300-101.exam

Number: 300-101
Passing Score: 800
Time Limit: 120 min
File Version: 1.0

Cisco

300-101

Implementing Cisco IP Routing (ROUTE)

Version 1.0

Exam A

QUESTION 1
You have a router that is running both OSPF and RIP. You have configured this router to perform mutual redistribution between the two protocols. The following conditions exist:

- The S0/0 interface, which is configured for RIP, is routing for the 172.16.5.0/24 network.
- The S0/1 interface, which is configured for OSPF, is routing for the 172.16.6.32/28 network.

Users in the RIP domain are unable to connect to devices in the OSPF domain.

What must be done to allow the OSPF routes to be redistributed into the RIP domain? (Choose two. Each correct answer is part of the solution.)

A. Create a static route that points to 172.16.6.0/24 with a next hop of null0.
B. Execute the passive-interface command on S0/0.
C. Create a loopback address on the router
D. Redistribute static routes into RIP.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The OSPF domain has a different mask than the RIP domain, and they are on the same major network. The OSPF domain's mask is also longer than the RIP domain's mask. Therefore, the RIP domain will not advertise routes learned from OSPF and redistributed into RIP. To solve this problem, you can create a static route to the major (classful) network 172.16.6.0/24, which includes all of the subnets in the OSPF domain, set the destination as null0, and then redistribute static routes into RIP. The following commands would enable this process:

```
router1(config)# ip route 172.16.5.0 255.255.255.0 null0
router1(config)# router rip
router1(config-router)# redistribute static
router1(config-router)# default metric 1
```

You should include the metric as well to ensure redistribution. This will allow the 172.16.5.0/24 network to be advertised to the RIP domain and, when the frames arrive at the null0 interface, will ensure the routing table of the router will have routes to the specific subnets of the OSPF domain.

You should not execute the passive-interface command. This would prevent the interface from advertising either RIP or OSPF routes, and would only allow RIP updates inbound. This would not solve the problem and will create additional problems when the router is unable to advertise RIP routes to the other routers in the RIP domain.

You should not create a loopback address on the router. Loopback addresses are logical addresses that can be created and used as the source of routing updates. Under normal circumstances, if routing updates are sourced from a physical interface and the interface goes down, the route will be removed from the routing tables. Since a loopback interface cannot go down, it provides the advantage of keeping a route in the tables even if the physical interface that services the route goes down. Loopback interfaces are of no help in solving the redistribution problem.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes >

Redistributing Between Classful and Classless Protocols: EIGRP or OSPF into RIP or IGRP

QUESTION 2
You are the network administrator for a large software organization. You designed the LAN in the organization's main building for connecting the internal LAN to a WAN as shown below:

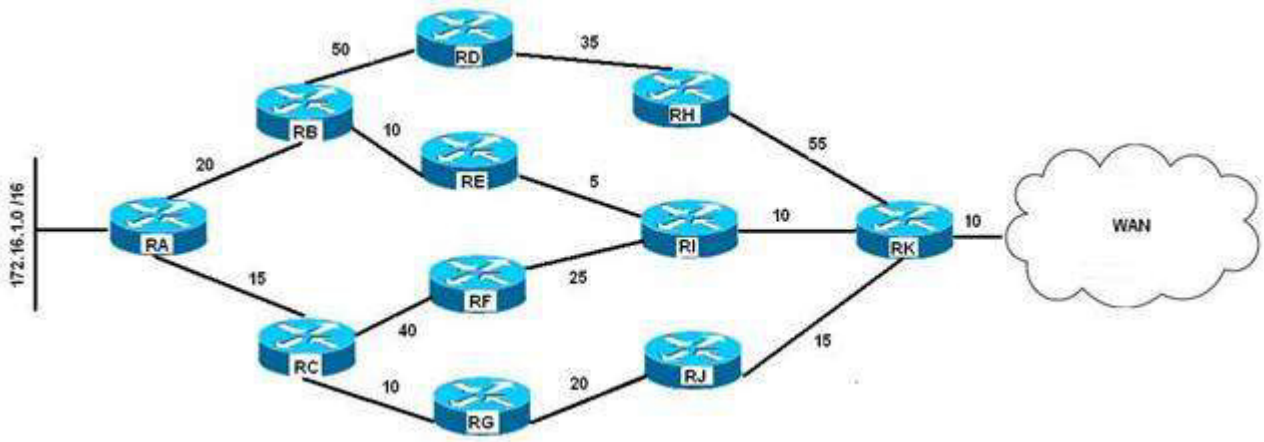

You have configured EIGRP with the variance parameter set to 3 on all the routers to enable unequal load balancing from the 172.16.1.0 network to the WAN. The delay configured on each of the routers is shown in the LAN diagram, and the K values are set as follows:

K1 = 0
K2 = 0
K3 = 1
K4 = 0
K5 = 0

Which of the following paths are entered into the routing tables as a result of the unequal load balancing configured on the routers? (Choose all that apply.)

A. RA-RB-RD-RH-RK
B. RA-RB-RE-RI-RK
C. RA-RC-RF-RI-RK
D. RA-RC-RG-RJ-RK

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The only path is entered in the routing table as a result of the unequal load balancing configured on the routers:

RA-RB-RE-RI-RK

In EIGRP networks, bandwidth and delay are the default factors for calculating the metric/cost for a given route. Additional factors such as load and reliability can be considered in the computation of the EIGRP metric, as given in the following formula:

Metric = [K1 * bandwidth + (K2 * bandwidth) / (256 - load) + K3 * delay] * [K5 / (reliability + K4)]

In this case, only the K3 value has a non-zero value. This implies that only delay is taken into consideration to calculate the metric of the shortest path from 172.16.1.0 network to the WAN. The path with the lowest metric,

which is delay in this scenario, is the shortest path, and is therefore entered automatically in the routing table. The total delay and the corresponding metric for the three best paths are given as follows:

Path	Total Delay	Scaled EIGRP Delay	Metric
RA-RB-RE-RI-RK	55	14080	14080
RA-RB-RD-RI-RK	100	25600	25600
RA-RC-RG-RJ-RK	70	17920	17920

In the given table, the path RA-RB-RE-RI-RK has the lowest metric of 14080. This is the shortest path, so it would be entered in the routing table even if variance were not enabled. In this scenario variance is set to 3, which enables unequal load balancing among those paths that have a metric less than three times the least metric for the given route. Three times the least metric in this scenario is 42240 (14080 x 3). This implies that paths between the 172.16.1.0 network and the WAN having a metric less than 42240 participate in the load balancing. On metric values alone, those paths would appear in the routing tables. However, to be eligible to be a feasible successor the reported distance of the path must be less than the feasible distance (current best path). None of the paths, with the exception of RA-RB-RE-RI-RK meet that requirement.

The path RA-RB-RD-RH-RK is not entered in the routing table as a result of the unequal load balancing. The scaled EIGRP delay for this path is 43520 (170 x 256), which is more than three times the least metric available from the 172.16.1.0 network to the WAN (42240). In addition, the reported distance for this path is more than the feasible distance. Therefore, the path RA-RB-RD-RH-RK is not used for balancing the load from the 172.16.1.0 network to the WAN and does not appear in the routing tables.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP load balancing

References:
Cisco > Support > Technology Support > IP > IP Routing > Design > Design Technotes > How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP? > Document ID: 13677
Cisco > Support > Technology Support > IP > IP Routing > Design > Design Technotes > How Does Load Balancing Work? > Document ID: 5212
Cisco > Support > Technology Support > IP > IP Routing > Technology Information > Technology Whitepaper > Enhanced Interior Gateway Routing Protocol > Document ID: 16406 > Feasible Distance, Reported Distance, and Feasible Successor

QUESTION 3
Examine the exhibit by pressing the Exhibit(s) button.

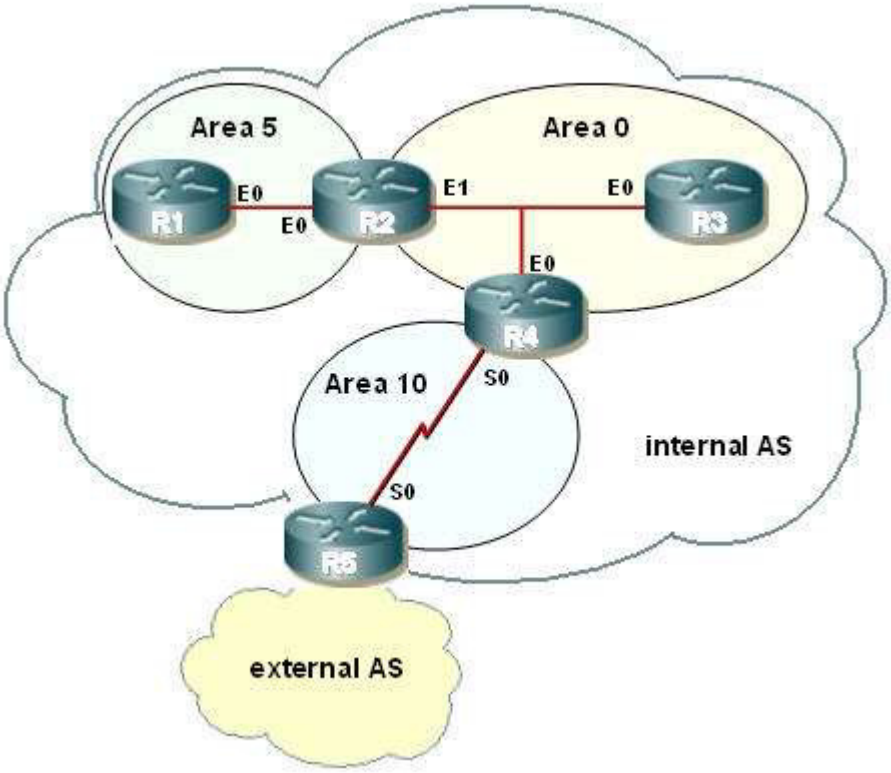

You are to configure R1 to belong to area 5. This area does not accept routes from the external AS or summary routes from any other internal areas. Refer to the IP addressing below.

R1 - int E0 - 192.168.5.1/24
R2 - int E0 - 192.168.5.2/24
R2 - int E1 - 192.168.0.2/24
R3 - int E0 - 192.168.0.3/24

Which configuration commands are required to correctly configure R1?

A. R1(config)# router ospf 10
R1(config-router)# area 5 no-summary stub
R1(config-router)# network 192.168.5.0 0.0.0.255 area 5
B. R1(config)# router ospf 5
R1(config-router)# area 5 stub
R1(config-router)# network 192.168.5.0 0.0.0.255 area 5
C. R1(config)# router ospf 10
R1(config-router)# area 5 stub
R1(config-router)# network 192.168.5.0 255.255.255.0 area 5
D. R1(config)# router ospf 5
R1(config-router)# area 5 stub no-summary
R1(config-router)# network 192.168.5.0 255.255.255.0 area 5

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
All routers within a stub area must be configured as stub, or adjacencies will not form. Besides the command to enable OSPF and the command to identify the area, the only other required command identifies the area as a

stub. At the area border router (ABR), R2, the no-summary keyword is required. The following commands are required to configure R1:

R1(config)# router ospf 5
R1(config-router)# area 5 stub
R1(config-router)# network 192.168.5.0 0.0.0.255 area 5

A totally stubby area does not accept any external network LSAs (Type 5) or any inter-area summary LSAs (Types 3 and 4) from entering the area. Use the area stub command with the no-summary keyword on the ABR only to configure a totally stubby area.

The correct syntax for the area stub command is shown below:

Router(config-router)# area area-id stub [no-summary]

Note that the optional no-summary keyword is used only on ABRs to block summary link advertisements into the stub area. This option creates a totally stubby area. It is very important to configure the command consistently on all routers within the area. OSPF sends its stub status (on or off) in its hello packets.

If two neighbors have conflicting stub status, they will not form an adjacency, and you end up with no OSPF communication over that link.

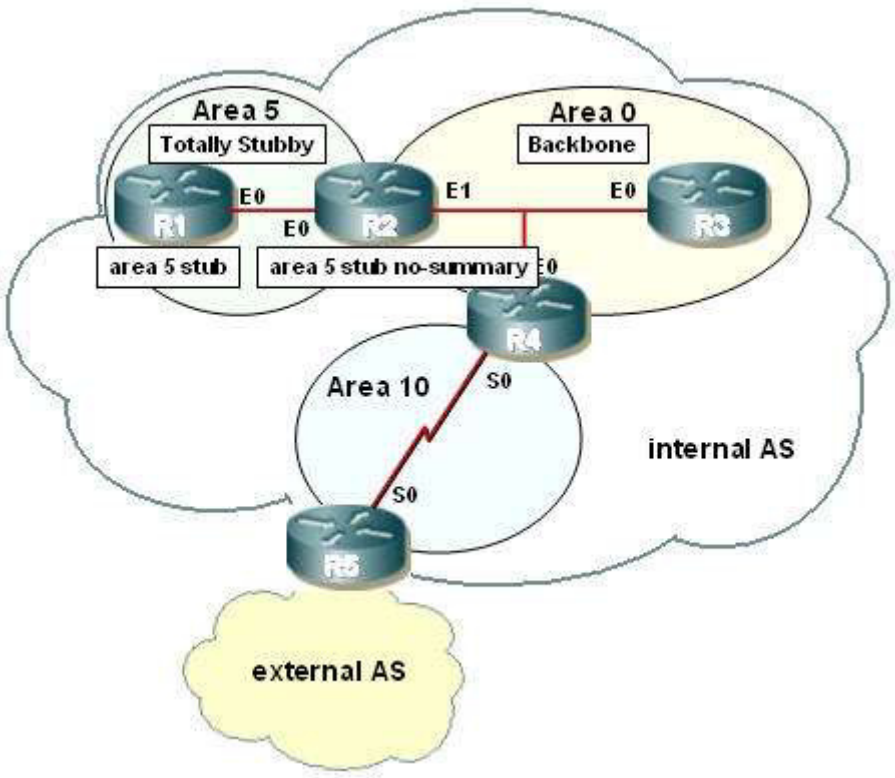

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco > Home > Support > Technology Support > IP Routing > Design > Design Technotes > What Are OSPF Areas and Virtual Links? > What Are Areas, Stub Areas, and Not-So-Stubby Areas?

Cisco IOS Master Command List, Release 12.4 > a through b > area stub

QUESTION 4
Which of the following commands will display information about Type 4 LSAs?

A. show ip ospf database external
B. show ip ospf database asbr-summary
C. show ip ospf database summary
D. show ip ospf database router

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The command show ip ospf database asbr-summary will display information about Type 4 LSAs. These LSAs provide next-hop information for areas that are receiving Type 5 or external LSAs. Consider the following sample output of the show ip ospf database asbr-summary command:

```
Router# show ip ospf database asbr-summary
OSPF Router with id(192.168.239.128) (Process ID 566)
Displaying Summary ASB Link States(Area 0.0.0.0)
LS age: 1463
Options: (No TOS-capability)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 172.16.245.63 (AS Boundary Router address)
Advertising Router: 172.16.241.75
LS Seq Number: 80000072
Checksum: 0x3548
Length: 28
Network Mask: 0.0.0.0 TOS: 0 Metric: 1
```

The output shows that the router that sent this LSA is at 172.16.241.75. The router functioning as the ASBR is at 172.16.245.63. The advertising router, located at 172.16.241.75, is broadcasting that its best metric to reach the ASBR at 172.16.254.63 is 1.

The command show ip ospf database external will not display information about Type 4 LSAs. It will display information about Type 5 LSAs, or External Link LSAs, instead of ASBR summary links, which are Type 4 LSAs.

The command show ip ospf database summary will not display information about Type 4 LSAs. It will display information about summary links, or Type 3 LSAs, that are generated by an ABR, not summary links generated by an ASBR.

The command show ip ospf database router will not display information about Type 4 LSAs. It will display information about router links, or Type 1 LSAs, instead of ASBR summary links, which are Type 4 LSAs.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe OSPF packet types

References:
Cisco > Cisco IOS IP Routing: OSPF Command Reference > show ip ospf database

QUESTION 5
Which command should be executed on all ABRs in an area to configure it as a totally stubby area?

A. Router(config-router)# area process-id stub [no-summary]
B. Router(config-router)# area area-id [no-summary] stub
C. Router(config-router)# area area-id stub [no-summary]
D. Router(config-ospf)# area router-id [no-summary] stub

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct syntax for the area stub command to configure a totally stubby area is shown below:

Router(config-router)# area stub [no-summary]

Note that the optional no-summary keyword is used only on area border routers (ABRs) to block summary link advertisements into the stub area. This option creates a totally stubby area. All internal routers in the area need only the stub keyword without the no summary keyword.

It is very important to configure the command consistently on all routers within the area. OSPF sends its stub status (on or off) in its hello packets. If two neighbors have conflicting stub status, for example, if one indicates that a stub is present and the other indicates that no stub is present, they will not form an adjacency, and you end up with no OSPF communication over that link.

The other options are either using incorrect syntax or being executed at an incorrect prompt. The area stub command should be executed at the OSPF router configuration prompt.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco > Home > Support > Technology Support > IP Routing > Design > Design Technotes > What Are OSPF Areas and Virtual Links? > What Are Areas, Stub Areas, and Not-So-Stubby Areas?
Cisco > Cisco IOS IP Routing: OSPF Command Reference > area stub

QUESTION 6
You have configured a BGP network with several routers in the same autonomous system (AS). There are three possible routes from router A to router B in the network. The following conditions exist:
- All three routes have the same weight
- All three routes were originated locally through the use of the network command
- The bgp default local-preference 50 command is executed for all three routes
- All three routes have different lists of AS through which they travel

Which of the following parameters is used to select the best path among the three routes?

A. Weight
B. MED
C. LOCAL_PREF
D. AS_Path

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
The AS_Path parameter is used to select the best path among the three routes. To select the best path from router A to router B, BGP analyses various BGP attributes that are set during the configuration of the network. The key BGP attributes and the order in which they are checked are as follows:

1. Weight - highest weight is selected
2. LOCAL_PREF - highest LOCAL_PREF is selected
3. Locally originated routes - local routes are selected
4. AS_PATH - shortest AS_PATH is selected
5. Origin type - lowest origin type is selected
6. Multi-exit Discriminator (MED) - lowest MED is selected

Because the weight attribute is same for all three routes, BGP checks the value of the LOCAL_PREF attribute. However, this attribute is also same for the three routes because the bgp default local-preference 50 command was executed for the three routes, which sets the value of the LOCAL_PREF attribute to 50 for those routes.

BGP then checks whether any of the routes were locally originated using either the network or aggregate commands. As stated in the scenario, all three routes were locally originated with the network command during BGP configuration. Consequently, BGP analyzes the value of the AS_PATH attribute. This attribute refers to a list of AS numbers that are traversed by a particular route. The route with the shortest AS_PATH is selected as the best path.

The weight attribute is not used to select the best path in this case. The weight attribute for all three routes is the same. If this attribute were different for the three routes, then the route with the highest weight would be considered the best path.

The MED attribute is not used to select the best path in this case. The MED, or multi-exit discriminator, specifies the route into an AS that has more than one entry points. A route with the lowest MED is selected as the best path. However, in this case, the MED attribute is not considered because the AS_PATH attribute is different for the three routes. If the AS_PATH attribute for the three routes were the same, then the MED attribute would have been considered.

The LOCAL_PREF attribute is not used to select the best path. The LOCAL_PREF attribute is checked if the weight attribute for the routes is same. The LOCAL_PREF attribute refers to the local preference, which specifies the route that has preference to exit the AS for a given destination network. The route with the highest LOCAL_PREF value is selected as the best path. However, the bgp default local-preference 50 command was executed for all three routes. Hence, this attribute is not considered to select the best path between the BGP routers A and B.

Objective:
Layer 3 Technologies
Sub-Objective:
Explain BGP attributes and best-path selection

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design TechNotes > BGP Best Path Selection Algorithm

QUESTION 7
Examine the following output of the show ip ospf interface command.

```
Router43# show ip ospf interface brief
Interface PID Area IP Address/Mask Cost State Nbrs F/C
Se0/0/0 1 0 172.16.1.1/30 50 P2P 1/1
Fa0/0 1 0 10.0.0.5/24 1 BDR 1/1
Fa0/1 1 11 10.1.2.1/24 1 DR 0/0
```

What would be the effect of executing the auto-cost reference bandwidth 2000 command on Router43 in router OSPF mode?

A. the cost of the Serial interface would increase to 20
B. the cost of the FastEthernet interfaces would increase to 2000
C. the cost of the Serial interface would increase to 647
D. the cost of the FastEthernet interfaces would increase to 20

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If the auto-cost reference bandwidth 2000 command is executed in router OSPF mode it will result in a cost to the FastEthernet interfaces of 20. The formula for arriving at the cost is:

reference bandwidth / interface bandwidth = cost

The default reference bandwidth for FastEthernet is 100 Mbps. If the reference bandwidth is set at 2000 Mbps using the auto-cost reference command, and the FastEthernet interface has a bandwidth of 100 Mbps, the resulting cost is 20 (2000 / 100 = 20).

The auto-cost reference bandwidth command is executed in router OSPF mode to affect all interfaces. Alternatively, the cost of each interface can be set separately with the ip ospf cost command issued in interface configuration mode. The two commands can also be used in combination: you can set all interfaces with the auto-cost reference bandwidth command, and then set a single interface to a different cost with the ip ospf cost command.

The command would not result in the cost of the Serial interface increasing to 20 or to 647. With a reference bandwidth of 2000 Mbps and interface bandwidth of 1544 kbps (the default bandwidth of a serial interface), the resulting cost would be 1294.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF path preference

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > OSPF Cost
Cisco > Cisco IOS IP Routing: OSPF Command Reference > show ip ospf interface
Cisco > Cisco IOS IP Routing: OSPF Command Reference > auto-cost

QUESTION 8
You instructed your assistant to configure redistribution of OSPF routes into EIGRP on Router 9. The routes are not being advertised to EIGRP and you are troubleshooting the problem. The EIGRP process ID is 100 and the OSPF process ID is 20. When you ask your assistant what commands were executed, you are shown the following:

Router9(config)# router eigrp 100
Router9(config-router)# redistribute ospf 20

What is the problem?

A. no metric was configured
B. the process IDs are incorrect
C. the redistribute command is executed at the interface configuration prompt
D. the redistribute command is executed at the global configuration prompt

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:
Explanation:
The problem is that the metric was not configured. Some routing protocols require that a metric be provided for the redistributed routing protocol or route redistribution will not occur successfully. RIP and EIGRP both require that a metric be provided. IS-IS and OSPF do not have this requirement.

When you redistribute traffic into EIGRP without specifying a metric, then the default metric applied is zero, the route will be treated as unreachable, and the route will not be advertised. The addition of the metric parameter as shown below would solve this issue:

Router9(config)# router eigrp 100
Router9(config-router)# redistribute ospf 20 metric 10000 100 255 1 1500

In this example, 1000 is the bandwidth, 100 is the delay, 255 is the reliability, 1 is the load, and 1500 is the MTU.

The process IDs are correct in the original scenario, and the command was executed in the correct context.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Cisco > Home > Support > Technology Support > IP > IP Version 6 > Configure > Configuration Examples and Technotes > Redistributing Routing Protocols

QUESTION 9
Refer to the following exhibit.

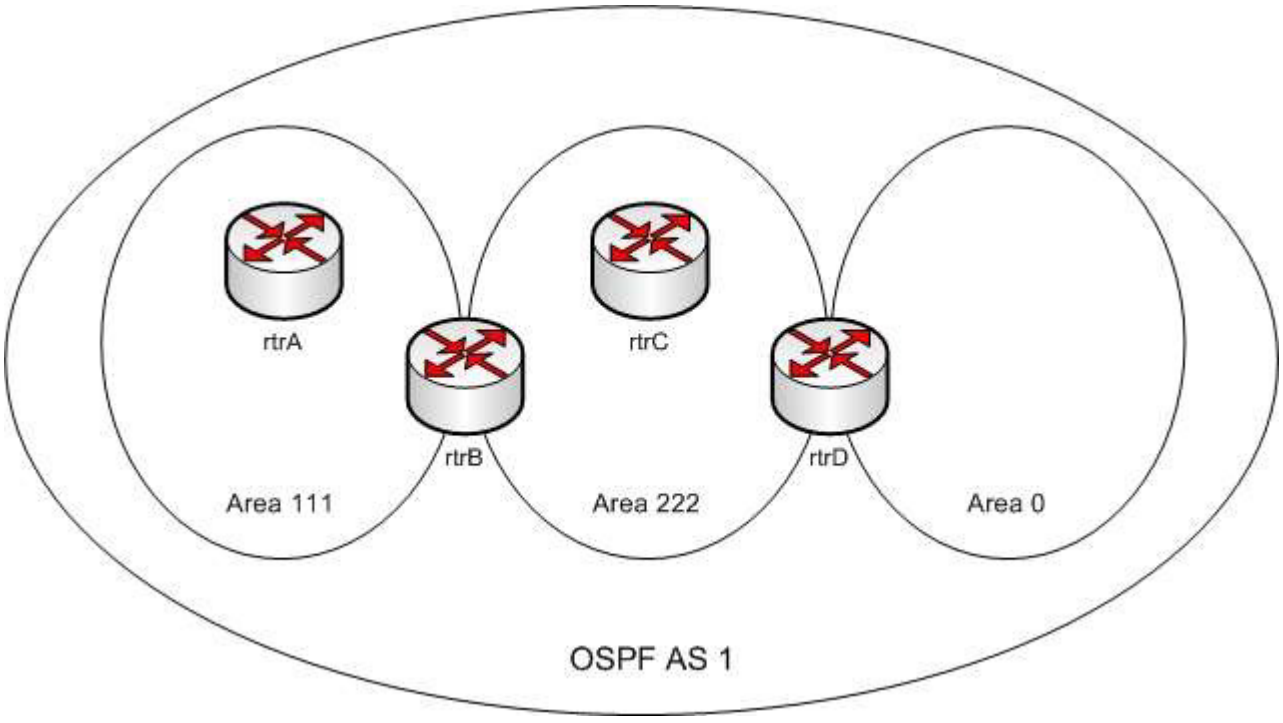

You executed the following commands on all three routers in OSPF AS 1:

- The ipv6 cef command in the global configuration mode
- The interface serial command in the global configuration mode
- The ipv6 address command in the interface configuration mode
- The ipv6 ospf area command in the interface configuration mode

You run the show ipv6 traffic command and observe that IPv6 packets are not being exchanged between the OSPF routers.

Which of the following commands should be configured on the routers to fix the problem?

A. ipv6 enable
B. ip address
C. ipv6 router ospf
D. ipv6 unicast-routing

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The ipv6 unicast-routing command should be used on all of the routers to rectify the problem. The ipv6 unicast-routing command allows the forwarding of IPv6 packets. You should execute the ipv6 unicast-routing command in the global configuration mode.

A sample configuration to enable OSPF for IPv6 on the S0/1 interface of rtrA is as follows:

```
rtrA(config)# ipv6 unicast-routing
rtrA(config)# ipv6 cef
rtrA(config)# interface S0/1
rtrA(config-if)# ipv6 enable
rtrA(config-if)# ipv6 address 2001:77A2::1:1:1/64
rtrA(config-if)# ipv6 ospf 1 area 111
```

The ipv6 enable command is not required if an IPv6 address has been configured on an interface. If it is executed with no IPv6 addresses configured, the interfaces will use the link local IPv6 addresses that each interface generates automatically.

The ip address command is not required to fix the problem because this command is used to specify an IPv4 address to a router interface. The use of this command depends on the type of tunneling mechanism used. In this case, no tunneling mechanism is being used.

The ipv6 router ospf command does not rectify the problem because this command is used to enter the router configuration mode for OSPF for IPv6. Using this command is optional and does not affect the activation of OSPF for IPv6 on the routers.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF for IPv6

References:
Cisco Press > Articles > Network Technology > General Networking > Cisco Self-Study: Implementing Cisco IPv6 Networks (IPV6) > Configuration Exercise: Configuring an IPv6
Cisco Press > Articles > Network Technology > General Networking > Cisco Self-Study: Implementing Cisco IPv6 Networks (IPV6) > Configuring IPv6 on Cisco IOS Software
Cisco IOS IPv6 Command Reference > ipv6 summary-address eigrp Through mpls ldp router-id > ipv6 unicast-routing

Cisco > Support > Technology Support > IP > IP Version 6 (IPv6) > Configure > Configuration Examples and Technotes > Sample Configuration for OSPFv3

QUESTION 10
The exhibit contains portions of RouterA's BGP configuration and IP routing table.

```
!
router bgp 65100
neighbor 192.168.12.34 remote-as 65101
network 172.16.0.0
no synchronization
auto-summary
!

RouterA# show ip route
:
o    172.16.16.0/24 [110/128] via 10.1.2.3 00:24:16, Serial0
o    172.16.24.0/24 [110/144] via 10.1.2.3 00:24:16, Serial1
:
```

Which IP network addresses, that were not learned using BGP, will be present in BGP advertisements from RouterA?

A. 172.16.0.0/16
B. 172.16.16.0/24
C. 172.16.24.0/20
D. No IGP networks will be advertised because synchronization is disabled.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The auto-summary command can affect which networks, identified by using the network command, will be advertised. Using the existing BGP configuration, the router will not announce the 172.16.16.0/24 subnet. Instead, it will announce the classful address 172.16.0.0/16 when the IP routing table maintained by the IGP contains any subnet of that classful address.

The network command directly affects what network is advertised in BGP. If the network command does not also include a network mask, and if auto-summary is enabled, the classful address of 172.16.0.0/16 is advertised any time that the router learns about a 172.16.0.0 subnet via its Interior Gateway Protocol (IGP), such as OSPF or EIGRP. In the exhibit, the routing table does contain entries of the 172.16.16.0/24 and 172.16.24.0/24 subnets that were learned by using the IGP.

If auto-summary is disabled by using the no auto-summary command, only networks in the routing table that are exact matches to the network commands are advertised. For example, to have the router announce only the 172.16.16.0/24 subnet learned via its IGP, you should alter the network command's IP address and include the subnet mask as follows:

network 172.16.16.0 mask 255.255.255.0

A combination of network statements and route statements can be used to advertise a subset of networks that exist. Examine the output shown below:

router bgp 68410

network 192.168.24.0 255.255.252.0
neighbor 172.16.8.5 remote-as 68441
ip route 192.168.24.0 255.255.252.0 null 0

The router is configured to advertise a summary route to the network 192.168.24.0 255.255.252.0. Consider the following networks:

192.168.24.0/24
192.168.25.0/24
192.168.26.0/24
192.168.32.0/24

If this router was connected to those networks, and received a packet destined for 192.168.25.1, it would successfully route the packet because the summary address (where the summarization is the result of the mask 255.255.252.0) is designed to include all of the subnets above except for 192.168.32.0/24. Therefore, all subnets except 192.168.32.0/24 will be advertised by the network and ip route statements with the summary mask.

Note: Whenever changes are made to a routing policy or to an access list that is used by a routing policy, the change will not be reflected in the routing tables of the receiving routers until the BGP session has been cleared with the clear ip bgp command.

The BGP synchronization rule specifies that networks will not be advertised or used via iBGP unless it also has been learned through an IGP. If synchronization is disabled, iBGP will advertise a network without also learning it through an IGP.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify eBGP (IPv4 and IPv6 address families)

References:
Cisco IOS Master Command List > a through b > BGP Commands: A through B > auto-summary (BGP)
Cisco > Cisco IOS IP Routing: BGP Command Reference > router bgp
Cisco > Cisco IOS IP Routing: BGP Command Reference > network (BGP and multiprotocol BGP)

QUESTION 11
Which conditions will prevent two EIGRP routers from becoming neighbors? (Choose two.)

A. Their K-values do not match.
B. Their hold times do not match.
C. Their AS numbers do not match.
D. Their hello intervals do not match.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
EIGRP routers will not become neighbors if the K-values do not match or if the autonomous system (AS) numbers do not match. They also will not become neighbors if EIGRP is not enabled for the proper networks on the local and remote routers. However, routers can become neighbors if their hello intervals and hold times do not match.

The AS number is designed to control the routers with which a router can communicate. If the AS numbers do not match, EIGRP will not exchange routes between the two routers by design and definition.

The K-values are flags that state whether a certain metric component, such as Load, is used. They must match

because they regulate how the metric values are calculated. If one router is just using bandwidth and delay to calculate its metric, and another is using bandwidth, delay, and load; they could make contradictory routing decisions that would lead to a routing loop. Because of this possibility, EIGRP requires that the K-values must match before it will allow the routers to exchange routes.

EIGRP does not require that the hello and hold times match. Although this flexibility can be helpful, it can also lead to unforeseen problems if they are accidentally mismatched. The hello interval is the amount of time in seconds to wait before sending another hello packet. The hold time is the amount of time in seconds to wait before declaring a link to be down.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP neighbor relationship and authentication

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > Introduction to EIGRP > How does EIGRP work?

QUESTION 12
Click the Exhibit(s) button to view an EIGRP network. The partial output of the show running-config command on the rtrB router is as follows:

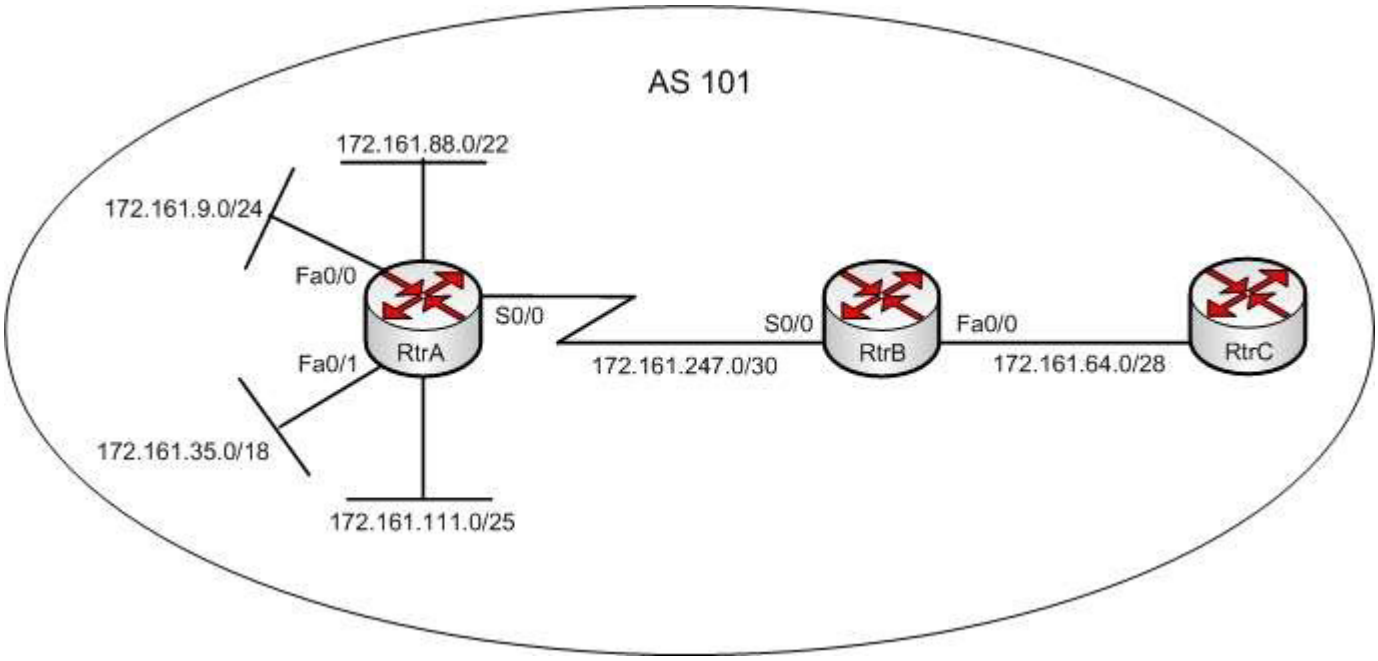

```
rtrB# show running-config
.
.
.
ip prefix-list blk_A deny 172.161.0.0/16 ge 24 le 30
ip prefix-list blk_A permit 0.0.0.0/0 le 32
.
.
router eigrp 101
network 172.161.247.0 0.0.0.255
network 172.161.64.0 0.0.0.15
distribute-list prefix blk_A out
auto-summary
!
<output omitted>
```

Which of the following subnets are blocked through the Fa0/0 interface of rtrB while sending updates to rtrC? (Choose all that apply.)

A. 172.161.9.0/24
B. 172.161.35.0/18
C. 172.161.64.0/28
D. 172.161.88.0/22
E. 172.161.111.0/25
F. 172.161.247.0/30

Correct Answer: AEF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The 172.161.9.0/24, 172.161.111.0/25 and 172.161.247.0/30 subnets are blocked through the Fa0/0 interface of rtrB while sending updates to rtrC. The following lines in the output create an IP prefix list named blk_A:

```
ip prefix-list blk_A deny 172.161.0.0/16 ge 24 le 30
ip prefix-list blk_A permit 0.0.0.0/0 le 32
```

The blk_A list blocks the subnets that exactly match the first 16 most significant bits as 172.161.0.0. The ge keyword indicate that the subnet mask for the 172.161.0.0 subnets must be greater than or equal to 24 bits. Similarly, the le keyword indicates that the mask for the 172.161.0.0 subnets should be less than or equal to 30 bits. Therefore, all subnets of 172.161.0.0 network with masks 24, 25, 26, 27, 28, 29, and 30 are blocked.

The second line permits all other routes to be passed on. The subnets that match the blk_A prefix list are 172.161.9.0/24, 172.161.111.0/25, 172.161.247.0/30, and 172.161.64.0/28.

The line distribute-list prefix blk_A out indicates that the distribute-list command applies the blk_A prefix list to all the outgoing interfaces. This implies that if rtrB receives an update about the 172.161.9.0/24, 172.161.111.0/25, 172.161.247.0/30 or 172.161.64.0/28 subnets, they are blocked. In this case, the 172.161.64.0/28 is not blocked through the Fa0/0 interface to rtrC because it is directly connected.

The 172.161.35.0/18 and 172.161.88.0/22 subnets are not blocked through the Fa0/0 interface of rtrB to rtrC. This is because both these subnets are outside the range of prefix masks 24 through 30; hence, these two subnets are allowed through the Fa0/0 interface.

Objective:
Layer 3 Technologies
Sub-Objective:

Configure and verify filtering with any protocol

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Technology Information > Technology Technotes > Filtering Routing Updates on Distance Vector IP Routing Protocols
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > distribute-list in
Cisco > Cisco IOS IP Routing: BGP Command Reference > ip prefix-list

QUESTION 13
Examine the exhibit.

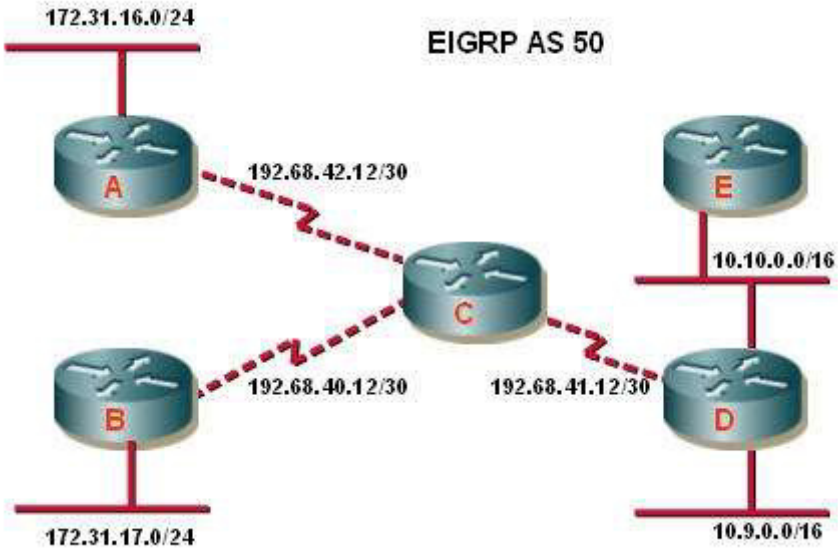

What additional EIGRP configuration is required to ensure that all destination networks are reachable if all routers are running pre- 15.0 versions of the IOS?

A. The eigrp stub receive only command should be executed on routers A and B.
B. A static route to 10.10.0.0/16 via the interface to router D should be configured at router C.
C. The no auto-summary router configuration command should be executed on router C.
D. The passive interface command should be executed on routers A and B.
E. The no auto-summary command should be executed on routers A and B.

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To ensure the full network is reachable, routers A and B must advertise their networks without first summarizing them to the class B 172.31.0.0/16 address in updates to router C. Otherwise, router C would incorrectly assume that it has two paths to the 172.31.0.0 classful network: one via router A and the other via router B. Therefore, routers A and B should be configured with the no auto-summary command so that they advertise 172.31.16.0/24 and 172.31.17.0/24, respectively. Starting with version 15, EIGRP auto summarization is disabled by default

Summarization is beneficial in most cases. It reduces the number of routes in the neighboring router tables and effectively contains EIGRP queries. The problem with discontiguous networks (or subnets) using EIGRP is that EIGRP will automatically summarize on the classful network boundary. By configuring the router to disable automatic summarization with the no auto-summary command, the routers will be able to see all of the individual subnets, not just a summary. The no auto-summary command must be issued from router

configuration mode as shown below:

router(config-router)# no auto-summary

Note that auto summarization is effective only on directly connected routes. For example, in the scenario exhibit, router C does not need to have auto summarization disabled in order to advertise the subnets to routers D and E. Since those subnet routes were learned via a route advertisement, they will be advertised to routers D and E without summarization.

In some situations, it may be necessary to turn off auto summarization globally while still summarizing specific networks. If you need to manually summarize a set of networks, the following command when executed in EIGRP configuration mode can summarize those specific networks while auto summarization is disabled:

ip summary-address [eigrp as-number] [address] [mask]

For example:

```
router10(config)# int Ethernet0/0
router10(config-if)# ip summary-address eigrp
```

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Summarization and Auto-summarization in EIGRP

QUESTION 14
You are configuring EIGRP on a spoke router in a hub-and-spoke topology. In an effort to keep the routing table small, the hub router has been configured to send only a default route to the remote routers.

What command would you use on the spoke routers to enable them to send only connected and summary routes to the hub router, and prevent the hub router from sending a query to the spoke router when a route is lost elsewhere?

A. eigrp stub
B. eigrp stub static
C. eigrp passive
D. eigrp stub receive-only

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The eigrp stub command is used to configure a router to send only connected and summary routes to its neighboring router. For example, examine the following output of the show ip route command that was executed on a router configured as a stub router:

```
router10#show ip route
C 172.16.5.0/24 is directly connected, Serial 0
D 192.168.7.0/24 [90/16523564] via 172.16.4.1, 00:21:20, Serial 1
D 172.16.0.0/16 is a summary, 00:21:23, Null 0
C 172.16.4.0/24 is directly connected, Serial 2
```

The routes that will be advertised are 172.16.5.0/24, 172.16.4.0/24, and the summary route 172.16.0.0/16. The first two is directly connected routes, and the last is the summary route that is auto configured by the EIGRP process.

When the stub feature is enabled on a router, the router will announce itself as a stub router. Neighbor routers will not query a stub router for alternate routes when a route is lost elsewhere in the network. The EIGRP stub feature works well in hub-and-spoke topologies when the goal is to minimize the amount of EIGRP bandwidth and processing associated with the spoke router. The eigrp stub command has the following syntax:

eigrp stub [receive-only | connected | static | summary]

When you do not specify any keywords with the command, connected and summary are used by default.

- receive-only: Prevents the router from sending any connected or summary routes.
- connected: Instructs the router to send connected routes.
- static: Instructs the router to send static routes that were redistributed by using the redistribute static command.
- summary: Instructs the router to send summary routes.

These parameters can be combined to resolve various problems, as seen in the following image:

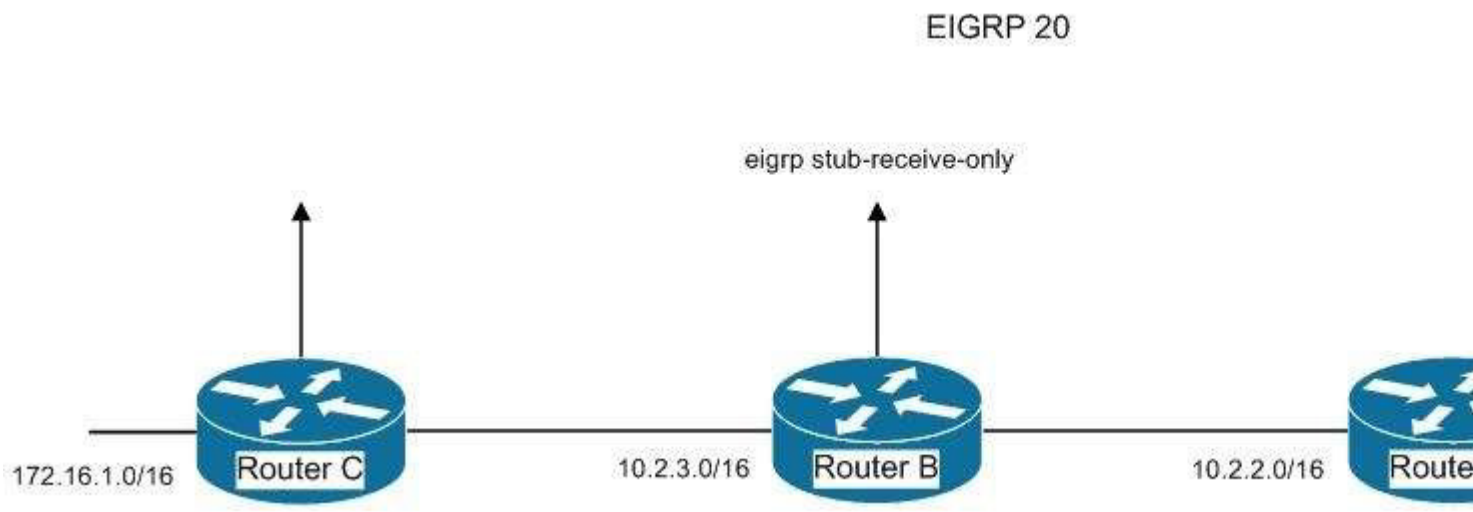

Router A is not receiving the route to the 172.16.1.0/16 network because Router B, which stands between Router A and C, is configured with the eigrp stub-receive-only command. This is resulting in hosts from the corporate office being unable to connect to hosts in the 172.16.0.0/16 network. If there were a legitimate reason to keep Router B configured with the eigrp stub-receive-only command, the problem could be solved by executing the following command set on Router A:

routerA(config)# router eigrp 20
routerA(config-router)# ip summary-address eigrp 20 172.16.0.0 255.255.0.0
routerA(config-router)# eigrp stub connected summary

This command set would create a summary address for the 172.16.0.0/16 network and then advertise it to the corporate office as a result of using the eigrp stub connected summary command. The inclusion of the connected parameter ensures that the directly connected networks will also be advertised, to ensure that hosts in the corporate office can reach the 172.16.0.0/16 network.

The eigrp stub static command instructs the router to send static routes that were redistributed by using the redistribute static command. Examine the EIGRP configuration shown below:

```
<output omitted>
ip route 10.4.4.0 255.255.255.0 10.4.3.10
Route eigrp 200
No auto-summary
Redistribute static 1000 1 255 1 1500
```

```
Network 10.4.1.0 0.0.0.3.
Network 10.4.2.0 0.0.0.255
Eigrp stub static
```

With this configuration, the router would not advertise any of the networks defined in the network statements, but would only advertise the static route configured with the line ip route 10.4.4.0 255.255.255.0 10.4.3.10.

Eigrp passive is not a valid Cisco command.

Eigrp stub receive-only will cause the router to not advertise any routes. The router will only receive updates.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP stubs

References:
Cisco IOS Master Command List, Release 12.4 > e through h > eigrp stub

QUESTION 15
OSPF area border routers (ABRs) advertise a default route to stub and totally stubby areas.

Which command is the BEST command to configure a cost of 25 for the default route advertised to area 1?

A. Router(config-router)# area 1 cost 25
B. Router(config-router)# area 1 default 25
C. Router(config-router)# area 1 default-cost 25
D. Router(config-router)# area 1 default-route-cost 25

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct answer is area 1 default-cost 25. Even though another option (area 1 default 25) is a configurable abbreviation for the command, the more correct answer explicitly specifies the default-cost parameter. The correct syntax for the area default-cost command is shown below:

Router(config-router)# area area-id default-cost cost

If you have multiple border routers between two areas, you might prefer one exit-point router over the other for that area. By configuring one with a lower cost than the other, it will become the preferred exit point. If that router or its links were to fail, then the routers interior to the area would route through the second-best exit point. You could also set the default costs to values that are close to achieve better load balancing. The default default-cost is 1. Please see the network shown in the graphic.

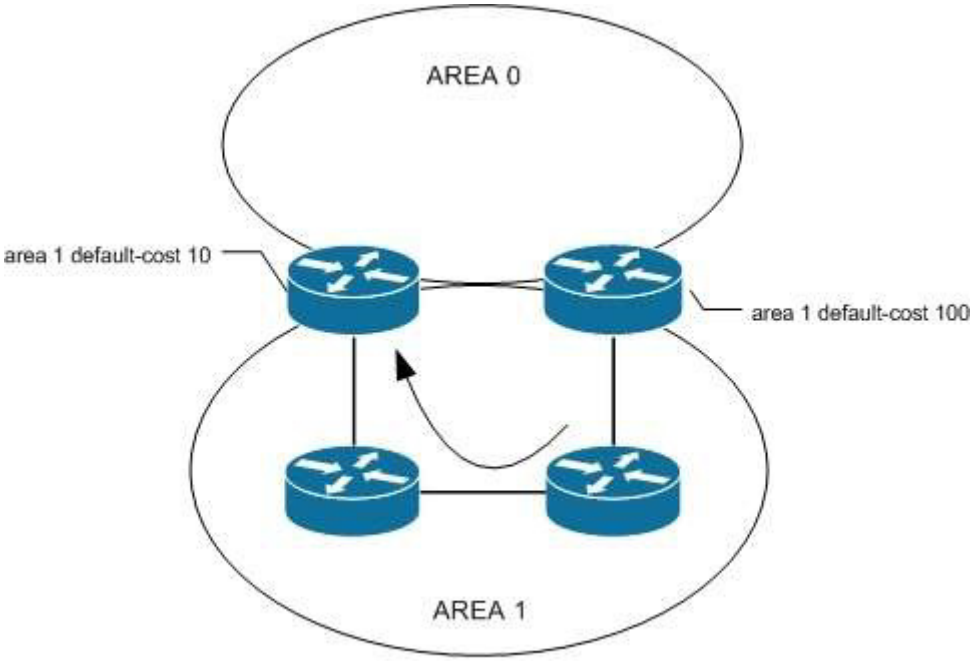

All traffic will follow the path indicated by the curved arrow to the preferred ABR.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF path preference

References:
Cisco IOS Master Command List, Release 12.4 > a through b > area default-cost

QUESTION 16
You need to manually assign IPv6 addresses to the interfaces on an IPv6-enabled router. While assigning addresses, you need to ensure that the addresses participate in neighbor discovery and in stateless auto-configuration process on a physical link.

Which of the following addresses can be assigned to the interfaces?

A. FEC0:0:0:1::1/64
B. FE80::260:3EFF:FE11:6770/10
C. 2001:0410:0:1:0:0:0:1/64
D. 2002:500E:2301:1:20D:BDFF:FE99:F559/64

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The FE80::260:3EFF:FE11:6770/10 address can be assigned to an interface of the IPv6-enabled router. This address is a link-local address as it has the prefix FE80::/10. Link-local addresses can be configured for an interface either automatically or manually.

Link-local addresses are IPv6 unicast addresses that are configured on the interfaces of an IPv6-enabled router. With link-local addresses, the nodes can connect to a network (local link) and communicate with other nodes. In addition, these addresses participate in the neighbor discovery protocol and the stateless auto-

configuration process.

The FEC0:0:0:1::1/64 address should not be used for the interfaces because this address is a site-local address. Site-local addresses are IPv6 equivalent addresses to IPv4's private address classes. These addresses are available only within a site or an intranet, which typically is made of several network links.

You should not use the 2001:0410:0:1:0:0:0:1/64 and 2002:500E:2301:1:20D:BDFF:FE99:F559 addresses for the interfaces. These two addresses are global unicast addresses as they fall in the range from 2000::/3 and to E000::/3. A global address is used on links that connect organizations to the Internet service providers (ISPs).

Objective:
Layer 3 Technologies
Sub-Objective:
Identify IPv6 addressing and subnetting

References:
Cisco > Understanding IPv6 Link Local Address

QUESTION 17
Which of the following commands allows a Cisco router to obtain an IP address from a DHCP server?

A. Router(config-if)# ip address dhcp
B. Router(config)# ip address dhcp
C. Router(dhcp-config)# ip address dhcp
D. Router(config)# address dhcp
E. Router(dhcp-config)# address dhcp

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The ip address dhcp command when issued from interface configuration mode will allow a router to obtain an IP address for that interface from a DHCP server.

In this scenario, the router is acting as a DHCP client, not a server, so the command would not be issued from dhcp-config mode. In addition, the IP address is being assigned to an interface on the router, not the router as a whole so the command would not be entered at global config mode.

The most common situation in which a router interface might be set as a DHCP client is to enable one DHCP router to obtain configuration options from another router providing this service.

Consider an example where RouterA is connected to RouterB. RouterA contains a complete DHCP configuration including the options (DNS server, domain name). RouterB is connected to RouterA through its FastEthernet0 interface. The following configuration would allow RouterB to issue a different set of addresses than RouterA while importing the options from Router A. The configuration of RouterB is below as shown in the partial output of the show run command:

```
hostname RouterB
!
ip dhcp-excluded-address 40.0.0.1
ip dhcp pool B
     network 40.0.0.0 255.255.255.0
     default-router 40.0.0.1
     import all
!
interface fastethernet0
     ip address dhcp
```

Note that for this configuration to function properly, the FastEthernet0 interface on RouterB must be set as a DHCP client.

The command router(config)# ip address dhcp is incorrect because it is executed at the global configuration prompt. The command must be executed in interface configuration mode.

The command router(dhcp-config)# ip address dhcp is incorrect because it is executed at the DHCP configuration prompt. The command must be executed in interface configuration mode.

The command router(config)# address dhcp is incorrect because it is missing the ip part of the command.

The command router(dhcp-config)# address dhcp is incorrect because it is missing the ip part of the command and it is executed at the DHCP configuration prompt. It must be executed in interface configuration mode.

Objective:
Layer 3 Technologies
Sub-Objective:
Identify, configure, and verify IPv4 addressing and subnetting

References:
Cisco > Cisco IOS IP Addressing Services Command Reference > ip address dhcp

QUESTION 18
Which statements about BGP policy-based routing are true? (Choose two.)

A. BGP policy-based routing is performed on a router's inbound interface.
B. A BGP administrator can use policy-based routing to alter the final destination of the packet.
C. BGP policy-based routing will select the next-hop of the packet based on its source address.
D. BGP policy-based routing can be used to alter the path selection for a packet in a downstream AS.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
BGP policy-based routing is performed on a router's inbound interface. BGP policy-based routing will select the next-hop of the packet based on its source address. It does this through the use of route maps.

Below is a partial output of the show run command executed on a router that has a BGP configuration that uses a route map to alter the local preference of a route (172.16.0.0/16) to 90 if it is advertised from the BGP neighbor at 10.5.5.1:

```
Neighbor 10.5.5.1 route-map test in
!
Access-list 99 permit 172.16.0.0 0.0.255.255
!
Route-map test permit 10
Match ip address 99
Set local-preference 90
```

The effect of this route map can be seen in the following partial output of the **show ip bgp** command executed on the same router:

```
Router5# show ip bgp

BGP table version is 5, local router ID is 10.5.5.34
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
*> 10.0.0.0 10.5.5.1
* 20.0.0.0 10.5.5.35 100 <output omitted>
*> 11.0.0.0 10.5.5.35 100
> 172.16.0.0/16 10.5.5.35 100
*> 172.16.0.0/16 10.5.5.1 90
```

The above output indicates that the local preference for the route to 172.16.0.0/16 is 90 ONLY if it comes from 10.5.5.1, but not if the same route is advertised from 10.5.5.35.

Route maps can be used to influence a part of the routing table without affecting the rest of the table. Consider an example where a router had two routes to every network in the table, and it prefers Neighbor A as the next hop for all routes. If you desired to change the next hop for one of the routes to Neighbor B without affecting the others, you could use route maps to take two different approaches, altering different attributes, which would arrive at the same result. The approaches would be:

- Apply a route map to Neighbor B incoming that would set the local preference to 200 (default is 100) for the route. Local preference values determine the path used to exit the AS. A higher value is preferred.
- Apply a route map to Neighbor A such that it advertises the route with a MED of 200. Med values determine the preferred path into the AS. A lower value is preferred. The default is 0.

Either of these approaches would result in the next hop for the network hanging to Neighbor B without affecting the others,

Policy-based routing does not alter the destination address of the packet. It can only alter the path to that final destination.

The BGP routing policy in one AS cannot determine the BGP routing policy in another AS.

Objective:
Layer 3 Technologies
Sub-Objective:
Identify suboptimal routing

References:
Cisco > Home > Support > Technology Support > IP Routing > Design > Design Technotes > Route-Maps for IP Routing Protocol Redistribution Configuration

QUESTION 19
Which statements in regards to route filtering are true? (Choose two.)

A. Network security is the primary role of route filtering.
B. If no route filter exists for an interface, the packet is processed normally.
C. Route filtering on an interface cannot filter routes that originate from the same router.

D. The distribute-list command enables the administrator to filter redistributed routes.
E. The network keyword of the passive-interface command enables you identify the routes to advertise.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Distribute lists are used to filter inbound, outbound, or redistributed routing updates. Instead of using the passive-interface command, distribute lists enable you to selectively control which routes are processed.

If no distribute list is associated with the interface, the routing update packets are processed normally.

If a distribute list is associated with an interface, the routing update is compared to the access list that was specified in the distribute list. If a match is found to a permit statement, then the packet is forwarded. If a match is found to a deny statement, the packet is discarded. If no match is found, the implicit deny statement at the end of the access list will drop the packet.

Network security is not the primary role of route filtering. Its primary function is to reduce unnecessary routing update traffic.

Route filtering on an interface can filter routes that originate from the same router.

The network keyword of the passive-interface command does not enable you identify the routes to advertise.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify filtering with any protocol

References:
Cisco > Home > Support > Technology Support > IP Routing > Design > Design Technotes > Filtering Routing Updates on Distance Vector IP Routing Protocols
Cisco > Cisco IOS IP Configuration Guide: Configuring IP Routing Protocol-Independent Features > Filtering Routing Information

QUESTION 20
By default, how often are EIGRP hello packets sent on a LAN?

A. 5 seconds
B. 10 seconds
C. 30 seconds
D. 60 seconds

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The EIGRP default hello time over a LAN or high-speed dedicated WAN link is five seconds.

On multipoint circuits with less than T1 bandwidth, EIGRP hello packets are sent every 60 seconds.
EIGRP sets the default hello interval to five seconds to ensure that it can quickly sense if connectivity to a neighboring router has been cut. If a router does not hear from a neighboring EIGRP router in 15 seconds, it will declare that neighbor as no longer reachable.

The five-second hello interval is shorter than the default values for OSPF hellos (10 seconds), RIP updates (30), or IGRP updates (90). As a result, EIGRP senses network faults faster by default than do other protocols.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP neighbor relationship and authentication

References:
Internetworking Technology Handbook > Enhanced Interior Gateway Routing Protocol (EIGRP) > Underlying Processes and Technologies

QUESTION 21
With respect to modifying an OSPF router ID to a loopback address, which of the following statements are true?

A. OSPF is not as reliable if a loopback interface is configured.
B. Using a loopback address avoids wasting an additional IP address.
C. A loopback interface is not always active, and it can go "down" like a real interface.
D. The loopback address does not automatically appear in the routing table of neighboring OSPF routers, so it cannot be pinged from other routers unless you include it with a network statement on the router local to the loopback interface.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A loopback address does not automatically appear in neighboring routers' routing tables, so it cannot be pinged for network troubleshooting.

A work-around for this problem is to add a network statement under OSPF that advertises the loopback address network so that other routers will know how to reach your loopback.

A loopback address is an IP address assigned to a loopback interface, which is a logical interface on a router that behaves like a physical interface. Their advantage is that, unlike physical interfaces, logical interfaces do not go down.

For example:
Router(config)# interface loopback 0
Router(config-if)# ip address 172.17.1.1 255.255.255.0

In the example, a loopback IP address is used by OSPF to provide its router ID. This type of address is preferred because it is assumed to be more stable than a router ID tied to a physical interface. The traditional problem with a router ID tied to a physical interface is that if the physical interface were to go down, the router would have to change its router ID to some other value. That would cause the OSPF neighbor relationships to reset and change values in the link-state advertisements (LSAs), causing a disruption to the OSPF area.

With this consideration in mind, OSPF is more reliable when using a loopback interface than using a physical interface.

Using a loopback address does not avoid wasting an additional IP address. The address must still be unique.

A loopback interface is always active, and it cannot go "down" as a physical interface can.

Objective:
Layer 3 Technologies
Sub-Objective:

Configure and verify OSPF operations

References:
Cisco > IP Routing: OSPF Configuration Guide > Configuring OSPF > Forcing the Router ID Choice with a Loopback Interface

QUESTION 22
Which of the following commands need to be configured on a RIPng router prior to enabling this routing protocol?

A. ipv6 rip enable
B. ipv6 multicast-routing
C. ipv6 unicast-routing
D. ipv6 router rip

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The ipv6 unicast-routing command should be used before enabling RIPng on a router. This command should be executed in global configuration mode of the router. IPv6 can then be enabled by using the ipv6 enable command on any of the interfaces of the router. The ipv6 unicast-routing command allows you to forward IPv6 unicast datagrams.

Routing Information Protocol Next Generation (RIPng) allows routers to learn about routes in an autonomous system. RIPng is an extension of the RIPv2 protocol to provide support IPv6 for future adherence.

The similarities between RIPv2 and RIPng are as follows:
- Both protocols use User Datagram Protocol (UDP).
- Both use distance vector algorithm to find the best route.
- Both of them measure the metric in terms of hops.
- Both have the same maximum hop count of 15 for valid routes.

The differences between RIPv2 and RIPng are as follows:
- RIPv2 learns IPv4 routes, whereas RIPng learns IPv6 routes
- RIPv2 supports automatic summarization as IPv4 defines classful addresses, whereas RIPng does not support automatic summarization
- RIPv2 uses UDP port 520, whereas RIPng supports port 521
- RIPv2 requires authentication for RIP packets, whereas RIPng does not require RIP-specific authentication as IPv6 has an in-built IPsec authentication

The ipv6 rip enable command should not be used because this command allows you to enable IPv6 RIP routing process on the interfaces of a router.

You should not use the ipv6 multicast-routing command prior to enabling IPv6 on the router. This command is used after IPv6 is enabled on one or more interfaces of the router to allow multicast forwarding using Protocol Independent Multicast (PIM) and Multicast Listener Discovery (MLD) on all the IPv6-enabled interfaces.

The ipv6 router rip command should not be used prior to enabling IPv6 because it allows you to enter the RIP for IPv6 router mode.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe RIPng

References:

Cisco > Configuring IPv6 Routing
Cisco > Cisco IOS IPv6 Command Reference > ipv6 unicast-routing

QUESTION 23
You are configuring BGP speakers RouterA and RouterB to authenticate one another. The following conditions exist:

- RouterA has an IP address of 192.168.5.3
- RouterB has an IP address of 192.168.5.2
- Both routers reside in AS 6550.

Which of the following commands will result in successful authentication?

A. neighbor 192.168.5.2 password routera executed on RouterA
neighbor 192.168.5.3 password routerb executed on RouterB
B. neighbor 192.168.5.2 password routerb executed on RouterA
neighbor 192.168.5.3 password routera executed on RouterB
C. neighbor 192.168.5.2 password routera executed on RouterA
neighbor 192.168.5.3 password routera executed on RouterB
D. neighbor 192.168.5.2 password routera executed on RouterA
E. neighbor 192.168.5.2 password routerb executed on RouterB

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The following command pair should be used to configure successful authentication:

neighbor 192.168.5.2 password routera executed on RouterA
neighbor 192.168.5.3 password routera executed on RouterB

When setting the keys for authentication, the keys must match. The keys do not need to be the names of either router, and should be a combination of letters numbers and symbols. In this example, both keys are set to the value routera.

The following two command pairs are incorrect because the keys do not match:

neighbor 192.168.5.2 password routera executed on RouterA
neighbor 192.168.5.3 password routerb executed on RouterB

and

neighbor 192.168.5.2 password routerb executed on RouterA
neighbor 192.168.5.3 password routera executed on RouterB

If you executed the debug ip bgp command to perform troubleshooting with either of these configurations in place, the error message you would see would be as follows:

```
%TCP-6-BADAUTH: Invalid MD5 digest from 192.168.5.3 (12293) to 192.168.5.2 (179)
```

In the error message, the numbers in parentheses are the port numbers used for the attempted communication.

The single commands would be incorrect because the key has only been configured on one end:

neighbor 192.168.5.2 password routera executed on RouterA

and

neighbor 192.168.5.2 password routerb executed on RouterB

If you executed the debug ip bgp command to troubleshoot with either of these configurations in place, you would see the following message:

```
%TCP-6-BADAUTH: No MD5 digest from 192.168.5.3 (12293) to 192.168.5.2 (179)
```

Objective:
Layer 3 Technologies
Sub-Objective:
Describe, configure, and verify BGP peer relationships and authentication

References:
Cisco IOS Master Command List, Release 12.4 > I through q > Cisco IOS IP Routing: BGP Command Reference > neighbor password

QUESTION 24
If you executed the show ip ospf database command, which keyword would you add to the command to produce the following output?

```
Router# show ip ospf database _____
OSPF Router with id(192.168.55.56) (Process ID 300)
Displaying Summary Net Link States(Area 0.0.0.0)
LS age: 1401
Options: (No TOS-capability)
LS Type: Summary Links(Network)
Link State ID: 10.199.240.0 (summary Network Number)
Advertising Router: 10.199.241.5
LS Seq Number: 80000072
Checksum: 0x84FF
Length: 28
Network Mask: 255.255.255.0 TOS: 0 Metric: 1
```

A. router
B. summary
C. network
D. external

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The output was produced with the summary keyword. When the show ip ospf database command is executed, any of several keywords can be used to specify the type of link-state advertisements (LSAs) to display. The output LS Type: Summary Links(Network) indicates that these are summary links. Summary LSAs are generated by an area border router (ABR) and will be displayed when you execute the summary keyword. These are Type 3 LSAs.

The router keyword was not used. If this keyword had been used, the LS Type line would have included Router Links instead of Summary Links. Router LSAs are Type 1 LSAs.

The network keyword was not used. If this keyword had been used the LS Type line would have included Network Links instead of Summary Links. Network LSAs are Type 2 LSAs.

The external keyword was not used. If this keyword had been used the LS Type line would have included AS External Links instead of Summary Links. External LSAs are Type 5 LSAs.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe OSPF packet types

References:
Cisco > Cisco IOS IP Routing: OSPF Command Reference > show ip ospf database

QUESTION 25
Which command can you use to specify that network 208.15.208.0 belongs to OSPF area 0?

A. router(config)# network 208.15.208.0 area 0
B. router(config-if)# ip ospf area 0
C. router(config)# network 208.15.208.0 255.255.255.0 area 0
D. router(config-router)# network 208.15.208.0 0.0.0.255 area 0

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You identify the area to which a network belongs with the network area command issued from router configuration mode:

router(config-router)# network address wildcard-mask area area-id

To enter router configuration mode, enter the command router ospf process ID in global configuration mode. For this command to be accepted and acted upon by the router, at least one interface on the router must have an IP address assigned and be up.

The command router(config)# network 208.15.208.0 area 0 is incorrect because it is executed in global configuration mode, as evidenced by the prompt router(config)#.

The command router(config-if)# ip ospf area 0 is incorrect. This command would be used to configure the router for OSPF and its area. It would also enter configuration mode for that particular process of OSPF so the user can enter additional commands that affect that process. However, this command is missing the process ID.

The command router(config)# network 208.15.208.0 255.255.255.0 area 0 is incorrect because it is executed in the wrong mode. It is entered in global configuration mode instead of OSPF configuration mode. It also has an incorrect mask. You must use a wildcard mask instead of a regular mask in the network statements for OSPF. In this case, the mask should be 0.0.0.255 instead of 255.255.255.0.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF operations

References:
Cisco : OSPF Commands > network area

QUESTION 26
Consider the partial output of the show ip bgp command:

```
RouterA# show ip bgp
BGP table version is 89, local router ID is 200.17.34.15
.
.
Network Next Hop Metric LocPrf Weight Path
*>i 10.62.7.0 10.62.7.78 0 100 0 1 i
*>e 192.177.1.0 10.62.7.115 100 75 50 1 2 3 5 i
h 61.80.3.0 10.62.7.44 0 100 0 1 2 3 5 9 i
*>i 200.17.56.0 200.17.56.101 0 100 0 i
*> 200.17.34.0 0.0.0.0 0 100 32768 i
```

Which of the following statements are TRUE about the given output? (Choose all that apply.)

A. The 10.62.7.0 route is learned by the router through an iBGP neighbor.
B. All five routes have been originated by an IGP.
C. The router is aware of the best path for the 61.80.3.0 destination.
D. There are four AS between the router and the 192.177.1.0 subnet.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The following statements are TRUE about the given output:
- The 10.62.7.0 route is learned by the router through an iBGP neighbor.
- All five routes have been originated by an IGP.

The show ip bgp command displays information about the BGP routing table, including origin type, metric, next-hop addresses for every route learned by BGP, router ID, local preference, and BGP path. In the output, the character i in the first entry of the 10.62.7.0 destination indicates that the route was learned by an iBGP neighbor. The * symbol at the beginning of the routes indicate that they are valid routes, while the > symbol indicate that the route is the current best route.

The i at the end of the entries under the Path column indicates that the routes have been originated by an interior gateway protocol (IGP). In the scenario output, all five routes have an i at the end of their respective entries. If the character e appears as the origin code, the routes are considered to have originated from an exterior gateway protocol (EGP). The origin code can also be the ? character, which implies that the origin of the route is unknown.

The output also displays the next-hop addresses for the routes. The 200.7.34.0 subnet is a local route, and hence has its next-hop address as 0.0.0.0.

The show ip bgp command also displays the local router's ID (RID), local preference, weight, and next-hop addresses for every route learned by BGP. In this case, the RID of RouterA is 200.17.34.15 and the local preference, weight, and next-hop address for the 10.62.7.0 network are 100, 0, and 10.62.7.78, respectively. The metric and the next-hop address for the BGP routes can also be viewed by using the show ip route bgp command, as follows:

```
RouterA# show ip route bgp
B 10.62.7.0 [200/0] via 10.62.7.78, 01:34:16
B 200.17.56.0 [200/0] via 10.62.7.78, 01:34:16
B 192.177.1.0 [20/100] via 10.62.7.115, 01:34:16
```

The BGP table version can also be displayed by using the show ip bgp neighbors and the show ip bgp summary commands. The show ip bgp neighbors command also displays the address, ASN, and RID of neighbors of the local router, as shown below:

```
RouterA# show ip bgp neighbors
```

```
BGP neighbor is 192.177.1.6, remote AS 200, external link
BGP version 17, remote router ID 200.17.34.15
BGP state = Established, table version = 16, up for 01:45:03
<output omitted>
```

The show ip bgp summary command displays the RID and the BGP table version, as shown in the following output:

```
RouterA# show ip bgp summary
BGP router identifier 200.17.34.15, local AS number 100
BGP table version is 17, main routing table version 18
<output omitted>
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.62.7.90 17 200 56 55 18 0 0 01:42:13 27
10.62.7.145 17 300 34 33 18 0 0 00:31:20 0
```

The router is not aware of the best path for the 61.80.3.0 route. The character h appears at the beginning of the entry for the 61.80.30 destination. This means that the route is in the history state currently and that the best route is not known.

There are not four AS between the router and the 192.177.1.0 subnet. In the output, the Path column for the 192.1771.1.0 subnet lists four AS numbers. The four AS numbers refer to the ASNs traversed by the route from RouterA to the 192.177.1.0 subnet. The first AS refers to the first neighbor of RouterA; the second AS refers to the neighbor of the first neighbor; and so on. The last AS in the column is the AS of the 192.177.1.0. This implies that there are three AS (1, 2, and 3) that exist between RouterA and the subnet.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify eBGP (IPv4 and IPv6 address families)

References:
Cisco > Cisco IOS IP Routing: BGP Command Reference > show ip bgp
Cisco > Cisco IOS IP Routing: BGP Command Reference > show ip route bgp
Cisco > Cisco IOS IP Routing: BGP Command Reference > show ip bgp summary

QUESTION 27
Which commands will display the other routers with which the local router has established an adjacency with, including hold time and uptime parameters?

A. show ip eigrp neighbors
B. show ip route
C. show adjacencies
D. show eigrp neighbors

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The show ip eigrp neighbors command will display the neighboring EIGRP routers with which the local router has established an adjacency. It will also display hold time and uptime statistics. In this case, the uptime statistic refers to how long the adjacency has been established. A sample output of the show ip eigrp neighbors command is below.

```
Router2# show ip eigrp neigh
IP-EIGRP neighbors for process 100
H Address Interface Hold Uptime SRTT RTO Q Seq
(sec) (ms) Cnt Num
```

```
1 10.20.0.1 Se1 11 22:37:26 28 200 0 2
0 10.10.0.1 Se0 13 22:38:09 19 200 0 4
```

The show ip route command simply displays the routing table and does not provide neighbor information.

The other commands are not valid IOS commands.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP neighbor relationship and authentication

References:
Cisco IOS IP Routing: EIGRP Command Reference > show ip eigrp neighbors

QUESTION 28
Which of the following statements is TRUE about the communication occurring between rtrA and rtrB in the given exhibit?

```
rtrA(config)# interface loopback 0
rtrA(config-interface)# ip address 10.144.1.1 255.255.255.255
rtrA(config-interface)# exit
rtrA(config)# router bgp 450
rtrA(config-router)# network 10.144.1.0 mask 255.255.255.0
rtrA(config-router)# neighbor 131.78.45.2 remote-as 450
rtrA(config-router)# neighbor 131.78.45.2 update-source loopback 0
```

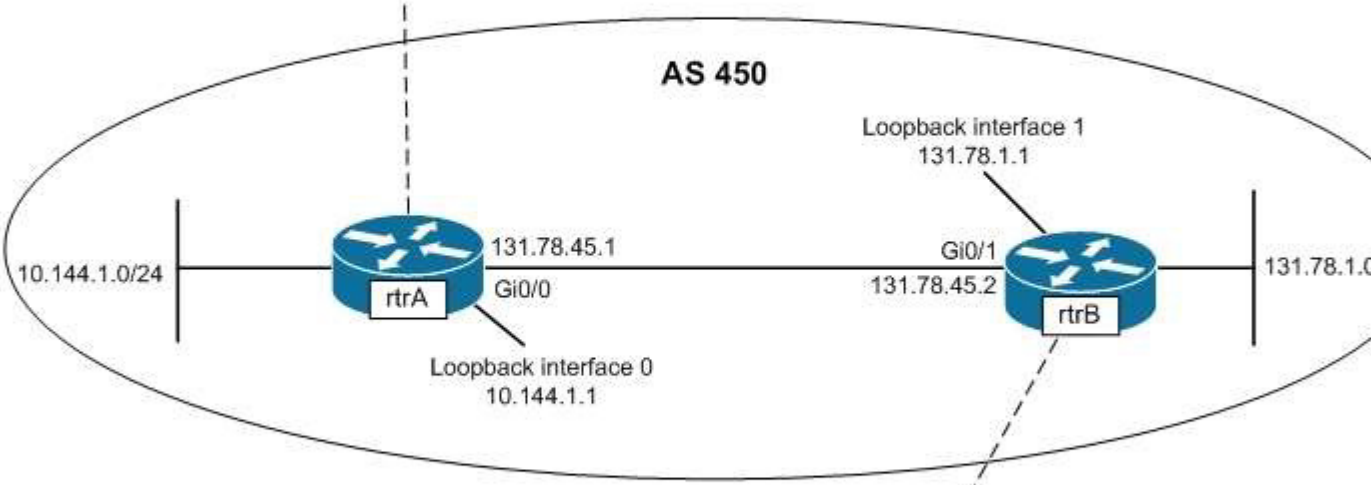


```
rtrB(config)# interface loopback 1
rtrB(config-interface)# ip address 131.78.1.1 255.255.255.255
rtrB(config-interface)# exit
rtrB(config)# router bgp 450
rtrB(config-router)# network 131.78.1.0 mask 255.255.255.0
rtrB(config-router)# neighbor 10.144.1.1 remote-as 450
```

A. The only loopback interface used in the communication is the loopback 0 interface of rtrA.
B. The only loopback interface used in the communication is the loopback 1 interface of rtrB.

C. Both loopback 0 and loopback 1 interfaces are used for the communication between rtrA and rtrB.
D. Neither loopback 0 nor loopback 1 interface is used for the communication between rtrA and rtrB.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The only loopback interface used in the communication is the loopback 0 interface of rtrA. The configuration on the rtrA router indicates that BGP is enabled on rtrA in the autonomous system number (ASN) 450. The neighbor 131.78.45.2 remote-as 450 command establishes a connection with the rtrB interface having the 131.78.45.2 address. The Gi0/1 interface of rtrB has the address 131.78.45.2, which is directly connected to the Gi0/0 interface (132.78.45.1) of rtrA. The next line, neighbor 131.78.45.2 update-source loopback 0, specifies that the 131.78.45.2 interface (Gi0/1) of rtrB communicates with the loopback 0 interface on rtrA.

In the configuration of rtrB, the neighbor 10.144.1.1 remote-as 450 command establishes a neighboring relationship with the interface having the address10.144.1.1. The loopback 0 interface of rtrA has the address 10.144.1.1. The loopback 1 interface on rtrB is assigned an IP address but not used in establishing BGP connections between rtrA and rtrB

Loopback 1 interface of rtrB only would only be used in the communication between rtrA and rtrB if you configured rtrA and rtrB as follows:

```
rtrA(config)#router bgp 450
rtrA(config-router)# neighbor 131.78.1.1 remote-as 450

rtrB(config)#router bgp 450
rtrB(config-router)# neighbor 131.78.45.1 remote-as 450
rtrB(config-router)# neighbor 131.78.45.1 update-source loopback 1
```

Both loopback 0 and loopback 1 interfaces are NOT used for communication between rtrA and rtrB. Only the loopback 0 interface of rtrA is used. Both of the loopback interfaces would be used in the communication between rtrA and rtrB only if you changed the configuration of rtrA and rtrB, as given below:

```
rtrA(config)# router bgp 450
rtrA(config-router)# neighbor 131.78.1.1 remote-as 450
rtrA(config-router)# neighbor 131.78.1.1 update-source loopback 0

rtrB(config)#router bgp 450
rtrB(config-router)# neighbor 10.144.1.1 remote-as 450
rtrB(config-router)# neighbor 10.144.1.1 update-source loopback 1
```

Because the loopback 0 interface of rtrA is used in communication, is incorrect to state that neither loopback 0 nor loopback 1 interface is used. To ensure that neither of the loopback interfaces are be used for communication, you would configure rtrA and rtrB as follows:

```
rtrA(config)# router bgp 450
rtrA(config-router)# neighbor 131.78.45.2 remote-as 450

rtrB(config)# router bgp 450
rtrB(config-router)# neighbor 131.78.45.1 remote-as 450
```

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify eBGP (IPv4 and IPv6 address families)

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > BGP Case

Studies > eBGP Multihop
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > neighbor update-source
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > neighbor remote-as

QUESTION 29
Which command can you use to display the topological database maintained by an OSPF router?

A. show ip ospf topology
B. show ip ospf database
C. show ip ospf [process-id]
D. show ip ospf border-routers

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct answer is show ip ospf database. Partial output is shown below:

```
Router# show ip ospf database
OSPF Router with id(192.168.239.75) (Process ID 352)
Displaying Router Link States(Area 0.0.0.0)
LinkID ADV Router Age Seq# Checksum Link count
172.16.21.6 172.16.21.6 1461 0x80002CFB 0x69BC 10
172.16.21.5 172.16.21.5 1146 0x800009D2 0xA2B8 4
172.16.1.2 172.16.1.2 1649 0x80000A98 0x4CB6 6
172.16.1.1 172.16.1.1 1136 0x800009B6 0x5F2C 1
172.16.1.5 172.16.1.5 1415 0x80002BC 0x2A1A 9
172.16.65.6 172.16.65.6 1785 0x80001947 0xEEE1 5
172.16.241.5 172.16.241.5 1136 0x8000007C 0x7C70 6
172.16.27.6 172.16.27.6 1415 0x80000548 0x8641 9
172.16.70.6 172.16.70.6 1666 0x80000B97 0xEB84 2
```

Issuing the show ip ospf database command will show you a summary of the database; however, to obtain details you must use a keyword with the command, such as router, network, summary, or external.

The following commands are available to verify OSPF configurations:
- show ip route - displays known routes and from which protocol the routes were discovered for all routing protocols.
- show ip ospf - displays the number of times the SPF algorithm has run and the default Link State Update (LSU) interval.
- show ip ospf database - displays the router ID, the OSPF process ID, and the contents of the topological database.

There is no show ip ospf topology command.

The show ip ospf [process-id] command can be used to view the number of times the SPF algorithm has been executed, but not to view the database.

The show ip ospf border-routers command display the ABRs and the routes to them, but not the contents of the database.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF path preference

References:
Cisco IOS Master Command List, Release 12.4T > sa ipsec through show ip route dhcp > show ip ospf database

QUESTION 30
Which command can be used to view the number of times the SPF algorithm has been executed?

A. show ip ospf
B. show ip ospf interface
C. show ip ospf database
D. show ip ospf neighbor

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The show ip ospf command can be used to view the number of times the SPF algorithm has been executed, as shown in the last line of the following output:

```
Router# show ip ospf
Routing Process "ospf 10" with ID 192.42.110.21
Supports only single TOS(TOS0) route
It is an area border and autonomous system boundary router
Redistributing External Routes from,
     igrp 200 with metric mapped to 2, includes subnets in redistribution
     rip with metric mapped to 2
     igrp 2 with metric mapped to 100
     igrp 32 with metric mapped to 1
Number of areas in this router is 3
Area 192.42.132.0
     Number of interfaces in this area is 2
     Area has simple password authentication
     SPF algorithm executed 10 times
```

The show ip ospf interface command can be used to view neighbor adjacencies. A partial output of the command is shown below. It will not show the number of times the SPF algorithm has been executed.

```
Router1# show ip ospf interface ethernet 1
Ethernet1 is up, line protocol is up
Internet Address 192.168.5.1/24, Area 0
Process ID 1, Router ID 192.168.45.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.10.1, Interface address 192.168.5.4
Backup Designated router (ID) 192.168.45.1, Interface address 10.10.10.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 172.16.10.1 (Designated Router)
Suppress hello for 0 neighbor(s)
```

The show ip ospf neighbor command can also be used to view neighbor adjacencies, although its output is slightly different from the show ip ospf interface command. A partial output of the show ip ospf neighbor command is shown below. It also does not show the number of times the SPF algorithm was executed.

```
Router# show ip ospf neighbor
ID Pri State Dead Time Address Interface
10.199.199.145 1 FULL/DR 0:00:31 192.168.80.37 Ethernet0
172.16.49.1 1 FULL/DROTHER 0:00:33 172.16.49.1 Fddi0
172.16.48.200 1 FULL/DROTHER 0:00:33 172.16.49.200 Fddi0
10.199.199.137 5 FULL/DR 0:00:33 172.16.49.189 Fddi0
```

The show ip ospf database command does not show the number of times the SPF algorithm has executed. It shows the contents of OSPF database. Partial output is shown below:

```
Router# show ip ospf database
OSPF Router with ID(192.168.1.11) (Process ID 1)
Router Link States(Area 0)
LinkID ADV Router Age Seq# Checksum Link count
192.168.1.8 192.168.1.8 1381 0x8000010D 0xEF60 2
192.168.1.11 192.168.1.11 1460 0x800002FE 0xEB3D 4
192.168.1.12 192.168.1.12 2027 0x80000090 0x875D 3
192.168.1.27 192.168.1.27 1323 0x800001D6 0x12CC 3
Net Link States(Area 0)
```

You can make the command output more specific by using parameters with the show ip ospf database command. For example, to view only Type 5 LSAs in the database, you would execute the show ip ospf database external command. Since all Type 5 LSAs are from external networks, this keyword will trim the output to only those types of LSAs. When Type 5 (or external) routes are placed in the database, the next hop address will be 0.0.0.0, which makes it appear as if it is a default route. What this really means is that any traffic that needs to go to that external network will be sent to the router that originated the advertisement (the ASBR).

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF operations

References:
Cisco IOS IP Routing: OSPF Command Reference > OSPF Commands: show ip ospf through T > show ip ospf

QUESTION 31
Which command shows a list of neighboring routers, their priorities, and their current state?

A. show ip ospf
B. show ip protocol
C. show ip ospf database
D. show ip ospf neighbor [detail]

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The show ip ospf neighbor [detail] command will display the OSPF information that is known about OSPF neighbors and the OSPF operating state with them.

The commands below can be used to monitor and verify OSPF operation:

- show ip ospf - shows the number of times the SPF algorithm has run and the default LSU interval.
- show ip protocol - displays information about timers, filters, metric, etc. for the entire router.
- show ip ospf database - shows the router ID, the OSPF process ID, and the contents of the topological database.

These commands do not show details about OSPF neighbors.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF neighbor relationship and authentication

References:
Cisco > Cisco IOS IP Routing Protocols Command Reference > IP Routing Protocol-Independent Commands: S through T > show ip ospf neighbor

QUESTION 32
Consider the partial output of the show ip route eigrp command:

```
rtrA# show ip route eigrp

Gateway of last resort is not set

15.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D 15.11.78.0/24 [90/5345354] via 10.10.70.41, 01:43:05, S0/0
D 15.200.16.0/24 [90/1723780] via 10.10.78.23, 00:50:37, S0/0
[90/1723780] via 10.10.19.40, 01:04:58, S0/0
[90/1723780] via 10.10.70.41, 01:20:37, S0/0
D 15.90.4.0/16 [90/4869420] via 10.10.19.40, 01:13:17, S0/0

172.161.0.0/16 is variably subnetted, 6 subnets, 3 masks
D 172.161.50.0/24 [90/4531003] via 10.10.70.41, 00:53:10, S0/1
D 172.161.98.1/30 [90/1723695] via 10.10.78.23, 01:27:03, S0/1
D 172.161.11.0/27 [90/1723695] via 10.10.19.45, 00:56:17, S0/1
[90/1723695] via 10.10.19.40, 00:50:58, S0/1
D 172.161.2.74/27 [90/6356189] via 10.10.70.41, 01:36:31, S0/1
D 172.161.4.47/30 [90/88258329] via 10.10.78.23, 01:44:20, S0/1
D 172.161.150.6/24 [90/3285083] via 10.10.70.41, 02:13:55, S0/1
D*EX 0.0.0.0/0 [170/2645987] via 10.10.70.41, 00:05:12, Ethernet0/0
[170/2645987] via 10.10.70.23, 00:05:12, Ethernet0/0
```

Which of the following destination subnets have equally load-balanced routes? (Choose all that apply.)

A. 172.161.4.47/30
B. 172.161.11.0/27
C. 15.200.16.0/24
D. 15.11.78.0/24
E. 0.0.0.0/0

Correct Answer: BCE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The 172.161.11.0/27 and 15.200.16.0/24 networks have equally load-balanced routes. A default route, 0.0.0.0/0, has been configured for load balancing as well. These three subnets are each load balanced on multiple routes. The output entry for the 172.161.11.0/27 subnet is as follows:

```
D 172.161.11.0/27 [90/1723695] via 10.10.19.45, 00:56:17, S0/1
[90/1723695] via 10.10.19.40, 00:50:58, S0/1
```

This subnet can be reached by rtrA through two routes: 10.10.19.45 and 10.10.19.40 next-hop addresses. Both these routes have the same metric (1723695), and so are equally load balanced.

In the output, the 15.200.16.0/24 subnet has three equal-metric routes: 10.10.78.23, 10.10.19.40, and 10.10.70.41. These three routes are equally used to balance the load from rtrA to the 15.200.16.0/24 subnet.

The default route 0.0.0.0/0 is load balanced through two interfaces, as shown in the output:

```
D*EX      0.0.0.0/0 [170/2645987] via 10.10.70.41, 00:05:12, Ethernet0/0
          [170/2645987] via 10.10.70.23, 00:05:12, Ethernet0/0
```

This load balancing of the default route could be tested by using traceroute to any IP address not represented in the routing table and verifying the path taken.

Subnets 172.161.4.47/30 and 15.11.78.0/24 are not participating in load balancing. In the given output, there is a single route (line) for both of these subnets. The rtrA router uses the route through the next-hop 10.10.78.23 to reach the 172.161.4.47/30 destination subnet. Similarly, rtrA uses the next-hop 10.10.70.41 to transmit packets to the 15.11.78.0/24 subnet.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify EIGRP load balancing

References:
Cisco IOS IP Routing: Protocol-Independent Command Reference > show ip route
Cisco > Support > Technology Support > IP > IP Routing > Technology Information > Technology Whitepaper > Enhanced Interior Gateway Routing Protocol > Document ID: 16406 > Load Balancing

QUESTION 33
Routers R1 and R2 are being added to the network shown in the exhibit.

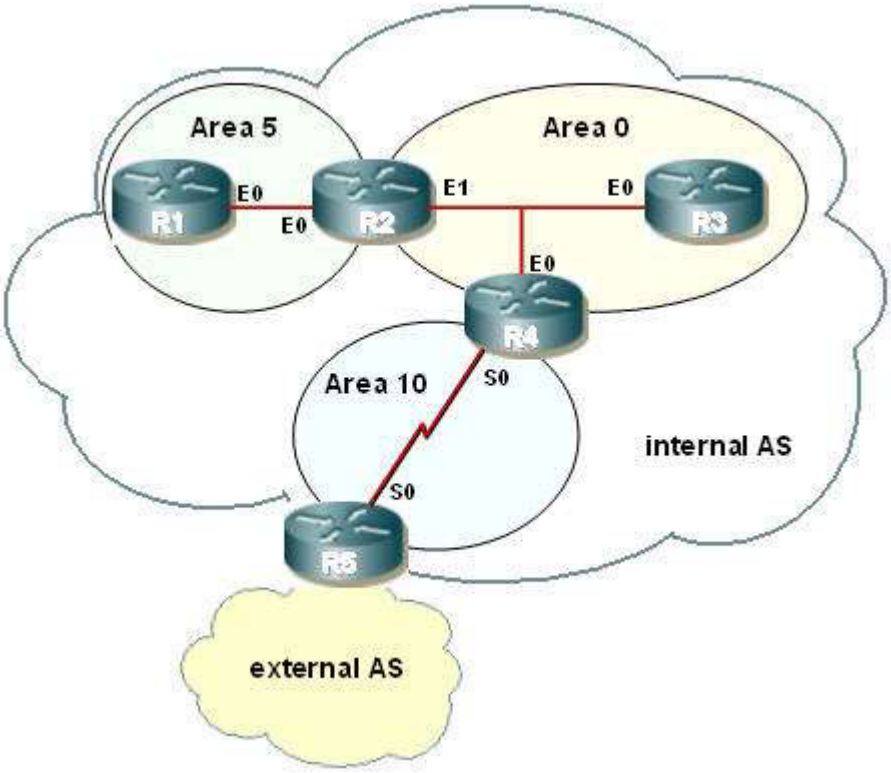

The addresses of their respective interfaces have already been configured as follows:

- R1: E0 192.168.4.5/30
- R2: E0 192.168.4.6/30
- R2: E1 192.168.72.6/30

You have been assigned to complete the following as a part of implementing OSPF area 5:

- The E0 interface on R1 should be in area 5.
- The E0 interface on R2 should be in area 5.
- The mask used with the OSPF configuration should only include the addresses for R1 and R2.
- Area 5 should not allow any external or inter-area routes (except for the default route).

Which commands are required to accomplish this set of requirements? (Choose all that apply.)

A. R1# configure terminal
R1(config)# router OSPF 1
R1(config-router)# network 192.168.4.4 0.0.0.3 area 5
R1(config-router)# area 5 stub
R1(config-router)# end
R1# copy running-config startup-config

B. R1# configure terminal
R1(config)# router OSPF 1
R1(config-router)# network 192.168.4.4 0.0.0.3 area 5
R1(config-router)# area 5 stub no-summary
R1(config-router)# end
R1#copy running-config startup-config

C. R1# configure terminal
R1(config)# router OSPF 1
R1(config-router)# network 192.168.4.4 0.0.0.4 area 5
R1(config-router)# area 5 stub
R1(config-router)# end
R1# copy running-config startup-config

D. R2# configure terminal
R2(config)# router OSPF 1
R2(config-router)# network 192.168.4.4 0.0.0.3 area 5
R2(config-router)# area 5 stub no-summary
R2(config-router)# end
R2# copy running-config startup-config

E. R2# configure terminal
R2(config)# router OSPF 1
R2(config-router)# network 192.168.4.4 0.0.0.3 area 0
R2(config-router)# area 0 stub no-summary
R2(config-router)# end
R2# copy running-config startup-config

F. R2# configure terminal
R2(config)# router OSPF 1
R2(config-router)# network 192.168.4.4 0.0.0.3 area 5
R2(config-router)# area 5 stub
R2(config-router)# end
R2# copy running-config startup-config

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The following set of commands will configure R1 properly and satisfy the requirements:

```
R1# configure terminal
R1(config)# router OSPF 1
R1(config-router)# network 192.168.4.4 0.0.0.3 area 5
R1(config-router)# area 5 stub
R1(config-router)# end
R1# copy running-config startup-config
```

The configure terminal command enters global configuration mode, from which the router ospf 1 command can be executed to enable OSPF process 1. The network command allows the192.168.4.4/30 network to join OSPF area 5 and uses a wildcard mask (0.0.0.3) that only includes the E0 interfaces on R1 and R2. The area 5 stub command configures R1 as an internal router in a totally stubby area, which is necessary because no external or inter-area routes are allowed. The final two commands exit OSPF configuration mode and save the configuration.

The following set of commands will configure R2 properly and satisfy the requirements:

```
R2# configure terminal
R2(config)# router OSPF 1
R2(config-router)# network 192.168.4.4 0.0.0.3 area 5
R2(config-router)# area 5 stub no-summary
R2(config-router)# end
R2# copy running-config startup-config
```

The configure terminal command enters global configuration mode, from which the router ospf 1 command can be executed to enable OSPF process 1. The network command allows the192.168.4.4/30 network to join OSPF area 5, and uses a wildcard mask (0.0.0.3) that only includes the E0 interfaces on R1 and R2. The area 5 stub no-summary command configures R2 as an area border router (ABR) in a totally stubby area, which is necessary because no external or inter-area routes are allowed. The final two commands exit OSPF configuration mode and save the configuration.

The wildcard mask on both network statements, 0.0.0.3, is the wildcard equivalent of a 255.255.255.252 mask (/30). When used with the network address 192.168.4.4, this mask will only allow two addresses in the area, 192.168.4.5 and 192.168.4.6, as per the scenario requirements.

The command set that executes the area 5 stub no-summary command on router R1 is incorrect because R1 is an internal router and does not require the no-summary keyword. The no-summary keyword is only required on the ABR when configuring a totally stubby area.

The command set that executes the network 192.168.4.4 0.0.0.4 area 5 command on router R1 has the wrong wildcard mask.

The command set that executes the network 192.168.4.4 0.0.0.3 area 0 command on router R2 is incorrect because the area should be area 5, not area 0.

The command set that executes the area 5 stub command on router R2 is incorrect becauseR2 is an ABR router and requires the no-summary keyword when configuring a totally stubby area.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco > Home > Support > Technology Support > IP Routing > Design > Design Technotes > What Are OSPF Areas and Virtual Links? > What Are Areas, Stub Areas, and Not-So-Stubby Areas?
Cisco > Cisco IOS IP Routing: OSPF Command Reference > area stub

QUESTION 34
You need to configure eBGP on the rtrA and rtrB routers, as shown in the following image:

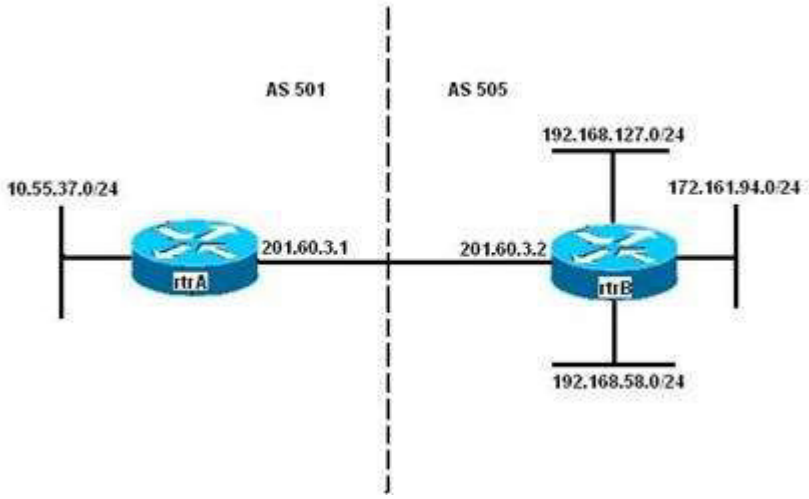

You have configured eBGP on rtrA through the following commands:

```
rtrA(config)# router bgp 501
rtrA(config)# neighbor 201.60.3.2 remote-as 505
```

While configuring eBGP on rtrB, you need to ensure that updates about the 192.168.58.0/24 and the 192.168.127.0/24 subnets are sent to rtrA with a metric of 300. In addition, rtrB should send updates about the 172.161.94.0/24 subnet are sent with a metric of 500.

Which of the following set of commands would NOT be part of the set used to correctly configure eBGP on rtrB?

A. access-list 1 permit 192.168.0.0 0.0.255.255
access-list 2 permit 172.161.94.0 0.0.0.255
router bgp 505

neighbor 201.60.3.1 remote-as 501
neighbor 201.60.3.1 route-map change_parameters in

B. access-list 1 permit 192.168.0.0 0.0.255.255
access-list 2 permit 172.161.94.0 0.0.0.255
router bgp 505
neighbor 201.60.3.1 remote-as 501
neighbor 201.60.3.1 route-map change_parameters out

C. route-map change_parameters permit 10
match ip-address 2
set metric 500

D. route-map change_parameters permit 20
match ip-address 1
set metric 300

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The following command set would NOT be used because it only applies the access list route-map change_parameters inbound instead of outbound, as would be required:

```
access-list 1 permit 192.168.0.0 0.0.255.255
access-list 2 permit 172.161.94.0 0.0.0.255
router bgp 505
neighbor 201.60.3.1 remote-as 501
neighbor 201.60.3.1 route-map change_parameters in
```

The following commands should be used to configure BGP on rtrB as desired:

```
access-list 1 permit 192.168.0.0 0.0.255.255
access-list 2 permit 172.161.94.0 0.0.0.255
router bgp 505
neighbor 201.60.3.1 remote-as 501
neighbor 201.60.3.1 route-map change_parameters out

route-map change_parameters permit 10
match ip-address 2
set metric 500

route-map change_parameters permit 20
match ip-address 1
set metric 300
```

The following set of commands creates two standard access-lists numbered 1 and 2:

```
access-list 1 permit 192.168.0.0 0.0.255.255
access-list 2 permit 172.161.94.0 0.0.0.255
router bgp 505
neighbor 201.60.3.1 remote-as 501
neighbor 201.60.3.1 route-map change_parameters out
```

The ACL 1 allows the 192.168.58.0/24 and the 192.168.127.0/24 subnets, while the ACL 2 allows the 172.161.94.0/24 subnet. The neighbor route-map command specifies a route-map named change_parameters for the 201.60.3.1 BGP peer. The out keyword at the end of the command indicates that the route-map is applied only to the updates sent by rtrB, and not received by rtrB.

In the following command, the route map change_parameters is defined with the permit keyword. The permit keyword indicates that if a match occurs, the actions specified in the set sub-command are executed:

route-map change_parameters permit 10
match ip-address 2
set metric 500

In this case, this command checks if the IP address of the subnets advertised to rtrA is in the 172.161.94.0/24 subnet (specified by ACL 2). If the IP address matches, then the metric of those routes are set to 500.

In the following command, the route map change_parameters is defined with the permit keyword:

route-map change_parameters permit 20
match ip-address 1
set metric 300

In this case, this command checks if the IP address of the subnets advertised to rtrA is in the 192.168.58.0/24 or the 192.168.127.0/24 subnets (specified by ACL 1) If the IP address matches, then the metric of those routes are set to 300.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify eBGP (IPv4 and IPv6 address families)

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > BGP Case Studies > Route Maps
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > route-map

QUESTION 35
Which command can you use to display the area border routers (ABRs) and the routes to them?

A. show ip ospf dr
B. show ip opsf bdr
C. show ip ospf database
D. show ip ospf border-routers

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct answer is show ip ospf border-routers. The following commands are available to verify OSPF configurations:

- show ip ospf border-routers - displays internal ISPF routing table entries for an ABR.
- show ip ospf virtual-links - displays the current state of OSPF virtual links.
- show ip ospf - displays information about the router's role and each area to which the router is connected.
- show ip ospf database - displays the contents of the router's topological database. Note that a number of keywords can be used with the show ip ospf database command to get specific information.

The command show ip ospf dr is not correct because dr is not a parameter of the show ip ospf command.

The command show ip ospf bdr is not correct because bdr is not a parameter of the show ip ospf command.

Objective:
Layer 3 Technologies
Sub-Objective:

Configure and verify network types, area types, and router types

References:
Cisco > Cisco IOS IP Routing: OSPF Command Reference > show ip ospf border-routers

QUESTION 36
If the following protocols are redistributed into OSPF, which protocol will receive a metric of 1 if none is specified, rather than the default metric of 20?

A. EIGRP
B. RIP
C. IGRP
D. BGP

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Border Gateway Protocol (BGP) routes that are redistributed into OSPF will be marked with a metric of 1 if no other metric is specified. All other routing protocols will receive a metric of 20 when redistributed into OSPF.

A metric can be manually specified when redistributing the route, as shown below:

```
router5(config)# router ospf 10
router5(config-router)# redistribute bgp 120 metric 5
```

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Cisco Press > Articles > Network Technology > General Networking > Cisco OSPF Route Redistribution
Cisco > Support > Technology Support > IP > IP Routing > Design > Design Technotes > Redistributing Routing Protocols > Document ID: 8606

QUESTION 37
Examine the sample output of the show ip eigrp topology command.

```
Router2# show ip eigrp topology
IP-EIGRP Topology Table for process 100

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

P 10.10.0.0/16, 1 successors, FD is 2169856
via Connected, Serial0
P 10.0.0.0/8, 1 successors, FD is 2169856
via Summary (2169856/0), Null0
P 10.20.0.0/16, 1 successors, FD is 2169856
via Connected, Serial1
P 65.0.0.0/8, 1 successors, FD is 2297856
via 10.20.0.1 (2297856/128256), Serial1
P 192.168.10.0/24, 1 successors, FD is 2297856
via 10.10.0.1 (2297856/128256), Serial0
P 130.10.0.0/16, 1 successors, FD is 2297856
via 10.20.0.1 (2297856/128256), Serial1
P 150.10.0.0/16, 1 successors, FD is 2297856
via 10.10.0.1 (2297856/128256), Serial0
P 200.10.10.0/24, 1 successors, FD is 128256
via Connected, Loopback1
```

The network 65.0.0.0 is one of the advertised networks in the routing table. What does the value 128256 represent?

A. The advertised distance
B. The feasible distance
C. The administrative distance
D. The hop count

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The number 128256 after the advertisement for network 65.0.0.0 represents the advertised distance. The advertised distance is the metric that the neighboring router advertised to the local router.

The feasible distance is the metric that the local router would advertise to the next router. Feasible distance is represented by the number preceding the advertised distance number in the output.

The administrative distance is a number that represents the trustworthiness of a routing protocol. It allows a router to decide which routing protocol's route to use in the event that more than one protocol advertises a route to the same network. The administrative distance is not shown in the output of the show ip eigrp topology command.

Hop count is a simple metric that RIP uses to compare multiple routes to the same network.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify policy-based routing

References:
Cisco > Home > Support > Technology Support > IP Routing > Technology Information > Technology White Paper > Enhanced Interior Gateway Routing Protocol > Feasible Distance, Reported Distance, and Feasible

Successor
Cisco > Cisco IOS IP Routing: EIGRP Command Reference > show ip eigrp topology

QUESTION 38
Consider the following commands:

```
RouterA(config)# router ospf 10
RouterA(config-router)# redistribute eigrp 20 metric 30
```

What does the value of 30 represent?

A. It identifies the seed metric associated with OSPF routes that are redistributed into EIGRP.
B. It identifies the seed metric associated with EIGRP routes that are redistributed into OSPF.
C. It identifies the amount that the existing EIGRP metric will increment as it is redistributed into OSPF.
D. It specifies that routes that contain metrics of less than 30 will be redistributed from OSPF into EIGRP.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The value 30 represents the seed metric for routes that are redistributed from EIGRP into OSPF.

When configuring the OSPF process, the redistribute command is used to identify the source protocol, its AS or process ID, and several other optional parameters, such as metric. The default seed metric for all routing protocols except BGP is 20. When redistributing BGP, the default seed metric is 1.

It does not identify the seed metric associated with OSPF routes that are redistributed into EIGRP. The command is redistributing EIGRP into OSPF, not OSPF into EIGRP.

It does not identify the amount that the existing EIGRP metric will increment as it is redistributed into OSPF. A seed metric value is an absolute value not incremental.

It does not specify that routes that contain metrics of less than 30 will be redistributed from OSPF into EIGRP. It not used to filter routes.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > redistribute (ip)

QUESTION 39
Which of the following statements is NOT true about BGP peers?

A. eBGP peers use TCP to communicate
B. eBGP peers use port 179 by default
C. eBGP peers do not update the AS_Path attribute before sending updates to another eBGP peer
D. iBGP peers do not update the AS_Path attribute before sending updates to an iBGP peer

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
External BGP (eBGP) peers do update the AS_Path attribute before sending updates to another eBGP peer. This helps to maintain the path back to the source of the update.

eBGP peers use TCP to communicate, and they do so on port 179 by default.

Internal BGP (BGP) peers are routers that reside in the same AS. iBGP peers do not update the AS_Path attribute before sending updates to an iBGP peer. That is only done when an update is sent from an eBGP peer to another eBGP peer.

Objective:
Layer 3 Technologies
Sub-Objective:
Explain BGP attributes and best-path selection

References:
Home > About Cisco > Publications and Merchandise > The Internet Protocol Journal > Back issues > Volume 9,Number 1, March 2006 > Autonomous System Numbers > Exploring Autonomous System Numbers

QUESTION 40
Which parameter does EIGRP use to compute the bandwidth part of the metric?

A. The maximum bandwidth link in the path, in kilobits per second
B. The minimum bandwidth link in the path, in kilobits per second
C. The average bandwidth of all the links in the path, in kilobits per second
D. The average bandwidth of all the links in the path, in kilobytes per second

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The minimum bandwidth link, in kilobits per second, is used in the EIGRP metric calculation, because this is the limiting factor in the overall speed of delivery over the path.

BW = (10,000,000 / bandwidth in Kbps) x 256
Delay = (delay in microseconds / 10) x 256

The formula for calculating the EIGRP metric is shown below:

[K1 x BW + (K2 X BW) / (256 - load) + K3 x delay] X [K5 / (reliability + K4)]

You should note, however, that when K5 = 0 (as it is by default), a slightly different formula applies.

When K5 = 0, the EIGRP metric is [K1 X BW + (K2 X BW) / (256 - load) + K3 x delay]

By default, K1 = 1, K3 = 1, and K2, K4, and K5 = 0.

Therefore, the default EIGRP metric is BW + Delay, where "BW" and "Delay" are determined according to the formula above.

The final formula is shown below:

[10,000,000 / (bandwidth in Kbps) + (delay in microseconds) / 10] * 256

These usually are derived from the values listed in the show interfaces command.

Objective:

Layer 3 Technologies
Sub-Objective:
Describe and optimize EIGRP metrics

References:
Cisco > Home > Support > Technology Support > IP Routing > Technology Information > Technology White Paper > Enhanced Interior Gateway Routing Protocol > Using The Metrics

QUESTION 41
You have two routers connected to each other that are both running the EIGRP protocol. The routers have built a neighbor relationship and are exchanging routing information. You execute the following command on the EIGRP process on Router 1:

router1(config)# router eigrp 100
router1(config-router)# passive-interface

What will be the effect of this command?

A. Only routing advertisements from Router 1 to Router 2 will be prevented.
B. Only router advertisements to and from Router 1 will be prevented.
C. All hellos and routing updates will be prevented, and the neighbor relationship between Router 1 and Router 2 will be broken.
D. Hellos will be prevented, but routing updates will continue to be sent out.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The effect of the passive-interface command is dependent on the routing protocol. With RIP, the command prevents the sending of route updates, but does not prevent the reception of route updates. With EIGRP, the passive-interface command prevents both the sending and receiving of route updates, and also the sending of hellos. Without hello packets, the routers are unable to maintain the neighbor relationship, upon which all communications including route updates depend.

If the intent was to preventing routing updates from Router 1 to Router 2 while still allowing updates from Router 2 to Router 1, the routing updates must be filtered out and denied on Router 1 with a distribute list, as shown in the following command set:

router1(config)access-list 101 deny any
router1(config)#router eigrp 100
router1(config-router)distribute-list 101 out

Objective:
Layer 3 Technologies
Sub-Objective:
Troubleshoot passive interfaces

References:
Cisco IOS Master Command List, Release 12.4T > p through r > passive-interface
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > Filtering Routing Updates on Distance Vector IP Routing Protocols

QUESTION 42
You are the network administrator for a corporate organization. You changed the BGP configuration, then executed the following command on the rtrA router:

```
clear ip bgp 172.161.18.5 soft out
```

What is the result of this command?

A. The outbound session between rtrA and 172.161.18.5 is cleared and reset.
B. The inbound session between rtrA and 172.161.18.5 is cleared and reset.
C. The outbound session between rtrA and 172.161.18.5 is cleared.
D. The inbound session between rtrA and 172.161.18.5 is cleared.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The outbound TCP session between rtrA and 172.161.18.5 is cleared as a result of the given command. The given command is a variation of the clear ip bgp command.

The clear ip bgp command allows you to clear and reset the sessions or routing updates in BGP routers so that changes in the BGP configuration can take effect. You can use this command to clear and reset the sessions for all neighbors, a specific neighbor, or a group of neighbors. Use an asterisk (*) or the group name instead of the IP address to apply the command on all the neighbors of a router or a particular peer group, respectively.

For example, if you execute the clear ip bgp * command, all the sessions currently active are cleared and reset. If you use the clear ip bgp 172.161.18.5 command on rtrA, the current session between rtrA and its neighbor 172.161.18.5 is cleared and reset. Such a reset of sessions is known as hard reset. When hard resets are performed, the neighbor relationship is broken and must be reestablished.

The soft keyword, which is optional, indicates a soft reset. This keyword allows you to clear the BGP table without resetting the session. If you do not use this keyword, the sessions are cleared and then reset with a hard reset.

The out keyword specifies that the command should be applied to only outbound sessions. If you use the in keyword, the command is applied to only inbound sessions.

The outbound TCP session between rtrA and 172.161.18.5 is not cleared and reset by the given command. If the clear ip bgp 172.161.18.5 out command was used, then the outbound session between rtrA and 172.161.18.5 would be both cleared and reset.

The inbound TCP session between rtrA and 172.161.18.5 is not cleared and reset by the given command. If the clear ip bgp 172.161.18.5 in command were used, then the inbound TCP session between rtrA and 172.161.18.5 would be cleared and then reset.

The inbound TCP session between rtrA and 172.161.18.5 is not cleared by the given command. If the in keyword were used instead of the out keyword in the given command, the outbound TCP session between the rtrA and 172.161.18.5 would be cleared.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe, configure, and verify BGP peer relationships and authentication

References:
Cisco IOS IP Routing: BGP Command Reference > clear ip bgp

QUESTION 43
The network administrator has configured router R2 to redistribute a newly installed EIGRP network into their core OSPF network. The redistributed networks and subnets are not properly appearing in the routing tables of the other routers. The following output displays partial configuration for router R2:

```
router ospf 10
redistribute eigrp 50 metric 100 metric-type 1
network 192.16.31.0 0.0.0.255
```

What two modifications would correct the problem? (Choose two.)

A. Change the EIGRP AS number from 50 to 10
B. Change the AS number specified for OSPF to 50
C. Add the command network 10.0.0.0 0.0.0.255
D. Add the command network 10.0.0.0 255.255.255.0
E. Add the level-1-2 keyword to the redistribute command
F. Add the subnets keyword to the redistribute command
G. Change the command network 192.16.31.0 0.0.0.255 to include the area keyword and value

Correct Answer: FG
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The R2 router will not form adjacencies with neighboring routers in the area if the area IDs do not match. The area keyword in the network command is missing from the initial router R2 configuration. The correct command would be:

```
R2(config)# network 192.16.31.0 0.0.0.255 area 1
```

Secondly, the subnets keyword should be used in the redistribute command to ensure that all of the subnets in the 10.0.0.0/8 are redistributed into OSPF. For example, you would use the following commands to redistribute EIGRP autonomous system (AS) 50 networks and subnetworks into OSPF with a metric of 100 and advertise them as external Type 1 routes:

```
R2(config)# router ospf
R2(config-router)# redistribute eigrp 50 metric 100 metric-type 1
```

The complete syntax for the redistribute command when used in OSPF is as follows:

redistribute protocol [process-id] [metric metric-value] [metric-type type-value] [subnets]

The command parameters are:

- protocol - Identifies the source protocol, such as BGP, connected, EIGRP, IGRP, ISIS, OSPF, static, or rip.
- process-id - Depending on the routing protocol, identifies the source autonomous system number or process ID.
- metric - Identifies the seed metric for the redistributed route. The default is 0.
- metric-type - For OSPF, it identifies the redistributed routes as either external Type 1 or Type 2 routes. The default is Type 2.
- subnets - Optional keyword for use with OSPF to indicate that the scope of the networks to be redistributed also includes subnets.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > redistribute (ip)

QUESTION 44

A neighboring EIGRP router fails. Its advertised distance (AD) to network 10.10.10.0 was 2 and the feasible distance (FD) was 3.

Which route will be used to route packets destined for network 10.10.10.0 if the other routes have the following feasible and advertised distances respectively to the destination network?

A. FD-6
   AD-3
B. FD-4
   AD-1
C. FD-5
   AD-3
D. FD-4
   AD-3

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When EIGRP loses it best route, called the successor route, it will then use a feasible successor route, if available, to route the packets to that destination. To be considered a feasible successor, the advertised distance, which is the neighboring router's distance, needs to be less than the feasible distance, which is the local router's own metric.

In this scenario, the feasible distance is 3. The only available feasible successors are the ones that have the advertised distance/feasible distance of 1/4 and 2/4.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify policy-based routing

References:
Cisco > Home > Support > Technology Support > IP Routing > Technology Information > Technology White Paper > Enhanced Interior Gateway Routing Protocol > Feasible Distance, Reported Distance, and Feasible Successor

QUESTION 45
Which command can you use to display information about OSPF virtual links?

A. debug ip ospf adj
B. show ip ospf [process-id]
C. show ip ospf virtual-links
D. show ip ospf border-routers

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct answer is show ip ospf virtual-links. The show ip ospf virtual-links command displays the current state of OSPF virtual links, as shown below.

```
Router10# show ip ospf virtual-links
Virtual Link to router 192.168.15.7 is up
Transit area 0.0.0.1, via interface Ethernet1, Cost of using 10
Transmit Delay is 1 sec, State POINT_TO_POINT
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

The following additional commands are available to verify OSPF configurations: show ip ospf border-routers, debug ip ospf adj, and show ip ospf.

The show ip ospf border-routers command displays internal OSPF routing table entries for an ABR, as shown below.

```
router10#show ip ospf border-routers
Codes: i - Intra-area route, I-Inter-area route

Type Dest Address Cost NextHop Interface ABR ASBR Area SPF
i 2.2.2.2 10 192.1.1.199 Ethernet 2 TRUE FALSE 0 3
i 3.2.2.2 10 192.1.1.200 Ethernet 2 TRUE FALSE 0 3
```

The show ip ospf command displays information about the router's role and each area to which the router is connected, as shown below.

```
router10# show ip ospf
Routing Process "ospf 3" with ID 15.0.0.1 and Domain ID 15.20.0.1
Supports only single TOS(TOS0) routes
Supports opaque LSA
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 100 secs
Interface flood pacing timer 55 msecs
Retransmission pacing timer 100 msecs
Number of external LSA 0. Checksum Sum 0x0
Number of opaque AS LSA 0. Checksum Sum 0x0
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. 2 normal 0 stub 0 nssa
External flood list length 0
Area BACKBONE(0)
Number of interfaces in this area is 3
Area has message digest authentication
SPF algorithm executed 4 times
Area ranges are
Number of LSA 4. Checksum Sum 0x29BEB
Number of opaque link LSA 0. Checksum Sum 0x0
Number of DCbitless LSA 3
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
Area 172.16.40.0
Number of interfaces in this area is 0
Area has no authentication
SPF algorithm executed 1 times
Area ranges are
192.168.0.0/16 Passive Advertise
Number of LSA 1. Checksum Sum 0x44FD
Number of opaque link LSA 0. Checksum Sum 0x0
Number of DCbitless LSA 1
Number of indication LSA 1
Number of DoNotAge LSA 0
Flood list length 0
```

The debug ip ospf adj command displays information about the state of neighbor adjacencies, as shown below.

R3#debug ip ospf adj
OSPF adjacency events debugging is on

```
00:54:04: OSPF: Rcv pkt from 172.12.23.2, Ethernet0, area 0.0.0.1 : src not on
the same network
```

In the above example, either the IP address or the subnet mask is misconfigured on either this router or the neighbor.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco > Cisco IOS IP Routing Protocols Command Reference > IP Routing Protocol-Independent Commands: S through T > show ip ospf virtual-links

QUESTION 46
View the sample output of the debug ip eigrp command.

```
IP-EIGRP: Processing incoming REPLY packet
IP-EIGRP: Int 10.20.0.0/16 M 4294967295 - 1657856 4294967295 SM 4294967295 - 1657856 4294967295
IP-EIGRP: Int 65.0.0.0/8 M 4294967295 - 1657856 4294967295 SM 4294967295 - 1657856 4294967295
IP-EIGRP: Int 130.10.0.0/16 M 4294967295 - 1657856 4294967295 SM 4294967295 - 1657856 4294967295
```

What is the significance of the number 4294967295 as shown in the output?

A. It represents the unreachable metric for EIGRP.
B. It represents the administrative distance for EIGRP.
C. It represents a reachable metric for the given network.
D. It represents one of the link characteristics that EIGRP uses to calculate the metric.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The value 4294967295 in the debug ip eigrp output represents the unreachable metric for EIGRP. This means that the network has become unavailable and cannot be reached. In this output, the M represents the local metric, and the SM represents the metric that was reported by the neighbor that advertised the network to the local router.

The administrative distance (AD) for internal EIGRP is 90.

The link characteristics that are used in the EIGRP calculation are shown following the dash after the M and SM values (1657856 4294967295). By default, EIGRP only uses bandwidth and delay in its calculation.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe and optimize EIGRP metrics

References:
Cisco > Cisco IOS Debug Command Reference > debug h225 asn1 through debug ip ftp > debug ip eigrp

QUESTION 47
For a non-ISP autonomous system (AS), which combination of two conditions will require redistribution from BGP into Interior Gateway Protocol (IGP)? (Choose two.)

A. All routers run BGP.
B. Not all routers run BGP.
C. No knowledge of external routes is required inside the AS.
D. Routers inside the AS require knowledge of external routes.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
For non-ISP autonomous systems (AS), redistribution into IGP is required when BOTH of the following conditions exist:
- Not all routers run BGP

- Knowledge of external routes is required inside the AS

Note: Redistribution of any BGP routes into your IGP can cause serious problems, even if your internal routers can handle the load. This should be done rarely, if at all. If you do decide to do this, configure heavy filtering to allow only very few routes into OSPF or EIGRP so as not to overwhelm those protocols. For instance, do it only for a select group of networks for which optimal routing is critical.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Cisco > Support > Technology Support > IP > IP Routing > Design > Design Technotes > BGP Case Studies > Document ID: 26634 > Static Routes and Redistribution
Cisco > Cisco IOS IP Routing: BGP Configuration Guide, Release 12.4 > Cisco BGP Overview > Information About Cisco BGP > BGP Autonomous Systems

QUESTION 48
Your network has an OSPF area that connects to an EIGRP area at two points, Router A and Router B. You directed your assistant to set up these two routers in order to have traffic load-balanced between the two entry points. However, you discover that all traffic is going through Router A. When you view the results of the show run command for each device, you get the partial output shown below:

```
Hostname routerA
!
Router ospf 200
Redistribute eigrp 50 metric 20 metric-type 2 subnets
Network 192.168.8.0 255.255.255.0 area 0
<<output omitted>>


Hostname routerB
!
Router ospf 200
Redistribute eigrp 50 metric 50 metric-type 2 subnets
Network 192.168.8.0 255.255.255.0 area 0

<<output omitted>>
```

What action should be performed to make traffic use both routes to the EIGRP area?

A. change the metric for EIGRP to 50 on Router A
B. change the metric for EIGRP to 45 on Router B
C. change the metric type to Type 1 on Router A
D. change the metric type to Type 1 on Router B

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You should change the metric for EIGRP to 50 on Router A. The metric can be defined when configuring the redistribution of one routing protocol into another. A lower metric will cause traffic to be routed in that direction. Therefore, to make the paths from the two routers equal, you should raise the metric on Router A to 50 to match that of Router B.

You should not lower the metric on Router B to 45. It will still be a higher metric than that of A and traffic will still go in that direction.

You should not change the metric type on either router. The metric type determines whether the downstream OSPF routers should add their cost to the total cost to get to the ASBR when computing cost. In this scenario, Router A and Router B are both ASBRs. When set to Type 1, downstream OSPF routers do add their metric. With Type 2, they simply use the configured metric. If you want true load balancing, you should leave the metric type set to Type 2, since you have no information on the cost from the other routers to the ASBRs. However, when Type 1 is used, you must also take into consideration the metric from the other routers to the ASBR when determining the true cost to leave the OSPF area.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify redistribution between any routing protocols or routing sources

References:
Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > Redistributing Routing Protocols
Cisco > Support > Technology Support > IP > IP Routing > Design > Design Technotes > Route Selection in Cisco Routers > Document ID: 8651

QUESTION 49
Examine the exhibit.

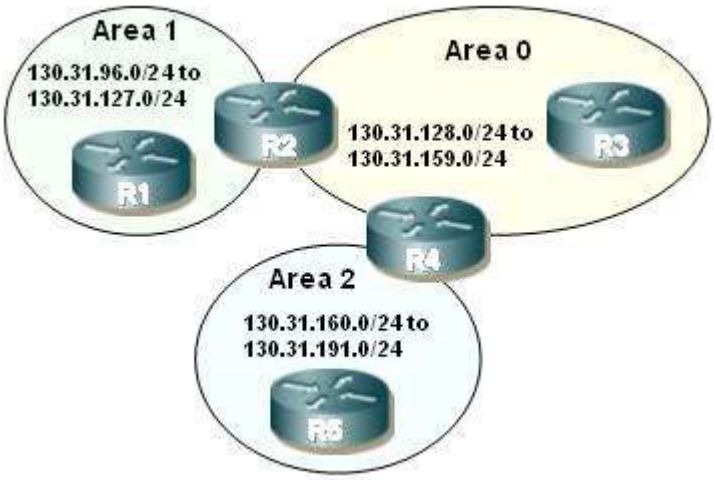

Router R2 has been configured with the following OSPF router command:

area 1 range 130.31.96.0 255.255.224.0

Which addresses listed will be summarized by R2 into area 0? (Choose all that apply.)

A. 130.31.128.0/23
B. 130.31.112.0/20
C. 130.31.130.0/24
D. 130.31.160.0/22
E. 130.31.104.0/21

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The command area 1 range130.31.96.0 255.255.224.0 is used to summarize the IP network addresses from 130.31.96.0/24 to 130.31.127.0/24 in area 1. Addresses 130.31.112.0/20 and 130.31.104.0/21 are both in that range of network addresses.

To determine if an address is a part of a summary, put the summary address and summary mask in binary format. Do the same with the address and the summary mask you are examining, as shown below:

130.31.96.0 10000010.00001111.01100000.00000000
130.31.112.0 10000010.00001111.01110000.00000000
255.255.224.0 11111111.11111111.11100000.00000000

If the address you are testing and the summary address agree to the point where the mask stops, then the test address is part of the summary. In this case, they agree through the 27th bit, so this address is a part of the summary. The same is true for the 130.31.104.0 address.

In OSPF, you can only configure summarization on the border routers. The summaries need to be of routes within a single area. You summarize the routes of an area so that routers in another area do not see the individual networks, just the summary. The correct command is:

area area id range ip-address mask

The area id parameter is the number of the area whose networks are being summarized. For example, in the network shown in the exhibit, to summarize the networks within area 1 to 130.31.96.0/19 you would configure router R2 with the command area 1 range 130.31.96.0 255.255.224.0. This would not affect the routing tables of the routers within area 1, but instead make the routing tables of areas 0 and 2 smaller. These other routers would only have the summary route listed instead of the individual networks. Router 3 would only see the summary route 130.31.96.0/19.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Cisco > Home > Support > Technology Support > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > OSPF and Route Summarization > Inter-Area Route Summarization
Cisco IOS Master Command List, Release 12.4 > a through b > area range

QUESTION 50
When an EIGRP router starts, it sends a hello packet out of all interfaces.

Which type of packet do neighboring routers send in response?

A. ACK
B. Hello
C. Query
D. Reply
E. Update

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When an EIGRP router starts, it sends hello packets out of each interface. Neighboring routers respond with

update packets. These update packets are sent reliably, and must be acknowledged with an ACK packet from the EIGRP router.

EIGRP makes neighbor relationships simple. If a router hears a hello from a new neighbor, it sends that neighbor updates for all routes that it knows. This is different from Open Shortest Path First (OSPF), which has a complex series of rules governing how neighbor relationships are formed and how databases are synchronized. When changes to the network occur in OSPF, updates packets route reliable change information only to the affected routes.

Queries and replies in EIGRP only occur when a router loses a route to a network and is actively seeking a replacement route.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe EIGRP packet types

References:
Internetworking Technology Handbook > Enhanced Interior Gateway Routing Protocol (EIGRP) > EIGRP Packet Types

QUESTION 51
You are using an aggregate static route to null 0 to redistribute static routes into BGP.

Which problem can result if the router loses access to one of these routes?

A. Black hole
B. Routing loop
C. Split horizon
D. Unstable BGP table

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If one of the aggregated routes is lost, the router will discard packets destined for that route. This condition is known as a black hole.

For example, suppose you have a number of subnets of range 11.1.0.0/16, all of which have 24 bit masks, such as 11.1.2.0/24. You aggregate them all to 11.1.0.0/16 and advertise that aggregate. If this router were to lose connectivity to one of the subnets, for example 11.1.3.0/24, then any traffic routed through this router to that subnet would never reach it, even if there were another valid path.

Split horizon is a loop avoidance mechanism that is by default always in effect, and is not affected by the loss of a subnet route that is part of an aggregate route.

BGP tables are not made unstable by the loss of the loss of a subnet route that is part of an aggregate route.

Routing loops would not occur simply from the loss of a subnet route that is part of an aggregate route.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Cisco > IP Routing: BGP Configuration Guide > BGP4 > Aggregating Route Prefixes Using BGP

QUESTION 52
You have configured BGP on both rtrA in AS 1 and rtrB in AS 2. There are two routes created using the network command between the two routers. One route traverses through AS 5 and AS 6 from rtrA to rtrB, while the other route traverses AS 7, AS 8, and AS 9 from rtrA to rtrB. Both routes use default values for the Weight and LOCAL_PREF attributes.

Which of the following attributes determines the BEST route between rtrA and rtrB routers?

A. Weight
B. LOCAL_PREF
C. Origin type
D. AS_PATH

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The AS_PATH attribute is used to determine the best path between the two routes. To select the best path from rtrA to rtrB, BGP analyzes attributes that are set for the two available routes during the configuration of the network. The key BGP attributes and the order in which they are checked are as follows:
1. Weight - highest weight is selected
2. LOCAL_PREF - highest LOCAL_PREF is selected
3. Locally originated routes - local routes are selected
4. AS_PATH - shortest AS_PATH is selected
5. Origin type - lowest origin type is selected
6. Multi-exit Discriminator (MED) - lowest MED is selected

The weight attribute is the first attribute to be checked while selecting the best BGP route. This attribute is relevant only to the local router on which it is set. The value of this attribute can be any number from 0 to 65535. The default values are 32768 for locally originated routes and 0 for other types of routes. Both routes in this case are originated locally and have the default weight values. Therefore, in this case, the weight attribute is not used to determine the best route.

BGP then checks the value of the LOCAL_PREF attribute, which refers to local preference. Local preference is a value indicates the route that is preferred to exit the AS to reach a given network. Routes with higher local preference are selected by BGP. You can set the local preference for a route to any value; however, if you do not, the route uses the default value of 100. Because both routes have the default LOCAL_PREF value, this attribute is not used to determine the best route.

Next BGP checks whether any of the routes are locally originated using the network, redistribute, or aggregate commands. As stated, both routes originated using the network command on the routers. Consequently, BGP analyzes the value of the AS_PATH attribute, which is a list of the AS numbers traversed by a particular route. The route with the shortest AS_PATH is selected as the best path. In this case, the route that consists of AS 5 and 6 is considered the best path because the AS_PATH value for this route is shorter than that for the route traversing AS 7, 8, and 9. The AS_PATH value for the route traversing AS 5 and 6 is [6 5 1], while the AS_PATH for the route traversing AS 7, 8, and 9 is [9 8 7 1].

The other options are incorrect because the corresponding attributes are same for both the routes; hence, those attributes are not considered while BGP determines the best path.

Objective:
Layer 3 Technologies
Sub-Objective:
Explain BGP attributes and best-path selection

References:

Internetworking Technology Handbook > BGP > BGP attributes

QUESTION 53
An OSPF area contains the following networks:

165.164.8.0 255.255.254.0
165.164.10.0 255.255.254.0
165.164.12.0 255.255.254.0
165.164.14.0 255.255.254.0

How can the route to these networks be summarized?

A. 165.164.8.0 255.255.240.0
B. 165.164.8.0 255.255.248.0
C. 165.164.10.0 255.255.252.0
D. 165.164.14.0 255.255.240.0

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Summarization is the process of advertising a network with a subnet mask such that it includes all of the subnets. For a simple example if you had two Class C networks, you could advertise them as a Class B network and it would encompass them both. Normally summarization should be implemented such that it summarizes ONLY the networks desired and no others (in the simple example it would possibly include other Class C networks). The process for arriving at the "best" summarization is a follows.

First, write the last octet that all networks share in common (third octet in this case) in binary form for each network:

165.164.8.0--00001000
165.164.10.0--00001010
165.164.12.0--00001100
165.164.14.0--00001110

The addresses have the first five bits in common; therefore, they can be summarized with the third octet 00001000 and a subnet mask of 255.255.248.0.

Another way of looking at it is that 165.164.8.0 255.255.248.0 covers the range of 165.164.8.0 through 165.164.15.255, the same range as all the component subnets.

None of the following possible answers is a valid range, nor do most of them cover the correct range of addresses:

165.164.8.0 255.255.240.0 is not a valid range. A 20-bit mask can only be on a subnet that is a multiple of 16, such as .16.0, .32.0, .48.0 etc. The subnet .8.0 is not a multiple of 16.

165.164.10.0 255.255.252.0 is not valid. A 22-bit mask requires a multiple of 4 in the third octet, and 10 is not a multiple of four. Even if it were a valid range, it does not cover the entire range of addresses that need to be summarized.

165.164.14.0 255.255.240.0 is not valid. The 20-bit mask is only usable on ranges that are multiples of 16 in the third octet, and 14 is not a multiple of 16. Even if the mask were valid, it could not cover the correct addresses.

When addresses are summarized the cost of the summary address will the highest cost of the component subnets. For example, in the partial sample output of the show ip route command below, there are three routes.

The output is from a router running OSPFv3, so the addresses are IPv6, but the concept is the same.

```
OI 2001:0D B 8:0:0:7/64 [110/20]
     via FE 80::A8BB:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:0D B 8:0:0:8/64 [110/100]
     via FE 80::A8BB:CCFF:FE 00:6F00, FastEthernet0/0
OI 2001:0D B 8:0:0:9/64 [110/40]
     via FE 80::A8BB:CCFF:FE 00:6F00, FastEthernet0/0
```

The routes have metrics (the second value in brackets, [administrative distance/cost]) of 20, 100, and 40. Therefore, the metric for the summarized route would be 100.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Cisco > Home > Support > Technology Support > IP Routing > Technology > Information Technology > White Papers > OSPF Design Guide

QUESTION 54
You have implemented OSPF for IPv6 for the following areas in OSPF AS 1:

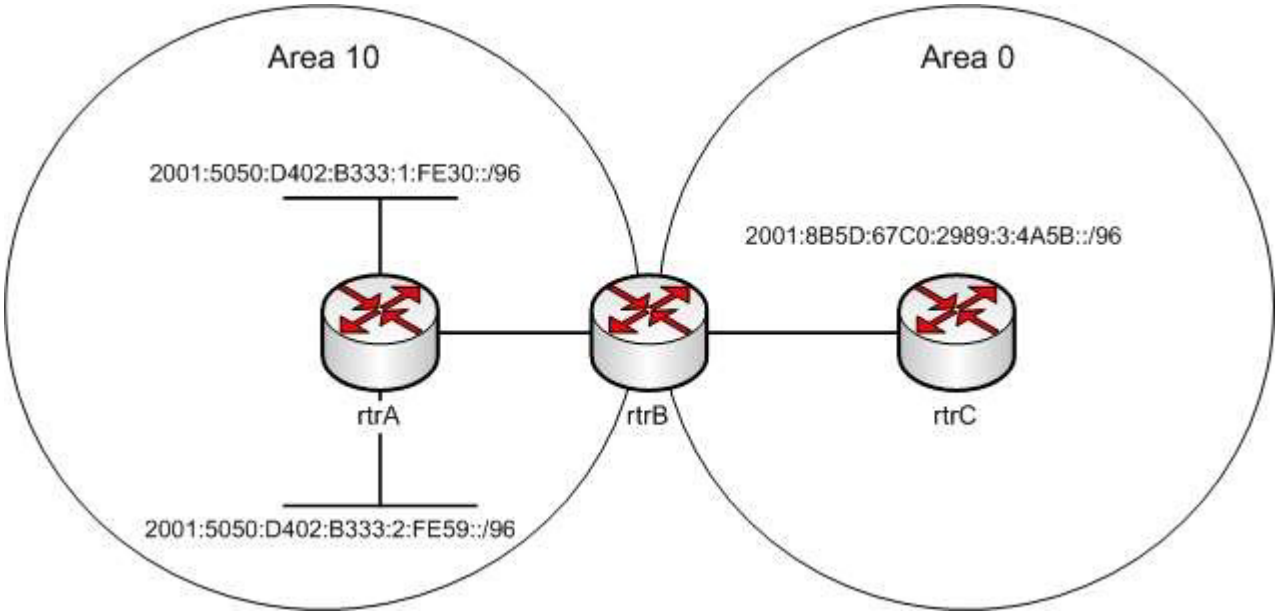

The cost from rtrB to the 2001:5050:D402:B333:1:FE30::/96 network is 80, while the cost from rtrB to the 2001:5050:D402:B333:2:FE59::/96 network is 130.

Which of the following area range cost commands should be executed on rtrB?

A. area 10 range 2001:5050:D402:B333::/64 cost 80
B. area 10 range 2001:5050:D402:B333::/64 cost 130
C. area 10 range 2001:5050:D402:B333::/64 cost 210
D. area 10 range 2001:5050:D402:B333::/64 cost 0

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The area 10 range 2001:5050:D402:B333::/64 cost 130 command should be executed on rtrB. This command defines an area range for an area border router (ABR) that has OSPF for IPv6 enabled on it. This command provides a summary route of the routes in an OSPF area to be distributed to another area.

The range keyword in the command provides the summary route. The cost keyword in the command specifies the cost of the summary route. The highest cost of the routes that are being summarized becomes the cost of the summary route. In this case, the cost from rtrB to the 2001:5050:D402:B333:1:FE30::/96 network is 80, and the cost from rtrB to the 2001:5050:D402:B333:2:FE59::/96 network is 130. The cost of the summary route is 130 as it is higher.

All the other options are incorrect because they do not specify the correct cost of the summary route.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF for IPv6

References:
Cisco Learning Home > Groups > CCNP R&S Study Group > Discussions > What would be the Metric for a Summary Route in OSPFv3
Cisco IOS IPv6 Configuration Guide, Release 12.4 > Implementing OSPF for IPv6 > How to Implement OSPF for IPv6

QUESTION 55
You are configuring Open Shortest Path First (OSPF) protocol for IPv6 on Router5. The router has two interfaces, which have been configured as follows:

S0/0 - 192.168.5.1/24
S0/1 - 10.0.0.6/8

You would like OSPF to route for IPv6 only on the S0/0 network and not route for IPv6 on the S0/1 network. The process ID you have chosen to use is 25. You do not want to apply an IPv6 address yet.

Which of the following command sets would enable OSPF for IPv6 as required?

A.
```
Router5(config)#ipv6 ospf 25
Router5(config)# network 192.168.5.0
```
B.
```
Router5(config)#ipv6 ospf 25
Router5(config)#router-id 192.168.5.1
```
C.
```
Router5(config)#ipv6 unicast-routing
Router5(config)#ipv6 router ospf 25
Router5(config-rtr)#router-id 1.1.1.1
Router5(config)#interface S0/0
Router5(config-if)#ipv6 ospf 25 area 0
```
D.
```
Router5(config)#ipv6 unicast-routing
Router5(config)#ipv6 ospf 25
Router5(config-rtr)#router-id 1.1.1.1
```

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct command sequence would be as follows:

```
Router5(config)# ipv6 unicast-routing
```

```
Router5(config)# ipv6 router ospf 25
Router5(config-rtr)# router-id 1.1.1.1
Router5(config)# interface S0/0
Router5(config-if)# ipv6 ospf 25 area 0
```

The first line enables IPv6 routing with the ipv6 unicast-routing command. The second line enables OSPF routing for IPv6 with the ipv6 router ospf command. The third assigns a necessary router ID (which was chosen at random) with the router-id command. The last two lines enable OSPF for area 0 on the proper interface.

The following command set is incorrect because it does not enable OSPF routing for IPv6, assign a necessary router ID, or enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 ospf 25
Router5(config)# network 192.168.5.0
```

This command set also displays incorrect use of the network command. The network command would be used with OSPF v2.

The following command set fails to enable OSPF routing for IPv6, assign a necessary router ID, or enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 ospf 25
Router5(config)# router-id 192.168.5.1
```

It also assigns the router ID under global configuration mode, rather than under router ospf 25 configuration mode as required.

The following command set fails to enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 unicast-routing
Router5(config)# ipv6 ospf 25
Router5(config-rtr)# router-id 1.1.1.1
```

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF for IPv6

References:
Cisco > Implementing OSPF for IPv6 > How to Implement OSPF for IPv6
Cisco > Cisco IOS IPv6 Command Reference > ipv6 unicast-routing
Cisco > Cisco IOS IPv6 Command Reference > ipv6 ospf area

QUESTION 56
Refer to the following exhibit, where three routers have EIGRP for IPv6 enabled on them:

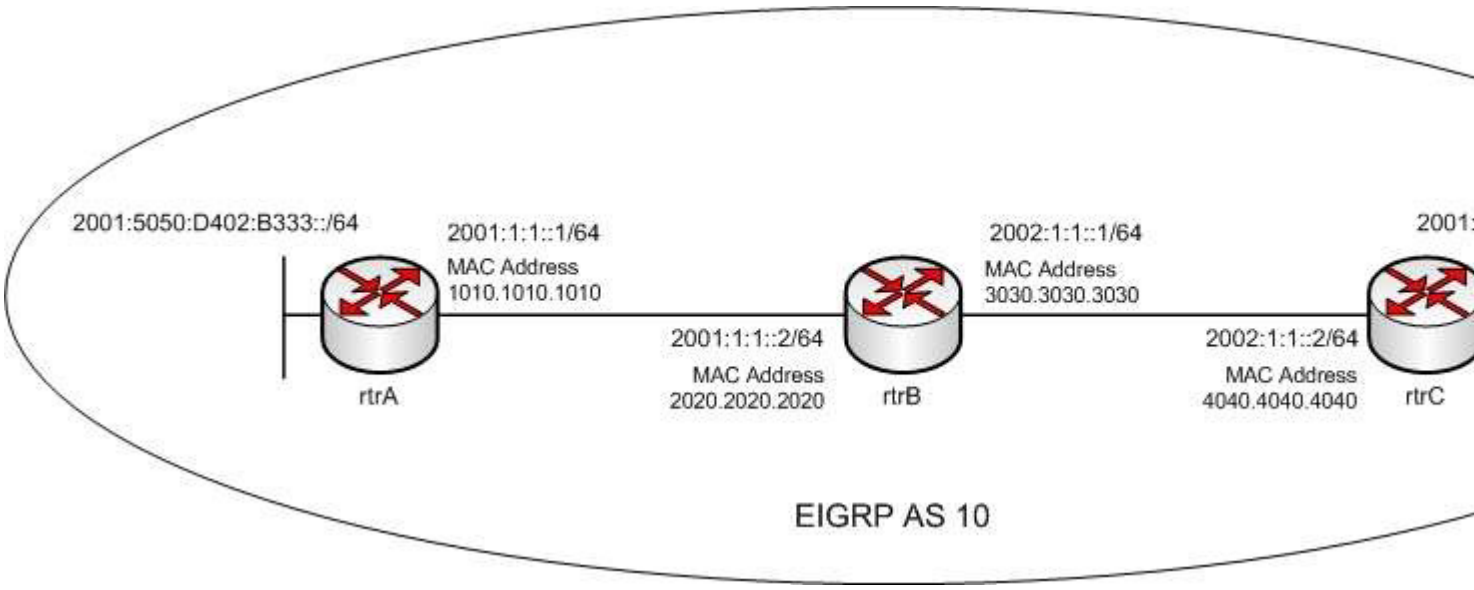

What is the next-hop address when rtrB advertises the 2001:5050:D402:B333::/64 IPv6 subnet to rtrC?

A. FE80::3030:3030:3030/64
B. FE80::3230:3030:3030/64
C. FE80::3030:30FF:FE30:3030/64
D. FE80::3230:30FF:FE30:3030/64

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The next-hop address when rtrB advertises the 2001:5050:D402:B333::/64 IPv6 subnet to rtrC is FE80::3230:30FF:FE30:3030/64. In routers with EIGRP for IPv6 enabled on them, the next-hop address is the IP address of the interface that advertises routes. The next-hop addresses should be link-local addresses. Link-local addresses are IPv6 unicast addresses that are automatically assigned to the router interfaces. These addresses have the FE80::/10 prefix and the EUI-64 standard interface address.

EUI-64 is an IEEE standard that allows the determination of an IPv6 address from the MAC address of an interface. The 64 most significant bits should be specified in the ipv6 address command. The 64 least significant bits are determined by using the EUI-64 standard. The steps to determine the 64 least significant bits are as follows:
1. Divide the 48-bit MAC address into two 24-bit parts.
2. Embed FFFE (16 bits) between the two parts resulting in a 64-bit address.
3. If required, toggle the seventh bit of the first octet in the address. In EUI-64 format, if the seventh bit is 0, then the address is local. If the seventh bit is 1, the address is global.

In this case, when rtrB advertises any route to rtrC, it advertises the interface with the MAC address 3030.3030.3030 as the next-hop. When the given steps are performed on the MAC address, it result in the EUI-64 standard address 3230.30FF.FE30:3030. This address is then appended to the FE80::/10 prefix. The resultant IPv6 link-local address of the interface is FE80::3230.30FF.FE30:3030/10.

The remaining three options are incorrect as their interface address is not in the EUI-64 standard.

Objective:
Layer 3 Technologies
Sub-Objective:

Identify IPv6 addressing and subnetting

References:
Cisco IPv6 Configuration Guide, Release 15.2 > IPv6 Neighbor Redirect Message
Cisco IPv6 Configuration Guide, Release 15.2 > IPv6 Unicast Routing > Aggregatable Global Address

QUESTION 57
You have configured OSPF on your network and enabled route summarization on an area border router (ABR) with the following command:

Router(config-router)# area 3 range 165.164.8.0 255.255.248.0

What does the 3 specify in this command?

A. The ID of the OSPF backbone
B. The number of networks summarized in the area
C. The ID of the area about which routes will be summarized
D. The ID of the area to which the summary route information will be sent

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The 3 in the area range command specifies the area that contains the routes that are to be summarized. In OSPF, you can only configure summarization on the border routers. The summaries need to be of routes within a single area. You summarize the routes of an area so that routers in another area do not see the individual networks, just the summary. The correct command syntax is shown below:

area number range ip-address mask

The number parameter is the number of the area whose networks are being summarized. For example, in the network shown in the graphic below, to summarize the networks within area 2 to 10.1.0.0/16, you would configure router A with the command area 2 range 10.1.0.0 255.255.0.0. This would not affect the routing tables of the routers within area 2, but instead make the routing tables of areas 0 and 1 smaller. These other routers would only have the summary route listed instead of the individual networks. Router C would only see the summary route 10.1.0.0/16.

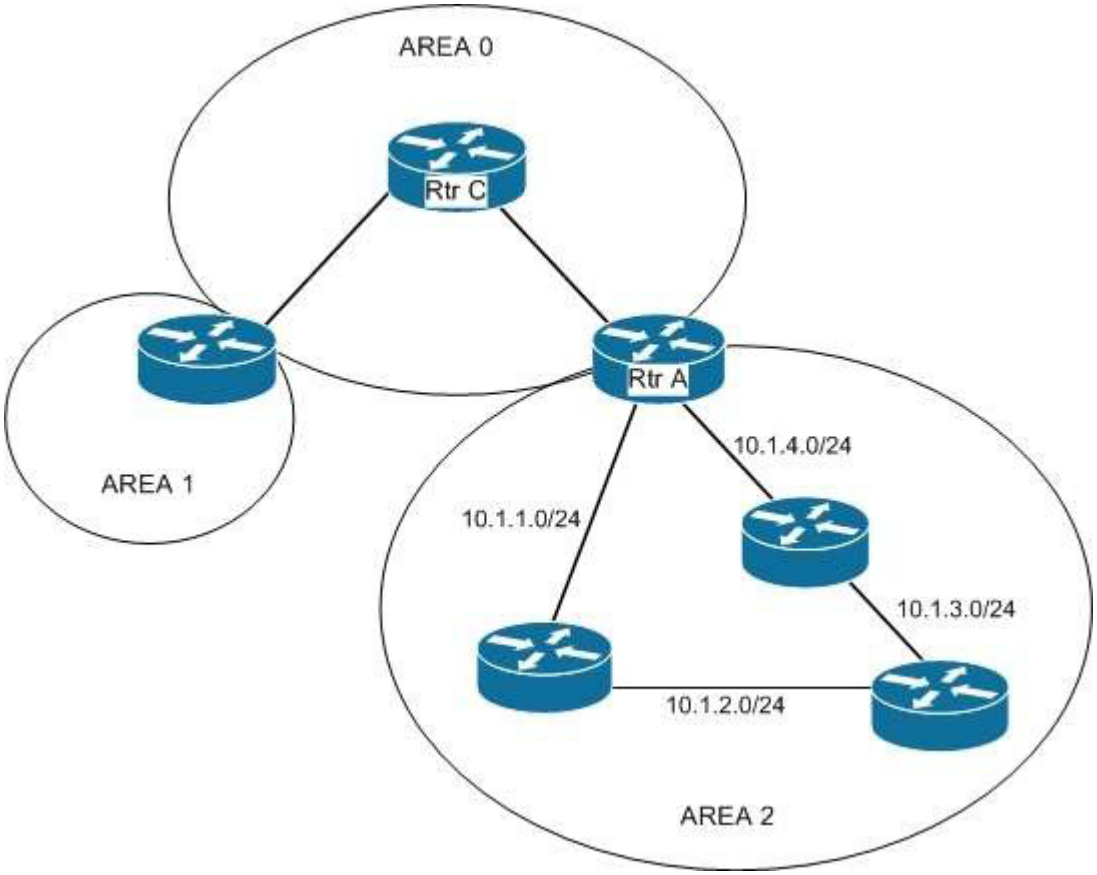

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify manual and autosummarization with any routing protocol

References:
Cisco IOS Master Command Reference > a through b > area range

QUESTION 58
Which of the following commands is used to verify the link-local, global unicast, and multicast addresses of an IPv6 router?

A. show ipv6 neighbors (only link-local addresses)
B. show ipv6 route
C. show ipv6 protocols
D. show ipv6 interface

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The show ipv6 interface command is used to verify the link-local, global unicast, and multicast addresses assigned to an IPv6-enabled router interface. The show ipv6 interface command displays information regarding that interface, such as the physical state, MTU, and IPv6 enable/disable state.

A partial output of the show ipv6 interface command on an IPv6-enabled router named rtrA is as follows:

```
rtrA# show ipv6 interface FastEthernet 0/1
FastEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::6339:7BFF:FE5D:A031/64
  Global unicast address(es):
    2001:7067:90D1:1::1, subnet is 2001:7067:90D1:1/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF5D:A031
  MTU is 1500 bytes
  <output omitted>
```

In the given sample output, you can see that the Fa0/1 interface of rtrA has the link-local address FE80::6339:7BFF:FE5D:A031/64 and the global unicast address 2001:7067:90D1:1::1. The global unicast address is not in EUI-64 format because when the ipv6 address command was issued, the eui64 keyword was not used. If EUI-64 format had been specified with the eui64 keyword, the global unicast address would have been 2001:7067:90D1:1:6339:7BFF:FE5D:A031.

An IPv6-enabled interface has not only a link-local and global unicast address, but also one or more multicast addresses. A multicast address is an IPv6 address that has the prefix FF00::/8. These addresses are assigned to interfaces of different nodes such that they appear as a logical group. This implies that when a packet is destined for a multicast address, that packet is delivered to all the interfaces that have the same multicast address. The various multicast groups are as follows:

- FF02::1 Indicates the group of all the nodes on the local segment
- FF02::2 Indicates the group of all the routers on the local segment
- FF02::1:FF00:0/104 Indicates a solicited-node multicast group for every unicast or anycast address assigned to the interface

You can also notice in the sample output that the Fa0/1 interface belongs to three multicast groups: FF02::1, FF02::2, and FF02::1:FF5D:A031. The first two multicast groups refer to the all-host and all-router multicast groups, respectively. The third group, FF02::1:FF5D:A031, is the solicited-node multicast address. This address is created for every unicast or anycast address. A solicited-node multicast address is determined by assigning the least significant 24 bits of the unicast address to the least significant 24 bits of the FF02::1:FF00:0 address.

The show ipv6 neighbors command displays the link-local /global unicast addresses of the neighbors, including other information such as state and the next-hop interface.

The show ipv6 route command is used to view the IPv6 routing table on the router. This command displays the prefixes, administrative distance, metric, and next-hop addresses for various IPv6 networks.

The show ipv6 protocols command is used to view the active routing protocols for IPv6 on the router. This command shows the interfaces, redistribution status, and summarization status about each of the routing protocols enabled on the router.

Objective:
Layer 3 Technologies
Sub-Objective:
Identify IPv6 addressing and subnetting

References:
Cisco IOS IPv6 Command Reference > show ipv6 eigrp topology through show ipv6 nat statistics > show ipv6 interface
Cisco IOS IPv6 Command Reference > show ipv6 nat translations through show ipv6 protocols > show ipv6 neighbors
Cisco IOS IPv6 Command Reference > show ipv6 nat translations through show ipv6 protocols > show ipv6 protocols
Cisco > Products & Services > Cisco IOS and NX-OS Software > Cisco IOS Technologies > IPv6 > Product Literature > White Papers > Cisco IOS IPv6 Multicast Introduction

Cisco > IPv6 Implementation Guide, Release 15.2M&T > Implementing IPv6 Multicast

QUESTION 59
Which method should you use to block all routing updates from being sent into a network through an interface?

A. Static route
B. Default route
C. Passive interface
D. Route-update filtering

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To stop all outbound routing updates from an interface, you can use the passive-interface command. The effect of the passive-interface command is dependent on the routing protocol running on the interface. For EIGRP, the router will not only stop sending routing updates, but also hellos, which means that it will not form a neighbor relationship with another EIGRP router on that interface. This is also the case with OSPF and IS-IS. With RIP, however, the router will continue to send hellos even as it stops sending routing updates and it will still receive routing updates.

There are numerous reasons to use the passive-interface command. For instance, suppose that you have a LAN with only end hosts on it and no other router. Additionally, there is no reason to send EIGRP packets, but you want EIGRP to advertise that the network can be reached. The combination of a network statement for that interface plus a passive-interface command would be appropriate.

Route-update filtering can block all routing updates if you prefer, but it is really intended for selective filtering of updates. If your goal is to block all updates, the passive-interface command is best.

Default routes and static routes can be used as ways around having to send routing updates out an interface. However, if your goal is to block updates, you should issue the passive-interface command.

Objective:
Layer 3 Technologies
Sub-Objective:
Troubleshoot passive interfaces

References:
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > passive-interface
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > How Does the Passive Interface Feature Work in EIGRP?

QUESTION 60
Which show command displays detailed information about a router's BGP connections to neighboring routers?

A. show ip bgp
B. show ip bgp summary
C. show ip bgp neighbors
D. show ip bgp connections

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The show ip bgp neighbors command will show you detailed information about all of the router's neighbors or peers. A sample of the show ip bgp neighbors output is shown below. The sample utilizes the ip address parameter, which is optional, but can used to limit the output to display information about only one neighbor:

```
Router15# show ip bgp neighbors 10.5.1.6

BGP neighbor is 10.5.1.6, remote AS 11, internal link
BGP version 4, remote router ID 10.1.5.6
BGP state = Active, table version 0
Last read 00:00:12, hold time is 180, keepalive interval is 60 seconds
Minimum time between advertisement runs is 30 seconds
Received 19 messages, 0 notifications, 0 in queue
Sent 17 messages, 0 notifications, 0 in queue
Inbound path policy configured
Route map for incoming advertisements is testing
Connections established 2; dropped 1
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.5.1.7, Local port: 11002
Foreign host: 10.5.1.6, Foreign port: 179

<output omitted>
```

In the above example, router15 has sent out a BGP open packet to the peer at 10.5.1.6 and is listening for a connection request from the peer. This can be determined by the line that says BGP state = Active. It can also be determined that router has established a TCP connection two times, as evidenced by the line Connections established 2.

The show ip bgp command displays the contents of the BGP routing table. It will not display detailed information about a router's BGP connections to neighboring routers.

The show ip bgp summary command displays a summary of the status of BGP connections. It will not display detailed information about a router's BGP connections to neighboring routers.

There is no show ip bgp connections command.

Objective:
Layer 3 Technologies
Sub-Objective:
Describe, configure, and verify BGP peer relationships and authentication

References:
Cisco IOS Master Command List, Release 12.4 > a through b > BGP Commands: show ip through T > show ip bgp neighbors

QUESTION 61
Examine the exhibit.

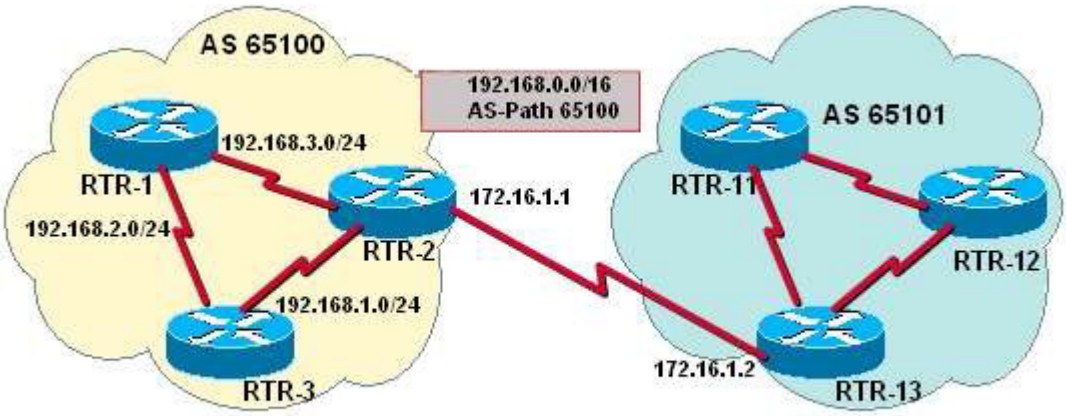

You have determined that RTR2 is not advertising the CIDR summary address 192.168.0.0 to the other routers in AS 65100.

Which set of configuration commands will enable the BGP router RTR2 to announce the network prefix 192.168.0.0/16 to the other routers in the AS 65100?

A.
```
router bgp 65100
neighbor 172.16.1.2 remote-as 65100
neighbor 192.168.3.2 remote-as 65100
network 192.168.3.0
```

B.
```
router bgp 65100
neighbor 172.16.1.2 remote-as 65101
neighbor 192.168.3.2 remote-as 65100
network 192.168.0.0
```

C.
```
router bgp 65100
neighbor 172.16.1.2 remote-as 65100
neighbor 192.168.3.2 remote-as 65100
network 192.168.0.0 mask 255.255.0.0
ip route 192.0.0.0 255.0.0.0 null 0
```

D.
```
router bgp 65100
neighbor 172.16.1.2 remote-as 65101
neighbor 192.168.3.2 remote-as 65100
network 192.168.0.0 mask 255.255.0.0
ip route 192.168.0.0 255.255.0.0 null 0
```

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Issuing the following commands will cause RTR2 to advertise the CIDR block 192.168.0.0/16 to the other routers by using BGP:

```
RTR2(config)# router bgp 65100
RTR2(config-router)# neighbor 172.16.1.2 remote-as 65101
RTR2(config-router)# neighbor 192.168.3.2 remote-as 65100
RTR2(config-router)# network 192.168.0.0 mask 255.255.0.0
RTR2(config-router)# ip route 192.168.0.0 255.255.0.0 null 0
```

The network command specifies the address that will be inserted into the BGP table. Without the mask

keyword, the classful network will be assumed. Because 255.255.0.0, or /16, is not the natural mask for any Class C address, the mask keyword must also be specified. Thus, 192.168.0.0 and 255.255.0.0 identify the desired address and mask of the 192.168.0.0/16 network prefix.

The router checks the IP forwarding table for an exact match before it advertises the route. Without a matching entry in the IP forwarding table, that route will not be advertised. RTR2 must be able to advertise a CIDR block and not the individual subnets. A static route is required because BGP requires that a match of the network prefix be present in the forwarding table when using the network command with the mask keyword. Therefore, to ensure an exact match for the identified prefix exists in the IP forwarding table, and to ensure that the prefix will always be advertised, a static route for 192.168.0.0/16 to null 0 is also required.

The syntax for the network command is shown below:

network network-number [mask network-mask] [route-map map-tag]

The parameters are:
- mask - This parameter is optional and identifies the network or subnetwork to advertise.
- route-map - This parameter is optional and identifies a preconfigured route-map that will be used to filter specific addresses from being advertised.

The following command set is missing the mask keyword in the network command and the command to create a static route to null 0. The address used in the network command is also incorrect. It should 192.168.0.0:

router bgp 65100
neighbor 172.16.1.2 remote-as 65100
neighbor 192.168.3.2 remote-as 65100
network 192.168.3.0

The following command set is missing the mask keyword in the network command and the command to create a static route to null 0:

router bgp 65100
neighbor 172.16.1.2 remote-as 65101
neighbor 192.168.3.2 remote-as 65100
network 192.168.0.0

The following command set uses an incorrect mask (255.0.0.0) in the command that creates the static route to null 0. It should be 255.255.0.0:

router bgp 65100
neighbor 172.16.1.2 remote-as 65100
neighbor 192.168.3.2 remote-as 65100
network 192.168.0.0 mask 255.255.0.0
ip route 192.0.0.0 255.0.0.0 null 0

Objective:
Layer 3 Technologies
Sub-Objective:
Describe, configure, and verify BGP peer relationships and authentication

References:
Internetworking Case Studies > Using the Border Gateway Protocol for Interdomain Routing > Controlling the Flow of BGP Updates > CIDR and Aggregate Addresses > Aggregation and Static Routes

QUESTION 62
Examine the following output.

```
!
router bgp 65100
neighbor 192.168.12.34 remote-as 65101
network 172.16.0.0
no auto-summary
!

RouterA# show ip route 172.16.31.0
o 172.16.31.0/24 [110/128] via 10.1.2.3 00:24:16, Serial0
```

You are investigating why router RouterA does not include the 172.16.0.0 network in its BGP advertisements. The output contains portions of RouterA's configuration and IP routing table.

Which statement correctly identifies the reason why this subnet does not appear in the BGP advertisements?

A. BGP synchronization is enabled by default.
B. The no auto-summary command was used.
C. The 172.16.31.0/24 network was learned through BGP.
D. The 10.1.2.3 IP address was not defined as a BGP neighbor.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The no auto-summary command affects how a network identified by using the network command is advertised. The way the router in the exhibit is configured, it will not announce the 172.16.31.0/24 subnet. The BGP router will announce the classful address 172.16.0.0/16 when the routing table, maintained by the IGP, contains an exact match to the network command.

The network command directly affects what network is advertised in BGP. If the command does not also include a network mask, and if auto-summary is enabled, then the classful address 172.16.0.0 is advertised any time that the router learns about a subnet of 172.16.0.0 via its IGP such as OSPF or EIGRP. The routing table does contain an entry for the 172.16.31.0/24 subnet that was learned through the IGP. However, auto-summary is disabled with the no auto-summary command. Therefore, only networks in the routing table that are an exact match to the network commands are advertised.

If the configuration goal is to announce the 172.16.0.0/16 network any time one of its subnets is learned, such as 172.16.31.0/24, then you should enable auto-summary. If the goal is to announce only the 172.16.31.0/24 subnet learned through the IGP, then you should alter the network command's IP address and include the subnet mask.

The BGP synchronization rule specifies that networks will not be advertised or used via iBGP unless it also has been learned through an IGP. If synchronization is disabled, iGBP will advertise a network without also learning it through an IGP. It does not affect the summarization of routes.

The 172.16.31.0 network was learned through an IGP session with router 10.1.2.3.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify eBGP (IPv4 and IPv6 address families)

References:
Cisco IOS Master Command List, Release 12.4 > a through b > BGP Commands: A through B > auto-summary (BGP)
Cisco > Cisco IOS IP Routing: BGP Command Reference > router bgp

Cisco > Cisco IOS IP Routing: BGP Command Reference > network (BGP and multiprotocol BGP)

QUESTION 63
What does the passive-interface command do when implement with RIP? (Choose two.)

A. Allows an interface to receive routing update traffic
B. Prevents an interface from sending routing update traffic
C. Prevents an interface from sending any normal data traffic
D. Allows an interface to receive normal data traffic
E. Disables a router interface
F. Places a router interface in standby mode

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The effect of the passive-interface command is dependent on the routing protocol running on the interface. For EIGRP, the router will not only stop sending routing updates, but also hellos, which means that it will not form a neighbor relationship with another EIGRP router on that interface. This is also the case with OSPF and IS-IS. With RIP, however, the router will continue to send hellos even as it stops sending routing updates, and it will still receive routing updates. An example of using the passive-interface command is below. The command is issued from the router configuration mode.

```
Router(config-router)# passive-interface ethernet 0/0
```

The passive-interface command will even overrule a configuration that includes a distribute list that allows the advertisement of a network through the interface. Examine the partial output of the show run command taken from a router running EIGRP below:

```
router6#show run
!
router eigrp 100
network 10.16.18.0 0.0.255.255
network 10.16.19.0 0.0.255.255
passive-interface serial 0/0
distribute-list 50 out serial 0/0
!
Access-list 50 permit 10.16.8.0 0.0.255.255
```

In this case, although the distribute list allows the advertisement of the 10.16.8.0 network, the passive-interface command applied to the Serial 0/0 interface will disallow all outgoing and incoming updates.

The passive-interface command does not affect the transmission or reception of normal data traffic, only routing updates.

The passive-interface command does not disable the router interface. The shutdown command is used to disable a router interface.

The passive-interface command does not place the router in standby mode.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify loop prevention mechanisms

References:
Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design Technotes > How Does

the Passive Interface Feature Work in EIGRP?
Cisco > Cisco IOS IP Routing: Protocol-Independent Configuration Guide, Release 12.4 > Configuring IP Routing Protocol-Independent Features > Filtering Routing Information
Cisco > Cisco IOS IP Routing: Protocol-Independent Command Reference > passive-interface

QUESTION 64
As the network administrator, you need to develop a verification plan for an OSPF network. The OSPF network has several area routers, area border routers (ABRs), and autonomous system boundary routers (ASBRs).

Which LSA types should you expect ABRs to generate while verifying the OSPF network? (Choose two.)

A. Type 4
B. Type 3
C. Type 2
D. Type 5

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
ABRs generate Type 3 and Type 4 LSAs in an OSPF network. ABRs are those routers that exist between two OSPF areas, as shown in the following figure:

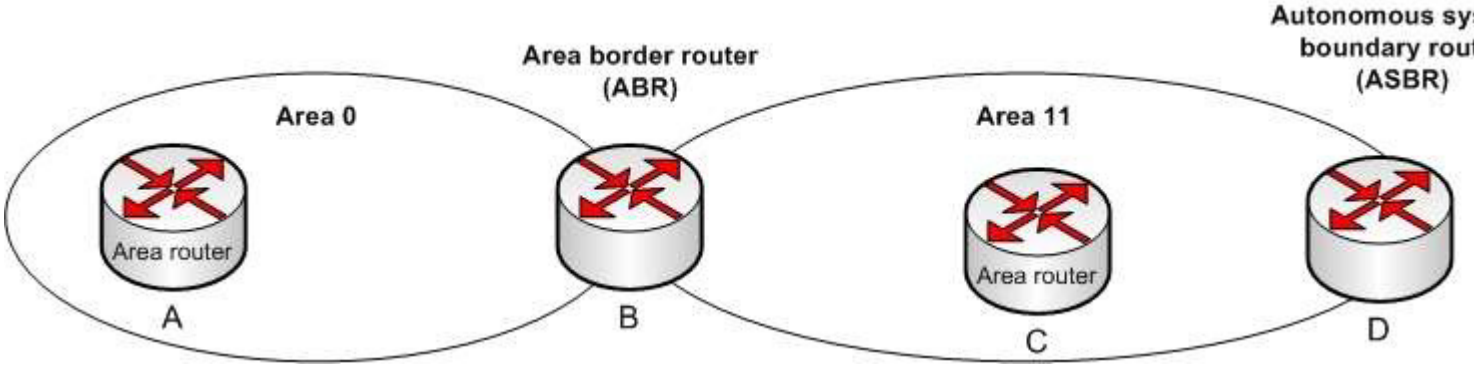

Type 3 and Type 4 LSAs are generated by ABRs to be flooded into other areas to and from the backbone area (area 0). Type 3 LSAs, or summary link advertisements, contain the list of networks known by one area. ABRs send Type 3 LSAs to the other OSPF areas in a given AS.

OSPF ABRs generates Type 4 LSAs to advertise the list of routes that point to an ASBR. These LSAs advertise the location of the ASBR.

Type 5 LSAs are not generated by an ABR. These LSAs are generated by ASBRs to describe routes redistributed into the area from other autonomous systems.

Type 2 LSAs are not generated by an ABR. A Type 2 LSA is generated only by the designated router (DR) of a segment to be sent to the other routers that belong to the same area as the DR. A DR is a router that has the highest OSPF priority on a segment. These advertisements are used by the DR to represent the routers that are connected to the network.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco Learning Home > Groups > CCNP R&S Study Group > Discussions > OSPF Level of Detail
Cisco > Support > Technology Support > IP > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > Link State Packets

QUESTION 65
You are troubleshooting an issue with the configuration of mGRE on the hub router in a hub-and-spoke configuration. Examine the output of the configuration of the tunnel interface on the hub router:

```
interface FastEthernet1/0
ip address 192.168.251.1 255.255.255.0

interface Tunnel0
bandwidth 1536
ip address 10.62.1.1 255.255.255.0
tunnel source FastEthernet1/0
```

Which of the following statements is true?

A. The tunnel destination must be specified on the tunnel interface
B. the tunnel mode gre multipoint command must be executed on the tunnel interface
C. the tunnel mode gre multipoint command must be executed on the physical interface
D. The tunnel destination must be specified on the physical interface

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The tunnel mode gre multipoint command must be executed on the tunnel interface. An mGRE configuration is one in which the tunnel is allowed to have multiple destinations. The distinguishing feature between an mGRE interface and a point-to-point GRE interface is the tunnel destination. While it is specified on a point-to-point GRE interface, it is not on an mGRE interface. Instead the command tunnel mode gre multipoint is executed on the tunnel interface. This allows the interface to use the Next Hop Routing protocol (NHRP) to discover the IP addresses of the other tunnel endpoints.

The tunnel destination is not specified on the tunnel interface using mGRE. Instead the command tunnel mode gre multipoint is executed on the tunnel interface.

The tunnel mode gre multipoint command must be executed on the tunnel interface, not the physical interface.

The tunnel destination is neither specified on the tunnel interface nor on the physical interface when using mGRE.

Objective:
VPN Technologies
Sub-Objective:
Configure and verify GRE

References:
Cisco > Dynamic Multipoint VPN (DMVPN) Design Guide (Version 1.1) > DMVPN Design and Implementation > mGRE Configuration
Cisco > Cisco IOS IP Mobility Command Reference > tunnel mode gre

QUESTION 66
You are planning the configuration of Easy Virtual Networking (EVN).

Which of the following statements is true of an interface that will be an EVN trunk?

A. It must support 802.1q encapsulation
B. The interface can also be configured for VRF-Lite
C. The interface will support OSPFv3
D. The interface can support RIP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The interface must be able to support 802.1q encapsulation. The EVN trunk carries the traffic of multiple virtual routing and forwarding (VRF) instances, with the traffic of each instance tagged with an ID called the virtual network tag. Since the VLAN ID field of an 802.1q encapsulated packet is used for this ID, the link must be one that supports 802.1q.

Easy Virtual networking is a technology that allows for the creation of separate networks with separate routing tables and routing instances using the same physical topology. The IP addressing for the networks can even overlap with no problem. The networks are kept separate using the network ID tags in a similar fashion to the way switches keep VLANs separate by using VLAN tags.

An EVN trunk interface cannot also be configured for VRF-Lite. VRF-Lite is an earlier technology that accomplishes the same goal, but lacks the simplicity of EVN.

Neither RIP nor OSPFv3 is supported in Easy Virtual Networking EVN at all.

Objective:
VPN Technologies
Sub-Objective:
Describe Easy Virtual Networking (EVN)

References:
Cisco > Easy Virtual Network Configuration Guide, Cisco IOS XE Release 3S > Overview of Easy Virtual Network

QUESTION 67
You have a DMVPN hub with the following configuration applied:

```
Hub645# show running-config
crypto isakmp policy 1
authentication pre-share
crypto isakmp key cisco47 address 0.0.0.0
!
crypto IPsec transform-set trans2 esp-des esp-md5-hmac
!
crypto ipsec profile vpnprof
set transform-set trans2
!
interface Tunnel0
bandwidth 1000
ip address 10.0.0.1 255.255.255.0
ip mtu 1416
ip nhrp authentication donttell
ip nhrp map multicast dynamic
ip nhrp network-id 99
ip nhrp holdtime 300
no ip split-horizon eigrp 1
no ip next-hop-self eigrp 1
delay 1000
tunnel source Ethernet0
tunnel mode gre multipoint
tunnel key 100000
tunnel protection ipsec profile vpnprof
interface Ethernet0
ip address 172.17.0.1 255.255.255.0
!
interface Ethernet1
ip address 192.168.0.1 255.255.255.0
!
router eigrp 1
network 10.0.0.0 0.0.0.255
network 192.168.0.0 0.0.0.255
!
```

What problem could occur if the bandwidth 1000 command were missing from the tunnel interface?

A. the tunnel interface will intermittently flap up and down
B. split horizon will prevent routing updates from traversing from spoke to spoke
C. congestion will develop in the tunnel interface
D. the IPsec association will fail

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In the absence of a bandwidth command on the tunnel interface, the default bandwidth on a tunnel interface is 9 Kbps. EIGRP will use 50% of that (4.5K), which is too low. This will cause problems with the maintenance of EIGRP neighbor relationships. From time to time this will cause the tunnel to flap up and then down as the relationships go up and down. When you execute the bandwidth command it has no real effect on the bandwidth of the link but it will allow EIGRP to use 50% of 10k or 5k for its purposes, leaving 4k for data. This will have little impact on the data while maintaining the neighbor relationships.

The bandwidth command will have no effect on split horizon. There will be no problems with split horizon, even though the output shows that it has been disabled on the tunnel interface with the no ip split-horizon eigrp 1

command.

The bandwidth command will not cause congestion on the link. It will only lower the bandwidth available to data from 4.5K to 4K.

The bandwidth command will not cause the IPsec association to fail. There is sufficient bandwidth for this process.

Objective:
VPN Technologies
Sub-Objective:
Describe DMVPN (single hub)

References:
Understanding Cisco Dynamic Multipoint VPN - DMVPN, mGRE, NHRP

QUESTION 68
The following configuration was applied to the router R66:

```
R66# show running-config
Building configuration...
Current configuration: 1072 bytes

<output omitted>

vrf definition red
vnet tag 3
<output omitted>

address-family ipv4
exit-address-family
<output omitted>

interface FastEthernet 1/0/0
vnet trunk
ip address 10.1.1.1 255.255.255.0
```

What is the interface ID and the IP address of the subinterface created to host the virtual network named red? (Choose two.)

A. FastEthernet1/0/0.3
B. FastEthernet0/0/0.red
C. FastEthernet0/0/3
D. 10.1.1.3
E. 10.1.1.1
F. 10.0.0.3

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The interface ID of the subinterface created to host the virtual network named red will be FastEthernet1/0/0.3, and the IP address will be 10.1.1.1.

When a virtual routing and forwarding (VRF) instance is defined, it will have a name and a tag number. The tag

number is used by the router to dynamically create a subinterface on the specified physical interface of the EVN trunk. The tag number is appended to the physical interface ID. Since the virtual network (vnet) trunk was defined as FastEthernet1/0/0, the subinterface for vrf red will be FastEthernet1/0/0.3. All subinterfaces on the trunk will use the same IP address as the physical interface defined as the trunk.

Easy virtual networking (EVN) is a technology that allows for multiple logical networks to use the same physical infrastructure. EVN trunks carry the traffic of multiple VRFs. While the subinterfaces dedicated to each VRF use the same IP address (that of the physical interface of the EVN trunk), no IP address conflicts ever occur because each VRF maintains its own routing and forwarding tables, and while on the trunk, each uses a VRF tag to separate the traffic from each VRF.

Objective:
VPN Technologies
Sub-Objective:
Describe Easy Virtual Networking (EVN)

References:
Cisco > Easy Virtual Network Configuration Guide, Cisco IOS XE Release 3S > Overview of Easy Virtual Network

QUESTION 69
Which command sets the OSPF priority value of a router interface to 10?

A. Router(config)# ospf priority 10
B. Router(config-if)# ospf priority 10
C. Router(config)# ip ospf priority 10
D. Router(config-if)# ip ospf priority 10

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The correct syntax for the ip ospf priority command is shown below:

Router(config-if)# ip ospf priority {number}

The number is a value from 0 to 255, and 1 is the default priority. A priority value of 0 means that the interface cannot be elected as the designated router (DR) or backup designated router (BDR). The higher the priority, the more preferred the router is when there is an election for DR and BDR for that network.

NOTE: The ip ospf priority command is entered in interface configuration mode, not router configuration mode.

All other options either use incorrect syntax or are executed at an incorrect prompt.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify network types, area types, and router types

References:
Cisco IOS Master Command List, Release 12.4 > i through k > ip ospf priority

QUESTION 70
When configuring a DMVPN solution, which of the following technologies makes it possible for the spoke routers to use dynamic IP addressing?

A. IPsec

B. mGRE
C. NHRP
D. Dynamic routing protocols

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Next Hop Resolution Protocol (NHRP) allows the spoke routers to register their IP addresses with the NHRP server, which is the hub router. It also allows the spoke routers to then learn the physical IP addresses of the other spoke routers from the hub router, allowing for GRE links to be built dynamically as needed between the spokes. This eliminates the need for the traffic to go through the hub router.

Dynamic Multipoint VPN (DMVPN) technology leverages the following associated technologies:
IPsec
mGRE
Dynamic routing protocols
NHRP
Cisco Express Forwarding

It makes it possible to build the hub router once, and add spokes later, making no additional changes to the hub. The spokes are able to register with the hub and dynamically build their own connections to other spokes using the IP addresses learned from the hub using NHRP. DMVPN also allows IPsec point-to-point GRE tunnels to be built to new spokes with no IPsec peering configuration. The multipoint GRE technology (mGRE) allows a single physical interface on the hub to be used for all spoke connections.

Finally, the routing protocols used by DMVPN allow the routers to share routing information, while Cisco Express Forwarding (CEF) is a switching technology that improves performance while reducing the load on the CPUs of the routers.

Objective:
VPN Technologies
Sub-Objective:
Describe DMVPN (single hub)

References:
Cisco > Dynamic Multipoint VPN (DMVPN) Design Guide (Version 1.1) > DMVPN Design Overview

QUESTION 71
You executed the following commands to assign an IPv6 link-local address to the Fa0/0 interface of the R1 router:

```
R1(config)# interface Fa0/0
R1(config-if)# ipv6 ospf 1 area 1
```

When you executed the show running-config command on the R1 router, you observed that OSPF for IPv6 is not running on the router.

Which of the following commands should be added to the interface configuration?

A. ipv6 router ospf
B. ipv6 enable
C. ipv6 ospf neighbor
D. ipv6 ospf cost

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
The ipv6 enable command should be used on R1 to enable IPv6. This command automatically provides an IPv6 link-local unicast address for the interface on which IPv6 is being configured. If an explicit IPv6 address were configured on the interface, the command would not be required.

The ipv6 router ospf command should not be used in the configuration because this command allows you to enter the router configuration mode for OSPF for IPv6.

The ipv6 ospf neighbor command is used to configure neighboring routers for OSPF.

The ipv6 ospf cost command should not be added to the configuration because this command allows you to specify the OSPF cost to send packets from a given interface.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF for IPv6

References:
Cisco > Cisco IOS IPv6 Command Reference > ipv6 enable

QUESTION 72
Consider the following diagram. All PVCs are active.

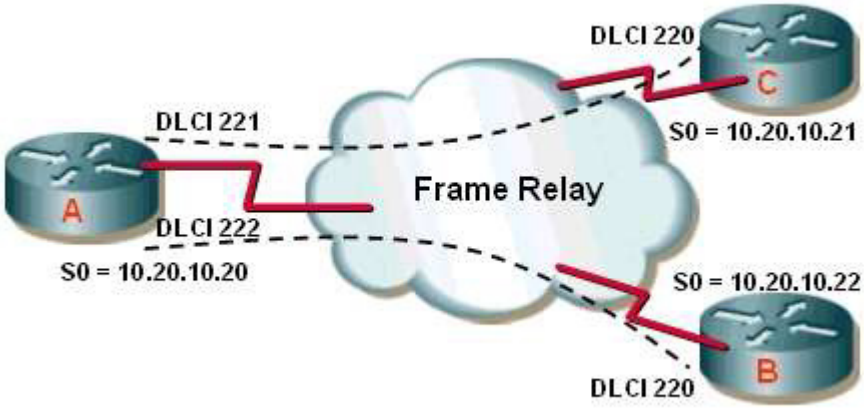

If the partial output of the show ip ospf neighbor command executed on Router A is as follows, which of the following statements is TRUE?

```
RouterA# show ip ospf neighbor

Neighbor ID Pri State Dead Time Address Interface
1.1.1.1 1 FULL/DROTHER 00:00:13 10.20.10.21 Serial0
2.2.2.2 1 FULL/DR 00:00:51 10.20.10.22 Serial0
```

A. Router C and Router B will fail to have all OSPF routes in their tables.
B. All routing tables will be populated correctly.
C. Router A will be the DR.

D. Router C will be the DR.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The output of the command shows that Router C and Router B will fail to have all OSPF routes in their tables. In a hub and spoke configuration, as depicted in the diagram, the hub router (Router A) should be the designated router (DR) or the source of updates to the other routers. However, Router B is the DR, as evidenced by the output of the show ip ospf neighbor command executed on Router A.

This situation could be rectified by setting Routers B and C with a priority of 0, which would disqualify them from being the DR. After that, all routes could be distributed from the hub, which would have visibility of all routes.

All routing tables will be populated correctly until the hub router is made the DR.

Neither Router A nor C will be the DR, since it is indicated that Router B is the DR in the output of the command.

Objective:
Layer 3 Technologies
Sub-Objective:
Configure and verify OSPF operations

References:
Cisco > Home > Support > Support Technology > Support > IP Routing > Configure > Configuration Examples and Technotes > Initial Configurations for OSPF over Frame Relay Subinterfaces
Cisco > Cisco IOS Wide-Area Networking Command Reference > frame-relay lapf n201 through fr-atm connect dlci > frame-relay map
Cisco > Cisco IOS IP Routing: OSPF Command Reference > ip ospf network

QUESTION 73
The following access lists are applied to an interface connecting two OSPF routers:

```
R2(config)#ipv6 access-list PERMIT_HTTP
R2(config-ipv6-acl)# permit tcp any any eq 80
R2(config-ipv6-acl)# deny ip any any log
```

What is the result?

A. the DR on the link will begin updating
B. the OSPF adjacency will go down
C. the last deny statement will fail to log traffic
D. the list will only permit IPv6 neighbor advertisements

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If this list is applied to the interface connecting two OSPF routers, the OSPF adjacency would go down. The deny ip any any log statement will deny the IPv6 link local addresses, which are used for the neighbor discovery process and by OSPF routers to establish neighbor adjacencies when directly connected.

By default, IPv6 access lists have a deny all at the end that does NOT include those addresses. However, when

you set an explicit deny all as shown in the scenario, you will block all traffic that is not specified by an earlier statement in the list.

The DR on the link, if present, will not begin updating because the adjacency will fail. It will then have no neighbor to update.

The last deny statement in the scenario will log any traffic it blocks, as indicated by the inclusion of the log keyword.

The list will NOT permit neighbor advertisements. These are always done in terms of link local addresses, which the explicit deny ip any any log statement at the end will block.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Security Configuration Guide: Access Control Lists, Cisco IOS XE Release 3S > IPv6 Access Control Lists
Cisco > Cisco IOS IPv6 Command Reference > ipv6 access-list
Cisco > Cisco IOS Security Command Reference: Commands M to R > permit (IP)
Cisco > Cisco IOS Security Command Reference: Commands D to L > deny (IP)

QUESTION 74
Which of the following IPv6 access list statements would permit SSH traffic from 2001:DB8:0:4::32 when applied to the VTY lines?

A. permit ipv6 2001:DB3:0:5::/48 any eq ssh
B. permit ipv6 2001:DB8:0:4::/64 any eq ssh
C. permit ipv6 host 2001:DB8:0:4::32 any eq 23
D. permit ipv6 2001:DE8:0:4:::/48 any eq 22

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The only statement that would allow SSH traffic from 2001:DB8:0:4::32 is permit ipv6 2001:DB8:0:4::/64 any eq ssh. It would match because it specifies the 2001:DB8:0:4:: subnet as a result of the /64 prefix. With that prefix, traffic must match in the first four hextets. Since the address 2001:DB8:0:5::32 matches in the first four hextets, it is allowed.

The statement permit ipv6 2001:DB3:0:5::/48 any eq ssh will not permit traffic from 2001:DB8:0:4::32. With a /48 subnet mask, the address must match in the first three hextets, and it does not do

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Catalyst 3750 Software Configuration Guide, Release 12.2(55)SE > Configuring IPv6 ACLs
Cisco > Cisco IOS IPv6 Command Reference > permit (IPv6)

QUESTION 75
An associate creates the following access list that she plans to apply to an interface on a router:

access-list 100 permit ip any any log

What type of traffic could cause this ACL to place a heavy load on the CPU of the router, and what command could be used to reduce the impact of the ACL? (Choose two.)

A. traffic that is CEF switched
B. traffic that is process switched
C. traffic that is fast switched
D. ip access-list log-update threshold
E. ip access-list logging interval
F. logging rate limit

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
There are two contributors to the CPU load increase from ACL logging: process switching of packets that match log-enabled access control entries (ACEs), and the generation and transmission of the log messages. To reduce the impact of process switched traffic, the ip access-list logging interval command can be used. The interval is specified in milliseconds and represents how often a single packet is process switched. While the messages in the generated log entries may not be as comprehensive after this command is executed, the counter values that are generated by the show access-list and show ip-access list commands will still be accurate.

Packets that are not process switched (CEF switched and fast switched) will examined or accounted for in the logging, so they are not the source of the problem.

The ip access-list log-update threshold command is used to configure how often syslog messages are generated and sent after the initial packet match. While this would be a beneficial command to run, as it addresses the second source of CPU congestion that is the sending of the syslog messages, that was not listed as a traffic type option. Therefore, this would not be a solution to the issue presented by packet switched traffic.

The logging rate limit command also will reduce the impact of log generation and transmission on the CPU, but again, it does not address the issue presented by process switched traffic.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Understanding Access Control List Logging
Cisco > Cisco IOS Security Command Reference: Commands D to L > ip-group

QUESTION 76
Which of the following commands enables Unicast Reverse Path forwarding in loose mode?

A. ip verify unicast source reachable-via rx
B. ip verify unicast source reachable-via any
C. ip verify unicast source reachable-via rx allow default
D. ip verify unicast source reachable-via allow default

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The command ip verify unicast source reachable-via any enables Unicast Reverse Path Forwarding (RPF) in loose mode. In loose mode, traffic is allowed if the source address is reachable via any interface on the router as indicted in the routing table. Unicast Reverse Path forwarding uses the source IP address when it validates the packet. Packets are validated when the source address is contained in the routing table and is reachable either via the ingress interface (strict mode) or via any interface (loose mode).

The command ip verify unicast source reachable-via rx enables Unicast RPF in strict mode, not loose mode. The rx keyword indicates the source must be reachable on the interface where the packet arrived.

The command ip verify unicast source reachable-via rx allow default enables Unicast RPF in strict mode. The inclusion of the allow default keyword indicates the source can be reachable via a default route to be accepted.

The command ip verify unicast source reachable-via allow default is syntactically incorrect. The allow default keyword cannot be present by itself. It must follow either the rx or any keywords.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Understanding Unicast Reverse Path Forwarding
Cisco > Cisco IOS Security Command Reference: Commands D to L > ip verify unicast source reachable-via

QUESTION 77
Examine the following access list:

```
access-list 110 deny ip 10.0.0.0 0.255.255.255 any
access-list 110 deny ip 172.16.0.0 0.15.255.255 any
access-list 110 deny ip 192.168.0.0 0.0.255.255 any
access-list 110 deny ip 208.0.0.0 0.255.255.255 any
```

Which statement is NOT designed to prevent IP spoofing attacks from packets that appear to be sourced from inside the network, but are actually sourced from outside the network?

A. access-list 110 deny ip 10.0.0.0 0.255.255.255 any
B. access-list 110 deny ip 172.16.0.0 0.15.255.255 any
C. access-list 110 deny ip 192.168.0.0 0.0.255.255 any
D. access-list 110 deny ip 208.0.0.0 0.255.255.255 any

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Infrastructure access control lists are designed to prevent spoofing attacks from packets that appear to be sourced from inside the network when they are in fact sourced from outside the network. There are two groups of address that should be blocked at the edge of the network:
- The private address space, which are called RFC 1918 addresses
- Certain "special use addresses" as defined in RFC 3330

The address 208.0.0.0 0.255.255.255 falls into neither of those categories.

The RFC 1918 addresses that should be blocked are:

10.0.0.0/24
172.16.0.0/16
192.168.0.0/16

The RFC 3330 addresses that should be blocked are:

0.0.0.0
127.0.0.0/8
192.0.2.0/24
224.0.0.0/4

For more information about these special use addresses, see RFC 3330.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Home > Support > Technology Support > IP > IP addressing services > Technology information > Technology white paper >Protecting Your Core: Infrastructure Protection Access Control Lists

QUESTION 78
Examine the following output of the show ip route command and the partial output of the show run command from the router R63:

```
R63#show ip route

10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks

C          10.2.1.0/24 is directly connected, Serial0/0
L          10.2.1.1/32 is directly connected, Serial0/0
     10.0.0.0/24 is subnetted, 1 subnets
S          10.10.10.0 is directly connected, Tunnel0
     10.11.0.0/24 is subnetted, 1 subnets
S          10.11.11.0 is directly connected, Ethernet0/0
S     0.0.0.0/0 [1/0] via 172.21.114.65, Ethernet0/1


R63#show run
<output omitted>

interface Serial0/0
ip address 10.2.1.1 255.255.255.0
ip verify unicast source reachable via rx
```

What will the router do with a packet with a source address of 192.168.5.5/24 and a destination address of 10.11.11.20/ 24 that arrives on the Serial0/0 interface?

A. forward it out the Ethernet0/0 interface
B. forward it out the Tunnel0 interface
C. drop the packet
D. forward it out the Ethernet0/1 interface

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
It will drop the packet. The partial output of the show run command shows that the ip verify unicast source reachable via rx command has been executed on the Serial 0/0 interface. This enables the Unicast Reverse Path Forwarding (Unicast RPF) feature. This feature prevents IP spoofing by verifying from the routing table that there is a valid return path to the source IP address. If there is not valid return path, you can assume the IP address has been spoofed. When the command ends in the keyword rx, it means that there must be a return path through the interface where the command was executed. This is called strict mode.

The packet arrived on the Serial0/0 interface. The routing table shows that there is no routing entry for the 192.168.5.0/24 network that leads back through the entry interface of Serial0/0. In fact, in this instance there is no routing table entry for that network leading to any interface. When this occurs, the router will drop the packet.

The router will not send the packet to either the Ethernet0/0 or the Tunnel0 interfaces because the destination network, 10.11.11.0/24, is not a reachable destination on those interfaces. Even if it were reachable, the Unicast Reverse Path Forwarding (Unicast RPF) feature will drop the packet because it has been spoofed.

It will not send the packet to the Ethernet0/1 interface. The Unicast Reverse Path Forwarding (Unicast RPF) feature will drop the packet because it has been spoofed. If the packet were not spoofed, it would be sent to the Ethernet0/1 interface because that is the interface used by the default route. Because there is no route in the table to the 10.11.11.0/24 network, it would be sent to the default route.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Cisco IOS Security Configuration Guide, Release 12.2 > Configuring Unicast Reverse Path Forwarding
Cisco > Configuring Unicast Reverse Path Forwarding

QUESTION 79
After an associate configured a DMVPN hub, you execute the following command on the hub router:

```
Router#show ip nhrp detail
10.1.1.2/8 via 10.2.1.2, Tunnel1 created 00:00:12, expire 01:59:47
Type: dynamic, Flags: authoritative unique nat registered used
NBMA address: 10.12.1.2
```

Which of the following statements is true of this output?

A. The NMBA address was statically configured
B. The NHRP information did not come from the NHS
C. The mapping was created through an NHRP registration request
D. The device at 10.1.1.2 is behind a NAT router

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The mapping was created through an NHRP registration request, as indicated by the flag setting registered. Next Hop Resolution Protocol (NHRP) can be used in place of static IP address to NBMA address mappings to allow the spoke routers in an mGRE hub-and-spoke configuration to discover one another's physical IP addresses.

When the output of the show nhrp detail command shows the registered flag listed, it means that the mapping was created dynamically and was learned through a registration request to the next hop server (NHS).

The mapping was not created statically. Had it been created statically, the Type field would not be listed as dynamic. It would say static.

The NHRP information DID come from the next hop server (NHS). That is indicated by the presence of the authoritative flag. The NHS is the next hop to the destination as indicated by the routing table.

The device at 10.1.1.2 is not necessarily behind a NAT router. The presence of the nat flag in the output indicates that the device at 10.1.1.2 supports the NHRP NAT extension type for supporting dynamic spoke-to-spoke tunnels to or from spokes behind a NAT router. This flag does not mean that the spoke (NHS client) is behind a NAT router.

Objective:
VPN Technologies
Sub-Objective:
Describe DMVPN (single hub)

References:
Home > Support > Product support > Cisco IOS and NX-OS software > Cisco IOS software releases 12.4 mainline > Configure > Feature Guides > NHRP

QUESTION 80
The following commands were executed on the perimeter router. The Fa1/0 interface in the router is the external interface.

```
router(config)# access-list 101 deny ip 10.0.0.0 0.255.255.255 any log
router(config)# access-list 101 deny ip 192.168.0.0 0.0.255.255 any log
router(config)# access-list 101 deny ip 172.16.0.0 0.15.255.255 any log
router(config)# interface fastEthernet 1/0
router(config-if)# ip access-group 101 in
```

What will be the effect of these commands?

A. all traffic will be blocked incoming
B. traffic sourced from private IP addresses will be blocked incoming
C. traffic destined for private IP addresses will be allowed incoming
D. no traffic will be blocked incoming

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
All traffic will be blocked incoming. While it appears on the surface that this list was designed to block incoming traffic sourced from private IP addresses, it is lacking a single permit statement. Due to the implied deny all at the end of the list, no traffic will be allowed incoming.

Blocking incoming traffic from private IP addresses is a way to prevent IP spoofing, since there should be no reason for traffic from private IP addresses to be incoming from the Internet. However, you need to include a permit statement at the end to allow all other traffic types.

Traffic destined for private IP addresses is not all that will be blocked by this command set. In fact, no traffic would be allowed. If there were a permit ip any any at the end of the list, then incoming traffic destined for private IP addresses would be allowed. This is probably not a great idea either, but if it a permit IP any were added at the end of the command set in the scenario, it would allow incoming traffic destined for private IP

addresses.

Objective:
Infrastructure Security
Sub-Objective:
Configure and verify router security features

References:
Cisco > Cisco IOS Security Command Reference: Commands A to C > access-list
Cisco > Cisco IOS Security Command Reference: Commands D to L > ip-group
Prevent IP spoofing with the Cisco IOS