

Cisco.Premium.642-883.by.VCEplus.192q

Number: 642-883 VCEplus
Passing Score: 800
Time Limit: 120 min
File Version: 18.7



Exam Code: 642-883

Exam Name: Deploying Cisco Service Provider Network Routing (SPROUTE)

Certification Provider: Cisco

Corresponding Certification: CCNP Service Provider

Website: www.vceplus.com

Free Exam: <https://vceplus.com/exam-642-883-ccnp-test/>

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in 642-883 exam products and you get latest questions. We strive to deliver the best 642-883 exam product for top grades in your first attempt.

Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Exam A

QUESTION 1

Routes that are received from an IBGP peer will be propagated to which other routers by default?

- A. to the EBGP peers only
- B. to the IBGP peers only
- C. to both EBGP and IBGP peers
- D. to no other peers

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

What are two characteristics of the multihomed customers to service providers connection option? (Choose two.)

- A. Multihomed customers must use a private AS number.
- B. The traffic load can be shared for different destination networks between service providers.
- C. Multihomed customers must receive a full routing table from the service providers.
- D. The routing methodology must be capable of reacting to dynamic changes. BGP is used to achieve this flexibility.
- E. Multihomed customers must use a provider-assigned address space.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3

An engineer is using configuration blocks to build policy-based routing. All other traffic must not be affected and the policy must be applied to GigabitEthernet1/0. Which block correctly implements traffic that is destined for 10.10.10.0/24 using a next hop of 192.168.55.1?

- A. access-list 110 permit ip any 10.10.10.0 0.0.0.255 route-map PBR permit 10 match ip address 110 set next-hop 192.168.55.1 interface GigabitEthernet1/0 ip

policy route-map PBR

- B. access-list 110 permit ip any 10.10.10.0 0.0.0.255 route-map PBR permit 10 match ip address 110 set next-hop 192.168.55.1 route-map PBR permit 20 interface GigabitEthernet1/0 ip policy route-map PBR
- C. access-list 110 permit ip any 10.10.10.0 255.255.255.0 route-map PBR permit 10 match ip address 110 set next-hop 192.168.55.1 interface GigabitEthernet1/0 ip policy route-map PBR
- D. access-list 110 permit ip any 10.10.10.0 255.255.255.0 route-map PBR permit 10 match ip address 110 set next-hop 192.168.55.1 route-map PBR permit 20 interface GigabitEthernet1/0 ip policy route-map PBR

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 4

Which option is a mechanism that can be implemented between two eBGP peers to communicate the routes each peer needs from the other?

- A. ORF
- B. ACL
- C. prefix list
- D. route map



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 5

Which high-availability mechanism is a detection protocol that is enabled at the interface and at the routing protocol levels?

- A. NSF
- B. SSO
- C. NSR
- D. BFD
- E. SDR

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/fs_bfd.html

QUESTION 6

Which configuration implements the most secured OSPF authentication mode on Cisco IOS XE?

- A. interface Gi0/0/0/0 ip ospf authentication message-digest ip ospf message-digest-key 1 md5 cisco
- B. router ospf 1 area 1
authentication message-digest interface GigabitEthernet0/0/0/0 message-digest-key 1 md5 cisco
- C. router ospf 1 area 1 interface GigabitEthernet0/0/0/0 authentication message-digest message-digest-key 1 md5 cisco
- D. router ospf 1 area 1 authentication interface GigabitEthernet0/0/0/0 ip ospf authentication-key cisco
- E. interface Gi0/0/0/0 ip ospf authentication ip ospf authentication-key cisco

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

QUESTION 7

What is the default OSPF seed metric and type?

- A. 10 and E1
- B. 10 and E2
- C. 20 and E1
- D. 20 and E2
- E. 0 and E1
- F. 0 and E2

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8

Which four statements are correct regarding IS-IS operations? (Choose four.)

- A. By default, Level 1 routers within an IS-IS area do not carry any routing information external to the area to which they belong. They use a default route to exit the area.
- B. Summarization should be configured on the Level 2 routers, which injects the Level 2 routes into Level 1.
- C. IS-IS supports "route leaking" in which selected Level 2 routes can be advertised by a Level 1/Level 2 router into Level 1.
- D. The IS-IS backbone is a contiguous collection of Level 1 capable routers, each of which can be in a different area.
- E. With IS-IS, an individual router is in only one area, and the border between areas is on the link that connects two routers that are in different areas.
- F. Cisco IOS XR Software supports multitopology for IPv6 IS-IS unless single topology is explicitly configured in IPv6 address-family configuration mode.

Correct Answer: ACEF

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.0/routing/configuration/guide/rc3isis.html#wp1148617

Single-Topology IPv6 Support

Single-topology IPv6 support on Cisco IOS XR software allows IS-IS for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is used, only narrow link metrics, also known as old-style type, length, value (TLV) arguments, may be employed. During single-topology operation, one shortest path first (SPF) computation per level is used to compute both IPv4 and IPv6 routes. Using a single SPF is possible because both IPv4 IS-IS and IPv6 IS-IS routing protocols share a common link topology.

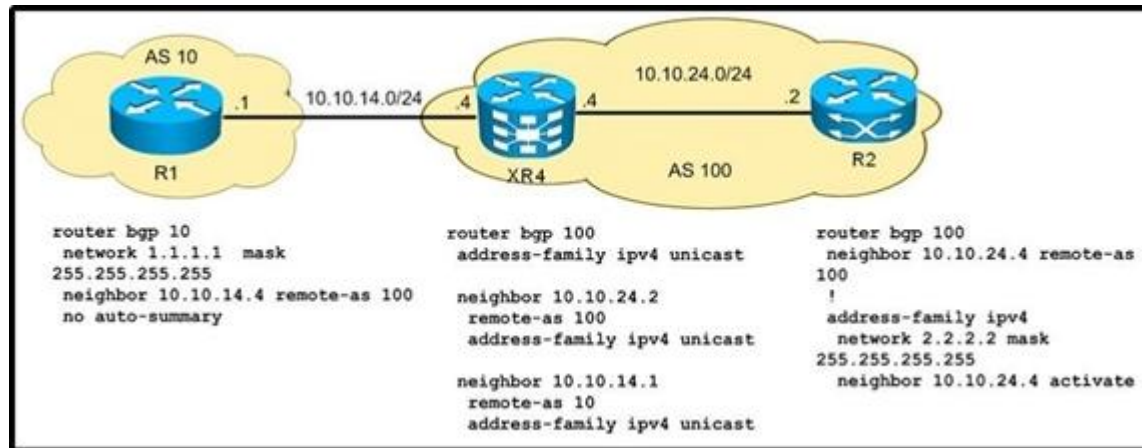
Because multitopology is the default behavior in the software, you must explicitly configure IPv6 to use the same topology as IPv4 in order to enable single-topology IPv6. Configure the singletopology command in ipv6 address family configuration submode of the IS-IS router stanza.

Multitopology IPv6 Support

Multitopology IPv6 support on Cisco IOS XR software differs from Cisco IOS software in that ISIS assumes that multitopology support is required as soon as it detects interfaces configured for both IPv6 and IPv4 within the IS-IS stanza. You must use the metric-style wide command to configure IS-IS to wide link metrics as multitopology link advertisements.

QUESTION 9

Refer to the exhibit. R2 does not mechanism that can be implemented between table. Which issue can be causing this problem?



- A. XR4 does not have next-hop-self for a R2 iBGP session.
- B. R2 does not have no auto-summary configured.
- C. XR4 drops any update coming from R1.
- D. IPv4 address-family is not enabled on R1.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 10

Which option is where BGP graceful restart is negotiated between peers?

- A. open message
- B. keep-alive message
- C. notification message
- D. update message

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 11

Refer to the exhibit. Which configuration is required on XR2 to enable MD5 IS-IS Authentication between R1 and XR2 using the password CISCO?



- A. Router isis ABC interface GigabitEthernet0/0/0/0 hello-password encrypted CISCO
- B. Router isis ABC interface GigabitEthernet0/0/0/0 hello-password hmac-md5 CISCO
- C. Router isis ABC lsp-password hmac-md5 CISCO
- D. Router isis ABC lsp-password encrypted CISCO

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

QUESTION 12

Hotspot

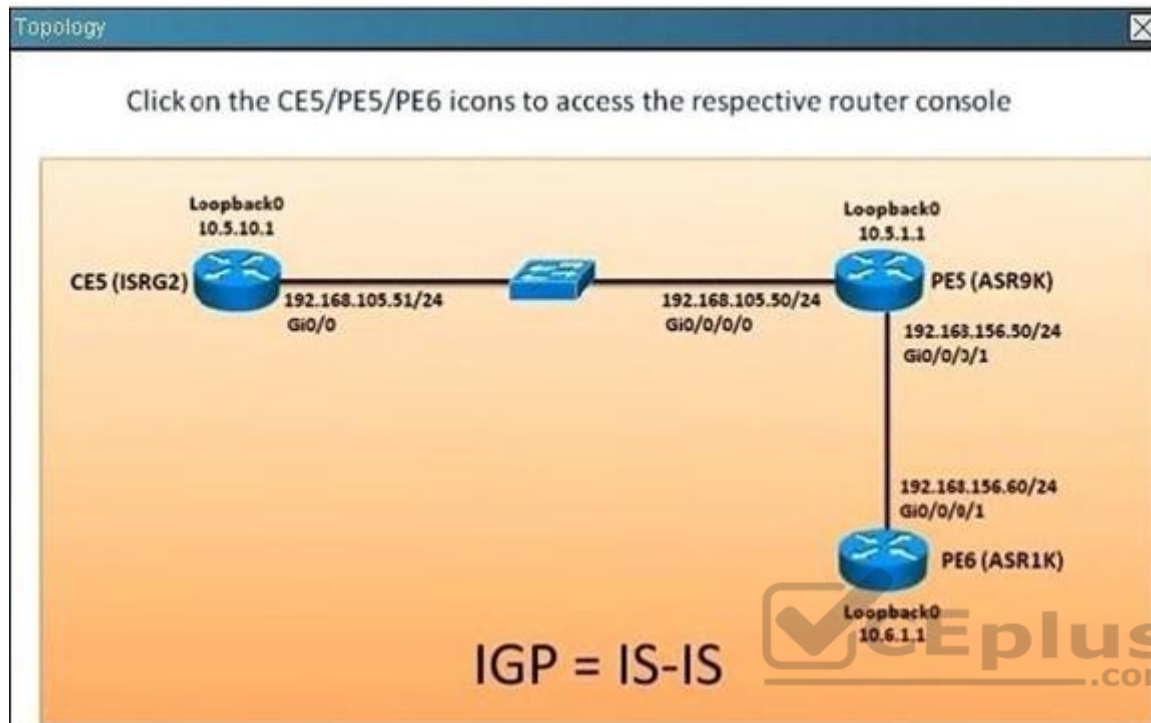
Instructions

- Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.
- From the network topology diagram, click on each of the router icon to gain access to the console of each router.
- No console or enable passwords are required.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.
- Not all the CLI commands or commands options are supported or required for this simulation.
- For example, the show running-config and the ping commands are **NOT** supported in this simulation.
- All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



Which router has the su 10.5.10.0/24 [20/115] entry pointing to Null0 in its routing table and why?
(Choose two.)

- A. CE5
- B. PE5
- C. PE6
- D. The router is summarizing 10.5.10.1/32 into 10.5.10.0/24
- E. The router is suppressing the 10.5.10.0/24 route
- F. The L2/L1 IS-IS router is blocking the 10.5.10.0/24 interarea route to the L1-only router

Correct Answer: DF

Section: (none)

Explanation

Explanation/Reference:

show clns route

QUESTION 13

Hotspot


Instructions

- Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.
- From the network topology diagram, click on each of the router icon to gain access to the console of each router.
- No console or enable passwords are required.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.
- Not all the CLI commands or commands options are supported or required for this simulation.
- For example, the show running-config and the ping commands are **NOT** supported in this simulation.
- All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



What is the configured IS-IS circuit-type between the CE5 and PE5 routers and between the PE5 and PE6 routers? (Choose two.)

- A. Circuittype L1 between CE5 and PE5 4
- B. Circuittype L1/L2 between CE5 and PE5
- C. Circuittype L2 between CE5 and PE5
- D. Circuittype L1 between PE5 and PE6
- E. Circuittype L1/L2 between PE5 and PE6
- F. Circuittype L2 between PE5 and PE6

Correct Answer: AF

Section: (none)

Explanation

Explanation/Reference:

#show cns neighbor

QUESTION 14

Hotspot


Instructions

- Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.
- From the network topology diagram, click on each of the router icon to gain access to the console of each router.
- No console or enable passwords are required.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.
- Not all the CLI commands or commands options are supported or required for this simulation.
- For example, the show running-config and the ping commands are **NOT** supported in this simulation.
- All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



There are how many IS-IS area(s) defined in this network Scenario?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

#show clns

QUESTION 15

Hotspot


Instructions

- Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.
- From the network topology diagram, click on each of the router icon to gain access to the console of each router.
- No console or enable passwords are required.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.
- Not all the CLI commands or commands options are supported or required for this simulation.
- For example, the show running-config and the ping commands are **NOT** supported in this simulation.
- All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



Which router(s) is/are IS-Type L1/L2 IS-IS router?

- A. CE5 only
- B. PE6 only
- C. CE5 and PE6 only
- D. PE5 and PE6 only
- E. CE5, PE5 and PE6

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

#show cns neighbor

QUESTION 16

When using the Cisco IOS XR route policy language to define a logical if-then-else condition, which logical operator has the highest precedence?

- A. AND
- B. OR
- C. NOT
- D. IS
- E. IN

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.0/routing/configuration/guide/rc3rpl.html

Boolean Operator Precedence

Boolean expressions are evaluated in order of operator precedence, from left to right. The highest precedence operator is not, followed by and, and then or. The following expression:

```
med eq 10 and not destination in (10.1.3.0/24) or community matches-any ([10..25]:35)
```

if fully parenthesized to display the order of evaluation would look like this:

```
(med eq 10 and (not destination in (10.1.3.0/24))) or community matches-any ([10..25]:35)
```

The inner not applies only to the destination test; the and combines the result of the not expression with the Multi Exit Discriminator (MED) test; and the or combines that result with the community test. If the order of operations are rearranged:

```
not med eq 10 and destination in (10.1.3.0/24) or community matches-any ([10..25]:35)
```

then the expression, fully parenthesized, would look like the following:

```
((not med eq 10) and destination in (10.1.3.0/24)) or community matches-any ([10..25]:35)
```

QUESTION 17

Refer to the Cisco IOS route map configuration exhibit.


```
route-map test permit 10
match ip address prefix-list PL1 PL2
match as-path APACL1
set local-preference 200
set metric 1000
!
route-map test permit 100
```

Which two statements are correct? (Choose two.)

- A. The match prefix-list condition is a logical OR: match prefix list PL1 OR PL2.
- B. All match conditions are logical OR: match prefix list PL1 OR PL2 OR match the APACL1 AS path access list.
- C. The three match conditions are logical AND: match prefix list PL1 AND PL2 AND match the APACL1 AS path access list.
- D. The local preference AND the metric will be set to 100 IF the route matches the PL1 OR PL2 prefix list AND the route must also match the APACL1 AS path access list.
- E. All routes that are not matched by the sequence 10 route map statement will be dropped.

Correct Answer: AD

Section: (none)

Explanation



Explanation/Reference:

http://www.routeralley.com/ra/docs/route_maps.pdf

When match criteria is contained within a single line, a logical OR is applied.

QUESTION 18

Refer to the Cisco IOS XR route policy exhibit.

```
route-policy SetLP
if med eq 10 then
set local-preference 200
endif
if local-preference eq 100 then
set weight 100
endif
if local-preference eq 200 then
set weight 200
endif
end-policy
```

If the original incoming routing update has an MED of 10 and a local preference of 100, how will the routing update be modified?

- A. The local preference will be set to 100, the MED will be set to 10, and the weight will be set to 100.
- B. The local preference will be set to 100, the MED will be set to 10, and the weight will be set to 200.
- C. The local preference will be set to 200, the MED will be set to 10, and the weight will be set to 100.
- D. The local preference will be set to 200, the MED will be set to 10, and the weight will be set to 200.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 19

Which two mandatory tasks must an IS-IS NSF-capable router perform for RP switchover?
(Choose two.)

- A. Relearn the available IS-IS neighbors.
- B. Reacquire the contents of the LSD.
- C. Reset peering with the available IS-IS neighbors.
- D. Keep the existing contents of the LSD.
- E. Rediscover DIS for each link segment.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

QUESTION 20

When creating iBGP multipaths which three criteria must be met by multiple paths to the same destination? (Choose three.)

- A. Router IDs must be the same on all routers.
- B. Each destination must have a different next-hop address.
- C. The destination AS-number must be different for each destination.
- D. Multi-exit discriminator attributes must be the same on all paths.
- E. Interior Gateway Protocol distance must be identical on each path.

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:



QUESTION 21

What can cause a single sourced iBGP route not to be selected as the best route?

- A. The BGP MED is 0.
- B. The BGP next-hop is unreachable.
- C. The BGP origin is incomplete.
- D. The BGP weight is 0.
- E. The BGP local preference is 0.
- F. BGP synchronization is disabled.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 22

Which two of the following are true regarding the BGP Prefix-Based outbound route filtering feature? (Choose two.)

- A. IP multicast routes are not supported.
- B. Outbound route filtering is configured only on a per-address family basis.
- C. Outbound route filtering can be configured for either iBGP or eBGP sessions.
- D. The outbound route filter can be defined in a Prefix list, Distribute list or Access lists.
- E. Outbound route filtering is more effective when a distance vector IGP is used.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

QUESTION 23

Which three conditions can make a transit autonomous system unable to forward packets from one neighboring autonomous system to another neighboring autonomous system? (Choose three.)

- A. sync enabled
- B. bgp next-hop not reachable
- C. no full-mesh IBGP and not using route reflectors or confederations
- D. non-directly connected IBGP neighbors
- E. non-directly connected EBGP neighbors
- F. EBGP to IBGP route redistribution configuration errors

Correct Answer: ABC

Section: (none)

Explanation

Explanation/Reference:

QUESTION 24

What are three characteristics of the local preference BGP attribute? (Choose three.)

- A. It is stripped in outgoing EBGp updates except with confederation peers.
- B. It is used to influence the incoming traffic.
- C. Its default local preference is 32768.
- D. It is used to select routes with equal weight.
- E. It can be used to ensure an AS-wide route selection policy.
- F. It is a Cisco-only BGP path attribute.

Correct Answer: ADE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 25

Which series of commands configures area 1 as an OSPF totally stubby area on a Cisco IOS XR router?

- A.

```
router ospfv3 1 router-id 10.10.10.1 area 0
interface GigabitEthernet 0/0/0/1 area 1 stub no-summary
interface GigabitEthernet 0/0/0/2
```
- B.

```
router ospfv3 1 router-id 10.10.10.1 default-information originate area 0
interface GigabitEthernet 0/0/0/1 area 1 stub
interface GigabitEthernet 0/0/0/2
```
- C.

```
ipv6 prefix-list default permit ::0/0 router ospfv3 1 router-id 10.10.10.1 default-information originate area 0 interface GigabitEthernet 0/0/0/1
distribute-list prefix-list default in area 1
interface GigabitEthernet 0/0/0/2
```
- D.

```
router ospfv3 1 router-id 10.10.10.1 area 0
interface GigabitEthernet 0/0/0/1 area 1
interface GigabitEthernet 0/0/0/2
no-summary passive
```
- E.

```
router ospfv3 1 router-id 10.10.10.1 area 0
interface GigabitEthernet 0/0/0/1 area 1 stub
interface GigabitEthernet 0/0/0/2 default-cost 20
```

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 26

Refer to the PE1 router routing table output exhibit.

```
RP/0/RSP0/CPU0:PE1#show route ipv4 isis
```

```
<output omitted>
i su 10.1.10.0/24 [115/30] via 0.0.0.0, 00:40:34, Null0
i L1 10.1.10.1/32 [115/30] via 192.168.101.11, 00:42:39, GigabitEthernet0/0/0/0
i L1 10.1.10.2/32 [115/24] via 192.168.112.21, 00:44:40, GigabitEthernet0/0/0/1
i L1 10.1.10.3/32 [115/32] via 192.168.113.22, 00:38:23, GigabitEthernet0/0/0/2
i L1 10.1.10.4/32 [115/22] via 192.168.114.23, 00:14:10, GigabitEthernet0/0/0/3
<output omitted>
```

What is causing the i su 10.1.10.0/24 [115/30] via 0.0.0.0, 00:40:34, Null0 entry on the PE1 router routing table?

- A. The PE1 router is receiving the 10.1.10.0/24 summary route from the upstream L1/L2 IS-IS router.
- B. The PE1 router has been configured to summarize the 10.1.10.x/32 IS-IS routes to 10.1.10.0/24.
- C. The 10.1.10.0/24 has been suppressed because IS-IS auto-summary has been disabled on the PE1 router.
- D. The 10.1.10.0/24 has been suppressed because of a route policy configuration on the PE1 router.
- E. The 10.1.10.0/24 has been suppressed because the more specific 10.1.10.x/32 IS-IS routes have been configured to leak into the IS-IS non-backbone area.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1
N2 - OSPF NSSA external type 2, E1 - OSPF external type 1
E2 - OSPF external type 2, E - EGP, i - ISIS, L1 - IS-IS level-1
L2 - IS-IS level-2, ia - IS-IS inter area
su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local

QUESTION 27

For a Cisco IOS XR router, under which BGP configuration group can default-originate be added?

- A. session-group
- B. af-group
- C. neighbor-group
- D. peer-session template

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 28

Which of the following is a characteristic of dual-multihomed connectivity between an enterprise network and the service provider network or networks?

- A. An enterprise network that is connected to two or more different service providers with two or more links per service provider and using BGP to exchange routing updates with the service providers.
- B. Each service provider announces a default route on each of the links that connect to the customer with a different metric.
- C. An enterprise network announces a default route to each service provider.
- D. Load balancing can be achieved using the maximum-paths command.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 29

Which BGP attribute is a set of generic tags that can be used to signal various routing policies between BGP routers?

- A. AS path
- B. MED
- C. weight
- D. communities
- E. route tags

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801475b2.shtml

QUESTION 30

Refer to the Cisco IOS-XR route policy exhibit.

```
route-policy setcomm
if community matches-any (11:11, 44:44) then
set community (55:55) additive
elseif community matches-any (22:22) then
set community (77:77) additive
endif
```

Which statement correctly describes this route policy?

- A. The pass action is required after each of the set community statements to make this route policy functional.
- B. If a route has both the 11:11 and 22:22 communities (or 44:44 and 22:22), the router adds the 55:55 community only.
- C. If a route contains the 11:11 or 44:44 community, the router adds the 55:55 community and continues. Additionally, if the same route also contains the 22:22 community, the router also adds another 77:77 community to the same route.
- D. If a route only has the 22:22 community, then no community will be added by the router.

E. If a route only has the 11:11 or 44:44 community, then no community will be added by the router.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 31

An engineer wants to define a route policy that sets MED to a given value in case the extended community attributes of a network prefix match at least one of the values: 65000:10, 65000:20, or 65000:30. Which configuration accomplishes this task?

- A. extcommunity-set rt my-list
65000:10
65000:20 65000:30
end-set
route-policy my-list (\$med)
if extcommunity rt matches-within my-list then set met \$med
- B. extcommunity-set rt my-list
65000:10
65000:20 65000:30
end-set

route-policy my-list (\$med)
if extcommunity rt matches-any my-list then set met \$med
- C. extcommunity-set rt my-list
65000:10
65000:20 65000:30
end-set

route-policy my-list (\$med)
if extcommunity rt matches-within my-list or med is \$med is then set met \$med
- D. extcommunity-set rt my-list
65000:10
65000:20 65000:30
end-set

route-policy my-list (\$med)
if extcommunity rt matches-any my-list or med is \$med then pass

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 32

Which of the following is used by an IS-IS router to detect other IS-IS neighbor routers and to form adjacencies?

- A. ESH
- B. ISH
- C. IIH
- D. PSNP
- E. CSNP

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

ISIS actually features three different hello types An ES Hello (ESH) is send by all End Systems, and all IS devices listen for this Hello An IS Hello (ISH) announces the presence of an IS - An IS Hello is sent by all IS devices, and End Systems listen for these hellos Finally, an IS-to-IS Hello (IIH) is used by an IS to discover other ISes and to form adjacencies with them A router will send an IIH to another router on the link to form or maintain an adjacency, but it will still send an ISH as well in case there are end systems located on that segment

QUESTION 33

When using the show bgp ipv6 unicast summary command to verify the IPv6 BGP session status with the IPv6 BGP peers, you noticed the "St/PfxRcd" status for one of the IPv6 BGP peers is in the "Active" state. What does the "Active" state indicate?

- A. The IPv6 BGP session has been established with the IPv6 BGP peer.
- B. The router is in the process of sending BGP routing updates to the IPv6 BGP peer.
- C. The router is in the process of establishing the IPv6 BGP session with the IPv6 BGP peer.
- D. The router is exchanging BGP notification messages with its IPv6 BGP peer.

Correct Answer: C

Section: (none)

Explanation**Explanation/Reference:****QUESTION 34**

Which RPLs can be used to conditionally advertise a default route based on the prefix 1.1.1.1 being in the routing table?

- A. route-policy RPL_Default_Originate if rib-has-route in (1.1.1.1/32)
pass else drop endif end-policy
- B. route-policy RPL_Default_Originate if rib-has-route in (1.1.1.1/32)
drop else pass endif end-policy
- C. route-policy RPL_Default_Originate if route in (1.1.1.1/32)
pass else drop endif end-policy
- D. route-policy RPL_Default_Originate if not rib-has-route in (1.1.1.1/32)
pass else drop endif end-policy

Correct Answer: A

Section: (none)

Explanation**Explanation/Reference:****QUESTION 35**

Which statement best describes the function of a BGP af-group?

- A. Use templates in different neighbors so it makes the BGP implementation easier.
- B. Group different address family neighbors inside BGP.
- C. Create a group of neighbors that can inherit address family configuration inside BGP.
- D. Substitute the class maps inside BGP on Cisco IOS XR devices.

Correct Answer: B

Section: (none)

Explanation**Explanation/Reference:**

QUESTION 36

When configuring IPv4 and IPv6 IS-IS routing on Cisco IOS XR routers, which three statements are correct? (Choose three.)

- A. By default, a single SPF is used for both IPv4 and IPv6, so the IPv4 and IPv6 topology should be the same.
- B. By default, the IS-IS router type is Level 1 and Level 2.
- C. All IS-IS routers within the same IS-IS area must be configured with the same IS-IS routing process instance ID.
- D. By default, metric-style narrow is used.
- E. By default, the IS-IS interface circuit type is Level 1 and Level 2.
- F. The area IS-IS address-family configuration command is used to specify the IS-IS area address.

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:



```
is-type {level-1 | level-1-2 | level-2-only}
```

Example:

```
RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only
```

(Optional) Configures the system type (area or backbone router).

- By default, every IS-IS instance acts as a **level-1-2** router.
- The **level-1** keyword configures the software to perform **Level 1** (intra-area) routing only. Only **Level 1** adjacencies are established. The software learns about destinations inside its area only. Any packets containing destinations outside the area are sent to the nearest **level-1-2** router in the area.
- The **level-2-only** keyword configures the software to perform **Level 2** (backbone) routing only and the router establishes only **Level 2** adjacencies, either with other **Level 2-only** routers or with **level-1-2** routers.
- The **level-1-2** keyword configures the software to perform both **Level 1** and **Level 2** routing. Both **Level 1** and **Level 2** adjacencies are established. The router acts as a border router between the **Level 2** backbone and its **Level 1** area.

The default metric style for single topology is narrow metrics. However, you can use either wide metrics or narrow metrics. How to configure them depends upon how single-topology is configured. If both IPv4 and IPv6 are enabled and single-topology is configured, the metric style is configured in the address-family ipv4 stanza. You may configure the metric style in the address-family ipv6 stanza, but it will be ignored in this case. If IPv6 only is enabled and single topology is configured, then the metric style is configured in the address-family ipv6 stanza.

| | |
|---|--|
| <pre>circuit-type {level-1 level-2-only level-1-2}</pre> <p>Example: RP/0/RP0/CPU0:router(config-isis-if)# circuit-type level-1-2</p> | <p>(Optional) Configures the type of adjacency.</p> <ul style="list-style-type: none"> The default circuit type is the configured system type (configured through the is-type command). Typically, circuit type needs to be configured when the router is configured only level-1-2 and you want to constrain an interface to form only level-1 or level-2-only adjacencies. |
|---|--|

QUESTION 37

In which network environment is IS-IS adjacency check important?

- A. in a multitopology environment where there are different instances of IS-IS running on the same router
- B. in an IPv4/IPv6 environment and running single-topology IS-IS
- C. when a level L1/L2 IS-IS router is neighboring with a Level 1 only or Level 2 only router
- D. when IS-IS neighbors are in an NBMA environment like over Frame Relay
- E. when IS-IS neighbors are in a broadcast environment like an Ethernet LAN

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Disabling IPv6 Protocol-Support Consistency Checks

Perform this task to disable protocol-support consistency checks in IPv6 single-topology mode.

For single-topology IS-IS IPv6, routers must be configured to run the same set of address families. IS-IS performs consistency checks on hello packets and will reject hello packets that do not have the same set of configured address families. For example, a router running IS-IS for both IPv4 and IPv6 will not form an adjacency with a router running IS-IS for IPv4 or IPv6 only. In order to allow adjacency to be formed in mismatched address-families network, the adjacency-check command in IPv6 address family configuration mode must be disabled.

Entering the no adjacency-check command can adversely affect your network configuration. Enter the no adjacency-check command only when you are running IPv4 IS-IS on all your routers and you want to add IPv6 IS-IS to your network but you need to maintain all your adjacencies during the transition. When the IPv6 IS-IS configuration is complete, remove the no adjacency-check command from the configuration.

QUESTION 38

Hotspot

Instructions ✕

Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.

From the network topology diagram, click on each of the router icon to gain access to the console of each router.

No console or enable passwords are required.

There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Not all the CLI commands or commands options are supported or required for this simulation.

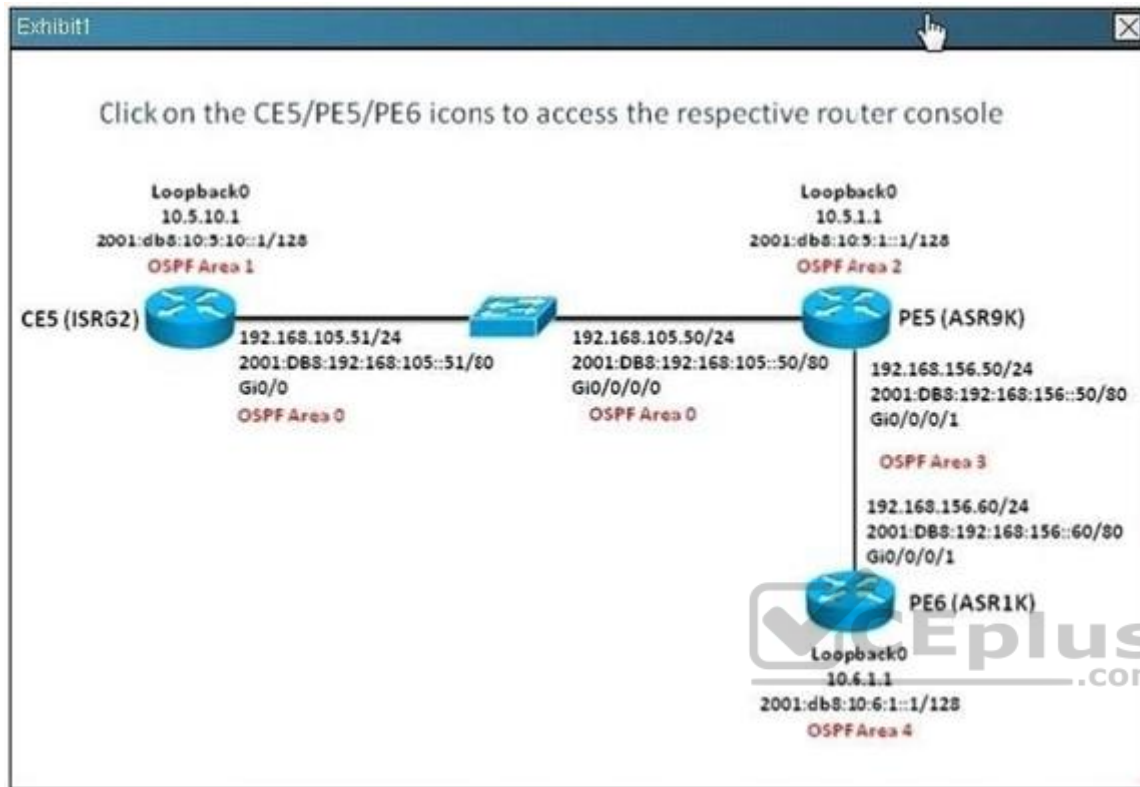
For example, the show running-config and the ping commands are **NOT** supported in this simulation.

All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario ✕

Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



Which statement is correct regarding the DR election process In this OSPF implementation?

- A. Both the CE5 and PE6 routers are DROTHERs
- B. PE5 is the OSPFv2 DR for the GigabitEthernet link between PE5 and PE6
- C. PE6 is the OSPFv2 and OSPFv3 DR for the GigabitEthernet link between PE5 and PE6
- D. PE6 is the DR for the GigabitEthernet link between PE5 and PE6 because it has a higher OSPF priority

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:


```
# show ip ospf neighbor  
# show ipv6 ospf neighbor
```

QUESTION 39

Hotspot

Instructions

Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.

From the network topology diagram, click on each of the router icon to gain access to the console of each router.

No console or enable passwords are required.

There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Not all the CLI commands or commands options are supported or required for this simulation.


For example, the show running-config and the ping commands are **NOT** supported in this simulation.

All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



Which three statements are correct regarding the OSPF operations? line (Choose three.)

- A. PE5 OSPFv2 and OSPFv3 router id is 10.5.1.1
- B. The OSPF virtual-link cost is 1
- C. Area 3 is a transit area
- D. Area 4 is using MD5 authentication
- E. Area 1 IPv4 and IPv6 networks are not reachable from PE6

Correct Answer: CDE

Section: (none)

Explanation

Explanation/Reference:

```
#show ip protocols
#show ip route ospf
#show ip ospf interfaces
#show ip ospf neighbors
#show ip ospf database
#show ip ospf border-routers
#show ip ospf
#show ip route
#show ip protocols
```

QUESTION 40

Hotspot

Instructions

Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.

From the network topology diagram, click on each of the router icon to gain access to the console of each router.

No console or enable passwords are required.

There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Not all the CLI commands or commands options are supported or required for this simulation.


For example, the show running-config and the ping commands are **NOT** supported in this simulation.

All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



On the PE5 router, there are how many IPv4 and IPv6 OSPF neighbors and how many IPv4 OSPF interfaces? (Choose three)

- A. Two ipv4 ospf neighbors
- B. Three ipv4 ospf neighbors
- C. Two ipv6 ospf neighbors
- D. Three ipv6 ospf neighbors
- E. Three ipv4 ospf interfaces
- F. Four ipv4 ospf interfaces

Correct Answer: AEF

Section: (none)

Explanation

Explanation/Reference:

show ip ospf neighbor
show ipv6 ospf neighbor

QUESTION 41

Hotspot

Instructions

Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.

From the network topology diagram, click on each of the router icon to gain access to the console of each router.

No console or enable passwords are required.

There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Not all the CLI commands or commands options are supported or required for this simulation.

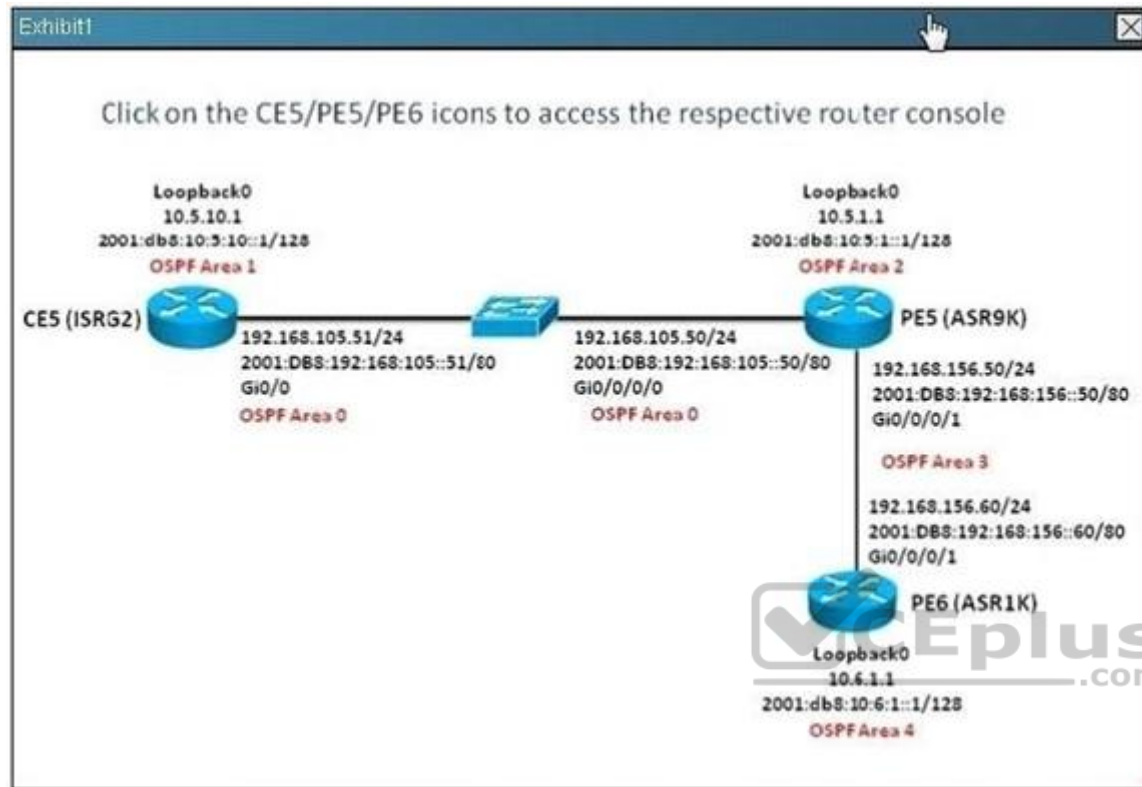
For example, the show running-config and the ping commands are **NOT** supported in this simulation.

All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario

Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.



How many total OSPF virtual links are configured in this multiarea OSPF network which supports both OSPFv2 and OSPFv3 for IPv4 and IPv6 routing?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

#Show running-config | include virtual - check how many router having "virtual " entry

QUESTION 42

Which high-availability routing feature requires the neighbor router to support the graceful restart capability?

- A. BFD
- B. NSR
- C. NSF
- D. MTR

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

On Cisco IOS XR software, NSF minimizes the amount of time a network is unavailable to its users following a route processor (RP) failover. The main objective of NSF is to continue forwarding IP packets and perform a graceful restart following an RP failover.

When a router restarts, all routing peers of that device usually detect that the device went down and then came back up. This transition results in what is called a routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. NSF helps to suppress routing flaps in NSF-aware devices, thus reducing network instability.

NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following an RP failover. When the NSF feature is configured, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards while the standby RP assumes control from the failed active RP during a failover. The ability of line cards to remain up through a failover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to NSF operation.

When the Cisco IOS XR router running IS-IS routing performs an RP failover, the router must perform two tasks to resynchronize its link-state database with its IS-IS neighbors. First, it must relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the link-state database for the network. The IS-IS NSF feature offers two options when configuring NSF:

IETF NSF

Cisco NSF

If neighbor routers on a network segment are NSF aware, meaning that neighbor routers are running a software version that supports the IETF Internet draft for router restartability, they assist an IETF NSF router that is restarting. With IETF NSF, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a failover.

In Cisco IOS XR software, Cisco NSF checkpoints (stores persistently) all the state necessary to recover from a restart without requiring any special cooperation from neighboring routers. The state is recovered from the neighboring routers, but only using the standard features of the IS-IS routing protocol. This capability makes Cisco NSF suitable for use in networks in which other routers have not used the IETF standard implementation of NSF

QUESTION 43

In comparing IS-IS with OSPF, a Level-1-2 IS-IS router is similar to which kind of OSPF router?

- A. ASBR on a normal OSPF area
- B. ASBR on NSSA
- C. ABR on totally stubby OSPF area
- D. ABR on stubby OSPF area
- E. ABR on a normal OSPF area

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 44

What is recursive lookup in BGP and how does it work?

- A. The router looks up the EBGp route and the EBGp next hop to reach a destination in the remote AS. Then the router looks up the route to reach the EBGp next hop using the IBGP.
- B. The router looks up the IBGP route and the IBGP next hop to reach a destination in the remote AS. Then the router looks up the route to reach the IBGP next hop using the EBGp.
- C. The router looks up the BGP route and the BGP next hop to reach a destination in the remote AS. Then the router looks up the route to reach the BGP next hop using the IGP.
- D. The router looks up the route and the next hop to reach a destination in the remote AS using the IGP. Then the router looks up the route to reach the next hop using BGP.
- E. The router perform three routing lookups to determine the route to reach a destination in the remote AS. The first lookup is done using EBGp, the second lookup is done using IBGP, and the third lookup is done using the IGP.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

A few different approaches are available to deal with iBGP and synchronization. We may turn on the synchronization option on our routers and wait for the IGP to have a route for the destination before it's advertised to peers. Another option is to simply use a full mesh, so that iBGP convergence isn't an issue.

Clearly that isn't going to happen when a network's core needs to scale: it will implement something like reflectors that cause iBGP's full mesh to be broken. The real alternative, if you don't enable synchronization, is to use route recursion. A recursive route lookup uses the BGP next-hop attribute to actually make a different route lookup. The IGP can use the destination network instead of the AS-path to determine where it gets sent. Even if the iBGP hasn't converged, the routers will still know how to get to that network, since it will exist in the router it was advertised from, who will know the next-hop.

QUESTION 45

A static default route has been inserted into the configuration of a Cisco IOS XE router. Which option redistributes the route into the local OSPF process?

- A. router ospf 1 redistribute static subnets
- B. router ospf 1 redistribute static
- C. router ospf 1 redistribute static subnets default-information originate always
- D. router ospf 1 network 0.0.0.0 area 0 redistribute static subnets

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 46

What is function of the RP/0/RSP0/CPU0:PE1(config-ospf)#distance Cisco IOS-XR command?

- A. To modify the administrative distance of the OSPF routes
- B. To modify the default seed metric of the OSPF external routes
- C. To modify the OSPF default reference bandwidth
- D. To modify the OSPF cost

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.8/routing/command/reference/rr38ospf.pdf

distance (OSPF)

To define an administrative distance, use the **distance** command in the appropriate mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address wildcard-mask* [*access-list-name*]]

no distance *weight ip-address wildcard-mask* [*access-list-name*]

| Syntax Description | | |
|-------------------------|--|---|
| <i>weight</i> | | Administrative distance. Range is 10 to 255. Used alone, the <i>weight</i> argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. Table 1 lists the default administrative distances. |
| <i>ip-address</i> | | (Optional) IP address in four-part, dotted-decimal notation. |
| <i>wildcard-mask</i> | | (Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the <i>mask</i> argument instructs the software to ignore the corresponding bit in the address value. |
| <i>access-list-name</i> | | (Optional) Name of an IP access list to be applied to incoming routing updates. |

QUESTION 47

Refer to the exhibit.

```
router ospf 100
bfd minimum-interval 2
bfd multiplier 20
area 0
interface gigabitEthernet 0/3/0/1
interface gigabitEthernet 0/3/0/2
end
```

Which Cisco IOS XR configuration is missing to complete the configuration task of enabling BFD with only the OSPF peer over the gi0/3/0/1 interface in area 0?

- A. bfd fast-detect also needs to be enabled globally under router ospf 100. RP/0/RSP0/CPU0:P1(config-ospf)#bfd fast-detect
- B. bfd fast-detect also needs to be enabled for the OSPF area under area 0. RP/0/RSP0/CPU0:P1(config-ospf-ar)#bfd fast-detect
- C. bfd fast-detect also needs to be enabled for the OSPF interface under area 0 interface gi0/3/0/1. RP/0/RSP0/CPU0:P1(config-ospf-ar-if)#bfd fast-detect

- D. bfd fast-detect also needs to be enabled globally on the router. RP/0/RSP0/CPU0:P1(config)#bfd fast-detect
- E. bfd fast-detect also needs to be enabled on the gi0/3/0/1 interface under interface gi0/3/0/1. RP/0/RSP0/CPU0:P1(config-if)#bfd fast-detect

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

bfd fast-detect

To enable Bidirectional Forwarding Detection (BFD) to detect failures in the path between adjacent forwarding engines, use the bfd fast-detect command in the appropriate configuration mode. To return the software to the default state in which BFD is not enabled, use the no form of this command.

bfd fast-detect [disable | ipv4] no bfd fast-detect

QUESTION 48

Which configuration is used when matching a prefix-list <name> in a route map?

- A. match ip address prefix-list mylist
- B. match ip address mylist
- C. match ip next-hop mylist
- D. ip local-policy route-map mylist
- E. ip-policy route-map mylist



Correct Answer: A


Section: (none)

Explanation

Explanation/Reference:

QUESTION 49

Refer to the exhibit.



Routers B and C in transit AS 200 are the exit points toward AS 100 and AS 300. Router B is running Cisco IOS XR and Router C IOS XE. Internally, AS 200 runs OSPF only for internal routing. Which two methods can an engineer use to ensure proper reachability between AS 100 and AS 300 for the needed routes? (Choose two.)

- A. Router B:
router ospf 200 redistribute bgp 200 tag 20 router bgp 200 address-family ipv4 unicast redistribute ospf 200 metric-type 2
- B. Router B:
router ospf 200 redistribute bgp 200 router bgp 200 address-family ipv4 unicast redistribute ospf 200
- C. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200 router bgp 200 address-family ipv4 redistribute ospf 200
- D. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200 router bgp 200 address-family ipv4
redistribute ospf 200 match external 2
- E. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200 router bgp 200 address-family ipv4
redistribute ospf 200 route-map match-tag match external route-map match-tag deny 10 match tag 200

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 50

Which three statements are true regarding the OSPF router ID? (Choose three.)

- A. The OSPF routing process chooses a router ID for itself when it starts up.
- B. The router-id command is the preferred procedure to set the router ID.
- C. If a loopback interface is configured, its address will always be preferred as the router ID over any other methods.
- D. After the router ID is set, it does not change, even if the interface that the router is using for the router ID goes down. The router ID changes only if the router reloads or if the OSPF routing process restarts.
- E. In OSPF version 3, the OSPF router ID uses a 128-bit number.

Correct Answer: ABD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 51

When redistributing EIGRP routes into OSPF as type E2 external OSPF routes, what is the default OSPF seed metric?

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

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

QUESTION 52

Which reserved AS number or range of numbers is used for backward compatibility between old BGP peers using 16-bit AS number and new BGP peers using 32-bit AS number?

- A. AS 65001 to 65535
- B. AS 65512 to 65535
- C. AS 12345
- D. AS 23456
- E. AS 64001

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6554/ps6599/4byte_asnios.pdf

New Reserved AS#

AS_TRANS = AS #23456

2-byte placeholder for a 4-byte AS number

Used for backward compatibility between OLD and NEW BGP speakers

QUESTION 53

A network engineer wants to set the local preference to 100 for the specific BGP community sets of 100:1, 101:1, and 102:1. Which option can the engineer use to implement this through RPL, where the BGP community set list is made by the name of lowpri-communities?

- A. route-policy ch-pref if community matches lowpri-communities then set local-preference 100
endif end-policy
- B. route-policy ch-pref if community eq lowpri-communities then set local-preference 100
endif end-policy
- C. route-policy ch-pref if community matches-any lowpri-communities then
set local-preference 100
endif end-policy
- D. route-policy ch-pref if community neq lowpri-communities then set local-preference 100
endif end-policy

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 54

In Cisco IOS and Cisco IOS XE Software images, when redistributing routes from other routing protocols into OSPF, what is a common reason why some of the routes might not be redistributed into OSPF?

- A. The OSPF external metric type (E1 or E2) is not defined.
- B. The OSPF seed metric is not defined.
- C. The OSPF level (Level 1, Level 2, or Level-1-2) to which the routes will be redistributed into is not defined.
- D. The subnets option in the redistribute command is missing.

Correct Answer: D

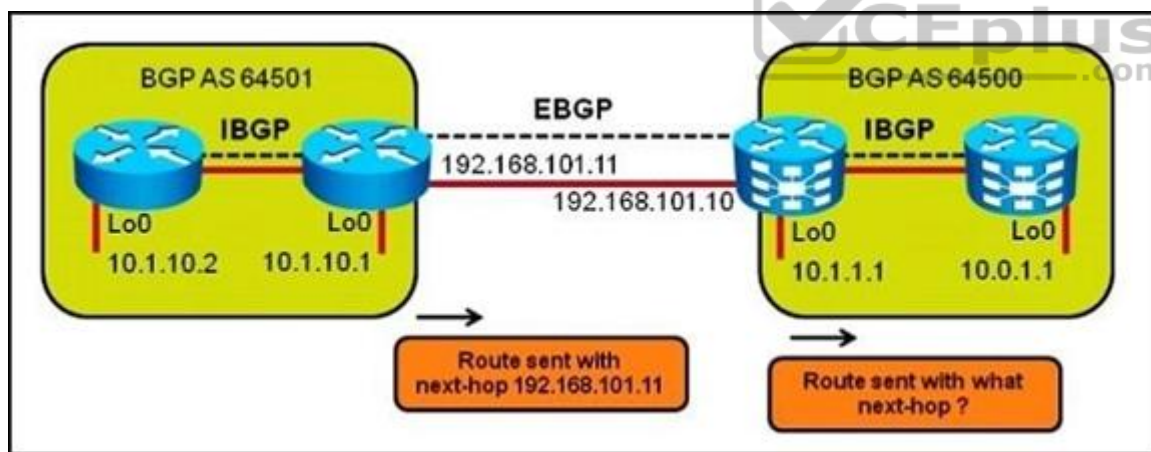
Section: (none)

Explanation

Explanation/Reference:

QUESTION 55

Refer to the network diagram in the exhibit.



Assuming the IBGP session within AS 64500 was established using the loopback 0 interface between the two routers, by default, what will be the next hop of the routes from AS 64501 when the routes appear on the router running IBGP only in AS 64500?

- A. 192.168.101.11

- B. 192.168.101.10
- C. 10.1.1.1
- D. 10.0.1.1
- E. 10.1.10.1

Correct Answer: A

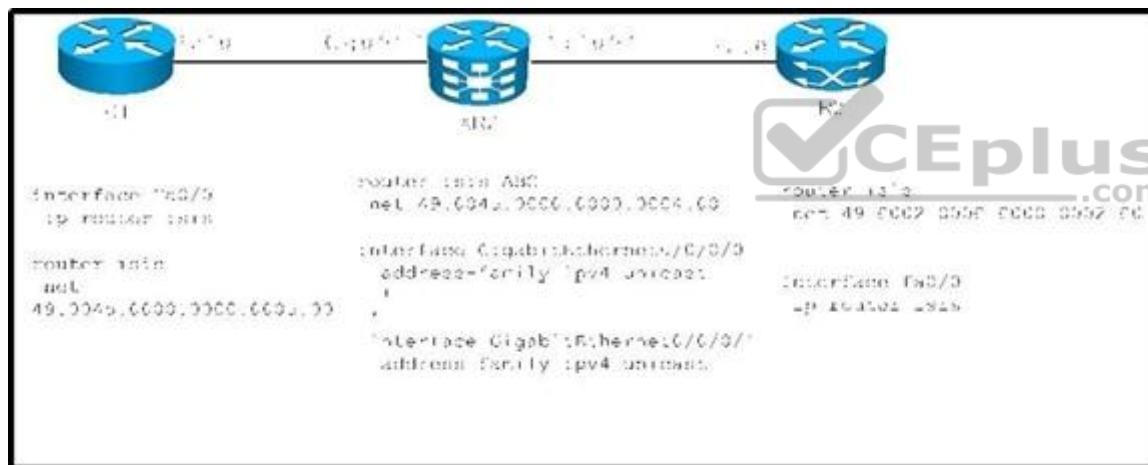
Section: (none)

Explanation

Explanation/Reference:

QUESTION 56

Refer to the exhibit.



Which configuration on XR2 provides R1 with only the minimum routing information to reach the rest of the network?

- A. Router isis ABC interface GigabitEthernet0/0/0/1 circuit-type level-2-only
- B. Router isis ABC interface GigabitEthernet0/0/0/0 circuit-type level-1-2
- C. Router isis ABC interface GigabitEthernet0/0/0/1 circuit-type level-1
- D. Router isis ABC interface GigabitEthernet0/0/0/0 circuit-type level-1

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 57

When configuring Cisco IOS XR route policy nesting, which command is used within a route policy to call another route policy?

- A. Apply
- B. Continue
- C. Jump
- D. Go to
- E. Call

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/routing/command/reference/rr37plcy.html

Examples

In the following example, the policy CustomerIn applies the route-policy SetLocalPref to conditionally set the local preference on a route. The parameters 20, 30, 40, and 50 are passed to the parameterized policy SetLocalPref, where the local preference is set to:

- 20, if the community 217:20 is present in the route
- 30, if the community 217:30 is present in the route
- 40, if the community 217:40 is present in the route
- 50, if the community 217:50 is present in the route

```
RP/0/RP0/CPU0:router(config)# route-policy SetLocalPref ($lp0, $lp1, $lp2, $lp3, $lp4)
RP/0/RP0/CPU0:router(config-rpl)# if community matches-any ($lp0:$lp1)then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp1
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp2) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp2
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp3) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp3
RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp4) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp4
```

```
RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# route-policy CustomerIn($cust) RP/0/RP0/CPU0:router(configrpl)# apply SetLocalPref ($cust, 20, 30, 40, 50) RP/0/RP0/
CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# route-policy Cust_217
RP/0/RP0/CPU0:router(config-rpl)# apply CustomerIn(217)
RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

QUESTION 58

A network engineer wants to test a path between any two points in the network for multicast traffic to identify congestion points in the network. Which command in a Cisco IOS XE router shows multicast traffic drops?

- A. mstat
- B. mtrace
- C. mrinfo
- D. ping

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 59

What are three common problems that can cause a BGP neighbor state to toggle between the idle state and the active state? (Choose three.)

- A. BGP network command misconfiguration
- B. route policy misconfiguration
- C. AS number misconfiguration
- D. route map misconfiguration
- E. BGP neighbor peering to wrong IP address
- F. IGP routing problem: not able to reach the source IP address of the BGP open packet

Correct Answer: CEF

Section: (none)

Explanation

Explanation/Reference:

QUESTION 60

Refer to the Cisco IOS XE IS-IS configuration exhibit.

```
interface GigabitEthernet0/0/0
ip address 192.168.104.40 255.255.255.0
ip router isis
ipv6 address 2001:DB8:192:168:104::40/80
ipv6 enable
ipv6 router isis
!
interface GigabitEthernet0/0/1
ip address 192.168.134.40 255.255.255.0
ip router isis
ipv6 address 2001:DB8:192:168:134::40/80
ipv6 enable
ipv6 router isis
!
router isis 1
net 49.0004.0100.0400.1001.00
summary-address 10.4.10.0 255.255.255.0
!
```



What are two problems with the configuration that are causing the IPv4, or the IPv6, or the IPv4 and IPv6 IS-IS operations to fail? (Choose two.)

- A. The configuration is missing the interface gi0/0 and interface gi0/0 commands under router isis 1. B. The configuration is missing the address-family ipv4 unicast and address-family ipv6 unicast commands under router isis 1.
- B. IPv6 unicast routing has not been enabled globally on the Cisco IOS XE router.
- C. The gi0/0 and gi0/1 interfaces are not participating in the router isis 1 routing instance.
- D. Multitopology IS-IS must be enabled to support both IPv4 and IPv6.
- E. Another router isis instance must be enabled to support IPv6.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/switches/metro/me3400/software/release/12.2_50_se/configuration/guide/swipv6.pdf

When using user-network interface (UNI) or enhanced network interface (ENI) ports for any IPv6-related features, you must first globally enable IP routing and IPv6 routing on the switch by entering the ip routing ipv6 unicast-routing

QUESTION 61

When implementing OSPF, which type of networks require DR/BDR election?

- A. point-to-point networks
- B. multi-access broadcast networks
- C. non-broadcast multi-access networks (Hub and Spoke Frame Relay) using point-to-multipoint OSPF network type
- D. All networks type

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 62

Refer to the Cisco IOS XR route policy exhibit.

```
route-policy new
  if community matches-any test then
    set local-preference 10
  elseif community matches-every foo then
    set local-preference 20
  else
    pass
  endif
end-policy
```

Which two statements are correct? (Choose two.)

- A. "test" and "foo" are references to the community sets.
- B. The "match-any" option indicates matching either the standard community or extended community.
- C. Routes that match both the "test" and "foo" conditions will have their local preference set to 20.

D. All non-matching routes will be permitted.

Correct Answer: AD


Section: (none)

Explanation

Explanation/Reference:

QUESTION 63

Refer to the exhibit. Which two configuration options can be used to optimize the IS-IS network scenario? (Choose two.)



- A. Change the R1 and R2 IS type to Level 2.
- B. Change the R1 and R2 IS type to Level 1.
- C. Change the gi0/0 interface IS-IS circuit type on R1 and R2 to Level 2 only.
- D. Change the gi0/1 interface IS-IS circuit type on R1 and R2 to Level 1.
- E. Change the IS type for all the routers to Level-1-2.
- F. Change the IS-IS network type for all the routers to point-to-point.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 64

Which option describes the configuration for a Cisco IOS XR OSPF router to give it the highest chances of becoming DR?

- A. ip ospf priority 0 under the OSPF interface statement
- B. ip ospf priority 0 under the OSPF area statement
- C. ip ospf priority 255 under the OSPF interface statement
- D. ip ospf priority 255 under the OSPF area statement
- E. Leave priority unchanged, but assign a high-numbered IP address on a loopback interface.
- F. Leave priority unchanged, but assign a high-numbered IP address on any physical interface.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 65

A network engineer is configuring IS-IS in Cisco IOS XR. Where is BFD configured?

- A. RP/0/RSP0/CPU0:router(config-isis-if)#
- B. RP/0/RSP0/CPU0:router(config)#
- C. RP/0/RSP0/CPU0:router(config-bfd)#
- D. RP/0/RSP0/CPU0:router(config-isis-if-af)#



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 66

When configuring BGP on Cisco IOS XR Software, which address-family is enabled by default?

- A. IPv4 unicast
- B. IPv6 unicast
- C. VPNv4
- D. IPv4 unicast and IPv6 unicast

- E. IPv4 unicast and IPv6 unicast and VPNv4
- F. No address-family is enabled by default.

Correct Answer: F

Section: (none)

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.8/routing/command/reference/rr38bgp.pdf An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

QUESTION 67

Under which two conditions does an enterprise require BGP peering with its upstream Internet provider? (Choose two.)

- A. The enterprise has two exit points that are geographically separated and routing must be controlled so that delay variations are reduced.
- B. The enterprise has one exit point and must announce the IP class it has been assigned from its ISP.
- C. The enterprise has two Layer 2 links to the upstream ISP in the same location and wants to achieve redundancy.
- D. The enterprise needs a simple and efficient way to ensure proper routing for its AP subnets on its single ISP link.
- E. The enterprise is multihomed to two ISPs and has its own IP space AS number.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 68

How can you prevent multihomed customers with connections to two service providers from acting as a transit AS?

- A. Enable BGP synchronization on all the customer routers
- B. Use MED to influence the inbound traffic from the ISPs
- C. Use static routing to the ISPs
- D. Use an AS-path access-list to filter the BGP updates to the ISPs
- E. Use conditional advertisements when sending BGP updates to the ISPs

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 69

Which option describes how an engineer must configure route summarization on a Cisco IOS XE OSPF-enabled network of routers?

- A. area range command under the router OSPF process of the ABR
- B. area range command under the router OSPF process of the router wanting to originate the summary route from
- C. summary-address command under the OSPF process of the ABR
- D. summary-address command under the OSPF process of the router wanting to originate the summary route from
- E. network statement for the aggregate address under the router OSPF process on the ABR
- F. network statement for the aggregate address under the router OSPF process of the router wanting to originate the summary route from

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 70

What are two consequences of having constant link flaps, resulting in the OSPF neighbor adjacencies going up and down repeatedly? (Choose two.)

- A. routes getting into the "Stuck In Active" state
- B. constant flooding of LSAs
- C. OSPF route dampening to occur
- D. many SPF recalculations
- E. routing loops may temporarily be introduced into the network

Correct Answer: BD

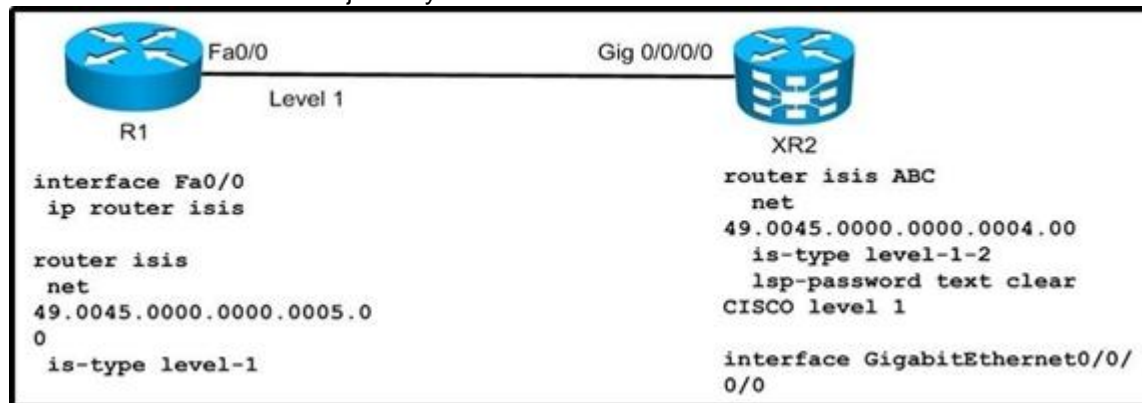
Section: (none)

Explanation

Explanation/Reference:

QUESTION 71

Refer to the exhibit. IS-IS adjacency is not established between XR2 and R1. Which action fixes this issue?



- A. unify IS-IS process IDs on each router
- B. configure on R1 under IS-IS is-type level-1-2
- C. configure on R1 under IS-IS area-password CISCO
- D. configure on XR2 under IS-IS under Gi0/0/0/0 address-family ipv4 unicast

Correct Answer: D

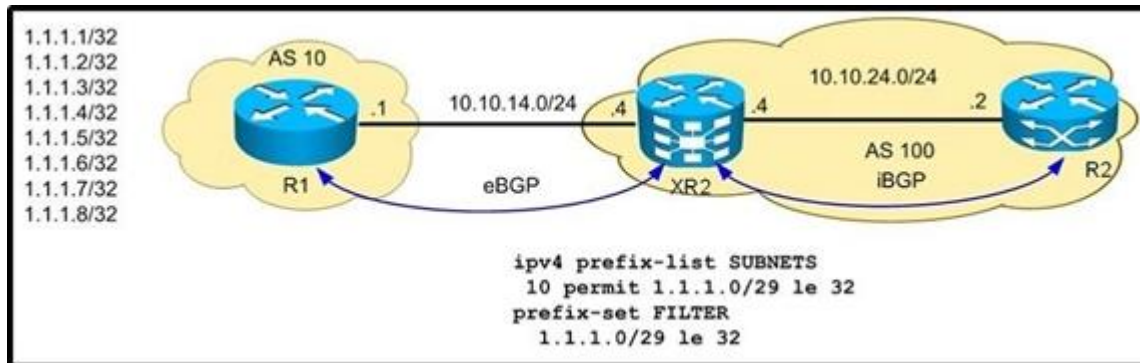
Section: (none)

Explanation

Explanation/Reference:

QUESTION 72

Refer to the exhibit. R1 is advertising subnets 1.1.1.X/32 via eBGP. XR4 must accept only the first seven subnets in its BGP table. Which configuration must occur on XR2 to meet these parameters?



- A. route-policy INBOUND if destination in SUBNETS then pass endif
router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast route-policy INBOUND in
- B. route-policy INBOUND if destination in FILTER then pass endif
router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast route-policy INBOUND in
- C. router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast
route-policy FILTER in
- D. router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast
route-policy SUBNETS in

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 73

When using the Cisco IOS-XR show bgp command to examine the BGP table, the Metric value being displayed is used to represent which BGP attribute?

- A. Weight
- B. Local Preference
- C. MED
- D. Cost Community

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 74

What are two ways to advertise networks into BGP? (Choose two.)

- A. using the neighbor router BGP command
- B. using a route policy in Cisco IOS XR Software or using a route map in Cisco IOS Software or Cisco IOS XE Software
- C. using route redistribution into BGP
- D. using the network router BGP command
- E. enabling an interface to run BGP using the interface router BGP command

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:



QUESTION 75

Refer to the Cisco IOS show command output shown in the exhibit.

```

R1# show ip bgp 10.2.10.1
BGP routing table entry for 10.2.10.1/32, version 32
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    2
  64500 64500 64502
    192.168.103.30 from 192.168.103.30 (10.3.1.1)
      Origin IGP, localpref 100, weight 100, valid, external
  64500 64500 64502
    192.168.134.40 (metric 11) from 192.168.134.40 (10.4.1.1)
      Origin IGP, localpref 100, weight 200, valid, external, best

R1# show ip bgp
BGP table version is 89, local router ID is 11.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*> 10.2.10.1/32    192.168.201.20          200 64500 64500 64500 64502 i
                   192.168.101.10         100 64500 64500 64502 i

```

Which of the following statements is correct?

- A. The best path to 10.2.10.1/32 was chosen because of the AS path.
- B. The best path to 10.2.10.1/32 was chosen because of the weight.
- C. The best path to 10.2.10.1/32 was chosen because of the local preference.
- D. The best path to 10.2.10.1/32 was chosen because of the MED.
- E. The best path to 10.2.10.1/32 was chosen because of the route origin.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 76

On Cisco IOS XR Software, which set of commands is used to enable the gi0/0/0/1 interface for OSPF in area 0?

- A. interface gi0/0/0/0 ip address 10.1.1.1 255.255.255.0
!
router ospf 1

- network 10.1.1.1 0.0.0.0 area 0
- B. interface gi0/0/0/0 ip address 10.1.1.1 255.255.255.0
!
router ospf 1
network 10.1.1.1 255.255.255.255 area 0 C. router ospf 1 area 0
interface GigabitEthernet0/0/0/1
- C. interface gi0/0/0/0 ip address 10.1.1.1 255.255.255.0 ip ospf 1 area 0
- D. router ospf 1 address-family ipv4 unicast
interface GigabitEthernet0/0/0/1 area 0
- E. router ospf 1 address-family ipv4 unicast interface GigabitEthernet0/0/0/1 area 0

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 77

An engineer is working for a service provider who is strategically important to the Internet data traffic. The CTO is promoting a core infrastructure upgrade to connect via optical links network service providers. Which three concerns must be addressed while planning the new BGP features and functionalities? (Choose three.)

- A. BGP suboptimal routing
- B. IPv4 exhaustion
- C. routers and switches hardware capabilities
- D. route dampening of current customers
- E. multihoming policies of new customers
- F. IGP reconvergence time on the inter-AS links
- G. BGP blackholing when acting as transit autonomous system

Correct Answer: ACG

Section: (none)

Explanation

Explanation/Reference:

QUESTION 78

When implementing LDP, what is liberal label retention mode?

- A. To reduce the convergence time, the LSR will retain all the received labels in its LIB even if all the neighbor LSRs go down temporary
- B. The LSR can assign it's own label for each destination network even though it has not been assigned a next hop label from the neighbor LSR
- C. The LSR will store the label received from the downstream LSRs in its LIB even if the downstream LSRs are not the next hop for the destination
- D. The LSR will not perform PHP if it is operating in liberal label retention mode

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

By default, LDP accepts labels (as remote bindings) for all prefixes from all peers. LDP operates in liberal label retention mode, which instructs LDP to keep remote bindings from all peers for a given prefix. For security reasons, or to conserve memory, you can override this behavior by configuring label binding acceptance for set of prefixes from a given peer.

The ability to filter remote bindings for a defined set of prefixes is also referred to as LDP inbound label filtering

QUESTION 79

Which option can a network specialist use to configure connected route redistribution inside VRF "TEST" on Cisco IOS XR and allow only the prefix 10.10.10.0/24?

- A. route-policy ALLOW-CONN if destination in PERMIT_PREFIX then pass
else drop end-policy
prefix-set PERMIT_PREFIX 10.10.10.0/24 end-set router bgp 65000 vrf TEST rd 65000:10000 address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN
- B. route-policy ALLOW-CONN if source in PERMIT_PREFIX then pass else drop end-policy
prefix-set PERMIT_PREFIX 10.10.10.0/24 end-set router bgp 65000 vrf TEST rd 65000:10000 address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN
- C. route-policy ALLOW-CONN if protocol is connected and source in PERMIT_PREFIX then pass else drop end-policy
prefix-set PERMIT_PREFIX 10.10.10.0/24 end-set router bgp 65000 vrf TEST rd 65000:10000 address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN"
- D. route-policy ALLOW-CONN if route-type is local and destination in PERMIT_PREFIX then pass else drop end-policy
prefix-set PERMIT_PREFIX 10.10.10.0/24 end-set router bgp 65000 vrf TEST rd 65000:10000 address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 80

Which three BGP configuration groupings are supported on Cisco IOS XR Software? (Choose three.)

- A. peer-group
- B. af-group
- C. bgp-group
- D. session-group
- E. neighbor-group
- F. as-group

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:

Commands relating to a peer group found in Cisco IOS Release 12.2 have been removed from Cisco IOS XR software. Instead, the af-group, session-group, and neighbor-group configuration commands are added to support the neighbor in Cisco IOS XR software:

-The af-group command is used to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have the same address family configuration are able to use the address family group name for their address family-specific configuration. A neighbor inherits the configuration from an address family group by way of the use command. If a neighbor is configured to use an address family group, the neighbor will (by default) inherit the entire configuration from the address family group. However, a neighbor will not inherit all of the configuration from the address family group if items are explicitly configured for the neighbor. - The session-group command allows you to create a session group from which neighbors can inherit address family-independent configuration. A neighbor inherits the configuration from a session group by way of the use command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the session group's entire configuration. A neighbor does not inherit all the configuration from a session group if a configuration is done directly on that neighbor.

-The neighbor-group command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address family groups. This additional flexibility can create a complete configuration for a neighbor. Once a neighbor group is configured, each neighbor can inherit the configuration through the use command. If a neighbor is configured to use a neighbor group, the neighbor (by default) inherits the neighbor group's entire BGP configuration.

-However, a neighbor will not inherit all of the configuration from the neighbor group if items are explicitly configured for the neighbor. In addition, some part of the neighbor group's configuration could be hidden if a session group or address family group was also being used.

QUESTION 81

You are implementing OSPF as the IGP using a single OSPF area design. The router memory usage for OSPF is too high. Which two methods can lower the OSPF memory usage? (Choose two.)

- A. Enable OSPF synchronization
- B. Implement multi-area OSPF
- C. Enable OSPF sham links to reduce the number of LSAs
- D. Implement route summarization on the ABRs
- E. Enable route leaking between Level 1 and Level 2 areas

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 82

Why can using the ip tcp path-mtu-discovery command improve BGP convergence?

- A. Single packet sizes in TCP sessions are limited.
- B. Smaller MSS sizes may reduce BGP convergence times.
- C. BGP is allowed to use a larger TCP window size.
- D. BGP is enabled to fragment its large update packets.
- E. The BGP memory requirements on routers are reduced.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 83

What is a key benefit of implementing the bgp dmzlink-bw command?

- A. establishes a secure BGP link between eBGP and iBGP peers
- B. enables the unequal cost path load balancing feature for external BGP links
- C. allows BGP to use link bandwidth as a metric when computing best route selection
- D. increases BGP security by dynamically allocating community strings to eBGP routes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 84

What is a key benefit of BGP dynamic update peer groups?

- A. Routing updates to the same destination are grouped to increase BGP efficiency.
- B. Newly configured BGP neighbors have peer group template configurations dynamically applied.
- C. Dynamic update groups use iBGP neighbor information to automatically calculate route reflector cluster configurations.
- D. Neighbors in a peer group are no longer required to share the same outbound routing policies.
- E. BGP configurations are automatically optimized by routers which dynamically create BGP peer groups.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



QUESTION 85

How does the extended community cost feature influence the BGP best path selection?

- A. alters the BGP AS exit path selection by adding the link cost to the local preference
- B. acts as a best path "tie breaker" when multiple IGP equal cost paths occur
- C. reflects the bandwidth of links entering the local AS from eBGP neighbors (in the MED attribute)
- D. selects the BGP route with the highest attached extended community cost value
- E. inserts the cost attribute after the MED attribute comparison, forcing best path route selection if all other preferred route selection criteria are equal

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 86

What is the maximum number of parallel routes that can be injected into the routing table when using the iBGP multipath load sharing feature?

- A. iBGP multipath can support up to 4 parallel paths.
- B. iBGP multipath can support up to 6 parallel paths.
- C. None, multipath is supported on eBGP only when used in conjunction with the update-source command.
- D. The number of parallel paths injected into the routing table will depend on the IGP currently in use within the iBGP domain.
- E. The maximum number of parallel paths injected into the routing table is specified by BGP MED and ranges from 2 to 256 parallel paths.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 87

What is a key benefit of implementing the bgp dmzlink-bw command?

- A. establishes a secure BGP link between eBGP and iBGP peers
- B. enables the unequal cost path load balancing feature for external BGP links
- C. allows BGP to use link bandwidth as a metric when computing best route selection
- D. increases BGP security by dynamically allocating community strings to eBGP routes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 88

In a Transit AS, how do the internal routers within the Transit AS forward packets destined for the external networks using a scalable solution?

- A. using the default route
- B. using the IGP routes where the external networks are redistributed into the IGP by the edge routers
- C. using the EBGP routes where the external networks are redistributed into the IBGP by the edge routers

D. using the IBGP routes, then using recursive lookup based on IGP information to resolve the BGP next-hop

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 89

Which two statements are true about an EBGP session or an IBGP session? (Choose two.)

- A. IBGP uses AS-Path to detect routing information loops within the AS.
- B. EBGP routes have a default Admin Distance of 20 and IBGP routes have a default Admin Distance of 200.
- C. No BGP attributes are changed in EBGP updates except for the next-hop attribute if next-hop-self is configured.
- D. Routes learned from an EBGP peer not advertised to another EBGP peer to prevent routing information loops.
- E. IBGP uses split horizon to prevent routing information loops; routes learned from an IBGP peer are not advertised to another IBGP peer.

Correct Answer: BE

Section: (none)

Explanation



Explanation/Reference:

QUESTION 90

Which AS path access list is used by a multihomed customer to only announce their own address space to their service providers to prevent the multihomed customer from becoming a transit AS?

- A. ip as-path access-list permit.*
- B. ip as-path access-list permit^\$
- C. ip as-path access-list permit _\$
- D. ip as-path access-list permit _^
- E. ip as-path access-list permit local-as ip as-path access-list deny *
- F. ip as-path access-list deny.* ip as-path access-list permit *

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



| CHAR | USAGE |
|------|---|
| ^ | Start of string |
| \$ | End of string |
| [] | Range of characters |
| - | Used to specify range (i.e. [0-9]) |
| () | Logical grouping |
| . | Any single character |
| * | Zero or more instances |
| + | One or more instance |
| ? | Zero or one instance |
| , | Comma, open or close brace, open or close parentheses, start or end of string, or space |

Some commonly used regular expressions include:

| Expression | Meaning |
|------------|----------------------------------|
| .* | Anything |
| ^\$ | Locally originated routes |
| ^100_ | Learned from AS 100 |
| _100\$ | Originated in AS 100 |
| _100_ | Any instance of AS 100 |
| ^[0-9]+\$ | Directly connected ASes |

QUESTION 91

Refer to the route policies exhibit.

```
route-policy one
end-policy
!
route-policy two
pass
end-policy
!
route-policy three
drop
end-policy
!
route-policy four
set weight 100
end-policy
!
route-policy five
pass
drop
pass
end-policy
!
```



Which five route policies will cause the routes to be dropped or passed? (Choose five)

- A. route-policy one will cause the routes to be dropped.
- B. route-policy two will cause the routes to be dropped.
- C. route-policy three will cause the routes to be dropped.
- D. route-policy four will cause the routes to be dropped.
- E. route-policy five will cause the routes to be dropped.
- F. route-policy one will cause the routes to be passed.
- G. route-policy two will cause the routes to be passed.
- H. route-policy three will cause the routes to be passed.
- I. route-policy four will cause the routes to be passed.

J. route-policy five will cause the routes to be passed.

Correct Answer: ACEGI

Section: (none)

Explanation

Explanation/Reference:

QUESTION 92

What are two purposes of the BGP scan-time command? (Choose two.)

- A. to tune the BGP process which walks the BGP table and confirms the reachability of next hops
- B. to allow faster detection of downed BGP peers
- C. to improve BGP convergence time
- D. to tune the BGP update interval
- E. to decrease the effects of unstable routes by increasing the route suppression time

Correct Answer: AC

Section: (none)

Explanation



Explanation/Reference:

<http://www.networkers-online.com/blog/2008/12/bgp-performance-tunning-convergence-stabilityscalabilityand-nsf-part-2/>

Background BGP scanner

It is responsible for BGP housekeeping by scanning both the BGP RIB and the IP RIB and cleaning and sorting things out.

BGP monitors the next hop of the installed routes to verify next-hop reachability and to select, install, and validate the BGP best path. By default, the BGP scanner polls the RIB for this information every 60 seconds.

During the 60 second time period between scan cycles, IGP instabilities or other network failures can cause temporarily black holes and routing loops.

NOTE With Cisco IOS the default timer is 60 seconds for the IPv4 address family and 15 seconds for the VPNv4 address family in order to optimize the VPNs routing table convergence.

This timer can be controlled via the following command:

```
Router(config-router)#bgp scan-time <5-60>
```

QUESTION 93

Which two statements about a transit AS are correct? (Choose two.)

- A. A transit AS has eBGP connection(s) to only one external AS.
- B. Routes between ASs are always exchanged via eBGP.
- C. A transit AS uses an IGP like OSPF or ISIS to propagate the external networks within the transit AS.
- D. Core routers within a transit AS normally use default routing to reach the external networks.
- E. iBGP sessions can be established between non directly connected routers.

Correct Answer: BE


Section: (none)

Explanation

Explanation/Reference:

QUESTION 94

Refer to the exhibit. Which configuration is correct for XR4 for an e-BGP session with R1?



- A. router bgp 100 neighbor 10.10.14.1 remote-as 100
- B. router bgp 100 neighbor 10.10.14.1 remote-as 10 address-family ipv4 unicast
- C. router bgp 100 address-family ipv4 unicast neighbor 10.10.14.1 remote-as 100 address-family ipv4 unicast
- D. router bgp 100 address-family ipv4 unicast neighbor 10.10.14.1 remote-as 100

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference: