



300-410^{Q&As}

Implementing Cisco Enterprise Advanced Routing and Services (ENARSI) (Include 2022 Newest Simulation Labs)

Pass Cisco 300-410 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4lead.com/300-410.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Cisco
Official Exam Center

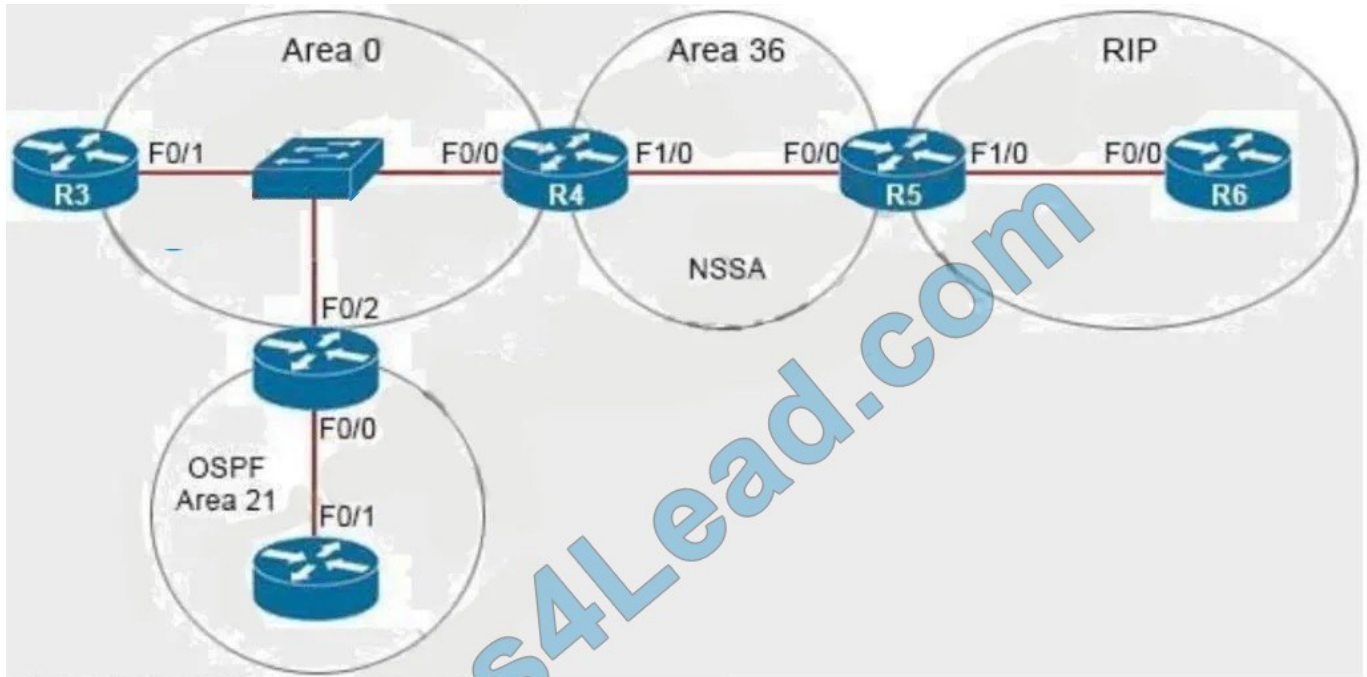
-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Refer to the exhibit. The network engineer configured the summarization of the RIP routes into the OSPF domain on R5 but still sees four different 172.16.0.0/24 networks on R4. Which action resolves the issue?



```
R5# show ip ospf 1 | begin Area 36
Area 36
Number of interfaces in this area is 2
It is a NSSA area
Area has no authentication
SPF algorithm last executed 00:32:46.376 ago
SFF algorithm executed 13 times
Area ranges are
172.16.0.0/16 Passive Advertise
```

- A. R5(config)#router ospf 1 R5(config-router)#no area R5(config-router)#summary-address 172.16.0.0 255.255.252.0
- B. R4(config)#router ospf 99 R4(config-router)#network 172.16.0.0 0.255.255.255 area 56 R4(config-router)#area 56 range 172.16.0.0 255.255.255.0
- C. R4(config)#router ospf 1 R4(config-router)#no area R4(config-router)#summary-address 172.16.0.0 255.255.252.0
- D. R5(config)#router ospf 99 R5(config-router)#network 172.16.0.0 0.255.255.255 area 56 R5(config-router)#area 56 range 172.16.0.0 255.255.255.0

Correct Answer: A

Area 36 is a NSSA so R5 is an ASBR so we can summarize external routes using the "summaryaddress" command. The command "area area-id range" can only be used on ABR so it is not correct. The summarization must be done on the

ASBR which is R5, not R4 so the correct answer must be started with "R5(config)#router ospf 1".



Note: The "no area" command is used to remove any existing "area ..." command (maybe "area 56 range ..." command).

QUESTION 2

You have enabled RIPng on one of the interfaces of a router with the basic configuration. You have assigned an address to that interface using the ipv6 address command.

Which of the following statements should appear in the output of the show running-config command executed on the router? (Choose all that apply.)

- A. ipv6 unicast-routing
- B. ipv6 enable
- C. ipv6 rip enable
- D. ipv6 router rip
- E. ipv6 unnumbered
- F. ipv6 prefix-list

Correct Answer: AC

The ipv6 unicast-routing and ipv6 rip enable statements should appear in the output of the show running-config command.

The ipv6 unicast-routing command is one of the basic IPv6 commands that needs to be executed on any router for IPv6 processing. This command is executed in the global configuration mode to allow IPv6 packet forwarding on the router. When it has been executed the ipv6 unicast-routing statement will appear in the output of the show run command.

The ipv6 rip enable command allows you to enable RIPng on a router interface. You should execute this command to create a RIPng routing process. When it has been executed the ipv6 rip enable statement will appear in the output of the show run command.

The ipv6 router rip command allows you to work with RIPng routing process by entering the router configuration mode for RIPng. It will only appear if modifications have been made to the RIPng routing process, which is not mentioned in this case.

The commands to configure a router with RIPng is as follows:

```
rtrA(config)# ipv6 unicast-routing rtrA(config)# interface Fa0/1 rtrA(config-if)# ipv6 rip rip process enable rtrA(config-if)#  
ipv6 address 2001:1:1:1::1/64
```

Important note: in the command set above, the command that enables the RIP process on interface Fa0/1 is executed before the command assigning the IPv6 address. The order of execution of those two commands does not matter. However, if a configuration file is copied and pasted into a router, then the order in which the statements appear does matter.

For example, if the partial configuration below were pasted into a router, the IPv6 RIP process 56 would not be enabled on Fa0/0:

```
ip unicast routing interface Fa0/0 ipv6 rip 56 enable ipv6 address 2001:1:1:1::1/64
```



The system would reject the ipv6 rip 56 enable command because an IPv6 address is not yet present. If the commands were reversed in the file, the system would accept the ipv6 rip 56 enable command.

In the scenario, the ipv6 enable command does not appear in the show running-config output. This command enables IPv6 routing on a router interface that has not been assigned an IPv6 address. In this case, an IPv6 address is explicitly assigned to the router interface by using the ipv6 address command. Therefore, the ipv6 enable command is not required.

The ipv6 unnumbered interface type command does not appear in the show running-config output. This command will allow you to enable IPv6 without assigning an IPv6 address to a router interface.

Objective:

Layer 3 Technologies

Sub-Objective:

Describe RIPng

References:

Cisco IPv6 Implementation Guide, Release 15.2MandT > Implementing RIP for IPv6 > How to Implement RIP for IPv6 > Enabling the IPv6 RIP Process Cisco IPv6 Implementation Guide, Release 15.2MandT > Implementing RIP for IPv6 >

Configuration Examples for IPv6 RIP > Examples: IPv6 RIP Configuration Cisco IOS IPv6 Command Reference > ipv6 ospf dead-interval through ipv6 split-horizon eigrp > ipv6 rip enable

Cisco IOS IPv6 Command Reference > ipv6 ospf dead-interval through ipv6 split-horizon eigrp > ipv6 router rip Cisco > Cisco IOS IPv6 Command Reference > ipv6 prefix-list Cisco IOS IPv6 Command Reference > ipv6 summary-address

through mpls ldp router-id > ipv6 unnumbered

QUESTION 3

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

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 4

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

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.

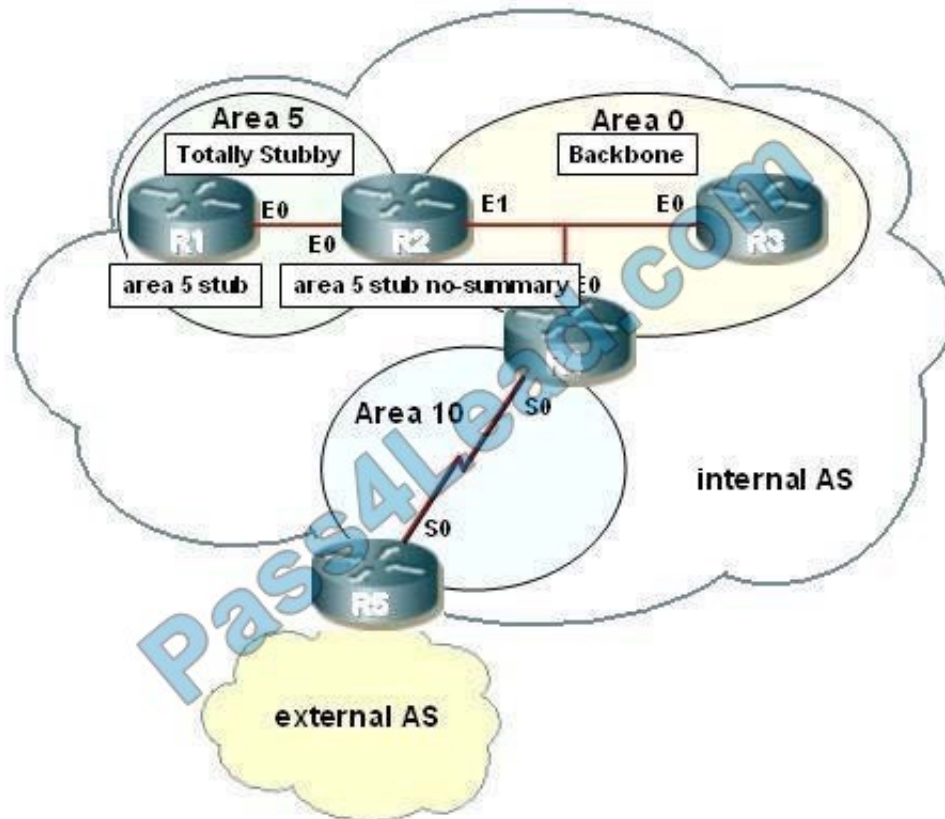


The correct syntax for the area stub command is shown below:

```
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 5

When the auth keyword is used in the snmp-server host command, which of the following must be configured with an authentication mechanism?

- A. the interface
- B. the host
- C. the user
- D. the group

Correct Answer: C

The auth keyword specifies that the user should be authenticated using either the HMAC-MD5 or HMAC-SHA algorithms. These algorithms are specified during the creation of the SNMP user.

For example, the following command creates a user named V3User that will be a member of the SNMP group V3Group and will use HMAC-MD5 with a password of MyPassword:

```
snmp-server user V3User V3Group v3 auth md5 MyPassword
```

The authentication mechanism is not configured on the interface. All SNMP commands are executed at the global configuration prompt.

The authentication mechanism is not configured at the host level. The version and security model (authentication, authentication and encryption, or neither) are set at the host level.

The authentication mechanism is not configured at the SNMP group level. The group level is where access permissions like read and write are set. This is why a user account must be a member of a group to derive an access level, even if it

is a group of one.

Objective:

Infrastructure Services

Sub-Objective:

Configure and verify SNMP

References:

Configuring SNMP Support > Understanding SNMP > SNMP Versions Cisco IOS Network Management Command Reference > snmp-server engineID local through snmp trap link- status > snmp-server host

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

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

Refer to the exhibit.

```
BRANCH-RTR#
router eigrp 100
 network 10.4.31.0 0.0.0.7
 network 10.100.100.1 0.0.0.0
 distribute-list route-map FILTER-IN in FastEthernet0/0
 eigrp router-id 10.100.100.1
!
 ip prefix-list 102 seq 10 permit 10.1.1.100/32
!
 route-map FILTER-IN deny 10
  match ip address prefix-list 102
!
```

A junior engineer updated a branch router configuration. Immediately after the change, the engineer receives calls from the help desk that branch personnel cannot reach any network destinations. Which configuration restores service and continues to block 10.1.1.100/32?

A. route-map FILTER-IN deny 5



B. ip prefix-list 102 seq 15 permit 0.0.0.0/32 le 32

C. ip prefix-list 102 seq 5 permit 0.0.0.0/32 le 32

D. route-map FILTER-IN permit 20

Correct Answer: D

By using “deny” keyword in a route-map, we can filter out the prefix specified in the prefix-list.

But there is an implicit “deny all” statement in the prefix-list so we must permit other prefixes with “permit” keyword in the route-map.

QUESTION 8

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

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:

```
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 9

Which EIGRP packet type is sent as a multicast when a new route is discovered, and sent as a unicast to synchronize topology tables when neighbors initialize?

- A. ACK
- B. Hello
- C. Update
- D. Replies
- E. Queries

Correct Answer: C

EIGRP update packets are sent as a multicast when a new route is discovered, and sent as a unicast to synchronize topology tables when a neighboring router initializes.

Whenever EIGRP only needs to communicate with a single neighbor, it sends a unicast to that neighbor instead of the standard multicast. In this case, it unicasts a packet to update a new EIGRP router on the network with the information that

all other routers on that network already know.

Hellos for neighbor discovery and maintenance are always multicasts.

ACKs are hellos without data, and are always unicast.

Queries are always multicast.

Replies to queries are always unicast.

Objective:

Layer 3 Technologies



Sub-Objective:

Describe EIGRP packet types

References:

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

QUESTION 10

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

When EIGRP loses its 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 11

Which attribute eliminates LFAs that belong to protected paths in situations where links in a network are connected through a common fiber?

- A. Shared Risk Link Group (SRLG)-disjoint



- B. linecard-disjoint
- C. lowest-repair-path-metric
- D. interface-disjoint

Correct Answer: A

LFA Tie-Breaking Rules

When there are multiple candidate LFAs for a given primary path, EIGRP uses a tie-breaking rule to select one LFA per primary path per prefix. A tie-breaking rule considers LFAs that satisfy certain conditions or have certain attributes. EIGRP uses the following four attributes to implement tie-breaking rules:

–

Interface-disjoint—Eliminates LFAs that share the outgoing interface with the protected path.

–

Linecard-disjoint—Eliminates LFAs that share the line card with the protected path.

–

Lowest-repair-path-metric—Eliminates LFAs whose metric to the protected prefix is high. Multiple LFAs with the same lowest path metric may remain in the routing table after this tie-breaker is applied.

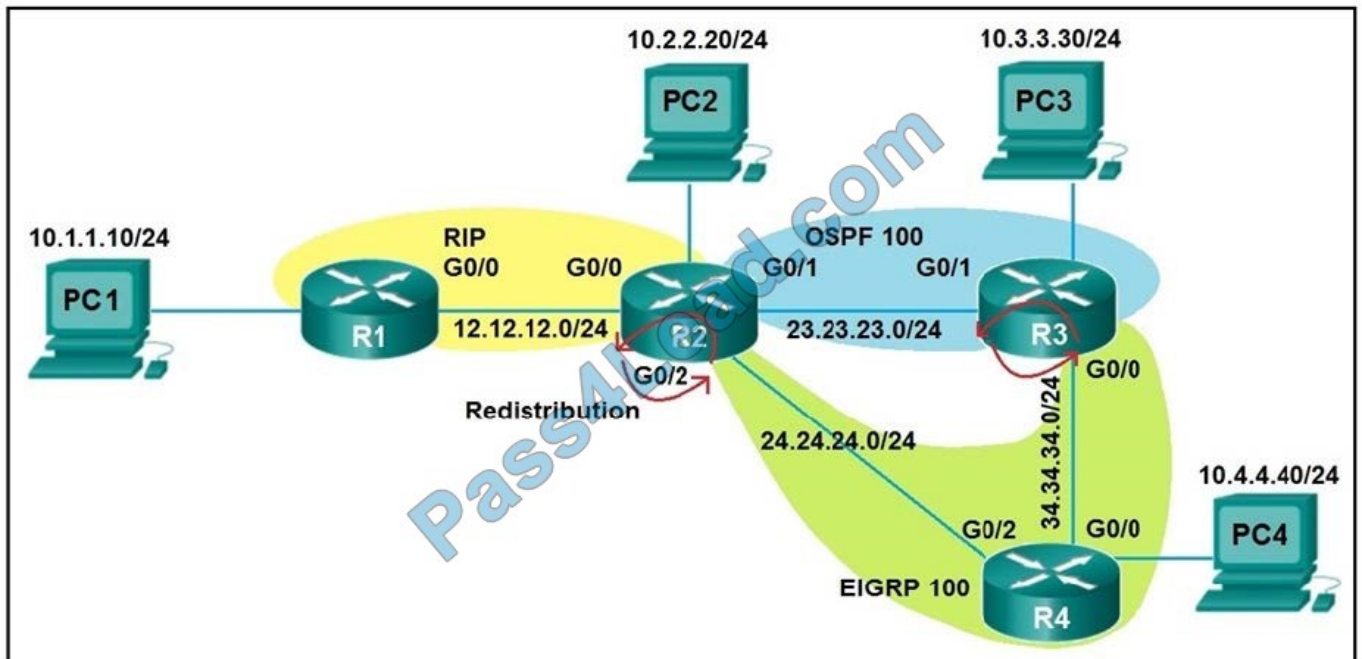
–

Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

Source : https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-3s/asr1000/ire-xe-3s-asr1000/ire-ipfrr.html

QUESTION 12

Refer to the exhibit. After redistribution is enabled between the routing protocols; PC2, PC3, and PC4 cannot reach PC1. Which action can the engineer take to solve the issue so that all the PCs are reachable?



- A. Set the administrative distance 100 under the RIP process on R2.
- B. Filter the prefix 10.1.1.0/24 when redistributed from OSPF to EIGRP.
- C. Filter the prefix 10.1.1.0/24 when redistributed from RIP to EIGRP.
- D. Redistribute the directly connected interfaces on R2.

Correct Answer: A

QUESTION 13

When determining if a system is capable of support, what is the minimum time spacing required for a BFD control packet to receive once a control packet is arrived?

- A. Desired Min TX Interval
- B. Detect Mult
- C. Required Min RX Interval
- D. Required Min Echo RX Interval

Correct Answer: C

QUESTION 14



Refer to the exhibit. The OSPF routing protocol is redistributed into the BGP routing protocol, but not all the OSPF routes are distributed into BGP. Which action resolves the issue?

Router# show ip route

```
2.0.0.0/24 is subnetted, 1 subnets
C    2.2.2.0 is directly connected, Ethernet0/0
C    3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
131.108.0.0/24 is subnetted, 2 subnets
O    131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
```

Router# show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.2.2.0/24	0.0.0.0	0	32768	?	
*> 131.108.1.0/24	2.2.2.2	84	32768	?	
*> 131.108.2.0/24	2.2.2.2	74	32768	?	

- A. Include the word external in the redistribute command
- B. Use a route-map command to redistribute OSPF external routes defined in an access list
- C. Include the word internal external in the redistribute command
- D. Use a route-map command to redistribute OSPF external routes defined in a prefix list.

Correct Answer: C

If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default. You can use the internal keyword along with the redistribute command under router

to redistribute OSPF intra- and inter-area routes.

Use the external keyword along with the redistribute command under router bgp to redistribute OSPF external routes into BGP.



-> In order to redistribute all OSPF routes into BGP, we must use both internal and external keywords. The full command would be (suppose we are using OSPF 1):

```
redistribute ospf 1 match internal external
```

Note: The configuration shows match internal external 1 external 2. This is normal because OSPF automatically appends "external 1 external 2" in the configuration. In other words, keyword external = external 1 external 2. External 1 = O E1

and External 2 = O E2.

QUESTION 15

Which feature drops packets if the source address is not found in the snooping table?

- A. IPv6 Source Guard
- B. IPv6 Destination Guard
- C. IPv6 Prefix Guard
- D. Binding Table Recovery

Correct Answer: A

[300-410 PDF Dumps](#)

[300-410 VCE Dumps](#)

[300-410 Practice Test](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4lead.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4lead, All Rights Reserved.