT_HOP is accessible
- E. paths that are marked as synchronized in the show ip bgp output
- F. paths that are marked as not synchronized in the show ip bgp output

Answer: B,C,F

QUESTION NO: 299

Refer to the exhibit. EIGRP had converged in AS 1 when the link between router R1 and R2 went down. The console on router R2 generated the following messages: *Mar 20 12:12:06: %DUAL-5-NBRCHANGE: IP-EIGRP 1: Neighbor 10.1.4.3 (Serial0) is down: stuck in active
*Mar 20 12:15:23: %DUAL-3-SIA: Route 10.1.1.0/24 stuck-in-active state in IP-EIGRP 1. Cleaning up

The network administrator issued the show ip eigrp topology active command on R2 to check the status of the EIGRP network. Which statement best describes the reason for the error messages? Select the best response.



- A. Incorrect bandwidth configuration on router R3 prevents R2 from establishing neighbor adjacency.
- B. Incorrect bandwidth configuration on router R5 prevents R2 from establishing neighbor adjacency.
- C. Router R3 did not reply to the query about network 10.1.1.0/24 sent by router R2 .
- D. Router R5 did not reply to the query about network 10.1.1.0/24 sent by router R2 .

Answer: C

QUESTION NO: 300

Refer to the exhibit. A network administrator has configured DHCP services on the router as shown. DHCP clients connected to the FastEthernet0/0 interface are working properly. DHCP clients connected to the FastEthernet0/1 interface are not receiving addresses. Which two statements contain recommendations that will solve the problem? (Choose two.)

```
WanRouter# show running-config
<output omitted>
ip dhcp excluded-address 10.0.0.1 10.0.0.5
!
!
ip dhcp pool Central
  network 10.0.0.0 255.255.255.0
  domain-name Central
  dns-server 10.0.0.2
  netbios-name-server 10.0.0.2
!
!
interface FastEthernet0/0
  ip address 10.0.0.1 255.255.255.0
  duplex auto
  speed auto
!
interface FastEthernet0/1
  ip address 10.10.0.1 255.255.255.0
  duplex auto
  speed auto
!
<output omitted>
```

ActualTests

- A. The network shown in the output under the ip dhcp pool Central command should be changed to network 10.10.0.0 with a mask of 255.255.255.0.
- B. A second DHCP pool for network 10.10.0.0/24 should be configured.
- C. An ip dhcp excluded-address global configuration command for network 10.10.0.0/24 should be issued.
- D. The ip helper-address 10.0.0.1 command should be issued so that the address can be added to the FastEthernet0/0 configuration.
- E. The ip helper-address 10.0.0.1 command should be issued so that the address can be added to the FastEthernet0/1 configuration.

Answer: B,C

QUESTION NO: 301

Which two statements are true about IS-IS routing? (Choose two.)

- A. Based on the default timers, OSPF detects a failure faster than IS-IS does.
- B. OSPF default timers permit more tuning than IS-IS does.
- C. IS-IS is more efficient than OSPF in the use of CPU resources.
- D. OSPF is more scalable than IS-IS because of its ability to identify normal, stub, and NSSA areas.
- E. IS-IS and OSPF are both Open Standard, link-state routing protocols which support VLSM.

Answer: C,E

Explanation:

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.

Differences between IS-IS and OSPF

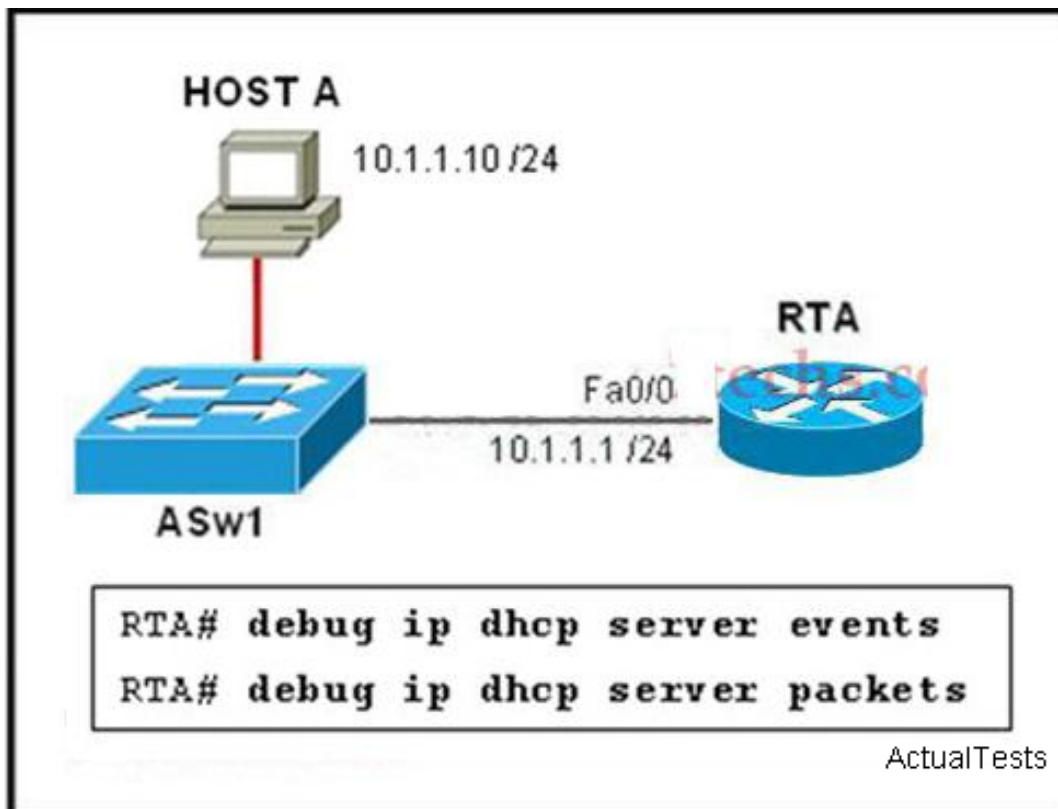
Although IS-IS and OSPF share many common features, they do have quite a few differences:

- * Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.
- *
- * In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
- * IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers instead of using IP packets, like OSPF does.
- * IS-IS allows for the preempting of DRs, where OSPF does not.
- * OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.

The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION NO: 302

Refer to the exhibit. Router RTA has been configured as a DHCP server. The two debug commands will generate output on RTA when Host A requests an IP address. Which set of DHCPD debug messages is in the correct sequence?



- A. DHCPD: Sending DHCPOFFER to client
DHCPD: DHCPDISCOVER received from client
DHCPD: DHCPREQUEST received from client
DHCPD: Sending DHCPACK to client
- B. DHCPD: DHCPDISCOVER received from client
DHCPD: DHCPREQUEST received from client
DHCPD: Sending DHCPOFFER to client
DHCPD: Sending DHCPACK to client
- C. DHCPD: DHCPDISCOVER received from client
DHCPD: Sending DHCPOFFER to client
DHCPD: DHCPREQUEST received from client
DHCPD: Sending DHCPACK to client
- D. DHCPD: DHCPREQUEST received from client
DHCPD: Sending DHCPOFFER to client
DHCPD: DHCPDISCOVER received from client
DHCPD: Sending DHCPACK to client
- E. DHCPD: Sending DHCPACK to client
DHCPD: DHCPDISCOVER received from client
DHCPD: Sending DHCPOFFER to client
DHCPD: DHCPREQUEST received from client
- F. DHCPD: DHCPDISCOVER received from client
DHCPD: Sending DHCPACK to client
DHCPD: Sending DHCPOFFER to client
DHCPD: DHCPREQUEST received from client

Answer: C

Explanation:

The following example shows a combination of DHCP server events and decoded receptions and transmissions:

```
Router# debug ip dhcp server events
Router# debug ip dhcp server packets
```

```
DHCPD:DHCPDISCOVER received from client 0b07.1134.a029 through relay 10.1.0.253.
DHCPD:assigned IP address 10.1.0.3 to client 0b07.1134.a029.
DHCPD:Sending DHCPOFFER to client 0b07.1134.a029 (10.1.0.3).
DHCPD:unicasting BOOTREPLY for client 0b07.1134.a029 to relay 10.1.0.253.
DHCPD:DHCPREQUEST received from client 0b07.1134.a029.
DHCPD:Sending DHCPACK to client 0b07.1134.a029 (10.1.0.3).
DHCPD:unicasting BOOTREPLY for client 0b07.1134.a029 to relay 10.1.0.253.
DHCPD:checking for expired leases.
```

Note that for this question, the correct order of events are highlighted above.

Reference:

http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_h1.html#wp1020307

QUESTION NO: 303

What is the MAC address that comes from the multicast address 239.255.0.1?

- A. 01-00-5e-7f-00-01
- B. 01-01-ef-ff-00-01
- C. 00-00-00-7f-00-01
- D. 10-00-ef-ff-00-01

Answer: A

Explanation:

The point of this question is transform multicast IP address into MAC address.


According to IPv4 multicast standards, the MAC destination multicast address begins with 0100:5e and is appended by the last 23 bits of the IP address. In this question the IP destination address is 239.255.0.1, so we can transform it into 0100:5e 7f:0001.

QUESTION NO: 304

Refer to the exhibit. EIGRP has been configured on all routers in the network. The command metric weights 0 0 1 0 0 has been added to the EIGRP process so that only the delay metric is

used in the path calculations.

Which router will R1 select as the successor and feasible successor for Network A?



Select the best response.

- A. R4 becomes the successor for Network A and will be placed in the routing table. R2 becomes the feasible successor for Network A.
- B. R4 becomes the successor for Network A and will be included in the routing table. No feasible successor will be selected as the advertised distance from R2 is higher than the feasible distance.
- C. R2 becomes the successor and will be placed in the routing table. R4 becomes the feasible successor for Network A.
- D. R2 becomes the successor and will be placed in the routing table. No feasible successor will be selected as the reported distance from R4 is lower than the feasible distance.

Answer: B

Explanation:

The point of the question is DUAL of EIGRP.

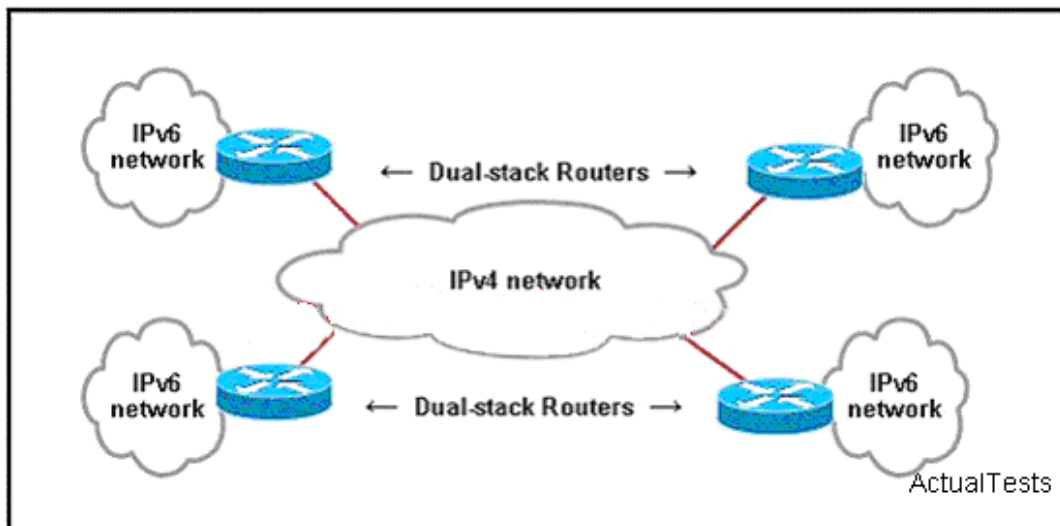
$FD=5+10+5=20$

R4 is successor

No FS, because no other router's AD is lower the FD

QUESTION NO: 305

Refer to the exhibit. Which interoperability technique implemented on the dual-stack routers would allow connectivity between IPv6 sites across automatic created tunnels using the 2002::/16 prefix? Select the best response.



- A. Dual Stack
- B. NAT-PT
- C. 6to4 tunnel
- D. GRE tunnel
- E. ISATAP tunnel

Answer: C

QUESTION NO: 306

If a metric is not specified for routes that are redistributed into OSPF, the default metric that is assigned to the routes is 20, except for redistributed BGP routes. What is the metric that is assigned to redistributed BGP routes?

- A. 200
- B. 0
- C. 1
- D. 10

Answer: C

Explanation:

If a metric is not specified, OSPF puts a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml#ospf

QUESTION NO: 307

During a redistribution of routes from OSPF into EIGRP, an administrator notices that none of the OSPF routes are showing in EIGRP. What are two possible causes? (Choose two.)

- A. incorrect distribute lists have been configured
- B. missing ip classless command
- C. CEF not enabled
- D. no default metric configured for EIGRP

Answer: A,D

QUESTION NO: 308

Which IS-IS router is equivalent to an ABR in OSPF?

- A. Level 2-3
- B. Level 1-2
- C. Level 2
- D. Level 0
- E. Level 3
- F. Level 1

Answer: B

Explanation:

The operation of IS-IS is similar to that of OSPF. IS-IS allows the routing domain to be partitioned into areas. IS-IS routing takes place at Level 1 and Level 2 within an autonomous system.

Area Border Router (ABR):

Routers that have interfaces attached to multiple areas, maintain separate LSDBs for each area to which they connect, and route traffic destined to or arriving from other areas is called Area Border Router(ABR).

Level 1-2: These routers learn about paths both within and between areas. They are equivalent to area border routers (ABRs) in OSPF

QUESTION NO: 309 DRAG DROP

Drop

Cisco 642-902: Practice Exam

Drag the appropriate IS-IS routing term to its correct definition.

Level 0 routing	Between ES's and IS's on the same subnet. OSI routing begins at this level.
Level 1 routing	Routing between separate domains. Similar to BGP.
Level 1/2 routing	Between IS's on the same area. Also called area-routing.
Level 2 routing	Called inter-area routing.
Level 3 routing	

ActualTests

Answer:

Drag the appropriate IS-IS routing term to its correct definition.

Level 0 routing	Level 0 routing
Level 1 routing	Level 3 routing
Level 1/2 routing	Level 1 routing
Level 2 routing	Level 2 routing
Level 3 routing	

ActualTests

Explanation:

Drag the appropriate IS-IS routing term to its correct definition.

Level 0 routing	Level 0 routing
Level 1 routing	Level 3 routing
Level 1/2 routing	Level 1 routing
Level 2 routing	Level 2 routing
Level 3 routing	

ActualTests
ActualTests

QUESTION NO: 310

LAB

Trebuchet Enterprises recently completed merging with Oxybeles Endeavors. The two companies have been using separate routing protocols on their corporate networks, and an immediate solution is required for the two companies to begin sharing data. A boundary router, Blockade, has been established to perform mutual redistribution of route information between the two networks. Configure route redistribution from EIGRP into IS-IS and from IS-IS into EIGRP on the boundary router per the following requirements:

1. Seed metric for EIGRP must have the following characteristics:
Bandwidth=512 Kbps

Name : Trebuchet	Name : Blockade	Name : Oxybeles
S0/0 : 10.211.211.5/30	S0/0 : 192.168.209.6/30	S0/1 : 192.168.209.5/30
Loopback1 : 1.1.1.1/32	S0/1 : 10.211.211.6/30	Loopback1 : 3.3.3.3/32
Secret Password : cisco	Secret Password : cisco	Secret Password : cisco

You will need to scroll this window and the problem statement window to view the entire problem.

To configure the router, click the picture of the host that is connected to the router via a serial console cable. The cable is represented in the diagram by a dotted line.

The Tab key

ActualTests

1. Seed metric for EIGRP must have the following characteristics:

Bandwidth=512 Kbps

2. Delay=100

3. Reliability=255

4. Load=1

5. MTU=1500

Seed metric for IS-IS must be set to 50.

Only redistribute Level-1 IS-IS routes into EIGRP.

Redistribute EIGRP routes into IS-IS as Level-1 routes.

Explanation:

```
Blockade# conf t
```

```
Blockade(config)# router isis
```

```
Blockade(config-router)# redistribute eigrp 100 level-1 metric 50
```

```
Blockade(config)# router eigrp 100
```

```
Blockade(config-router)# redistribute isis level-1 metric 512 10 255 1 1500
```

```
Blockade(config-router)# redistribute connected
```

```
Blockade# copy run start
```

Redistribution is the process of allowing routing information known in one routing protocol to be shared with another routing protocol. It should be noted here that routing protocols are the only items that can be redistributed into. You can redistribute such items as connected interfaces, static routes, and default routes into a routing protocol.

There are two types of redistribution available: one-way redistribution and mutual redistribution . One-way redistribution occurs when routing information contained in a routing protocol is shared with another routing protocol, but the other routing protocol doesn't share its routing information with the sharing protocol. Mutual redistribution is the process of two routing protocols sharing their routing information with each other.

IGRP and EIGRP IGRP and EIGRP are Cisco proprietary routing protocols. Both use a composite metric made up of the following:

- * Bandwidth
- * Delay
- * Reliability
- * Load

IGRP is a classful routing protocol, whereas EIGRP is a classless routing protocol. Remember: If IGRP and EIGRP are running on the same router and both of them are using the same autonomous system number, they will automatically redistribute with each other. If they do not have the same autonomous system number, you will need to manually redistribute between them.

These are the only IGP's we will discuss in this study guide that behave in this manner. R1# conf t

Enter configuration commands, one per line. End with CNTL/Z. R1(config)# router eigrp 100

R1(config-router)# redistribute isis level-2 metric 100 100 200 1 1500 R1(config-router)# ^Z R1#

This configuration will redistribute only the IS-IS Level 2 routes into EIGRP AS 100. If we wanted all IS-IS routes to redistribute into EIGRP AS 100, we would need to use the following

configuration: R1# conf t Enter configuration commands, one per line. End with CNTL/Z.

R1(config)# router eigrp 100 R1(config-router)# redistribute isis level-1-2 metric 100 100 200 1

1500 R1(config-router)# ^Z R1#

That's really all there is to basic IS-IS redistribution. So far, we have described basic redistribution into all of the IGP's covered in this study guide. We're now going to look at redistributing connected interfaces, static routes, and default routes into the different routing protocols. We will then conclude the redistribution section with a look at VLSM to FLSM redistribution and incorporating filters into redistribution.

QUESTION NO: 311

Which two networks does the SW1 device have feasible successors for? (choose two)

Refer to the exhibit. Bigbids incorporated is a worldwide auction provider. The network user EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question

Show commands

```

-SW1# show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001)/ID(172.29.10.1)

Codes P - Passive, A - Active, U - Update, Q - Query, R - Reply,
      r - reply Status, s - sia Status

P 172.29.3.128/25, 2 successors, FD is 30720, semo 9
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (30720/28160), FastEthernet0/3
P 10.19.0.0/24, 1 successor, FD is 156160, semo 16
  via 172.29.3.2 (156160/128256), FastEthernet0/3
  via 172.29.10.2 (157720/155160), FastEthernet0/1
P 172.29.10.0/24, 1 successors, FD is 28160, semo 1
  via Connected, FastEthernet0/1
P 172.29.0.0/30, 1 successors, FD is 20514560, semo 15
  via 172.29.1.1 (20514560/205122000), FastEthernet0/2
  via 172.29.10.2 (20516120/20513560), FastEthernet0/1
P 172.29.1.0/24, 1 successors, FD is 28160, semo 2
  via Connected, FastEthernet0/2
P 172.29.2.0/24, 1 successors, FD is 30720, semo 8
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (30720/30720), FastEthernet0/3
P 172.29.3.0/25, 1 successor, FD is 28160, semo 3
  via Connected, FastEthernet0/3
-SW1#
  
```

ActualTests

- A. 172.29.1.0/24
- B. 10.19.0.0/24
- C. 172.29.0.0/30
- D. 172.29.2.0/24

Answer: B,C

QUESTION NO: 312

Refer to the exhibit. BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's questions.

```

CoreSwitch1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.29.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.29.3.128/25, 2 successors, FD is 30720, serno 9
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (30720/28160), FastEthernet0/3
P 10.19.0.0/24, 1 successors, FD is 156160, serno 16
  via 172.29.3.2 (156160/128256), FastEthernet0/3
  via 172.29.10.2 (157720/155160), FastEthernet0/1
P 172.29.10.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet0/1
P 172.29.0.0/30, 1 successors, FD is 20514560, serno 15
  via 172.29.1.1 (20514560/205122000), FastEthernet0/2
  via 172.29.10.2 (20516120/20513560), FastEthernet0/1
P 172.29.1.0/24, 1 successors, FD is 28160, serno 2
  via Connected, FastEthernet0/2
P 172.29.2.0/24, 1 successors, FD is 30720, serno 8
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (33280/30720), FastEthernet0/3
P 172.29.3.0/25, 1 successors, FD is 28160, serno 3
  via Connected, FastEthernet0/3

```

ActualTests

CoreSwitch1#

Which three EIGRP routes will be installed for the 172.17.3.128/25 and 172.17.2.0/24 networks?
(Choose three)

- A. 172.17.3.128/25 [90/28160] via 172.17.1.2,01:26:35, FastEthernet0/2
- B. 172.17.3.128/25 [90/30720] via 172.17.3.2,01:26:35, FastEthernet0/3
- C. 172.17.3.128/25 [90/30720] via 172.17.10.2,01:26:35, FastEthernet0/1
- D. 172.17.2.0/24 [90/30720] via 172.17.10.2,02:10:11, FastEthernet0/1
- E. 172.17.2.0/24 [90/28160] via 172.17.10.2,02:10:11, FastEthernet0/1
- F. 172.17.2.0/24 [90/33280] via 172.17.3.2,02:10:11, FastEthernet0/3

Answer: B,C,D

QUESTION NO: 313

Refer to the exhibit. BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's questions.


```

CoreSwitch1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.29.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.29.3.128/25, 2 successors, FD is 30720, serno 9
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (30720/28160), FastEthernet0/3
P 10.19.0.0/24, 1 successors, FD is 156160, serno 16
  via 172.29.3.2 (156160/128256), FastEthernet0/3
  via 172.29.10.2 (157720/155160), FastEthernet0/1
P 172.29.10.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet0/1
P 172.29.0.0/30, 1 successors, FD is 20514560, serno 15
  via 172.29.1.1 (20514560/205122000), FastEthernet0/2
  via 172.29.10.2 (20516120/20513560), FastEthernet0/1
P 172.29.1.0/24, 1 successors, FD is 28160, serno 2
  via Connected, FastEthernet0/2
P 172.29.2.0/24, 1 successors, FD is 30720, serno 8
  via 172.29.10.2 (30720/28160), FastEthernet0/1
  via 172.29.3.2 (33280/30720), FastEthernet0/3
P 172.29.3.0/25, 1 successors, FD is 28160, serno 3
  via Connected, FastEthernet0/3

```

ActualTests

CoreSwitch1#

Which three networks is the router at 172.17.10.2 directly connected to? (Choose three)

- A. 172.17.0.0/30
- B. 172.17.1.0/24
- C. 172.17.2.0/24
- D. 172.17.3.0/25
- E. 172.17.3.128/25
- F. 172.17.10.0/24

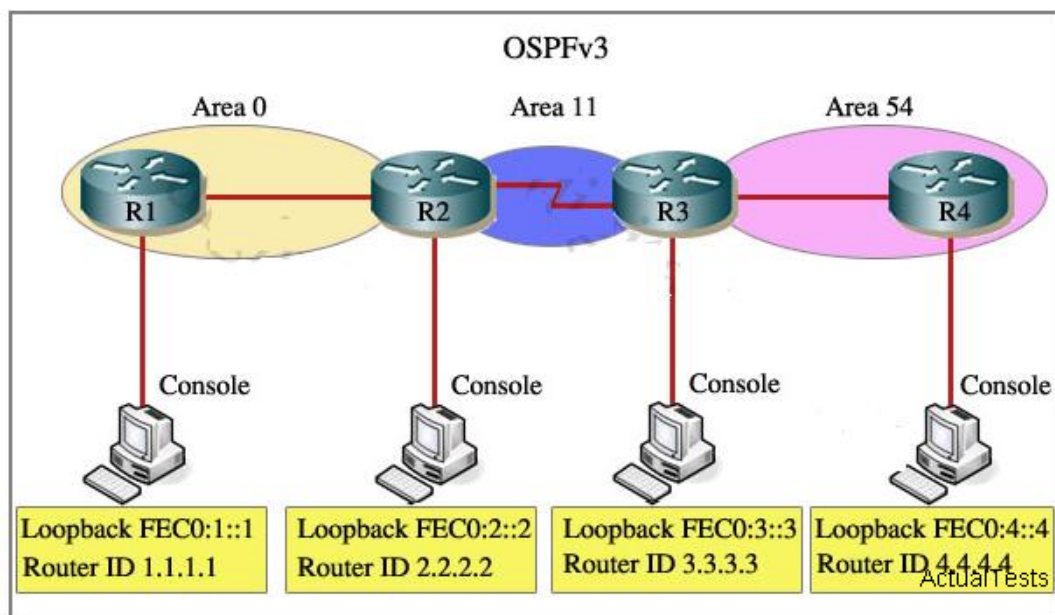
Answer: C,E,F

QUESTION NO: 314

LAB. (IPv6 OSPF Virtual Link Simulation)

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.



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

Explanation:

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
!
ActualTests
<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
!
ActualTests
<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
```

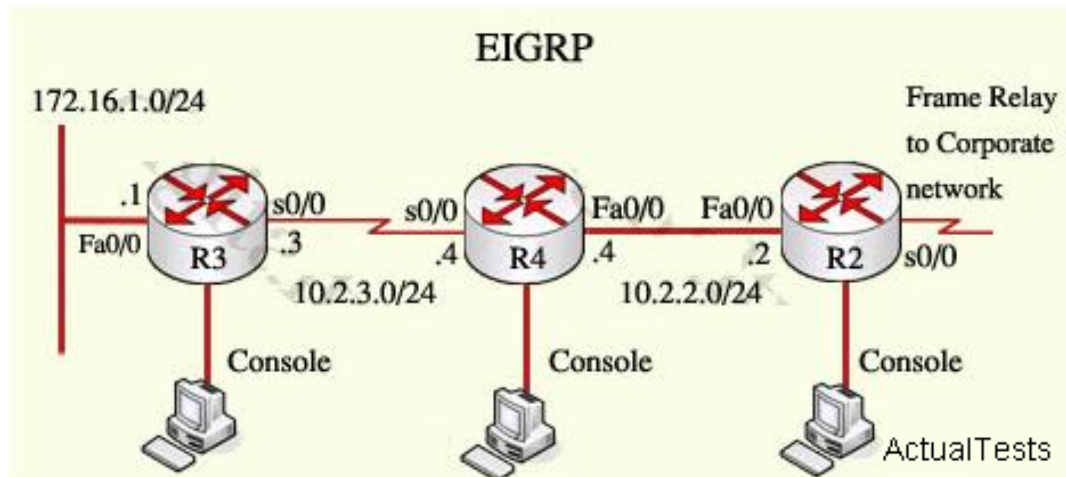
You should check the configuration of R4, too. If it has the command of "area ... virtual link ..." then remove it.

After finishing the configuration don't forget to ping between R1 and R4 to make sure they work well!

QUESTION NO: 315

LAB (EIGRP Stub Simulation)

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 stability and lower the memory usage and broadband 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.

Explanation:

First we have to figure out why R3 and R4 can not communicate with each other. Use the show running-config 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
!
ActualTests
<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, and the receive-only keyword will 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, Serial0/1
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, 00:00:02, 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 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 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
```

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

```

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

```

(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, FastEthernet0/1
C    10.2.3.0/24 is directly connected, FastEthernet0/1
    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 if in your real exam, 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.

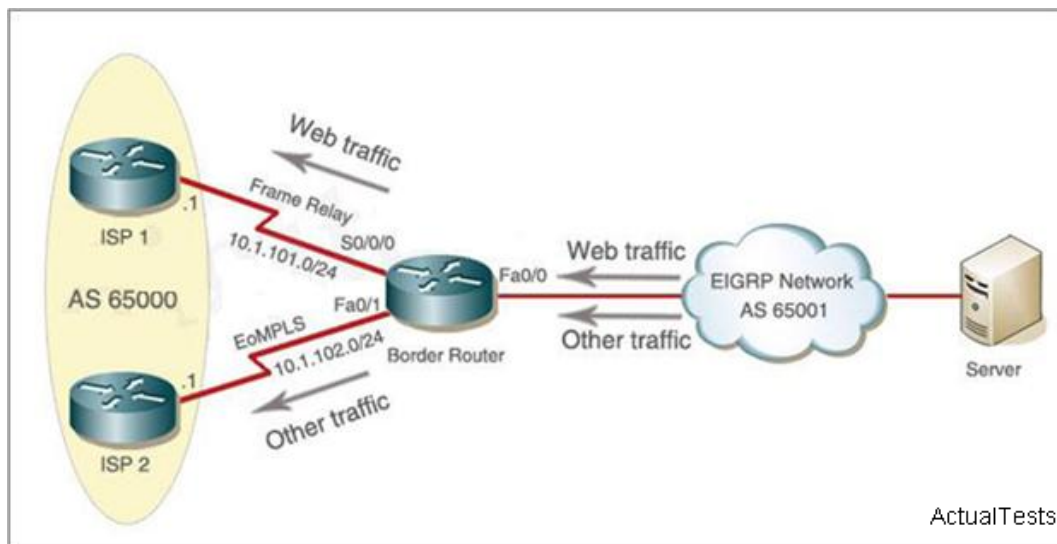
Finally don't forget to use the `copy running-config startup-config` command on routers R3 and R4 to save the configurations.

```
R4(config-if)#end R4#copy running-config startup-config
```

QUESTION NO: 316

Company 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.



Explanation:

All the HTTP traffic from the Server should go through Frame Relay link and all the other traffic should go through EoMPLS link.

The only router you are able to administrate is the Border Router, from the server you may only send HTTP traffic. As the other people mentioned, actually it is not a BGP lab. You are not able to execute the command "router bgp 65001" For the solution I did the following:

1) Access list that catches the HTTP traffic: `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. You do not know the exact IP address of the server, so the ACL catches all the source addresses. 2) Route map that sets the next hop address to be ISP1 and permits the rest of the traffic: `route-map pbr permit 10 match ip address 101 set ip next-hop 10.1.101.1 route-map pbr permit 20` 3) Apply the route-map on the interface to the server: `int fa0/0 ip policy route-map pbr` When you send traffic from the server the route map should match, you may check with: `show route-map`

QUESTION NO: 317

Which option is true regarding the synchronization rule ?

- A. Do not use or internally advertise a route until the route is learned from a source other than BGP.
- B. Do not use or advertise a route until the route is learned from a BGP peer.
- C. Do not use or advertise routes marked PARTIAL.
- D. Wait until a CONFIRM message is received before using routes from BGP neighbors.

Answer: A

QUESTION NO: 318

One of the most important characteristics of OSPF is Multiple areas. Which statement best describes why this feature is such an important enhancement to earlier routing protocols?

- A. The network domain, when divided into areas, allows for the use of both IANA classful addressing and private addressing.
- B. The use of multiple areas allows for the use of prioritization.
- C. All computation is kept within the area, with minimum communication between the areas, allowing the network to scale to larger sizes.
- D. It is easier to implement security.

Answer: C

QUESTION NO: 319

When learning a new route, if a LSA received is not found in the topological database, what will an internal OSPF router do?

- A. The sequence numbers are checked, and if the LSA is valid it is entered into the topology database.
- B. The LSA is placed in the topological database and an acknowledgement is sent to the transmitting router.
- C. The LSA is dropped and a message is sent to the transmitting router.
- D. The LSA is flooded immediately out of all the OSPF interfaces, except the interface from which the LSA was received.

Answer: D

QUESTION NO: 320

IS-IS creates and maintains adjacencies and neighbor relations by using the Hello protocol. For the following options, which are Hello packets used in IS-IS?

- A. A generic Hello
- B. Point-to-point Level 1, point-to-point Level 2, LAN Level 1, LAN Level 2
- C. Point-to-point Level 1, point-to-point Level 2, LAN
- D. Point-to-point, LAN Level 1, LAN Level 2

Answer: D

QUESTION NO: 321

The following exhibit shows ipv6 route output. What would the metric be for a summary route that summarizes all three OSPFv3 routes displayed?

```
OI 2001:OD B 8:0:0:7::/64 [110/20]
    via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:OD B 8:0:0:8::/64 [110/100]
    via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:OD B 8:0:0:9::/64 [110/40]
    via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0
```

- A. 160
- B. 140
- C. 120
- D. 100

Answer: D

Explanation:

The cost of the summarized routes is the highest cost of the routes being summarized. In fact, in the old RFC 1583 standard, the cost of the summary route was the cost of the lowest metric. But when OSPF was updated in RFC 2178 and RFC 2328, the summary route should have the same cost as the highest-cost summarized route. In this case, the highest-cost is 100 according to the second entry.

QUESTION NO: 322

Which statement is correct based upon the following output from the show command on RT1?

```
RT1# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
Link Local Address FE80:218:B9FF:FE12:2CD1, Interface ID 4
Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.1
Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State Drother, Priority 1
Designated Router (ID) 10.1.3.1, local address FE80::218:B9FF:FECD:BEF0
Backup Designated router (ID) 10.1.2.1, local address FE80::218:B9FF:FE92:28D8
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:04
Index 1/3/3, flood queue length 0
Next 0x(0) /0x0(0) /0x0(0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 10.1.3.1 (Designated Router)
  Adjacent with neighbor 10.1.2.1 (Backup Designated Router)
Suppress help for 0 neighbor(s)
```

- A. OSPFv3 establishes neighbor adjacencies by using global IPv6 addresses.
- B. OSPFv3 establishes neighbor adjacencies by using Link-local addresses.
- C. RT1 owns a subnet mask of 64 bits.
- D. OSPFv3 establishes neighbor adjacencies by using IPv4 addresses.

Answer: B

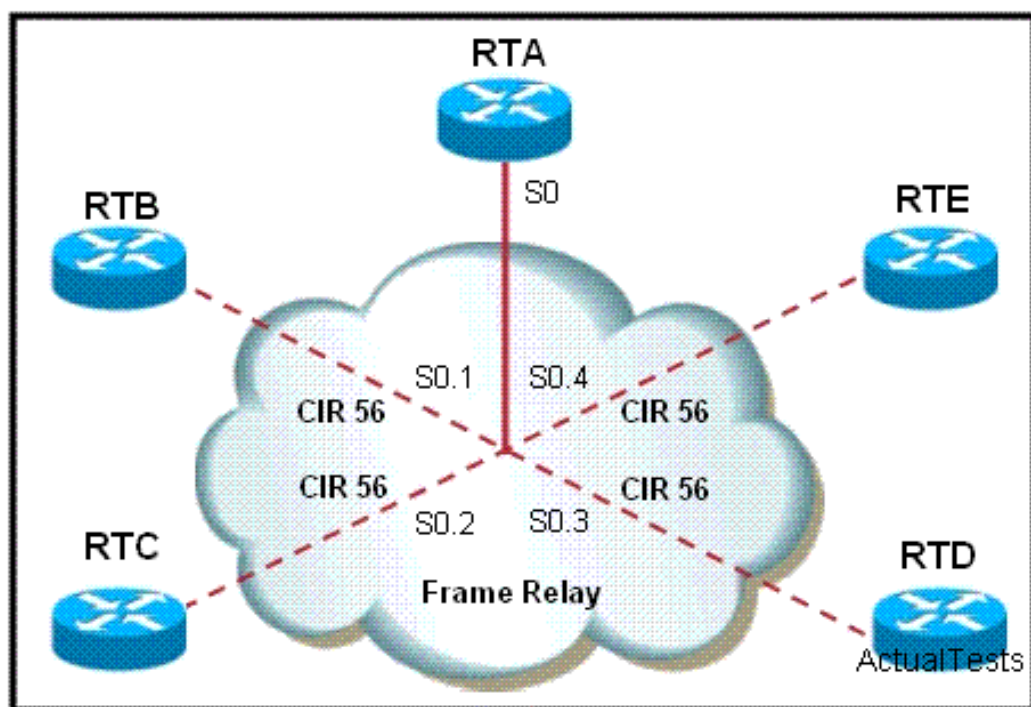
Explanation:

The point of this question is the OSPFv3.

As shown information, we notice that the router with both ipv6 address and ipv4 address and both of them are in ospf process, so in OSPFv3, we can find that, the DR BDR address are all link local address, so in OSPFv3 it establish neighbor with link local address.

QUESTION NO: 323

Refer to the exhibit. What must be done on router RTA to ensure that EIGRP works effectively in a Frame Relay multipoint environment?



- A. Issue the command bandwidth 224 on the physical interface.
- B. Issue the command bandwidth 56 on the physical interface.
- C. Issue the command bandwidth 224 on each subinterface.
- D. Issue the command bandwidth 56 on each subinterface.

Answer: A

QUESTION NO: 324

Which three route filtering statements are true? (Choose three.)

- A. After the router ospf 10 and passive-interface s0/0 commands have been issued , the s0/0 interface will not send any OSPF updates, but will receive routing updates on that interface.
- B. When you use the passive-interface command with RIPv2, multicasts are sent out the specified interface.
- C. After the router rip and passive-interface s0/0 commands have been issued, the s0/0 interface will not send any RIP updates, but will receive routing updates on that interface.
- D. When you use the passive-interface command with EIGRP, hello messages are not sent out the specified interface.
- E. After the router eigrp 10 and passive-interface s0/0 commands have been issued, the s0/0 interface will not send any EIGRP updates, but will receive routing updates on that interface.
- F. When you use the passive-interface command with OSPF, hello messages are not sent out the specified interface.

Answer: C,D,F

QUESTION NO: 325

Examine the following options, Multicasting supports applications that communicate:

- A. many-to-one
- B. one-to-one
- C. one-to-many
- D. many-to-many

Answer: C

QUESTION NO: 326

Refer to the exhibit. On the basis of the partial configuration, which two statements are correct? (Choose two.)

```
<output omitted>
!
router rip
  distribute-list 2 out ethernet 0
  distribute-list 1 out
!
access-list 1 permit 10.0.0.0 0.255.255.255
access-list 2 permit 10.0.1.0 0.0.0.255
!
<output omitted>
```

ActualTests

- A. Only routes 10.0.1.0/24 will be sent out all interfaces.
- B. Only routes matching 10.0.0.0/8 will be advertised out interfaces other than Ethernet 0.
- C. Only routes matching 10.0.0.0/8 will be advertised out Ethernet 0.
- D. Only routes matching 10.0.1.0/24 will be advertised out Ethernet 0.
- E. Only routes 10.0.1.0/24 will be allowed in the routing table.
- F. All routes will be advertised out interfaces other than Ethernet 0.

Answer: B,D

QUESTION NO: 327

What is the meaning of the default value of the EIGRP variance command of 1?

- A. Only the path that is the feasible successor should be used.
- B. Load balancing is disabled on this router.
- C. The router only performs equal-cost load balancing on all paths that have a metric greater than 1
- D. The router performs equal-cost load balancing.

Answer: D

Explanation:

The point of the question is about the balance configuration of EIGRP. If variance is 1, it means that it support Equal cost path.

QUESTION NO: 328

The following exhibit shows ipv6 route output. What would the metric be for a summary route that summarizes all three OSPFv3 routes displayed?

```

OI 2001:OD B 8:0:0:7::/64 [110/20]
   via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:OD B 8:0:0:8::/64 [110/100]
   via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:OD B 8:0:0:9::/64 [110/40]
   via FE 80::A 8B B:CCFF:FE 00:6F00, FastEthernet0/0

```

ActualTests

- A. 160
- B. 140
- C. 120
- D. 100

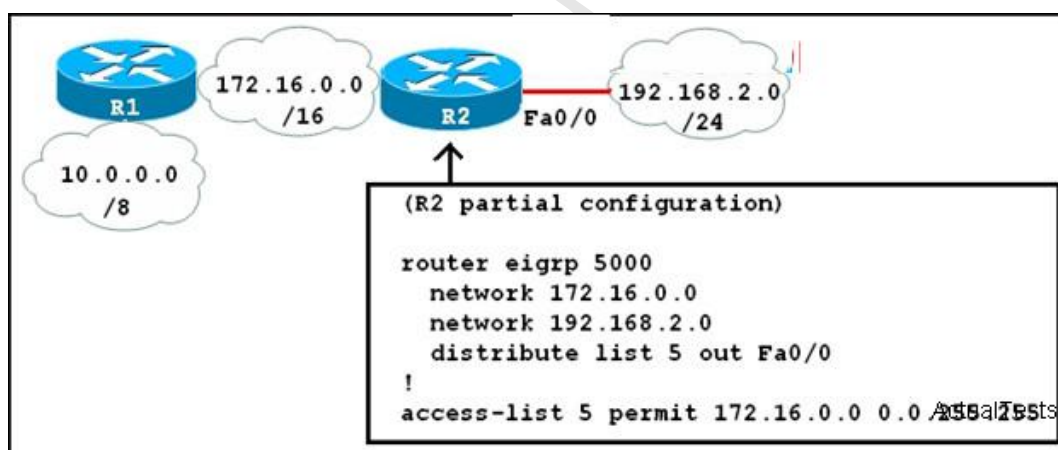
Answer: D

Explanation:

The cost of the summarized routes is the highest cost of the routes being summarized. In fact, in the old RFC 1583 standard, the cost of the summary route was the cost of the lowest metric. But when OSPF was updated in RFC 2178 and RFC 2328, the summary route should have the same cost as the highest-cost summarized route. In this case, the highest-cost is 100 according to the second entry.

QUESTION NO: 329

Refer to the exhibit. R1 and R2 have been configured to share routing information via EIGRP. What will be the result of the configuration section shown for R2?



- A. Any routes learned by R2 from the interface tied to the 172.16.0.0 network will not be advertised to neighbors on the 192.168.2.0 network.
- B. Only routes learned by R2 from the interface tied to the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.
- C. Only the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.

D. All networks, except the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.

Answer: C

QUESTION NO: 330

Look at the following exhibit. Which of the following correctly states the routes to be redistributed into OSPF? (Choose two)

```
<output omitted>
!
router ospf 10
  redistribute rip route-map rip-in
!
<output omitted>
!
route-map rip-in permit 10
  match ip address 10 20
  set metric 100
  set metric-type type-1

route-map rip-in deny 20
  match ip address 30

route-map rip-in permit 30
  set metric 200
  set metric-type type-2
!
access-list 10 permit 10.0.10.0 0.0.0.255
access-list 20 permit 192.168.1.0 0.0.255
access-list 30 permit 10.0.0.0 0.255.255.255
```

- A. The network 10.0.10.0/24 will be allowed and assigned a metric of 200
- B. All networks except 10.0.0.0/8 will be allowed and assigned a metric of 200
- C. The network 172.16.0.0/16 will be allowed and assigned a metric of 200
- D. The network 192.168.1.0 will be allowed and assigned a metric of 100

Answer: C,D

QUESTION NO: 331

Refer to the exhibit. Which statement is correct regarding the router configuration shown?

```
<output omitted>
!
interface FastEthernet 0/0
  ip address 192.168.23.1 255.255.255.0
  ip helper-address 172.16.31.248
  ip helper-address 172.16.18.11
  ip helper-address 172.16.21.172
!
<output omitted>
```

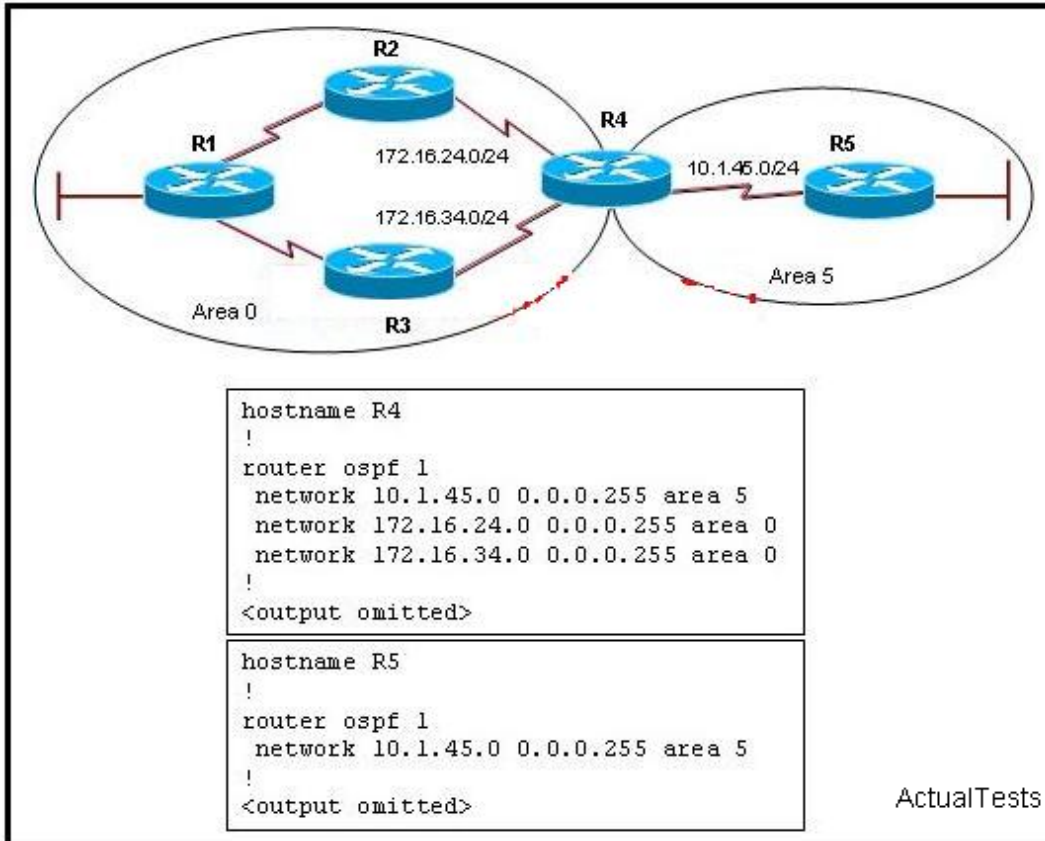
ActualTests

- A. The router will rebroadcast any DHCP requests received by this interface out to all interfaces included in the ip helper-address commands.
- B. The router will rebroadcast any UDP broadcasts received by this interface out to all interfaces included in the ip helper-address commands.
- C. The router will unicast any DHCP requests received by this interface out to all addresses included in the ip helper-address commands.
- D. This configuration will not do anything because the ip helper-address command needs to also be applied to the 192.168.23.1 IP address.

Answer: C

QUESTION NO: 332

Refer to the exhibit. What additional commands should be used to configure OSPF area 5 as a Totally Stubby area?



- A. area 0 stub on routers R4 and R5
- B. area 5 stub on routers R4 and R5
- C. area 5 stub no-summary on routers R4 and R5
- D. area 0 stub no-summary on router R4 and area 5 stub no-summary on router R5
- E. area 5 stub no-summary on router R4 and area 5 stub on router R5

Answer: E

Explanation:

The point of the question is the configuration of totally stub.

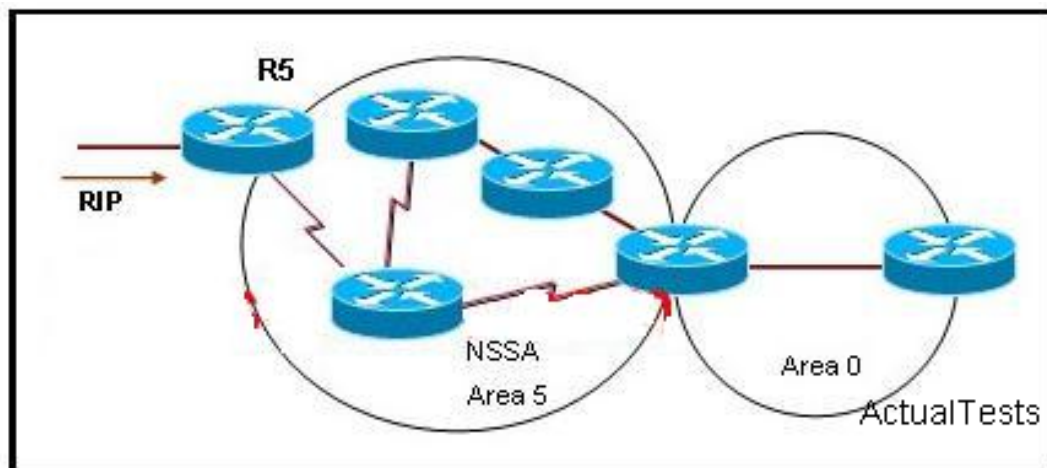
From this picture, we can find that R4 is an ABR, network 10.1.45.0/24 is in area 5. We have two commands to configure totally stub.

First: area 5 stub no-summary, it's on ABR.

Second: area xxx stub ,it's on other routers which belong to totally stub except ABR.

QUESTION NO: 333

Refer to the exhibit. OSPF is configured on all routers in the network. Area 5 is configured as an NSSA area. The RIPv2 routes are redistributed into the OSPF domain on router R5. What two types of LSAs will be originated by router R5? (Choose two.)



- A. type 1 Router LSA
- B. type 2 Network LSA
- C. type 3 Network Summary LSA
- D. type 4 ASBR Summary LSA
- E. type 5 AS External LSA
- F. type 7 NSSA External LSA

Answer: A,F

Explanation:

The point of the question is the LSA types in different kinds of areas.

NSSA area has 4 kinds of LSA : 1, 2, 3, 7

There is no LSA3 because it is produced by ABR to advertise routing entries between OSPF areas. There is no LSA2 because there are serial interface but no fast Ethernet interface, no DR elected, no LSA2. LSA2 is produced by DR.

QUESTION NO: 334


Which routing protocol will continue to receive and process routing updates from neighbors after the passive-interface router configuration command is entered?

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

Answer: B

QUESTION NO: 335

Given the network diagram, which address would successfully summarize only the networks seen?



- A. 192.168.0.0/24
- B. 192.168.8.0/20
- C. 192.168.8.0/21
- D. 192.168.12.0/20
- E. 192.168.16.0/21
- F. These networks cannot be summarized.

Answer: C

QUESTION NO: 336

Why should iBGP sessions be fully meshed within a Transit AS?

- A. BGP requires redundant TCP sessions between iBGP peers.
- B. Routes learned via eBGP are never propagated to other iBGP peers.
- C. A full mesh allows for optimal routing within the Transit AS.
- D. Routes learned via iBGP are never propagated to other eBGP peers.
- E. Routes learned via iBGP are never propagated to other iBGP peers.

Answer: E

Explanation:

Any two routers that have formed a TCP connection in order to exchange BGP routing information are called peers or neighbors. It is important to remember that the BGP peers will never become established unless there is IP connectivity between the two peers.

BGP does not advertise routes learned by way of IBGP peers to other IBGP peers. If BGP did, BGP routing inside the AS would present a dangerous potential for routing loops. For IBGP routers to learn about all BGP routes inside the AS, they must connect to every other IBGP router in a full IBGP mesh. This full mesh needs to be only logical, not physical. In other words, as long as the IBGP peers can connect to each other using TCP/IP, a logical full mesh can be created even if the routers are not directly connected

QUESTION NO: 337

Which three characteristics apply to IS-IS but not to OSPF? (Choose three.)

- A. uses a default IOS metric of 10 on each interface
- B. uses an on-demand circuit to reduce the hello and LSA flooding across switched WAN links, such as ISDN
- C. uses a DIS and a backup DIS to present the pseudo-node on the LAN
- D. encapsulates PDUs directly into a data-link frame
- E. uses stubby areas to improve network scalability
- F. runs PRC (Partial Route Calculations) to calculate IP reachability information

Answer: A,D,F

Explanation:

The configuration of OSPF is based on a central backbone, Area 0, with all other areas being physically attached to Area 0. Because of this, certain design constraints will inevitably exist. A good, consistent IP addressing structure is necessary when this type of hierarchical model is used. It is used to summarize addresses into the backbone, and reduce the amount of information that is carried in the backbone and advertised across the network. In comparison, IS-IS also has a hierarchy with Level 1 and Level 2 routers. With IS-IS, the area borders lie on the links. However, significantly fewer link-state PDUs (LSPs) are used. Therefore, many more routers, up to 1000, can reside in a single area. This capability makes IS-IS more scalable than OSPF. IS-IS allows a more flexible approach to extending the backbone. Adding Level 2 routers can extend the backbone. This process is less complex than with OSPF.

With regard to CPU use and the processing of routing updates, IS-IS is more efficient. Not only are there fewer LSPs to process, as compared to OSPF LSAs, but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive.

Both OSPF and IS-IS are link-state protocols and therefore provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and types of routers. Based on the default timers, IS-IS will detect a failure quicker than OSPF and therefore should converge more rapidly. If there are many neighbors and adjacencies to consider, the convergence time depends on the processing power of the router. IS-IS is typically less CPU

intensive than OSPF

QUESTION NO: 338

Which three statements are correct about the differences in IS-IS and OSPF? (Choose three.)

- A. IS-IS is more CPU-intensive than is OSPF.
- B. For greater fine tuning there are more IS-IS timers.
- C. IS-IS LSP contains TLV fields and OSPF LSU contains the LSAs.
- D. New additions to the protocol are easily implemented in OSPF but not with IS-IS.
- E. OSPF has more area types than does IS-IS.

Answer: B,C,E

Explanation:

The configuration of OSPF is based on a central backbone, Area 0, with all other areas being physically attached to Area 0. Because of this, certain design constraints will inevitably exist. A good, consistent IP addressing structure is necessary when this type of hierarchical model is used. It is used to summarize addresses into the backbone, and reduce the amount of information that is carried in the backbone and advertised across the network. In comparison, IS-IS also has a hierarchy with Level 1 and Level 2 routers. With IS-IS, the area borders lie on the links. However, significantly fewer link-state PDUs (LSPs) are used. Therefore, many more routers, up to 1000, can reside in a single area. This capability makes IS-IS more scalable than OSPF. IS-IS allows a more flexible approach to extending the backbone. Adding Level 2 routers can extend the backbone. This process is less complex than with OSPF.

With regard to CPU use and the processing of routing updates, IS-IS is more efficient. Not only are there fewer LSPs to process, as compared to OSPF LSAs, but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive.

Both OSPF and IS-IS are link-state protocols and therefore provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and types of routers. Based on the default timers, IS-IS will detect a failure quicker than OSPF and therefore should converge more rapidly. If there are many neighbors and adjacencies to consider, the convergence time depends on the processing power of the router. IS-IS is typically less CPU intensive than OSPF.

The timers in IS-IS allow more tuning than OSPF. There are more timers to adjust, and therefore finer granularity can be achieved. By tuning the timers, convergence time can be significantly decreased. However, this speed may be at the expense of stability, so a compromise may have to be made. A network engineer should understand the implications of adjusting these timers.

QUESTION NO: 339

Which two provide intra-area routing services? (Choose two.)

- A. L1 ES
- B. L2 IS
- C. L2 ES
- D. L1 IS
- E. L1/L2 IS

Answer: D,E

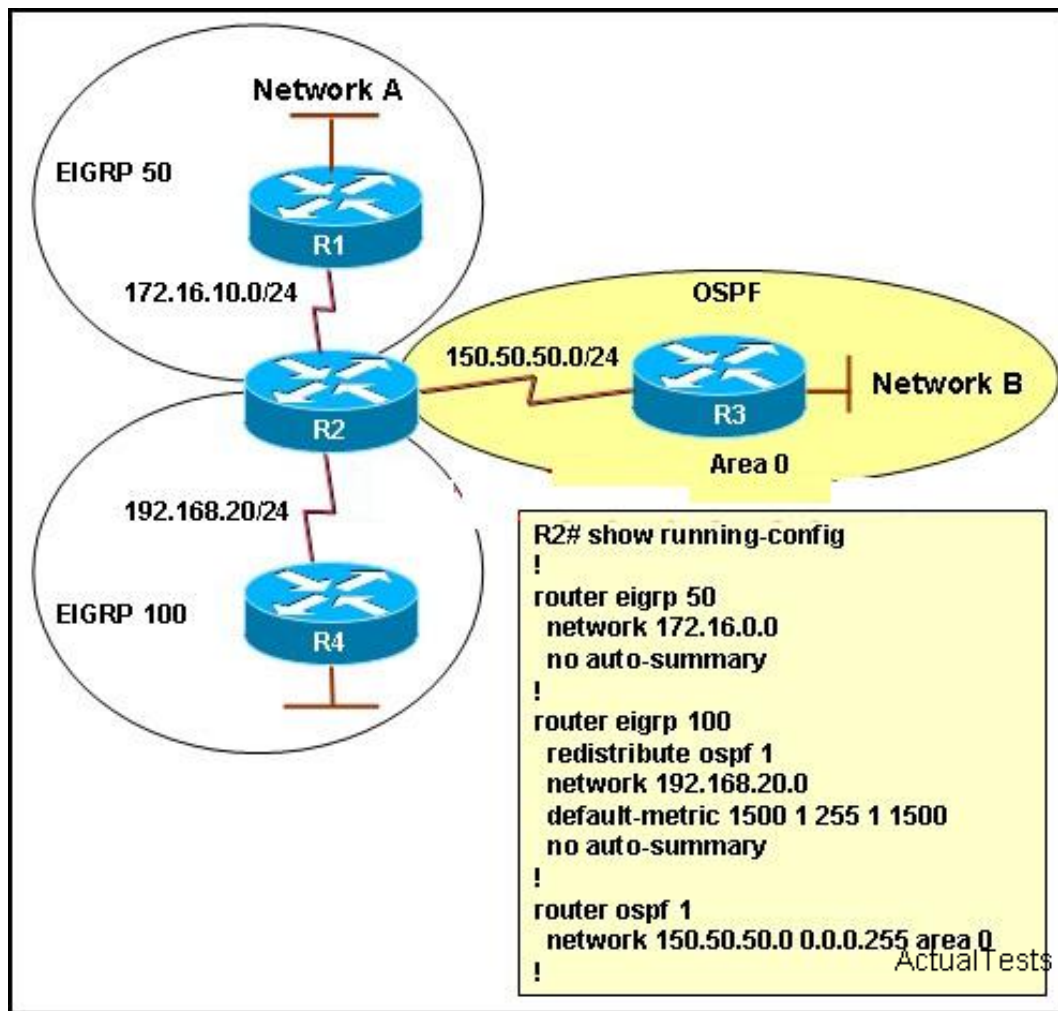
Explanation:

An IS-IS network is termed a domain, analogous to an autonomous system (AS) in TCP/IP. The following shows IS-IS within the domain as a two-level hierarchy: Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area. L1 routers are also referred to as station routers because they enable stations to communicate with each other and the rest of the network. A contiguous group of L1 routers defines an area. The L1 routers maintain the L1 link-state PDU database (LSPD), which defines the picture of the area itself and the exit points to neighboring areas. Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only. Level 1 and Level 2 (L1L2) Intermediate ISs, closely equivalent to area border routers (ABRs) in OSPF, route between areas and the backbone. They participate in the L1 intra-area routing and the L2 inter-area routing.

L1L2 routers support an L1 function to communicate with the other L1 routers in their area and maintain the L1 LSP information in an L1 LSPD. They inform other L1 routers that they are an exit point for the area. They also support an L2 function to communicate with the rest of the backbone and maintain an L2 topology database separately from their L1 LSPD. IS-IS does not share the concept of an Area 0 with OSPF. An IS-IS domain appears as a set of distinct areas interconnected by a chain of L2 routers, weaving their way through and between the Level 1 areas.

QUESTION NO: 340

Refer to the exhibit. The routing protocols EIGRP and OSPF have been configured as indicated in the exhibit. Given the partial configuration of router R2, which network will be present in the routing table of R4?



- A. Network B
- B. Network A and Network B
- C. Network A
- D. neither Network A nor Network B

Answer: A

Explanation:

In this exhibit the OSPF domain is redistributed into the EIGRP 100 domain so Network B will present into Router R4. However, the Network A network will not be seen on router R4 (The bottom router which is improperly labeled Network B) because EIGRP 50 was not redistributed into EIGRP 100.

QUESTION NO: 341

When policy-based routing (PBR) is being configured, which three criteria can the set command specify? (Choose three.)

- A. all interfaces in the path toward the destination
- B. all networks in the path toward the destination

- C. adjacent next hop router in the path toward the destination
- D. all routers in the path toward the destination
- E. all interfaces through which the packets can be routed
- F. type of service and precedence in the IP packets

Answer: C,E,F

Explanation:


Configuring Policy-Based Routing (PBR):

You can configure PBR by following these steps. Some of the steps may be omitted depending on your application for PBR. For this example, note the set condition options listed in step 3 (answers are bolded).

Reference: <http://www.ciscopress.com/articles/article.asp?p=102092&rl=1>

QUESTION NO: 342

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?



```

R1# debug eigrp packet hello
EIGRP Packets debugging is on
(HELLO)
R1#
Nov 20 08:07:33.131: EIGRP: Sending HELLO on Serial0
Nov 20 08:07:33.135: AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
Nov 20 08:07:35.327: EIGRP: Received HELLO on Serial0 nbr 192.168.1.1
Nov 20 08:07:35.331: AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/rely 0/0

R1# show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H Address Interface Hold Uptime SRTT RT0 Q Seq Type
(sec) (ms) Cnt Num
0 192.168.1.1 Se0 13 00:24:47 1 3000 0 1
  
```

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

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

Answer: B

QUESTION NO: 343

Which three of the following protocols will be forwarded to a host specified by the ip helper-address interface configuration command if the configuration has not been modified by the ip forward-protocol udp global configuration command? (Choose three.)

- A. BOOTP
- B. TFTP
- C. ARP
- D. DNS
- E. proxy-ARP
- F. FTP

Answer: A,B,D

QUESTION NO: 344

Refer to the exhibit. Which two statements are true about the partial configuration that is shown? (Choose two.)


```
WanRouter# show running-config
<output omitted>
ip dhcp excluded-address 10.0.0.1 10.0.0.5
!
!
ip dhcp pool Central
  network 10.0.0.0 255.255.255.0
  domain-name Central
  dns-server 10.0.0.2
  netbios-name-server 10.0.0.2
!
!
interface FastEthernet0/0
  ip address 10.0.0.1 255.255.255.0
  duplex auto
  speed auto
!
interface FastEthernet0/1
  ip address 10.10.0.1 255.255.255.0
  duplex auto
  speed auto
!
ActualTests
```

- A. The first DHCP client to connect to the FastEthernet 0/1 interface will receive the IP address 10.10.0.1.
- B. DHCP requests received on the FastEthernet 0/1 interface will be forwarded to 10.0.0.2.
- C. The first DHCP client to connect to the FastEthernet 0/0 interface will receive the IP address 10.0.0.1
- D. The first DHCP client to connect to the FastEthernet 0/0 interface will receive the IP address 10.0.0.6.
- E. Hosts connected to the FastEthernet0/1 interface will not receive DHCP replies from the router.

Answer: D,E

QUESTION NO: 345

A company has a Frame Relay WAN with one central-site router and 100 branch office routers. A partial mesh of PVCs exists: one PVC between the central site and each of the 100 branch routers. Which of the following could be true about the number of EIGRP neighborships?

- A. A partial mesh totaling 100: one between the central-site router and each of the 100 branches.
- B. A full mesh - $(101 * 100) / 2 = 5050$ -One neighborship between each pair of routers.

- C. 101-One between each router (including the central site) and its nearby PE router.
- D. None of the answers is correct.

Answer: A

Explanation:

EIGRP forms neighborships only when two routers can communicate directly over a data link. As a result, with Frame Relay, EIGRP neighborships occur only between routers on the ends of a PVC, so in this case, 100 neighborships exist.

QUESTION NO: 346

A network design shows that R1 has four different possible paths from itself to the Data Center subnets. Which one of the following commands is most likely to show you all the possible next-hop IP addresses for these four possible routes?

- A. show ip eigrp topology
- B. show ip eigrp topology all-links
- C. show ip route eigrp
- D. show ip route eigrp all-links
- E. show ip eigrp topology all-learned

Answer: B

Explanation:

Of the five options, the show ip route eigrp all-links and show ip eigrp topology all-learned are not valid commands. Both show ip eigrp topology and show ip route eigrp can show at most successor and feasible successor routes. However, show ip eigrp topology all-links shows also nonfeasible successor routes, making it more likely to show all possible neighbors.

QUESTION NO: 347

Router R1 lists 4 routes for subnet 10.1.1.0/24 in the output of the show ip eigrp topology all-links command. The variance 200 command is configured, but no other related commands are configured. Which of the following rules are true regarding R1's decision of what routes to add to the IP routing table? Note that RD refers to reported distance and FD to feasible distance.

- A. Adds all routes for which the metric is $\leq 200 * \text{the best metric among all routes}$
- B. Adds all routes because of the ridiculously high variance setting
- C. Adds all successor and feasible successor routes

D. Adds all successor and feasible successor routes for which the metric is $\leq 200 \times$ the best metric among all routes

Answer: D

Explanation:

EIGRP considers only successor and feasible successor routes. Each of those routes must have metrics such that variance \times metric is less than the best route's metric; the best route's metric is called the feasible distance (FD).

QUESTION NO: 348

Router R1 has been configured for EIGRP. The configuration also includes an ACL with one line-access-list 1 permit 10.100.32.0 0.0.15.255-and the EIGRP configuration includes the distribute-list 1 in command. Which of the following routes could not be displayed in the output of the show ip eigrp topology command as a result?

- A. 10.10.32.0/19
- B. 10.10.44.0/22
- C. 10.10.40.96/27
- D. 10.10.48.0/23
- E. 10.10.60.0/30

Answer: D,E

Explanation:

The two listed commands correctly configure EIGRP route filtering such that prefixes matched by the ACL's permit clause will be allowed. All other prefixes will be filtered due to the implied deny all at the end of the ACL. The ACL permits numbers in the range 10.10.32.0-10.10.47.255, which leaves 10.10.48.0 and 10.10.60.0 unmatched by the permit clause.

QUESTION NO: 349

The command output that follows was gathered from router R1. If correctly referenced by an EIGRP distribution list that filters outbound Updates, which of the following statements is true about the filtering of various prefixes by this Prefix list?

```
R1#sh ip prefix-list
ip prefix-list question: 4 entries
seq 5 deny 10.1.2.0/24 ge 25 le 27
seq 15 deny 10.2.0.0/16 ge 30 le 30
```

seq 20 permit 0.0.0.0/0

- A. Prefix 10.1.2.0/24 will be filtered due to clause 5.
- B. Prefix 10.1.2.224/26 will be filtered due to clause 5.
- C. Prefix 10.2.2.4/30 will be filtered due to clause 15.
- D. Prefix 10.0.0.0/8 will be permitted.
- E. Prefix 0.0.0.0/0 will be permitted.

Answer: B,C,E

Explanation:

Sequence number 5 matches prefixes from 10.1.2.0-10.1.2.255, with prefix lengths between 25-27, and denies (filters) those prefixes. This results in answer A being incorrect, because the prefix length (/24) is not in the correct range. Clause 15 matches prefixes from 10.2.0.0-10.2.255.255, with prefix length exactly 30, matching answer C. Clause 20 matches only prefix 0.0.0.0 with length /0, so only a default route would match this entry. As a result, 10.0.0.0/8 does not match any of the three clauses.

QUESTION NO: 350

R1 has correctly configured EIGRP to filter routes using a route map named question. The configuration that follows shows the entire route map and related configuration. Which of the following is true regarding the filtering action on prefix 10.10.10.0/24 in this case?

```
route-map question deny 10
match ip address 1
route-map question permit 20
match ip address prefix-list fred
!
access-list 1 deny 10.10.10.0 0.0.0.255
ip prefix-list fred permit 10.10.10.0/23 le 25
```

- A. It will be filtered due to the deny action in route map clause 10.
- B. It will be allowed because of the double negative (two deny references) in clause 10.
- C. It will be permitted due to matching clause 20's reference to prefix-list fred.
- D. It will be filtered due to matching the implied deny all route map clause at the end of the route map.

Answer: C

Explanation:

When used for route filtering, the route map action (permit or deny) defines the filtering action, and any referenced match commands' permit or deny action just defines

whether the prefix is matched. By not matching ACL 1 with a permit action, EIGRP does not consider a match to have occurred with clause 10, so it moves to clause 20. The prefix list referenced in clause 20 has a permit action, matching prefixes from 10.10.10.0-10.10.11.255, with prefix lengths from 23-25. Both criteria match the prefix in question, making answer C correct.

QUESTION NO: 351

An engineer has typed four different single-line prefix lists in a word processor. The four answers show the four different single-line prefix lists. The engineer then does a copy/paste of the configuration into a router. Which of the lists could match a subnet whose prefix length is 27?

- A. ip prefix-list fred permit 10.0.0.0/24 ge 16 le 28
- B. ip prefix-list barney permit 10.0.0.0/24 le 28
- C. ip prefix-list wilma permit 10.0.0.0/24 ge 25
- D. ip prefix-list betty permit 10.0.0.0/24 ge 28

Answer: B,C

Explanation:

Answer A is invalid-the ge value must be larger than /24 in this case, so the command is rejected. Answer B implies a prefix length range from 24-28, inclusive. Answer C implies a range from 25-32 inclusive, because no le parameter exists to limit the prefix length lower than the full length of an IPv4 subnet mask. The same logic applies with answer D, but with a range from 28-32, so this final list could not match prefix lengths of /27.

QUESTION NO: 352

An engineer plans to configure summary routes with the ip summary-address eigrp asn prefix mask command. Which of the following, when added to such a command, would create a summary that includes all four of the following subnets: 10.1.100.0/25, 10.1.101.96/27, 10.1.101.224/28, and 10.1.100.128.25?

- A. 10.1.0.0 255.255.192.0
- B. 10.1.64.0 255.255.192.0
- C. 10.1.100.0 255.255.255.0
- D. 10.1.98.0 255.255.252.0

Answer: B

Explanation:

The 10.1.0.0/18 implies a range from 10.1.0.0-10.1.63.255, which includes none of the four subnets. 10.1.64.0/18 implies a range from 10.1.64.0-10.1.127.255, which includes all subnets. 10.1.100.0/24 implies range 10.1.100.0-10.1.100.255, which leaves out two of the subnets. Finally, 10.1.98.0/22 does not actually represent a summary-instead, 10.1.96.0/22 represents range 10.1.96.0-10.1.99.255, with 10.1.98.0 as listed in answer D being an IP address in that range. As such, IOS would actually accept the command, and change the parameter from 10.1.98.0 to 10.1.96.0, and would not include the four listed subnets.

QUESTION NO: 353

R1 has 5 working interfaces, with EIGRP neighbors existing off each interface. R1 has routes for subnets 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24, with EIGRP integer metrics of roughly 1 million, 2 million, and 3 million, respectively. An engineer then adds the ip summary-address eigrp 1 10.1.0.0 255.255.0.0 command to interface Fa0/0. Which of the following is true?

- A. R1 loses and then reestablishes neighborships with all neighbors.
- B. R1 no longer advertises 10.1.1.0/24 to neighbors connected to Fa0/0.
- C. R1 advertises a 10.1.0.0/16 route out Fa0/0, with metric of around 3 million (largest metric of component subnets).
- D. R1 advertises a 10.1.0.0/16 route out Fa0/0, with metric of around 2 million (median metric of component subnets).

Answer: B**Explanation:**

The ip summary-address command does reset neighborships, but only on the interface under which it is configured. After those neighborships come up, R1 will advertise the summary route, but none of the subordinate routes inside that summary. The summary route will use a metric equal to the metric of the lowest metric subordinate route, approximately 1,000,000 in this case.

QUESTION NO: 354

In a lab, R1 connects to R2, which connects to R3. R1 and R2 each have several working interfaces, all assigned addresses in class A network 10.0.0.0. Router R3 has some working interfaces in class A network 10.0.0.0, and others in class B network 172.16.0.0. The engineer experiments with the auto-summary command on R2 and R3, enabling and disabling the command in various combinations. Which of the following combinations will result in R1 seeing a route for 172.16.0.0/16, instead of the individual subnets of class B network 172.16.0.0?

- A. auto-summary on R2 and no auto-summary on R3

- B. auto-summary on R2 and auto-summary on R3
- C. no auto-summary on R2 and no auto-summary on R3
- D. no auto-summary on R2 and auto-summary on R3

Answer: B,D

Explanation:

R2 has interfaces only in class A network 10.0.0.0, so the auto-summary setting has no effect. R3 has interfaces in both class A network 10.0.0.0 and class B network 172.16.0.0, so auto-summary causes R3 to summarize all subnets of 172.16.0.0/16 as a summary route when advertising to R2.

QUESTION NO: 355

Router R1 exists in an Enterprise that uses EIGRP as its routing protocol. The show ip route command output on router R1 lists the following phrase: "Gateway of last resort is 1.1.1.1 to network 2.0.0.0". Which of the following is most likely to have caused this output to occur on R1?

- A. R1 has been configured with an ip default-network 2.0.0.0 command.
- B. R1 has been configured with an ip route 0.0.0.0 0.0.0.0 1.1.1.1 command.
- C. R1 has been configured with an ip route 2.0.0.0 255.0.0.0 1.1.1.1 command.
- D. Another router has been configured with an ip default-network 2.0.0.0 command.
- E. Another router has been configured with an ip route 2.0.0.0 255.0.0.0 1.1.1.1 command.

Answer: D

Explanation:

The phrase quoted in the question means that R1 is using its route for class A network 2.0.0.0 to decide where to send packets by default. R1's route for network 2.0.0.0 must have 1.1.1.1 as its next-hop router. This phrase occurs when EIGRP has learned a route for class A network 2.0.0.0 that has been flagged as a candidate default route by another router. The router flagging a route as a candidate default route, using the ip default-network command, does not actually use the route as its default route.

QUESTION NO: 356

Enterprise Router R1 connects an Enterprise to the Internet. R1 needs to create and advertise a default route into the Enterprise using EIGRP. The engineer creating the implementation plan has chosen to base this default route on the ip route command, rather than using ip default-network. Which of the following is not a useful step with this style of default route configuration?

- A. Create the default route on R1 using the ip route 0.0.0.0 0.0.0.0 outgoinginterface command.

- B. Redistribute the statically configured default route.
- C. Disable auto-summary.
- D. Configure the network 0.0.0.0 command.
- E. Ensure R1 has no manually configured summary routes using the ip summary-address eigrp command.

Answer: C,E

Explanation:

With the suggested configuration style, the static route must first be configured statically, as shown in answer A. Then, either this route must be redistributed as a static route into EIGRP (answer B), or pulled into EIGRP by virtue of the network 0.0.0.0 EIGRP subcommand answer D). The other two options have no effect on default route creation and advertisement.

QUESTION NO: 357

A router has been configured with the commands `router ospf 9`, `network 172.16.1.0 0.0.0.255 area 8`, and `network 172.16.0.0 0.0.255.255 area 9`, in that order. No other OSPF-related commands have been configured. The answers list the IP addresses that could be assigned to this router's Fa0/0 interface. Which answers list an IP address/prefix length that would cause the router to put Fa0/0 into area 9? (Choose two.)

- A. 172.16.0.1/23
- B. 172.16.1.1/26
- C. 172.16.1.1/24
- D. 172.16.0.255/23
- E. None of the other answers is correct.

Answer: A,D

Explanation:

The wildcard mask is used for matching the prefix only, and not the prefix length. As such, `172.16.1.0 0.0.0.255` matches all addresses that begin with 172.16.1, and `172.16.0.0 0.0.255.255` matches all addresses that begin 172.16. Also, OSPF reviews the network command with the most specific wildcard masks (wildcard masks with the most binary 0's) first, so an interface IP address beginning with 172.16.1 matches the command that references area 8.

QUESTION NO: 358

Which of the following is true about an OSPF area border router (ABR)?

- A. The ABR must have multiple interfaces connected to the backbone area.
- B. An ABR is a router with two interfaces, each connected to a different nonbackbone area.
- C. The only requirement to be considered an ABR is at least one interface connected to the backbone area.
- D. An ABR must have at least one interface in the backbone area plus at least one other interface in a nonbackbone area.

Answer: D

Explanation:

ABRs, by definition, connect the backbone area to one or more nonbackbone areas. To perform this function, a router must have at least one interface assigned to the backbone area, and at least one interface assigned to a nonbackbone area.

QUESTION NO: 359

Which of the following can either directly or indirectly identify all the interfaces for which both 1) OSPF has been enabled and 2) OSPF is not passive? (Choose two.)

- A. show ip ospf database
- B. show ip ospf interface brief
- C. show ip protocols
- D. show ip route ospf
- E. show ip ospf neighbors

Answer: B,C

Explanation:

First, for the two correct answers: show ip ospf interface brief explicitly lists all OSPF-enabled interfaces that are not passive. show ip protocols lists either the details of the configured network commands, or if configured using the ip ospf area command, it lists the interfaces on which OSPF is enabled. This command also lists the passive interfaces, so armed with interface IP address information, the list of OSPF-enabled nonpassive interfaces could be derived. Of the three wrong answers, show ip ospf database does not list enough detail to show the OSPF-enabled interfaces. show ip route ospf lists only routes learned with OSPF, so if no routes use a particular OSPF-enabled interface as an outgoing interface, this command would not indirectly identify the interface. Finally, an interface may be OSPF-enabled but with no neighbors reachable on the interface, so the show ip ospf neighbor command may not identify all OSPF-enabled interfaces.

QUESTION NO: 360

Router R1 directly connects to subnet 10.1.1.0/24 with its Fa0/0 interface. R1 can ping four other working OSPF routers in that subnet. R1 is neither the designated router (DR) nor backup DR (BDR). OSPF is working correctly on all five routers. Which of the following is true on R1? (Choose two.)

- A. The show ip ospf neighbors command lists two neighbors off Fa0/0.
- B. The show ip ospf neighbors command lists four neighbors off Fa0/0.
- C. The show ip ospf neighbors command lists two neighbors off Fa0/0 in the FULL state.
- D. The show ip ospf neighbors command lists two neighbors off Fa0/0 in the DISCO state.

Answer: B,C

Explanation:

On a LAN, the non-DR routers form fully adjacent neighborships with only the DR and BDR, giving R1 two neighbors in the FULL state. The other two neighbors settle into the 2WAY state.

QUESTION NO: 361

Routers R1 and R2 are OSPF neighbors using their Fa0/0 interfaces, respectively, using default settings for all timers. An engineer adds the ip ospf hello-interval 6 command to R1's Fa0/0 configuration. Which of the following is true regarding the results from this change? (Choose 2)

- A. The show ip ospf neighbor command on R1 lists the revised Hello timer.
- B. The show ip ospf interface brief command on R1 lists the revised Hello timer.
- C. The R1-R2 neighborship fails due to Hello timer mismatch.
- D. The show ip ospf interface command on R1 lists the revised Hello timer.

Answer: C,D

Explanation:

The show ip ospf interface command displays a router's OSPF Hello Interval setting for each enabled interface. The other listed commands do not display the timer. Also, OSPF routers do need to have matching Hello timers to become neighbors, so the neighborship would fail.

QUESTION NO: 362

Routers R1 and R2, OSPF neighbors in area 0 over their Fa0/0 interfaces (respectively), currently both successfully use OSPF MD5 authentication. The OSPF configuration includes the area 0 authentication command under the router ospf 1 command. Which of the following commands must have been configured on R1's Fa0/0 interface? (Choose two.)

- A. ip ospf authentication null
- B. ip ospf authentication message-digest
- C. ip ospf authentication-key whatever-it-is
- D. ip ospf message-digest-key 1 md5 whatever-it-is
- E. ip ospf md5 1 key whatever-it-is

Answer: B,D

Explanation:

The area 0 authentication command tells R1 to use simple text password authentication on all interfaces in area 0 unless overridden by an interface subcommand. So, R1 must have configured the ip ospf authentication message-digest on its Fa0/0 interface, enabling MD5 authentication instead. The other correct answer is the command that correctly configures the MD5 authentication key. Of the two incorrect answers that list an authentication key, ip ospf authentication-key defines the clear-text key, and the other is not a valid IOS command.

QUESTION NO: 363

Which of the following settings do not prevent two potential OSPF neighbors from becoming neighbors?

- A. The interface used to connect to that neighbor being passive in the OSPF process
- B. Duplicate OSPF router IDs
- C. Mismatched dead timers
- D. IP addresses of 10.1.1.1/24 and 10.2.2.2/24
- E. Mismatched OSPF process IDs

Answer: E

Explanation:

For OSPF, Router IDs must be unique; the interfaces must not be passive; the dead timers must match; and the primary IP addresses must be in the same subnet, with the same subnet mask. However, the process IDs, found on the router ospf process-id command, do not have to match.

QUESTION NO: 364

A company has a Frame Relay WAN with one central-site router and 100 branch office routers. A partial mesh of PVCs exists: one PVC between the central site and each of the 100 branch routers. All routers use point-to-point subinterfaces and one subnet per PVC. Which of the following is true about OSPF in this design?

- A. The central site router has 100 fully adjacent neighborships with the 100 branches.
- B. The central site router has neighborships with all branch routers, but fully adjacent neighbors with only two branches.
- C. The central site router has a neighborship with the Frame Relay switch.
- D. None of the other answers is correct.

Answer: A

Explanation:

Frame Relay is a Layer 2 service and as such does not participate in customer routing protocols. Because the design uses a separate subnet per PVC, and one point-to-point subinterface per PVC/subnet, OSPF will use a point-to-point network type.

That means that the two routers on either end of a PVC will become neighbors, and become fully adjacent, meaning the central site router will have 100 fully adjacent neighborships.

QUESTION NO: 365

A network design shows area 1 with three internal routers, area 0 with four internal routers, and area 2 with five internal routers. Additionally, one ABR (ABR1) connects areas 0 and 1, plus a different ABR (ABR2) connects areas 0 and 2. How many Type 1 LSAs would be listed in ABR2's LSDB?

- A. 6
- B. 7
- C. 15
- D. 12
- E. None of the other answers is correct.

Answer: D

Explanation:

As an ABR connected to areas 0 and 2, ABR2 will have LSDB entries for both area 0 and area 2. In Area 0, ABR2 learns Type 1 LSAs from the four routers internal to area 0, plus ABR1, and plus 1 for the area 0 Type 1 LSA ABR2 creates for itself. In area 2, ABR2 learns 1 each for the five routers internal to area 2, plus the 1 Type 1 LSA ABR2 created for itself inside area 2. The total is 12.

QUESTION NO: 366

A network planning diagram shows a large internetwork with many routers. The configurations show that OSPF has been enabled on all interfaces, IP addresses correctly configured, and OSPF

workinG. For which of the following cases would you expect a router to create and flood a Type 2 LSA?

- A. When OSPF is enabled on a LAN interface, and the router is the only router connected to the subnet.
- B. When OSPF is enabled on a point-to-point serial link, and that router has both the higher router ID and higher interface IP address on the link.
- C. When OSPF is enabled on a Frame Relay point-to-point subinterface, has the lower RID and lower subinterface IP address, and otherwise uses default OSPF configuration on the interface.
- D. When OSPF is enabled on a working LAN interface on a router, and the router has been elected BDR.
- E. None of the other answers is correct.

Answer: E

Explanation:

OSPF creates a Type 2 LSA for a subnet when the router interface connected to the subnet calls for the election of a designated router (DR), and at least two routers have discovered each other and elected a DR. Then, the DR creates and floods the Type 2 LSA. IOS by default does not elect a DR on point-to-point topologies. It does on router LAN interfaces. One answer states that one router only exists in the subnet, so it does not actually find a second router and elect a DR. In the other case, a DR and BDR have been elected, but the router described in the answer is the BDR, not the DR. So, none of the other answers is correct.

QUESTION NO: 367

A verification plan shows a network diagram with branch office Routers B1 through B100, plus two ABRs, ABR1, and ABR2, all in area 100. The branches connect to the ABRs using Frame Relay point-to-point subinterfaces. The verification plan lists the output of the show ip ospf database summary 10.100.0.0 command on a router B1, one of the branches. Which of the following is true regarding the output that could be listed for this command?

- A. The output lists nothing unless 10.100.0.0 has been configured as a summary route using the area range command.
- B. If 10.100.0.0 is a subnet in area 0, the output lists one Type 3 LSA, specifically the LSA with the lower metric when comparing ABR1's and ABR2's LSA for 10.100.0.0.
- C. If 10.100.0.0 is a subnet in area 0, the output lists two Type 3 LSAs, one each created by BR1 and ABR2.
- D. None, because the Type 3 LSAs would exist only in the ABR's LSDBs.

Answer: C

Explanation:

Each ABR, by definition, creates a single Type 3 LSA to represent a subnet known in one area to be advertised into another area. Assuming 10.100.0.0 is a subnet in area 0, both ABR1 and ABR2 would advertise a Type 3 LSA into area 100. The show ip ospf database summary command specifically lists type 3 network summary LSAs.

QUESTION NO: 368

Which of the following OSPF messages contains entire complete LSAs used during the database exchange process?

- A. LSR
- B. LSAck
- C. LSU
- D. DD
- E. Hello

Answer: C**Explanation:**

The Database Description (DD) packet lists a short LSA header but not the entire LSA. The Link State Request (LSR) packet asks the neighbors for a copy of an LSA. The Link State Update (LSU) holds the LSAs. LSAck simply acknowledges received LSAs, and Hello is used for neighbor discovery and neighbor state maintenance.

QUESTION NO: 369

Routers R1, R2, R3, and R4 connect to the same 10.10.10.0/24 LAN-based subnet. OSPF is fully working in the subnet. Later, R5, whose OSPF priority is higher than the other four routers, joins the subnet. Which of the following are true about the OSPF database exchange process over this subnet at this point? (Choose two.)

- A. R5 will send its DD, LSR, and LSU packets to the 224.0.0.5 all-DR-routers multicast address.
- B. R5 will send its DD, LSR, and LSU packets to the 224.0.0.6 all-DR-routers multicast address.
- C. The DR will inform R5 about LSAs by sending its DD, LSR, and LSU packets to the 224.0.0.6 all-SPF-routers multicast address.
- D. The DR will inform R5 about LSAs by sending its DD, LSR, and LSU packets to the 224.0.0.5 all-SPF-routers multicast address.

Answer: B,D**Explanation:**

Because the subnet was stable before R5 arrived, the other routers will have elected a DR and BDR. OSPF does not preemptively elect a new DR nor BDR, so R5 will be neither (DROther). As a result, R5's messages to the DR will be sent to the 224.0.0.6 all-DR-routers multicast address, and the DR's messages directed to R5 will be sent to the 224.0.0.5 all-SPF-router address.

QUESTION NO: 370

R1 is internal to area 1, and R2 is internal to area 2. Subnet 10.1.1.0/24 exists in area 2 as a connected subnet off R2. ABR ABR1 connects area 1 to backbone area 0, and ABR2 connects area 0 to area 2. Which of the following LSAs must R1 use when calculating R1's best route for 10.1.1.0/24?

- A. R2's Type 1 LSA
- B. Subnet 10.1.1.0/24's Type 2 LSA
- C. ABR1's Type 1 LSA in area 0
- D. Subnet 10.1.1.0/24's Type 3 LSA in Area 0
- E. Subnet 10.1.1.0/24's Type 3 LSA in Area 1

Answer: E

Explanation:

R1, internal to area 1, can use LSAs only in the area 1 LSDB. R2's Type 1 LSA exists only in area 2's LSDB. The Type 2 LSA for subnet 10.1.1.0/24, if one exists, also only exists in area 2's LSDB. R1 will use ABR1's Type 1 LSA in area 1 to calculate the possible intra-area routes inside area 1, but R1 will use ABR1's Type 1 LSA in area 1. Finally, the Type 3 LSA, created for 10.1.1.0/24, and flooded into area 1, is also needed to calculate the metric.

QUESTION NO: 371

Which of the following LSA types describes topology information that, when changed, requires a router in the same area to perform an SPF calculation? (Choose two.)

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5
- F. 7

Answer: A,B

Explanation:

OSPF builds the SPF tree based on the topology information Type 1 and Type 2 LSAs. Changes therefore require another SPF run. Changes to the other LSA types do not require an SPF calculation.

QUESTION NO: 372

The following output was taken from Router R3. A scan of R3's configuration shows that no bandwidth commands have been configured in this router. Which of the following answers lists configuration settings could be a part of a configuration that results in the following output? (Choose two.)

```
R3#show ip ospf interface brief
Interface PID Area IP Address/Mask Cost State Nbrs F/C
Se0/0/0.2 3 34 10.10.23.3/29 647 P2P 1/1
Se0/0/0.1 3 34 10.10.13.3/29 1000 P2P 1/1
Fa0/0 3 34 10.10.34.3/24 20 BDR 1/1
```

- A. An auto-cost reference-bandwidth 1000 command in router ospf mode
- B. An auto-cost reference-bandwidth 2000 command in router ospf mode
- C. An ip ospf cost 1000 interface S0/0/0.1 command in router ospf mode
- D. An auto-cost reference-bandwidth 64700 command in router ospf mode

Answer: A,B

Explanation:

Because none of the interfaces have a bandwidth command configured, the only commands that can influence the OSPF cost are the auto-cost reference-bandwidth router subcommand and the ip ospf cost interface subcommand. To give the output shown in the question, either the interface cost could be set directly on all three interfaces listed. Alternatively, the reference-bandwidth could be set to cause one of the interface costs to be as shown in the output, with the other two interfaces having their costs set directly.

For the wrong answers, the ip ospf cost interface s0/0/0.1 router subcommand does not exist-instead, it is an interface subcommand. An auto-cost of 64700, used as the numerator in the ref-bw/bandwidth cost calculation, does not result in any of the three listed interface costs.

For the two correct answers, with a default bandwidth of 1544 (Kbps) on the serial subinterfaces, a reference bandwidth of 1000 (Mbps) implies the math $1,000,000 / 1544$, for an Interface cost of 647. With a default bandwidth of 100,000 Kbps (100 Mbps) on Fa0/0, a reference bandwidth of 2000 (MBps) implies math

of 2,000 / 100 = 20.

QUESTION NO: 373

Which of the following LSA types describe information related to topology or subnets useful for calculating routes for subnets inside the OSPF domain? (Choose three.)

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5
- F. 7

Answer: A,B,C

Explanation:

OSPF uses Types 1, 2, and 3 for calculating routes internal to the OSPF domain. OSPF uses types 4, 5, and 7 for external routes redistributed into the OSPF domain.

QUESTION NO: 374

Router B1, an internal router in area 1, displays the following output. The only two ABRs connected to area 1 are performing Type 3 LSA filtering. Which of the following answers is true based on the information in the output from B1?

```
R1# show ip route 10.1.0.0 255.255.0.0 longer-prefixes
! Legend lines omitted for brevity
10.0.0.0/8 is variably subnetted, 17 subnets, 3 masks
O 10.1.2.0/24 [110/658] via 10.10.13.1, 00:00:32, Serial0/0/0.1
O IA 10.1.1.0/24 [110/658] via 10.10.23.2, 00:41:39, Serial0/0/0.2
O IA 10.1.3.0/24 [110/658] via 10.10.23.2, 00:41:39, Serial0/0/0.2
```

- A. A Type 3 LSA for 10.2.2.0/24 was filtered by both ABRs.
- B. A Type 3 LSA for 10.1.2.0/24 was not filtered by both ABRs.
- C. A Type 3 LSA for 10.1.3.0/24 was not filtered by at least one ABR.
- D. A Type 3 LSA for 10.1.1.0/24 filtered by both ABRs.

Answer: C

Explanation:

The output lists all R1's routes for subnets within the range of 10.1.0.0-10.1.255.255, whose prefix lengths are longer than /16. One answer lists subnet

10.2.2.0/24, which is not in this range, so the output cannot be used to confirm nor deny whether the subnet was filtered. R1's route for 10.1.2.0/24 is an intra-area route by virtue of not listing an IA code by the route; Type 3 LSA filtering only filters Type 3 LSAs, which routers use to calculate interarea routes, so the output tells us nothing about any filtering of 10.1.2.0/24. The output shows a single interarea route for 10.1.3.0/24, so at least one ABR has flooded a Type 3 LSA for this route. Additionally, the output confirms that at least one ABR flooded a type 3 LSA for 10.1.3.0/24, or the output would not show an IA route for 10.1.3.0/24. So, the type 3 LSA for 10.1.3.0/24 was not filtered by both ABRs.

QUESTION NO: 375

The following command output was gathered from Router R1, an ABR between areas 0 (backbone) and area 1. In this internetwork, area 0 contains all the subnets of class A network 10.0.0.0. R1's OSPF process has a distribute-list prefix question in command configured. Assuming the subnets listed in the answers actually exist in area 0, which of the following occurs on router R1?

```
R1#sh ip prefix-list
ip prefix-list question: 4 entries
seq 5 deny 10.1.2.0/24 ge 25 le 27
seq 15 deny 10.2.0.0/16 ge 30 le 30
seq 20 permit 0.0.0.0/0 le 32
```

- A. R1 will not create/flood a type 3 LSA for subnet 10.1.2.0/26 into area 1.
- B. R1 will not create/flood a Type 3 LSA for subnet 10.1.2.0/24 into area 1.
- C. R1 will not have an OSPF route for subnet 10.1.2.0/26 in its IP routing table.
- D. R1 will not have an OSPF route for subnet 10.1.2.0/24 in its IP routing table.

Answer: C**Explanation:**

When referenced from a distribute list, OSPF filters routes from being added to that router's IP routing table but has no impact on the flow of LSAs. As such, neither A nor B is correct. An OSPF distribute-list command does attempt to filter routes from being added to the IP routing table by OSPF, so the two answers that mention the IP routing table might be correct. Sequence number 5 matches prefixes from 10.1.2.0-10.1.2.255, with prefix lengths between 25-27, and denies (filters) those prefixes. So, the prefix list will match 10.1.2.0/26 with the first line, with a deny action. The 10.1.2.0/24 subnet does not match the first line of the prefix list, but it does match the third line, the match all line, with a permit action. Because 10.1.2.0/26 is matched by a deny clause, this route is indeed filtered, so it is not added to R1's IP routing table. 10.1.2.0/24, matched with a permit clause, is allowed and would be in the IP routing table.

QUESTION NO: 376

Use the same scenario as the previous question, with one change. Instead of the distribute-list prefix question in command configured on R1, R1's OSPF process has an area 1 filter-list prefix question in command configured. Again assuming that the subnets listed in the answers actually exist in area 0, which of the following occurs on router R1?

```
R1#sh ip prefix-list
ip prefix-list question: 4 entries
seq 5 deny 10.1.2.0/24 ge 25 le 27
seq 15 deny 10.2.0.0/16 ge 30 le 30
seq 20 permit 0.0.0.0/0 le 32
```

- A. R1 will not create/flood a type 3 LSA for subnet 10.1.2.0/26 into area 1.
- B. R1 will not create/flood a Type 3 LSA for subnet 10.1.2.0/24 into area 1.
- C. R1 will not have an OSPF route for subnet 10.1.2.0/26 in its IP routing table.
- D. R1 will not have an OSPF route for subnet 10.1.2.0/24 in its IP routing table.

Answer: A

Explanation:

When referenced from an area filter-list command, OSPF filters Type 3 LSAs created on that router, preventing them from being flooded into area 1 (per the configuration command). As an ABR, R1 would calculate intra-area routes to these area 0 subnets, so this filtering will have no effect on R1's routes. Sequence number 5 matches prefixes from 10.1.2.0-10.1.2.255, with prefix lengths between 25-27, and denies (filters) those prefixes. So, the prefix list will match 10.1.2.0/26 with the first line, with a deny action. The 10.1.2.0/24 subnet does not match the first line of the prefix list because the prefix length does not match, but it does match the third line, the match all line, with a permit action. By matching subnet 10.1.2.0/26 with a deny action, the filter-list does prevent R1 from flooding a Type 3 LSA for that subnet. By matching 10.1.2.0/24 with a permit action, R1 does not filter the Type 3 LSA for that subnet.

QUESTION NO: 377

R1, an ABR between backbone area 0 and area 1, has intra-area routes in area 0 for 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24. These routes have metrics of 21, 22, and 23, respectively. An engineer then adds the area 0 range 10.1.0.0 255.255.0.0 command under the OSPF process of R1. Which of the following is true? (Choose two.)

- A. R1 loses and then re-establishes neighborships with all neighbors.
- B. R1 no longer advertises 10.1.1.0/24 to neighbors into area 1.
- C. R1 advertises a 10.1.0.0/16 route into area 1 with a metric of 23 (largest metric).
- D. R1 advertises a 10.1.0.0/16 route into area 1 with metric of 21 (lowest metric).

Answer: B,D

Explanation:

The area range command does not cause a failure in neighborships. Because at least one intra-area subordinate subnet of 10.1.0.0/16 exists in R1, R1 both creates a summary route for 10.1.0.0/16 and stops advertising LSAs for the (three) subordinate subnets. By default, the metric of the summary is the metric of the lowest-metric component subnet.

QUESTION NO: 378

The following output exists on Router R1, a router internal to area 1. What can you determine as true from the output of the show ip ospf database summary command?

Routing Bit Set on this LSA

LS age: 124

Options: (No TOS-capability, DC, Upward)

LS Type: Summary Links (Network)

Link State ID: 10.1.0.0 (summary Network Number)

Advertising Router: 1.1.1.1

LS Seq Number: 80000001

Checksum: 0x878F

Length: 28

Network Mask: /22

TOS: 0 Metric: 11

- A. The LSA was created by an ABR due to an area range command.
- B. The LSA was created by an ASBR due to a summary-address command.
- C. If created by an area range command, the best metric for a subordinate subnet on that ABR must have been 11.
- D. None of the other answers is correct.

Answer: D

Explanation:

The show ip ospf database summary command lists only Type 3 LSAs. The summary - address command creates Type 5 LSAs on ASBRs, ruling out one answer.

The output does not specify whether the LSA was created as a summary route; all references to the word "summary" refer to Type 3 Summary LSAs. If created by an area range command, the metric defaults to be the best metric of all subordinate subnets, but it may also be explicitly set,

ruling out another of the possible answers. In short, this LSA may represent a route summarized by the area range command, but that fact cannot be proved or disproved by the output as shown..

QUESTION NO: 379

Router R1, an ASBR connected to the Internet and to backbone area 0, has been configured with a default-information originate command. Which of the following is true about the effects of this configuration command?

- A. R1 will always create and flood a default route into the OSPF domain.
- B. R1 will create and flood an LSA for prefix/length 0.0.0.0/0 into the OSPF domain if R1's IP routing table has a route to 0.0.0.0/0.
- C. R1 will set a flag on the LSA for the subnet between itself and one of the ISPs, noting this subnet as a default network, regardless of whether R1 has a default route.
- D. R1 will set a flag on the LSA for the subnet between itself and one of the ISPs, noting this subnet as a default network, but only if R1 has a route to 0.0.0.0/0.

Answer: B

Explanation:

Without the always parameter, the default-information originate command generates an LSA for a default route, with prefix 0.0.0.0/0, but only if its own IP routing table has a route for 0.0.0.0/0. It does not flag another LSA as being used as a candidate default route.

QUESTION NO: 380

Which of the following is true about routers internal to a totally NSSA area?
(Choose two.)

- A. Routers cannot redistribute external routes into the area.
- B. Routers should have zero Type 3 LSAs in their LSDBs.
- C. Routers should have zero Type 5 LSAs in their LSDBs.
- D. Routers should learn default routes from the ABRs attached to the area.

Answer: C,D

Explanation:

Both types of NSSA stubby areas allow the redistribution of external routes into the area, but these routes are advertised as Type 7 LSAs. As a totally NSSA area, the ABR should flood no Type 5 LSAs into the area and flood no Type 3 LSAs into the area, except for the Type 3 LSAs used to advertise the default route into the area. As such, a router internal to a totally stubby area should see zero Type 5 LSAs, and a small number of Type 3 LSAs for the default route(s) advertised by the ABR(s).

QUESTION NO: 381

ABR R1 has been configured with an area 1 stub no-summary command. Which stubby area type is area 1?

- A. Stub
- B. Totally stubby
- C. NSSA
- D. Totally NSSA

Answer: B

Explanation:

The stub keyword means either a stub area or totally stubby. The no-summary command means the area is totally stubby.

QUESTION NO: 382

Which of the following answers can be verified as true based on the following command output from Router R1?

```
R1#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 4.4.4.4 is up
Run as demand circuit
DoNotAge LSA allowed.
Transit area 1, via interface FastEthernet0/1, Cost of using 3
```

- A. R1 is configured with an area 0 virtual-link 4.4.4.4 cost 3 command.
- B. The ping 4.4.4.4 command on R1 must currently be successful.
- C. R1's Fa0/0 OSPF cost is 3.
- D. 4.4.4.4 is known to R1 based on a Type 1 LSA in area 1.

Answer: D

Explanation:

The answer with area 0 virtual-link 4.4.4.4 cost 3 is incorrect because the show command output lists a transit area of 1, but the answer's area parameter refers to area 0 as the transit area. (There is also no cost parameter on the area virtual-link command.) The RID of the router on the other end of the virtual link, 4.4.4.4 per the show command output, does not have to be pingable for the virtual link to work. The cost of the virtual link is 3, but that cost is calculated as the cost to reach the other router through the transit area, so the command output listed with the question can not be used to

predict Fa0/0's OSPF interface cost alone. However, because the output lists area 1 as the transit area, and because the neighbor RID is listed as 4.4.4.4, R1 will use the area 1 LSDB entries to calculate the cost to reach 4.4.4, a process that will include the area 1 Type 1 LSA for RID 4.4.4.4.

QUESTION NO: 383

Several links have been broken so that for the next day or two, what was formerly a contiguous area 0 has been broken into two parts. However, both parts of area 0 have working links into area 1 using routers with RID 1.1.1.1 and 2.2.2.2. Which answers list the command on the router with RID 1.1.1.1 to create a virtual link to help solve this temporary problem?

- A. area 0 virtual-link 2.2.2.2
- B. area 1 virtual-link 2.2.2.2
- C. area 0 source-rid 1.1.1.1 dest-rid 2.2.2.2
- D. virtual-link transit-area 1 RID 2.2.2.2

Answer: B**Explanation:**

The area virtual-link command defines the virtual link, with the transit area-the area through which the virtual link passes-listed as the first parameter. The other parameter is the RID of the other router. Two of the wrong answers are not IOS commands.

QUESTION NO: 384

Router R1 connects to a Frame Relay cloud using a multipoint subinterface, with ten PVCs associated with the subinterface. What command would make this router not use a DR and require static OSPF neighbor definition?

- A. ip ospf network broadcast
- B. ip ospf network non-broadcast
- C. ip ospf network point-to-multipoint
- D. ip ospf network point-to-multipoint non-broadcast

Answer: D**Explanation:**

Of the four types listed, only point-to-multipoint nonbroadcast does not use a DR but does require the static definition of neighbors.

QUESTION NO: 385

Router R1 connects to a Frame Relay cloud using a multipoint subinterface, with ten PVCs associated with the subinterface. What command would make this router not use a DR, and dynamically discover OSPF neighbors?

- A. ip ospf network broadcast
- B. ip ospf network non-broadcast
- C. ip ospf network point-to-multipoint
- D. ip ospf network point-to-multipoint non-broadcast

Answer: C

Explanation:

Of the four types listed, only point-to-multipoint does not use a DR and dynamically discovers neighbors.

QUESTION NO: 386

Ten routers, R1 through R10, connect in a partial mesh over Frame Relay. For the mesh, R1 and R2 have PVCs connected to all other routers, but Routers R3 through R10 act as branch routers, with only two PVCs—one to R1 and one to R2. The routers use IP subnet 10.1.1.0/24, with addresses 10.1.1.1, 10.1.1.2, and so on, through 10.1.1.10, respectively. The routers all use Inverse ARP to learn Frame Relay mapping information. All routers use a multipoint subinterface with network type point-to-multipoint nonbroadcast. A co-worker's implementation plan lists lots of configuration commands related to this design. The design states that all hosts should be able to ping all other hosts. Which commands are required for proper functioning of OSPF in this case? (Choose two.)

- A. frame-relay map commands on R3-R10 referencing the other routers in this group.
- B. Nine OSPF neighbor commands on each router.
- C. Nine OSPF neighbor commands on R1 and R2, with only two such commands on R3-R10.
- D. R1 and R2 with ip ospf priority 1 commands to ensure they become DR and BDR.
- E. R3-R10 with ip ospf priority 0 commands to ensure they do not become DR or BDR.

Answer: A,C

Explanation:

For routers to use their OSPF routes in a multipoint design, each router needs mapping to each other router in the Frame Relay subnet. In this case, R3-R10 all need frame-relay map commands to define the mapping to other routers with which they do not have a PVC. This network type requires static definition of neighbors, but the neighbor relationships match the PVC topology, so R3-R10 need only two neighbor commands. This OSPF network type does not use a DR, so the ip ospf

priority commands have no effect and would be unnecessary.

QUESTION NO: 387

Which of the following answers is the least likely reason for an engineer to choose to use route redistribution?

- A. To exchange routes between merged companies
- B. To give separate control over routing to different parts of one company
- C. To support multiple router vendors
- D. To knit together an OSPF area if the area becomes discontinuous

Answer: D

Explanation:

The three incorrect answers list typical reasons for using route redistribution. The correct answer- the least likely reason among the answers for using route redistribution- lists a problem for which an OSPF virtual link is often used. Route redistribution could be attempted to solve a problem with a discontinuous OSPF area, but the redistribution completely changes the LSAs that would have otherwise been known and could have negative impacts on route summaries and cause routing loops, and have other problems as well.

QUESTION NO: 388

For a router to successfully redistribute routes between OSPF and EIGRP, which of the following are true? (Choose two.)

- A. The router must have one routing protocol configured, but configuration for both routing protocols is not necessary.
- B. The router must have at least one working link connected to each routing domain.
- C. The redistribute command must be configured under EIGRP to send the routes to OSPF.
- D. The redistribute command should be configured under OSPF to take routes from EIGRP into OSPF.

Answer: B,D

Explanation:

For a router to redistribute routes between two routing protocols, the router must have both routing protocols configured, have a working link into each routing domain, and configure redistribute commands under each routing process. The redistribute command, issued in routing protocol configuration mode, pulls routes into that routing process, from another routing process as referenced on the redistribute command.

QUESTION NO: 389

Process EIGRP 1 is redistributing routes from process OSPF 2. Which two of the following methods may be used to set the metrics of the redistributed routes?

(Choose 2)

- A. Let the metrics default.
- B. Set the metric components using the redistribute command's metric keyword.
- C. Set the metric components using the default-metric router subcommand.
- D. Set the integer (composite) metric using the redistribute command's metric keyword.

Answer: B,C

Explanation:

Because the metrics come from a different routing protocol than EIGRP, the metric must be set. The metric must be set with five components; EIGRP will then use those components as it would for an internal route. The metric components may be set as listed in the two correct answers, plus using a route-map as referenced by the redistribute command.

QUESTION NO: 390

Examine the following excerpt from the show ip eigrp topology 10.2.2.0/24 command on router R1. Which answer can be verified as definitely true based on this output?

External data:

Originating router is 10.1.1.1

AS number of route is 1

External protocol is OSPF, external metric is 64

Administrator tag is 0 (0x00000000)

- A. R1 is the router that redistributed the route.
- B. R1's metric to reach subnet 10.2.2.0/24 is 64.
- C. The route was redistributed on a router that has a router ospf 1 command configured.
- D. R1 is redistributing a route to prefix 10.2.2.0/24 into OSPF.

Answer: C

Explanation:

This output is the external data section of a detailed view of an EIGRP topology table entry for an external route. This output confirms that this route was redistributed into EIGRP. If R1 were the redistributing router, the output would include the phrase "(this system)"; this example does not include that notation. The output means

that on the router that did the redistribution, the route was redistributed from OSPF process 1, and the OSPF metric was 64. R1's metric is not based on the OSPF metric of the route.

QUESTION NO: 391

Router R1 has a connected route for 10.1.1.0/24 off interface Fa0/0. Interface Fa0/0 has been enabled for OSPF due to a router ospf 1 and network 10.0.0.0 0.0.0.255 area 0 command. R1 also has EIGRP configured, with the redistribute ospf 1 metric 1000 100 10 1 1500 command configured under EIGRP.

Which one of the following is true?

- A. R1 will not redistribute 10.1.1.0/24 into EIGRP, because R1 knows it as a connected route and not as an OSPF route.
- B. For any OSPF routes redistributed into EIGRP, the metric components include a value equivalent to 1 Mbps of bandwidth.
- C. For any OSPF routes redistributed into EIGRP, the metric components include a value equivalent to 100 microseconds of delay.
- D. No subnets of network 10.0.0.0 will be redistributed due to the omission of the subnets parameter.

Answer: B

Explanation:

The redistribute ospf command will attempt to redistribute OSPF routes and connected routes from interfaces on which OSPF is enabled. The metric components include 1000 Kbps (or 1 Mbps), 100 tens-of-microseconds (or 1000 microseconds), 10 for the loading, 1 for the reliability, and 1500 for MTU. The EIGRP version of the redistribute command does not include a subnets option.

QUESTION NO: 392

Process OSPF 1 is redistributing routes from process OSPF 2. Which of the following methods may be used to set the metrics of the redistributed routes? (Choose two.)

- A. Let the metrics default.
- B. Use each redistributed route's OSPF metric using the redistribute command's metric transparent keywords.
- C. Set the metric using the default-metric router subcommand.
- D. Redistribution is not allowed between two OSPF processes.

Answer: A,C

Explanation:

Because the routes come from OSPF and feed into OSPF, the metrics can be set with the usual tools, or the metric can default. When taking routes from OSPF into another OSPF process, the default metric is taken from the source route's OSPF cost. Alternatively, the metric can be set for all routes, regardless of the route source, using the default-metric OSPF subcommand. The metric transparent keywords cannot be used for an OSPF redistribute command.

QUESTION NO: 393

Examine the following excerpt from the show ip ospf database asbr-summary command on router R1 (RID 1.1.1.1). Which answer can be verified as definitely true based on this output?

LS Type: Summary Links (AS Boundary Router)
Link State ID: 9.9.9.9 (AS Boundary Router address)
Advertising Router: 3.3.3.3
LS Seq Number: 8000000D
Checksum: 0xE43A
Length: 28
Network Mask: /0
TOS: 0 Metric: 100

- A. The output describes the contents of a Type 5 LSA.
- B. 3.3.3.3 identifies a router as being the router performing redistribution.
- C. R1's metric for its best route to reach the router with RID 9.9.9.9 is 100.
- D. The router with RID 3.3.3.3's metric for its best route to reach the router with RID 9.9.9.9 is 100.

Answer: D

Explanation:

This command lists the output of Type 4 Summary ASBR LSAs. The LSID identifies the redistributing ASBR (9.9.9.9). The advertising router is the ABR that created and flooded the LSA (3.3.3.3), and the metric is the ABR's best metric route to reach the ASBR.

QUESTION NO: 394

Router R1 sits inside OSPF area 1. Router R2 redistributes an E1 route into OSPF for prefix 2.2.2.0/24, with external metric 20. Router R22 redistributes an E2 route for the same prefix/length, external metric 10. Under what conditions will R1 choose as its best route the route through R22?

- A. R1 will always choose the route through R22.
- B. As long as R1's best internal OSPF cost to reach R22 is less than 10.
- C. As long as R1's best internal OSPF cost to reach R22 is less than 20.
- D. R1 will never choose the route through R22 if the E1 route through R2 is available.

Answer: D

Explanation:

Routers add internal and external costs for E1 routes and use only external costs for E2 routes, so the cost for the route through R22 will always be lower. However, for a given prefix/length, OSPF always prefers intra-area routes first, then interarea, then E1, and finally, E2, all regardless of metric.

QUESTION NO: 395

Router R1 has been configured with the redistribute ospf 1 route-map fred command under router eigrp 1. The route map named fred needs to be configured to match routes to determine which routes are redistributed into EIGRP. Which of the following answers lists an item that cannot be matched by route map fred?

- A. Subnet number
- B. Next-hop router IP address of the route
- C. Whether the route is an E1 or E2 route
- D. The route's tag
- E. The number of router hops between the router and the subnet

Answer: E

Explanation:

Because OSPF does not use hop count as a metric, the information about the number of hops is not available in OSPF routes in the IP routing table. The other answers list items that can be matched with the route map match subcommand.

QUESTION NO: 396

Router R1 refers to route-map fred when redistributing from EIGRP into OSPF. The entire route-map is listed next. Which of the following answers must be true based on the configuration as shown?

```
route-map fred deny 10  
match ip address one  
route-map fred deny 20
```

```
match ip address two
route-map fred permit 100
```

- A. The third route map clause will allow any routes not already filtered by the first two clauses.
- B. Routes permitted by ACL "two" will be redistributed.
- C. Routes denied by ACL "one" will be redistributed.
- D. All routes will be filtered.

Answer: A

Explanation:

The deny clauses in the route map mean that the route map will filter routes matched by that clause. The permit or deny action of the referenced ACLs just defines whether the route is matched. So, routes permitted by ACL "two" will be matched and then filtered due to the route-map clause deny action. Routes denied by ACL "one" simply do not match the route map clause numbered 10; such routes may or may not be redistributed depending on the next two clauses. Clause number 100 does not have a match command, meaning it matches all routes not otherwise matched, with a permit action, allowing these routes to be redistributed.

QUESTION NO: 397

On Router R1, process EIGRP 1 is redistributing routes from process OSPF 2, calling route-map fred with the redistribute ospf 2 command. R1 has learned intra-area routes for 10.1.1.0/24 and 10.1.2.0/24 in part due to the Type 2 LSAs known for each subnet. The route map filters route 10.1.1.0/24 and allows 10.1.2.0/24 through. Which of the following commands on router R1 list subnet 10.1.1.0/24? (Choose two.)

- A. show ip route
- B. show ip eigrp topology
- C. show ip ospf topology
- D. show ip eigrp topology 10.1.1.0/24

Answer: A,C

Explanation:

The problem states that R1 has learned OSPF intra-area routes for 10.1.1.0/24, so show ip route will display that subnet. As an intra-area route based on a Type 2 LSA, the show ip ospf topology command lists the summary of the LSAs, including the 10.1.1.0 subnet number for that Type 2 LSA. However, because the redistribution filtering discards subnet 10.1.1.0/24, this value will not be included in the EIGRP topology table.

QUESTION NO: 398

Router R1 is redistributing between two OSPF processes. Given the configuration shown, which includes all commands in the route map named fred, which of the following answers is true regarding the redistribution into OSPF process 1?

```
router ospf 1
redistribute ospf 2 external 2 route-map fred
!
route-map fred permit 10
match ip address 1
set metric-type type-1
```

- A. No routes are redistributed because a route cannot be both E1 and E2.
- B. Only OSPF E2 routes in the OSPF 2 domain will be considered for redistribution.
- C. Inside the OSPF 2 domain, any formerly E2 routes will become E1 routes.
- D. Routes permitted by ACL 1 will be redistributed, regardless of whether the routes are E1 or E2 routes.

Answer: B

Explanation:

The external 2 parameters on the redistribute command act as matching logic; only routes from the source routing protocol (in this case OSPF 2) that match this extra logic will be considered for redistribution by this redistribute command. The set metric-type type-1 route-map subcommand sets the route type as it is injected into the destination routing protocol (in this case OSPF 1); this logic is not used for matching the source routes. The routes permitted by ACL 1 will be redistributed, but only those that are also E2 routes from the (source) OSPF 2 domain. The redistribute function will not change the attributes of routes inside a single routing domain, but only in the destination routing domain (OSPF 1), so the configuration has no effect on the OSPF 2 routes that remain in OSPF 2.

QUESTION NO: 399

Which of the following is not true regarding IOS default settings for administrative distance?

- A. EIGRP internal: 90
- B. OSPF external: 110
- C. EIGRP external: 90
- D. RIP: 120

E. OSPF internal: 110

Answer: C

Explanation:

EIGRP, by default, sets a different AD for internal (90) and external (170) routes. The rest of the answers are accurate regarding default settings.

QUESTION NO: 400

A network includes a RIPv2 domain, an EIGRP domain, and an OSPF domain. Each pair of routing domains has multiple routers redistributing routes between the pair of domains. The design requires that the redistribution configuration avoid matching based on prefix/length because of the trouble in maintaining such configurations.

Which one of the following tools can be used in all three routing domains to attempt to prevent domain loops? (This book uses the term domain loop to refer to the long routes that might be chosen for routes when redistribution exists—for example, a route may forward packets from the EIGRP domain, to the OSPF domain, back to EIGRP, and then to subnet X in the RIP domain.)

- A. Setting route tags
- B. Setting the default administrative distance differently for internal and external routes
- C. Setting administrative distance differently per route
- D. Setting metrics much higher for all external routes than for all internal routes

Answer: A

Explanation:

All the answers list reasonable options in some cases, but the only feature listed that is useful with all three routing protocols is the route tag feature. RIPv2 does not support the concept of differentiating between internal and external routes, so the two answers that suggest setting administrative distance (AD) based on the route type (internal or external) could not be used in all three routing domains, as required by the question. All three routing protocols support setting route tags and setting the AD per route; however, because RIPv2 cannot match based on the route type (internal/external), the option to set the route tags is the only option that applies to all three routing domains.

QUESTION NO: 401

A co-worker is developing an implementation plan for a design that uses OSPF 2 and RIPv2 routing domains, with two routers redistributing between the two domains.

The co-worker asks your help in choosing how to prevent domain loops by setting administrative distance. Assuming all other related settings use defaults, which of the following would solve the domain loop problem?

- A. The distance ospf intra-area 80 inter-area 80 OSPF subcommand
- B. The distance ospf external 80 OSPF subcommand
- C. The distance ospf intra-area 180 inter-area 180 OSPF subcommand
- D. The distance ospf external 180 OSPF subcommand

Answer: D

Explanation:

AD can be used to prevent the domain loop problem with two routing domains by making each routing protocol's AD for internal routes be better (lower) than the other routing protocol's AD for external routes. RIP uses AD 120 for all routes, with no distinction of internal or external. As such, OSPF's internal default AD settings of 110 meet the requirement that OSPF's internal AD (110) is better than RIP's external (120). However, RIP's default of 120 is not better than OSPF's default for externals (110), so the distance ospf external 180 command changes that setting to meet both requirements. The three wrong answers, while syntactically valid, do not help meet the requirements.

QUESTION NO: 402

Router R1 sets a route tag for subnet 10.1.1.0/24 when redistributing from OSPF into EIGRP. Which of the following unit is assigned to the route tag?

- A. Kilobits/second
- B. Tens of microseconds.
- C. Cost.
- D. Hop count.
- E. No unit is assigned.

Answer: E

Explanation:

Route tags are unitless integers that can be given to a route and even passed between different routing protocols by routers that perform redistribution.

QUESTION NO: 403

Policy-Based Routing (PBR) has been enabled on Router R1's interface F0/0. Which of the following is true regarding how PBR works? (Choose two.)

- A. Packets entering F0/0 will be compared based on the PBR route map.
- B. Packets exiting F0/0 will be compared based on the PBR route map.
- C. IOS ignores the PBR forwarding directions when the packet matches a route map deny clause.
- D. IOS ignores the PBR forwarding directions when the packet matches a route map permit clause.

Answer: A,C

Explanation:

PBR supports processing packets on an interface, for the inbound direction only. The referenced route map causes PBR to attempt policy routing of packets that match a permit clause in the route map.

QUESTION NO: 404

Examine the following configuration on Router R1. R1's show ip route 172.16.4.1 command lists a route with outgoing interface S0/1/1. Host 172.16.3.3 uses telnet to connect to host 172.16.4.1. What will Router R1 do with the packets generated by host 172.16.3.3 because of the telnet, assuming the packets enter R1's F0/0 interface?

(Choose two.)

```
interface FastEthernet 0/0
ip address 172.16.1.1 255.255.255.0
ip policy route-map Q2
!
route-map Q2 permit
match ip address 101
set interface s0/0/1
!
access-list 101 permit tcp host 172.16.3.3 172.16.4.0 0.0.0.255
```

- A. The packet will be forwarded out S0/0/1, or not at all.
- B. The packet will be forwarded out S0/0/1 if it is up.
- C. The packet will be forwarded out S0/1/1 if it is up.
- D. The packet will be forwarded out S0/1/1 if it is up, or if it is not up, out s0/0/1.
- E. The packet will be forwarded out S0/0/1 if it is up, or if it is not up, out s0/1/1.

Answer: B,E

Explanation:

Packets created by Telnet use TCP, so the packet will match ACL 101 with a permit action. So, PBR will match the only route map clause shown in the configuration, with that permit route map clause listing a set command. The set command lists S0/0/1 as the outgoing interface, and without a default parameter. So, Router R1 will first attempt to forward the packet based on the set

command (interface S0/0/1), but if the interface is down, then try to forward based on the IP routing table (interface S0/1/1).

QUESTION NO: 405

The following output occurs on Router R2. Which of the following statements can be confirmed as true based on the output from R2?

```
R2# show ip policy
Interface Route map
Fa0/0 RM1
Fa0/1 RM2
S0/0/0 RM3
```

- A. R2 will forward all packets that enter Fa0/0 per the PBR configuration.
- B. R2 will use route map RM2 when determining how to forward packets that exit interface Fa0/1.
- C. R2 will consider using PBR for all packets exiting S0/0/0 per route map RM3.
- D. R2 will consider using PBR for all packets entering S0/0/0 per route map RM3.

Answer: D

Explanation:

The output from the show ip policy command shows the interfaces on which PBR has been enabled, and the name of the route map enabled for PBR on each interface. For the purposes of this question, the output tells us the interfaces on which PBR has been enabled. Two answers mention packets exiting the interface, so these answers cannot be correct, because PBR applies to packets entering an interface. For the two interfaces that mention inbound packets, one suggests that all packets will be forwarded per the PBR configuration; some may not be forwarded per PBR, depending on the configuration of the route map. The correct answer specifically mentions that PBR will consider all packets with PBR, which is the more accurate statement about PBR operations.

QUESTION NO: 406

Which of the following are examples of traffic that can be created as part of an IP Service-Level Agreement operation? (Choose two.)

- A. ICMP Echo
- B. VoIP (RTP)
- C. IPX
- D. SNMP

Answer: A,B

Explanation:

The IP SLA feature focuses on IP traffic, so IOS does not include Novell's older IPX protocol as part of IP SLA. IP SLA uses SNMP MIBs to store statistics, but it does not use SNMP as an operation.

QUESTION NO: 407

The following configuration commands exist only in an implementation plan document. An engineer does a copy/paste of these commands into configuration mode on Router R1. Which of the following answers is most accurate regarding the results?

```
ip sla 1
icmp-echo 1.1.1.1 source-ip 2.2.2.2
ip sla schedule 1 start-time now life forever
```

- A. The SLA operation will be configured but will not start until additional commands are used.
- B. The SLA operation is not completely configured so it will not collect any data.
- C. The SLA operation is complete and working, collecting data into the RTTMON MIB.
- D. The SLA operation is complete and working but will not store the data in the RTTMON MIB without more configuration.

Answer: C

Explanation:

The three lines shown create the operation number (first command), define the operation (second command), and start the operation (third command). All commands are correct. After the operation is started, IP SLA stores the data in the RTTMON MIB-no additional configuration necessary.

QUESTION NO: 408

The following output occurs on Router R1. IP SLA operation 1 uses an ICMP echo operation type, with default frequency of 60 seconds. The operation pings from address 1.1.1.1 to address 2.2.2.2. Which of the following answers is true regarding IP SLA and object tracking on R1?

```
R1# show track
Track 2
IP SLA 1 state
State is Up
```

3 changes, last change 00:00:03
Delay up 45 secs, down 55 secs
Latest operation return code: OK
Latest RTT (milliseconds) 6
Tracked by:
STATIC-IP-ROUTING 0

- A. The tracking return code fails after the SLA operation results in an ICMP echo failure three times.
- B. The tracking return code fails after the SLA operation results in an ICMP echo failure one time.
- C. After the tracking object fails, the tracking object moves back to an up state 45 seconds later in all cases.
- D. After moving to a down state, the tracking object moves back to an OK state 45 seconds after the SLA operation moves to an OK state.

Answer: D

Explanation:

The up timers on the tracking object defines how long to wait, when in a down state, after seeing the IP SLA object transition to an OK state. Similarly, the down timer defines how long to wait, when in an OK state, after seeing the IP SLA object move to a down state, before moving the tracking object to a down state.

QUESTION NO: 409

Which of the following are considered private IPv4 addresses? (Choose two.)

- A. 192.16.1.1
- B. 172.35.1.1
- C. 225.0.0.1
- D. 127.0.0.1
- E. 10.1.1.1

Answer: B,E

Explanation:

The private IPv4 address space consists of Class A network 10.0.0.0, Class B networks 172.16.0.0-172.31.0.0, and the 256 Class C networks that begin 192.168.

QUESTION NO: 410

Class C network 200.1.1.0/24 was allocated to an ISP that operated primarily in Asia. That ISP then assigned this entire Class C network to one of its Asian customers. Network

200.1.2.0/24 has yet to be assigned to any ISP. Which of the following is most likely to be true?

- A. 200.1.2.0/24 could be assigned to any registrar or ISP in the world.
- B. 200.1.2.0/24 will be assigned in the same geography (Asia) as 200.1.1.0/24.
- C. 200.1.2.0/24 cannot be assigned as public address space.
- D. Routers inside North American ISPs increase their routing table size by 1 as a result of the customer with 200.1.1.0/24 connecting to the Internet.

Answer: B

Explanation:

ICANN and IANA manage the assignment of public IPv4 address space such that large address blocks (often called CIDR blocks) exist in a particular geography or are assigned to particular ISPs. As such, Internet routers can more easily create summary routes to help keep the routing table small in the Internet. 200.1.2.0/24 would likely also be allocated to some registrar, ISP, or customer in Asia. Because of the large route summaries, in this case possibly a summary for 200.0.0.0/8, routers in North America would not see an increase in the size of their routing tables.

QUESTION NO: 411

Router R1, in ASN 11, learns a BGP route from BGP peer R22 in ASN 22. R1 and then uses BGP to advertise the route to R2, also in ASN 11. What ASNs would you see in the BGP table on R2 for this route?

- A. 22
- B. 11
- C. 1
- D. None

Answer: A

Explanation:

The router in ASN 22, R22, advertises the BGP update with (at least) 22 in the AS_Path Path Attribute (PA). When R1 advertises the route to R2, also in ASN 11, R1 does not add an ASN. As a result, R2's AS_Path has at least ASN 22 and not ASN 11.

QUESTION NO: 412

Which of the following are most likely to be used as an ASN by a company that has a registered public 16-bit ASN? (Choose two.)

- A. 1
- B. 65,000
- C. 64,000
- D. 64,550

Answer: A,C

Explanation:

The public range of 16-bit BGP ASNs is 1 through 64,495.

QUESTION NO: 413

Which of the following statements is true about a router's eBGP peers that is not also true about that same router's iBGP peers?

- A. The eBGP peer neighborhood uses TCP.
- B. The eBGP peer uses port 180 (default).
- C. The eBGP peer uses the same ASN as the local router.
- D. The eBGP peer updates its AS_Path PA before sending updates to this router.

Answer: D

Explanation:

The question asks which answers are true about the eBGP peer but also not true about an iBGP peer. Both iBGP and eBGP use TCP port 179. An eBGP peer uses a different ASN than the local router, by definition, making that answer incorrect.

The correct answer refers to the fact that an eBGP peer adds its own ASN to the BGP AS_Path PA before sending routing information to another router, whereas iBGP peers do not.

QUESTION NO: 414

Which of the following is the primary motivation for using BGP between an Enterprise and its ISPs?

- A. To influence the choice of best path (best route) for at least some routes
- B. To avoid having to configure static routes
- C. To allow redistribution of BGP routes into the IGP routing protocol
- D. To monitor the size of the Internet BGP table

Answer: A

Explanation:

Although using BGP does avoid some static configuration at the Enterprise and the ISP, the primary reason to consider using BGP in the Enterprise is to influence and react to Path Attributes for the purpose of choosing the best path. Typically, engineers do not redistribute BGP routes into the IGP due to scalability problems. And although it may be interesting to monitor the size of the Internet BGP table, it is not a primary motivation for choosing to use BGP on a router.

QUESTION NO: 415

The following terms describe various design options for Enterprise connectivity to the Internet. Which of the following imply that the Enterprise connects to two or more ISPs? (Choose two.)

- A. Single Homed
- B. Dual Homed
- C. Single Multihomed
- D. Dual Multihomed

Answer: C,D

Explanation:

The terms "homed" makes reference to a single homed ISP, and "multihomed" to multiple ISPs. The terms "single" and "dual" refer to the number of connections to each ISP.

QUESTION NO: 416

Enterprise Router R1, in ASN 1, connects to ISP Router I1, ASN 2, using eBGP. The single serial link between the two routers uses IP addresses 10.1.1.1 and 10.1.1.2, respectively. Both routers use their S0/0 interfaces for this link. Which of the following commands would be needed on R1 to configure eBGP? (Choose two.)

- A. router bgp 2
- B. router bgp 1
- C. neighbor 10.1.1.2 remote-as 2
- D. neighbor 10.1.1.2 Update-source 10.1.1.1
- E. neighbor 10.1.1.2 Update-source S0/0

Answer: B,C

Explanation:

The router bgp command lists the local ASN, and the neighbor remote-as command lists the neighbor's ASN. Because the neighbor relationship uses the IP addresses on the common link,

the routers do not need to identify the update source interface, because each will default to use their S0/0 interfaces (in this case) as the update source.

QUESTION NO: 417

Enterprise Router R1, in ASN 1, connects to ISP Router I1, ASN 2, using eBGP. There are two parallel serial links between the two routers. The implementation plan calls for each router to base their BGP TCP connection on their respective loopback1 interfaces, with IP addresses 1.1.1.1 and 2.2.2.2, respectively. Which of the following commands would not be part of a working eBGP configuration on Router R1?

- A. router bgp 1
- B. neighbor 2.2.2.2 remote-as 2
- C. neighbor 2.2.2.2 update-source loopback1
- D. neighbor 2.2.2.2 multihop 2

Answer: D

Explanation:

Three of the commands list valid commands. The neighbor 2.2.2.2 multihop 2 command is syntactically incorrect; it should be neighbor 2.2.2.2 ebgp-multihop 2 .

QUESTION NO: 418

The following output, taken from a show ip bgp command on Router R1, lists two neighbors. In what BGP neighbor state is neighbor 1.1.1.1?

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down
State/PfxRcd
1.1.1.1 4 1 60 61 26 0 0 00:45:01
0
2.2.2.2 4 3 153 159 26 0 0 00:38:13
1
```

- A. Idle
- B. Opensent
- C. Active
- D. Established

Answer: D

Explanation:

The show ip bgp command lists the BGP neighbor state in the last column of output, listing the literal state, unless in an established state. In that state, the output lists the number of prefixes learned from the neighbor, so a numeric value implies an established state.

QUESTION NO: 419

The following output was taken from the show ip bgp command on Router R2. In this case, which of the following commands are most likely to already be configured on R2? (Choose two.)

BGP router identifier 11.11.11.11, local AS number 11

...

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down
State/PfxRcd

1.1.1.1 4 1 87 87 0 0 0 00:00:06 Idle

(Admin)

2.2.2.2 4 3 173 183 41 0 0 00:58:47

2

- A. router bgp 11
- B. neighbor 1.1.1.1 remote-as 11
- C. neighbor 2.2.2.2 prefix-limit 1
- D. neighbor 1.1.1.1 shutdown

Answer: A,D

Explanation:

The output lists R2's local ASN as ASN 11, a value that is configured in the router bgp asn command. The line for neighbor 1.1.1.1 lists that router's ASN as 1, so a neighbor 1.1.1.1 remote-as 1 command should exist on R2 instead of the neighbor 1.1.1.1 remote-as 11 command. The state for neighbor 1.1.1.1 lists "Idle (Admin)," implying that the neighbor 1.1.1.1 shutdown command has been configured. The other answer lists a nonexistent command.

QUESTION NO: 420

Which of the following answers is most true about the BGP Update message?

- A. It lists a set of path attributes, along with a list of prefixes that use those PAs.
- B. It lists a prefix/length, plus the PA settings for that prefix.
- C. It lists withdrawn routes, but never in the same Update message as newly advertised routes.
- D. A single Update message lists at most a single prefix/length.

Answer: A

Explanation:

The BGP Update message lists a set of PAs, plus any prefixes/lengths that use those PAs. It can also list withdrawn routes in the same Update message as newly advertised routes. It can also list multiple prefixes in a single Update message.

QUESTION NO: 421

The following output occurs on Router R1. Which of the following cannot be determined from this output?

```
R1# show ip route 180.1.1.0 255.255.255.240
Routing entry for 180.1.1.0/28
Known via "bgp 2", distance 20, metric 0
Tag 3, type external
Last update from 192.168.1.2 00:10:27 ago
Routing Descriptor Blocks:
* 192.168.1.2, from 192.168.1.2, 00:10:27 ago
Route metric is 0, traffic share count is 1
AS Hops 2
Route tag 3
```

- A. The type of BGP peer (iBGP or eBGP) that advertised this route to R1
- B. R1's ASN
- C. The next-hop router's ASN
- D. The AS_Path length

Answer: C

Explanation:

The "Known via" text refers to the local router's (R1's) router bgp command, which identifies the local router's ASN. The rest of the output does not identify the neighboring ASN, nor the rest of the AS_Path details. It does list that the route is external, with the text "type external", and the AS Hops (which is the AS_Path length).

QUESTION NO: 422

The following line of output was extracted from the output of the show ip bgp command on Router R1. Which of the following can be determined from this output?

```
Network Next Hop Metric LocPrf Weight Path
* 130.1.1.0/28 1.1.1.1 0 1 2 3 4 i
```

- A. The route is learned from an eBGP peer.
- B. The route has no more than three ASNs in the AS_Path.
- C. The route is the best route for this prefix.
- D. None of these facts can be positively determined by this output.

Answer: A

Explanation:

The third character in each line for each router is either blank, meaning the route is an eBGP route, or an "i," meaning an iBGP-learned route. The contents of the AS_Path can be determined (1, 2, 3, 4), but the answer about AS_Path does not suggest 4 ASNs. The best route for each prefix has a ">" in the second character, and this route does not.

QUESTION NO: 423

Router R1 has eBGP connections to I1 and I2, routers at the same ISP. The company that owns R1 can use public address range 130.1.16.0/20. The following output lists all the IP routes in R1's routing table within this range. Which of the following answers would cause R1 to advertise the 130.1.16.0/20 prefix to its eBGP peers? (You should assume default settings for any parameters not mentioned in this question.)

```
R1# show ip route 130.1.16.0 255.255.240.0 longer-prefixes
! lines omitted...
O 130.1.16.0/24 [110/3] via 10.5.1.1, 00:14:36, FastEthernet0/1
O 130.1.17.0/24 [110/3] via 10.5.1.1, 00:14:36, FastEthernet0/1
O 130.1.18.0/24 [110/3] via 10.5.1.1, 00:14:36, FastEthernet0/1
```

- A. Configure R1 with the network 130.1.16.0 mask 255.255.240.0 command.
- B. Configure R1 with the network 130.1.16.0 mask 255.255.240.0 summaryonly command.
- C. Redistribute from OSPF into BGP, filtering so that only routes in the 130.1.16.0/20 range are redistributed.
- D. Redistribute from OSPF into BGP, filtering so that only routes in the 130.1.16.0/20 range are redistributed, and create a BGP summary for 130.1.16.0/20.

Answer: D

Explanation:

The network command will take the route from the IP routing table and put the equivalent into the BGP table, if that exact route exists. The output does not show a route for 130.1.16.0/20, so the network 130.1.16.0 mask 255.255.240.0 command does not match a specific route. The other answer with a network command is syntactically incorrect. Redistribution without aggregation would redistribute the three routes, but all three subordinate routes would be advertised into eBGP. By also using BGP route summarization, a single route for 130.1.16.0/20 can be

advertised.

QUESTION NO: 424

R1 in ASN 1 with loopback1 address 1.1.1.1 needs to be configured with an iBGP connection to R2 with loopback2 IP address 2.2.2.2. The connection should use the loopbacks. Which of the following commands is required on R1?

- A. neighbor 1.1.1.1 remote-as 1
- B. neighbor 2.2.2.2 remote-as 2
- C. neighbor 2.2.2.2 update-source loopback1
- D. neighbor 2.2.2.2 ibgp-multihop 2
- E. neighbor 2.2.2.2 ibgp-mode

Answer: C**Explanation:**

R1 needs to be configured with router bgp 1 , neighbor 2.2.2.2 remote-as 1 , and neighbor 2.2.2.2 update-source loopback1 . The neighbor 2.2.2.2 ibgp-multihop 2 and neighbor 2.2.2.2 ibgp-mode commands are simply unsupported commands. The neighbor 1.1.1.1 remote-as 1 command has correct syntax and is used as a command in R2's configuration but not on R1. The neighbor 2.2.2.2 remote-as 2 command has correct syntax but with the wrong ASN (2 instead of 1).

QUESTION NO: 425

The following output occurred as a result of the show ip bgp command on Router R1. The output shows all BGP table entries on R1. How many iBGP-learned routes exist on this router?

```
*>i181.0.0.0/8 10.100.1.1 0 100 0 1 2 111 112 i
*>i182.0.0.0/8 10.100.1.1 0 100 0 1 2 222 i
*>i183.0.0.0/8 10.100.1.1 0 100 0 1 2 i
*>i184.0.0.0/8 10.100.1.1 0 100 0 1 2 i
*> 192.135.250.0/28 192.168.1.6 0 3 4 i
```

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

Answer: D

Explanation:

The small letter "i" in the third character position implies the route was learned with iBGP. Of the five lines, four have an "i" in the third column.

QUESTION NO: 426

The following output on Router R1 lists details of a BGP route for 190.1.0.0/16. Which of the following is true based on this output? (Choose 2)

```
R1# show ip bgp 190.1.0.0/16
```

```
BGP routing table entry for 190.1.0.0/16, version 121
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

```
Advertised to update-groups:
```

```
1
```

```
1 2 3 4
```

```
1.1.1.1 from 2.2.2.2 (3.3.3.3)
```

```
Origin IGP, metric 0, localpref 100, valid, internal, best
```

- A. R1 has a neighbor 1.1.1.1 command configured.
- B. R1 has a neighbor 2.2.2.2 command configured.
- C. The show ip bgp command lists a line for 190.1.0.0/16 with both an > and an i on the left.
- D. R1 is in ASN 1.

Answer: B,C

Explanation:

The line reading "1.1.1.1 from 2.2.2.2..." implies the BGP RID of the neighbor is 1.1.1.1, with neighbor ID-the IP address on the local router's neighbor command-of 2.2.2.2. The end of the output shows that the route is internal (iBGP learned) and is best, so both the > and i will be displayed for this route by the show ip bgp command. Finally, the output does not identify the local ASN, although it does list the AS_Path of the route (1, 2, 3, 4).

QUESTION NO: 427

A company uses Routers R1 and R2 to connect to ISP1 and ISP2, respectively, with Routers I1 and I2 used at the ISPs. R1 peers with I1 and R2; R2 peers with I2 and R1. Assuming as many default settings as possible are used on all four routers, which of the following is true about the next-hop IP address for routes R1 learns over its iBGP connection to R2?

- A. The next hop is I2's BGP RID.
- B. The next hop is I2's IP address used on the R2-I2 neighbor relationship.
- C. The next hop is R2's BGP RID.
- D. The next hop is R2's IP address used on the R1-R2 neighbor relationship.

Answer: B

Explanation:

By default, when a router advertises an iBGP route, it leaves the Next-Hop PA unchanged.

By default, R2's next hop for routes learned from I2 will be I2's IP address used on the R2-I2 neighbor relationship.

QUESTION NO: 428

A company uses Routers R1 and R2 to connect to ISP1 and ISP2, respectively, with Routers I1 and I2 used at the ISPs. R1 peers with I1 and R2; R2 peers with I2 and R1. R1 and R2 do not share a common subnet, relying on other routers internal to the Enterprise for IP connectivity between the two routers. Which of the following could be used to prevent potential routing loops in this design? (Choose 2)

- A. Using an iBGP mesh inside the Enterprise core
- B. Configuring default routes in the Enterprise pointing to both R1 and R2
- C. Redistributing BGP routes into the Enterprise IGP
- D. Tunneling the packets for the iBGP connection between R1 and R2

Answer: A,C

Explanation:

The Enterprise core routers need to know which exit point (R1 or R2) is best; the correct answers supply those routes to the routers internal to the company. Note that redistribution from BGP into the IGP is not recommended, but it does defeat this particular problem.

QUESTION NO: 429

R1 is currently advertising prefixes 1.0.0.0/8, 2.0.0.0/8, and 3.0.0.0/8 over its eBGP connection to neighbor 2.2.2.2 (R2). An engineer configures a prefix list (fred) on R1 that permits only 2.0.0.0/8 and then enables the filter with the neighbor R2 prefix-list fred out command. Upon exiting configuration mode, the engineer uses some show commands on R1, but no other commands. Which of the following is true in this case?

- A. The show ip bgp neighbor 2.2.2.2 received-routes command lists the three original prefixes.
- B. The show ip bgp neighbor 2.2.2.2 advertised-routes command lists the three original prefixes.

- C. The show ip bgp neighbor 2.2.2.2 routes command lists the three original prefixes.
- D. The show ip bgp neighbor 2.2.2.2 routes command lists only 2.0.0.0/8.
- E. The show ip bgp neighbor 2.2.2.2 advertised-routes command lists only 2.0.0.0/8.

Answer: B

Explanation:

The show ip bgp neighbors 2.2.2.2 advertised-routes command does list the postoutbound-filter BGP Update; however, the user did not issue a clear command, so the filter has not yet taken effect. As such, the output still lists the original three prefixes as if the filter had not yet been applied.

QUESTION NO: 430

Which of the following three BGP filtering methods enabled with the neighbor command will filter BGP prefixes based on the prefix and prefix length? (Choose 3)

- A. A neighbor distribute-list out command, referencing a standard ACL
- B. A neighbor prefix-list out command
- C. A neighbor filter-list out command
- D. A neighbor distribute-list out command, referencing an extended ACL
- E. A neighbor route-map out command

Answer: B,D,E

Explanation:

The neighbor distribute-list out command refers to an ACL, but for the ACL to match on both prefix and prefix length, the ACL must be an extended ACL. The neighbor filter-list command refers to an AS-path filter and cannot match based on prefix/length

QUESTION NO: 431

Which of the following commands causes a router to bring down BGP neighbor relationships?

- A. clear ip bgp *
- B. clear ip bgp 1.1.1.1
- C. clear ip bgp * soft
- D. clear ip bgp 1.1.1.1 out

Answer: A,B

Explanation:

The router resets the BGP neighborship when performing a hard reset of the peer.

QUESTION NO: 432

An engineer is preparing an implementation plan in which the configuration needs to influence BGP's choice of best path. Which of the following is least likely to be used by the configuration in this implementation plan?

- A. Weight
- B. Origin code
- C. AS_Path
- D. Local_Pref

Answer: B**Explanation:**

Weight and Local_Pref were created for the purpose of giving engineers tools to influence the BGP best path choice. AS_Path was created for loop avoidance, but AS_Path length can also be manipulated (for instance, with AS_Path prepend) to influence the best path choice. Although the Origin PA can be changed by configuration for the purpose of influencing the best path decision, the intent of this PA is to identify the source from which the route was introduced into BGP. Additionally, the best path algorithm considers the Origin PA after the other PAs listed in the answers, making Origin the least useful of these answers for influencing path choice.

QUESTION NO: 433

Router R1 learns two routes with BGP for prefix 200.1.0.0/16. Comparing the two routes, route 1 has a longer AS_Path Length, bigger MED, bigger Weight, and smaller Local Preference. Which of the following is true about Router R1's choice of best path for this prefix?

- A. Route 1 is the best route.
- B. Route 2 is the best route.
- C. The routes tie as best, but one will be picked to be placed in the routing table based on tiebreakers.
- D. Neither route is considered best.

Answer: A**Explanation:**

Of the items listed in the question, Weight is the first one considered in the best path algorithm, with a bigger weight being better. As a result, Route 1 is the better route of the two.

QUESTION NO: 434

Router R1 learns two routes with BGP for prefix 200.1.0.0/16. Comparing the two routes, route 1 has a shorter AS_Path Length, smaller MED, the same Weight, and smaller Local Preference. Which of the following is true about Router R1's choice of best path for this prefix?

- A. Route 1 is the best route.
- B. Route 2 is the best route.
- C. The routes tie as best, but one will be picked to be placed in the routing table based on tiebreakers.
- D. Neither route is considered best.

Answer: B

Explanation:

Of the items listed in the question, Weight is the first one considered in the best path algorithm, and it is a tie. The next item considered, Local Preference, uses bigger-is-better logic, so Route 2 will be considered best.

QUESTION NO: 435

An engineer has been told to create an implementation plan to influence the choice of best BGP route on a single router using the Weight feature. The sole Enterprise Internet-connected router, Ent1, has neighbor relationships with Routers ISP1 and ISP2, which reside inside two different ISPs. The goal is to prefer all routes learned from ISP1 over ISP2 using Weight. Which of the following answers lists a configuration step that would not be useful for achieving these goals? (Choose two.)

- A. Configuring the neighbor weight command on Ent1.
- B. Having the ISPs configure the neighbor route-map out command on ISP1 and ISP2, with the route map setting weight.
- C. Configuring the set weight command inside a route map on Router Ent1.
- D. Configuring a prefix list to match all class C networks.

Answer: B,D

Explanation:

Weight, a Cisco-proprietary feature of BGP on Cisco routers, cannot be transmitted in a BGP Update, so setting Weight on an outbound route map at the ISPs will have no effect. Also, the goals call for setting Weight for all routes from an ISP to the same number, so creating a prefix list to match a subset of reachable prefixes, in this case all class C networks, is not useful. However, two methods of configuring Weight do exist: the neighbor weight command and configuring an

inbound route map with a set weight command in the route map.

QUESTION NO: 436

The following output on Router R1 lists details of a BGP route for 190.1.0.0/16. Which of the following is true based on this output? (Choose two.)

```
R1# show ip bgp 190.1.0.0/16
BGP routing table entry for 190.1.0.0/16, version 121
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Advertised to update-groups:
 1
 1 2 3 4
1.1.1.1 from 2.2.2.2 (3.3.3.3)
Origin IGP, metric 0, localpref 100, valid, internal, best
```

- A. R1 has a neighbor 1.1.1.1 command configured.
- B. R1 has a neighbor 2.2.2.2 command configured.
- C. The show ip bgp command lists a line for 190.1.0.0/16 with both an ">" and an "i" on the left.
- D. R1 is in ASN 1.

Answer: B,C

Explanation:

The line reading 1.1.1.1 from 2.2.2.2... implies the BGP RID of the neighbor is 1.1.1.1, with neighbor ID-the IP address on the local router's neighbor command-of 2.2.2.2. The end of the output shows that the route is internal (iBGP learned), and the output lists the word "best," so the show ip bgp command will display both the > and i for this route. Finally, the output does not identify the local ASN, although it does list the AS_Path of the route (1, 2, 3, 4).

QUESTION NO: 437

An Enterprise router, Ent1, displays the following excerpt from the show ip bgp command. ENT1 has an eBGP connection to an ISP router with address 3.3.3.3 and an iBGP connection to a router with address 4.4.4.4. Which of the following is most likely to be true?

```
Network Next Hop Metric LocPrf Weight Path
```

```
*> 3.3.3.3 0 0 1 1 1 1 2
```

```
18 i
```

- A. The Enterprise likely uses ASN 1.

- B. The neighboring ISP likely uses ASN 1.
- C. The route has been advertised through ASN 1 multiple times.
- D. Router Ent1 will add another ASN to the AS_Path before advertising this route to its iBGP peer (4.4.4.4).

Answer: B

Explanation:

The output shows the results of AS_Path prepending. The repetitive 1's cannot mean that the route has been advertised into and out of the same ASN repeatedly because loop prevention would have prevented such an advertisement. With AS_Path prepending, the neighboring ASN typically adds its own ASN to the end of the AS_Path (as listed on the left of the output).

QUESTION NO: 438

The following line of output was gathered on Enterprise Router Ent1 using the command show ip route. Which of the following answers is most likely to be true, based on this output?

B 128.107.0.0 [20/10] via 11.11.11.11, 00:02:18

- A. This router has set the Weight of this route to 10.
- B. This router's BGP table lists this route as an iBGP route.
- C. This router's MED has been set to 10.
- D. This router's BGP table lists an AS_Path length of 10 for this route.

Answer: C

Explanation:

The command lists the administrative distance as the first number inside the square brackets and the MED values as the second number in brackets. The AD of 20 implies an eBGP route instead of iBGP. The output says nothing about the Weight or AS_Path length.

QUESTION NO: 439

Which of the following is the shortest valid abbreviation for FE80:0000:0000:0000:0010:0000:0000:0123?

- A. FE80::10::123
- B. FE8::1::123
- C. FE80:0:0:0:10::123
- D. FE80::10:0:0:123

Answer: D

Explanation:

Inside a quartet, any leading 0s can be omitted, and one sequence of one or more quartets of all 0s can be replaced with "::". The correct answer replaces the longer three-quartet sequence of 0s with ::.

QUESTION NO: 440

An ISP has assigned prefix 3000:1234:5678::/48 to Company1. Which of the following terms would typically be used to describe this type of public IPv6 prefix?

- A. Subnet prefix
- B. ISP prefix
- C. Global routing prefix
- D. Registry prefix

Answer: C

Explanation:

The name of the prefix generally represents the group to which the prefix is given, with the exception of the term global routing. IANA assigns a prefix to a registry (registry prefix). The registry may assign a subset of that range as a prefix to an ISP (ISP prefix). That ISP then subdivides that range of addresses into prefixes and assigns a prefix to one of its customers (site prefix, also called global routing prefix). The Enterprise network engineers then further subdivides the range, often with prefix length 64, into subnet prefixes.

QUESTION NO: 441

Which of the following answers lists either a protocol or function that can be used by a host to dynamically learn its own IPv6 address? (Choose two.)

- A. Stateful DHCP
- B. Stateless DHCP
- C. Stateless autoconfiguration
- D. Neighbor Discovery Protocol

Answer: A,C

Explanation:

IPv6 supports stateful DHCP, which works similarly to IPv4's DHCP protocol to dynamically assign the entire IP address. Stateless autoconfiguration also allows for the assignment by finding the prefix from some nearby router and

calculating the Interface ID using EUI-64 format. Stateless DHCP simply supplies the DNS server IP addresses, and NDP supplies Layer 2 mapping information.

QUESTION NO: 442

Which of the following is helpful to allow an IPv6 host to learn the IP address of a default gateway on its subnet?

- A. Stateful DHCP
- B. Stateless RS
- C. Stateless autoconfiguration
- D. Neighbor Discovery Protocol

Answer: D

Explanation:

Stateless autoconfiguration only helps a host learn and form its own IP address, but it does not help the host learn a default gateway. Stateless RS is not a valid term or feature. Neighbor Discovery Protocol (NDP) is used for several purposes, including the same purpose as ARP in IPv4, plus to learn configuration parameters such as a default gateway IP address.

QUESTION NO: 443

Which of the following answers lists a multicast IPv6 address?

- A. 2000::1:1234:5678:9ABC
- B. FD80::1:1234:5678:9ABC
- C. FE80::1:1234:5678:9ABC
- D. FF80::1:1234:5678:9ABC

Answer: D

Explanation:

Global unicast addresses begin with 2000::/3, meaning the first three bits match the value in hex 2000. Similarly, unique local addresses match FD00::/8, and link local addresses match FE80::/10 (values that begin with FE8, FE9, FEA, and FEB hex). Multicast IPv6 addresses begin FF00::/8, meaning the first two hex digits are F.

QUESTION NO: 444

Router R1 has two LAN interfaces and three serial interfaces enabled for IPv6. All the interfaces use link local addresses automatically generated by the router. Which of the following could be the

link local address of R1's interface S0/0?

- A. FEA0::200:FF:FE11:0
- B. FE80::200:FF:FE11:1111
- C. FE80::0213:19FF:FE7B:0:1
- D. FEB0::211:11FF:FE11:1111

Answer: B

Explanation:

When created automatically, link local addresses begin FE80::/64, because after the prefix of FE80::/10, the device builds the next 54 bits as binary 0s. Statically assigned link local addresses simply need to confirm to the FE80::/10 prefix. As a result, only two answers are candidates with a beginning quartet of FE80. Of these, only one has only hex 0s in the second, third, and fourth quartets, making answer B the only valid answer.

QUESTION NO: 445

Router R1 has the following configuration. Assuming R1's F0/0 interface has a MAC address of 0200.0011.1111, what IPv6 addresses will R1 list for interface F0/0 in the output of the show ipv6 interface brief command?

```
interface f0/0
ipv6 address 2345:0:0:8::1/64
```

- A. 2345:0:0:8::1
- B. 2345:0:0:8:0:FF:FE11:1111
- C. FE80::8:0:FF:FE11:1111
- D. FE80:0:0:8::1

Answer: A,C

Explanation:

The ipv6 address command does not list an eui-64 parameter, so R1 does not form its global unicast address using the EUI-64 format. However, it does form its link local address using EUI-64. The show ipv6 interface brief command lists both the global unicast and link local addresses in its output.

QUESTION NO: 446

Router R1 lists the following output from a show command. Which of the following is true about R1?

```
R1# show ipv6 interface f0/0
FastEthernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::213:19FF:FE12:3456
No Virtual link-local address(es):
Global unicast address(es):
2000::4:213:19FF:FE12:3456, subnet is 2000:0:0:4::/64 [EUI]
Joined group address(es):
FF02::1
FF02::2
FF02::1:FF:12:3456
```

- A. R1's solicited node multicast address is FF02::1:FF:12:3456.
- B. R1's 2000::4:213:19FF:FE12:3456 address is a global unicast with all 128 bits statically configured.
- C. Address FF02::2 is R1's solicited node multicast.
- D. R1's solicited node multicast, not listed in this output, would be FF02::213:19FF:FE12:3456.

Answer: A

Explanation:

The group addresses listed in the output are the all IPv6 hosts address (FF02::1), the all IPv6 routers address (FF02::2), and the solicited node address that is based on R1's global unicast address (FF02::1:FF:12:3456). Also, R1's global unicast address is listed correctly in answer B, but the "[EUI]" notation implies that R1 derived the interface ID portion using EUI-64 conventions.

QUESTION NO: 447

Which of the following features work the same in both RIP-2 and RIPng? (Choose three.)

- A. Distance Vector Logic
- B. Uses UDP
- C. Uses RIP-specific authentication
- D. Maximum useful metric of 15
- E. Automatic route summarization

Answer: A,B,D

Explanation:

RIP-2 and RIPng both use UDP, both use Distance Vector logic, and both use the same metric, with the same maximum (15) and same metric that means infinity (16). RIPng does not perform automatic route summarization because IPv6 has no concept of a classful network. RIPng also uses the built-in IPv6 authentication mechanisms, rather than a RIP-specific authentication such as RIP-2.

QUESTION NO: 448

Router R1 currently has no configuration related to IPv6 or IPv4. The following configuration exists in a planning document, intended to be used to copy/paste into Router R1 to enable RIPng and IPv6 on interfaces F0/0 and S0/0/0. No other related configuration exists. Which of the following is true about RIPng on R1 after this configuration has been pasted into R1?

```
ipv6 unicast-routing
interface f0/0
ipv6 rip one enable
ipv6 address 2000::1/64
interface s0/0/0
ipv6 address 2001::/64 eui-64
ipv6 rip one enable
```

- A. RIPng will be enabled on no interfaces.
- B. RIPng will be enabled on one interface.
- C. RIPng will be enabled on two interfaces.
- D. RIPng will advertise about prefixes connected to S0/0/0 and F0/0, but only send Updates on one interface.

Answer: B

Explanation:

That the configuration will be copied/pasted into a router means that the order of the commands matters. In this case, that the `ipv6 rip one enable` command precedes the `ipv6 address` command on interface `f0/0` means that IOS will reject the first of these commands, therefore not enabling RIPng on `F0/0`. The correct order listed under `S0/0/0` means that RIPng will be enabled on `S0/0/0`. As a result, RIPng on R1 will advertise about `S0/0/0`'s connected IPv6 prefixes, and send Updates on `S0/0/0`, but will do nothing related for `F0/0`.

QUESTION NO: 449

Router R1 currently has no configuration related to IPv6 or IPv4. The following configuration exists in a planning document intended to be used to copy/paste into Router R1 to enable EIGRP for IPv6 on interfaces F0/0 and S0/0/0. No other related configuration exists. Assuming F0/0 and S0/0/0 reach an up/up state, which of the following is true about EIGRP for IPv6 on R1 after this configuration has been pasted into R1?

```
ipv6 router eigrp 1
```

```
ipv6 unicast-routing
interface f0/0
ipv6 address 2000::1/64
ipv6 eigrp 1
interface s0/0/0
ipv6 address 2001::/64 eui-64
ipv6 eigrp 1
```

- A. EIGRP works on F0/0 and S0/0/0 without further configuration.
- B. EIGRP works with the addition of one command: a no shutdown command in EIGRP router configuration mode.
- C. EIGRP works with the addition of one command: an eigrp router-id command in EIGRP router configuration mode.
- D. EIGRP for IPv6 needs at least two more configuration commands before it works on R1.

Answer: D

Explanation:

Because the question states that no IPv4 configuration exists, EIGRP for IPv6 cannot derive a 32-bit EIGRP router ID. Before EIGRP will work, R1 needs to define an EIGRP router ID (using the eigrp router-id command) and enable EIGRP using the no shutdown router subcommand.

QUESTION NO: 450

Router R1 connects to Router R2 over an Ethernet LAN with both routers using their F0/0 interfaces. R1 learns a route from R2 using EIGRP for IPv6. That route lists F0/0 as the outgoing interface with R2 as the next hop. The configuration excerpt shows all relevant configuration on R2's F0/0 interface. Which of the following is true about R1's route?

```
interface f0/0
mac-address 1111.1111.1111
ipv6 address 2000::/64 eui-64
ipv6 address 2001::1/64
```

- A. The next hop is 2000::1311:11FF:FE11:1111
- B. The next hop is FE80::1311:11FF:FE11:1111
- C. The next hop is FE80::5111:11FF:FE11:1111
- D. The next hop is 2001::1

Answer: B

Explanation:

EIGRP uses the link local address as the next hop for routing protocols. Based on R2's MAC address, R2's link local address on F0/0 will be FE80::1311:11FF:FE11:1111. This value is derived by splitting the MAC, inserting FFFE, and flipping bit 7, making the initial hex 11 become hex 13.

QUESTION NO: 451

Which of the following are true of both OSPFv2 and OSPFv3? (Choose two.)

- A. The method of choosing an OSPF router ID
- B. Verification checks that must be validated before two routers can become OSPF neighbors
- C. Support for route tags
- D. Support for multiple instances per interface

Answer: A,C

Explanation:

OSPFv3 supports multiple OSPF instances per interface, whereas OSPFv2 does not. Also, each version requires a different set of requirements be met before becoming neighbors, most notably that OSPFv3 does not require neighboring OSPFv3 routers to be in the same subnet.

QUESTION NO: 452

Router R1 currently has no configuration related to IPv6 or IPv4. The following configuration exists in a planning document, intended to be used to copy/paste into Router R1 to enable OSPFv3 on interfaces F0/0 and S0/0/0. No other related configuration exists. Assuming F0/0 and S0/0/0 reach an up/up state, which of the following is true about OSPFv3 on R1 after this configuration has been pasted into R1?

```
ipv6 router ospf 1
ipv6 unicast-routing
interface f0/0
ipv6 address 2000::1/64
ipv6 ospf 1 area 1
interface s0/0/0
ipv6 address 2001::/64 eui-64
ipv6 ospf 1 area 0
```

- A. OSPF works on F0/0 and S0/0/0 without further configuration.
- B. OSPF works with the addition of one command: a no shutdown command in OSPF router configuration mode.

- C. OSPF works with the addition of one command: an router-id command in OSPF router configuration modE.
- D. OSPFv3 needs at least two more configuration commands before it works on R1.

Answer: C

Explanation:

Because the question states that no IPv4 configuration exists, OSPFv3 cannot derive a 32-bit OSPF router ID. Before OSPFv3 works, R1 needs to define an OSPF router ID (using the OSPF router-id command).

QUESTION NO: 453

The following output occurs on Router R1, which runs both EIGRP for IPv6 and OSPFv3, with redistribution from EIGRP into OSPF configured with the redistribute eigrp 1 metric 25 command. Interface S0/0/1 has been enabled for EIGRP ASN 1. Which of the following should be true of redistribution in this case?

```
D 2000::/64 [90/1422516]
via FE80::213:19FF:FE7B:5026, Serial0/0/1
C 2000:0:0:1::/64 [0/0]
via Serial0/0/1, directly connected
L 2000:1:213:19FF:FE7B:5004/128 [0/0]
via Serial0/0/1, receive
```

- A. Route 2000::/64 will be redistributed.
- B. Route 2000:0:0:1::/64 will be redistributed.
- C. Route 2000:1:213:19FF:FE7B:5004/128 will be redistributed.
- D. No routes will be redistributed because of the omission of the subnets parameter of the redistribute command.

Answer: A

Explanation:

The redistribute command does not have a subnets option for IPv6 because OSPFv3 has no concept of IPv6 classful networks nor their subnets. As configured, the redistribute command redistributes only EIGRP-learned routes. If the includeconnected parameter had been included, the connected route for 2000:0:0:1::/64 would have also been redistributed. Local routes are never redistributed.

QUESTION NO: 454

Router R1 has been configured with an ipv6 route 2000::/64 S0/0/0 64 command. Which of the following does the 64 at the end of the command represent?

- A. Metric
- B. Administrative distance
- C. Timeout (seconds)
- D. Prefix length
- E. Interface ID

Answer: B

Explanation:

The only configurable item after the interface that does not first list a keyword is the administrative distance parameter.

QUESTION NO: 455

An enterprise has plans to start adding IPv6 support. For the first year, the IPv6 will be in small pockets spread around the existing large IPv4 network, with occasional IPv6 traffic while applications teams test IPv6-enabled servers and applications. Which of the following tools would be most appropriate?

- A. Native IPv6
- B. Point-to-point tunnels
- C. Multipoint tunnels
- D. NAT-PT

Answer: C

Explanation:

Native IPv6 makes the most sense when the IPv6 deployment is pervasive, with traffic loads heavy or at least steady. Point-to-point tunnels work best when IPv6 is needed in only a subset of sites but also when the traffic should be somewhat regular or with higher volume. Multipoint tunnels also work well when IPv6 is needed in a subset of sites, but it is more appropriate when the traffic is more occasional and lower volume. Finally, NAT-PT is useful when an IPv4-only host needs to communicate with an IPv6-only host.

QUESTION NO: 456

An enterprise has plans to start adding IPv6 support. The initial deployment requires support from some IPv6-only devices that need to access servers that support only IPv4. Which of the following tools would be most appropriate?

- A. Native IPv6
- B. Point-to-point tunnels
- C. Multipoint tunnels
- D. NAT-PT

Answer: D

Explanation:

Native IPv6 makes the most sense when the IPv6 deployment is pervasive, with traffic loads heavy or at least steady.. Point-to-point tunnels work best when IPv6 is needed in only a subset of sites, but also when the traffic should be somewhat regular or with higher volume. Multipoint tunnels also work well when IPv6 is needed in a subset of sites, but it is more appropriate when the traffic is more occasional and lower volume. Finally, NAT-PT is useful when an IPv4-only host needs to communicate with an IPv6-only host.

QUESTION NO: 457

A client host uses IPv4 to communicate with one server and IPv6 to communicate with another. Which of the following IPv6 coexistence features is likely at work on the host?

- A. Native IPv6
- B. Point-to-point tunnels
- C. Multipoint tunnels
- D. NAT-PT
- E. Dual stacks

Answer: E

Explanation:

Dual stacks means that the host runs both IPv4 and IPv6. A host must run both if the host is to send packets of each protocol. The host may use a multipoint tunnel, but the other three answers list features applicable to routers, but not hosts.

QUESTION NO: 458

The following configuration exists on a router on one end of an IPv6 tunnel. Which type of tunnel is created by this configuration?

```
interface loopback 1
ip address 1.1.1.1 255.255.255.255
interface tunnel 2
ipv6 address 2000::1::/64
```

```
tunnel source loopback 1
tunnel destination 2.2.2.2
tunnel mode ipv6ip
ipv6 eigrp 1
```

- A. Automatic 6to4
- B. Manually configured tunnel
- C. ISATAP
- D. GRE

Answer: B

Explanation:

A manually configured tunnel explicitly defines the destination IPv4 address, as does a GRE IPv6 tunnel. The other two types listed in the answer do not. Additionally, a manually configured tunnel uses mode ipv6ip , whereas GRE tunnels use mode gre.

QUESTION NO: 459

An engineer is reviewing another engineer's sample configuration for a GRE tunnel used to pass IPv6 traffic. The tunnel has not yet been configured on the router. Which of the following commands is not required for the configuration to pass IPv6 traffic?

- A. tunnel source
- B. tunnel destination
- C. tunnel mode
- D. All these commands are required.

Answer: C

Explanation:

IOS defaults to use GRE encapsulation mode on tunnel interfaces, so the tunnel mode command is not required. The mode will default to GRE. The tunnel source and tunnel destination commands are required.

QUESTION NO: 460

Which of the following IPv6 tunneling mechanisms support IPv6 IGP routing protocols? (Choose two.)

- A. Automatic 6to4
- B. Manually configured tunnel

- C. ISATAP
- D. GRE

Answer: B,D

Explanation:

The two point-to-point tunneling methods-manually configured tunnel and GRE-support IPv6 IGPs. The two multipoint tunneling methods-automatic 6to4 and ISATAP-do not.

QUESTION NO: 461

The following configuration exists on a router on one end of an IPv6 tunnel. Although the configuration added so far is correct, the configuration is incomplete.

Which type of tunnel is most likely to be intended by the network engineer?

```
interface loopback 1
ip address 192.168.1.1 255.255.255.255
interface tunnel 2
ipv6 address 2002:C0A8:101::1/64
tunnel source loopback 1
```

- A. Automatic 6to4
- B. Manually configured tunnel
- C. ISATAP
- D. GRE

Answer: A

Explanation:

An automatic 6to4 tunnel does not use a tunnel destination command on the tunnel interface; ISATAP tunnels also do not use this command. However, automatic 6to4 tunnels use IPv6 addresses that begin 2002::/16 and have the tunnel's source IPv4 address imbedded as the second and third octets of the IPv6 address, whereas ISATAP tunnels do not. Because C0A5:101 hex equals 192.168.1.1 in dotted decimal, this configuration represents the almost completed configuration for an automatic 6to4 tunnel.

QUESTION NO: 462

The answers each list a tunnel method and two consecutive IPv6 address quartets. Which answers identify a tunneling method that relies on an IPv4 address to be embedded into an

IPv6 address, within the correct quartets listed? (Choose two.)

- A. Automatic 6to4, quartets 2 and 3
- B. Automatic 6to4, quartets 7 and 8
- C. ISATAP, quartets 2 and 3
- D. ISATAP, quartets 7 and 8

Answer: A,D

Explanation:

ISATAP uses a modified EUI-64 format, which adds the IPv4 address, in hex, into quartets 7 and 8. Automatic 6to4 tunnels use address range 2002::/16, with the next two quartets (second and third quartets) used to store the hex version of an IPv4 address.

QUESTION NO: 463

Router R1 uses MAC address 1111.1111.1111 for its Fa0/0 interface. An engineer sees the following configuration in the output of a show running-config command. Then, the engineer issues a show ipv6 interface brief command. What global unicast IPv6 address does this command display for interface tunnel 1?

```
interface loopback 1
ip address 192.168.1.1 255.255.255.255
interface tunnel 1
tunnel source loopback 1
tunnel destination 192.168.1.2
tunnel mode ipv6ip isatap
ipv6 address 2000::/64 eui-64
```

- A. 2000::1311:11FF:FE11:1111
- B. 2000::C0A5:101
- C. 2000:C0A5:101::
- D. 2000::5EFE:C0A5:101

Answer: D

Explanation:

The combination of the configured eui-64 parameter on the ipv6 address command, and the tunnel mode of isatap , tells the router to use modified EUI-64 rules. These rules start with the configured 64 bit prefix (2000::/64 in this case), adding 0000:5EFE as the fifth and sixth quartets. The last two quartets are taken from the tunnel source command's referenced IPv4 address. In this case, 192.168.1.1 converts

to C0A8:0101, making the last answer correct.

QUESTION NO: 464

Router R1 sits at an Enterprise branch office, using the Internet for its only connectivity back to the rest of the Enterprise. Which of the following is not a benefit of using an IPsec tunnel for packets sent through the Internet, between R1 and the rest of the Enterprise?

- A. Privacy
- B. Authentication
- C. Allows using an IGP between R1 and the Enterprise
- D. Secure communications

Answer: C

Explanation:

IPsec tunnels make for more secure communications, including encryption and authentication. However, it does not support IGP communications across the tunnel.

QUESTION NO: 465

Router R1 sits at an Enterprise branch office, using both the Internet and a leased line to another Enterprise router for its two connectivity options back into the rest of the Enterprise network. The engineer planning for this branch decided to use the leased line for all Enterprise traffic, unless it fails, in which case the Internet connection should be used to pass traffic to the Enterprise. Which of the following is most likely to be useful on the branch router? (Choose two.)

- A. IPsec tunnel
- B. GRE tunnel
- C. Floating static route
- D. An IGP

Answer: A,C

Explanation:

An IPsec tunnel would be useful to allow the packet to pass over the Internet and into the Enterprise. The GRE tunnel would only be needed if an IGP is also needed, and for this design, an IGP is not required. Instead, a floating static default route would work fine, with the static route sending traffic over the IPsec tunnel but only when the private leased line fails.

QUESTION NO: 466

Router R1, a branch router, connects to the Internet using DSL. The engineer plans to use a configuration with a dialer interface. The answers list a feature and interface on which the feature could be configured. Which combinations accurately describe the interface under which a feature will be configured?

- A. PPP on the ATM interface
- B. VPI/VCI on the dialer interface
- C. IP address on the ATM interface
- D. CHAP on the dialer interface

Answer: D

Explanation:

The ATM details, like VPI/VCI, will be configured under the ATM interface. PPP (including CHAP) and Layer 3 details will be configured under the dialer interface.

QUESTION NO: 467

Router R1, a branch router, connects to the Internet using DSL. Some traffic flows through a GRE and IPsec tunnel, over the DSL connection, and into the core of an Enterprise network. The branch also allows local hosts to communicate directly with public sites in the Internet over this same DSL connection. Which of the following answers defines how the branch NAT config avoids performing NAT for the Enterprise directed traffic but does perform NAT for the Internet-directed traffic?

- A. By not enabling NAT on the IPsec tunnel interface
- B. By not enabling NAT on the GRE tunnel interface
- C. By configuring the NAT-referenced ACL to not permit the Enterprise traffic
- D. By asking the ISP to perform NAT in the cloud

Answer: C

Explanation:

The NAT configuration acts only on packets permitted by a referenced ACL. As a result, the ACL can permit packets destined for the Internet, performing NAT on those packets. The ACL also denies packets going to the Enterprise, meaning that the router does not apply NAT to those packets.

QUESTION NO: 468

Router R1, a branch router, connects to the Internet using DSL. Some traffic flows through a GRE and IPsec tunnel, over the DSL connection, destined for an Enterprise network. Which of the following answers best describes the router's logic that tells the router, for a given packet, to apply GRE encapsulation to the packet?

- A. When the packet received on the LAN interface is permitted by the ACL listed on the tunnel gre acl command under the incoming interface
- B. When routing the packet, matching a route whose outgoing interface is the GRE tunnel interface
- C. When routing the packet, matching a route whose outgoing interface is the IPsec tunnel interface
- D. When permitted by an ACL that was referenced in the associated crypto map

Answer: B

Explanation:

As for the correct answer, the process of routing a packet out a GRE tunnel interface triggers the GRE encapsulation action. As for the incorrect answers: There is no tunnel gre acl command. There is no IPsec tunnel interface. Finally, one answer refers to logic that would describe a router's logic when determining whether to encapsulate a packet into an IPsec tunnel.

QUESTION NO: 469

Which of the following must match for two directly connected routers running OSPF to establish a neighbor adjacency? (Select 2 choices.)

- A. area IDs
- B. process IDs
- C. router IDs
- D. hello timers

Answer: A,D

QUESTION NO: 470

An engineer is trying to summarize the following networks using the "ip summary-address eigrp" command:

- 10.8.88.0/25
- 10.8.89.48/29
- 10.8.64.96/27

Which network and subnet mask below would be the smallest EIGRP summary address to include all three subnets?

Choose the best answer.

- A. 10.8.64.0 255.255.224.0
- B. 10.8.64.0 255.255.128.0
- C. 10.8.64.0 255.255.192.0
- D. 10.8.0.0 255.255.192.0

Answer: A

QUESTION NO: 471

Look at the following configuration below and select the statements that are true.

```
RD2(config)#router ospf 100
RD2(config-router)# distance 180 10.1.10.1 0.0.0.0 match-me
RD2(config-router)#!
RD2(config-router)#ip access-list standard match-me
RD2(config-std-nacl)# permit host 172.16.0.0
```

Choose two.

- A. OSPF will have an AD of 180 for routes from 10.1.10.1 in the 172.16.0.0/16 range.
- B. The neighbor RID is 10.1.10.1
- C. OSPF will have an AD of 180 for all routes.
- D. The local RID is 10.1.10.1

Answer: A,B