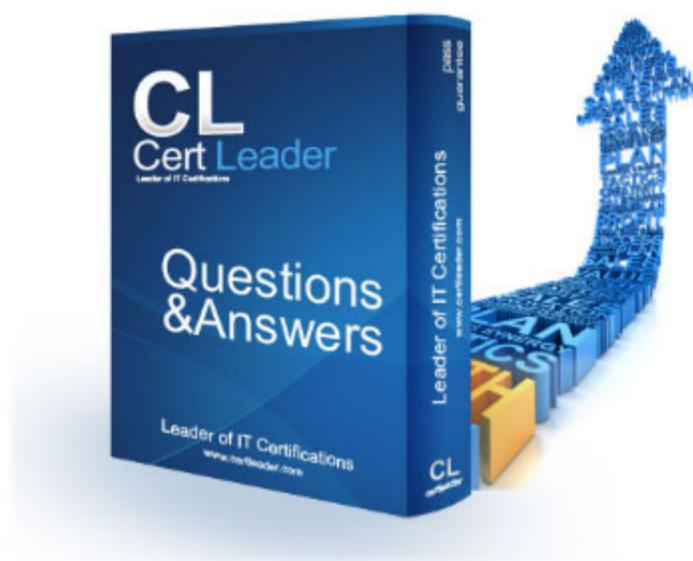


642-883 Dumps

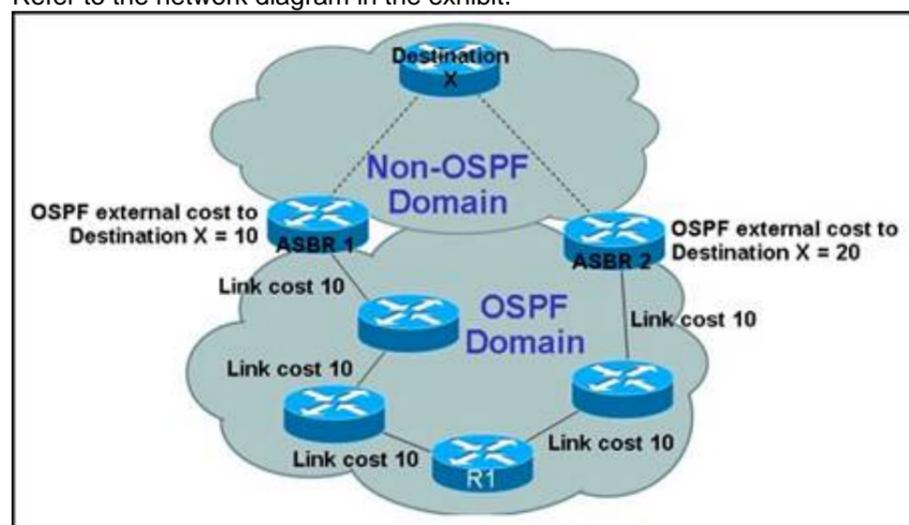
Deploying Cisco Service Provider Network Routing (SPROUTE)

<https://www.certleader.com/642-883-dumps.html>



NEW QUESTION 1

Refer to the network diagram in the exhibit.



If both ASBRs are advertising the external Destination X network as OSPF E2 route, what is the best path for the R1 router to reach Network X?

- A. R1 will use the path via ASBR 2 as the best path.
- B. R1 will use the path via ASBR 1 as the best path.
- C. R1 will load balance between two equal cost paths via ASBR 1 and ASBR 2.
- D. R1 will see two equal costs and will choose the path through the ASBR with the lower OSPF router ID.

Answer: B

NEW QUESTION 2

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

Instructions

Enter the proper CLI commands and analyze 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 command options are supported or required for this simulation.

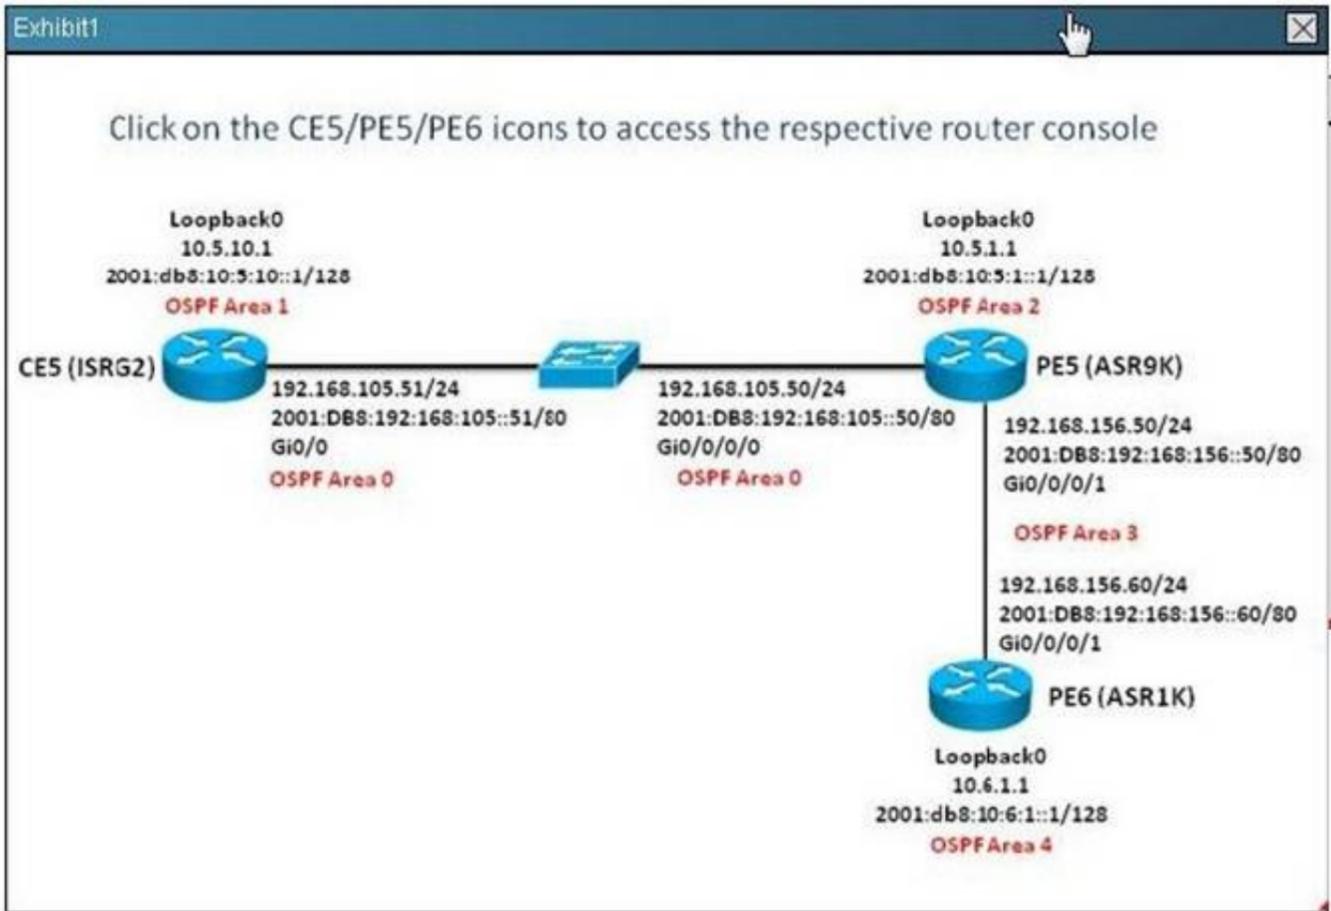
For example, the show running-config and the ping commands are **NOT** supported in this simulation.

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

Scenario

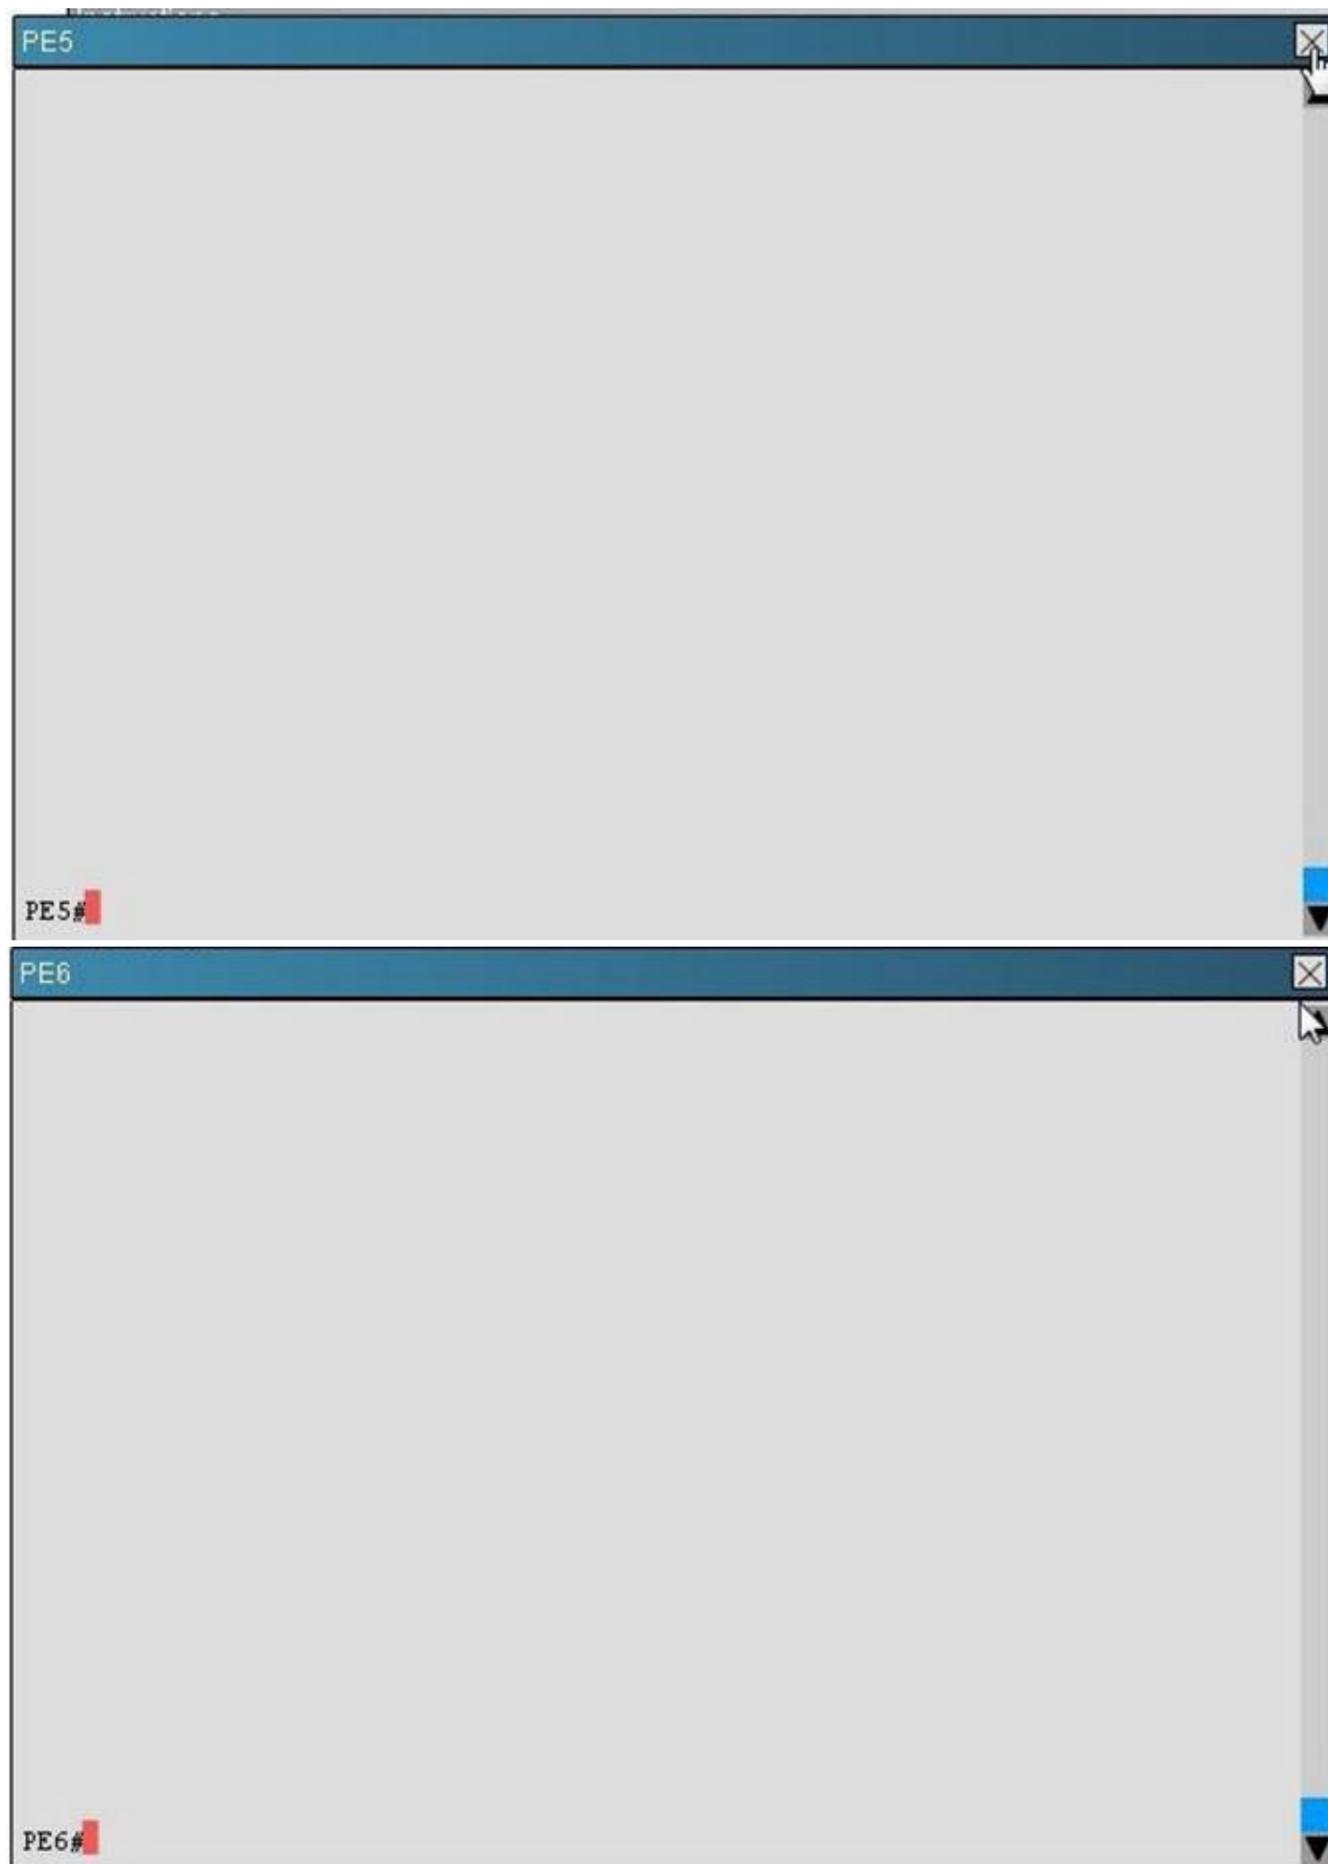
Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

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.



CE5

CE5#



- 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

Answer: A

Explanation: # show ip ospf neighbor
show ipv6 ospf neighbor

NEW QUESTION 3

Which configuration allows incoming traffic on the FastEthernet 0/1 interface of a router to be policy-routed via the next-hop 192.168.10.1, when it is reachable and otherwise be routed normally?

- A. interface Fa0/1 ip policy route-map nh-reachable route-map nh-reachable permit 10 set ip next-hop verify-availability route-map nh-reachable permit 20 set ip next-hop 192.168.10.1
- B. interface Fa0/1 ip policy route-map nh-reachable route-map nh-reachable permit 10 set ip next-hop verify-availability 192.168.10.1 route-map nh-reachable permit 20
- C. interface Fa0/1 ip policy route-map nh-reachable route-map nh-reachable permit 10 set ip next-hop 192.168.10.1 track route-map nh-reachable permit 20
- D. interface Fa0/1 ip policy route-map nh-reachable route-map nh-reachable permit 10 set ip next-hop 192.168.10.1 route-map nh-reachable permit 20 set default interface

Answer: B

NEW QUESTION 4

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.

Answer: B

Explanation:

```
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
```

C:\Documents and Settings\user-nwz\Desktop\1.JPG

NEW QUESTION 5

DRAG DROP

Drag the BGP attributes on the left to the boxes on the right in the correct order. The attribute that is used first during the BGP route selection process should be the top box on the right, and the attribute that is used last should be the bottom box on the right.

| | |
|------------------|--------|
| MED | Target |
| Weight | Target |
| AS-Path Length | Target |
| Local Preference | Target |

Answer:

Explanation: Weight Local-Preference
As-Path MED

NEW QUESTION 6

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)#

Answer: A

NEW 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

Answer: D

NEW QUESTION 8

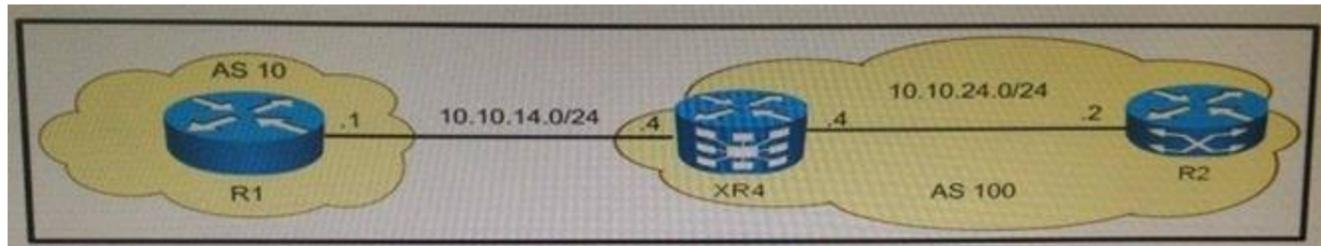
Which option can a network specialist use to configure connected route redistribution inside VRF "TEST" on Cisco IOS XR and allow only the prefix 10.10.10.0/24?

- A. route-policy ALLOW-CONN if destination in PERMIT_PREFIX then passsedepend-policy prefix-set PERMIT_PREFIX 10.10.10.0/24 end-setrouter 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 passsedepend-policy prefix-set PERMIT_PREFIX 10.10.10.0/24 end-setrouter 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 passsedepend-policy prefix-set PERMIT_PREFIX 10.10.10.0/24 end-setrouter 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 passsedepend-policy prefix-set PERMIT_PREFIX 10.10.10.0/24 end-setrouter bgp 65000 vrf TEST rd 65000:10000 address-family ipv4 unicast redistribute connected route-policy ALLOW-CONN

Answer: A

NEW QUESTION 9

Refer to the exhibit.



Which configuration is correct for RX4 for an e-BGP session with R1?

- A. router bgp 100 address-family ipv4 unicast! neighbor 10.10.14.1 remote-as 10 address-family ipv4 unicast route-policy POLICY in
- B. router bgp 100 address-family ipv4 unicast! neighbor 10.10.14.1 remote-as 100
- C. router bgp 100! neighbor 10.10.14.1 remote-as 100
- D. router bgp 100! neighbor 10.10.14.1 remote-as 10 address-family ipv4 unicast

Answer: A

NEW QUESTION 10

Which three valid session commands are supported by peer session templates? (Choose three.)

- A. local-as
- B. as-override
- C. inherit peer-session
- D. inherit peer-policy
- E. disable-connected-check
- F. route-reflector-client

Answer: ACE

NEW QUESTION 10

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

Answer: D

Explanation: 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

NEW QUESTION 11

An engineer is working on routers within AS 100. Which regular expression can be used in an AS path access list to match locally originated routes from AS 100?

- A. ^100\$
- B. ^\$
- C. _100\$
- D. ^[0-9]*\$

Answer: B

NEW QUESTION 16

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-prefif community matches lowpri-communities then set local-preference 100endifend-policy
- B. route-policy ch-prefif community eq lowpri-communities then set local-preference 100endifend-policy
- C. route-policy ch-prefif community matches-any lowpri-communities then set local-preference 100endifend-policy
- D. route-policy ch-prefif community neq lowpri-communities then set local-preference 100endifend-policy

Answer: C

NEW QUESTION 18

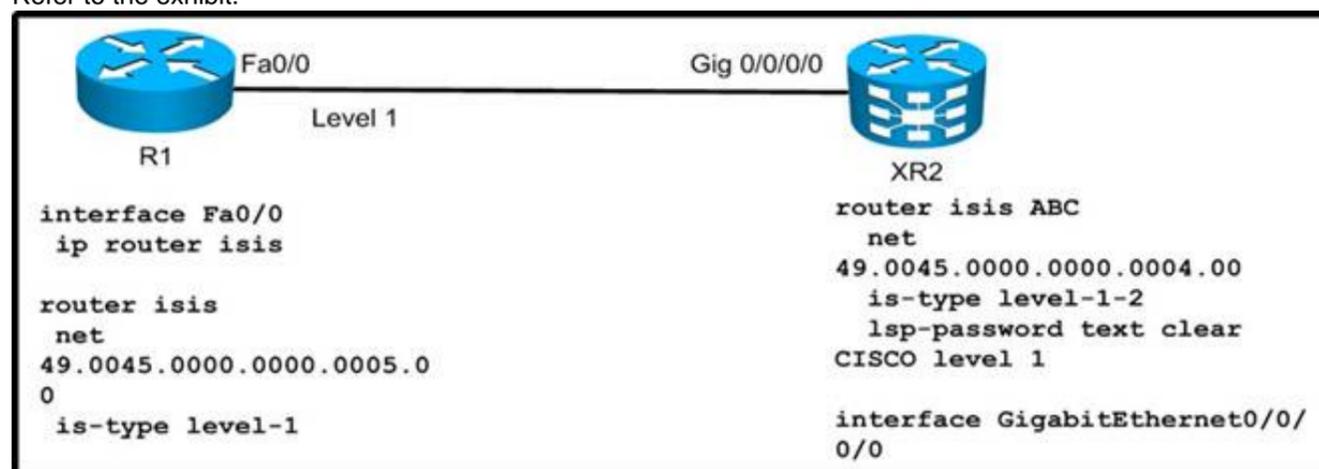
Which option is used by BGP for basic authentication between neighbors?

- A. TCP option 19 and MD5 hash
- B. TCP option 19 and SHA hash
- C. UDP option 19 and MD5 hash
- D. UDP option 19 and SHA hash

Answer: A

NEW QUESTION 21

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

Answer: D

NEW QUESTION 22

In a Cisco IOS XR OSPF NSF operation, which option is the result of a segment without NSF-capable peers?

- A. NSF is disabled on all segmentson any linecard that has a non-NSF-capable neighbor.
- B. NSF capabilities for that segment are disabled.
- C. NSF operates in unidirectional mode for that segment.
- D. NSF is disabled globally on the router.

Answer: B

NEW QUESTION 23

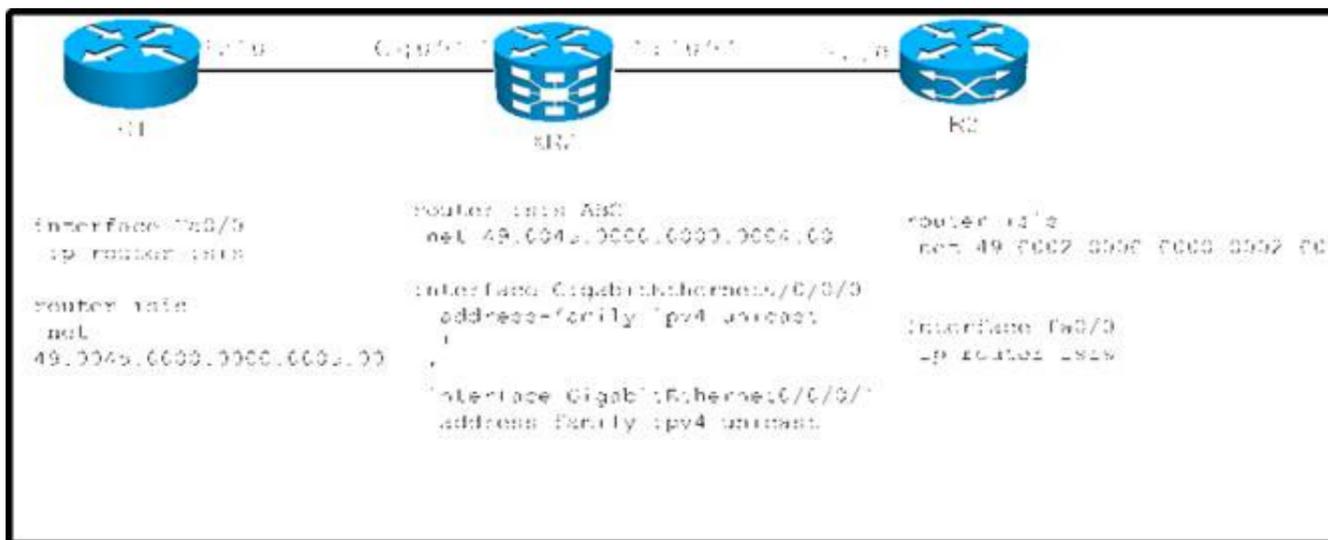
Which configuration implements the most secured OSPF authentication mode on Cisco IOS XE?

- A. interface Gi0/0/0/ip ospf authentication message-digest ip ospf message-digest-key 1 md5 cisco
- B. router ospf 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 ip ospf authentication ip ospf authentication-key cisco

Answer: A

NEW QUESTION 24

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 1919 ABCinterface GigabitEthernet0/0/0/1 circuit-type level-2-only
- B. Router 1919 ABCinterface GigabitEthernet0/0/0/0 circuit-type level-1-2
- C. Router 1919 ABCinterface GigabitEthernet0/0/0/1 circuit-type level-1
- D. Router 1919 ABCinterface GigabitEthernet0/0/0/0 circuit-type level-1

Answer: D

NEW QUESTION 29

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

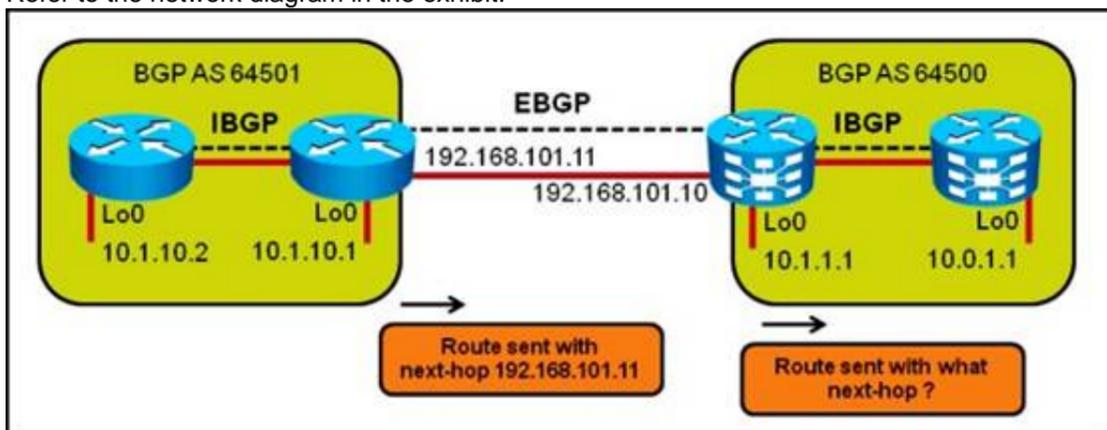
- A. NSF
- B. SSO
- C. NSR
- D. BFD
- E. SDR

Answer: D

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

NEW QUESTION 34

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

Answer: A

NEW QUESTION 36

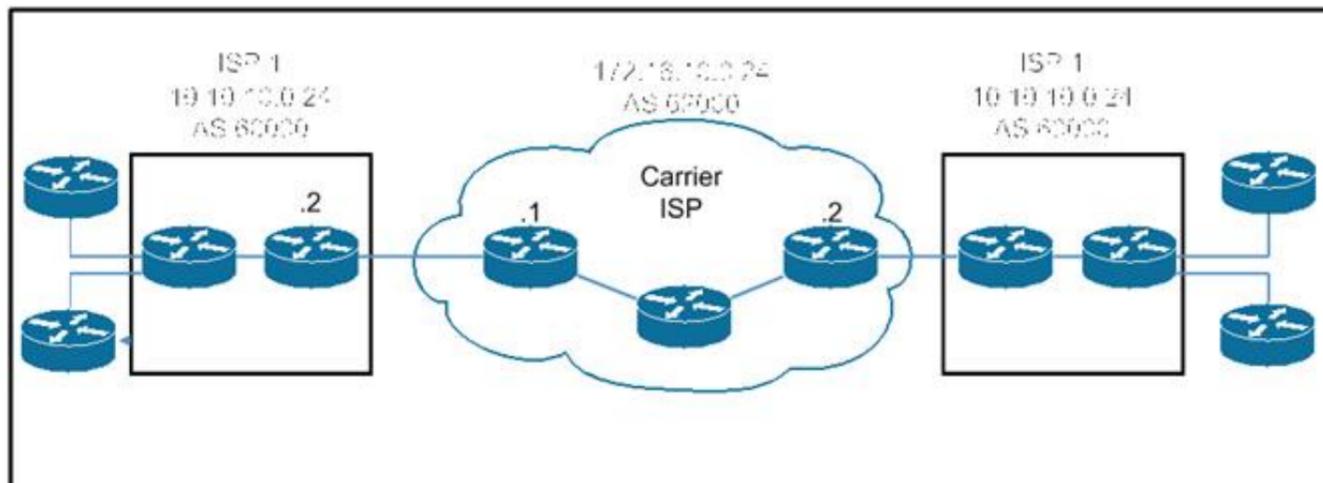
A Cisco IOS XR router is a member in OSPF 1 and EIGRP 100 domains, and needs to redistribute OSPF learned routes into EIGRP. Which configuration achieves this goal?

- A. router eigrp 100address-family ipv4 redistribute ospf 1
- B. router eigrp 100redistribute ospf 1 route-policy OS_INT0_EIGroute-policy OS_INT0_EIG set eigrp-metric 100 10 255 1 155
- C. router eigrp 100address-family ipv4redistribute ospf 1 route-policy OS_INT0_EIGroute-policy OS_INT0_EIG set eigrp-metric 100 10 255 1 155
- D. router eigrp 100default-metric 100 1 255 1 1500redistribute ospf 1

Answer: C

NEW QUESTION 40

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`

Answer: AD

NEW QUESTION 42

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.

Answer: C

NEW QUESTION 43

When redistributing routes into OSPF, which kind of routes will be redistributed by default on Cisco IOS XR Software but will not be automatically redistributed by default on Cisco IOS Software and Cisco IOS XE Software?

- A. Type 1 external routes
- B. Type 2 external routes
- C. subnetted routes
- D. /32 host routes
- E. tagged static routes

Answer: C

Explanation:

Table 12 Redistribution Configurations for OSPF in Cisco IOS and Cisco IOS XR

| Cisco IOS | Cisco IOS XR |
|---|--|
| <pre>router ospf 1 redistribute connected subnets redistribute static subnets route-map map</pre> | <pre>router ospf lab redistribute connected redistribute static policy</pre> |

You no longer need the **subnets** keyword in Cisco IOS XR, and you declare an RPL policy rather than a route map. If you want to redistribute only classful routes into OSPF in Cisco IOS XR, you can use the **redistribute** command with the **classful** keyword. The use of policies rather than route maps is covered in the "Border Gateway Protocol (BGP)" section and also in the existing Cisco IOS XR documentation on RPL. The mainpiece of information to get from this discussion is that redistribution into OSPF on Cisco IOS XR is a policy attach point.

C:\Documents and Settings\user-nwz\Desktop\1.JPG

NEW QUESTION 46

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.
- B. They use a default route to exit the area.
- C. Summarization should be configured on the Level 2 routers, which injects the Level 2 routes into Level 1.
- D. IS-IS supports "route leaking" in which selected Level 2 routes can be advertised by a Level 1/Level 2 router into Level 1.
- E. The IS-IS backbone is a contiguous collection of Level 1 capable routers, each of which can be in a different area.
- F. 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.
- G. Cisco IOS XR Software supports multitopology for IPv6 IS-IS unless single topology is explicitly configured in IPv6 address-family configuration mode.

Answer: ACEF

Explanation: 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 single-topology 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 IS-IS 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.

NEW QUESTION 47

A customer is concerned about rerouting attacks on the IS-IS domain spanning the company. Which keychain management configuration is valid to authenticate IS-IS link-state information and provide a first response to this type of attack for Cisco IOS XR instances?

- A. `configure key chain isis-keys key 1 key-string myP@ssw0rd cryptographic-algorithm MD5 send-lifetime 18:05:00 may 31 2014 infinite accept-lifetime 18:05:00 may 31 2014 infinite router isis 1 isp-password keychain isis-keys`
- B. `key chain cisco-xr key 10 key-string myP@ssw0rd! interface GigabitEthernet 3/1/1 ip address ip_address subnet_mask ip router isis isis authentication mode md5 level-1 isis authentication key-chain cisco-xr level-1`
- C. `configure key chain isis-keys! key 8 cryptographic-algorithm MD5 send-lifetime 18:05:00 may 31 2014 infinite accept-lifetime 18:05:00 may 31 2014 infinite router isis 1 isp-password keychain isis-keys`
- D. `configure key chain isis-keys key-string myP@ssw0rd send-lifetime 18:05:00 may 31 2014 infinite accept-lifetime 18:05:00 may 31 2014 infinite router isis 1 isp-password keychain isis-keys`

Answer: A

NEW QUESTION 48

A network engineer implements IS-IS authentication in a router. When the engineer implements the same password on the router's peer, the output below is displayed with the debug IS-IS authentication information:

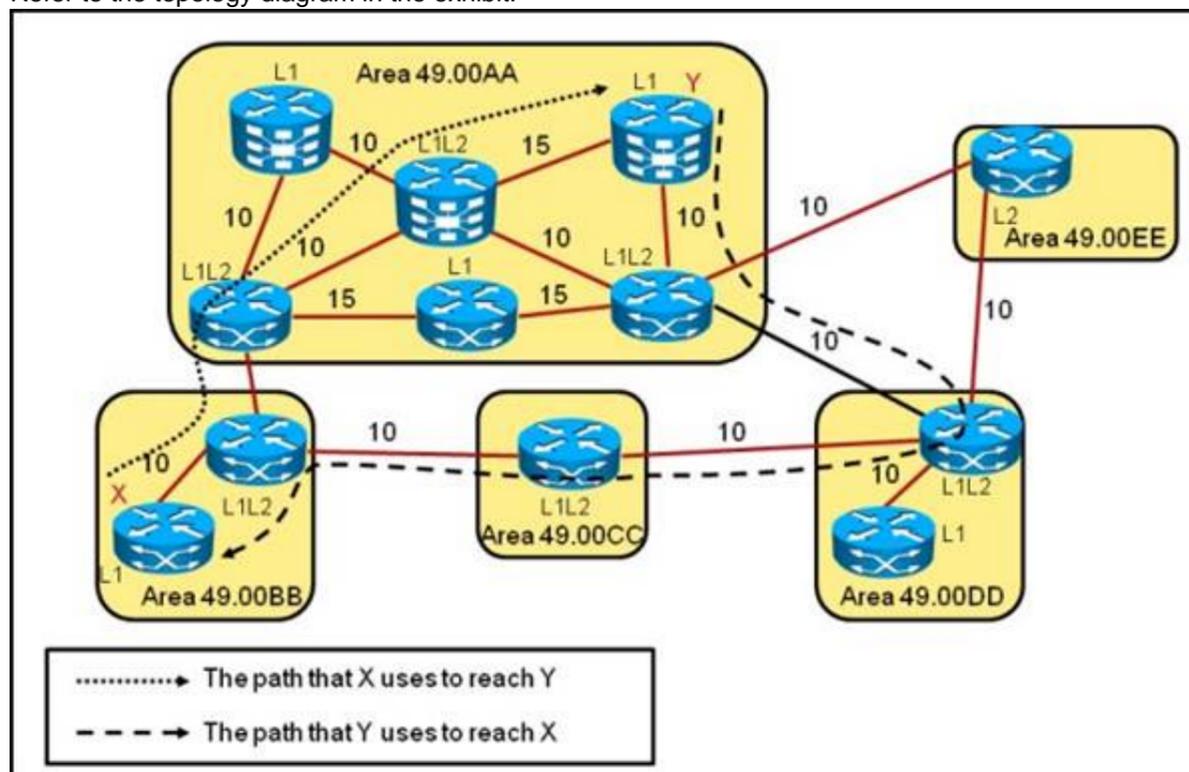
* Nov 6 13:23:46:967: ISIS-Authinfo: No auth TLV found in received packet Which option is one reason for this output?

- A. Authentication was not configured on the router that sent the update.
- B. The length of the password is not the same.
- C. There is an error with the NET address in one of the routers.
- D. The authentication is set the same way, but the password does not match.

Answer: A

NEW QUESTION 51

Refer to the topology diagram in the exhibit.



Which IS-IS feature could be implemented so that the return path for the packets from router Y in area 49.00AA to router X in area 49.00BB will use the more optimal path?

- A. Enable route leaking to pass Level 2 information into the Level 1 routers.
- B. Change the area 49.00AA type from a stub area to a regular area.
- C. Change the IS-IS administrative distance on router Y in area 49.00AA.
- D. Change the IS-IS metric type from narrow to wide on all IS-IS routers.

Answer: A

Explanation: <http://www.cisco.com/application/pdf/paws/13796/route-leak.pdf>

The ISIS routing protocol allows for a twolevel hierarchy of routing information. There can be multiple Level 1 areas interconnected by a contiguous Level 2 backbone. A router can belong to Level 1, Level 2, or both. The Level 1 linkstate database contains information about that area only. The Level 2 linkstate database contains information about that level as well as each of the Level 1 areas. An L1/L2 router contains both Level 1 and Level 2 databases. It advertises information about the L1 area to which it belongs into L2. Each L1 area is essentially a stub area. Packets destined for an address that is outside of the L1 area are routed to the closest L1/L2 router to be forwarded on to the destination area. Routing to the closest L1/L2 router can lead to suboptimal routing when the shortest path to the destination is through a different L1/L2 router.

Route leaking helps reduce suboptimal routing by providing a mechanism for leaking, or redistributing, L2 information into L1 areas. By having more detail about inter area routes, an L1 router is able to make a better choice with regard to which L1/L2 router to forward the packet.

NEW QUESTION 53

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.

Answer: A

NEW QUESTION 56

Refer to the exhibit.

```

Router A
router isis L2_isis
 net 47.0000.0000.0000.0001.00
 metric-style wide
!
interface Ethernet0/0
 ip address 172.16.0.1 255.255.255.252
 ip router isis L2_isis

Router B
router isis L1_isis
 net 47.0001.0000.0000.0001.00
 is-type level-1-2
 metric-style wide
!
interface Ethernet0/1
 ip address 172.16.0.2 255.255.255.252
 ip router isis L1_isis
!
interface Ethernet0/2
 ip address 172.16.0.10 255.255.255.252
 ip router isis L1_isis

Router C
router isis L1_isis
 net 47.0002.0000.0000.0001.00
 is-type level-1
 metric-style wide
!
interface Ethernet0/2
 ip address 172.16.0.9 255.255.255.252
 ip router isis L1_isis
    
```

A network engineer is asked to verify a multiarea IP IS-IS configuration before implementing it. Which statement is true?

- A. IS-type Level 1 must be configured on router B.
- B. Area addresses should be common.
- C. Wide metrics should be used only for router A.
- D. The system ID should be unique.

Answer: B

NEW QUESTION 57

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.

Answer: BDE

| | |
|---|--|
| <p>Explanation: <code>is-type {level-1 level-1-2 level-2-only}</code></p> <p>Example: <code>RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only</code></p> | <p>(Optional) Configures the system type (area or backbone router).</p> <ul style="list-style-type: none"> 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. |
|---|--|

C:\Documents and Settings\user-nwz\Desktop\1.JPG

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.

| | |
|---|--|
| <p><code>circuit-type {level-1 level-2-only level-1-2}</code></p> <p>Example: <code>RP/0/RP0/CPU0:router(config-isis-if)# circuit-type level-1-2</code></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. |
|---|--|

C:\Documents and Settings\user-nwz\Desktop\1.JPG

NEW QUESTION 62

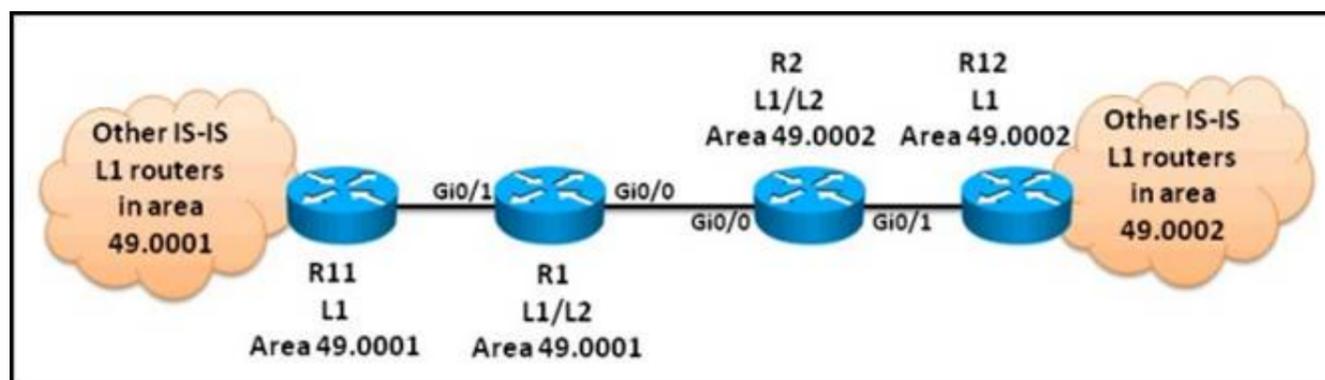
What are three common problems that can cause a BGP neighbor state to toggle between the idle state and the active state? (Choose three.)

- 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

Answer: CEF

NEW 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.

Answer: CD

NEW QUESTION 65

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

Answer: C

NEW QUESTION 66

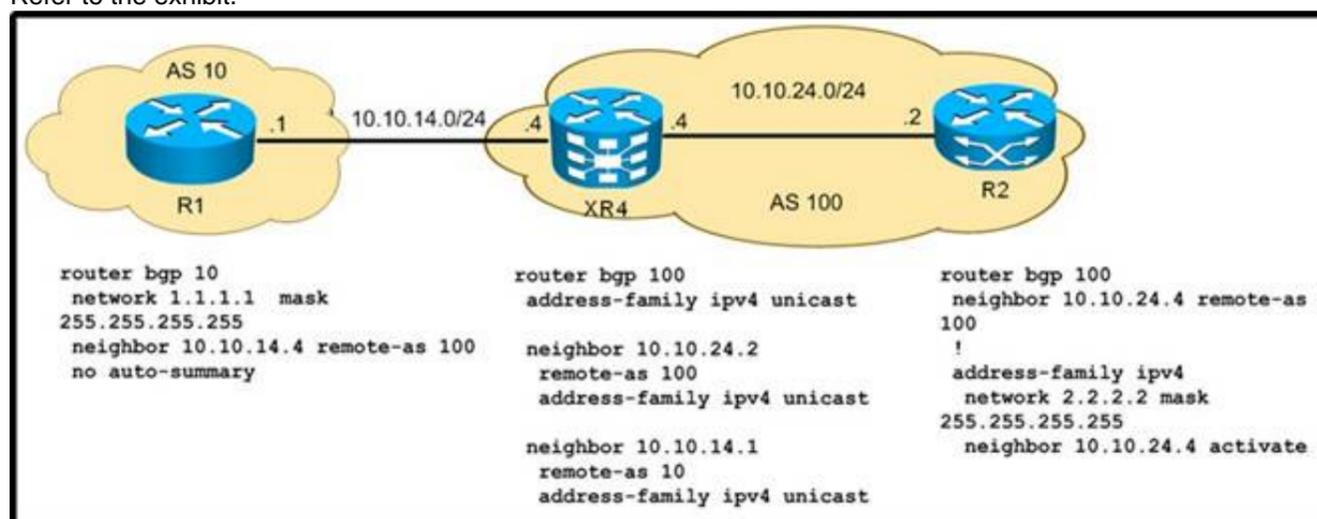
A network engineer wants to set the local preference to 100 for the community set list lowpri-communities. Which option can the engineer use to implement this through RPL?

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

Answer: B

NEW QUESTION 68

Refer to the exhibit.



R2 does not have R1 1.1.1.1/32 subnet in its BGP 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.

Answer: C

NEW QUESTION 72

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.
- C. IPv6 unicast routing has not been enabled globally on the Cisco IOS XE router.
- D. The gi0/0 and gi0/1 interfaces are not participating in the router isis 1 routing instance.
- E. Multitopology IS-IS must be enabled to support both IPv4 and IPv6.
- F. Another router isis instance must be enabled to support IPv6.

Answer: CD

Explanation: 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

NEW QUESTION 73

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.

Answer: D

NEW QUESTION 77

DRAG DROP

Drag the BGP attributes on the left to match the correct description on the right.

| | |
|------------------|--|
| Weight | Only used within an AS and has a default value of 100. |
| MED | Can use the prepending feature to influence inbound traffic flow. |
| Local Preference | The lower of this value is the more desirable value. Used to influence the incoming traffic flow from the neighbor autonomous systems. |
| AS-Path | Proprietary to Cisco and not sent to any BGP peers. It is local to the router only. |

Answer:

Explanation: Only used within an AS and has a default value of 100 – Local Preference Can use the prepending feature to influence inbound traffic flow – AS_Path

The lower of this value is the more desirable value .Used to influence the incoming traffic flow from the neighbor autonomous systems. -- MED
Proprietary to cisco and no sent to any BGP peers. Its local to router only -Weight

NEW QUESTION 81

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

Answer: A

NEW QUESTION 84

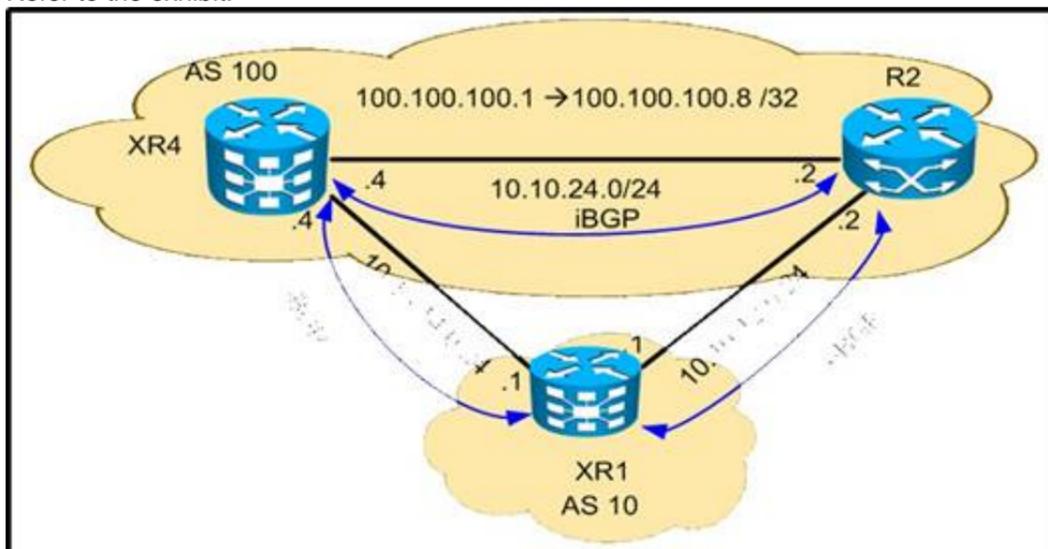
Which two OSPF network scenarios require OSPF virtual link configuration? (Choose two.)

- A. to connect an OSPF non-backbone area to area 0 through another non-backbone area
- B. to connect an NSSA area to an external routing domain
- C. to connect two parts of a partitioned backbone area through a non-backbone area
- D. to enable route leaking from Level 2 into Level 1
- E. to enable route leaking from Level 1 into Level 2
- F. to enable OSPF traffic engineering

Answer: AC

NEW QUESTION 86

Refer to the exhibit.



XR1 must have XR4 as the primary exit point for only the first three subnets received from AS 100 and may not impact the rest of the subnets. Which configuration on XR1 is correct?

- A. prefix-set PREFER 100.100.100.0/30 le 32 route-policy LPif destination in PREFER then set local-preference 200 endif
- B. route-policy LPif destination in (100.100.100.0/30 le 32) then set local-preference 200 else pass endif
- C. prefix-set PREFER 100.100.100.0/29 le 32 route-policy LPif destination in PREFER then set local-preference 200 else pass endif
- D. route-policy LPif destination in (100.100.100.0/29 le 32) then set local-preference 200 else set local-preference 100 endif

Answer: B

NEW QUESTION 90

A service provider is about to purchase more IPv4 address space. Which organization can facilitate this purchase?

- A. APNIC
- B. IETF
- C. USANOG
- D. IEEE
- E. NANOG

Answer: A

NEW QUESTION 95

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.

Answer: BE

NEW QUESTION 99

Cisco IOS XE has OSPF area authentication configured for MD5 and interface authentication configured for clear text. Which configuration takes precedence?

- A. null
- B. area
- C. MD5
- D. interface

Answer: D

NEW QUESTION 102

DRAG DROP

Drag and drop the terms on the left onto the appropriate definition on the right.

| | |
|-------------------|---|
| NSP | Is a network that solely purchases transit from other networks to reach the Internet |
| Tier 1 ISP | peers with some networks, but purchases IP transit or pays settlements to reach some portion of the Internet |
| Tier 2 ISP | peers with other ISPs of the same Tier level to form the backbone of the Internet |
| Tier 3 ISP | provides or sells access and bandwidth into its backbone infrastructure and network access points |
| NAP | interconnects network service providers |

Answer:

Explanation: Is a network that solely purchases transit from other networks to reach the Internet – Tier 3 ISP
Peers with some networks, but purchases IP transit or pays settlements to reach some portion of the Internet– Tier 2 ISP
Peers with other ISPs of the same Tier level to form the backbone of the Internet – Tier 1 ISP
Provides or sells access and bandwidth into its backbone infrastructure and network access points – NSP
Interconnects network service providers – NAP

NEW QUESTION 107

Refer to the IS-IS configuration exhibit.

```
router isis 1
net 49.0001.0100.0100.1001.00
address-family ipv4 unicast
!
address-family ipv6 unicast
single-topology
!
interface gi0/0/0/0
address-family ipv4 unicast
!
address-family ipv6 unicast
```

This is the typical IS-IS configuration of the routers in an AS using IS-IS as the IGP. This AS is in the transition phase of integrating IPv6 into the network. During this transition phase, some of the routers within the AS might be running IPv4 only, some might be running IPv6 only, and others might be running both IPv4 and IPv6. To avoid any black holes for the IPv6 traffic, which configuration change can be made?

- A. Disable IS-IS adjacency checks.
- B. Enable IPv6 adjacency over IPv4 IS-IS peering.
- C. Enable multi-topology IS-IS.
- D. Disable the IPv4 unicast address-family.
- E. Enable IS-IS wide metric to support the single-topology mode.

Answer: C

NEW QUESTION 109

Refer to the exhibit.

```

Router 1
router ospf 1
!
interface Ethernet0/0
 ip address 172.10.23.1 255.255.255.0
 ip ospf 1 area 0
 bfd interval 50 min_rx 50 multiplier 3

Router2
router ospf 1
!
interface Ethernet0/0
 ip address 172.10.23.2 255.255.255.0
 ip ospf 1 area 0
 bfd interval 50 min_rx 50 multiplier 3

```

Which Cisco IOS XE command should be added to this configuration for the BFD protocol to operate successfully?

- A. Router1(config-if)#bfd echoRouter2(config-if)#bfd echo
- B. Router1(config-ospf-ar-if)#bfd fast-detectRouter2(config-ospf-ar-if)#bfd fast-detect
- C. Router1(config-router)#bfd all-interfacesRouter2(config-router)#bfd all-interfaces
- D. Router1(config- bfd)#echo-enableRouter2(config- bfd)#echo-enable

Answer: C

NEW QUESTION 113

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

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

Answer: D

NEW QUESTION 114

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 continue
- D. Additionally, if the same route also contains the 22:22 community, the router also adds another 77:77 community to the same route.
- E. If a route only has the 22:22 community, then no community will be added by the router.
- F. If a route only has the 11:11 or 44:44 community, then no community will be added by the router.

Answer: B

NEW QUESTION 116

Which option describes a function of an OSPF virtual link?

- A. Create redundant path between an area and the backbone.
- B. Provide a backdoor link for traffic to flow through between areas.
- C. Connect the two parts of a partitioned backbone through a nonbackbone area.
- D. Create a virtual connection between two routers from different areas.

Answer: C

NEW QUESTION 120

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.

Answer: C

NEW QUESTION 123

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

Answer: A

NEW QUESTION 124

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

Answer: A

NEW QUESTION 127

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

Answer: AC

Explanation: <http://www.networkers-online.com/blog/2008/12/bgp-performance-tuning-convergence-stability-scalabilityand-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>

NEW QUESTION 131

An engineer is deciding what BFD implementation would be most appropriate for a legacy and unstable Packet Over SONET link. Which solution should be suggested?

- A. aggressive timers with the lowest possible multiplier
- B. slow timers with the lowest possible multiplier
- C. aggressive timers with a very high multiplier
- D. averages of a sequence of ping test response times to the remote peer

Answer: A

NEW QUESTION 132

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

Answer: B

NEW QUESTION 135

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 change
- E. BGP is used to achieve this flexibility.
- F. Multihomed customers must use a provider-assigned address space.

Answer: BD

NEW QUESTION 138

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

Answer: CD

NEW QUESTION 141

A network specialist is working for a large enterprise that has a BGP-enabled multihomed core infrastructure. Recently, the ISP that provides Internet connectivity to the company had to temporarily disable one of the BGP sessions as a result of routers blackholing BGP prefixes in transit. Which option can the specialist use to avoid this incident in the future?

- A. filtering local prefixes outbound
- B. using only one router as primary BGP speaker
- C. permitting only local prefixes outbound
- D. filtering external prefixes inbound

Answer: C

NEW QUESTION 145

Refer to the show command output in the exhibit.

```

Router#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR

Gateway of last resort is not set

 4.0.0.0/24 is subnetted, 1 subnets
 D    4.4.4.0 [90/409600] via 161.108.0.4, 00:49:24, Ethernet0/0
 5.0.0.0/24 is subnetted, 1 subnets
 C    5.5.5.0 is directly connected, Loopback0
162.108.0.0/16 is variably subnetted, 2 subnets, 2 masks
 C    162.108.10.0/24 is directly connected, Serial1/0
 C    162.108.4.0/22 is directly connected, Serial2/0
 C    161.108.0.0/16 is directly connected, Ethernet0/0
Router#
Router#show ip bgp
BGP table version is 6, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
* 11.1.1.0/24     132.108.10.1       0      100     0 1 i
* 1131.108.0.0    132.108.10.1       0      100     0 1 i
*>1161.108.0.0   4.4.4.4            0      100     0  i
Router#

```

For which reason will this router drop all traffic that is destined to the 1.1.1.0/24 network?

- A. The 1.1.1.0/24 route is not synchronized.
- B. The BGP next hop for reaching the 1.1.1.0/24 network is not reachable.
- C. The metric of the 1.1.1.0/24 route is set to 0.
- D. The weight of the 1.1.1.0/24 route is set to 0.
- E. The 1.1.1.0/24 route is an incomplete route.
- F. The IBGP split-horizon rule is preventing the router to use the IBGP route.

Answer: B

NEW QUESTION 150

Which three potential issues can cause an OSPF adjacency to not come up? (Choose three.)

- A. wrong area configured on OSPF peers
- B. wrong authentication key configured on OSPF peers

- C. different OSPF instance numbers configured on OSPF peers
- D. different OSPF instances and VRFs configured with the command: router ospf<number1> vrf <number2>
- E. having the no passive-interface statement configured for the interfaces on which OSPF adjacencies are expected to form
- F. Layer 2 connectivity problems on the shared medium of the participating OSPF-enabled routers

Answer: ABF

NEW QUESTION 152

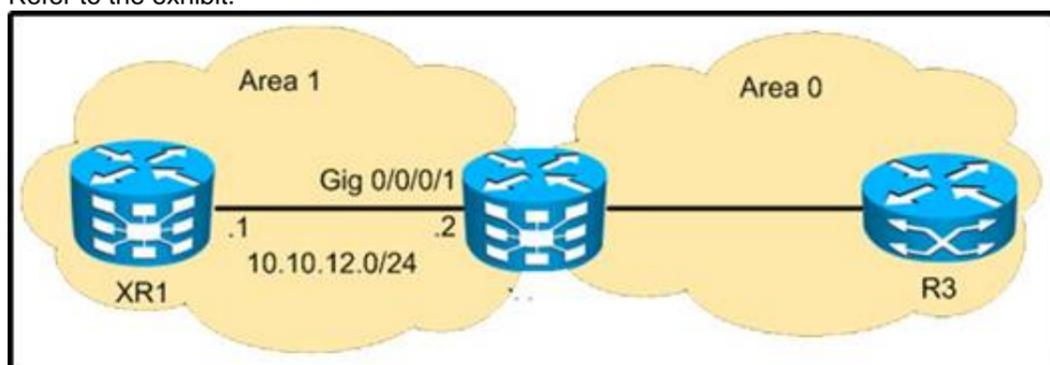
Nonstop Forwarding works with the Stateful Switchover feature in Cisco IOS and IOS XR software. Which three restrictions for OSPF are true? (Choose three.)

- A. OSPF NSF for virtual links is not supported.
- B. OSPF NSF for virtual links is supported.
- C. OSPF NSF for sham links is not supported.
- D. OSPF NSF for sham links is supported.
- E. OSPF NSF supports NSF/SSO for IPv6 traffic only.
- F. OSPF NSF supports NSF/SSO for IPv4 traffic only.

Answer: ACF

NEW QUESTION 156

Refer to the exhibit.



Assume that XR2 and R3 have been configured correctly for OSPF area 0. Which command is needed to configure XR2 as an ABR between area 0 and area 1?

- A. router ospf 1address-family ipv4interface GigabitEthernet0/0/0/1area 1
- B. router ospf 1area 1interface GigabitEthernet0/0/0/1
- C. router ospf 1area 1network 10.10.12.0 0.0.0.255
- D. router ospf 1interface GigabitEthernet0/0/0/1area 1

Answer: B

NEW QUESTION 157

An EIGRP domain is redistributed into an OSPF NSSA area. On which router would this redistribution occur?

- A. stub border router
- B. summary router
- C. autonomous system boundary router
- D. backbone router

Answer: C

NEW QUESTION 158

Which option is a tool that is used to ensure that BGP AS does not become a transit AS?

- A. as-path filter-list
- B. local-preference
- C. ttl-security
- D. confederations

Answer: A

NEW QUESTION 159

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

Answer: C

Explanation: 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

NEW QUESTION 162

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 route
- E. RP/0/RSP0/CPU0:P1(config)#bfd fast-detect
- F. 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

Answer: C

Explanation: 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

NEW QUESTION 164

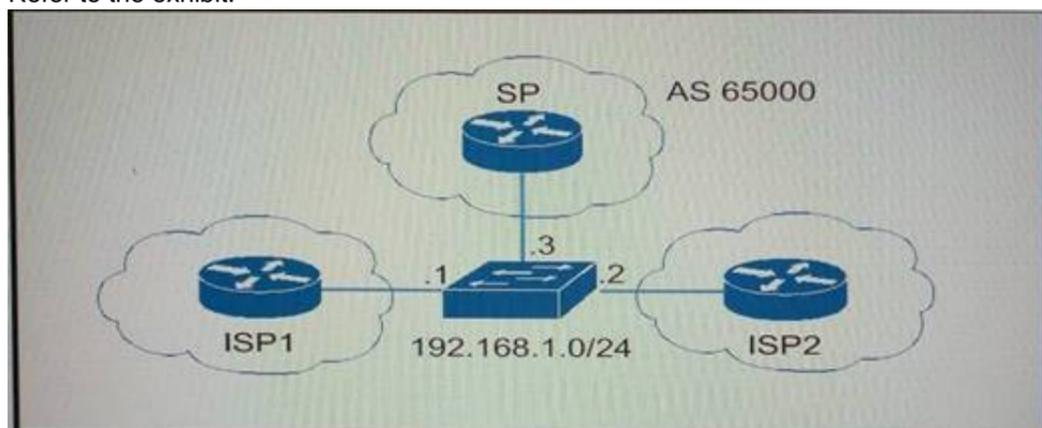
When redistributing EIGRP routes into OSPF as type E2 external OSPF routes, what is the default OSPF seed metric?

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

Answer:

NEW QUESTION 165

Refer to the exhibit.



A service provider (AS65000) is interconnected to two upstream providers (ISP 1 neighbor IP 192.168.1.1 and ISP 2 neighbor IP 192.168.1.2) via a single link in a peering exchange. Which option can an engineer use to implement nonproprietary policies to make ISP 1 the preferred link for incoming and outgoing traffic from the local SP?

- A. route-map outgoing permit 10set-as -path prepend 65000 65000 65000route-map incoming permit 10set local-preference 1000router bgp 65000neighbor 192.168.1.1 route- map incoming outneighbor 192.168.1.1 route-map outgoing in
- B. route-map outgoing permit 10set metric 100route-map incoming permit 10set weight 100router bgp 65000neighbor 192.168.1.1 route-map incoming inneighbor 192.168.1.1 route-map outgoing out
- C. route-map outgoing permit 10set-as -path prepend 65000 65000 65000route-map incoming permit 10set local-preference 1000router bgp 65000neighbor 192.168.1.1 route- map incoming inneighbor 192.168.1.2 route-map outgoing out
- D. set-as -path prepend 65000 65000 65000route-map incoming permit 10set weight 100router bgp 65000neighbor 192.168.1.1 route-map incoming inneighbor 192.168.1.1 route-map outgoing out

Answer: A

NEW QUESTION 166

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

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.

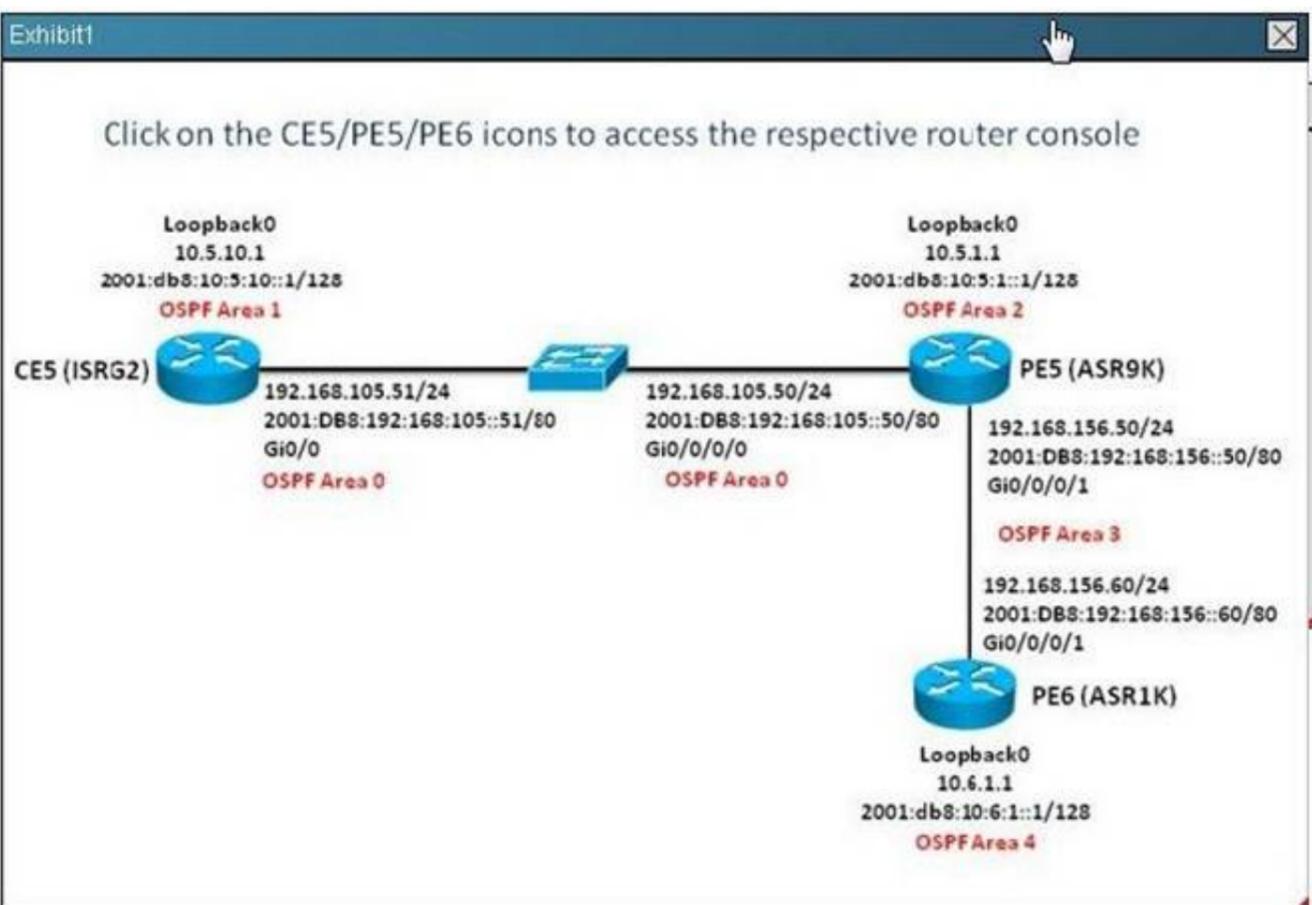
For example, the show running-config and the ping commands are **NOT** supported in this simulation.

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

Scenario

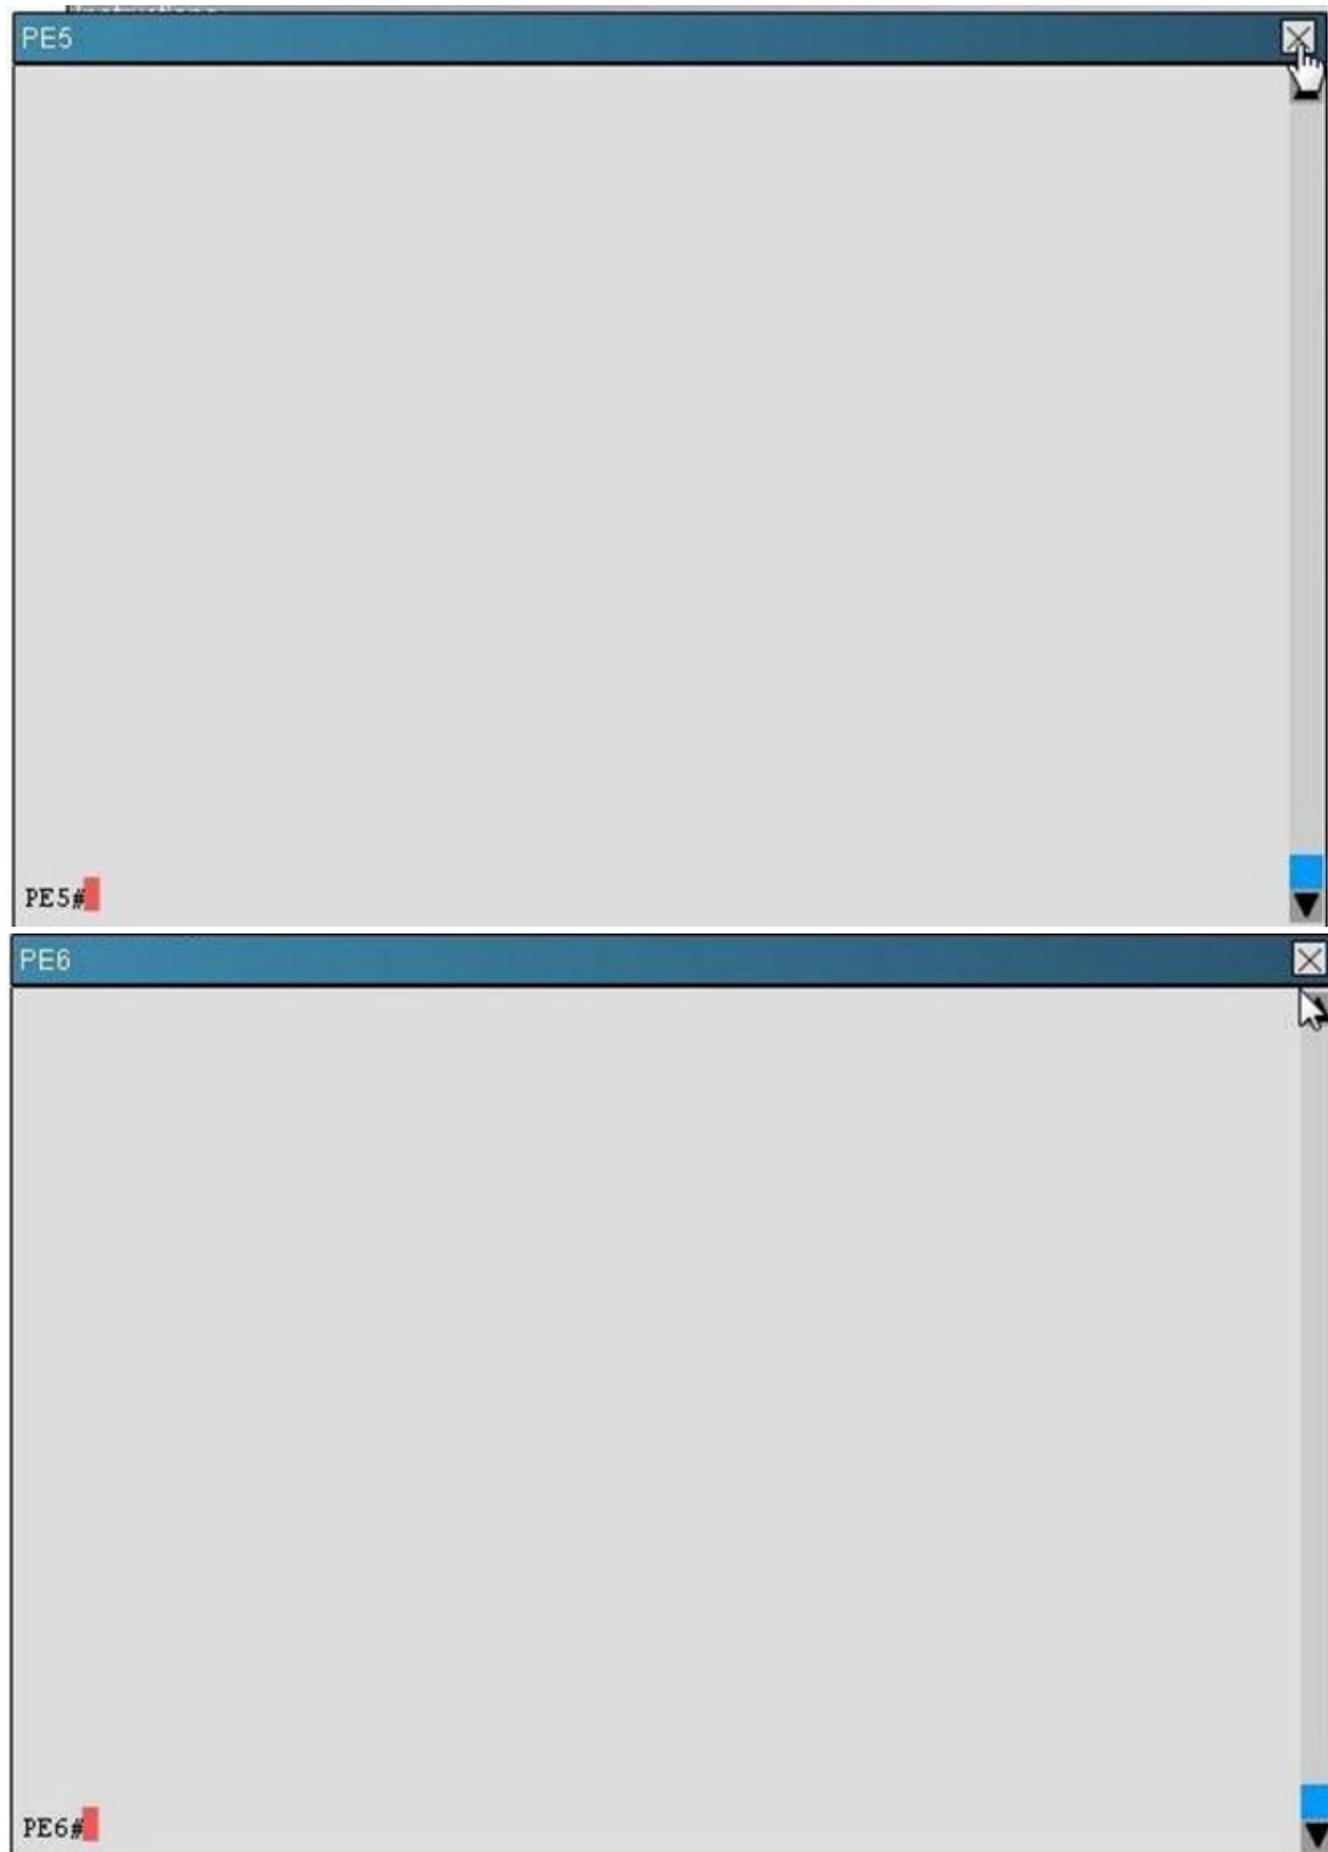
Referring to the multiarea IPv4 and IPv6 OSPF network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

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.



CE5

CE5#



- 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

Answer: CDE

Explanation: #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

NEW QUESTION 168

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

Answer: D

NEW QUESTION 170

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.

Answer: BE

NEW QUESTION 171

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

Answer: BD

NEW QUESTION 175

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

Answer: A

NEW QUESTION 176

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 642-883 Exam with Our Prep Materials Via below:

<https://www.certleader.com/642-883-dumps.html>