642-885 Dumps

Deploying Cisco Service Provider Advanced Routing (SPADVOUTE)

https://www.certleader.com/642-885-dumps.html

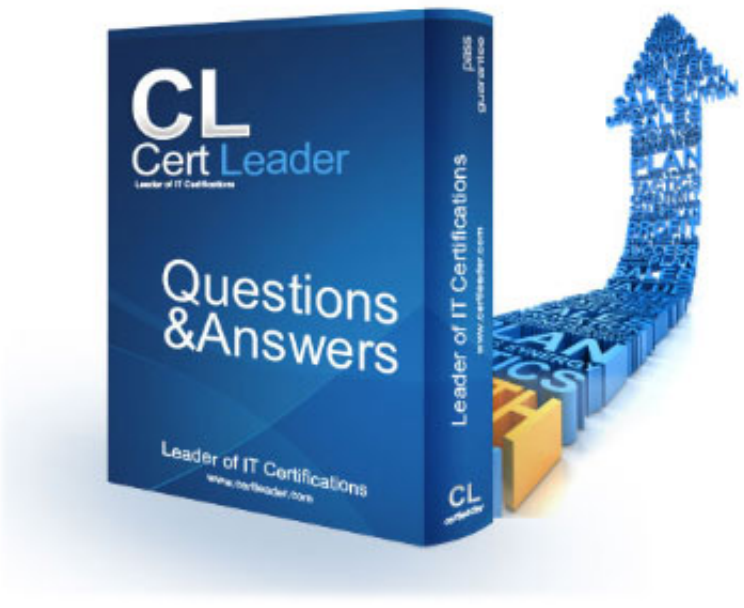

NEW QUESTION 1
Which four operations are components of MSDP in interdomain multicast setup? (Choose four.)

A. Multiple domains can have a single statically defined RP.
B. RPs interconnect between domains with UDP connections to pass source active messages.
C. RPs interconnect between domains with TCP connections to pass source active messages.
D. RPs send source active messages for internal sources to MSDP peers.
E. Source active messages are Peer-RPF checked before accepting or forwarding.
F. RPs learn about external sources via source active messages and may trigger (S,G) joins on behalf of local receivers.
G. MSDP connections typically parallel PIM-SM connections.

Answer: CDEF

NEW QUESTION 2
When implementing IP SLA icmp-echo probes on Cisco IOS-XE routers, which two options are available for IPv6? (Choose two.)

A. flow-label
B. hop-limit
C. DSCP
D. traffic-class
E. TOS

Answer: AD

NEW QUESTION 3
Given the IPv6 address of 2001:0DB8::1:800:200E:88AA, what will be its corresponding the solicited-node multicast address?

A. FF01::1:200E:88AA
B. FF01::1:FF0E:88AA
C. FF01:0DB8::1:800:200E:88AA
D. FF02::1:FF0E:88AA
E. FF02::1:200E:88AA
F. FF02:0DB8::1:800:200E:88AA

Answer: D

Explanation:
IPv6 nodes (hosts and routers) are required to join (receive packets destined for) the following multicast groups:
•All-nodes multicast group FF02:0:0:0:0:0:0:1 (scope is link-local)
•Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 for each of its assigned unicast and anycast addresses
IPv6 routers must also join the all-routers multicast group FF02:0:0:0:0:0:0:2 (scope is link- local).
The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast or anycast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast and anycast address to which it is assigned. The IPv6 solicited- node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast or anycast address (see Figure 2). For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages

NEW QUESTION 4
Refer to the exhibit.

```
Router A:
 interface GigabitEthernet 0/0/0/0
  ipv4 address 10.6.1.1 255.255.255.252
 interface loopback 0
  ipv4 address 10.0.1.1 255.255.255.255
router msdp
 peer 10.0.1.2

Router B:
interface GigabitEthernet 0/0/0/0
  ipv4 address 10.6.1.2 255.255.255.252
 interface loopback 0
  ipv4 address 10.0.1.2 255.255.255.255
router msdp
 peer 10.0.1.1
```

Router A and Router B are connected via GigabitEthernet interfaces, but they are unable to form an MSDP neighborship. Which two components must be addressed when fixing the MSDP peering issue? (Choose two.)

A. An msdp default peer is configured on both routers.
B. A BGP process on each router is present so that MSDP can peer and carry updates.
C. The router interfaces are PIM-enabled to transport MSDP updates.
D. The connect-source attribute is configured with a host route under the MSDP process.
E. The MSDP peering on both routers specifies an origin ID so that it can peer.
F. The router A loopback interface configures the correct subnet mask.

Answer: DF

NEW QUESTION 5
Refer to the exhibit.

224.10.0.1
224.138.0.1
225.10.0.1
225.138.0.1
226.10.0.1
226.138.0.1
227.10.0.1
227.138.0.1
228.10.0.1
228.138.0.1
229.10.0.1
229.138.0.1
230.10.0.1
230.138.0.1
231.10.0.1
231.138.0.1
232.10.0.1
232.138.0.1
233.10.0.1
233.138.0.1
234.10.0.1
234.138.0.1
235.10.0.1
235.138.0.1
236.10.0.1
236.138.0.1
237.10.0.1
237.138.0.1
238.10.0.1
238.138.0.1
239.10.0.1
239.138.0.1

The following multicast IP addresses map to which multicast MAC address?

A. 01:00:5E:8A:00:01
B. 01:00:5E:0A:00:01
C. 01:00:5E:7A:00:01
D. 01:00:5E:05:00:01

Answer: B

NEW QUESTION 6
Refer to the exhibit.

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. If a certain command or command option is not supported, please try to use a different command that is supported.

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 and PE5 routers
and interpret the supported CLI commands outputs to answer the four multiple choice questions.

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

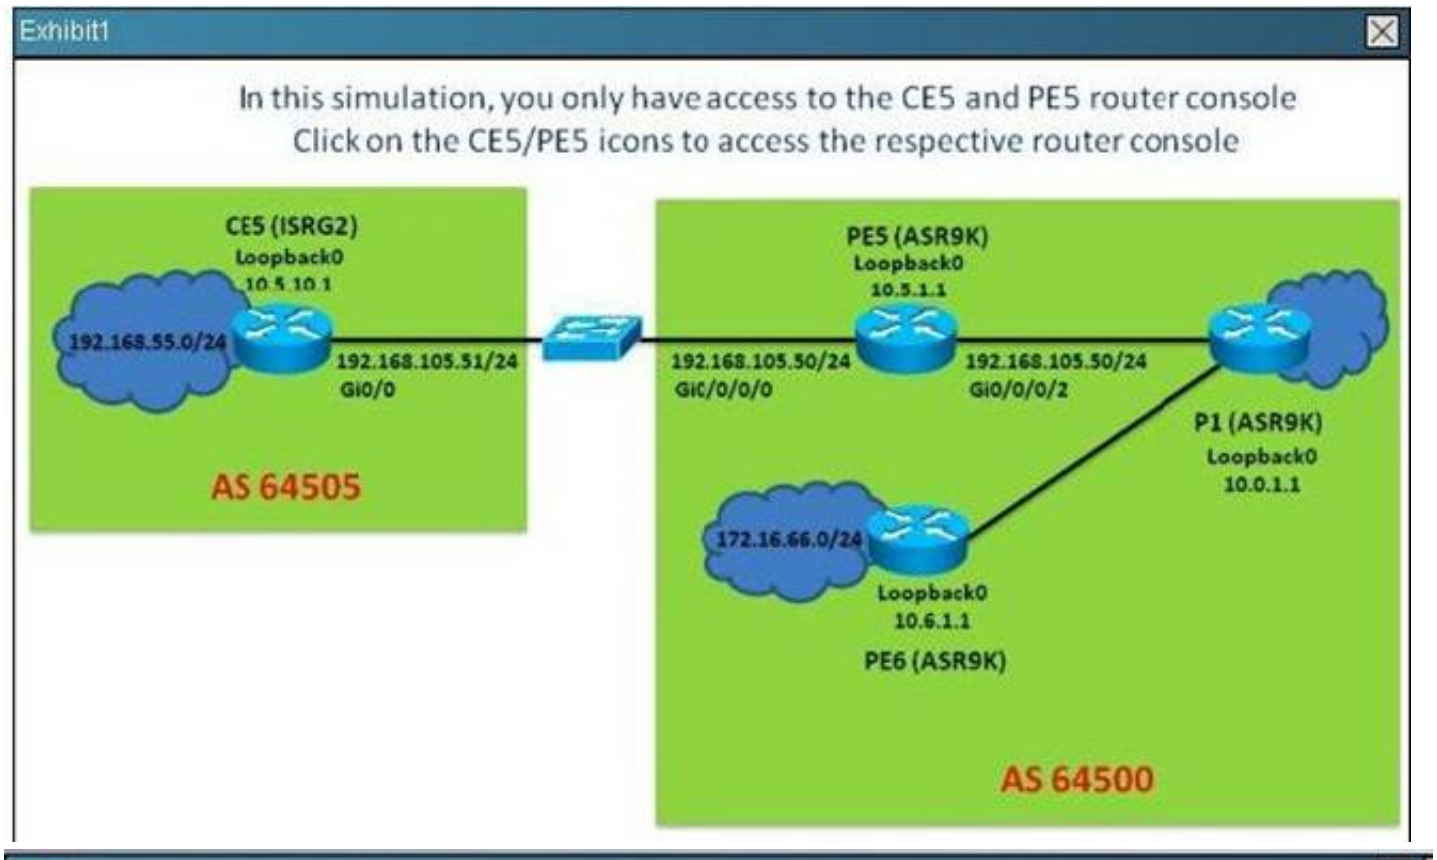

Which three statements regarding the BGP operations are correct? (Choose three)

A. PE5 is the route reflector with P1 and PE6 as its client
B. PE5 is using the IS-IS route to reach the BGP next-hop for the 172.16.66.0/24 prefix
C. PE5 has BGP route dampening enabled
D. The BGP session between PE5 and P1 is established using the loopback interface andnext-hop-self
E. The BGP session between PE5 and CE5 is established using the loopback interface

Answer: ACD

NEW QUESTION 7
After configuring the tunnel interface as shown in the exhibit, no IPv6 traffic is passed over the IPv4 network.

```
interface Tunnel0
ipv6 address 2001:db8:3::1/64
tunnel source GigabitEthernet0/0
tunnel destination 209.165.201.6
tunnel mode ipv6ip
```

Which additional configuration is required to pass the IPv6 traffic over the IPv4 network?

A. Configure an IPv4 address on the tunnel0 interface
B. Configure an IPv6 static route to send the required IPv6 traffic over the tunnel0 interface
C. The tunnel destination should be pointing to an IPv6 address instead of an IPv4 address
D. The tunnel0 interface IPv6 address must use the 2002::///16 prefix

Answer: B

NEW QUESTION 8
Which command set implements BGP support for NSF/SSO on Cisco IOS XE between a PE and a route reflector?

A. On RR:router bgp 300no synchronizationbgp log-neighbor-changesbgp graceful-restart restart-time 120 bgp graceful-restart stalepath-time 360 bgp graceful-restartneighbor 10.20.20.2 remote-as 200neighbor 10.20.20.2 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.2 activateneighbor 10.20.20.2 send-community both neighbor 10.20.20.2 route-reflector-client exit-address-familyOn PE:router bgp 300no synchronizationbgp log-neighbor-changesbgp graceful-restart restart-time 120 bgp graceful-restart stalepath-time 360 bgp graceful-restartneighbor 10.20.20.1 remote-as 300neighbor 10.20.20.1 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.1 activateneighbor 10.20.20.1 send-community both exit-address-family!
B. On RR:router bgp 300no synchronizationbgp log-neighbor-changesbgp graceful-restart restart-time 120 bgp graceful-restart stalepath-time 360 bgp graceful-restartneighbor 10.20.20.2 remote-as 200neighbor 10.20.20.2 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.2 activateneighbor 10.20.20.2 send-community both neighbor 10.20.20.2 route-reflector-client exit-address-familyOn PE:router bgp 300no synchronizationbgp log-neighbor-changes neighbor 10.20.02.1 remote-as 300neighbor 10.20.20.1 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.1 activateneighbor 10.20.20.1 send-community both exit-address-family!
C. On RR:router bgp 300no synchronizationbgp log-neighbor-changesbgp graceful-restart restart-time 120 bgp graceful-restart stalepath-time 360 bgp graceful-restartneighbor 10.20.20.2 remote-as 200neighbor 10.20.20.2 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.2 activateneighbor 10.20.20.2 send-community both neighbor 10.20.20.2 route-reflector-client exit-address-familyOn PE:router bgp 300no synchronizationbgp log-neighbor-changes neighbor 10.20.20.1 remote-as 300neighbor 10.20.20.1 update-source Loopback0 neighbor 10.20.20.1 ha-mode ssono auto-summary!address-family vpnv4 neighbor 10.20.20.1 activateneighbor 10.20.20.1 send-community both exit-address-family!
D. On RR:router bgp 300no synchronizationbgp log-neighbor-changes neighbor 10.20.20.2 remote-as 200neighbor 10.20.20.2 update-source Loopback0 neighbor 10.20.20.2 ha-mode ssono auto-summary!address-family vpnv4 neighbor 10.20.20.2 activateneighbor 10.20.20.2 send-community both neighbor 10.20.20.2 route-reflector-client exit-address-familyOn PE:router bgp 300no synchronizationbgp log-neighbor-changes neighbor 10.20.20.1 remote-as 300neighbor 10.20.20.1 update-source Loopback0 neighbor 10.20.20.1 ha-mode ssono auto-summary!address-family vpnv4 neighbor 10.20.20.1 activateneighbor 10.20.20.1 send-community both exit-address-family!
E. On RR:router bgp 300no synchronizationbgp log-neighbor-changesneighbor 10.20.20.2 remote-as 200neighbor 10.20.20.2 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.2 activateneighbor 10.20.20.2 send-community both neighbor 10.20.20.2 route-reflector-client exit-address-familyOn PE:router bgp 300no synchronizationbgp log-neighbor-changesbgp graceful-restart restart-time 120 bgp graceful-restart stalepath-time 360 bgp graceful-restartneighbor 10.20.20.1 remote-as 300neighbor 10.20.20.1 update-source Loopback0 no auto-summary!address-family vpnv4 neighbor 10.20.20.1 activateneighbor 10.20.20.1 send-community both exit-address-family!

Answer: A

NEW QUESTION 9
Which technology is categorized as multicast ASM and multicast SSM?

A. IP telephony
B. video conferencing
C. IPTV
D. live streaming

Answer: D

NEW QUESTION 10
A network engineer for an ISP wants to reduce the number of iBGP adjacencies. A merge is taking place with another ISP network, so the network engineer needs to make both ASNs look like a single network for the Internet. Which BGP technology is most suitable?

A. route reflector
B. confederation
C. clustering
D. peer group

Answer: B

NEW QUESTION 10
When implementing source-based remote-triggered black hole filtering, which two configurations are required on the edge routers that are not the signaling router? (Choose two.)

A. A static route to a prefix that is not used in the network with a next hop set to the Null0 interface
B. A static route pointing to the IP address of the attacker
C. uRPF on all external facing interfaces at the edge routers
D. Redistribution into BGP of the static route that points to the IP address of the attacker
E. A route policy to set the redistributed static routes with the no-export BGP community

Answer: AC

Explanation:
Source-Based RTBH Filtering
With destination-based black holing, all traffic to a specific destination is dropped after the black hole has been activated, regardless of where it is coming from. Obviously, this could include legitimate traffic destined for the target. Source-based black holes provide the ability to drop traffic at the network edge based on a specific source address or range of source addresses.
If the source address (or range of addresses) of the attack can be identified (spoofed or not), it would be better to drop all traffic at the edge based on the source address, regardless of the destination address. This would permit legitimate traffic from other sources to reach the target. Implementation of source-based black hole filtering depends on Unicast Reverse Path Forwarding (uRPF), most often loose mode uRPF.
Loose mode uRPF checks the packet and forwards it if there is a route entry for the source IP of the incoming packet in the router forwarding information base (FIB). If the router does not have an FIB entry for the source IP address, or if the entry points to a null interface, the Reverse Path Forwarding (RPF) check fails and the packet is dropped, as shown in Figure
2. Because uRPF validates a source IP address against its FIB entry, dropping traffic from specific source addresses is accomplished by configuring loose mode uRPF on the external interface and ensuring the RPF check fails by inserting a route to the source with a next hop of Null0.
This can be done by using a trigger device to send IBGP updates. These updates set the next hop for the source IP to an unused IP address that has a static entry at the edge, setting it to null as shown in Figure 2.

Figure 2. Source-Based Black Hole Filtering

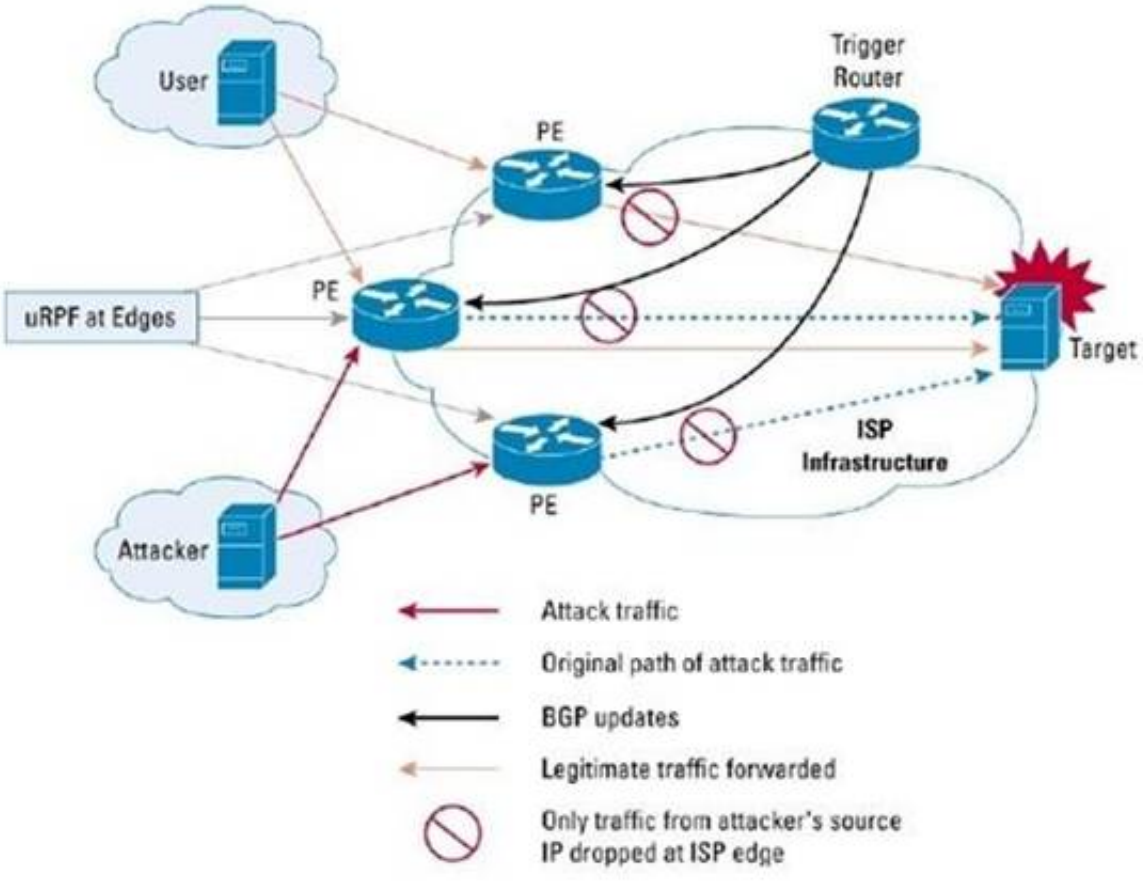

In this way, traffic that is entering the edge network sourced from a host that has a route pointing to null will result in a uRPF drop.

NEW QUESTION 11
Which IPv6 mechanism occurs between a provider edge router and the customer premises equipment router to allow an ISP to automate the process of assigning a block of IPv6 addresses to a customer for use within the customer network?

A. Router Advertisement
B. DHCPv6 Prefix Delegation
C. DHCPv6 Lite
D. Stateful DHCPv6

Answer: B

Explanation:
http://www.cisco.com/en/US/tech/tk872/technologies_configuration_example09186a0080b 8a116.shtml

NEW QUESTION 16
When implementing high-availability stateful switchover BGP routing, in which situation would Cisco NSR be required?

A. On the PE routers connecting to the CE routers which are not NSF aware or are not NSF capable
B. On the PE routers connecting to the CE routers which support graceful restart
C. On the PE routers connecting to the CE routers which are incapable of performing stateful switchover operations because the CE routers are only NSF aware but not NSF capable
D. On the PE routers connecting to the CE routers which are incapable of performingstateful switchover operations because the CE routers are only NSF capable

but not NSF aware
E. On the service provider core P routers which are also NSF aware
F. On the service provider core P routers which are also NSF capable

Answer: A

NEW QUESTION 17
The IPv6 2002::/16 prefix is used in which kind of implementations?

A. 6 RD
B. 6 to 4
C. NAT 64
D. IPv6 Multicast

Answer: B

NEW QUESTION 22
An engineer is providing DNS for IPv6 over a currently working IPv4 domain. Which three changes are needed to offer DNS functionality for IPv6? (Choose three.)

A. Define a new record that stores the 128-bit IPv6 address.
B. Expand the existing IP address record to allow for 128 bits.
C. Define the IPv6 equivalent of the in-addr.arpa.com domain of the IPv4 PTR.
D. Modify the in-addr.arpa.com domain of the IPv4 PTR.
E. Change the query messages.
F. Transport IPv6 query messages by using UDP.
G. Transport IPv6 query messages by using TCP.

Answer: ACE

NEW QUESTION 26
Which three statements regarding NAT64 operations are correct? (Choose three.)

A. With stateful NAT64, many IPv6 address can be translated into one IPv4 address, thus IPv4 address conservation is achieved
B. Stateful NAT64 requires the use of static translation slots so IPv6 hosts and initiate connections to IPv4 hosts.
C. With stateless NAT64, the source and destination IPv4 addresses are embedded in the IPv6 addresses
D. NAT64 works in conjunction with DNS64
E. Both the stateful and stateless NAT64 methods will conserve IPv4 address usage

Answer: ACD

Explanation:
Stateful NAT64-Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers
Stateful NAT64 multiplexes many IPv6 devices into a single IPv4 address. It can be assumed that this technology will be used mainly where IPv6-only networks and clients (ie. Mobile handsets, IPv6 only wireless, etc...) need access to the IPv4 internet and its services.
The big difference with stateful NAT64 is the elimination of the algorithmic binding between the IPv6 address and the IPv4 address. In exchange, state is created in the NAT64 device for every flow. Additionally, NAT64 only supports IPv6-initiated flows. Unlike stateless NAT64, stateful NAT64 does `not' consume a single IPv4 address for each IPv6 device that wants to communicate to the IPv4 Internet. More practically this means that many IPv6- only users consume only single IPv4 address in similar manner as IPv4-to-IPv4 network address and port translation works. This works very well if the connectivity request is initiated from the IPv6 towards the IPv4 Internet. If an IPv4-only device wants to speak to an IPv6-only server for example, manual configuration of the translation slot will be required, making this mechanism less attractive to provide IPv6 services towards the IPv4 Internet. DNS64 is usually also necessary with a stateful NAT64, and works the same with both stateless and stateful NAT64
Stateless NAT64-Stateless translation between IPv4 and IPv6 RFC6145 (IP/ICMP Translation Algorithm) replaces RFC2765 (Stateless IP/ICMP Translation Algorithm (SIIT)) and provides a stateless mechanism to translate a IPv4 header into an IPv6 header and vice versa. Due to the stateless character this mechanism is very effective and highly fail safe because more as a single-or multiple translators in parallel can be deployed and work all in parallel without a need to synchronize between the translation devices.
The key to the stateless translation is in the fact that the IPv4 address is directly embedded in the IPv6 address. A limitation of stateless NAT64 translation is that it directly translates only the IPv4 options that have direct IPv6 counterparts, and that it does not translate any IPv6 extension headers beyond the fragmentation extension header; however, these limitations are not significant in practice.
With a stateless NAT64, a specific IPv6 address range will represent IPv4 systems within the IPv6 world. This range needs to be manually configured on the translation device. Within the IPv4 world all the IPv6 systems have directly correlated IPv4 addresses that can be algorithmically mapped to a subset of the service provider's IPv4 addresses. By means of this direct mapping algorithm there is no need to keep state for any translation slot between IPv4 and IPv6. This mapping algorithm requires the IPv6 hosts be assigned specific IPv6 addresses, using manual configuration or DHCPv6.
Stateless NAT64 will work very successful as proven in some of the largest networks, however it suffers from some an important side-effect: Stateless NAT64 translation will give an IPv6-only host access to the IPv4 world and vice versa, however it consumes an IPv4 address for each IPv6-only device that desires translation -- exactly the same as a dual- stack deployment. Consequentially, stateless NAT64 is no solution to address the ongoing IPv4 address depletion.Stateless NAT64 is a good tool to provide Internet servers with an accessible IP address for both IPv4 and IPv6 on the global Internet. To aggregate many IPv6 users into a single IPv4 address, stateful NAT64 is required. NAT64 are usually deployed in conjunction with a DNS64. This functions similar to, but different than, DNS- ALG that was part of NAT-PT. DNS64 is not an ALG; instead, packets are sent directly to and received from the DNS64's IP address. DNS64 can also work with DNSSEC (whereas DNS-ALG could not).

NEW QUESTION 29
Which two actions result when a network administrator attempts to ping an IPv6 host on the LAN? (Choose two.)

A. ARP is used to determine the MAC address of the destination host.
B. Neighbor Discovery is used to determine the MAC address of the destination host.
C. Neighbor Solicitation messages are sent out by the source host to determine the data link-layer address of the destination host.
D. Neighbor Advertisement messages are sent by the source host to announce its presence on the local link.
E. Router Solicitation messages are sent out on a specific multicast address to request the data link-layer address of the target device.
F. Router Solicitation messages are sent to the local router on the network segment to request data link-layer information about the destination host.

Answer: BC

NEW QUESTION 31
An SP core is running PIM on the network. Multicast groups in this networkare in the 232.0.0.0/8 range. Which commandenables multicast routing operations without using an RP?

A. ip pim autorp
B. ip pim ssm default
C. ip pim bidir-enable
D. ip pim register-source

Answer: B

NEW QUESTION 36
A service providerrequests more details about the recent Inter-AS MPLS VPN Option B configuration that was recently deployed. Consider this configuration:
router bgp 3717
address-family vpnv4 unicast retain route-target all
commit
!
Which option describes why this particular command is needed?

A. The ASBRcan have many working customer VRFs, so this configuration ensures the coexistence of all the route-target extended communities that belong to the all ASBR- terminated customer VRFs.
B. When implementing the Inter-AS Option B MPLS VPN solution, all the route targets that are transmitted over the Inter-AS links need an ASBR local database to forward thecustomer traffic correctly.
C. The Inter-AS Option B design implements VPNv4 communication over the Inter-AS link, hence the requirement for a route-target association for each customer VPN connected across two or more ASs.
D. In the Inter-AS Option B design, no local VRF is maintained on the ASBR routers,so the default behavior of the operating system is to deny any route-target extended community that is encoded in the incoming iBGP update
E. This configuration permits VPNv4 communication by accepting the iBGP updates even if no route targets are configured locally.

Answer: D

NEW QUESTION 41
Which statement is correct regarding using the TTL threshold to define the delivery boundaries of multicast traffic?

A. If a packet TTL is less than the specified TTL threshold, the packet is forwarded out of the interface
B. If a packet TTL is greater or equal to the specified TTL threshold, the packet is forwarded out of the interface
C. If a packet TTL is equal to the specified TTL threshold, the packet is dropped
D. When a multicast packet arrives, the TTL threshold value is decremented by 1. If the resulting TTL threshold value is greater than or equal to 0, the packet is dropped

Answer: B

NEW QUESTION 45
In which four ways does DHCPv6 differ from DHCPv4? (Choose four.)

A. DHCPv6 uses the same message types as DHCPv4.
B. DHCPv4 functions without external protocols.
C. A host discovers a DHCPv6 server by using a DHCP Discover packet.
D. A hosts discovers a DHCPv6 server by using a DHCP Solicit packet.
E. A DHCPv6 server replies with a DHCP Offer packet.
F. A DHCP server replies with a DHCP Advertise message.
G. An IPv6 host can request multiple addresses at the same time from a DHCPv6 server.
H. An IPv6 host can request only one IP address at a time from a DHCPv6 server.

Answer: BDFG

NEW QUESTION 50
Refer to the exhibit.

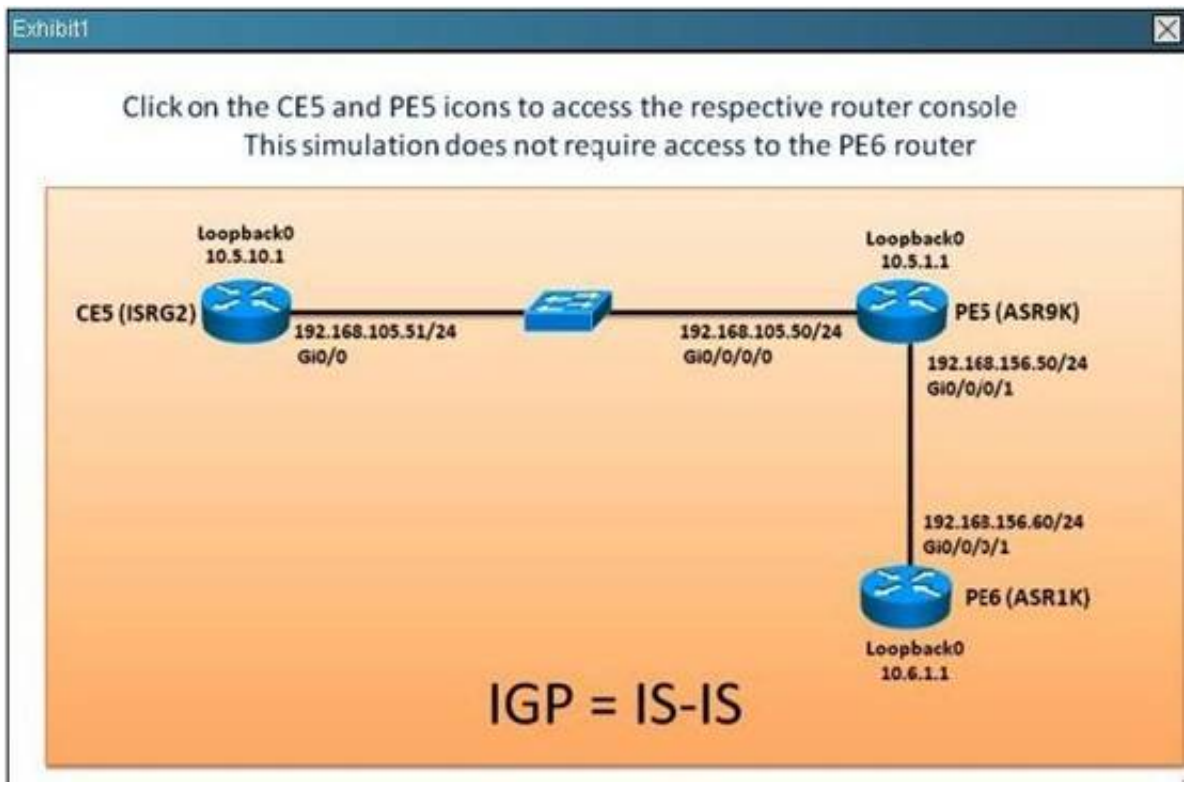

Which three statements are correct regarding the various multicast groups? (Choose three.)

A. Currently there is no source sending traffic to the 224.1.1.1 multicast group
B. PE5 has a Null OILforthe (*,224.0.1.40) entry
C. PE5 has a Null OILforthe (*,224.1.1.1) entry
D. CE5 has joined the 224.0.1.40 multicast group
E. CE5 has a Null OILforthe (*,224.1.1.1) entry

Answer: CDE

Explanation:
#show ip mroute

NEW QUESTION 52
To which three IP multicast groups can a multicast MAC address "01-00-5E-4D-62-B1" listen? (Choose three.)

A. 231.205.98.177
B. 231.205.99.177
C. 239.77.98.177
D. 239.205.99.177
E. 224.205.98.177
F. 224.205.99.177

Answer: ACE

NEW QUESTION 55
A network architect is responsible for the company's multicast network domain design. Which multicast component acts as a meeting place for sources and receivers?

A. multicast shared tree
B. multicast distribution point
C. multicast rendezvous point
D. multicast source tree

Answer: C

NEW QUESTION 59
Which multicast routing protocol supports dense mode, sparse mode and bidirectional mode?

A. DVMRP
B. MOSPF
C. PIM
D. MP-BGP
E. MSDP

Answer: C

NEW QUESTION 63
When configuring PIM operations, what is the effect of setting the SPT threshold to infinity?

A. The multicast source to the RP path will never switch over to the shortest path tree
B. All the PIM routers will have more (S,G) states, thus consuming more router resources
C. The receivers will be able to immediately switch over to the shortest path tree after receiving the first multicast packets on the shared tree via the RP
D. The last-hop routers will never switch over to the shortest path tree and will always remain on the shared tree

Answer: D

NEW QUESTION 65
R1 is designated as the PIM RP within the SP core. Which two configuration parameters must be used to enable and activate R1 as the BSR and RP for the core environment? (Choose two.)

A. ip pim send-rp-announce loopback0 scope 16
B. ip pim bsr-candidate loopback0
C. ip pim send-rp-discovery loopback0 scope 16
D. ip pim rp-candidate loopback0
E. ip pim send-RP-announce loopback0 scope 16 group-list 1

Answer: BD

NEW QUESTION 70
Refer to the exhibit.

```
router bgp 65123
 bgp graceful-restart
```

Which statement correctly explains the bgp graceful-restart command?

A. This command is used to enable NSR and is entered on the NSR-capable router, and also on any NSR-aware peer
B. This command is used to enable NSF and is entered on the NSF-capable router, and also on any NSF-aware peer
C. This command is only required on the NSF-capable routers to enable BGP graceful restart with the BGP peers
D. This command is only required on the NSF-aware routers to enable BGP graceful restart with the BGP peers
E. This command is only required on the NSR-capable routers to enable BGP graceful restart with the BGP peers

Answer: B

Explanation:
Graceful restart is supported in recent versions of Cisco IOS software (12.0S) and is supported in Cisco IOS XR software. Graceful restart is the mechanism by

which BGP routing peers avoid changes to their forwarding paths following a switchover. If the BGP peer has received this capability, it is aware that the device sending the message is nonstop forwarding (NSF)-capable. Both the NSF-capable router and its BGP peers (NSFaware peers) need to exchange the graceful restart capability in their OPEN messages, at the time of session establishment. If both peers do not exchange the graceful restart capability, the session will not be graceful restart-capable.
If the BGP session is lost during a Route Processor (RP) switchover or BGP process restart, the NSF-aware BGP peer marks all the routes associated with the NSF-capable router as stale; however, it continues to use these routes to make forwarding decisions for a set period of time. This functionality means that no packets are lost while the newly active RP is waiting for convergence of the routing information with its BGP peers.
After a failover event occurs, the NSF-capable router reestablishes the session with the BGP peer. In establishing the new session, it sends a new graceful restart message that identifies the NSF-capable router as having restarted. At this point, the routing information is exchanged between the two BGP peers. Once this exchange is complete, the NSF- capable device uses the newly received routing information to update the RIB and the Forwarding Information Base (FIB) with the new forwarding information. The NSF-aware device uses the network information to remove stale routes from its BGP table. The BGP protocol is then fully converged.
If a BGP peer does not support the graceful restart capability, it will ignore the graceful restart capability in an OPEN message but will establish a BGP session with the NSF- capable device. This functionality will allow interoperability with non-NSF-aware BGP peers (and without NSF functionality), but the BGP session with non- NSF-aware BGP peers will not be graceful restart-capable.

NEW QUESTION 72
Which four statements are correct regarding MSDP configurations and operations? (Choose four.)

A. The MSDP peers are also typically the RPs in respective routing domains.
B. SA messages are flooded to all other MSDP peers without any restrictions
C. On Cisco IOS, IOS-XE, and IOS-XR, the router can be configured to cache the SA messages to reduce the join latency
D. SA messages are used to advertise active sources in a domain
E. MSDP establishes neighbor relationships with other MSDP peers using TCP port 639
F. MSDP peerings on Cisco IOS, IOS-XE, and IOS-XR support MD5 or SHA1 authentication

Answer: ACDE

NEW QUESTION 76
Refer to the exhibit.

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.

Exhibit1

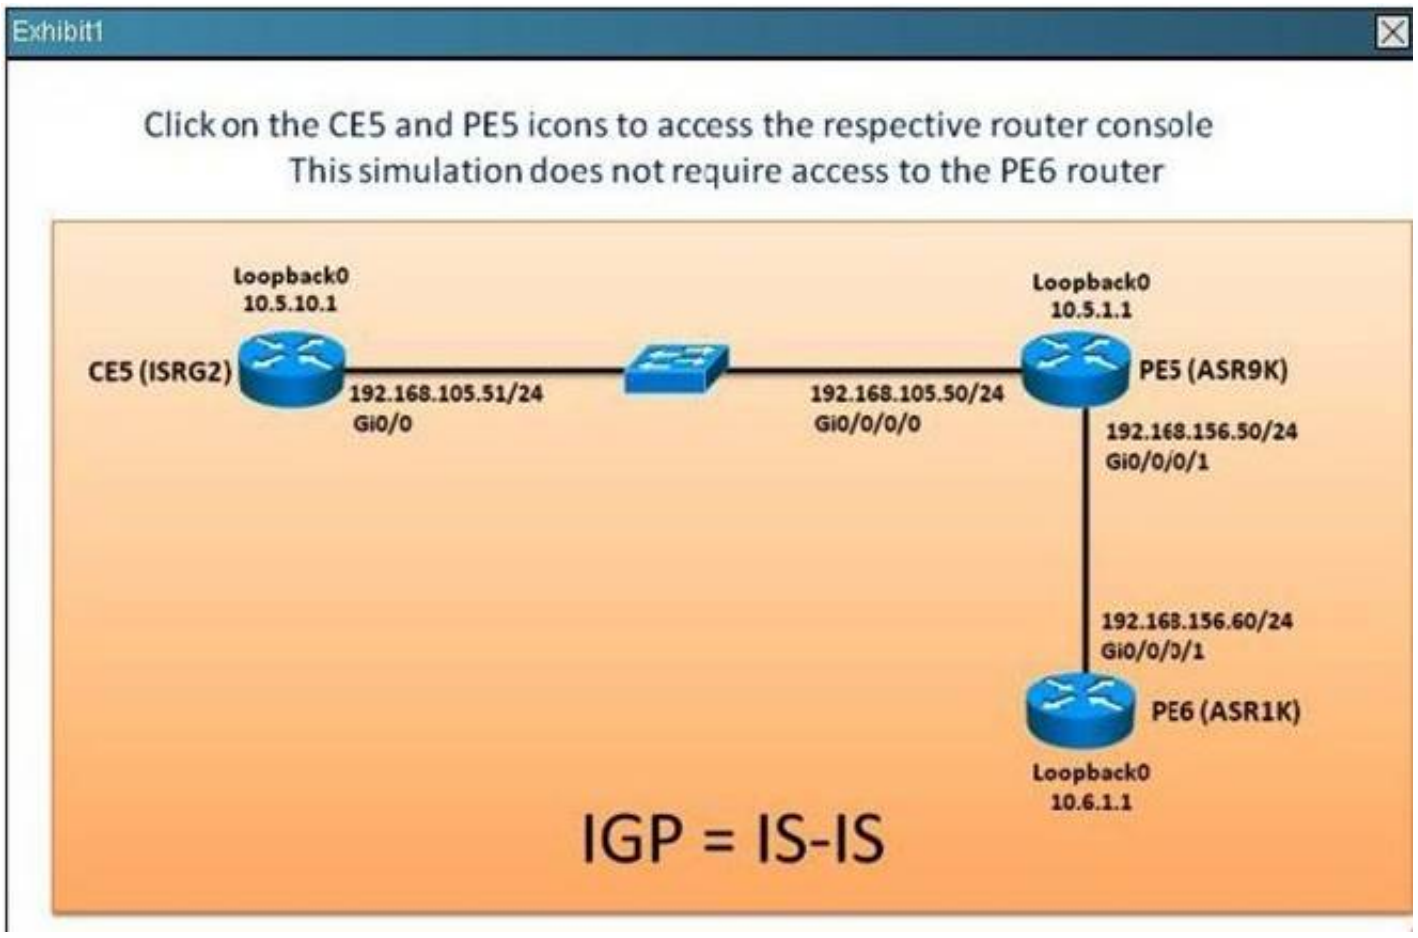

Which two statements are correct regarding the multicast operations on the router that is the RP? (Choose two.)

A. It is using IGMPv3
B. The IGMP query interval is set to 125 seconds
C. It is using the IPv4 unicast routing table to perform the RPF checks
D. Static multicast routes are configured on the RP

Answer: AC

Explanation:
#show ip mroute
#show ip pim interface
#show ip igmp group
#show ip pim neighbor

NEW QUESTION 77
Refer to the topology diagram shown in the exhibit and the partial configurations shown below.

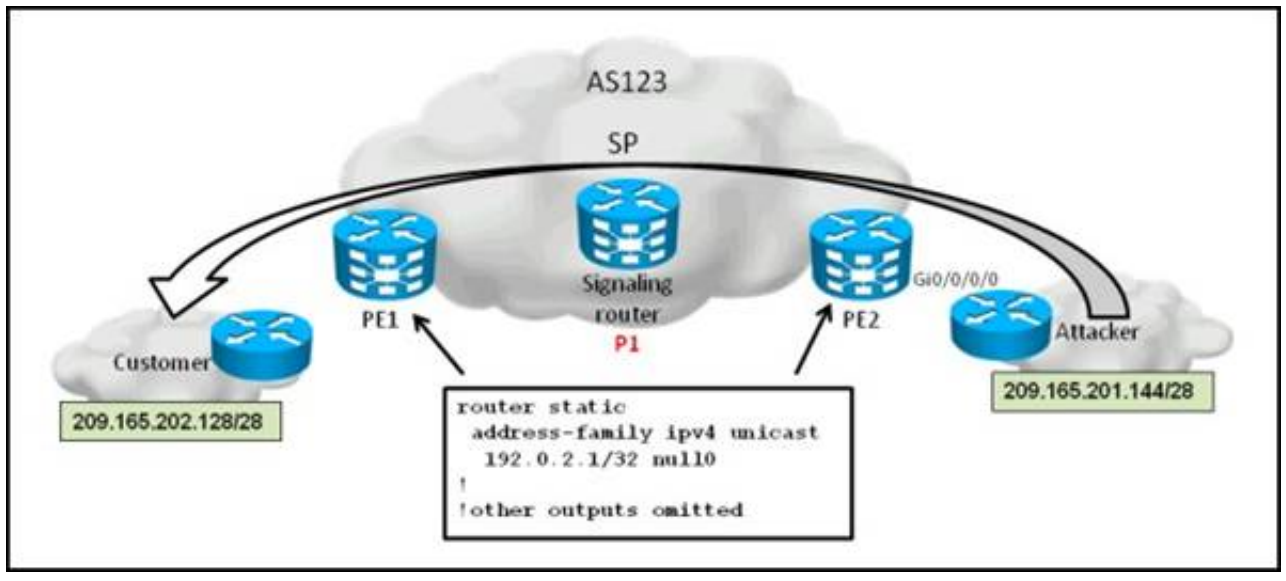

Once the attack from 209.165.201.144/28 to 209.165.202.128/28 has been detected, which additional configurations are required on the P1 IOS-XR router to implement source-based remote-triggered black hole filtering?
!
router bgp 123
address-family ipv4 unicast redistribute static route-policy test
!

A. router staticaddress-family ipv4 unicast 209.165.202.128/28 null0 tag 666192.0.2.1/32 null0 tag 667!route-policy test if tag is 666 thenset next-hop 192.0.2.1 endifif tag is 667 thenset community (no-export) endifend-policy!
B. router staticaddress-family ipv4 unicast 209.165.201.144/28 null0 tag 666192.0.2.1/32 null0 tag 667!route-policy test if tag is 666 thenset next-hop 192.0.2.1 endifif tag is 667 thenset community (no-export) endifend-policy!
C. router staticaddress-family ipv4 unicast 209.165.201.144/28 null0 tag 666192.0.2.1/32 null0!route-policy test if tag is 666 thenset next-hop 192.0.2.1set community (no-export) endifend-policy
D. router staticaddress-family ipv4 unicast 209.165.202.128/28 null0 tag 666192.0.2.1/32 null0!route-policy test if tag is 666 thenset next-hop 192.0.2.1set community (no-export) endifend-policy!

Answer: C

Explanation:
Source-Based RTBH Filtering
With destination-based black holing, all traffic to a specific destination is dropped after the black hole has been activated, regardless of where it is coming from. Obviously, this could include legitimate traffic destined for the target. Source-based black holes provide the ability to drop traffic at the network edge based on a specific source address or range of source addresses.
If the source address (or range of addresses) of the attack can be identified (spoofed or not), it would be better to drop all traffic at the edge based on the source address, regardless of the destination address. This would permit legitimate traffic from other sources to reach the target. Implementation of source-based black hole filtering depends on Unicast Reverse Path Forwarding (uRPF), most often loose mode uRPF. Loose mode uRPF checks the packet and forwards it if there is a route entry for the source IP of the incoming packet in the router forwarding information base (FIB). If the router does not have an FIB entry for the source IP address, or if the entry points to a null interface, the Reverse Path Forwarding (RPF) check fails and the packet is dropped, as shown in Figure 2. Because uRPF validates a source IP address against its FIB entry, dropping traffic from specific source addresses is accomplished by configuring loose mode uRPF on the external interface and ensuring the RPF check fails by inserting a route to the source with a next hop of Null0. This can be done by using a trigger device to send IBGP updates. These updates set the next hop for the source IP to an unused IP address that has a static entry at the edge, setting it to null as shown in Figure 2.
Figure 2. Source-Based Black Hole Filtering

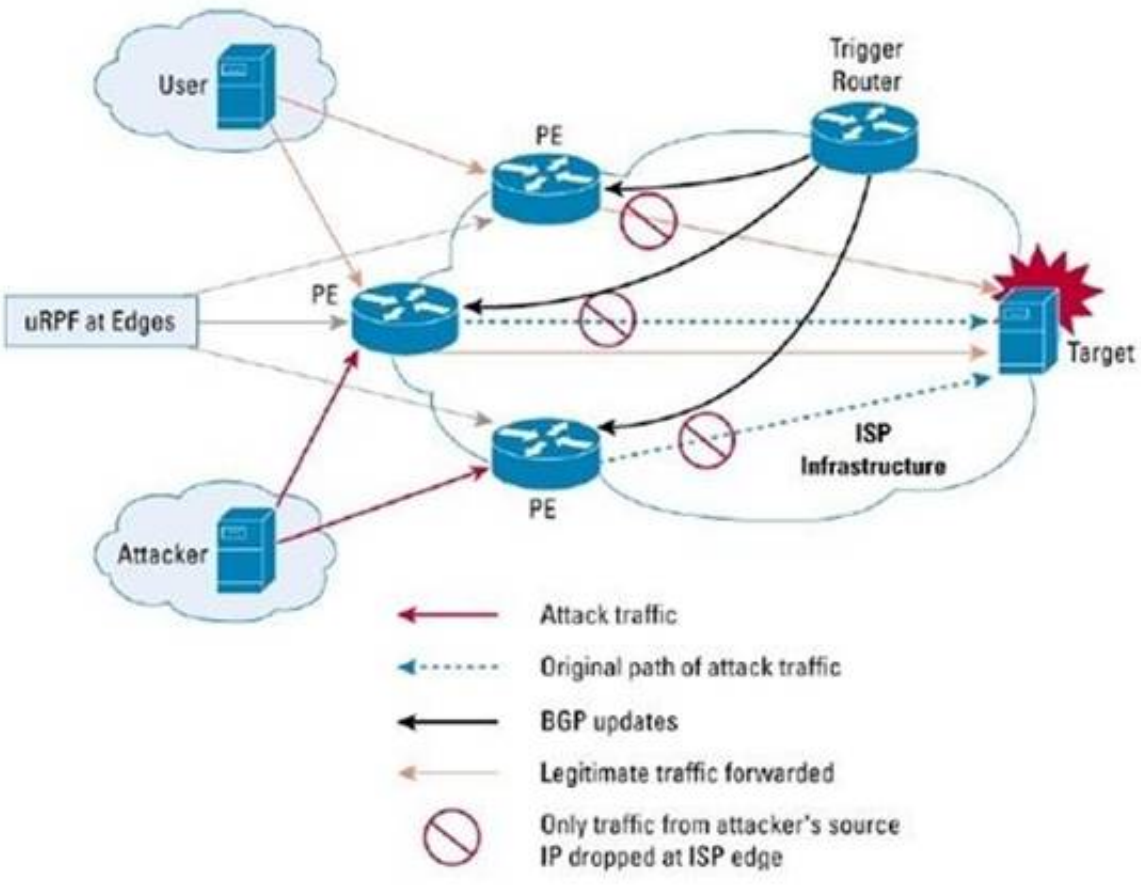

In this way, traffic that is entering the edge network sourced from a host that has a route pointing to null will result in a uRPF drop.

NEW QUESTION 78
In Cisco IOS-XR, the ttl-security command is configured under which configuration mode?

A. RP/0/RSP0/CPU0:P2(config)#
B. RP/0/RSP0/CPU0:P2(config-bgp)#
C. RP/0/RSP0/CPU0:P2(config-bgp-nbr)#
D. RP/0/RSP0/CPU0:P2(config-bgp-af)#
E. RP/0/RSP0/CPU0:P2(config-bgp-nbr-af)#

Answer: C

Explanation:
http://packetlife.net/blog/2009/nov/23/understanding-bgp-ttl-security/

NEW QUESTION 81
Refer to the exhibit.

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.

Exhibit1

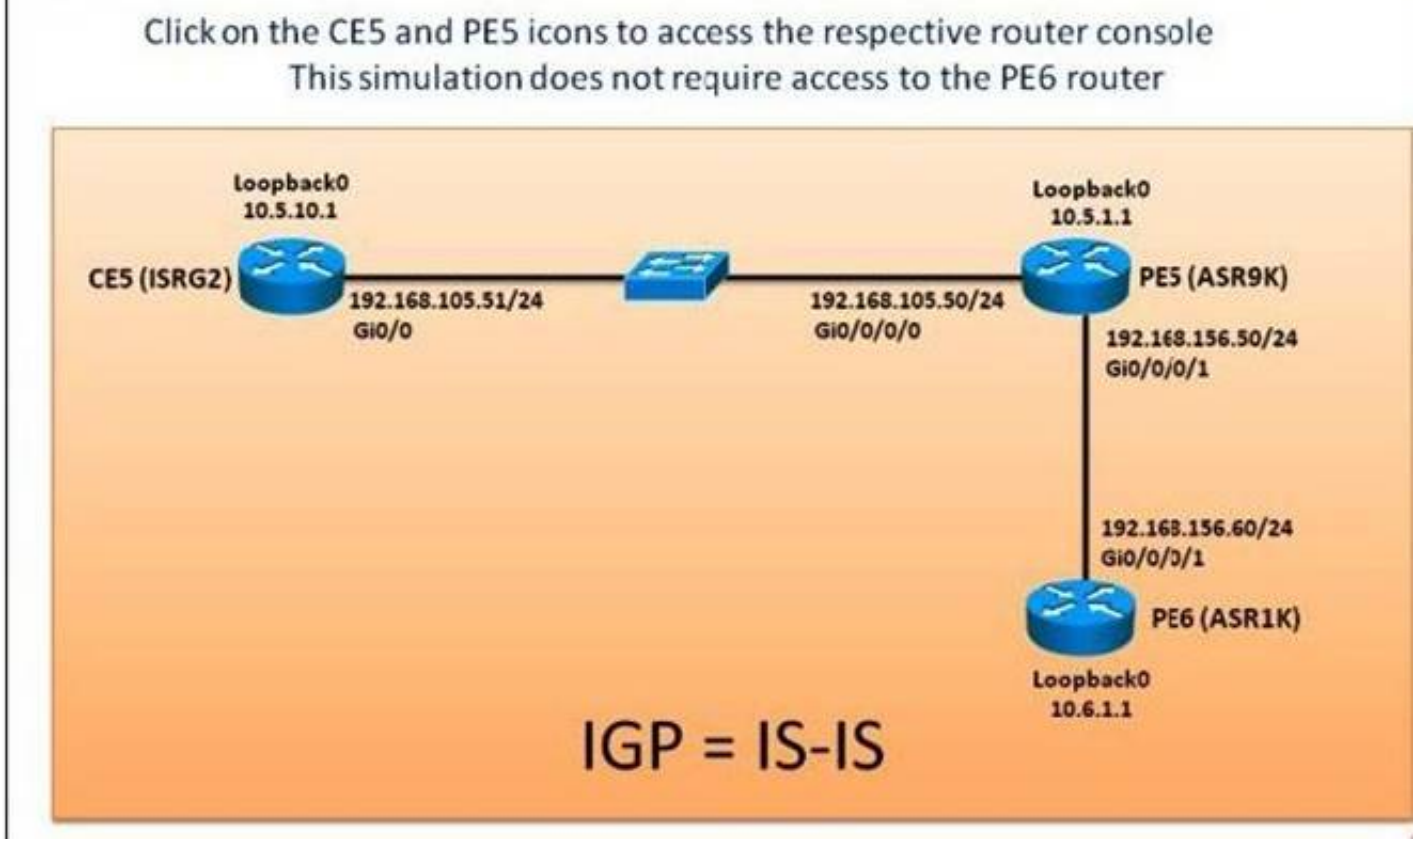

On the PE, which two statements are correct regarding the(192.168.156.60,224.1.1.1) entry? (Choose two,)

A. The RPF neighbor points towards the RP
B. The RPF neighbor is reachable overthe Gi0/0/0/1 interface
C. The OIL contains the GiO/0/0/0 interface
D. The IIL is Null

Answer: AC

Explanation:
#show ip mroute

NEW QUESTION 82
When a BGP route reflector receives an IBGP update from a non-client IBGP peer, the route reflector will then forward the IBGP updates to which other router(s)?

A. To the other clients only
B. To the EBGP peers only
C. To the EBGP peers and other clients only
D. To the EBGP peers and other clients and non-clients

Answer: C

NEW QUESTION 83

Which configuration would an engineer use to exchange IPv6 multicast routes via BGP with a neighbor that does not support thecorresponding Multicast SAFI onCisco IOS XE?

A. router bgp 100bgp router-id 209.165.201.10 no bgp default ipv4-unicastneighbor 2001:DB8::10 remote-as 201neighbor 2001:DB8::10 update-source GigabitEthernet 0/10 address-family ipv6 multicastneighbor 2001:DB8::10 activate network 2001:DB8:CDCD:1::/64exit-address-family
B. router bgp 100bgp router-id 209.165.201.10 no bgp default ipv4-unicastneighbor 2001:DB8::10 remote-as 201neighbor 2001:DB8::10 update-source GigabitEthernet 0/10 address-family ipv6neighbor 2001:DB8::10 translate-update ipv6 multicast unicast neighbor 2001:DB8::10 activateno synchronization exit address-familyaddress-family ipv6 multicast neighbor 2001:DB8::10 activate network 2001:DB8:CDCD:1::/64exit-address-family
C. router bgp 100bgp router-id 209.165.201.10 no bgp default ipv4-unicastneighbor 2001:DB8::10 remote-as 201neighbor 2001:DB8::10 update-source GigabitEthernet 0/10 address-family ipv6neighbor 2001:DB8::10 activate address-family ipv6 multicast neighbor 2001:DB8::10 activate network 2001:DB8:CDCD:1::/64exit-address-family
D. router bgp 100bgp router-id 209.165.201.10 no bgp default ipv4-unicastneighbor 2001:DB8::10 remote-as 201neighbor 2001:DB8::10 update-source GigabitEthernet 0/10 address-family ipv6neighbor 2001:DB8::10 translate-update ipv6 multicast unicast no synchronizationexit address-familyaddress-family ipv6 multicast neighbor 2001:DB8::10 activate network 2001:DB8:CDCD:1::/64exit-address-family
E. router bgp 100bgp router-id 209.165.201.10 no bgp default ipv4-unicastneighbor 2001:DB8::10 remote-as 201neighbor 2001:DB8::10 update-source GigabitEthernet 0/10 address-family ipv6neighbor 2001:DB8::10 send-labelneighbor 2001:DB8::10 override-capability-neg neighbor 2001:DB8::10 activateno synchronization exit address-familyaddress-family ipv6 multicast network 2001:DB8:CDCD:1::/64exit-address-family

Answer: B

NEW QUESTION 87
With IPv6 multicast, which feature can be used as a replacement method for static RP configuration?

A. PIM Snooping
B. MLD
C. MLD Snooping
D. Embedded RP
E. DHCPv6

Answer: D

NEW QUESTION 90
Which field in the IPv6 header can be used to set the DSCP value?

A. Flow Label
B. Type of Service
C. Traffic Class
D. Precedence
E. EXP

Answer: C

Explanation:
Traffic Class
The Traffic Class field is an 8 bit field that is used to signify the importance of the data contained within this specific packet. With IPv4, this information was signified with the TOS field and supported both IP precedence and Differentiated Services Code Point (DSCP). The Traffic Class field used with IPv6 supports DSCP solely; this specification uses the first 6 bits to indicate the Per Hop Behavior (PHB) of the contained data; these PHB's are defined in RFC 2474 and its additions.

NEW QUESTION 95
Refer to the exhibit. R2cannot to learn RP information from XR3. Which issue is the source of the problem?

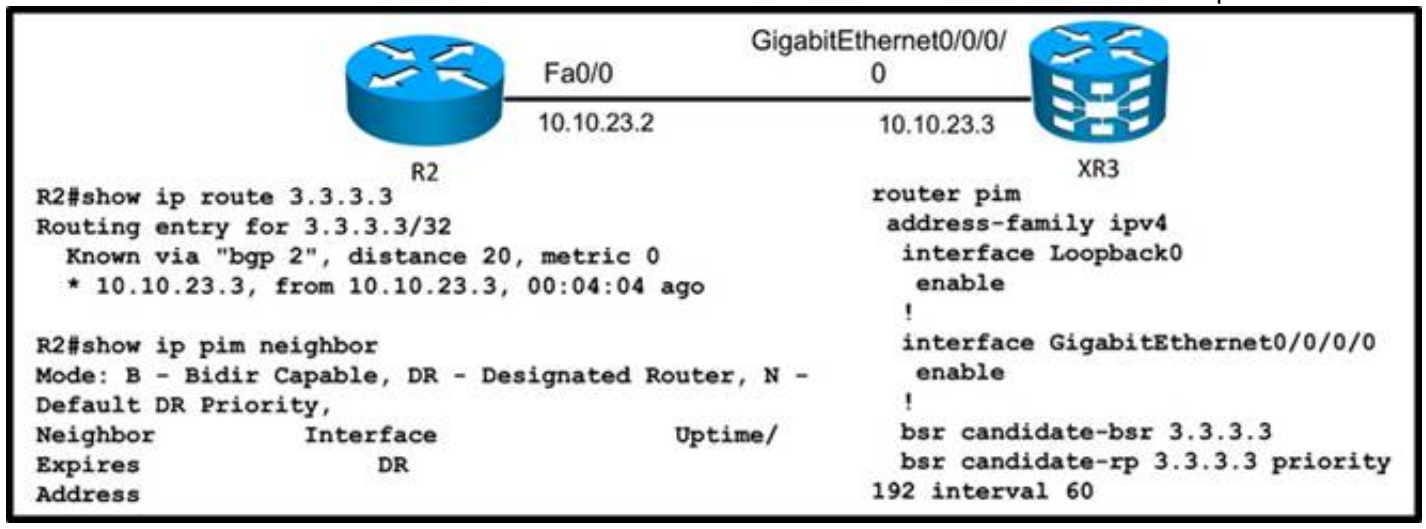

A. XR3 is not the DR.
B. Multicast routing is not enabled on the XR3 Giga0/0/0/0 interface.
C. R2 is learning the RP address via non-IGP routing protocol.
D. Multicast routing is not enabled on the XR3 Loopback0 interface.
E. BGP IPv4 MDT address family is not enabled on XR3.

Answer: D

NEW QUESTION 100
Refer to the exhibit.

```
interface loopback 0
 ipv4 address 10.0.0.1/24
 no shutdown
!
interface loopback 1
 ipv4 address 10.2.0.1/24
 no shutdown
!
ipv4 access-list acl1
 10 permit 224.11.11.11 0.0.0.0 any
!
ipv4 access-list acl2
 10 permit 224.99.99.99 0.0.0.0 any
!
multicast-routing
 interface all enable
!
router pim
 auto-rp mapping-agent loopback 0 scope 15 interval 60
 auto-rp candidate-rp loopback 0 scope 15 group-list acl1 interval 60 bidir
 auto-rp candidate-rp loopback 1 scope 15 group-list acl2 interval 60
!
end
```

Which three statements are correct regarding the Cisco IOS-XR configuration? (Choose three.)

A. This router, acting as the RP mapping agent, will send RP announcement messages to the 224.0.1.40 group
B. This router, acting as the RP mapping agent, will send RP discovery messages to the224.0.1.39 group
C. This router is the RP mapping agent only for the 224.11.11.11 and 224.99.99.99 multicast groups
D. This router is a candidate PIM-SM RP for the 224.99.99.99 multicast group
E. This router is a candidate PIM-BIDIR RP for the 224.11.11.11 multicast group
F. IGMPv3 is enabled on all interfaces
G. Other routers will recognize this router as the RP for all multicast groups with this router loopback 0 IP address

Answer: DEF

NEW QUESTION 101
Which statement is correct regarding MP-BGP?

A. MP-BGP can indicate whether an advertised prefix (NLRI) is to be used for unicast routing, multicast RPF checks or for both using different SAFIs.
B. MP-BGP uses a single BGP table to maintain all the unicast prefixes for unicast forwarding and all the unicast prefixes for RPF checks.
C. MP-BGP can be used to propagate multicast state information, which eliminates the need to use PIM for building the multicast distribution trees.
D. MP-BGP enables BGP to carry IP multicast routes used by MSDP to build the multicast distribution trees.

Answer: A

Explanation:
Protocol Independent Multicast
Protocol Independent Multicast (PIM) is a routing protocol designed to send and receive multicast routing updates. Proper operation of multicast depends on knowing the unicast paths towards a source or an RP. PIM relies on unicast routing protocols to derive this reverse-path forwarding (RPF) information. As the name PIM implies, it functions independently of the unicast protocols being used. PIM relies on the Routing Information Base (RIB) for RPF information. If the multicast subsequent address family identifier (SAFI) is configured for Border Gateway Protocol (BGP), or if multicast intact is configured, a separate multicast unicast RIB is created and populated with the BGP multicast SAFI routes, the intact information, and any IGP information in the unicast RIB. Otherwise, PIM gets information directly from the unicast SAFI RIB. Both multicast unicast and unicast databases are outside of the scope of PIM.
The Cisco IOS XR implementation of PIM is based on RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification. For more information, see RFC 4601 and the Protocol Independent Multicast (PIM): Motivation and Architecture Internet Engineering Task Force (IETF) Internet draft

NEW QUESTION 103
Which command set should be used for a 6to4 tunnel in a Cisco IOS XE router, considering the border interface with IPv4 address of 209.165.201.2?

A. interface Tunnel2002 ipv6 enableipv6 address 2002:D1A5:C902::1/128 tunnel source Ethernet0/0tunnel mode ipv6ip 6to4
B. interface Tunnel2002 ipv6 enableipv6 address 2002:D1A5:D902::1/128 tunnel source Ethernet0/0tunnel mode ipv6ip 6to4
C. interface Tunnel2002 ipv6 enableipv6 address 2002:D1A5:D902::1/128 tunnel source Ethernet0/0tunnel mode ipv6ip
D. interface Tunnel2002 ipv6 enableipv6 address 2002:D1A5:C902::1/128 tunnel source Ethernet0/0tunnel mode ipv6ip auto-tunnel
E. interface Tunnel2002ipv6 enableipv6 address 2002:D1A5:D902::1/128 tunnel source Ethernet0/0tunnel mode ipv6ip auto-tunnel

Answer: B

NEW QUESTION 104
Which two options areadvantages of an IPv6 dual-stack implementation in an enterprise environment? (Choose two.)

A. simplifies the route redistribution policies complexity
B. requires IPv6-to-IPv4 translation on the uplinks to the service providers
C. provides built-in support for Kerberos authentication
D. does not have to worry about NAT traversal
E. supports multicast properly

Answer: DE

NEW QUESTION 107
Which multicast group range is reserved for SSM?

A. 224.0.0.0/8
B. 225.0.0.0/8
C. 232.0.0.0/8
D. 239.0.0.0/8

Answer: C

Explanation:
PIM-SSM Operations
PIM in Source Specific Multicast operation uses information found on source addresses for a multicast group provided by receivers and performs source filtering on traffic.
•By default, PIM-SSM operates in the 232.0.0.0/8 multicast group range for IPv4 and ff3x::/32 (where x is any valid scope) in IPv6. To configure these values, use the ssm range command.
•If SSM is deployed in a network already configured for PIM-SM, only the last-hop routers must be upgraded with Cisco IOS XR software that supports the SSM feature.
•No MSDP SA messages within the SSM range are accepted, generated, or forwarded

NEW QUESTION 110
DRAG DROP

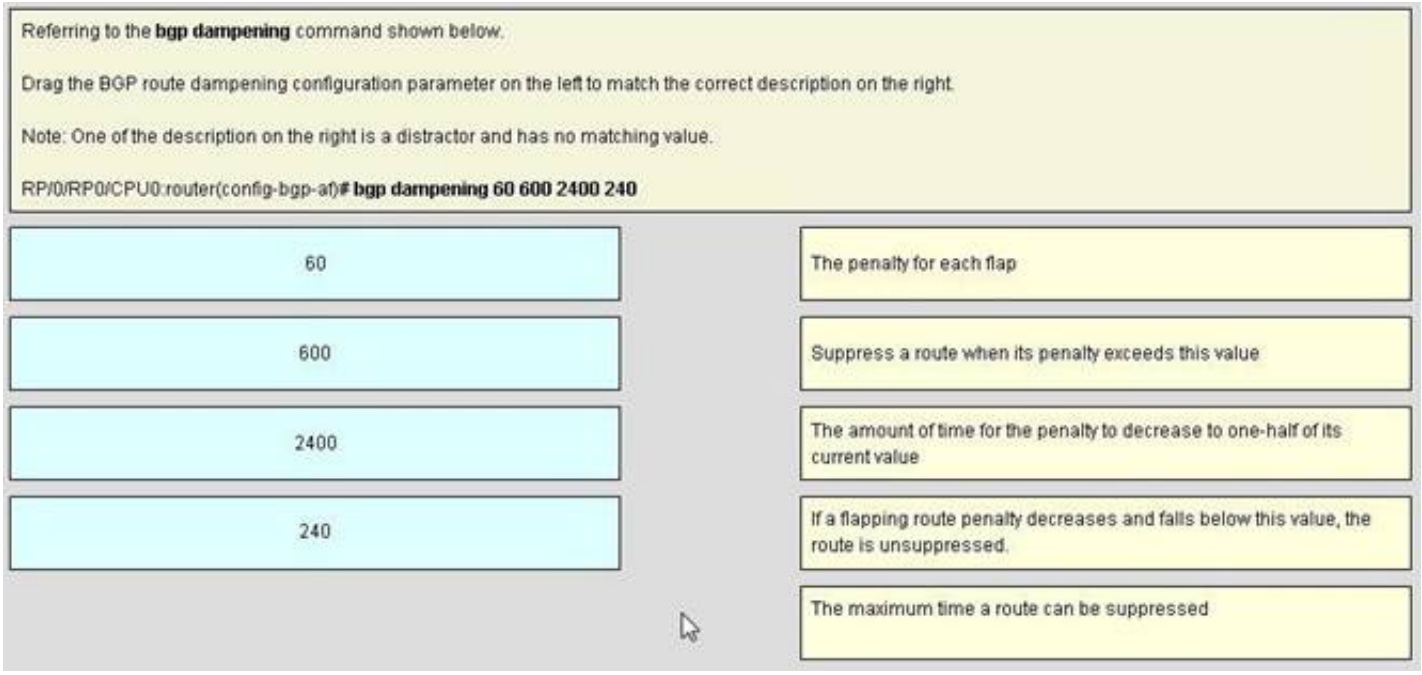

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The amount of time for the penalty to decrease to one-half of its current value - 60 Suppress a route when its penalty exceeds this value - 2400
If a flapping route penalty decreases and falls below this value , the route is unsuppressed
- 600
The maximum time a route can be suppressed – 240

bgp dampening

To enable Border Gateway Protocol (BGP) route dampening or change various BGP route dampening factors, use the **bgp dampening** command in address family configuration mode. To disable route dampening and reset default values, use the **no** form of this command.

bgp dampening [*half-life* [*reuse suppress max-suppress-time*] | **route-policy** *route-policy-name*]

no bgp dampening

Syntax Description

half-life	(Optional) Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). Penalty reduction happens every 5 seconds. Range of the half-life period is from 1 to 45 minutes.
reuse	(Optional) Value for route reuse if the flapping route penalty decreases and falls below the reuse value. When this happens, the route is unsuppressed. The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000.
suppress	(Optional) Maximum penalty value. Suppress a route when its penalty exceeds the value specified. When this happens, the route is suppressed. Range is 1 to 20000.
max-suppress-time	(Optional) Maximum time (in minutes) a route can be suppressed. Range is 1 to 20000. If the *half-life* value is allowed to default, the maximum suppress time defaults to 60 minutes.
route-policy *route-policy-name*	(Optional) Specifies the route policy to use to set dampening parameters.

SO bgp dampening 60 600 2400 240 is:
60 half life
600 reuse
2400 suppress
240 max-suppress-time

NEW QUESTION 114
What are three BGP configuration characteristics of a multihomed customer that is connected to multiple service providers? (Choose three.)

A. The multihomed customer can use local preference to influence the return traffic from the service providers
B. The multihomed customer announces its assigned IP address space to its service providers through BGP
C. The multihomed customer has to decide whether to perform load sharing or use a primary/backup implementation
D. The multihomed customer must use private AS number
E. The multihomed customer configures outbound route filters to prevent itself from becoming a transit AS

Answer: BCE

NEW QUESTION 118
Refer to the exhibit for the outputs from an ASR9K router.

```
RP/0/RSP0/CPU0:PE1#show route ipv6
Wed Oct 26 20:57:46.433 UTC

Codes: C - connected, S - static, R - RIP, 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 - 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, G  - DAGR
       A - access/subscriber, (!) - FRR Backup path

Gateway of last resort is not set

L    2001:db8:10:1:1::1/128 is directly connected,
     09:20:18, Loopback0
i L2 2001:db8:10:1:10::1/128
     [105/20] via fe80::eab7:48ff:fe2c:a180, 07:59:22, GigabitEthernet0/0/0/0
C    2001:db8:192:168:101::/80 is directly connected,
     1d05h, GigabitEthernet0/0/0/0
L    2001:db8:192:168:101::10/128 is directly connected,
     1d05h, GigabitEthernet0/0/0/0

RP/0/RSP0/CPU0:PE1#ping 2001:db8:10:1:10::1/128
Wed Oct 26 20:58:01.969 UTC
%Bad hostname or protocol not running
```

Why did the ping fail?

A. The ping command is missing the ipv6 option: ping ipv6 2001:db8:10:1:10::1/128

B. There is a problem with the IS-IS configurations
C. The fe80::eab7:48ff:fe2c:a180 next-hop is not reachable
D. The prefix length should be removed from the IPv6 address in the ping command: ping ipv6 2001:db8:10:1:10::1
E. IPv6 is not enabled on the Gi0/0/0/0 interface
F. The IPv6 neighbor discovery protocol is not enabled on the Gi0/0/0/0 interface

Answer: D

NEW QUESTION 121
What is determined by running the same hash algorithm on all PIMv2 routers?

A. The SPT from the RP to the multicast source
B. The SPT from the last hop router to the multicast source
C. Auto RP election
D. Which BSR to use for a particular multicast group
E. Which RP to use from a set of candidate RPs in the RP set

Answer: E

NEW QUESTION 126
Assume that the R1 router is enabled for PIM-SM and receives a multicast packet sourced from 172.16.1.100, and the R1 router has multicast receivers on the Gi0/1, Gi0/2, Gi0/3 and Gi0/4 interfaces.

R1 routing table:

172.16.1.0/24 via Gi0/1
172.16.2.0/24 via Gi0/2
172.16.3.0/24 via Gi0/3
0.0.0.0/0 via Gi0/4

The multicast packet from the 172.16.1.100 source must arrive on which interface on the R1 router for it to be forwarded out the other interfaces?

A. Gi0/1
B. Gi0/2
C. Gi0/3
D. Gi0/4
E. Gi0/1 or Gi0/2 or Gi0/3 or Gi0/4
F. Gi0/2 or Gi0/3
G. Gi0/1 or Gi0/4

Answer: A

NEW QUESTION 131
On Cisco IOS-XR, which BGP process can be distributed into multiple instances?

A. BGP process manager
B. BGP RIB process
C. BGP speaker process
D. BGP scanner process
E. BGP dampening process

Answer: C

Explanation:
Cisco IOS XR allows you to control the configuration of the number of distributed speakers and enables you to selectively assign neighbors to specific speakers. On the CRS-1 platform, multiple speaker processes up to 15 may be configured. However, configuring all the different speakers on the primary route processor simply adds to the load on the single RP.
Distributed speaker functionality is useful if Distributed Route Processor (DRP) hardware is available to take advantage of process placement. Later sections in this chapter depict distributed
BGP and placement of BGP process speakers on DRPs on a CRS-1 router.
In addition to the speaker process, BPM starts the bRIB process once BGP is configured. bRIB process is responsible for performing the best-path calculation based on partial best paths received from the speaker processes. The best route is installed into the bRIB and is advertised back to all speakers. The bRIB process is also responsible for installing routes

NEW QUESTION 136
Refer to the exhibit.

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. If a certain command or command option is not supported, please try to use a different command that is supported.

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 and PE5 routers
and interpret the supported CLI commands outputs to answer the four multiple choice questions.

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

Exhibit1

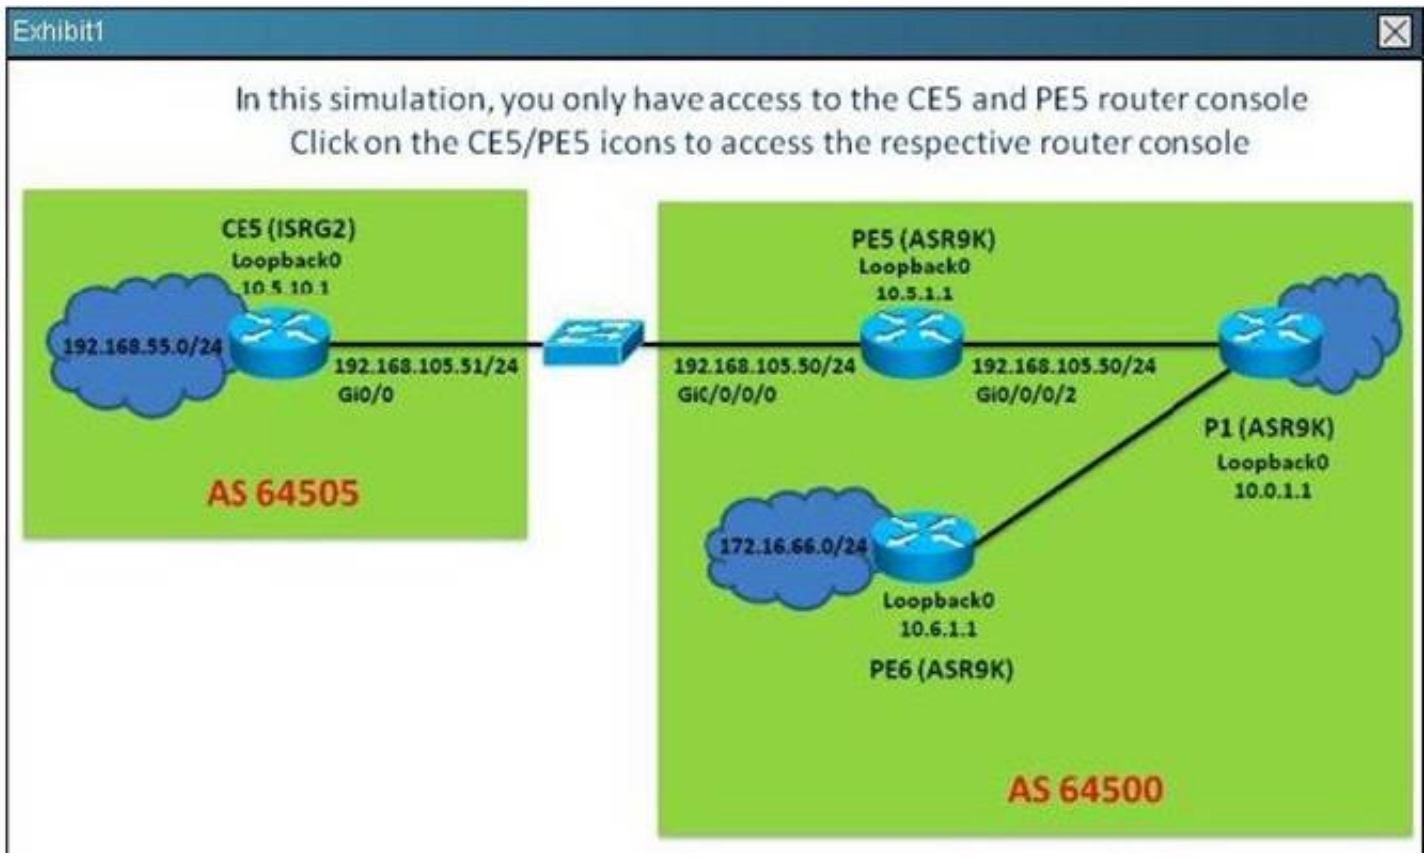

CE5

CE5#

Which two statements regarding the BGP peerlngs are correct? (Choose two)

A. On PE5,the incoming prefixes received from the 192.168.105.51 EBGP peer is limited to a maximum of 10 prefixes
B. On PE5, the "rplin" inbound route policy is applied to the 192.168.105.51 EBGP peer
C. On PE5, the "pass" outbound route policy is applied to the 192.168.105.51 EBGP peer
D. PE5 has one EBGP peer (CE5) and two IBGP peers (P1 and PE6)
E. PE5 has received a total of 60 prefixes from its neighbors

Answer: AE

Explanation:
#show ip bgp

NEW QUESTION 138
With PIM-SM operations, which four pieces of information are maintained in the multicast routing table for each (*,G) or (S,G) entry? (Choose four.)

A. RPF Neighbor
B. RP Set
C. Incoming Interface
D. OIL
E. DF priority
F. PIM SM state flags

Answer: ACDF

Explanation:
The following is sample output from the show ip mroute command for a router operating in sparse mode:
show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode
(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
Incoming interface: Tunnel0, RPF neighbor 10.3.35.1, Dvmrp Outgoing interface list:
Ethernet0, Forward/Sparse, 5:29:15/0:02:57
(198.92.46.0/24, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
Incoming interface: Tunnel0, RPF neighbor 10.3.35.1 Outgoing interface list:
Ethernet0, Forward/Sparse, 5:29:15/0:02:57

NEW QUESTION 141
Refer to the Cisco IOS configuration exhibit.

```
interface Gi0/0
 ip multicast boundary 1
!
access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40
```

Which statement is correct?

A. This configuration is typically configured on the boundary routers within a PIM SM domain to filter out malicious candidate-RP-announce and candidate-RP-discovery packets
B. This configuration is typically configured on the RPs within a PIM-SM domain to restrict the candidate-RP-announce packets

C. This configuration is typically configured on the mapping agents within a PIM-SM domain to restrict the candidate-RP-discovery packets
D. This configuration is typically configured on the MSDP peering routers within a PIM-SM domain to filter out malicious MSDP SA packets

Answer: A

NEW QUESTION 145
Refer to the exhibit.

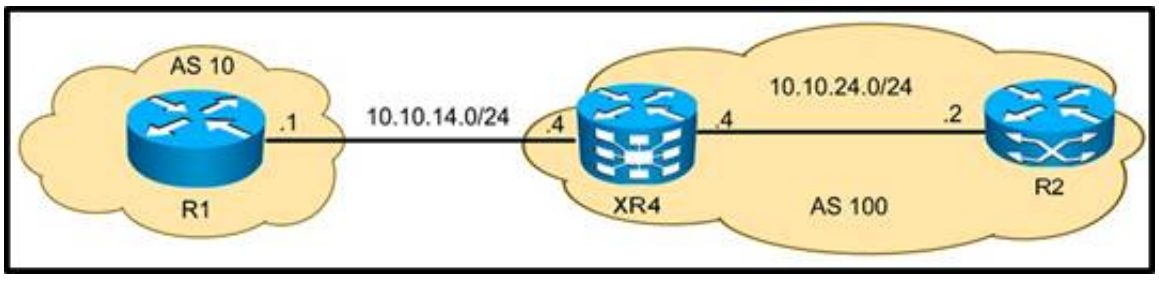

XR4 must protect itself from a DOS attack against its BGP process from R1 by using the TTL security feature. Which configuration achieves this goal?

A. router bgp 100neighbor 10.10.14.1 ttl-security
B. router bgp 100neighbor 10.10.14.1 ttl-security hops 1
C. router bgp 100neighbor 10.10.14.1 ttl-security hops 254
D. router bgp 100neighbor 10.10.14.1 ttl-security hops 255

Answer: A

NEW QUESTION 149
.....

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 642-885 Exam with Our Prep Materials Via below:

https://www.certleader.com/642-885-dumps.html